

The University of Sheffield

Doctoral Thesis

Investigating the Value of In-Situ Monitoring in the Laser Powder Bed Fusion of Nickel-Based Superalloys

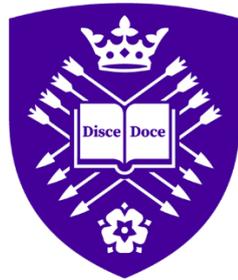
Author:

Alexander Sloane

Supervisors:

Prof. Katerina Christofidou

Prof. Iain Todd



University of Sheffield

A thesis submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy

Faculty of Engineering
School of Chemical, Materials, & Biological Engineering

October 24th, 2025

Declaration of Authorship

I, Alexander Sloane, the author, confirm that the Thesis is my own work. I am aware of the University's Guidance on the Use of Unfair Means (www.sheffield.ac.uk/ssid/unfairmeans). This work has not been previously presented for an award at this, or any other university. I confirm that:

- This work was done wholly while in candidature for a research degree at this university.
- Where I have consulted the published work of others, this is clearly attributed.
- Where I have quoted from the work of others the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all the main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself

You won't get anything done by planning

Karl Pilkington

Abstract

Laser Powder Bed Fusion (L-PBF) is an additive manufacturing method, in which 2D patterns are selectively melted into successive layers of metal powder by a high power laser, resulting in the creation of a 3D component. L-PBF offers a range of advantages over traditional manufacturing techniques: otherwise-impossible geometries, such as internal cooling channels; the ability to build complex systems as a single part; and rapid design-to-production timelines.

Several problems must be overcome before L-PBF technology can achieve widespread adoption, however. Rapid thermal cycling causes unpredictability during solidification, resulting in geometric distortions, anisotropy, porosity, and undesirable microstructures in as-built components.

In-situ monitoring methods, including real-time imaging of the melt pool with high-speed cameras, are a vital tool in solving these problems. Should engineers be able to detect and correct for defects during the build, a significant step forward will have been made towards mainstream L-PBF usage.

The aim of this project was to add to the existing literature on the use of in-situ monitoring in L-PBF, specifically in terms of the identification of pore formation and the prediction of microstructure, in nickel-based superalloys. Various builds were completed, ranging from powder-free, single line scans, to geometrically complex components. Infrared camera data was collected, alongside ex-situ data from SEM, EBSD, and XCT. These datasets were compiled and cross-referenced, to discover correlations and trends between melt pool images, and the resulting build.

The strengths and limitations of using melt pool images to predict pore and microstructure formation were revealed. While it was found to be challenging, with the equipment tested, to identify the creation of specific pores, it was found that distinct processing regions exist, in which the likelihood of pore formation was greater.

This research builds upon previous work by other authors, advancing L-PBF towards the predictability required for widespread adoption by industry.

Acknowledgments

This research project would not have been possible without the supervision and assistance of a large group of people, both at the University of Sheffield and at affiliated organisations, to whom I am grateful for all their help. In particular, I would like to thank the following:

Prof. Katerina Christofidou, for her ongoing, day-to-day support, in all matters large and small, and for providing me with all the necessary guidance, throughout the entire duration of my PhD.

Prof. Iain Todd, for securing the funding for my project, for assembling Sheffield's Additive Manufacturing research group, and for ensuring that all the behind-the-scenes aspects ran smoothly.

Dr Minh Phan, for trusting me as a collaborator on his research projects during my early years at the university, and thereby playing a crucial role in shaping my project's path in my final years.

Dr George Maddison, for all his guidance and assistance in operating the Additive Manufacturing machines in the university, and for his rapid responses whenever a fault arose.

Dr Yufeng Lai, for his instrumental support in designing and installing the in-situ monitoring equipment used throughout this project, and for his ongoing support in my use of the device.

Dr Robert Snell, for all the help in my earliest days in the research group, when I was faced with an overwhelming amount of theory and practice to learn, and he provided support and training.

Dr Ben Thomas, Dr Alex Goodall, and Dr Felicity Freeman, for all the direction, assistance, and encouragement that each brought, in the years we worked alongside each other at the university.

Dr Elaine Livera and Henry Saunders, for all the brainstorming, questions, advice, discussion, and non-work chats that happened between us as we sat in office H8 in the Hadfield building.

Hugh Banes and Dr Miguel Espadero, for lending their skills and knowledge in the domain of modelling, and for being always available in the domain of idea sense-checking.

Dr Jack Donoghue, for facilitating the large amount of EBSD sample characterisation that occurred at the University of Manchester, without which this project would be much smaller in scale.

Jonah Shrive, for helping me analyse the data that accompanied this EBSD, running his codes on my data, thereby allowing me to focus on the more relevant aspects of my experimental work.

Karel Mrzena, Milo Maguire, Mostafa Salem, Dr Robert Deffley, Dr Frances Livera, Dr Lucy Farquhar, Dr Josh Berry, and every other member of the Additive Manufacturing research group.

Dr Lewis Owen, Ben Jolly, Dr Frances Synnott, Ciaran Miles, Joe Smith, Karol Murgabia, Martin Tse, Matt Turton, Niels Schreiner, Tom Cole, Sophie Barwick, and all the Modern AlChEME group.

Dr Partha Paul, Dr Jaianth Vijayakumar, David McArthur, Prof. Chu Lun Leung, Dr Vincent Fernandez, and Dr Joseph Brunet, for running the XCT at the European Synchrotron Radiation Facility.

The team at the Mechanical Workshop at the University of Sheffield, for always carrying out EDM jobs, to remove my built samples from their base plates, without delay.

Boeing, for sponsoring my PhD project, providing the funds for the experimental work.

Sharon Brown, Jacob Whittle, William Robertson, Kamila Nowosad, Sian Odell, Sam Engel, Lakshmi Shenoy, Mimi Olukoya, and everyone else I've known professionally and socially at the University of Sheffield and University of Manchester, over my four years studying my PhD.

Table of Contents

1	Introduction.....	1
2	Literature Review	4
2.1	Introduction to L-PBF.....	4
2.2	Melt Pool Physics	12
2.3	Alloys and Microstructure.....	25
2.4	Defects.....	42
2.5	Factors Affecting the Build.....	70
2.6	In-Situ Process Monitoring.....	94
2.7	Literature Review Conclusions	134
3	Experimental Methods	135
3.1	Powder Alloys.....	135
3.2	L-PBF Machine.....	143
3.3	Thermal Camera.....	144
3.4	Metallographic Preparation	149
3.5	Microscopy	152
3.6	X-Ray Computed Tomography	153
3.7	Hardness Measurements.....	155
3.8	ThermoCalc Modelling.....	155
4	Camera Characterisation	157
4.1	Determining Camera Frame Edges.....	157
4.2	Determining Spatial Accuracy.....	160
4.3	Determining Spatial Resolution.....	163
4.4	Understanding the Effects of Motion Blur.....	166
4.5	Understanding the Effects of Camera Focus.....	170
4.6	Basic Evaluation of Camera	176
4.7	Accounting for Limitations of Emissivity	189
4.8	Comparing to Ex-Situ and Modelling Methods	204
4.9	Camera Capability Conclusions	211

5	In-Situ Monitoring and Microstructure	213
5.1	Experimental Setup.....	213
5.2	Data Gathering	214
5.3	Data Analysis.....	222
5.4	Regression Modelling.....	246
5.5	In-Situ Microstructure Conclusions.....	254
6	In-Situ Monitoring and Mesostructure	255
6.1	Parameter Selection	255
6.2	In-Situ Monitoring and Deliberate Porosity	262
6.3	Industrial Applications of L-PBF	286
6.4	In-Situ Mesostructure Conclusions	302
7	Conclusions.....	303
8	Future Work	305
8.1	In-Situ Monitoring Equipment	305
8.2	Microstructural Development.....	307
8.3	Mechanical Defect Development	309
8.4	Future Work Conclusions	310
9	Appendix.....	311
9.1	Literature Review	311
9.2	In-Situ Monitoring and the Microstructure.....	318
9.3	In-Situ Monitoring and the Mesostructure.....	337
10	Bibliography.....	340

Glossary

Additive Manufacturing (AM): A modern manufacturing process involving the layer-wise fusion of two-dimensional shapes to create a three-dimensional component.

Anisotropy: A condition where the mechanical properties of a component vary depending on the direction in which they are tested, often caused by unidirectional solidification.

Beam Trapping: A phenomenon in transition or keyhole melting modes where a portion of reflected light is reabsorbed within a vapour depression, deepening the melt pool.

Cellular Automata (CA): A microstructural modelling technique that uses discrete grids of cells to simulate the formation of grains and phases during solidification.

Closed Loop Control: An advanced system that uses in-situ sensors to detect defects in real-time and automatically adjusts machine parameters to correct them.

Computational Fluid Dynamics (CFD): A complex modelling method used to simulate mass flow, such as convection currents and vapour depressions, by solving the Navier-Stokes equations.

Constitutional Undercooling: A solidification mechanism where the liquid composition in front of the interface is altered by segregation, making further growth energetically favourable.

Denudation: A phenomenon where powder particles adjacent to a scan track are blown away by gas flows from the melt pool's vapour plume.

Electron Backscatter Diffraction (EBSD): A microscopic characterisation technique used to determine crystallographic information, including grain size, shape, and orientation.

Emissivity: A temperature-dependent property that dictates how effectively a body emits thermal radiation, which is essential for converting light intensity into temperature readings.

Epitaxial Growth: A common solidification process in L-PBF where the liquid at the bottom of a melt pool nucleates coherently with the underlying substrate, causing grains to extend through multiple layers.

Ex-Situ Data: Information gathered from a component after the build process is complete using laboratory equipment such as optical microscopes or X-ray scanners.

Grain Average Misorientation (GAM): A metric derived from EBSD that measures the internal strain of a grain by calculating the average misorientation between its pixels and their neighbours.

Hatch Spacing: The defined distance between adjacent parallel laser tracks within a single build layer.

Hot Isostatic Pressing (HIP): A post-processing method that uses high temperature and high-pressure inert gas to compact internal defects like pores and cracks.

In-Situ Monitoring: The use of sensors to collect real-time data on electromagnetic, acoustic, or electrical signals during the manufacturing process.

Keyhole Porosity: Voids created when a deep, unstable vapour depression collapses and traps metal vapour bubbles within the solidifying alloy.

Lack of Fusion (LoF) Porosity: Irregularly shaped voids formed when laser energy is insufficient to melt all powder particles or provide enough overlap between scan tracks.

Laser Powder Bed Fusion (L-PBF): A sub-branch of additive manufacturing where a high-power laser selectively melts patterns into successive layers of metallic powder.

Marangoni Force: A primary force within the melt pool caused by variations in surface tension that results in forced convective flow.

Melt Pool: The localized volume of molten metal created by the interaction between the laser beam and the powder feedstock.

Near-Net Shape (NNS): Components produced via L-PBF that require minimal material removal or subtractive machining before they are ready for end-use.

Recoil Pressure: The pressure exerted on the melt pool surface by the rapid volumetric expansion of metal molecules as they are vaporised.

Scanning Electron Microscopy (SEM): A microscopy method that uses a focused beam of high-energy electrons to image the sub-grain microstructure and other small features.

Segregation: The non-uniform spatial distribution of alloying elements that occurs during the transition from the liquid to the solid phase.

Spatter: The ejection of molten material or powder particles from the melt pool area, which then solidify and land elsewhere in the build chamber.

Staircase Effect: A geometrical inaccuracy inherent to layer-wise manufacturing where discrete layers appear as steps on upwards-facing slopes.

Superalloy: A category of alloys, often nickel-, cobalt-, or iron-based, designed to maintain mechanical integrity under extreme stress and high temperatures.

Thermal Gradient: The rate at which temperature decreases along a directional line extending from the peak temperature within a melt pool.

X-Ray Computed Tomography (XCT): A non-destructive characterisation technique that uses X-rays to create a 3D map of internal features such as pores and cracks.

1 - Introduction

Additive Manufacturing (AM) is a modern manufacturing technique, which involves the often-layerwise fusion of two-dimensional shapes, to form a three-dimensional component. Laser-Powder Bed Fusion (L-PBF) is a sub-branch of AM, in which successive layers of metallic powder are coated upon each other, with a laser selectively melting shapes into each new layer after coating, thereby transforming powder feedstock into solid parts.

L-PBF allows for the creation of components that would be difficult or fully impossible via traditional methods. Fuel nozzles in airplane engines, which typically consist of multiple separate parts brazed together, have been produced as a single component [3]. Curved cooling channels within industrial turbine blades, which cannot be machined due to their geometry, are straightforward to create with L-PBF [4]. Fine lattice structures within artificial bone implants, can be built to osseointegrate inside the human body [5].

L-PBF requires no tooling, or product-specific setup. Whereas an automotive factory may require years of preparation time before a newly-designed component is ready for manufacture, with L-PBF the preparation time may be measured in hours [6]. The per-component manufacturing cost with L-PBF can be unaffected by scale: producing one, custom-designed part is as economically feasible as mass producing one million [7]. This lack of requirement for large, complex factories, means that widespread adoption of L-PBF allows the potential for a worldwide democratization in manufacturing, with local workshops in any country being equally as capable at production as big corporations. L-PBF is even being considered as a reliable method of self-sustaining production for future colonies on the moon or Mars [8].

Before such ambitious futures can be realized, however, there remain many hurdles to overcome. Material within the L-PBF build chamber undergoes extreme temperature fluctuation, with material melting and solidifying multiple times per second, as the laser spot passes above it and near it. Pores, cracks, residual stresses, rough surfaces, anisotropy, and poor microstructures are all commonplace, whenever components more complex than cubes are attempted [9]. Enormous unpredictability is currently inherent to the L-PBF process, with parts produced under seemingly-identical conditions often displaying large differences in macro, meso, and micro-structural results. The aforementioned fuel nozzles and turbine blades, while impressive as showcase pieces, cannot meet the quality assurance requirements of the civilian marketplace [10].

Many research institutions are attempting to address these issues with L-PBF, using a wide range of experimental techniques. One of the most prominent set of tools in these institutions, are in-situ monitoring devices, which allow for real time feedback of the quality of the build process. Such devices are pivotal in the ultimate end goal of much current L-PBF research: the implementation of closed loop control systems inside production machines [11].

In such a system: in-situ monitoring sensors detect a range of electromagnetic, acoustic, or electrical signals; these signals are fed into a processor, which determines, in real time, whether the build is progressing as intended, or whether cracks, pores, or other such defects are present; and if defects are detected, machine parameters (such as the laser power, the laser speed, the distance in between laser tracks, or the pattern which the laser follows) are updated in order to prevent further defects, or even heal the defects of previous layers.

Developing and implementing closed loop control is an enormous research challenge, which encompasses several separate but related goals. This project deals specifically with in-situ monitoring, and relating detected signals to pore defects, to mechanical properties, and to microstructural features. The objectives of this work are as follows:

1. Evaluate the in-situ monitoring equipment at the University of Sheffield's AM workshop, to determine how accurately it records the laser-powder interactions inside the build chamber
2. Determine the extent to which melt pool images can be used to predict the formation of porosity within built components
3. Determine whether thermal signatures can reveal information about the microstructural and mechanical properties of the component

The Literature Review chapter presents a detailed overview of the exact processes occurring inside a L-PBF machine, inside the melt pool, and inside the solidifying alloy. The cause-and-effect patterns of defects are thoroughly discussed, as well as the many methods of designing geometries and selecting machine parameters to prevent them. Finally, various types of in-situ monitoring are examined, and the works of previous authors discussed, to better understand the strengths and limitations of the current technology, and where exactly further research is required.

The Experimental Methods chapter characterises the two nickel-based powders that will be experimented on in this work, and details the L-PBF machine, the in-situ monitoring device, and the rest of the laboratory equipment that will be used throughout.

Chapter Four explores the characterisation of the in-situ monitoring equipment, providing insight into the strength and limitations of the devices, and discovering the corrections and compensations that must be applied to results in order to make them as accurate as possible.

Chapter Five focuses on what in-situ monitoring can reveal about the microstructure. A range of simple parts are built, with various machine parameters, and their internal structure examined at the scale of the millimeter (for porosity and cracking), the micrometer (for grain size and shape), and the nanometer (for sub-grain cellular features). Correspondences are searched for, between the in-situ signals and these macro, meso, and microstructural results.

Chapter Six explores an attempt to deliberately introduce a variety of porosity defects into a set of test pieces. These pieces are both monitored in-situ during building, and undergo X-ray imaging ex-situ, to reveal pores within their internal structure. An industrially-relevant component, with a

complex geometry, is analysed in this manner, in order to demonstrate the differences between in-situ monitoring in the building of simple shapes, and in the building of real world components.

The Conclusions chapter will establish the value of this work within the framework of L-PBF research.

Finally, the Future Work chapter places the experimental findings in the context of existing literature, and provides recommendations on the ways in which the research can be expanded upon, with the aim of building towards the ultimate goal towards closed loop control.

2 - Literature Review

In order to fully understand and explain the current state of the art of Laser Powder Bed Fusion, the gaps in knowledge, and the aspects which require further research, a comprehensive literature review was undertaken before any experimental work was carried out. This chapter summarises the review, presenting the information that is relevant to the experiments.

The chapter is broken into seven sections:

1. Introduction to L-PBF - A broad overview of the technique, the strengths, the limitations, and the possible future of L-PBF
2. Melt Pool Physics - An in-depth exploration of the manner in which laser energy melts metallic powder, forms a melt pool, and re-solidifies into bulk alloy
3. Alloys and Microstructures - A review of the basics of metallurgy, followed by a discussion of metallurgy in L-PBF, and the specific alloys used in this project
4. Defects - A comprehensive breakdown of all the geometrical, surface quality, microstructural, and mechanical issues that can arise in L-PBF
5. Factors Affecting the Build - A discussion of the design considerations that engineers and technicians must make, when attempting to create a component that matches specifications
6. In-Situ Process Monitoring - An overview of the various signals that can be sensed during the L-PBF process, in order to provide feedback on the build's quality
7. Literature Review Conclusions - A summary of the literature gaps identified, which will be explored experimentally in further chapters

2.1 - Introduction to L-PBF

This section introduces the basic concepts upon which the rest of the literature review is based; defines the exact category of additive manufacturing being studied; and briefly explores the history, future, advantages, and disadvantages of the manufacturing technique.

2.1.1 - Additive Manufacturing

The standards organisation ASTM International defines Additive Manufacturing (AM) as “the process of joining materials to make objects from 3D model data, usually layer upon layer, as opposed to subtractive manufacturing methodologies, such as traditional machining” [12]. The world's first AM machine was created in 1983 by Charles Hull, who three years later founded the company 3D Systems, which began manufacturing and selling Hull's machines [13]. This original design used focused beams of ultraviolet light, to selectively cure photopolymer resin, in a process known as stereolithography [14]. Over the next several decades, a wide range of alternative AM processes and designs were developed, upon which numerous research organisations and commercial companies were founded. Since the turn of the millennium, standards and classification

systems have begun being applied to the ever-growing field. The term Additive Manufacturing was only formally adopted by the ASTM in 2009 [15].

Regardless of the exact process employed, all AM machines are intended to achieve the same underlying principle: the creation of a complex, three-dimensional part, without the need for additional tooling. One unifying requirement of all AM processes is that of the digitisation of solid objects: the ability of a computer to transform shapes into points and lines, and guide automated tools through the retransformation of these points and lines back into shapes. As a result, the advances of AM over the past forty years have gone hand-in-hand with the advancement of computing power of this period [16]. Today, the technology has matured to the extent that it is applicable to a growing number of industrial and end-user applications, most notably the aerospace, medical, and automotive sectors [17].

In order to bring some structure to the wide range of different AM processes, researchers classify machines by the medium that the feedstock exists in before it is joined to the part being produced (powder, liquid, solid, or wire filament), and the method used to join the feedstock to the part (heat, photons, electrons, or adhesive) [18]. As shown in Figure 2.1, the ASTM has classified all AM processes into seven broad categories [19], including Powder-Bed Fusion, the family of AM devices that will be addressed in this work. These devices operate through the selective melting of two-dimensional shapes, into successive layers of metal powder particles [20].

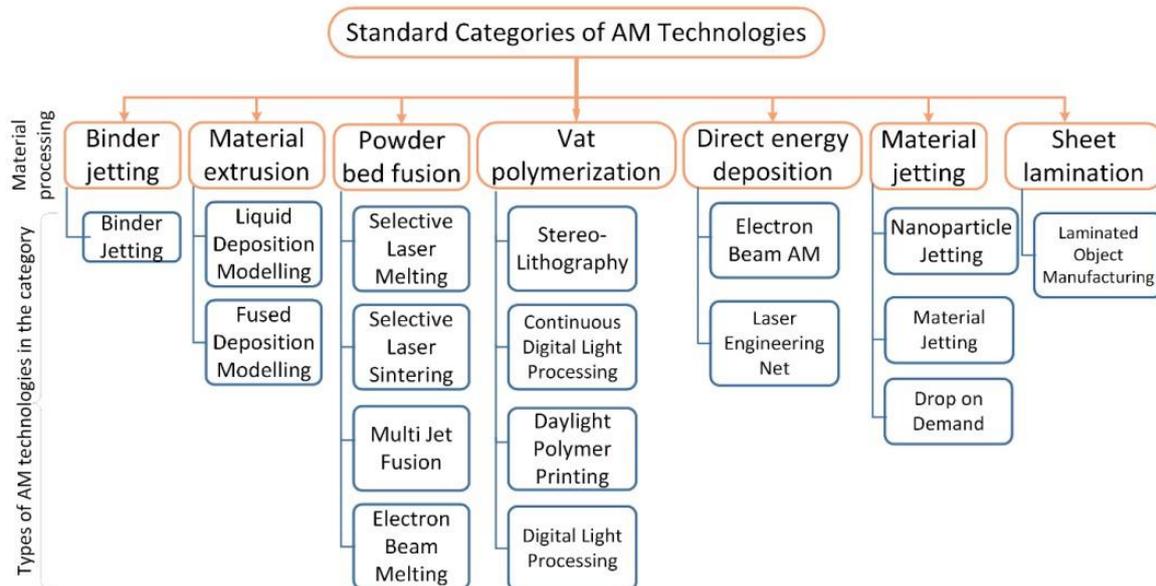


Figure 2.1

The seven categories of Additive Manufacturing, and their sub-categories [19]

Metallic Powder-Bed Fusion methods can be further broken down into two sub-categories, based on whether the energy used to selectively melt the powder particles is supplied by an electron beam, or a laser beam [21]. It is the latter of these categories, Laser Powder-Bed Fusion (L-PBF) that shall be focussed on from this point onwards.

2.1.2 - L-PBF Machines

As shown in Figure 2.2, all L-PBF machines incorporate the same basic components: a build chamber, into which a build plate is lowered vertically as a component is built upon it; a powder supply feed, which may exist as another chamber next to the build plate, or a container elsewhere in the machine; a spreading device, usually a rake or roller, which deposits thin, uniform layers of powder from the supply feed onto the build chamber, every time the latter sinks into its chamber; a laser, which supplies the energy required to melt a two-dimensional shape into each successive layer of powder; a set of scanning mirrors, which direct the laser beam across the shape being scanned; and an inert gas supplier which ensures that the whole system operates in an extremely-low-oxygen environment [9].

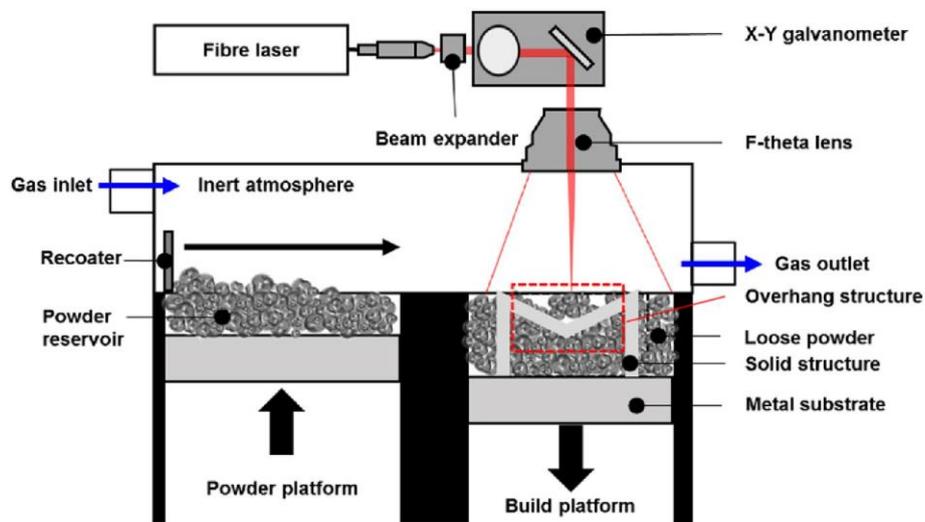


Figure 2.2

The basic components of any L-PBF machine [22]

In selecting which kind of laser is most appropriate for a L-PBF machine, the primary concern is the extent to which the wavelength produced can be absorbed by the intended metal powder particles. Most commercial machines use a neodymium-doped yttrium aluminium garnet (Nd:YAG) crystal as the gain medium, which produces light at a wavelength of $1.06\mu\text{m}$, of which 64% can be absorbed by iron atoms, for example, and 77% by titanium [23].

Even with an appropriate wavelength selected for the material in question, the photons of the laser beam may not be able to impart enough thermal energy into the powder particles to cause sufficient melting. In order to increase this likelihood, some L-PBF machines operate their lasers in pulsed mode, rather than as a continuous beam. The advantages of using short pulses (with durations in the range of tens of nanoseconds to tens of milliseconds) are twofold: the power of the laser during a pulse can be orders of magnitude greater than the power of an equivalent continuous beam, allowing for much greater melting of the powder beneath the laser spot; and due to the extremely short timescale, there is a diminished opportunity for heat conduction to melt powder not underneath the laser spot [24].

Another machine component that has major influence on the quality of a finished part is the powder spreading device. Horn et al. [25] studied the effects of altering recoater types in a L-PBF process. Three common recoaters were used: carbon fibre brushes, polymer lips, and a steel blade, as shown in Figure 2.3. It was found that, when building parts for which ideal laser parameter settings have been established, the steel blade provides the greatest accuracy and repeatability of powder layer spreading. For parts that may undergo swelling (excessive heat causes solidified surfaces to rise above the bed level after scanning), however, because either their ideal parameters are unknown or intentionally being experimented with, the softer brushes or polymer is more desirable. This is because contact with the steel blade will result in damaging particles being abraded into the powder bed, and also the deformed blade not spreading the powder uniformly for the rest of the build. The brushes and polymer, though less accurate overall, are far less likely to experience these issues.

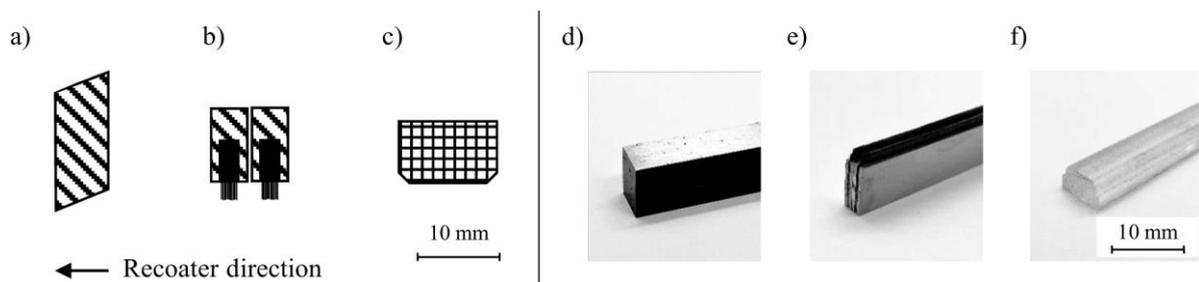


Figure 2.3

Cross-section schematics, and actual images, of the recoater types used by Horn: steel blade (a) and (d); carbon fibre brushes (b) and (e); polymer lips (c) and (f) [25]

The list of additional machine components which affect the quality of a L-PBF-produced part is extensive. Much could be written about: the galvanometers which control the mirrors that direct the laser beam; the lenses which shape the beam into a focussed spot; the accuracy of the piston beneath the build plate; and the rate of gas flow, which removes vapour plumes from the path of the laser; to name just a few of many examples [26]. Some of these components will be discussed later in this chapter, while others will have to be omitted for the sake of scope.

The majority of innovation being conducted by L-PBF machine manufacturers, attempts to meet two demands: increased part quality, and increased build rates. In the former category, the uses of ultra-short laser pulses and full vacuum systems (i.e. without even an inert gas present in the build chamber) are being explored. In the latter category, some manufacturers are introducing multiple lasers into the build chamber, increasing laser power, and increasing powder layer thickness, to skyrocket build rates from an industry-norm of approx. 50cm³/h, to greater than 150cm³/h [27].

2.1.3 - Advantages of L-PBF

The list of reasons why a manufacturing company may desire to produce parts through L-PBF, rather than with traditional, subtractive processes, are many. The most prominent will be discussed here.

Depending on the industry, process, and application, an extreme amount of material may be removed from the raw metal billets that a manufacturing company purchases, during subtractive machining. In the aerospace industry, for example, it is not uncommon for machining processes to remove 80 - 90% of the mass of forged titanium components, with these offcuts contributing up to 60% of the component's manufacturing costs [28]. L-PBF can produce Near-Net Shape (NNS) parts, which require minimal material removal before the as-built component is ready for end-use. Since the unused and unmelted powder that surrounds the part inside the build chamber can be often be recycled, the amount of metal waste through by L-PBF can be very small (although it should be noted that powder production wastage, and wastage from failed builds and supports, can be significant).

L-PBF methods are capable of producing certain shapes and designs that would be entirely impossible with traditional means; or at least impossible to create from a single component. The most obvious example of these are non-linear internal channels, as shown in Figure 2.4. Liquid rocket engines, which require high-pressure fluids to be delivered to a combustion chamber, and turbine blades, which require coolant to be circulated throughout their bulk, are two applications for which L-PBF is yielding many benefits [29, 30].

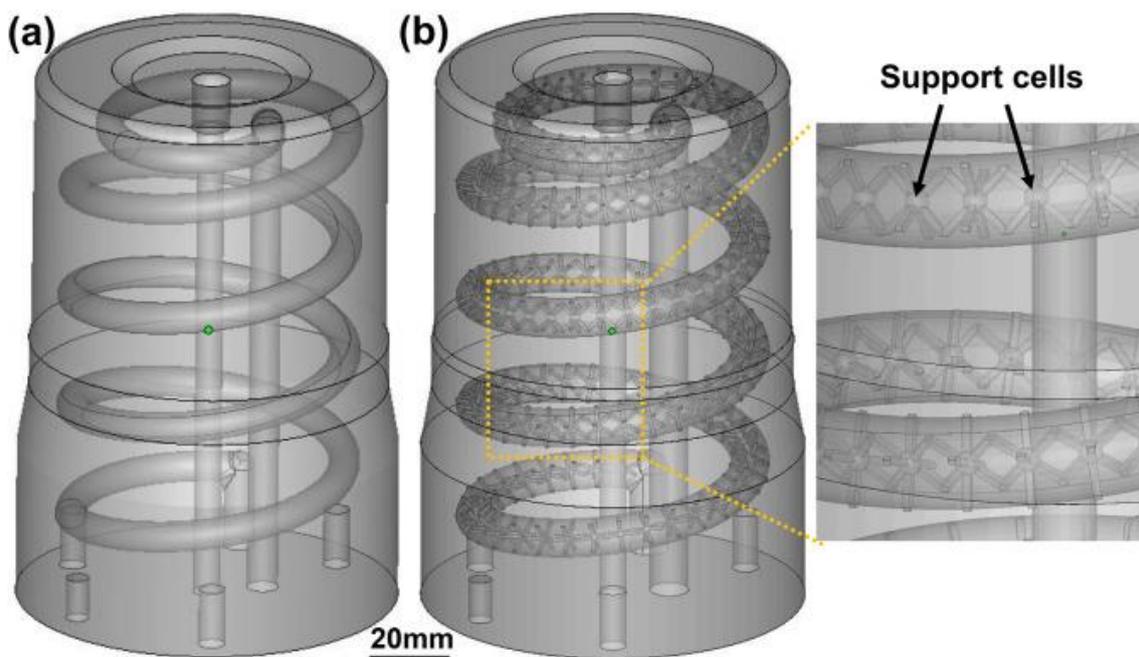


Figure 2.4

Curved internal cooling channels, with lattice cells for support, inside an injection mould produced by L-PBF. This geometry could not have been created as a single component by traditional methods. [31]

When using traditional methods, a manufacturing company may require years to acquire and configure all the tooling necessary to produce a new part, even if the part only deviates slightly from previous designs. With L-PBF, no such setup time exists. Though experimentation may be required to find the ideal machine parameters for a new part, and though there are design limitations on what geometries may be produced (discussed in detail later in this chapter), any L-PBF machine can begin producing any new part instantly, without the need for tool reconfiguration [32].

Another effect of the sometimes-enormous tooling requirements for traditional methods, as well as the long timeframes, is the fact that it can be economically infeasible for manufacturing companies to produce small batches of novel components. With traditional methods, there will always exist some minimum number of components that must be made and sold, to justify the up-front cost of acquiring and configuring the production machines. This minimum number also does not exist with L-PBF. Manufacturing companies using L-PBF can even operate profitably when producing parts that are customised for each individual customer, with batch sizes never exceeding one [33].

2.1.4 - Future of L-PBF

As of 2024, the global metal AM industry was worth over \$20 billion USD. This market value showed an 11% increase over the previous year, thanks largely to a 24% increase in the number of metal AM machines being sold [34]. In the field of medical engineering, where complex, intricate parts can be required to withstand the harsh environment of the human body, metal AM has shown substantial promise in recent years. L-PBF is being used to fabricate femoral hip stems that are custom-made to fit a patient’s exact body dimensions, and to create highly porous titanium alloy scaffolds to assist bone regeneration [35]. In the fields of both automotive and aerospace engineering, where lightweight components may need to be built to high tolerances, and survive extreme stresses and temperatures, L-PBF is proving adept at constructing complex designs as single components, from high-temperature superalloys [36].

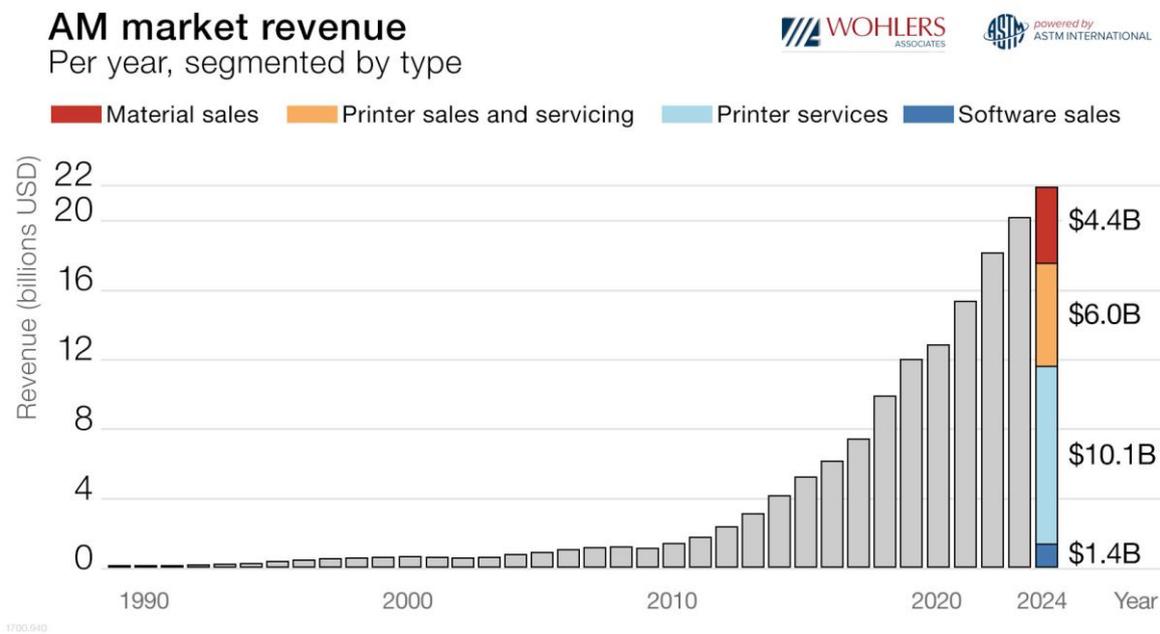


Figure 2.5
A breakdown of the Additive Manufacturing market, in 2024 [34]

As L-PBF continues to expand into multiple manufacturing industries, greater focus is being placed on the environmental sustainability of the method. Priarone et al. [37] conducted a review into the lifetime energy demand and carbon dioxide footprint of L-PBF parts, and compared these metrics to

traditionally-produced counterparts. It was concluded that, in their current state of relative-infancy, where failure rates are not insignificant, L-PBF methods require significantly more energy and output significantly more carbon dioxide than traditional methods, when comparing on a per-component basis. It was stressed, however, that as the process continues maturing, and the limitations and issues discussed later in this chapter are addressed, L-PBF has the potential to become a far more environmentally sustainable manufacturing method than traditional means.

As mentioned previously, achieving almost-complete powder recycling can reduce the material wastage in L-PBF to near-zero, and also significantly reduce the environmental impact of the process. In recent years, much work has been dedicated to determining the effects of powder recycling on powder particle properties. Gruber et al. [38] found that the spherical morphology of Inconel 718 powder particles was not noticeably altered after twenty recycles. In fact, the authors actually noted that the flowability of the powder improved after a number of recycles, as the finest particles found in unused powder were removed through repeated builds. Heiden et al. [39], however, conducting similar tests with 316L stainless steel powder, found opposing and less promising results: only minor changes to particle size distribution after thirty recycles, but significant changes in particle morphology that caused lower flowability, and also increased oxidation on particle surfaces.

The inclusion of multiple lasers in recent L-PBF machine designs (as shown in Figure 2.6) has begun introducing a range of potential benefits. The most obvious of these is an increased build rate. Tsai et al. [40] found that a three-laser system could reduce the amount of time needed to scan a single layer by over 60%, compared to a one-laser system. The same authors also found that, by re-scanning the outer contours of components with a following laser, the cooling profile of the metal could be altered in a way that resulted in a reduction in average surface roughness by up to 50%, thereby significantly reducing post-processing requirements. Similarly, Wong et al. [41] demonstrated an increase in build rate of 274% when comparing a four-laser system to a one-laser system, without any compromise on the part's mechanical properties.

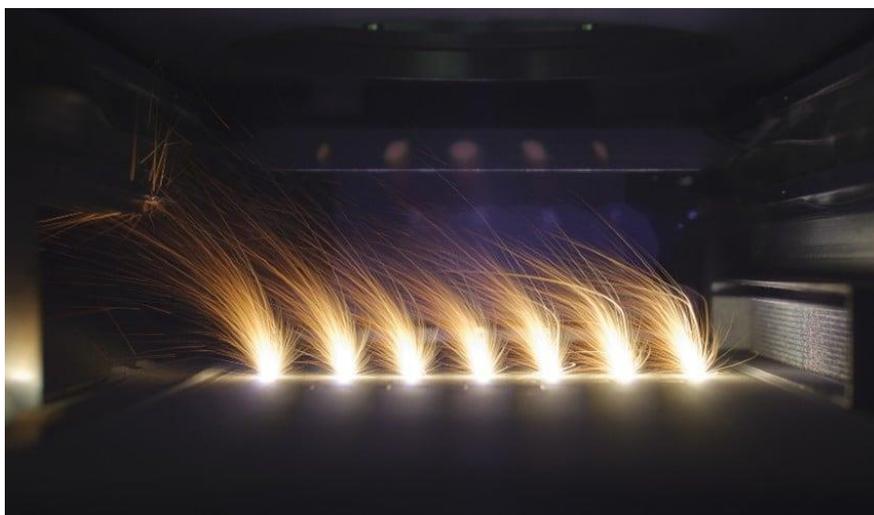


Figure 2.6

Seven independently-operating lasers, inside an Aconity3D L-PBF build chamber [42]

Another major development that is likely to soon significantly improve L-PBF processes, and which will be discussed in greater detail later in this chapter, is the introduction of Machine Learning. The high number of interdependent factors and machine parameters that affect the quality of a part, can be overwhelming for a human operator to understand. Machine Learning (ML) has the potential to interpret vast datasets of experimental results, compiled from researchers and workshops all across the world, and identify trends and cause-effect relationships that may allow for greater part quality control [43].

2.1.5 - Limitations of L-PBF

Despite the current and future promises of L-PBF, multiple, significant areas of concern continue to prevent the method from achieving widespread industry adoption. Fundamentally, the majority of the issues stem from the fact that the thermal cycling rates inside a L-PBF machine are orders of magnitude greater than for any comparable traditional method. Whereas a metal part being forged or cast may undergo cooling in the range of $10^0 - 10^1$ K/s, individual powder particles in a L-PBF build chamber may melt and solidify multiple times per second, with cooling rates above 10^6 K/s [44], as demonstrated in Figure 2.7. These extreme conditions, plus the difficulty in imparting the exact right amount of energy into a melt pool to make it a reasonable size and temperature, can often result in numerous defects existing in the as-built part. These defects will be discussed in detail later in the chapter.

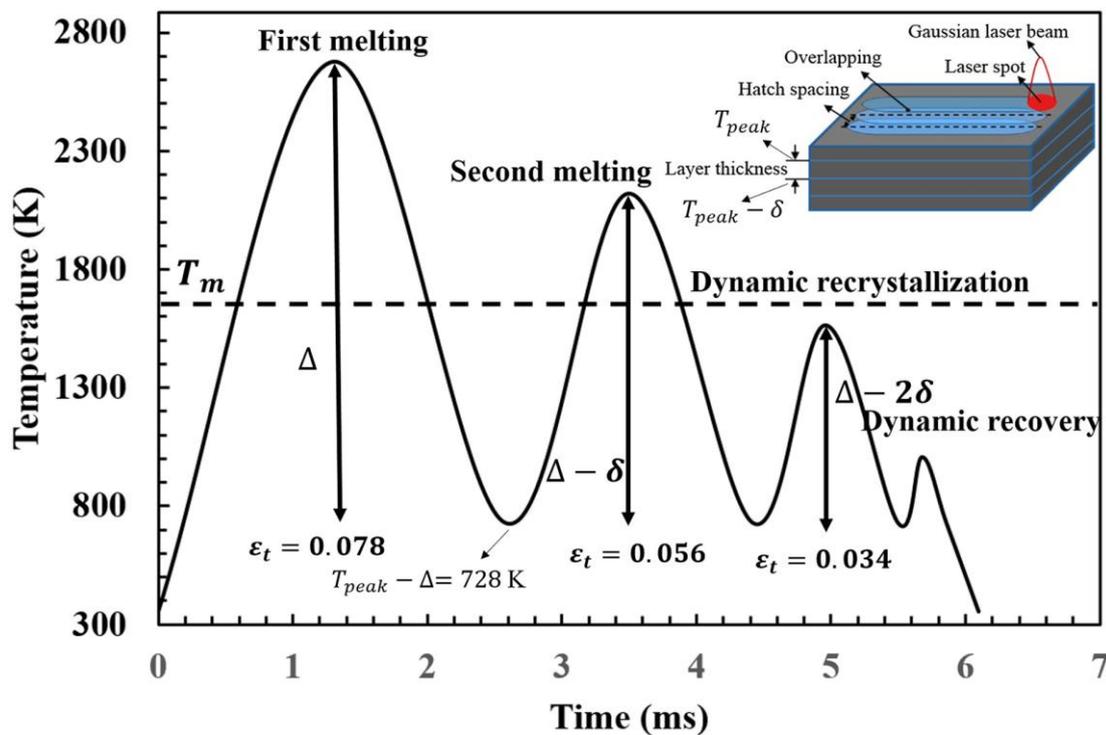


Figure 2.7

Sabzi et al. used analytical modelling to show how rapidly material oscillates above and below the melting temperature, and the thermal strain (ϵ_t) that results, when building in 316L steel [45]

In Section 2.1.3, a few bespoke examples were listed of L-PBF components being used in emerging and cutting-edge applications in the medical, automotive, and aerospace industries. Despite these, however, there exists a general consensus amongst experts in these fields that metal AM parts still do not, in general, meet their strict safety requirements [46].

Overall Equipment Effectiveness (OEE) is a standardised benchmark that calculates a number based on a combination of: the percentage of passable components a factory produces; the percentage of the maximum build rate achieved; and the percentage of machine uptime achieved. In high-tech manufacturing industries, an OEE of above 85% is the norm. In factories and workshops using L-PBF processes, however, OEE rests around 30% [47].

The repeatability and reproducibility of L-PBF processes is also a significant concern. Slotwinski et al. [48] recruited five laboratories to build cobalt-chrome tensile test pieces on identical L-PBF machines, and studied the elongation at failure of each. The standard deviation of the measurements was found to be 51% of the mean. Such variability demonstrates the difficulty of adopting metal AM to applications where extreme predictability is required.

2.2 - Melt Pool Physics

This section provides a concise overview of the complex physical phenomena occurring inside the most important unit of any L-PBF process: the melt pool. Processes will be discussed chronologically, beginning with the incidence of laser photons on powder particles, progressing through to powder melting, and finishing with solidification.

2.2.1 - Laser-Matter Interaction

When an electromagnetic wave strikes a particle, several processes result from the interaction. For the purposes of understanding L-PBF physics, there are three that are of interest. Firstly, some portion of the energy contained in the beam will be absorbed by the particle, and converted into other forms of energy, including thermal energy. Secondly, some portion of the wave's energy will be reflected, and continue propagating through space at the same wavelength as before the interaction. Thirdly, some portion of the thermal energy absorbed by the particle will be emitted as further, separate electromagnetic waves, that propagate outwards in all directions and at all wavelengths, with their wavelength distribution being dependent upon the particle's temperature [49].

The mechanisms through which electromagnetic waves are absorbed by particles and converted into thermal energy are manifold. Depending on the frequency of the light and the composition of the particle, energy may be absorbed via electron orbitals, nuclear vibration, rotations of molecular dipoles, and more. In the case of metallic particles, the majority of absorption occurs within the sea of free electrons, which are excited into higher energy states by the photon-electron interaction [50]. A detailed discussion of this process is beyond the scope of this work.

The ratio in which photons are absorbed and reflected by the matter they interact with largely depends on the morphology of the surface they strike. In terms of L-PBF, powder particles are considered ‘grey bodies’, due to them both absorbing and reflecting photons. The absorptivity of a powder layer is greater still than that of a single particle, due to the fact that photons reflected by one particle can then be absorbed by another, deeper within the layer [51]. Once enough thermal energy has been absorbed by the powder and/or bulk solid metal, and melting occurs, the new liquid surface will generally show an even greater absorptivity value [52].

Any attempt at a mathematical understanding of the L-PBF process requires a robust estimation of a powder’s average absorptivity as a starting point. Such an estimation, however, can be extremely difficult. Ye et al. [53] used calorimetry measurements to determine the effective absorptivity (A_e) of Ti-6Al-4V powder, where A_e is defined as the energy required for the observed temperature increase in the powder, divided by the energy input from the laser. It was found that the laser power (varied between 25W and 300W), laser speed (varied between 500 mm/s and 1500 mm/s), beam diameter (varied between 57 μm and 107 μm), and powder layer thickness (varied between 50 μm and 100 μm), each significantly affected the measured absorptivity, with A_e values ranging from as low as 0.3 to as high as 0.7, depending on the parameter combinations. These results are shown in Figure 2.8.

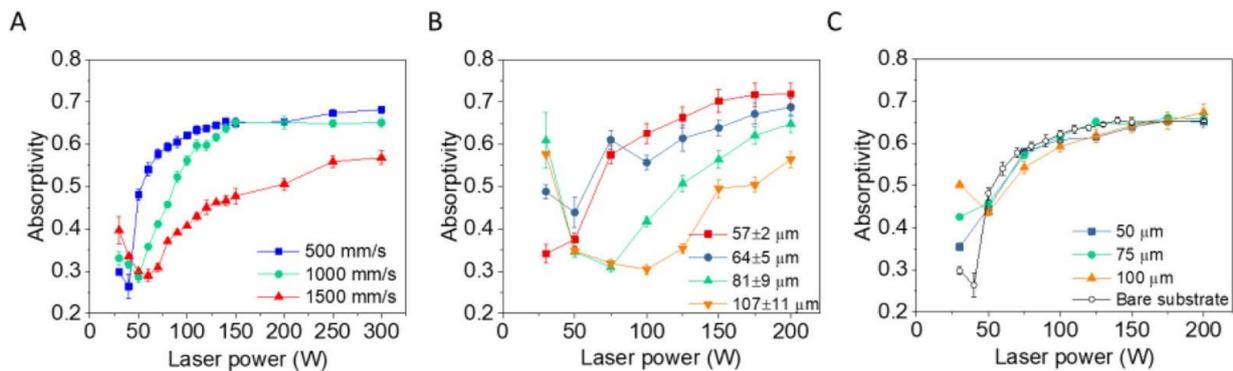


Figure 2.8

Ye showed how absorptivity in Ti-6Al-4V varies significantly with laser power, and also with laser speed (A), beam diameter (B), and powder layer thickness (C) [53]

Trapp et al. [54], using similar calorimetric methods for 316L stainless steel samples, also showed that absorptivity depends not only on feedstock, but on machine and laser settings. Depending on their setup parameters, A_e was found to vary from 0.3 to 0.8. They demonstrated, however, that the lower A_e values strongly correspond to lower laser energy input values, where melt pools are in conduction mode. As laser energy increases, and the melt pool transitions to keyholing mode, A_e values spike rapidly towards a steady state. A detailed discussion of these melting modes will follow in Section 2.2.2.

2.2.2 - Melt Pool Evolution

As powder particles absorb laser energy, and their temperature rises, a phase change from solid to liquid will occur. A semi-circular pool of molten metal will exist beneath and around the spot of laser incidence, with a depth less than or equal to its half-width. A large portion of the laser light that is reflected rather than absorbed, will be directed elsewhere in the build chamber. Conduction is the primary method by which heat is transferred away from the laser spot, into the surrounding liquid and solids [55]. Therefore, this type of melting regime is known as Conduction Mode.

Should the laser impart sufficient energy to raise the liquid temperature much further, a phase change from liquid to gas may occur. When metallic molecules are vaporised from the liquid surface, the matter undergoes a rapid volumetric expansion, thereby exerting a high pressure upon the surface it left behind, known as the recoil pressure. This process of laser-induced material removal is referred to as ablation [56]. For continuous lasers, the recoil pressure experienced by the melt pool surface is in the range of Megapascals [57]. For ultrashort pulsed laser, the recoil pressure can measure Gigapascals [58].

As the recoil pressure on the melt pool surface increases, a vapour depression appears, both deepening the pool and increasing the ratio between its depth and width, known as the aspect ratio. The dominant method of heat transfer outwards from the laser spot begins to change from conduction to convection [59]. As the pool deepens, the portion of reflected light that is directed elsewhere inside the build chamber will decrease, and the portion that is instead absorbed elsewhere inside the vapour depression increases. This phenomenon is called beam trapping. The depth of the melt pool will fluctuate stochastically, in response to the stochastic fluctuation of these portions of completely-reflected and reabsorbed photons. This melting regime is known as the Transition Mode.

As the laser energy increases still, and the vapour depression deepens into a thin, vertical tunnel, a stage is reached at which almost all of the reflected laser light is reabsorbed within the melt pool. In this melting regime, known as Keyhole Mode due to the melt pool's extreme aspect ratio, convection is the dominant method of heat transfer [60]. There exists a steady state, at which the machine parameters and material properties allow a melt pool in keyhole mode to maintain a consistent depth and volume [61]. Further increases in laser energy, however, will continue to deepen the melt pool and the vapour depression within it until the keyhole collapses, leaving behind a keyhole pore. This process will be explored in more detail later on in the chapter. A diagram of Conduction, Transition, and Keyhole Mode melt pools is presented in Figure 2.9.

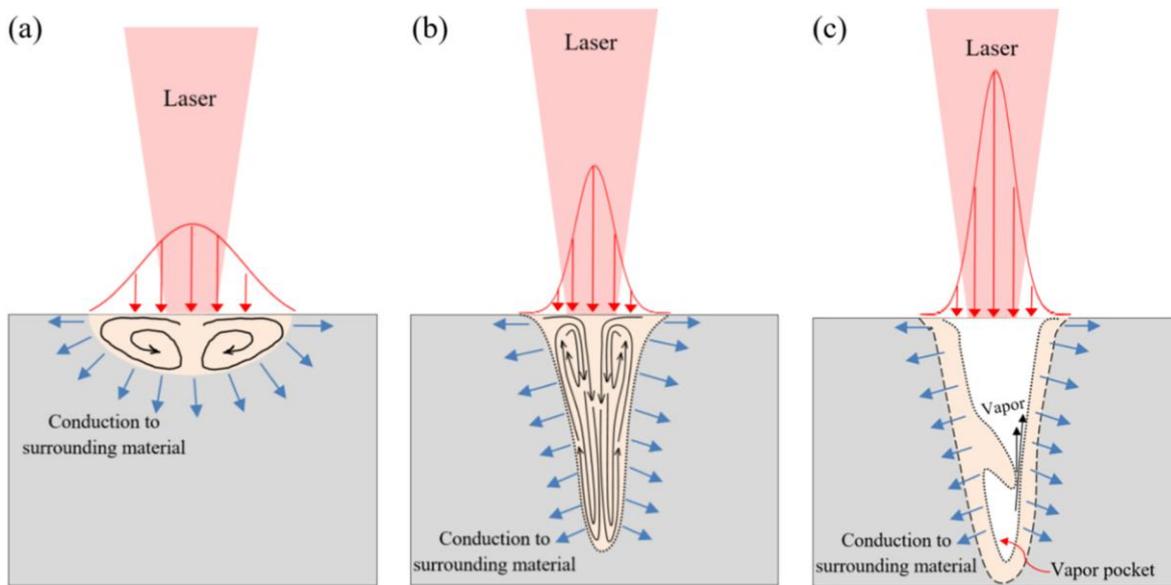


Figure 2.9

Conduction mode (a), Transition mode (b), and Keyhole mode (c) melt pools [62]

As demonstrated in Figure 2.10, four principle forces exist within a melt pool: buoyancy forces, resulting from spatial variations in temperatures; gravitational forces, which are constant throughout the volume; surface tension across the liquid-gas boundary; and the recoil pressure. The interplay of gravitational and buoyancy forces result in free convective flows within the melt pool, which have a negligible effect on melt pool dynamics [63]. Variations in surface tension across the melt pool surface, known as the Marangoni force, result in a forced convective flow, which has enormous influence on the pool's size and shape.

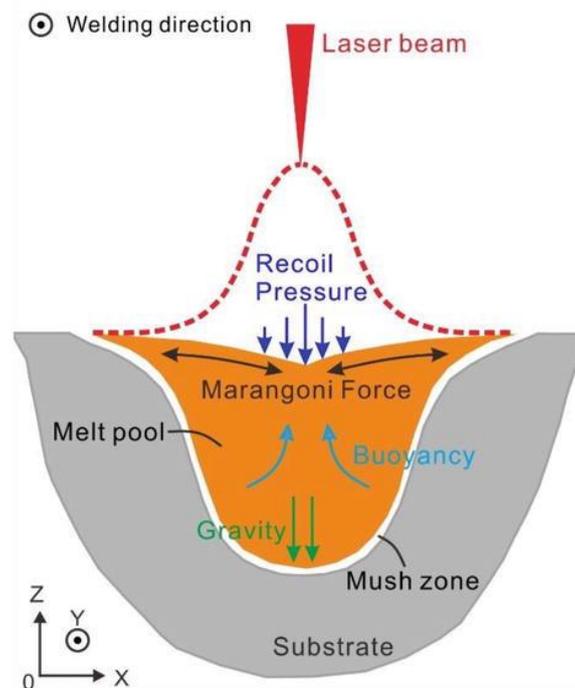


Figure 2.10

The four primary forces inside a melt pool: Recoil, Marangoni, Buoyancy, and Gravity [63]

In cases where surface tension forces decrease with temperature, liquid will flow outwards from the laser spot, across the melt pool's surface towards its edges. This surface flow creates an outward circulation within the pool, with mass rising in its centre and falling along its outer edges, thereby resulting in a wider, flatter pool morphology. In cases where surface tension increases with temperature, the reverse is true: liquid flowing inwards across the centre from the edges to the laser spot; an inward circulation forming within the pool; and narrower, deeper pool morphology [64]. The difference between these cases is demonstrated in Figure 2.11. It should be noted that in the vast majority of cases, surface tension decreases with temperature. Some exceptions to this rule have been found, however, namely in stainless steels treated with: high sulphur additions [65]; oxide fluxes [66]; and oxygen inclusions [67].

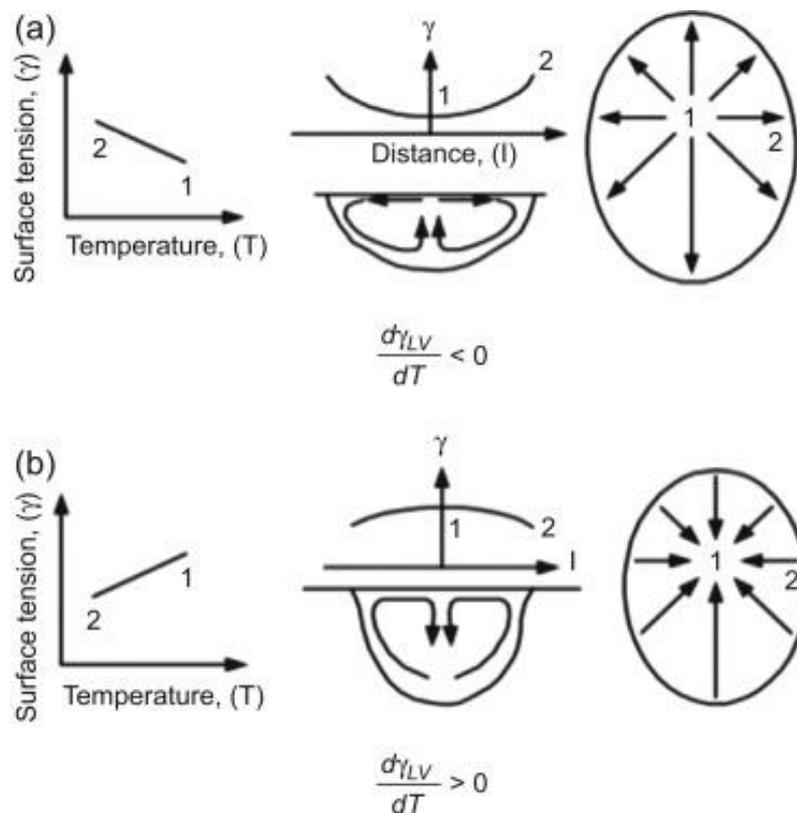


Figure 2.11

When surface tension decreases with rising temperature (a), liquid flows from the melt pool centre to its edges, resulting in wide, shallow pools. When surface tension increases with temperature, the opposite is true. [68]

The hottest part of the melt pool lies beneath the laser spot. The exact rate at which temperature decreases along any directional line extending from the peak temperature, depends on a multitude of machine and powder properties. Some factors which play significant roles in determining the temperature gradients and isotherm layouts inside a melt pool include: the position, size, and morphology of powder particles along the laser scanning path; the smoothness of the underlying bulk metal; the extent to which the melt pool overlaps with solidified metal from the previous parallel track; and the temperatures of the bulk metal and powder that surrounds the melt pool [69].

Lane et al. [70] used thermal imaging to measure the variation in isotherm width in a melt pool building cubes of Inconel 625. Temperatures along the melt pool lengths were measured; i.e. along the line of the laser direction of travel. It was found that, depending on the position of the melt pool within the scan track (and therefore the temperatures surrounding the melt pool), the amount of this line with temperatures above 1000°C may be as little as 0.5mm, or as much as 2.5mm. An example of Lane’s results is presented in Figure 2.12.

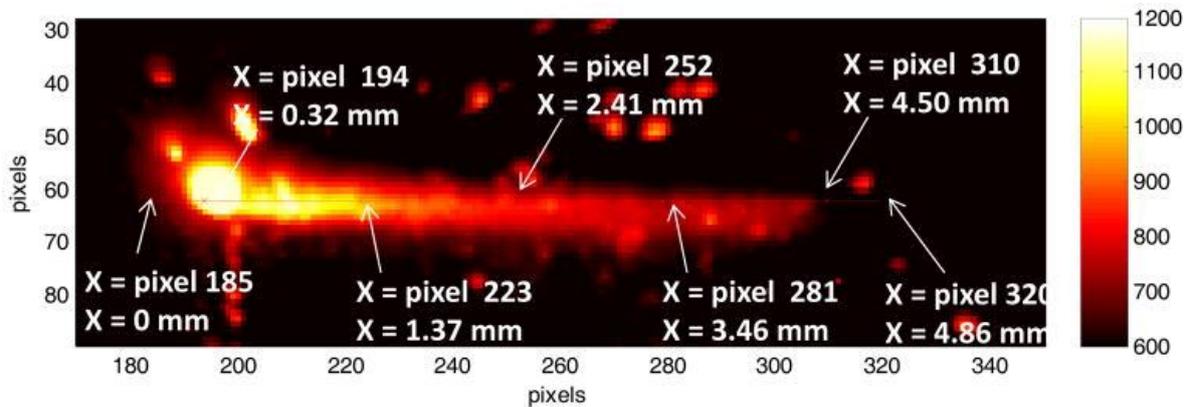


Figure 2.12

Lane tracked melt pool dimensions, and found huge variability in the length of line (X) above 1000 °C [70]

Hooper [71] similarly used thermal imaging cameras to determine temperature gradients and cooling rates within a Ti-6Al-4V test piece. It was found that maximum temperature gradients of approx. 20 K/μm existed just outside the laser spot, when the melt pool was scanning the piece’s outer perimeter. Further out from the laser spot, closer to the boundary of the liquid pool, the temperature gradient drops to a minimum of 5 K/μm, as the relatively-steady temperatures of the mushy zone is approached.

Furumoto et al. [72] used thermal imaging to investigate the extent to which temperature gradients within a melt pool are affected by the temperatures of the powder particles being assimilated into it. Single tracks of 18Ni maraging steel powder were scanned, and the temperatures of particles that lay in the melt pool’s path were recorded, in the moments before the laser spot crossed over them. It was found that the extent to which powder in front laser spot was pre-heated by conduction, greatly influenced the temperature gradient at the front of the melt pool.

As well as temperature gradients within the above-surface melt pool and adjacent powder and bulk metal, sub-surface temperature gradients are found within the underlying bulk. In conduction mode, the difference between above-surface and sub-surface temperature gradients within a melt pool are found to be inconsequential. In keyholing mode, however, sub-surface gradients are far greater than above-surface [73].

Crales et al. [74] showed the influence of the temperature of the surrounding material on a melt pool’s size and shape. At the end of a track, when the solidified metal from the adjacent track has been given the maximum amount of time to cool, the thermal conditions inside a melt pool are very

different from those at the beginning of a track, immediately after the laser spot changes direction, when the adjacent metal has been given the minimum amount of time to cool. The thermal conditions are so different, in fact, that the authors classified the two scenarios as Type I (at the start of a track) and Type II (at the end of a track). These melt pool types are shown in Figure 2.13.

In their experiments with Inconel 625, the authors found that, due to the increase of thermal conductivity with temperature, Type I melt pools were wider than Type II, and skewed heavily in shape towards the adjacent track. i.e. at the beginning of a track, the melt pool will extend into the hot metal beside it; whereas at the end of a track, no such extension into the cooled metal is found.

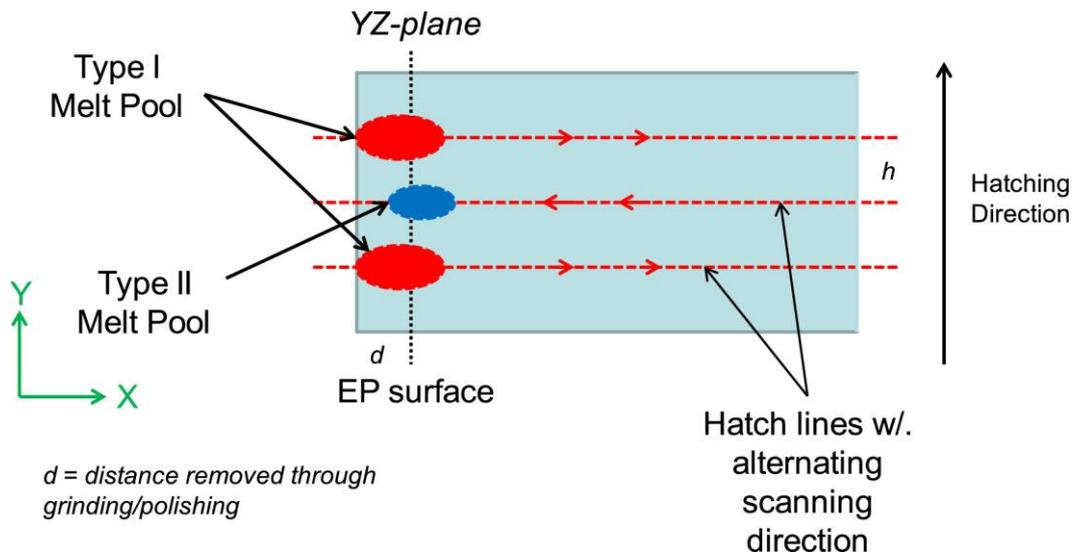


Figure 2.13

Criales defined Type I melt pools as those having just started a track (and therefore surrounded by hot material), and Type II melt pools as those ending a track (and therefore surrounded by cool material) [74]

In analysing the effects of laser power and speed on melt pool width, depth, and cross-sectional area in Inconel 718 samples, Scime et al. [75] also demonstrated the high variability of these metrics. Across a parameter space of 200mm/s to 1200mm/s, and 100W to 370W, melt pool widths were found to vary from 100 μm to 200 μm , melt pool depths varied from 50 μm to 300 μm , and melt pool areas varied from 0.005 mm^2 to 0.04 mm^2 . The standard deviations of these measurements, however, were found to range from 10% of the mean values, up to 35% of the mean, highlighting the extent to which melt pool morphology is a stochastic phenomenon.

2.2.3 - Solidification

Ultimately, across metallurgical manufacturing processes spanning from casting, to laser welding, to metal AM, it is the manner in which material is returned from the liquid to solid phase that determines the majority of a part's mechanical properties. A detailed discussion of the ways in which cooling rates and thermal gradients can result in different cells and grains depending on exact alloy types, will be presented in Section 2.3.1. The purpose of this current section is to provide an overview of the general solidification mechanics in L-PBF.

Just as powder particles to the fore of a melt pool are being rapidly heated, liquid at its rear is being cooled at rates unheard of in any other metallurgical process. In casting and forging processes, normal cooling rates range from 10^0 to 10^2 K/s [76, 77]. Cooling rates in the region of 10^3 and 10^5 K/s have been reported for Direct Energy Deposition, and laser welding processes, respectively [78, 79]. In L-PBF, cooling rates range from 10^6 to 10^7 K/s, with rates varying significantly depending on the precise location within the melt pool's tail [80].

Bertoli et al. [80] used a high-speed camera to record the scanning of single line tracks in 316L stainless steel. An inverse relationship was found to exist between the cooling rate, and the linear energy density (power/speed). This finding confirms the logical assumption that when a melt pool is scanning across a powder bed with greater velocity, the cooling rate will be greater too.

The powder packing density of nominally spherical particles used in L-PBF (i.e. the density of the powder compared to an equivalent bulk alloy) is in the range of 60% to 80% [81]. As a result, the volume of material before and after the melting and re-solidification is significantly different. Additionally, many alloys used in L-PBF undergo significant thermal expansion. The linear thermal expansion coefficient of solidified Inconel-718 is in the range of 1.0 to 2.0×10^{-5} K⁻¹ [82], that of 316 stainless steel is in the range of 1.5 to 2.0×10^{-5} K⁻¹ [83], and that of AlSi10Mg is in the range of 2.0 to 2.1×10^{-5} K⁻¹ [84]. Assuming a melt pool cross-section dimension of 100 μm for simplicity, being cooled by 1500 K, the aforementioned alloys could undergo a thermal contraction of up to 0.3 μm : possibly a sufficient for solidification cracking to occur, a topic which will be discussed in more detail in Section 2.4.4.

2.2.4 - Modelling

Computer-based simulations, designed to predict the melting and resolidification of metal heated by lasers, have been advancing alongside experimental L-PBF since the techniques inception. A thorough examination of such modelling techniques is far beyond the scope of this work. Instead, a brief overview of the various categories of modelling methods will be presented, roughly ranked in order of increasing complexity and accuracy.

2.2.4.1 - Analytical Models

Straightforward equations, which often assume uniform material properties across all temperatures, can be a useful first step in predicting the extent to which a moving laser beam can heat a substrate. Freeman [85] explores the use of several analytical models in predicting the widths and depths of melt pools in 17-4PH stainless steel. Most relevant to this project were the author's employment of the Rosenthal equation and the Eagar equation.

The Rosenthal equation assumes a steady-state, point heat source, moving across the surface of a solid plane, which experiences no phase or property changes regardless of temperature. The temperature of a point at X (direction of laser travel), Y, Z is calculated as:

$$T_{XYZ} = T_0 + \frac{AP}{2\pi\lambda} \exp\left(\frac{-V(X + \sqrt{X^2 + Y^2 + Z^2})}{2\alpha\sqrt{X^2 + Y^2 + Z^2}}\right)$$

where T_0 is the solid's ambient temperature [K], A is absorptivity, P is laser power [W], λ is thermal conductivity [J/s/m/K], V is laser speed [m/s], α is thermal diffusivity [m²/s], and X, Y, and Z are positional coordinates [m].

When experimentally validating the Rosenthal equation, Freeman found that the model slightly over-predicted melt pool dimensions at lower energy densities, during conduction mode melting. The dimensions were in the correct order of magnitude, and followed the correct trend. At higher energy densities, however, during keyhole mode melting, the model's inability to account for fluid flow or vapour depression, causes it to massively under-predict depth.

The Eagar equation advances the Rosenthal equation, by both implementing a Gaussian beam rather than a point heat source, and by introducing the time dimension. Like the Rosenthal equation, the Eagar equation assumes a solid plane with uniform properties across temperature. The model calculates the temperature of a point at X, Y, Z, and at a target time t_{max} as:

$$T_{XYZt_{max}} = T_0 + \int_{t=0}^{t=t_{max}} \frac{AP}{\pi\rho c\sqrt{4\pi\lambda}} \times \frac{t^{-1/2}}{2\lambda t + r^2} \times \exp\left(\frac{X^2 + Y^2 + 2XVt + v^2t^2}{4\lambda t + 2r^2} - \frac{Z^2}{4\lambda t}\right) dt$$

where T_0 is the ambient temperature [K], A is absorptivity, P is laser power [W], ρ is density [kg/m³], c is specific heat capacity [J/kg/K], λ is thermal conductivity [J/s/m/K], t is time [s], r is beam radius [m], V is laser speed [m/s], α is thermal diffusivity [m²/s], and X, Y, and Z are positional coordinates [m].

Compared to the Rosenthal results, the Eagar equation more accurately predicted the depths of melt pools in conduction mode. After transitioning to keyholing mode, however, the depths were again far greater than the model's prediction. Freeman's results are shown in Figure 2.14.

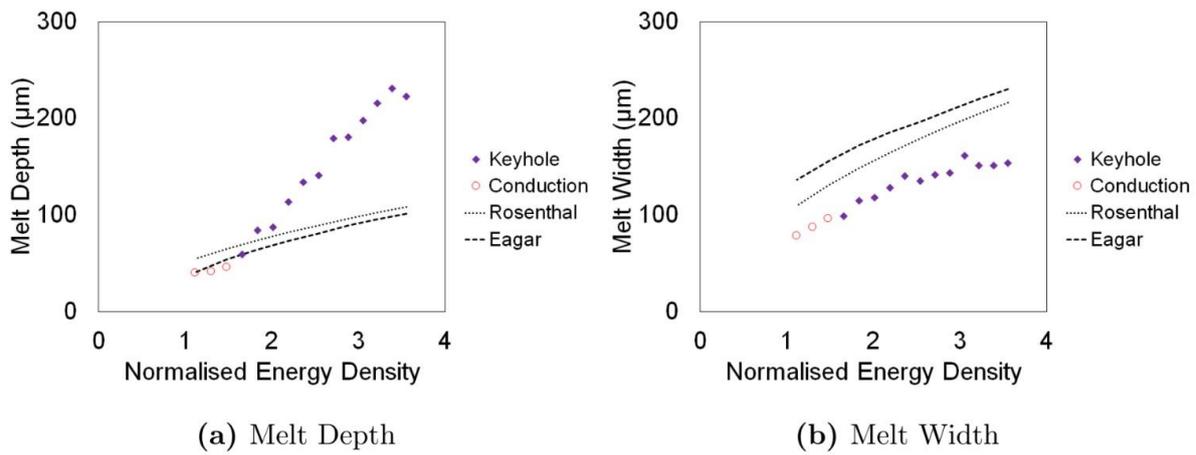


Figure 2.14

Both the Rosenthal and Eagar equation predict melt pool width relatively well, when compared to the experimental data (circles). Melt pool depth predictions are accurate for conduction mode, but greatly under-predict for keyhole mode. [85]

2.2.4.2 - Finite Element (FE) Models

Finite Element (FE) modelling techniques incorporate meshes composed of 1D points, 2D interfaces, and 3D volume elements. Each volume element is intended to represent a unit of the material being simulated, with various thermomechanical properties embedded within. By applying heat, stress, or a similar stimulus to one location within the mesh, the interactions between all the volume elements can be solved for over time, thereby predicting the heat, stress, or similar anywhere. An example of heat applied to such volume elements is shown in Figure 2.15.

FE models solve for heat and force transfer, but not for mass flow. As a result, convection currents, Marangoni flows, vapour depressions, and other such crucial phenomena are lost in the method.

Sarkar et al. [86] reviewed the current state of FE modelling in L-PBF, with an aim to assess the accuracy of various techniques, and the limitations that may prevent high fidelity to experimental measurements. The authors found that FE techniques are highly accurate at predicting strain and distortion at the length scale of the component, but suffer from the same issues as the analytical methods discussed previously, when it comes to predicting melt pool dimensions.

The authors point out that many FE models assume isotropy, limiting their accuracy in predicting certain types of failure or defects in alloys prone to directional solidification. The current processing power of computers still forces a trade-off between the high fidelity that requires fine meshes and accurate thermomechanical properties, and computational time. The authors point out several challenges that will need to be met by FE models in the future, such as recoater interference, sintering effects, and the complexities of multi-material builds.

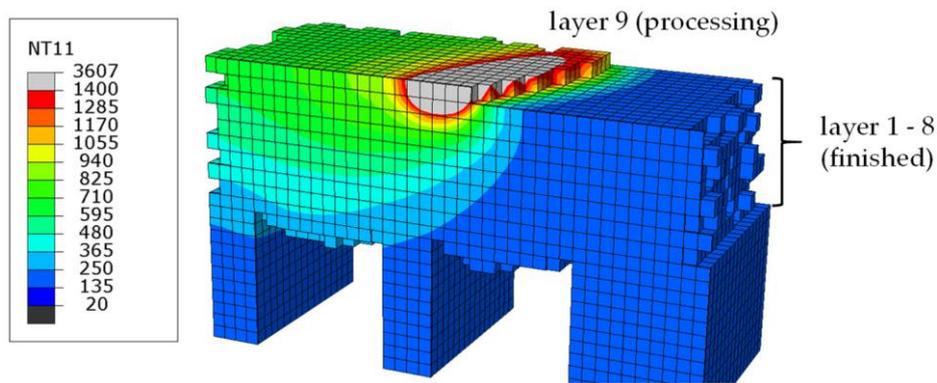


Figure 2.15

A snapshot of the FE results from Sarkar, showing the expected temperatures as the laser progresses across layer 9 [87]

2.2.4.3 - Computational Fluid Dynamics (CFD) Models

In order to simulate convection currents, Marangoni flows, vapour depressions, and other examples of mass flow, the Navier-Stokes equations must be solved and incorporated into the heat-flow model. These partial differential equations describe the motion of fluid, by relating terms like pressure, viscosity, and external forces.

Qiu et al. [88] created a CFD model designed to simulate the melting and flowing of regularly-arranged, $50\mu\text{m}$ diameter, Ti-6Al-4V powder particles, filling a $250\mu\text{m} \times 250\mu\text{m}$ bed, as shown in Figure 2.16. Their model simulated fluid flow, and the gas-liquid interface, by solving and coupling the Navier-Stokes equations with heat and energy conservation equations. The system accounted for surface tension, Marangoni flow, recoil pressure, drag force, and buoyancy forces.

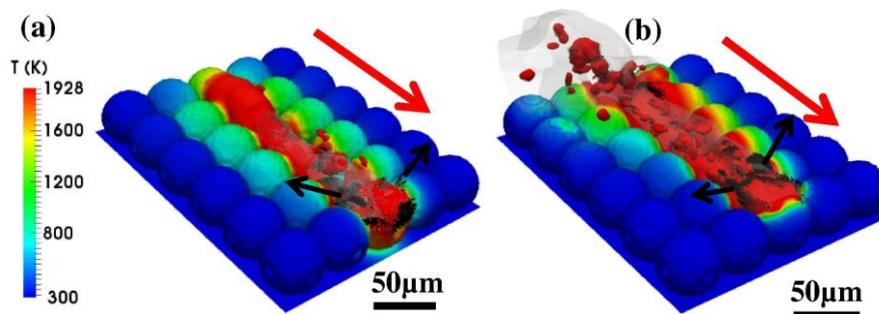


Figure 2.16

The heating effect of a 400W laser, passing over Ti-6Al-4V powder at 3500 mm/s (a) and 5000 mm/s (b), according to Qiu's CFD model [88]

The model showed strong fidelity to experimental measurements of laser-powder interactions and fluid motion, as provided via high-speed in-situ imaging, and ex-situ Scanning Electron Microscopy (SEM). The model did not account for laser reflection within the powder, or gas expansion due to heating.

The fact that such a small area was scanned in this simulation, points to the fundamental issue currently plaguing CFD: the extreme complexity, and high computational costs, that currently make the technique incapable of modelling component-scale systems.

2.2.4.4 - Microstructural Models

At the smallest scale, some modelling techniques attempt to predict the microstructure that will result from solidification in L-PBF. These systems can simulate the formation grains of various phases, and the boundaries between them.

Cellular Automata (CA) models are comparable to FE models, in the sense that grids of cells with distinct phases and orientations are defined, and the interactions between adjacent cells will change these states. Chen et al. [89], using CA to simulate the distribution of grain areas and orientations in L-PBF of IN-718, found that the predicted results showed high fidelity to experimentally obtained results, imaged via Electron Backscatter Diffraction (EBSD). These results are shown in Figure 2.17.

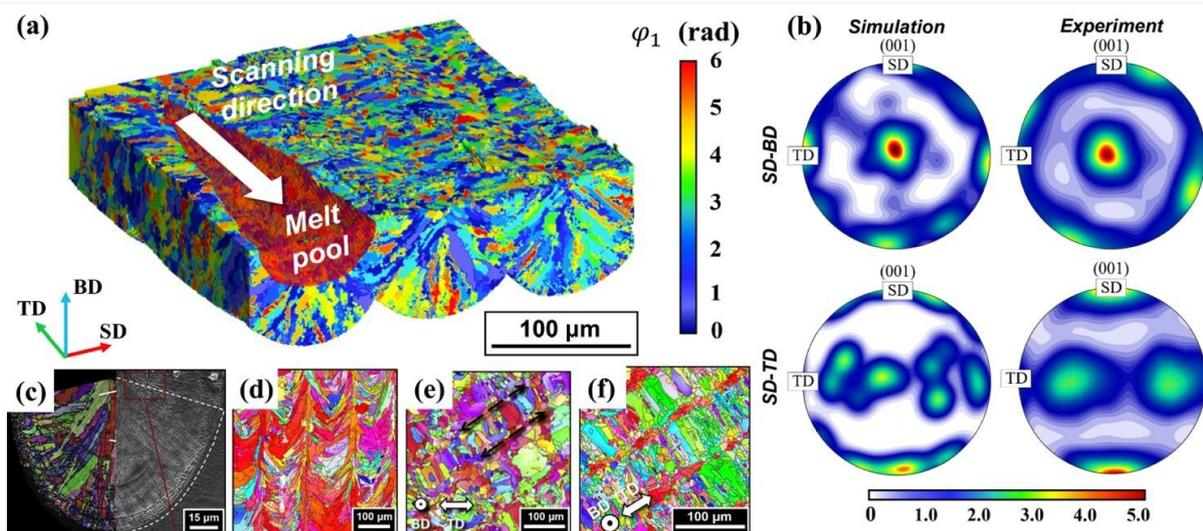


Figure 2.17

Chen's CA model: simulated grain growth across adjacent weld tracks (a); pole figures from simulations and EBSD results (b); simulated (c) and EBSD (d) grains in the XZ plane; and simulated (e) and EBSD (f) grains in the XY plane [89]

Phase Field (PF) models do not use discrete grids with of elements with distinct phases, but rather solve for a continuous interface. The boundary between different phases is mushy, with properties smoothly transitioning from one to the other. These models are extremely accurate at simulating the growth of sub-grain microstructural elements, albeit at enormous computation costs. Choi et al. [1] developed in 2024 a PF model, operating on a state-of-the-art high performance computer, that simulated the building of a316L steel cube of side length 2mm, in approximately 5 hours. This accelerated model ran significantly faster than models existing in literature at the time, and managed to reliably reproduce features seen in the experimental measurements. A snapshot of these results is shown in Figure 2.18.

Monte Carlo (MC) models use probabilistic methods to simulate grain growth, predicting phases and orientations based on energy minimisation within the solidified alloy. A mesh is defined, similarly to FE methods, and various combinations of phases and orientations are tested at each point, until a lowest-energy state is found. The method is best suited for post-process simulations, as it cannot easily simulate the solid-liquid interface. Rodgers et al. [2] used an FE model coupled with an MC model to predict the solidification of 300-series steel, built with complex scanning patterns. The coupled model managed to simulate microstructural development so accurately, that a previously unknown correlation between the number of re-melting cycles, and the mean size of columnar grains, was uncovered. An example of the microstructural results predicted by the method is shown in Figure 2.19.

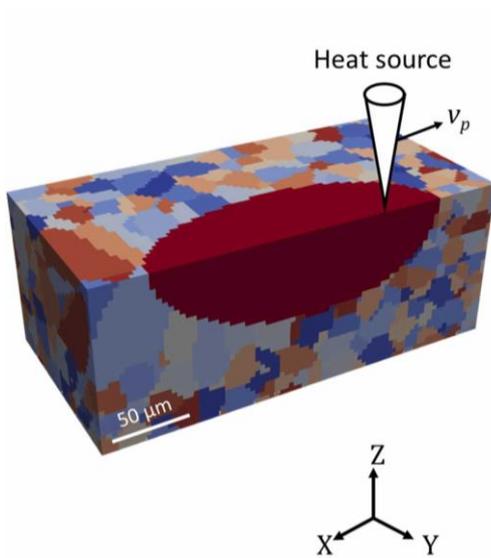


Figure 2.18
Choi's PF model, predicting grain growth in 316L steel [1]

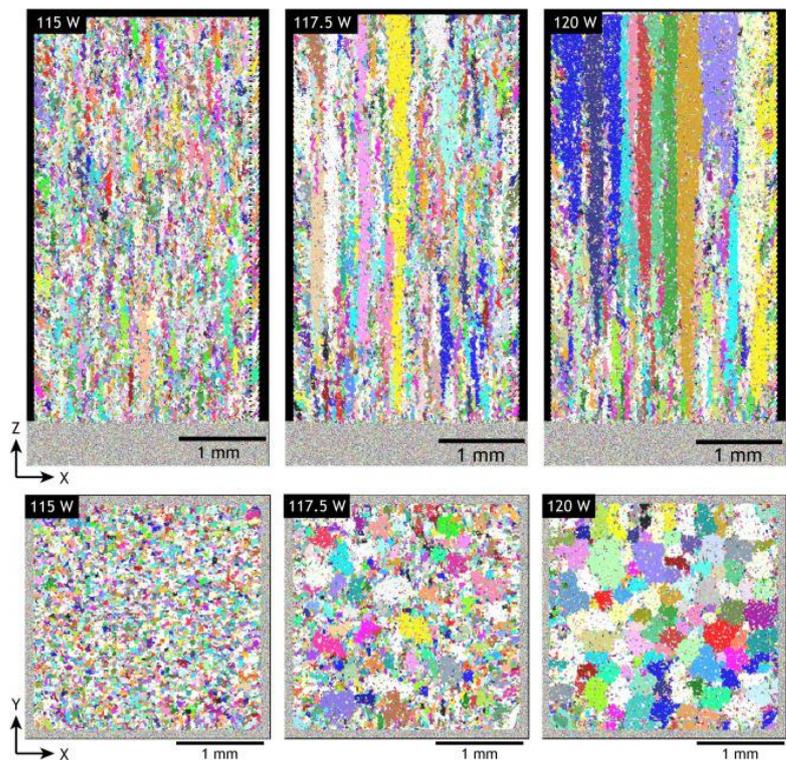


Figure 2.19
The microstructure predicted by Rodgers, at three different laser powers, in both XZ and XY [2]

2.3 - Alloys and Microstructure

This section introduces detailed discussions of the metallurgy of L-PBF: beginning with crystallographic and microstructural concepts that are applicable to any metal; then becoming more focused upon the alloys that are used in AM and L-PBF; and finally landing on the specific materials that will be experimented upon in this thesis.

2.3.1 - Microstructure Overview

As mentioned previously, it is the process of transitioning from the liquid phase to the solid phase that determines the majority of a metal's physical properties. In order to understand the strength, ductility, and hardness of an alloy, one must understand the thermodynamics of solidification; and in order to understand the thermodynamics of solidification, one must start with an understanding of Gibbs free energy.

2.3.1.1 - Thermodynamics of Solidification

Every material can, theoretically, exist in any of its phases at any temperature: gaseous, liquid, or any available form of solid. However, each phase has an associated energy state, known as the Gibbs free energy, that varies across temperature, and the material will tend to transition into the phase with the lowest energy state [90], as demonstrated in Figure 2.20. Once a solid material is heated to its melting point, its liquid phase becomes the more energetically-favourable phase to exist in. Should additional thermal energy be supplied to overcome the enthalpy of fusion barrier, the chemical bonds that maintain the solid's rigidity will break down, allowing it to flow in liquid form. Likewise, once a liquid material is cooled to its melting temperature, the solid phase becomes energetically-favourable. Should the material be given the opportunity to release additional thermal energy in the form of enthalpy of fusion, the chemical bonds reform, returning the solid's rigidity.

The form that these chemical bonds take varies based upon the material. In covalent network solids like silica, it is the interatomic/intermolecular covalent bonds that break and reform during the liquid-solid phase transitions [91]. Similarly, in ionic solids like salt, it is the ionic bonds that break and reform [92]. For solids held together by intermolecular forces, such as hydrogen bonding in ice [93], or van der Waals bonding in paraffin wax [94], it is these forces that breakdown during melting. In the case of metallic solids, where positively charged metal atoms are held in place via the electrostatic charge of a sea of delocalised electrons, melting occurs when these metallic bonds can no longer restrain the atoms, allowing them to flow freely [95].

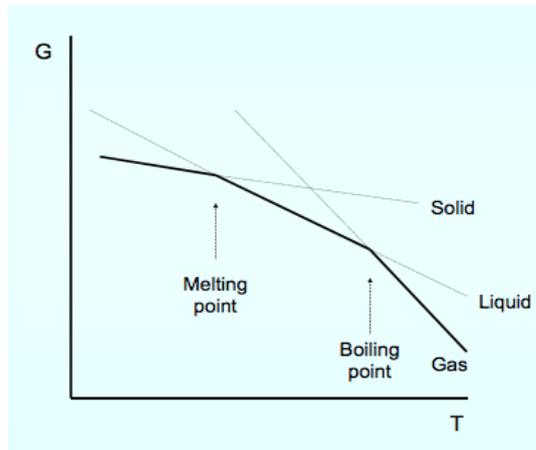


Figure 2.20

A simplified schematic, showing how the phase with the lowest level of Gibb's free energy varies across temperature [96]

More often than not, materials being produced and analysed in the field of metallurgy are alloys rather than pure metals. In such a case, no single melting point exists for the system, but rather a range of temperatures over which it is partially melted; i.e. solidified crystals will be suspended in liquid [97]. The upper limit of this range, above which the whole system is liquid, is called the liquidus temperature. The lower limit, below which the system is solid, is called the solidus temperature.

In order for a liquid alloy to fully solidify, it must transit through this region in between the liquidus and solidus temperatures, known as the mushy zone. As crystals of a certain chemical composition begin to solidify, the extraction of these atoms from the liquid bulk will alter the chemical composition of the remaining liquid [98]. This process, in which cooling spatially alters the composition of a material, is known as segregation, and will be discussed in more detail in this section. Figure 2.21 demonstrates this segregation in the form of a binary phase diagram.

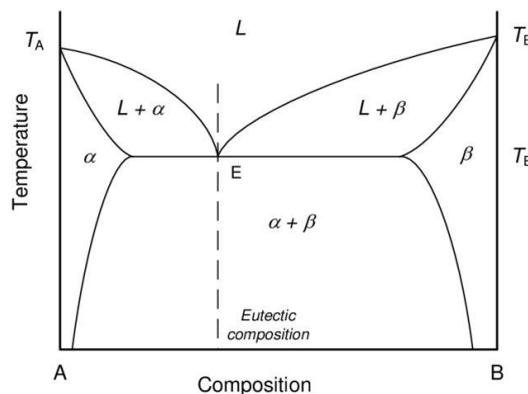


Figure 2.21

A simple phase diagram, showing the most energetically favourable phases across temperature for all compositions of elements A and B. Note how for all binary compositions apart from the Eutectic, there exists a temperature range (the mushy zone), inside which there will be liquid (L), and some solids (either α or β) [99]

2.3.1.2 - Nucleation

There are two mechanisms by which metallic atoms or molecules that exist in the liquid state can begin the initial transition to the solid state [100]. The first, which is very uncommon to observe in either nature or industry, involves the liquid atoms reforming their metallic bonds with each other; a process known as homogeneous nucleation. The second mechanism, which accounts for the vast majority of solidification in metal manufacturing, involves metal atoms solidifying onto a pre-existing solid, such as a container surface or foreign particle; a process known as heterogeneous nucleation.

The reason that heterogeneous nucleation is much more common than homogeneous, is because of competing factors in the Nucleation Free Energy equation:

$$\Delta G(r) = \frac{4}{3}\pi r^3 \Delta g_v + 4\pi r^2 \gamma$$

where $\Delta G(r)$ is the total change in Gibbs free energy [J] for liquid to nucleate into a solid sphere with radius r [m], Δg_v is the change in Gibbs free energy for a unit volume to solidify [J/m^3], and γ is the interfacial energy between the solid nucleus and the liquid [J/m^2].

Energy is released when liquid solidifies, so therefore Δg_v is negative (favourable to nucleation). Energy is required for the interface to exist, so therefore γ is positive (unfavourable to nucleation). Because volume is a product of r^3 , and surface area is a product of r^2 , their combination results in an initial positive spike, followed by a negative drop, as shown in Figure 2.22.

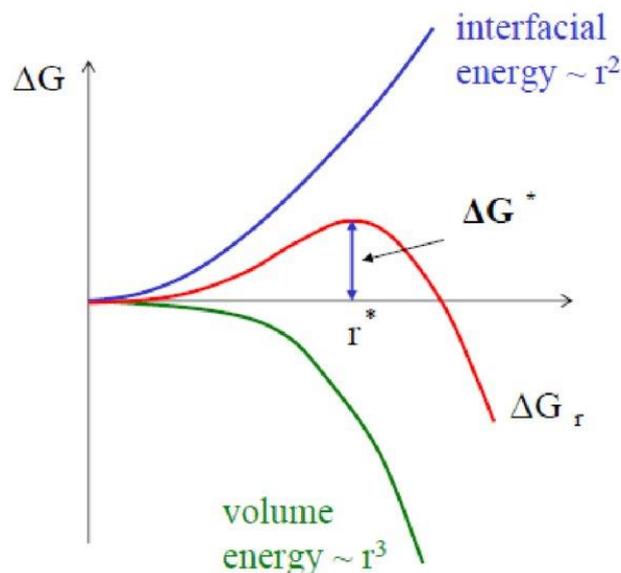


Figure 2.22

The total energy change of nucleation (red line), is a combination of the favourable volumetric term (green line), and the unfavourable surface term (blue line) [101]

It is therefore found that, at small nucleus sizes, it is energetically favourable for solidified particles to return to the liquid phase. It is only once a critical radius has been reached, that it becomes energetically favourable for the nucleus to grow, and for new particles to solidify onto it [102]. In the case of heterogeneous nucleation, the foreign particle or container boundary provides a surface, which reduces the critical radius by changing the energy balance of nucleation.

Due to the fact that homogeneous nucleation has an energy barrier, due to the critical nucleus size, it is found that metal atoms in a liquid state can be cooled to temperature below their melting point, yet still remain liquid. This process is called undercooling, and is a significant driving force of solidification, since atoms below their melting temperature are very likely to heterogeneously nucleate or bond onto pre-solidified material [103]. The extent of undercooling dictates how energetically favourable the liquid-solid phase transformation will be, and therefore the ease with which the metal solidifies.

Undercooling can present in two forms. Thermal undercooling refers to the temperature of liquid metal falling below its melting temperature, and is caused by rapid cooling or a lack of nucleation sites [104]. Constitutional undercooling refers to the composition of the liquid immediately in front of a solid-liquid interface being altered via segregation, in a manner which reduces the liquidus temperature of this new composition. This complex process will be discussed in more detail in this section.

It is possible to take advantage of the fact that heterogeneous nucleation is likely to occur on foreign particles, and introduce such particles into a liquid metal to alter the final, solidified makeup of the metal. This process is known as inoculation, and can take many forms [105].

Once a stable nucleus has formed, whether by homogeneous or heterogeneous nucleation, with or without an inoculant, further solidification onto the nucleus can occur. As metals are crystalline in their solid state, they require specific atomic arrangements to form. Therefore, the rate and morphology with which a nucleus or solid-liquid interface can grow, depends on the ability of liquid atoms to find a suitable location at the crystal interface [106].

In an alloy, where multiple metals and impurities are mixed together, the metal with the highest concentration is the solvent. The rest of the chemical contributions are called the solute. The solute to solvent ratio is the solute concentration. For any given chemical mixture, at any given temperature, assuming equilibrium (i.e. all atoms have been given infinite time to diffuse into their preferred positions) there will be two different solute concentrations: the liquid concentration, and the solid concentration. The liquid concentration refers to the overall solute concentration of the entire mixture, while the solid concentration refers to the quantity of solute which the solid solvent crystals will allow to be dissolved within their crystals. Any solute particles that exceed this concentration limit will be ejected from the solid (again, assuming equilibrium) [107].

The ratio of solid concentration to liquid concentration, is called the partition coefficient:

$$k = \frac{C_s}{C_l}$$

where C_s is the solid concentration, and C_l is the liquid concentration.

When applying the partition coefficient to the solidification of alloys specifically, it is useful to consider these concentrations at the exact solid-liquid interface. In this context, when $k < 1$ (as is the case in most circumstances), solute will be rejected into the liquid as the interface grows.

2.3.1.3 - Growth Mechanisms

Imagine a flat interface, uniformly growing as liquid atoms bond to it at a constant rate across its span. This type of solidification is called planar growth [108]. As the liquid transitions to solvent crystals, and the excess solute that cannot be supported inside these crystals remains in the liquid (the process of segregation), the solute concentration of this liquid immediately in front of the interface rises. As is visible from many phase diagrams, when the solute concentration of an alloy increases, the liquidus often temperature falls, requiring even lower temperatures for solidification to occur. Therefore, this solute-rich layer of liquid that exists in front of a planar interface can slow its growth.

However, slight perturbations will exist in the planar interface, caused by random temperature fluctuations, and the stochastic movement of atoms as they bond to the growing solid. It is almost inevitable, for most metallurgical processes, that these perturbations will extend slightly into the liquid, out into the region away from the interface where there is a lower solute concentration, and therefore a higher liquidus temperature, where solidification is more favourable. This is the mechanism of constitutional undercooling [109]. This mechanism is shown in Figure 2.23.

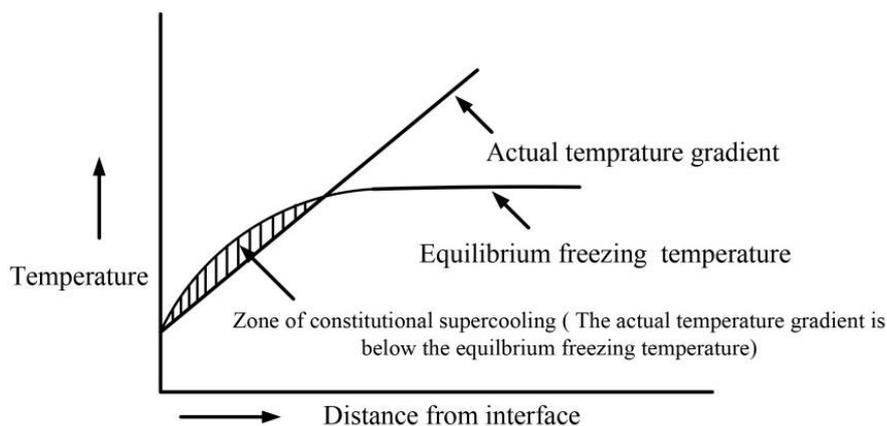


Figure 2.23

Constitutional undercooling: as distance from the solidification interface increases, both the liquid temperature and the liquid's freezing temperature increase. There exists a zone, slightly in front of the interface, into which if a random perturbation could extend, it will find itself at a temperature and solute concentration where further growth will be energetically favourable. [110]

These extensions, in the lower-solute region of liquid, will therefore grow more quickly, amplifying themselves into a series of parallel fingers, advancing into the liquid [111]. This type of interface is called cellular growth.

As the cellular interface advances, the edges of the parallel fingers will meet. Two factors will prevent a pair of colliding crystals from synthesising into a single crystal. Firstly, even if the interfaces both originated from the same nucleus, stochastic imperfections in their crystal structures will result in them having slightly different orientations, making a perfect bond impossible. Secondly, each interface will be pushing a high concentration of solute before it, which will solidify into different, solute-rich phases between the two crystals upon collision.

Figure 2.24 demonstrates entirely separate crystals (grains), with very different orientations growing and colliding. The boundaries between these grains will exhibit completely mismatched lattices, and, depending on the chemical composition, may exhibit very high solute concentrations.

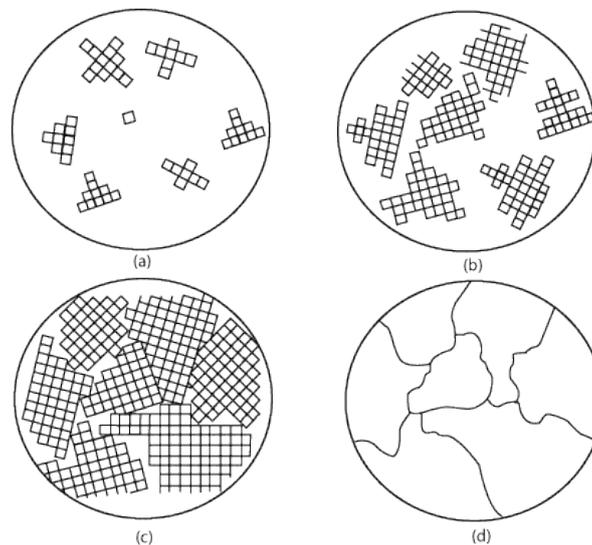


Figure 2.24

A simplified diagram showing how separate nucleation sites will result in lattices with different orientations, which will result in separate grains and their boundaries [112]

Whether an interface remains stable as planar growth, or breaks into the instability of cellular growth, depends on the temperature gradient of the liquid, and the growth rate of the interface. It is the contribution of constitutional undercooling, and the rate of atomic diffusion within the liquid, that prevents the growth rate to be completely dependent on the temperature gradient, as may initially seem logical. At higher growth rates, when the pointed tips of cellular interfaces extend rapidly into the liquid, it is even possible for these conical interfaces to become unstable, and for cellular growth to offshoot horizontally in the form of dendrites [113]. This type of growth is called columnar dendritic. Should the growth rate be great enough, secondary dendrites can grow from the sides of these primary dendrites. At greater growth rates still, tertiary dendrites have been found to protrude from the secondary.

These three types of interface growth - planar, cellular, and columnar - are different forms of the same underlying mechanism of directional solidification, propagating from a starting surface, usually a container boundary. A fourth, and distinctly different type of grain growth, is equiaxed growth. In cooling liquids where nucleation is abundant, possibly due to purposeful inoculation, many grains can grow throughout the melt, with their interfaces expanding roughly equally in every direction until a neighbouring grain is collided with. Due to the lack of directionality amongst the grains, metal parts composed largely of equiaxed grains exhibit very different mechanical properties than parts composed of any other grain types [114]. The effects of microstructure on mechanical properties will be discussed in more detail in this section.

The two most significant factors determining the type of grain growth (both in terms of grain morphology and average size) are the temperature gradient across the solid-liquid interface, and the velocity at which the interface progresses. Of course, each of these factors are both inter-related, and dependent on multiple other factors, such as the rate at which the solid conducts away heat, the thermal properties of the solid and liquid. Figure 2.25 demonstrates how temperature gradient (G) and growth rate (R) determine grain characteristics.

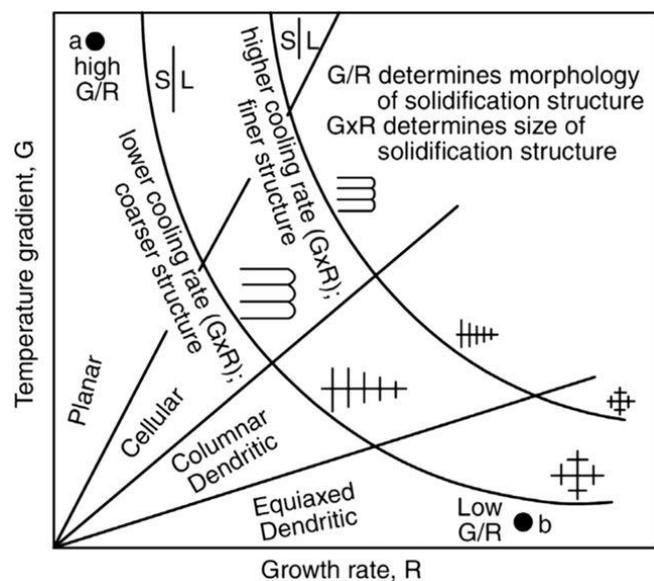


Figure 2.25

Temperature Gradient (G), and solidification velocity (here called Growth Rate, R) determine microstructure. Their ratio (G/R) determines grain morphology, while their product (GR , the cooling rate) determines grain size. [115]

It must be noted that the nomenclature used in literature can be inconsistent. Different authors may use different words to differentiate between: the different fingers growing parallel to each other during cellular or columnar growth; the different regions where the overall direction of this growth slightly shifted due to random perturbations; and the different zones in which solidification originated from clearly distinct nucleation sites. For the remainder of this work, the words used for these concepts will be cells, sub-grains, and grains, respectively.

As a general rule, smaller grains result in stronger materials, as a finer microstructure presents more boundaries and impediments for the movement of dislocations. It is not only the size of grains that determines how many boundaries a moving dislocation will be impeded by, but also their morphology. Columnar grains present many boundaries to dislocations moving perpendicular to the solidification direction, but very few boundaries to those moving in the solidification direction [116]. Therefore, components composed of many columnar grains are anisotropic: their mechanical properties vary based on the direction being tested. Components with mostly equiaxed grains are isotropic, where no such directional dependency exists. Depending on the component's purpose, and the orientation in which it will be used, engineers may subject it to processes to either increase or decrease isotropy. For example, rolling metal can squash grains in one direction and elongate them in another, decreasing isotropy. Alternatively, holding metal at elevated temperatures for prolonged periods can cause grains to recrystallize, increasing isotropy. This kind of heat treatment will be discussed in more detail in this section.

2.3.1.4 - Segregation

Regardless of their morphology, as grains grow, they eject solute into the liquid before them, thereby raising the liquid solute concentration. Under real-world, non-equilibrium conditions, where atoms do not have infinite time to diffuse into their preferred positions, as the liquid concentration rises, so too does the solid concentration, past the theoretical limit discussed previously [117].

It is therefore observed that the chemical composition of the beginning of a grain (in the case of cellular or columnar growth from a container boundary) or the centre of a grain (in the case of equiaxed growth), varies continuously towards the end/edge of it [106].

Assuming that the overall solute concentration of the liquid before solidification is known, and the equilibrium partition coefficient is known, then as an alloy solidifies along the length of its container, it is possible to predict using straightforward calculus, what the local (non-equilibrium) solid concentration will be at any point along the length, using the Scheil equation [118]:

$$C_s = kC_0(1 - f_s)^{k-1}$$

where C_s is the (non-equilibrium) solute concentration at a given point, k is the partition coefficient, C_0 is the overall solute concentration of the mixture, and f_s is the fractional distance of the point along the length of the container.

As the liquid concentration rises, a eutectic composition may be reached, at which the remaining composition of solute and solvent solidifies without further segregation [119]. Alternatively, solute-rich precipitates may form, consuming large quantities of the solute.

As solidification progresses from the start/centre of a grain towards the end/edge, the local solid concentration will rise non-linearly from equilibrium to some higher value. Once the local value has reached the liquid concentration, steady state solidification may occur, whereby a wave of higher

local liquid concentration will be pushed forth by the interface, but otherwise the overall liquid and local solid concentrations will be steady and equal until near the end of the grain [120]. This bow wave is demonstrated in Figure 2.26.

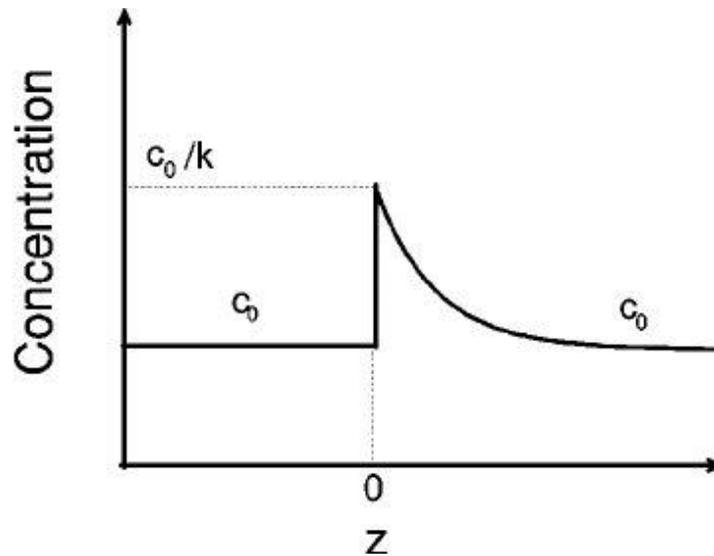


Figure 2.26

A simplified diagram showing the solute concentration in the solid ($Z < 0$) and liquid ($Z > 0$) at C_0 . A bow wave of high solute flows in front of the solidification interface ($Z=0$). The solute concentration immediately ahead of the interface is C_0/k . [121]

After the alloy has entirely solidified, the movement of atoms is not complete. It is very common for different solid phases to be energetically-preferred at different temperatures. For example, as solid Ti-6Al-4V is cooled below around 990°C, the atoms will begin moving and rearranging, transforming from a cubic crystal structure into a hexagonal crystal structure [122].

2.3.1.5 - Post-Solidification Treatments

Heat treatment can refer to a range of processes, all of which involve three stages: heating a component to some temperature between operating and solidus; holding the component at that temperature for a set time; and returning it to room temperature at a roughly-known cooling rate [123]. Heating the microstructure for prolonged periods can allow the crystal lattice sufficient time and temperatures to rearrange into lower-energy states. Dislocations will annihilate, and solute will become more homogeneously distributed throughout the alloy. Recrystallisation, where entirely new grains nucleate and grow, will eventually cause the existing microstructure to be entirely replaced with new, equiaxed grains. This low-energy, defect-free microstructure will have decreased strength and hardness, and increased ductility after the atomic rearrangement. Recrystallisation is demonstrated in Figure 2.27.

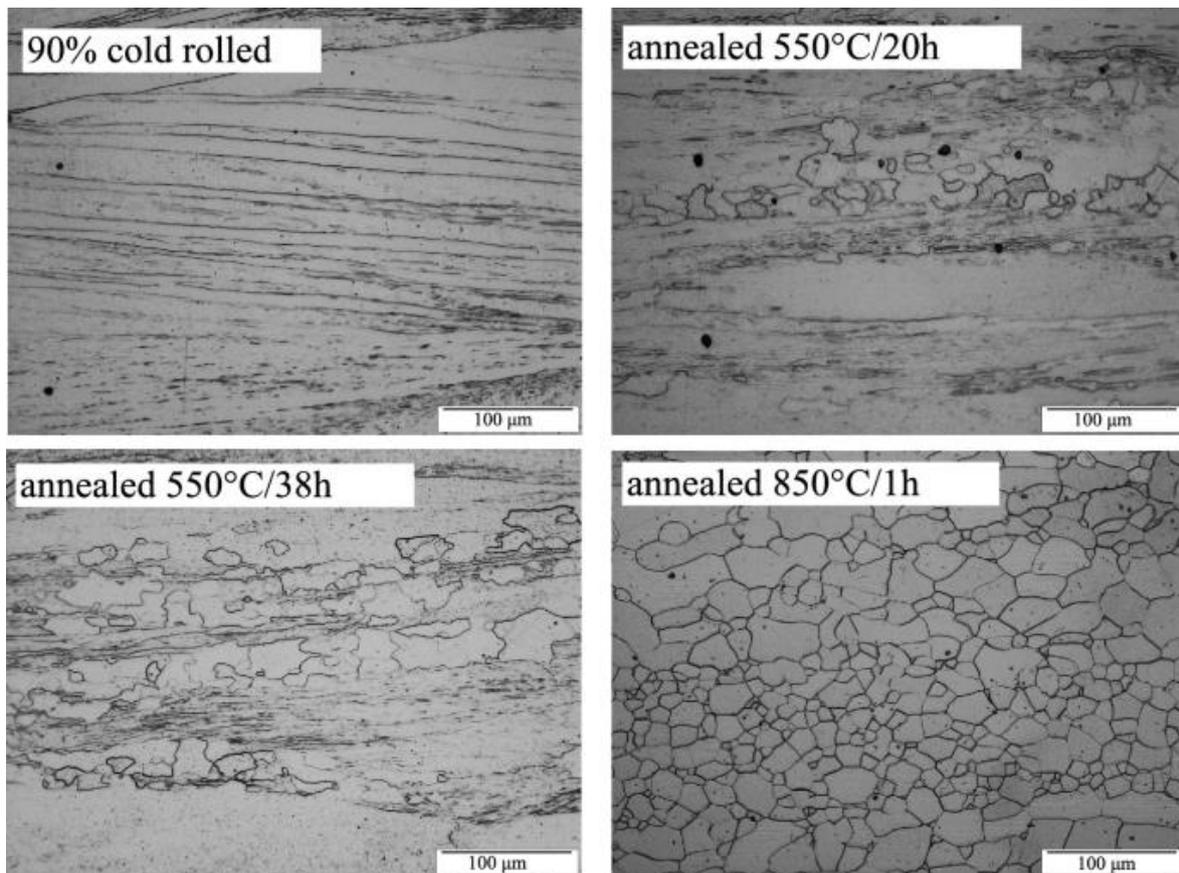


Figure 2.27

A Fe-based alloy that has been cold-rolled to a 90% height reduction, and then annealed according to one of three different methods. Note the partial recrystallization evident at 20h and 38h, and the full recrystallization at 850°C. [124]

The rate at which a metal is cooled determines the overall effect of the heating. Annealing processes involve a slow cooling, which allows the microstructure to gradually adapt to the decreasing temperatures [125]. As the metal cools, the lowest-energy phases will change, and the atoms are allowed time to continue diffusing into these. Once the metal is finally returned to room temperature, it will remain soft and relatively defect-free.

Quenching refers to a much more rapid cooling of the metal, which prevents atomic diffusion as the lowest-energy phases change [126]. The lattice remains locked in a stressed, high-energy state, often resulting in a microstructure that is strong and hard, but brittle. In order to combat this brittleness, combinations of quenching and annealing at various temperatures are often used in industry, in a process called tempering. An example heat treatment recipe, including annealing and quenching, is shown in Figure 2.28.

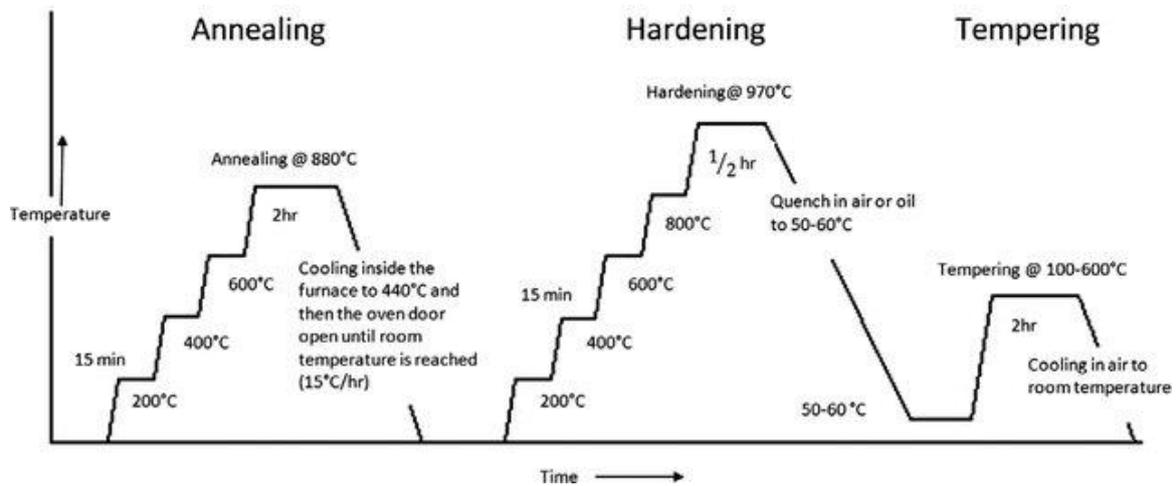


Figure 2.28

The heat treatment routine for a speciality tool steel, with each heating and cooling step designed to promote, remove, or retain certain microstructural elements [127]

During both solidification and heat treatment, as atomic bonds are formed and atoms move through diffusion and solute rejection, it is possible for phases high in solute, which are chemically different to the main lattice, to form and exist as separate grains. The main lattice is called the matrix, while these distinct grains are called precipitates [128]. When precipitates sit incoherently amongst the matrix (i.e. there is high misorientation between the two, especially if the lattices have different atomic arrangements), they act as significant barriers to dislocation movement. Therefore, in general, alloys with numerous, small precipitates are stronger and harder.

Precipitation strengthening refers to heat treatment processes that are intended to increase the amount of precipitates, while also optimising their distribution, in an alloy [129]. The method resembles tempering: the component is annealed until all elements are dissolved within the matrix; then quenched to trap the solute atoms in this now-high-energy state; then heated to some midway temperature, allowing the solute to diffuse together, into precipitates that can exist both between the boundaries of matrix grains, and inside matrix grains. The precipitation strengthening of an alloy surface can be seen in Figure 2.29.

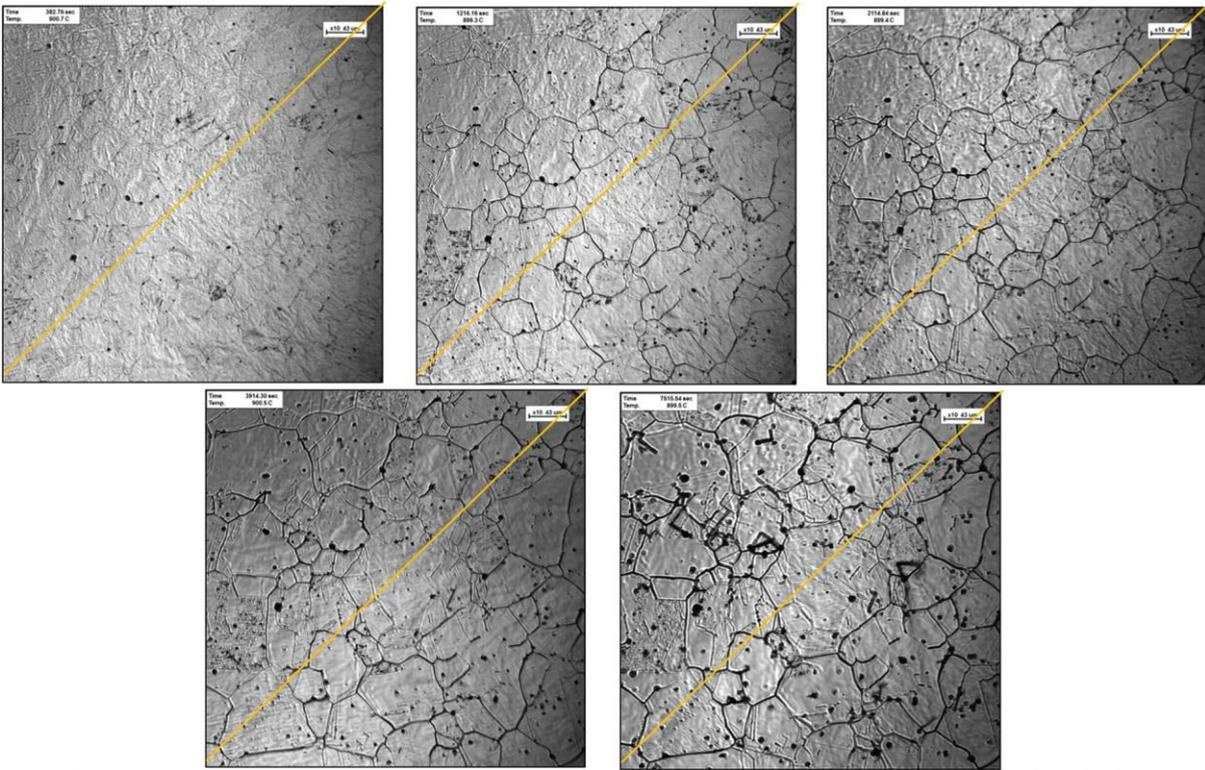


Figure 2.29

High carbon steel, heat treated to 900°C for 1 min (a), 15 mins (b), 30 mins (c), 1 hour (d), and 2 hours (e).

Note the formation of precipitates, here shown in black. [130]

2.3.2 - Metallurgy in Additive Manufacturing

The journey of an alloy being processed through L-PBF, begins with the creation of the powder, typically via atomisation, in which a stream of molten alloy is disintegrated into fine and rapidly-solidifying particles by means of gas, water, or plasma [131].

Gas atomisation (GA) involves the high-pressure melt stream being forced into the path of a high-velocity jet of inert gas, which breaks up the melt into particles of varying size, and carries them into a collection chamber. Water atomisation (WA) processes are similar, but with purified water used instead of inert gas [132]. In the case of plasma atomisation, the alloy enters the path of the high-temperature, high-velocity plasma jet while still in solid form, before being melted and disintegrated. Due to the complexity and expense of this latter technique, it is used in the production of AM powders far less frequently than GA or WA. Upon collection, the powder particles are passed through repeated sieving steps, in order to segregate them into batches of suitable size distribution.

Due to the greater cooling rates of WA, powder particles produced this way have a much more irregular morphology, as opposed to the sphere-like particles produced through GA. Fedina et al. [133], examining the differences in L-PBF tracks of GA and WA powder low-alloy steel, found that the irregular morphology of the latter had significant impacts on packing density, and flowability. Specifically, it was noted that the cross-sectional area of a single track was 31% lower for WA than GA, due to the both the lower packing density, and the inability of the mechanically interlocked

powder particles to flow into the melt pool. Images of the WA and GA powder particles used by Fedina are shown in Figure 2.30.

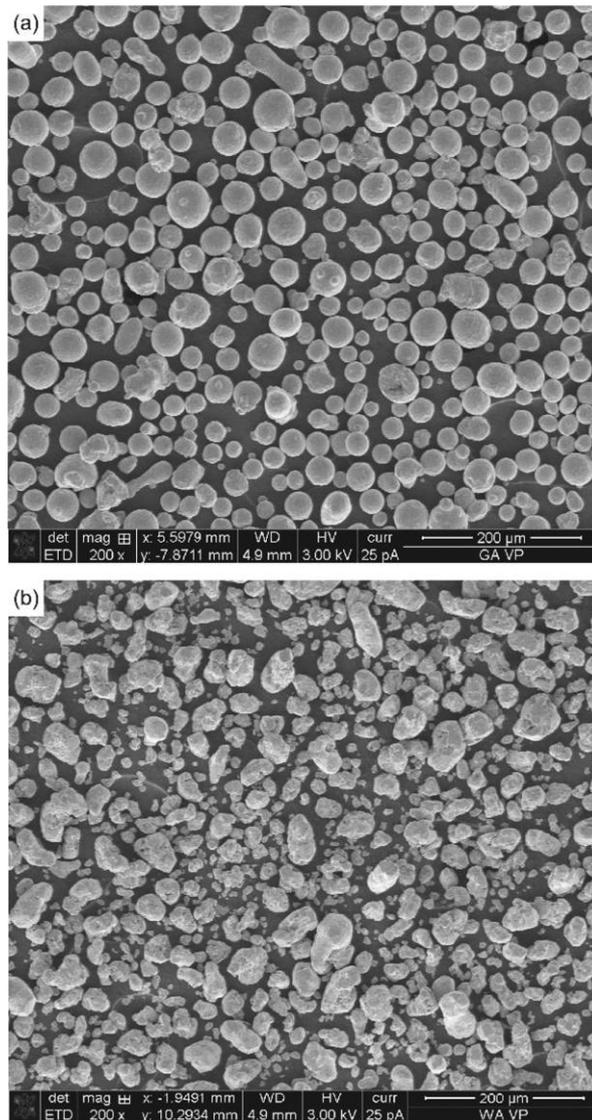


Figure 2.30

The Gas Atomised (top) and Water Atomised (bottom) low alloy steel [133]

2.3.2.1 - Nickel-Based Superalloys in L-PBF

The term ‘Superalloy’ is not strictly defined, but is generally understood to include alloys that are: rich in at least one of the elements nickel, cobalt, or iron; and maintain mechanical and microstructural integrity at high temperatures, under large stresses, and in chemically-harsh environments. Nickel-based superalloys are the industry-standard for many medical, aerospace, automotive, and power-production applications. Nickel forms a face-centred cubic (FCC) structure as the matrix phase, with a variety of other phases being possible depending on the alloy’s composition. Some of these phases contribute positively to a component’s thermomechanical properties, while others contribute negatively or neutrally. Following is a list of the phases commonly found in nickel-based superalloys, roughly ranked by their relevance in the context of L-

PBF. Their standard chemical compositions are given, although it should be noted that substitutional solutes are not uncommon. [134-136]

- The Gamma (γ) matrix phase has a primary composition of nickel, and an FCC structure. Depending on the alloy, Co, Cr, Mo, W, and Fe atoms may exist as substitutes.
- Gamma Prime (γ') is a strengthening phase, also with an FCC structure, that precipitates coherently within the matrix. The composition takes the form Ni_3Al or Ni_3Ti , and can be a significant source of high-temperature strength.
- Gamma Double Prime (γ'') is another strengthening phase, that precipitates in some nickel-based alloys with high iron content. The phase is body-centre tetragonal (BCT): a structure that can form coherently with the FCC matrix, especially when in the flattened disc-shaped form that the precipitate usually takes. The composition is Ni_3Nb , and provides excellent mid-temperature strength.
- The Eta (η) phase is chemically identical to γ' (Ni_3Ti), although with a hexagonal close-packed (HCP) structure that is incoherent with the matrix. η can precipitate either from existing γ' or from the matrix, often as a result of overaging via prolonged high-temperature heat treatment. These precipitates are commonly found at grain boundaries, and reduces ductility, toughness, and creep resistance.
- The Delta (δ) phase is chemically identical to γ'' (Ni_3Nb), but with an orthorhombic structure that is not coherent with the matrix. δ often forms from overaged γ'' , usually at grain boundaries. Its formation negatively impacts mechanical properties, by both reducing γ'' quantity, and also directly decreasing ductility and toughness via incoherence.
- Carbides are hard, ceramic, carbon-rich compounds that form either during solidification (primary carbides) or heat treatment (secondary carbides). They are composed of at least one carbon atom (C), and at least one metal atom (M). MC carbides are primary, have an FCC structure, and can increase alloy strength. M_6C , M_7C_3 , and M_{23}C_6 carbides are secondary, have a complex cubic structure, form along grain boundaries, and can reduce ductility.
- Topologically Close-Packed (TCP) phases are a group of phases that can form due to overaging or high stress. Their structures are complex (e.g. rhombohedral or hexagonal), and do not form coherently with the matrix. They precipitate at grain boundaries, and reduce alloy toughness and ductility both directly through their presence, and indirectly by consuming the extremely temperature resistance refractory elements (Nb, Mo, Cr, Ta, W, Re) that are required for γ' and γ'' . Common TCP phases are Sigma (σ), Mu (μ), Laves, and R.
- Oxides in nickel-based superalloys mostly form on component surfaces, due to high-temperature oxidation. Some oxides, such as Al_2O_3 (alumina) or Cr_2O_3 (chromia) can provide beneficial protection to components, while others such as NiO, TiO_2 , or Ta_2O_5 can be brittle, volatile, and damaging.

Many different compositions of nickel-based superalloy are processable via L-PBF, with each exhibiting a unique combination of phases, benefits, and challenges. Here follows a list of some of the most common of these alloys, which will not be experimentally examined in this project:

- Inconel 625 is a solid-solution strengthened alloy, with significant effects contributed by MC, Laves, δ , and M_6C phases. The powder is highly weldable and printable, and components exhibit high corrosion and oxidation resistance, and high ductility and creep strength. Care must be taken to reduce Laves, however, as these can increase brittleness and reduce strength. [137]
- Hastelloy X is solid-solution strengthened, with M_6C , $M_{23}C_6$, σ , and Laves significant. Components exhibit good as-built properties, and high strength and oxidation resistance in high-temperature environments. The alloy is susceptible to solidification cracking during printing, however, and an abundance of carbides may decrease creep strength. [138]
- Rene 41 is a precipitation-strengthened alloy, with γ' , MC, $M_{23}C_6$, and η phases significant. The γ' provides high-temperature creep and tensile strength, and age-hardening can be easily achieved via heat treatments. Components are extremely crack-prone however, with machine parameters requiring careful optimisation before printing. [139]
- CM247LC is a low-carbon (LC) alloy, in which γ' is the most abundant phase, and γ , MC, $M_{23}C_6$, and TCP phases are significant. The very high γ' content allows for exceptional high-temperature strength, and minor elemental additions lead to good oxidation resistance, and grain boundary strength. This alloy is also extremely crack-prone, with both optimisation and preheating often required before building. [140]
- Haynes 282 is a precipitation-strengthened alloy, with γ' , MC, $M_{23}C_6$, and TCP phases significant. The lower γ' content allows for excellent weldability and printability, while still offering high-temperature strength. The disadvantage of this lack of γ' , however, is that heat treatment is required for precipitations to strengthen the as-built component. [141]
- Inconel 738LC is a low-carbon, precipitation-strengthened alloy, with γ' , MC, $M_{23}C_6$, and TCP phases significant. Similarly to CM247LC, the γ' phase is extremely abundant, resulting in components that retain strength at high temperatures, but are highly prone to cracking. Pre-heating and parameter optimisation are required for successful printing. [142]

2.3.2.2 - Inconel 718 and ABD 900 AM

The two alloys which will be experimentally examined in this project are Inconel 718, and ABD 900 AM.

The as-built microstructure of IN-718 is rich in Laves and MC carbide phases. Heat treatment can dissolve the Laves, and lead to the precipitation of both γ' and γ'' , the latter of which is the primary strengthening mechanism [143]. A micrograph of these precipitates is shown in Figure 2.31. The

size and morphology of grains, in both the as-built and heat-treated forms, depends heavily on various factors. These factors will be discussed in more detail in Section 2.5.

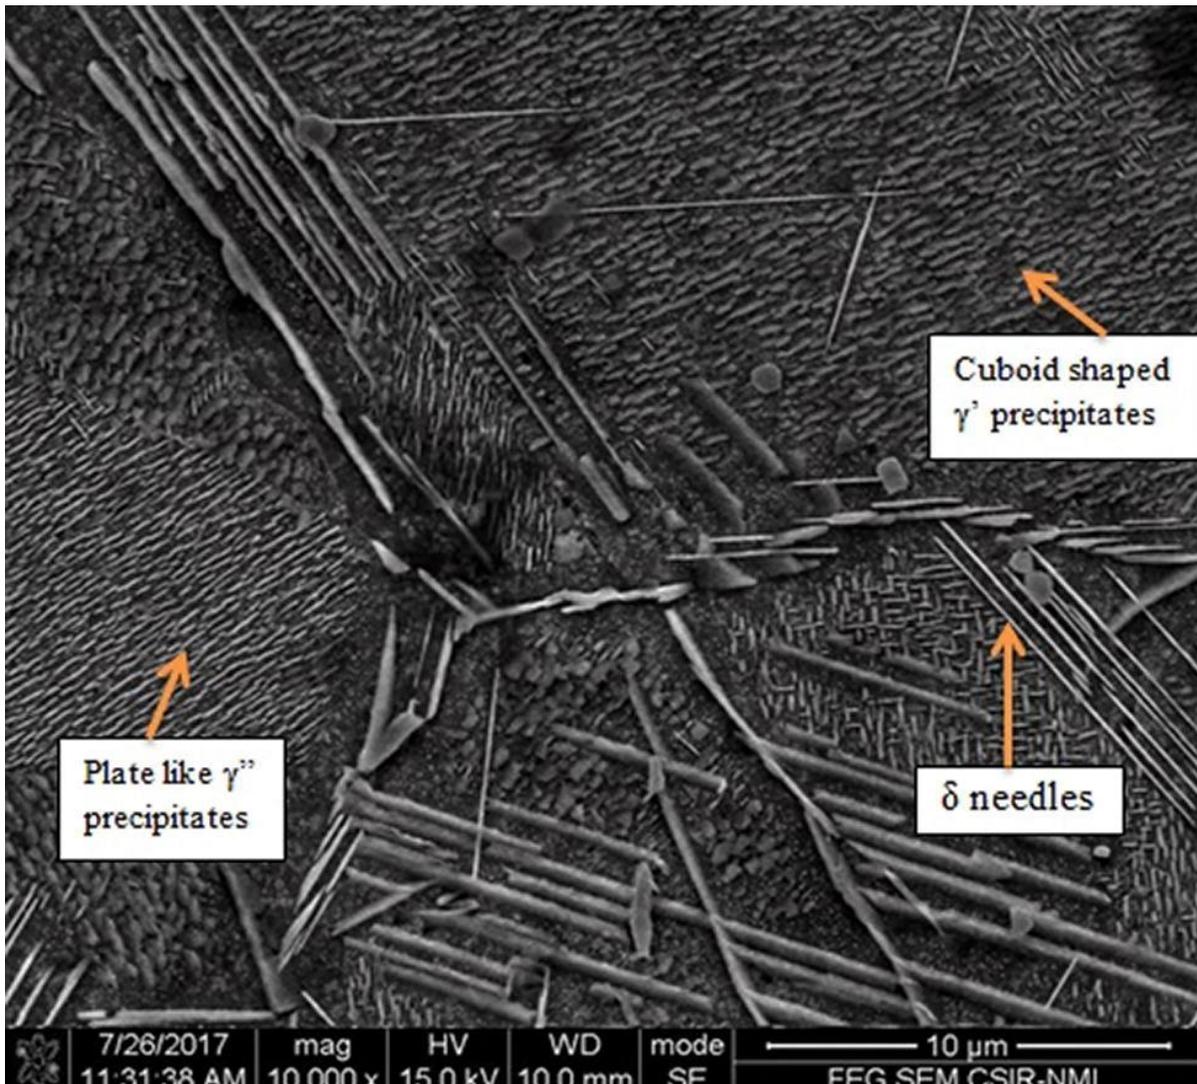


Figure 2.31

A micrograph of an IN-718 sample, showing γ' and γ'' precipitates, and over-aged δ needles [144]

The solidus and liquidus temperatures of IN-718 are 1255°C and 1337°C, respectively [82]. The thermal conductivity of IN-718 increases with temperature, as shown in Figure 2.32. Sweet et al. [145] determined the relationship between the two as:

$$k = 11.45 + 1.156 \times 10^{-2}T + 7.72 \times 10^{-6}T^2$$

where k is thermal conductivity [W/m/K] and T is temperature [°C]. At 25°C and 1200°C, these values are 11.7 W/m/K and 36.4 W/m/K, respectively.

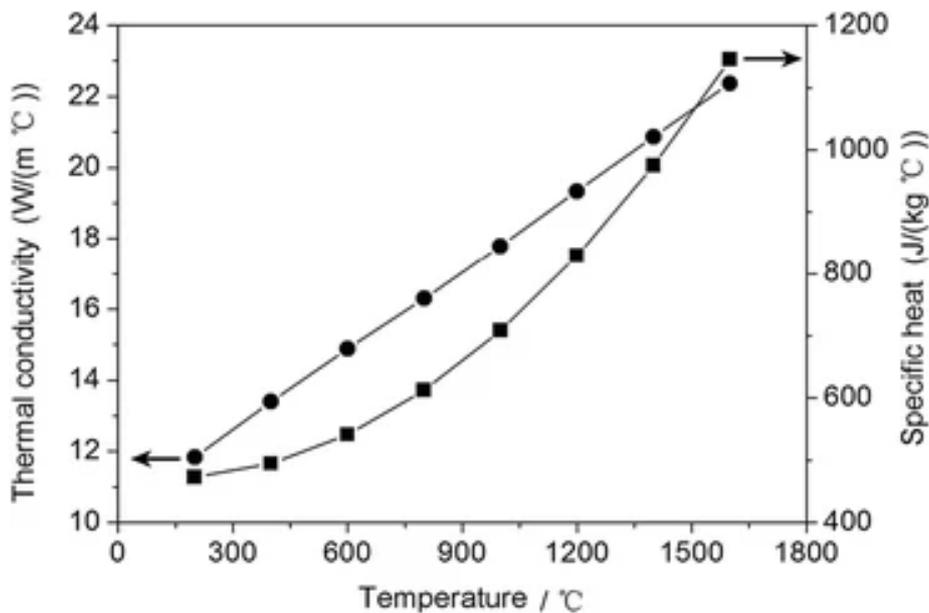


Figure 2.32

The thermal conductivity and specific heat capacity of IN-718 over temperature [146]

The alloy is commonly used in aerospace, automotive, and similar applications, where excellent strength and corrosion resistance are required at very high temperatures. Up to 50% of the weight of some turbojet engines are composed of IN-718. Its high weldability, and lack of solidification cracking, make it an ideal alloy for the purposes of L-PBF [147].

The most notable limitation of IN-718 in the context of L-PBF, is the fact that due to the strength that comes from its precipitates, heat treatment is required in order to fully realise the mechanical properties of components. Abundant Laves phases, which reduce ductility, as well as extreme anisotropy, must be addressed with annealing. Also, the fact that IN-718 exhibits both lower thermal conductivity, and greater elastic modulus, than many other alloys used in L-PBF, can result in higher residual stresses being left in the wake of the extreme thermal cycling that the process involves. Barros et al. [148] found also that residual stresses of up to 400 MPa in the transverse direction, and up to 600 MPa in the build direction, existed in as-built components. These stresses were significantly reduced via heat treatment.

Alloys By Design (ABD) is a computational alloy design approach, developed by *Alloyed*, a company spun out from Oxford University in 2017. ABD 900 AM was created via ABD, specifically for use in additive manufacturing [149].

ABD 900 AM exhibits excellent high-temperature strength, thanks to the large (approx. 35% by volume) γ' content. Minor amounts of carbides and borides are also present in the microstructure. The solidus and liquidus temperatures are 1305°C and 1380°C, respectively. At 25°C and 1200°C, the thermal conductivity of ABD 900 AM is 11.0 W/m/K and 30.1 W/m/K, respectively [150]. The thermal conductivity and specific heat capacity are shown in Figure 2.33.

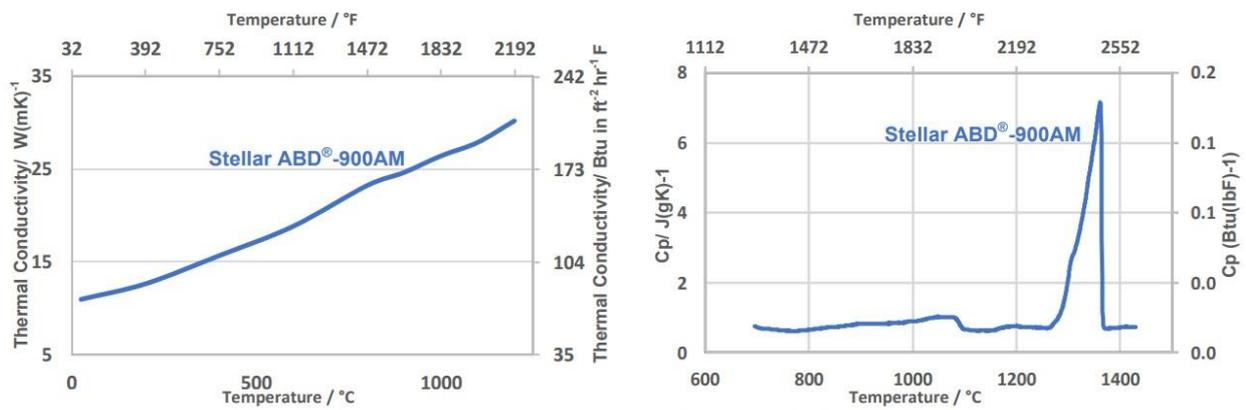


Figure 2.33

The thermal conductivity and specific heat capacity of ABD 900 AM [151]

Due to its very recent arrival on the market, the industrial usage of ABD 900 AM has not yet been realised to the extent that its designers intended. More research, characterisation, and qualification will be required before the alloy becomes widespread in critical components.

2.4 - Defects

This section will provide a comprehensive overview of the ways in which an as-built L-PBF component will deviate from its designer's intended specification, as well as some post-processing techniques that may reduce these deviations, or otherwise improve the properties of the component. For the purpose of clarity, defects will be divided into four broad categories: geometrical, surface, microstructural, and mechanical. In reality, of course, many defects could actually fall into multiple, or all, of these classes.

2.4.1 - Geometrical Defects

Geometrical defects are those which cause the component's size and shape to differ from the CAD model at a macroscopic scale, visible to the human eye. Errors may come in the form of geometric inaccuracy, dimensional inaccuracy, or thermal warping.

2.4.1.1 - Geometric Inaccuracies

Geometric inaccuracies occur when the L-PBF machine is unable to reproduce the form of the CAD model, irrespective of any other imperfections. The most unavoidable of these inaccuracies is the 'staircase effect', which presents as discrete layers being visible on upwards-facing slopes (upskins), and in an inevitable result of the layerwise fundamentals of the AM process [152]. In theory, the 'steps' of the staircase should equal the layer height. In practice, upskin surface roughness is often even greater. The staircase effect is demonstrated in Figure 2.34.

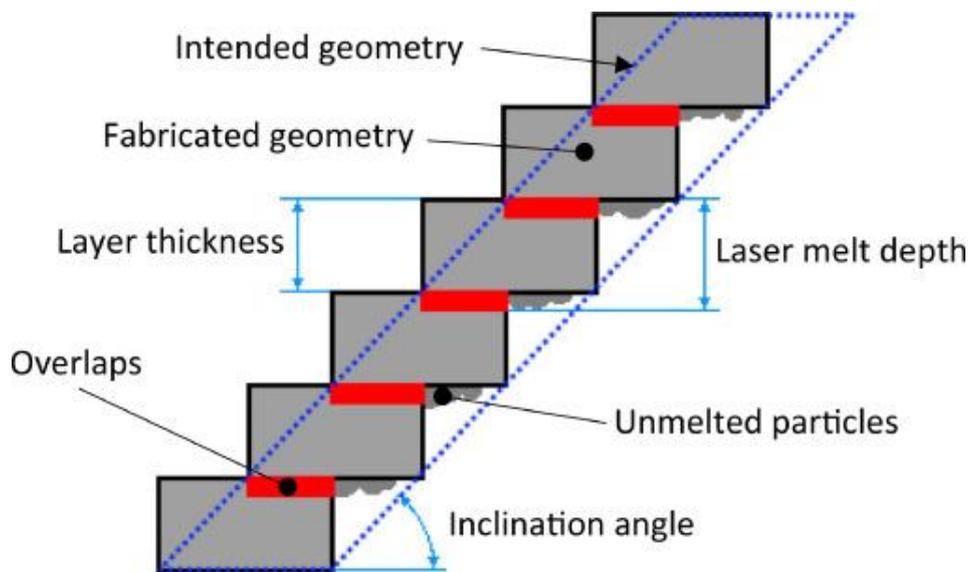


Figure 2.34

An XZ schematic, showing how the CAD geometry (blue line) can result in a staircase with a poor upskin surface [153]

Geometric inaccuracies can also be caused by machine errors. For example, should the galvanometers that angle the mirrors to aim laser beam direction become misaligned by a fraction of a degree, this error will be transformed to significant misstep on the powder bed surface [154]. Likewise, inaccuracies may arise from errors in the vertical displacement of the build platform, or from errors in beam shape or focus, via the collimator, beam expander, or f-theta lens [155].

2.4.1.2 - Dimensional Inaccuracies

Dimensional inaccuracies occur when the size of component features differ from their intended sizes, mainly due to thermal shrinkage. Assuming that each individual weld track meets its geometrical requirements in the instant after full solidification, the thermal contraction it undergoes as it cools from solidus to ambient temperatures will be significant. For example, a 20mm track of IN-718, with a coefficient of thermal expansion of $14 \times 10^{-6} / ^\circ\text{C}$ [82], cooling by 1200°C , if unconstrained would contract by 0.336mm; an error large enough to be noticed in certain high-end contexts.

Ning et al. [156] found that the extent of thermal shrinkage can be reduced via a reduction in line energy density, accomplished through use of lower power, or larger laser speed or hatch spacing. Likewise, Wang et al. [157] found that lowering volumetric energy density by increasing layer height decreased thermal shrinkage. Of course, altering these parameters with the intent of reducing thermal shrinkage will also effect build time and other factors. A geometrical compensation method to combat shrinkage is shown in Figure 2.35.

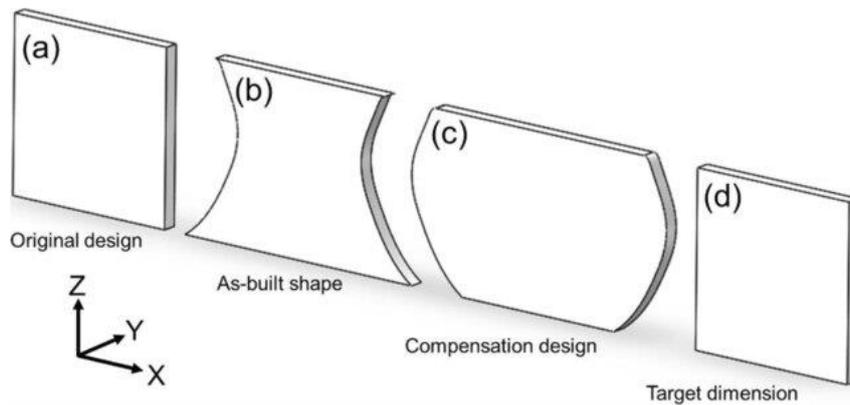


Figure 2.35

An exaggerated schematic, showing how an attempted rectangle (a) can contract due to thermal shrinkage (b), and so a compensated design (c) must be used to achieve the desired dimensions (d) [158]

Since thermal contraction due to metal cooling cannot be entirely avoided, technicians must instead design builds and L-PBF systems to compensate for its effects. Shrinkage is much more prevalent along the build direction, since each new layer is constrained laterally by previous layers. Therefore, it must be understood that the height of any previous layer beneath the build chamber's level will be slightly greater than that explained by build platform movement alone, and excess powder must be dosed each layer to compensate for this [159]. Also, technicians should be aware that for each scanned layer, thermal shrinkage will be more prevalent in areas with greater heat accumulation. Scanning patterns must therefore be designed to attempt uniform heating across the bed [160].

In an attempt to minimise thermal shrinkage (as well as reduce staircase surfaces and required supports), Paramita et al. [161] designed a novel approach to optimising build orientation. Their system analysed the faces of each surface of the components, and analytically determined build angles for best outcomes. Compared to non-optimised builds, their components achieved an average reduction of 34% in deviations from the intended CAD model.

2.4.1.3 - Thermal Warping

When thermal shrinkage cannot be constrained laterally, and the ensuing contraction is so significant that it forces previously solidified layers to curl upwards, the errors present as thermal warpage. Warping develops not just as a result of high thermal gradients within a single layer, but as a combined result of thermal gradients between multiple layers. Each successive layer contributes additional tensile stresses (due to lateral thermal contraction), pulling outer edges inwards with greater force as the build progresses, until elastic deformation occurs [162].

Reducing thermal warping is a complex process, which will be discussed in more detail in Section 2.5.3 on residual stresses. Various methods have shown promise, such as keeping scan lengths below 15mm [163], reducing line energy density [164], increasing layer height to reduce the number of layers in a build [159], and utilising scan patterns that minimise heat accumulation and thermal gradients [165].

2.4.1.4 - Post Processing Geometrical Defects

There is little that can be done to correct geometrical defects after a build is finished. It is impractical to attempt to reverse warpages or geometric inaccuracies. In a few specific cases, such as if a hole is misaligned due to thermal contraction, subtractive methods like CNC milling may provide improvement. In general, though, if any of the defects listed previously are significant enough to render the component unusable in its as-built state, it must be scrapped. CNC milling (and similar methods) can be highly effective in reducing surface quality defects, which will be discussed.

2.4.2 - Surface Quality Defects

Surface quality defects are those which cause the outer surfaces of a component to be more irregular than intended, either physically or chemically. There are multiple reasons why component surfaces will deviate geometrically from the ideal, smooth surfaces of a CAD model. Fundamentally, upskin surfaces will always suffer from the staircase effect, and vertical and downskin faces will suffer because weld tracks exist as semi-circular trenches, therefore making perfect joining of layer edges impossible.

Strategies have been identified to alleviate the effects of these unavoidable issues. Experimenting with AlSi10Mg, Krishnan et al. [166] determined a line energy density of 1.2 to 1.8 J/mm² to result in the best surfaces: lower energies led to uneven and porous surfaces, while at high energies surfaces were degraded due to track thermal distortion. Working with Ti-6Al-4V, Kruth et al. [167] found that using larger laser spot sizes (up to 200µm) significantly improved surface quality, due to a remelting of adjacent tracks. In samples of 316L steel, Yasa et al. [168] achieved a 90% reduction in the average roughness of components by entirely re-scanning each layer, as shown in Figure 2.36.

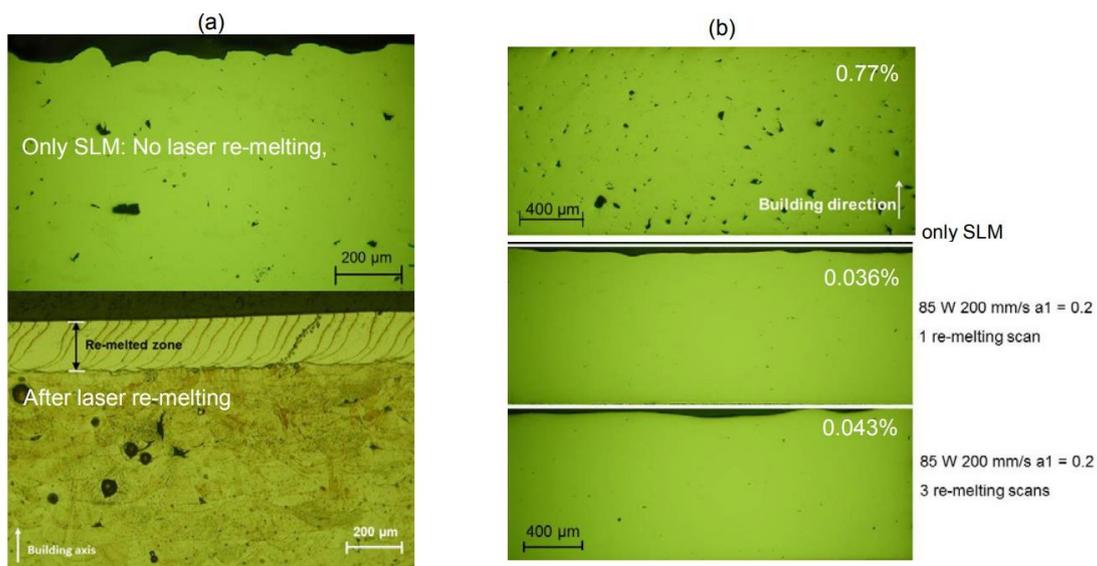


Figure 2.36

Yasa achieved a better surface finish (a) through re-melting, and also reduced porosity (b). Note how porosity actually increased when the layers were re-melted three times, rather than just once. [168]

When designing components, and positioning them upon the baseplate, engineers must be aware of the circumstances which both reduce and improve surface quality. In general, upskin slopes display better surface quality than downskin, as the former only exhibits the staircase effect, while the latter exhibits both the staircase effect, and the chaotic underside of melt pools building upon powder [169]. Also, in terms of non-sloped surfaces, side surfaces are generally of poorer quality than top surfaces, due to the tendency of the latter to incorporate partially-melted or sintered particles adhered to them [170]. It should also be noted that poor surface quality can become cumulative: defects on one layer can amplify defects on subsequent layers [171]. Therefore, components should be designed in a way to have their most error-prone features printed last.

2.4.2.1 - Edge Effects

It is often observed that the edges of each scanned layer rise above their centres. The exact reasons for this “edge effect” are difficult to measure directly, although several hypotheses have been suggested [172-174]:

1. The first track to be scanned in each part in each layer (which is very likely to be an edge, in most scan strategies) will feature larger melt pools than subsequent layers, due to the lack of adjacent solidified tracks to conduct heat laterally into.
2. Denudation is a phenomenon in which powder particles that lie adjacent to a track being scanned are blown away by gas flows, resulting from the melt pool’s vapour plume, and from convective currents within the inert atmosphere. Denudation is demonstrated in Figure 2.37. Since this process will not have affected particles in the path of very first track, its melt pools will be larger.
3. Assuming that wetting angles between solidified material and melt pools always remain equal for a given layer, melt pools in the first track (in which both sides of the melt pool contact the relatively-flat previous-layer surface below) will be pulled narrower and taller by surface tension than melt pools in the subsequent tracks (in which one side will contact the steep walls of an adjacent track).
4. At the end of each vector (which will also likely be at an edge), the laser switches off instantly, thereby causing the mass of molten material ahead of the vapour depression to be rapidly solidified in place, rather than flow around the depression to fill it from behind, as happens elsewhere in the track.

The edge effect can cause uneven surfaces, a lack of powder particles spread in its wake, or, in extreme cases, wear to the recoater blade. Kleszczynski et al. [175] proposes a reduction in laser power for the vectors likely to cause bulging, as a means of mitigating the effect.

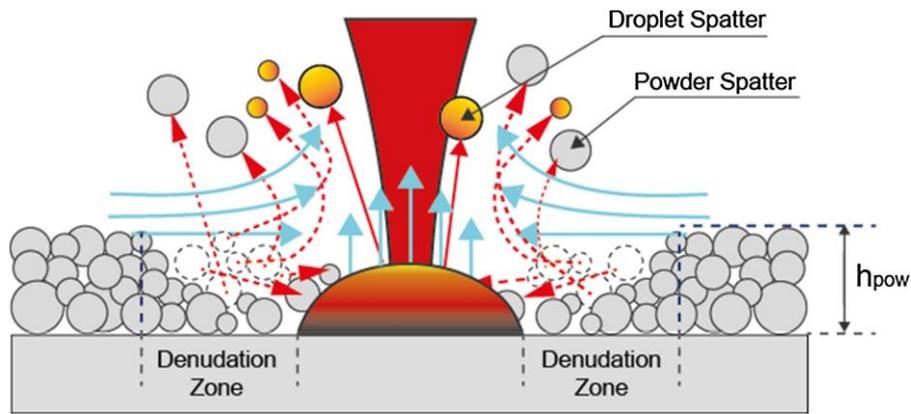


Figure 2.37

Denudation: powder being swept out of the paths of adjacent melt tracks, by the convection currents of this melt track [176]

2.4.2.2 - Periodic Oscillations

Researchers often find that, for their alloy and build geometry, there exists a range of energy densities in which laser scans result in continuous tracks, but above and below which periodic imperfections occur. Different authors use different terms to refer to these low and high energy defects. This work will refer to the periodic oscillations visible upon solidified tracks with insufficient energy density as “humping” or “balling”, and to those visible upon tracks scanned with too much energy as “rippling”.

At low energy densities, the laser cannot ensure a deep and wide melt pool that penetrates into the previous layer. Instead, newly-melted powder particles sit upon the solid layer, analogous to a cylinder of molten liquid. In order to minimise surface energy, stochastic perturbations along the cylinder’s surface become exaggerated, as it attempts to break up into discrete globules. Usually, the wetting between the molten and solid material will be sufficient to prevent a complete break up, and humping will occur: the track will solidify with alternating tall-wide sections, and short-narrow sections. Should the wetting not be sufficient, however, balling will occur: the track will degrade into a series of non-touching droplets [177]. This process is shown in Figure 2.38.

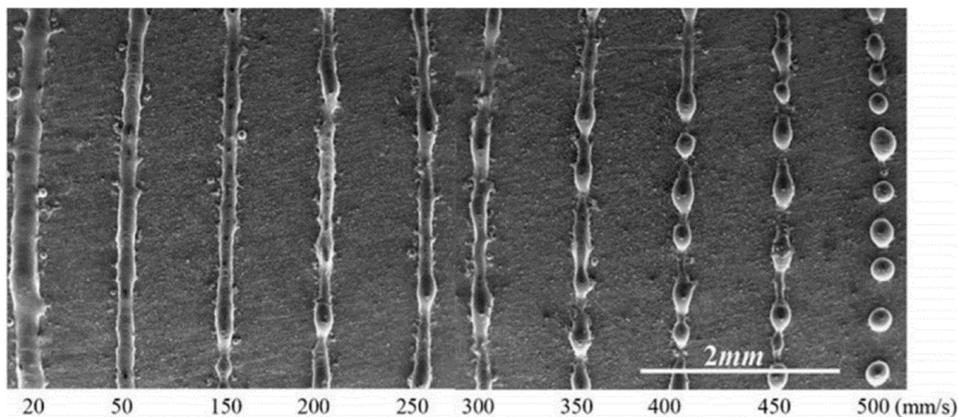


Figure 2.38

The effects of laser speed on the balling of 316L at 190W [178]

The wetting between the liquid inside a melt pool, and the underlying solid substrate, is both a crucial factor in preventing balling, and a complex process to predict or control. Some alloys (such as aluminium with high silicon content) inherently exhibit low surface tension, and therefore are less susceptible to balling. Others (such as titanium or steels), exhibit the opposite. Certain machine parameter settings have been shown to reduce surface tension and balling: preheating the baseplate, especially to temperatures close to the alloy's solidus [163]; using powder with smoother, more spherical surfaces [179]; and ensuring as low an oxygen content as possible within the powder and build chamber [180].

Tang et al. [181], working in 316L steel, determined that the prevalence of balling depends entirely on scanning speed. At higher speeds the defect occurs, even if laser power is increased to keep line energy density constant.

Gu et al. [182], also using 316L steel, divided the balling phenomenon into two categories. At lower laser powers, coarse balls form along the track, composed of unmelted powder particles, sintered together or joined with limited amounts of solidified liquid. These large balls contribute extreme structural weakness to the component, due to the poor bonding within them. At higher powers combined with higher speeds, micrometer-scale balls were formed due to the high surface tension forces described previously.

At higher energy densities, the vapour depression inside the melt pool will deepen, transforming the pool into a keyhole-shaped channel, and forcing the displaced liquid out to the melt pool periphery, as discussed in Section 2.2.2. Simultaneously, Marangoni forces create a difference in surface tension between the pool's centre (higher temperature, lower surface tension) and periphery (lower temperature, higher tension), which results in a radial flow of material, outwards from the melt pool centre. Combined, these effects induce a significant displacement of mass inside the melt pool.

This mass flow is never entirely steady. Size, shape, and placement of powder particles, fluctuations in laser power, and small differences in local temperature, all contribute to stochastic instability in the fluid mechanics of a travelling melt pool. This instability, combined with the effects of surface tension, result in the formation of periodic oscillations upon the molten material at the rear of the melt pool, much like the waves following a boat. Due to the rapid solidification of the rear of the melt pool, these ripples can be left frozen in its wake [183].

Caprio et al. [184], working with 316L steel, found that the frequency of these ripples depended on the penetration depth of the melt pool. As laser power was increased from 200W to 300W, the length, width, and depth of the melt pool all increased linearly, while the frequency of the oscillations following it decreased from 5.5 kHz to 3.5 kHz, with the rate of decrease slowing down as the power increased. This rippling is shown in Figure 2.39.

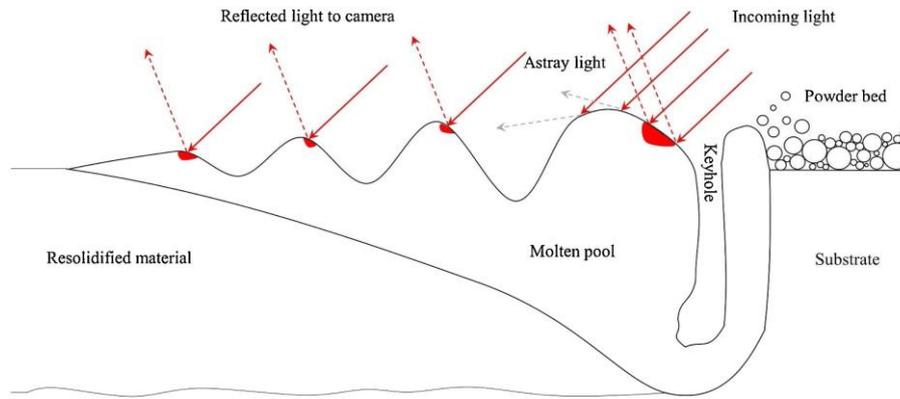


Figure 2.39

Caprio used the reflectivity of ripple peaks to estimate keyhole depth [184]

Because these ripples do not pose significant structural impact on components, and their harm is largely limited to minor aesthetic effects on outer surfaces and the final layer, they appear in very little literature. A gap exists in better understanding which machine and material parameters lead to larger and small ripples, in terms of both amplitude and frequency.

2.4.2.3 - Oxidation

Jia et al. [185] determined that for IN-718, oxidation can be classified into two categories. An External Oxidation Scale (EOS) is formed around the melt pool via chemical absorption of oxygen, leading to a rapid mass gain. After solidification, an Internal Oxidation Zone is formed via diffusion of oxygen through the oxide layer. It was found that as the laser volumetric energy density increased, the thickness of both of these regions decreased (EOS fell from $\sim 50 \mu\text{m}$, to $\sim 40 \mu\text{m}$, to $\sim 12 \mu\text{m}$; and IOZ fell from $\sim 80 \mu\text{m}$, to $\sim 55 \mu\text{m}$, to $\sim 30 \mu\text{m}$; as the VED rose from 70 J/mm^3 , to 110 J/mm^3 , to 130 J/mm^3). At the highest densities, the oxide layer was found to be dense and protective, consisting largely of Cr_2O_3 .

Lu et al. [186], experimenting with 316L steel, demonstrated that even with oxygen levels inside the build chamber as low as 500 ppm (with an oxygen partial pressure of 0.05 bar), this is still more than enough oxygen for oxides to readily form. At the high temperatures present inside the melt pool, the oxygen partial pressure required for the formation of oxides like Cr_2O_3 , MnO , and SiO_2 , can be as low as 10^{-12} bar. This work therefore further showcases the near-inevitability of some degree of oxidation in the L-PBF process.

The two major sources of oxygen for oxidation processes are the processing environment, and the powder feedstock. Even once a build chamber is purged with an inert gas like nitrogen or argon, small pockets of oxygen may remain, and leakage from poor seals may introduce additional oxygen. Oxygen will also exist within air entrapped amongst powder particles during storage or handling, and within air entrapped inside particles during powder production. To minimise this, it is recommended to continually flood the powder hopper with inert gas during building, and to use finer powder particles as these are less susceptible to containing pores of gas [187].

2.4.2.4 - Post Processing Surface Quality Defects

Modern 5-Axis Computer Numerical Control (CNC) machines have the ability to accurately and securely manoeuvre cutting tools onto almost any external surface on a component. Depending on part geometry, CNC machines can be extremely effectively at reducing surface roughness. Çevik et al. [188] built a set of Ti-6Al-4V samples, with an as-built Ra (mean height of the imperfections that cause deviation from perfect flatness) ranging from 6.98 μm to 9.69 μm . After CNC milling, this roughness was reduced to a range of 0.26 μm to 0.92 μm .

Electrical Discharge Machining (EDM) uses rapid, repeated, high-voltage sparks to erode a surface, with tiny volumes of metal, concentrated in a specific area, being melted and vaporised thousands of times per second. EDM is often the preferred method to cut a component from a baseplate after building, and can also be used to reduce surface roughness. Çevik et al. also experimented with EDM on their Ti-6Al-4V samples, but found that the Ra was only reduced to a range of 2.31 μm to 3.09 μm : a lot lower than the as-built roughness, but still significantly greater than the CNC machined surface roughness.

Shot peening involves small, spherical balls being blasted at a surface at high velocity. Roughness is improved not via material removal, but rather through plastic deformation that smooths out protrusions. Both the porosity and the microstructure of the sub-surface material can be affected as well, and significant residual stresses can be left within the component's outer regions. Lesyk et al. [189], working with IN-718, found that shot peening could reduce the Ra of as-built surfaces by 53% to 19%, depending on the shot velocity used. It was found, however, that Rz (maximum valley-peak deviations from a perfect surface) could be reduced by 20% at lower velocities, but increased by 50% at higher velocities, indicating the dangers of using the method incorrectly. Porosity was also found to be reduced by over 60% close to the impacted surface. These results are shown in Figure 2.40.

2.4.3 - Microstructural Defects

Microstructural defects are those which remain largely invisible, even if components are sectioned, polished, and examined under low-magnification optical microscopy, but which affect the composition, morphology, size, and bonding of grains, sub-grain structures, and intergranular material. These defects can be divided into four categories: microstructural anisotropy, segregation, non-equilibrium phases, and microstructural heterogeneity. Effective remedies exist for the majority of microstructural features, in the form of heat treatments, which will be discussed in Section 2.4.4.

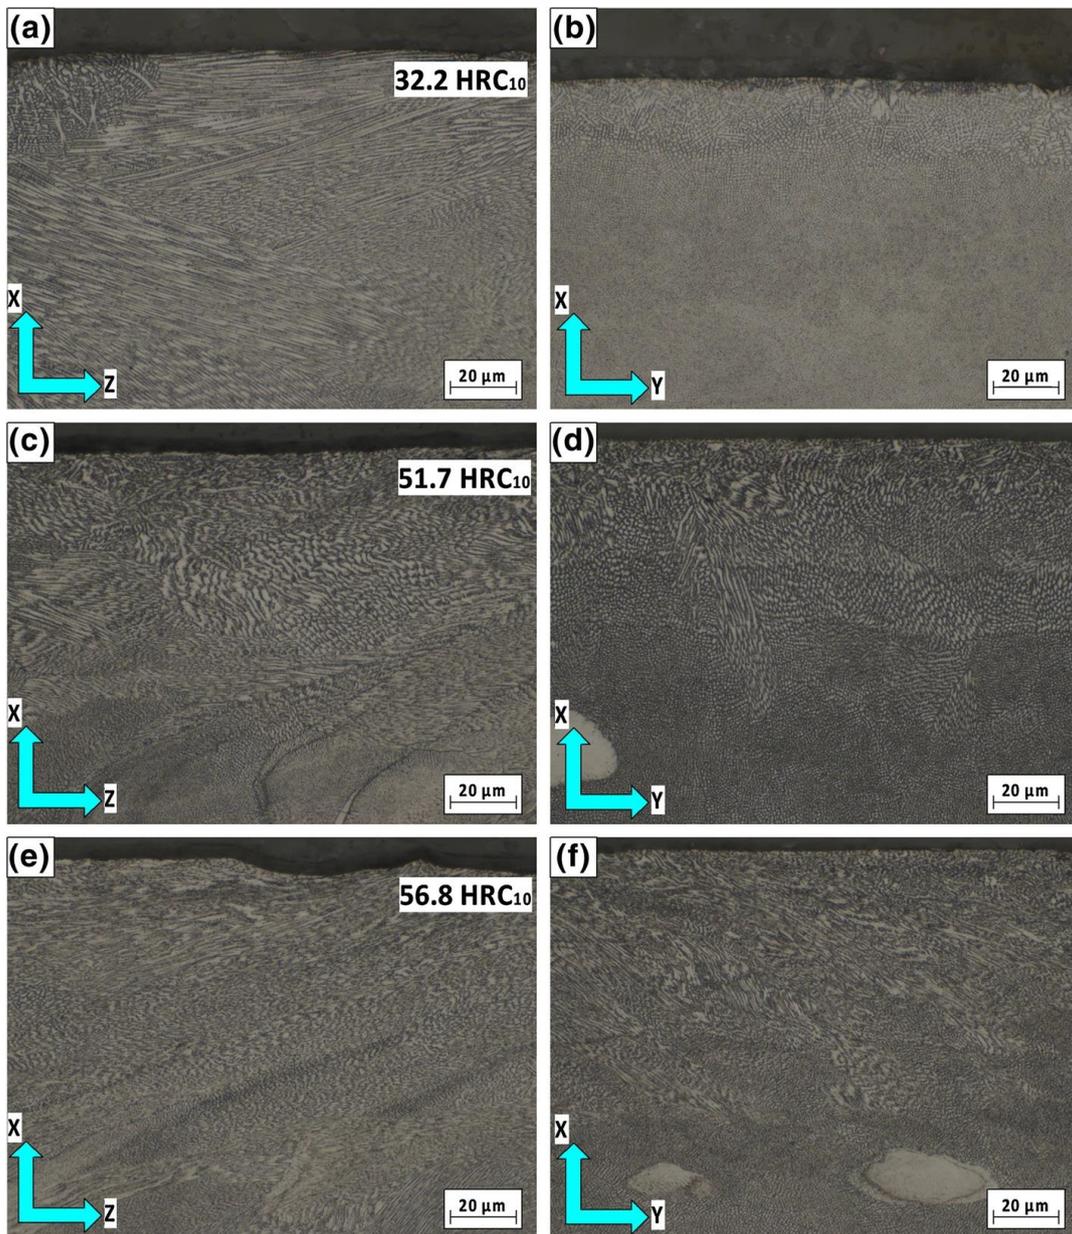


Figure 2.40

XZ (left column) and XY (right column) of as-built sample (a, b), and those subjected to shot peening at low (c, d), and high (e, f) velocity by Lesyk. Note how in the XZ images, the surface is smoothest at low velocity, but less so at high velocity. The surface micro-hardness (HRC), however, increases at high velocity. [189]

2.4.3.1 - Microstructural Anisotropy

Anisotropy can exist at all scales within components produced via L-PBF, from the centimetre to the micron. Extreme thermal gradients exist within every melt pool, from the vapour depression at its centre, towards the mushy zone at its periphery. The upper section of each melt pool, where the greatest thermal gradients exist laterally, are likely to be remelted and incorporated into the lower sections of melt pools on subsequent layers, where the greatest thermal gradients follow the build direction. As a result, the entire component will be constructed of a microstructure which rapidly solidified from one direction only.

This combination of extreme thermal gradients, and unidirectional solidification, often results in fine, columnar grains, oriented with their long axis along the build direction [190]. Epitaxial growth is also common: liquid at the bottom of a melt pool nucleating coherently with the underlying substrate, and therefore growing as a continuation of the underlying grains. Components can therefore become formed largely of columnar grains which extend through multiple layers, resulting in very significant differences in the strength and ductility of the build direction, versus the lateral direction. Epitaxial growth is shown in Figure 2.41. Methods that had been shown to prevent epitaxial growth include: adding inoculants to promote nucleation ahead of the solidification front; manipulating laser parameters to achieve a growth rate and thermal gradient that favours more equiaxed grains; and alternating between thick and thin powder layers, to alter the thermal gradient and therefore growth direction of grains between layers [191].

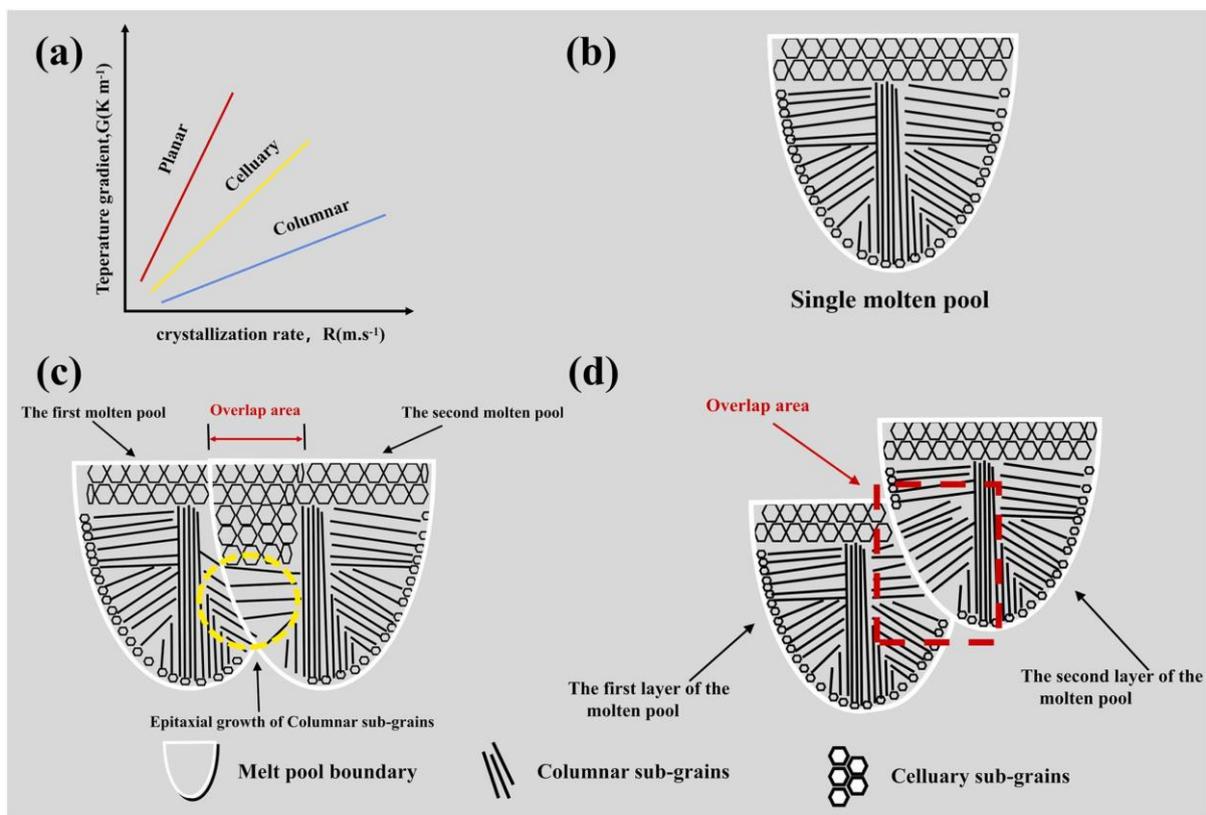


Figure 2.41

The GR diagram as presented previously (a); the grain morphologies found in a single melt pool, with cellular around the edges and columnar in the centre (b); and the epitaxial growth of columnar grains between adjacent tracks (c) and subsequent layers (d) [192]

2.4.3.2 - Segregation

As discussed in detail in a Section 2.3.1, segregation refers to the non-uniform distribution of alloying elements during solidification. Segregation can occur both on the macro-scale (the width of depth of a melt pool) and the micro-scale (the length of cells or even dendrite arms).

Elements like Nb and Mo in IN-718 [193], Si in AlSi10Mg [194], and Cr and Mo in 316L [195], tend to segregate preferentially to the liquid phase, leading to a solute-rich grain boundaries and dendrite arm boundaries. The effects of this microsegregation are complex. In IN-718, for example: the interiors of dendrites, rich in Al but depleted of Nb and Mo, favours the formation of γ' ; and the periphery of dendrites, rich in Nb and Mo, favours the formation of γ'' ; while interdendritic spaces, however, rich in Nb, Mo, and C, favour the formation of harmful Laves phases. Figure 2.42 shows the segregation of Nb during dendritic solidification.

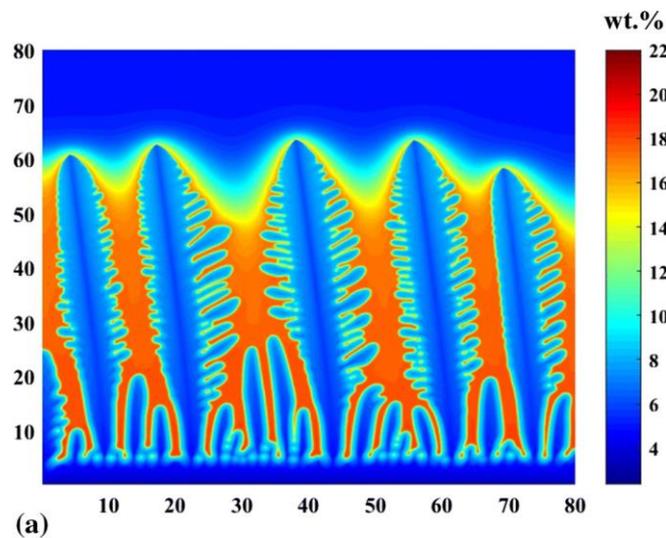


Figure 2.42

An integrated FE and PF simulation of IN-718 grain growth, showing the concentration of Nb. Note how it segregates from the growing dendrites into the liquid. [196]

Ghosh et al. [197] performed an uncertainty analysis to determine the factors that influenced segregation in the L-PBF of IN-718. It was found that the velocity of the solidification front, and the diffusivity of the liquid, were the two most significant factors, while thermal gradient was negligible. Microsegregation was found to increase with an increasing front velocity, and decrease with increasing diffusivity. This study suggests, therefore, that reducing laser speed may lead to a more homogeneous microstructure.

2.4.3.3 - Non-Equilibrium Phases

A related defect, also a direct consequence of the extreme solidification conditions found in L-PBF, is the formation of phases that exist outside equilibrium conditions either at the moment of solidification, or after the component's rapid cooling. As Song et al. [198] points out, for example, the Laves phase is not present on the IN-625 phase diagram, so therefore its presence in components indicates that precipitation is occurring outside of equilibrium conditions.

Non-equilibrium phases show a high propensity as nucleation sites for crack and micro-voids, and also result in high-density dislocations and residual stresses, due to uneven thermal contraction compared with surrounding equilibrium phases [199].

2.4.3.4 - Heterogeneity

Heterogeneity refers to distinct microstructures existing in a component over the macro-scale: different thermal conditions in different geometrical features resulting in different phases and grain morphologies. Conduction conditions, which affect heat accumulation, thermal gradients, and melt pool geometries, will vary significantly across any complex part, leading to microstructural heterogeneity. Using higher line energy density can increase heterogeneity [200], while using longer scan tracks (and therefore reducing heat accumulation) can decrease it [156].

2.4.3.5 - Post Processing Microstructural Defects

As discussed in Section 2.3.1, a metallic microstructure can be altered (through atom diffusion, stress relief, phase transformation, or recrystallization) by heating it up to some temperature between ambient and solidus, holding it there for some time, and cooling it. Essentially, all heat treatments can be described by four parameters: the furnace temperature, the time at that temperature, the cooling rate, and the gas pressure applied inside the furnace.

It is useful to categorise heat treatments by their intended purposes: stress relief; microstructural homogenisation and grain refinement; and precipitation hardening. The first of these, stress relief, involves components being placed for long durations in furnaces at relatively low temperatures. Residual stresses (examined further in Section 2.4.4) are reduced, without any drastic alteration to the microstructure.

Gruber et al. [201], experimenting with IN-718, demonstrated the effectiveness of higher temperature stress relief treatments, measured via the bending of a rectangular component after being cut from the base plate. The component's long, narrow geometry (60mm, 12mm, 3mm in X, Y, Z, respectively) caused the tip of the as-built part to displace by 3.2mm, after it was cut from the base plate. Heating to 1150°C for 6 hours before cutting, however, completely eliminated this displacement.

The second category, microstructural homogenisation and grain refinement, involves heating close to the recrystallization temperature; approximately half the solidus temperature. At these temperatures new grains begin to nucleate and grow, eventually entirely replacing the previous microstructure with a finer and more uniform grain distribution. The process is usually referred to as 'annealing', although when applied to steels, or when the cooling is rapid, the term 'normalising' may be used.

The third category, precipitation hardening, could be considered a subset of the second: heat treatments intended to alter the microstructure, specifically by allowing the nucleation and growth of phases that add to the mechanical properties. The process usually follows two steps. First components are 'solutionised', essentially an annealing step, where solutes are dissolved into the matrix to form a uniform microstructure. Components are then 'aged' at a lower temperature,

allowing precipitates to form and grow. Cooling, both in between the steps and after aging, must be rapid (quenching) to ensure the microstructure is retained.

Zhou et al. [202] designed such a two-step program for IN-718. Samples were first solutionised at 720°C for 8 hours, then aged at 620°C for 8 hours. These were compared to ‘direct aged’ samples, which were aged at 620°C for 8 hours without the annealing steps. While precipitation of ‘ and “ was evident in both sample sets, those that had been solutionised showed more uniform distributions. Differences in hardness increases were minimal (296 HV, 467 HV, and 477 HV, in the as-built, direct aged, and solutionised-aged samples, respectively).

As mentioned previously, another entire dimension exists to the application of heat treatments: coupling heat with high pressure. This process, known as Hot Isostatic Pressing (HIP), will be discussed further in Section 2.4.4.

2.4.4 - Mechanical Defects

Mechanical defects are those which have a direct effect on the mechanical properties of the component, and are usually (though not always) visible under low-magnification optical microscopy. Many of these defects originate as a consequence of defects listed previously, especially microstructural defects. Mechanical defects fall into the categories of: residual stress, cracking, delamination, spatter, and porosity.

2.4.4.1 - Residual Stress

Residual stresses are those which remain within a component, after it has been removed from the build chamber and has reached ambient temperature. All residual stresses in L-PBF are fundamentally caused by fact that whenever material undergoes thermal expansion and contraction, it is going to be constrained by adjacent material. This process occurs both on the scale of the layer (i.e. the contraction of a newly-printed layer being inhibited by previous layers), and of the melt track (i.e. the contraction of a newly-solidified vector being inhibited by adjacent vectors) [203].

In their review paper, Li at al. [204] examined the complex relationships between build parameters and residual stress. They found that: depending on the alloy, residual stress can either be increased or decreased by increasing inter-layer dwell time; laser power and scanning speed also very significantly affect residual stress, with greater line energy densities generally leading to higher stress; orienting a build upon the baseplate in a manner which maximises heat dissipation via conduction reduces residual stresses; preheating the build plate reduces thermal gradients, and therefore reduces thermal stresses; and scan strategies that minimise heat accumulation, such as a chessboard pattern, reduce thermal stress, albeit with the unintended consequence of an increased risk of cracking.

Evans et al. [205], working with IN-718, further explored the effects of scan strategy on residual stress. It was found that using short scan lengths and larger hatch distances allowed for a more even application of heat onto the printed layer, without an excessive accumulation concentrated at any point. Longitudinal stresses (parallel to the scan vector) were found to be the main contributor to residual stresses, and both longitudinal and transverse stresses were found to concentrate at the ends of the vectors.

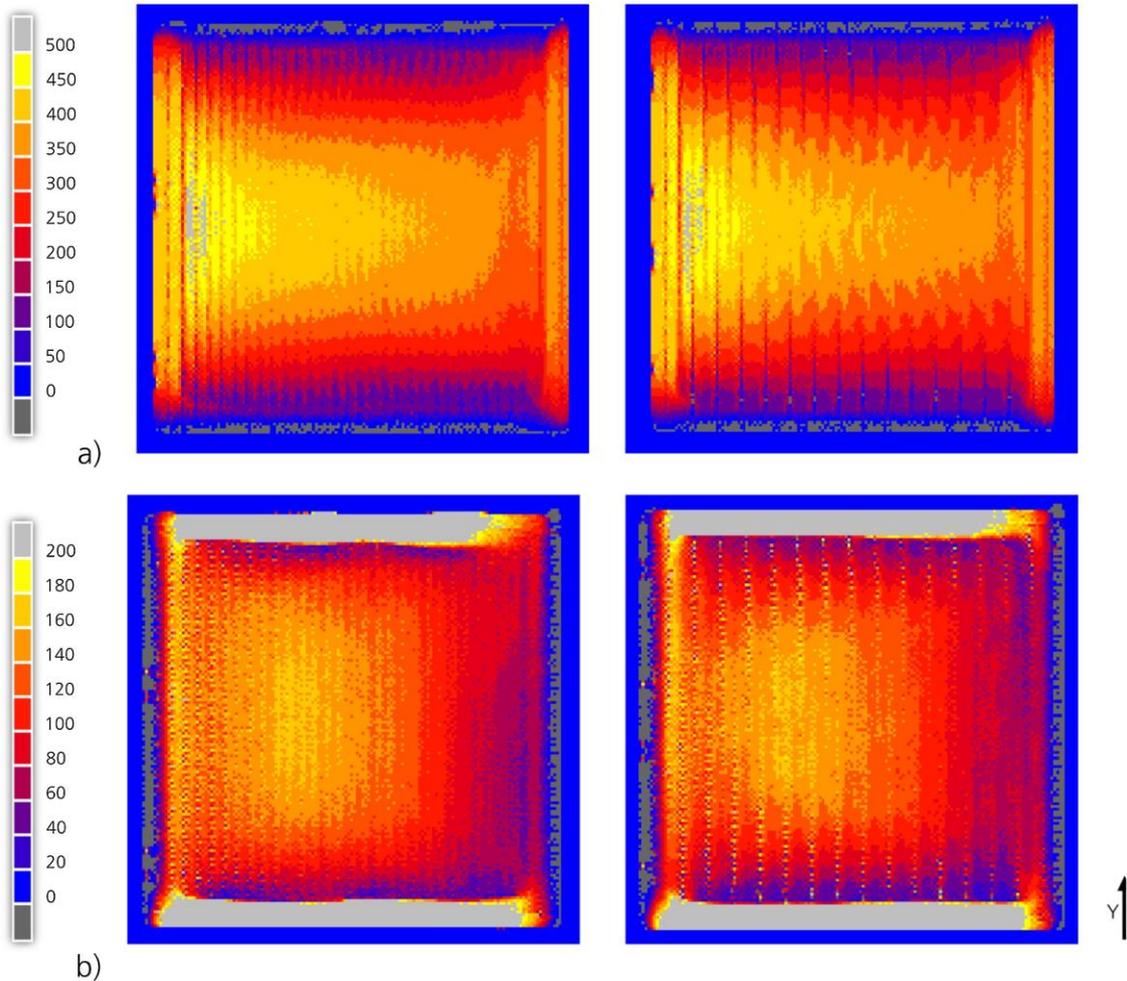


Figure 2.43

An FE model that accompanied Evans' experimental work, showing that normal stresses σ_{yy} (a) and σ_{xx} (b) are both more slightly more uniform when using unidirectional (left) scan strategies, rather than alternating (right) [205]

As pointed out by Bartlett et al. [206], solid-state phase transformations can significantly affect residual stresses. The rapid cooling of steels often results in a transform from austenite to martensite, accompanied by a 4% volumetric expansion, which counteracts thermal contraction to reduce the formation of residual stress. Similarly, the β to α martensitic transformation also incorporates a volumetric expansion, which can counteract thermal contraction. It must be noted, however, that these volumetric expansions are very small compared to the thermal contractions, and each martensitic transformation will exhibit its own inherent stresses.

The negative effects of excessive residual stresses are manifold. Component fatigue strength can be reduced, as sections of the material experiencing tensile strength are more susceptible to crack formation and propagation. The uneven distribution of tensile and compressive residual stresses within the build can also cause an anisotropy of mechanical properties. Geometric distortions are possible, both during the build and after the part is removed from the base plate. Using a base plate with both high rigidity, and high thermal conductivity, can reduce residual stresses [207].

2.4.4.2 - Cracking

Cracking refers to the formation of long, narrow voids within the component, as a result of the alloy being pulled apart under stress. Cracking can occur at various points within the L-PBF process, due to stresses originating from a variety of sources. Due to their morphology, cracks can easily propagate further under minimal load. Therefore, cracking is the most serious defect a L-PBF component can exhibit, in terms of a reduction in mechanical properties.

Solidification cracking occurs when an alloy that exhibits a wide freezing range (temperature difference between solidus and liquidus), and low ductility in the mushy zone, rapidly solidifies under conditions that produce columnar or dendritic grains. In the final stages of solidification, after a robust network of columns and dendrites have formed and are undergoing thermal contraction, the liquid remaining in the inter-columnar and inter-dendritic spaces may not be sufficient to fill the volume. Crack-like voids will therefore form along the boundaries between grains, columns, or dendrites [208]. Solidification cracking is demonstrated in Figure 2.44.

As Zhang et al. [209] describe in their review paper, two methods exist to combat solidification cracking. Processing parameters can be optimised to encourage slower cooling or more equiaxed growth. Also, powder feedstock can be modified with additives or inoculants to narrow the freezing range, or promote homogenous nucleation ahead of the solidification front.

Liquation cracking is related concept, but occurs under slightly different conditions, at slightly different locations. Within the Heat Affected Zone (HAZ) that surrounds a melt pool, if there exists solid material with a heterogeneous microstructure and a wide freezing range, certain constituents may be partially melted. This partial melting is especially prevalent at grain or sub-grain boundaries, or in inter-dendritic regions where solute concentration may be high, and therefore melting temperatures may be low [210]. As with solidification cracking, the mechanism that here results in crack formation is a combination of thermal contraction and insufficient liquid to fill voids. Liquation cracking is demonstrated in Figure 2.45.

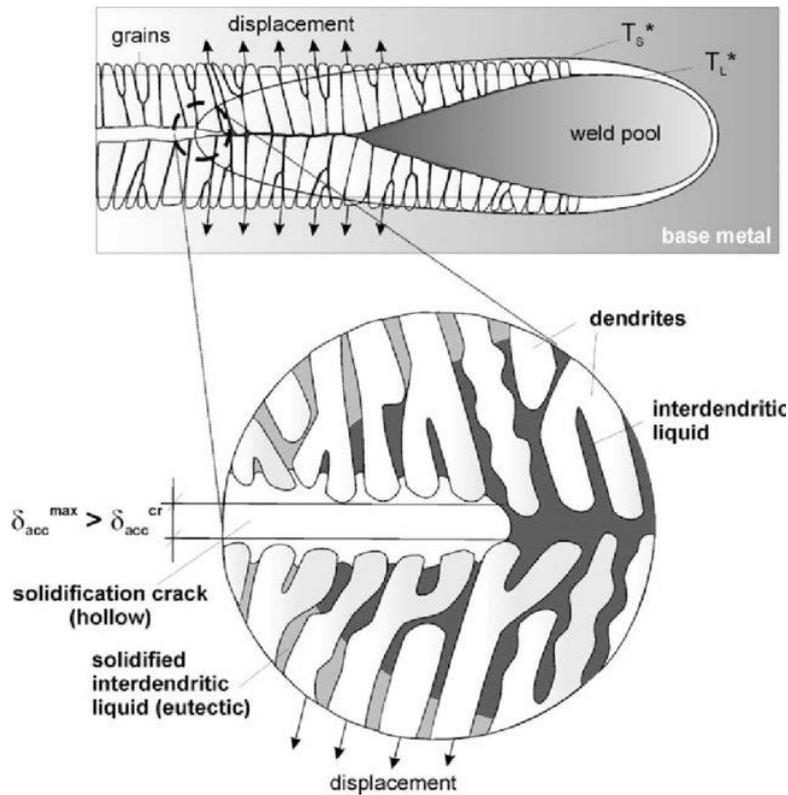


Figure 2.44

As grains attempt to thermally contract away from the melt pool centre, if the alloy's ultimate stress is overcome, solidification cracking will occur [211]

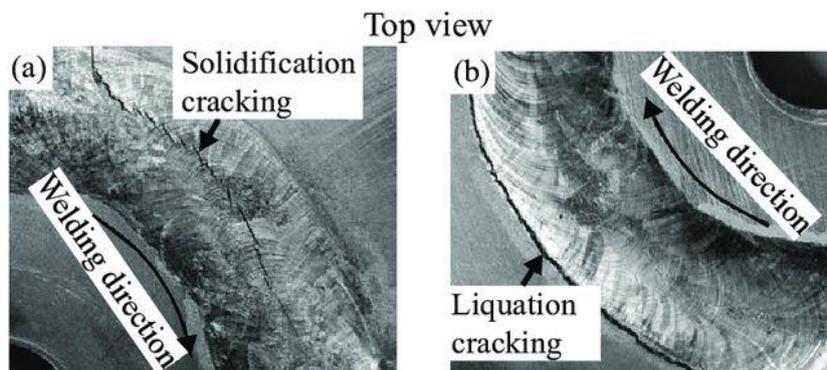


Figure 2.45

Two curved weld tracks: (a) shows central solidification cracking; (b) shows peripheral liquation cracking [212]

Liquation cracking is common in alloys containing elements which both segregate along boundaries, and can form a eutectic mixture (i.e. a composition which exhibits a much lower melting point than others in the alloy). In IN-718, for example, Nb solute becomes concentrated at grain boundaries, where it can form a NiNb eutectic, the lower melting temperature of which is susceptible to liquation. Similarly, in AISI 321 steel, a eutectic composition of TiC and austenite iron can cause liquation at grain boundaries [213]. Liquation cracking can be reduced via alloy design to reduce the quantity of segregating elements that form eutectic compositions.

Cold cracking also occurs due to thermal contraction, but in fully solidified material that is being cooled from solidus to ambient temperature. High residual stresses, especially those which exhibit significant heterogeneity throughout the component, can cause material to tear in tension. Building cubes of Mg-Gd-Zn-Zr alloy GZ151K, Deng et al. [214] demonstrated that cold cracking is attributed to the build-up of residual stress over the height of the build, until the tensile forces within the component exceed the alloy's ultimate tensile strength. Wei et al. [215], in cubes of Al-Fe-V-Si alloy, showed that cold cracking is most severe along melt pool boundaries, due to high concentrations of precipitates that exhibit different thermal expansion coefficients to the bulk. Methods of reducing cold cracking include: alloy design to reduce thermal contraction and increase ductility; reducing thermal gradients via scan strategy optimisation or beam shaping; and preheating of the base plate. Cold cracking is shown in Figure 2.46.

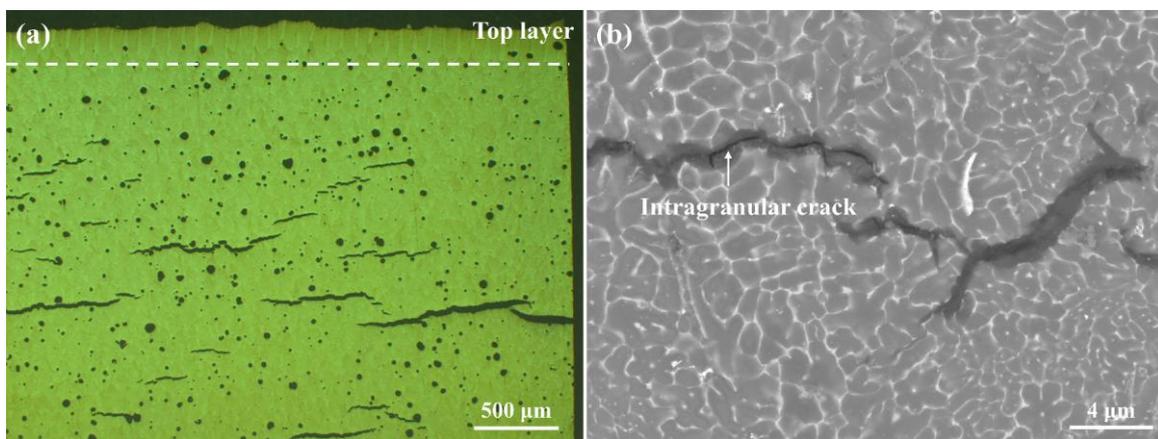


Figure 2.46

An XZ image (a) of an as-built GZ151K sample, and an XY subsurface image (b), both showing cold cracks that have propagated across many grains [214]

Ductility dip cracking is a phenomenon that occurs in some alloys, such as certain austenitic steels and nickel-based superalloys. A temperature-ductility graph of the alloys in question will show ductility increasing as temperature increases from ambient towards solidus temperature (T_s), at which point ductility falls precipitously as expected. A deviation to this expected curve, however, exists in the form of a sharp 'dip' in ductility, between approx. $0.5 T_s$ and approx. $0.8 T_s$. Proposed mechanisms for this dip in ductility include grain boundary sliding, precipitate-induced strain, and solute segregation, yet fundamental gaps in understanding still exist [216]. As solidified and cooling material passes through this dip in ductility, the stress it is experiencing will suddenly cause cracking. Ductility dip cracking is shown in Figure 2.47.

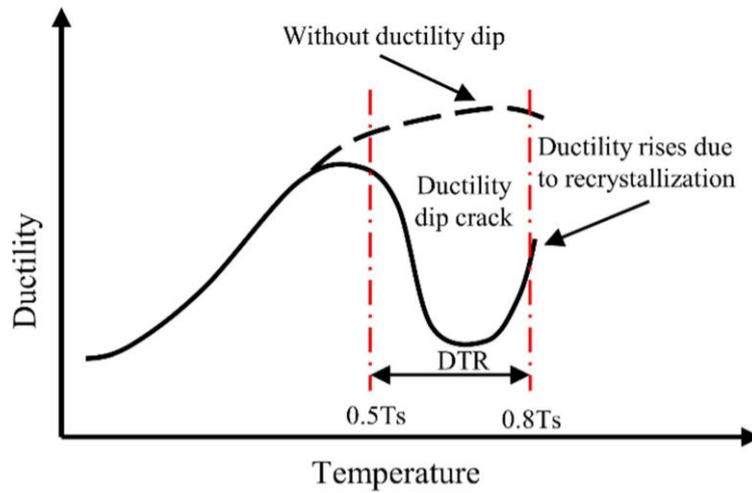


Figure 2.47

The ductility dip between $0.5T_s$ and $0.8T_s$ of a Nickel-based alloy. The rise in ductility due to recrystallization at higher temperatures is another proposed mechanism. [217]

2.4.4.3 - Delamination

Delamination refers to the splitting of a L-PBF component along its build layers, due to a combination of high residual stresses, and poor inter-layer bonding, as shown in Figure 2.48. Zhang et al. [218] demonstrated that it is the latter of these which is more crucial in determining whether a part will delaminate. The authors, building cubes of Ni-rich NiTi, demonstrated delamination occurring in low energy density regions (laser powers below 100W), with large hatch spacings (above $100\mu\text{m}$). These parameters resulted in sufficient lack of fusion porosity (a defect discussed in detail in this section), for the cubes to exhibit a tensile strength in their build direction that was overcome by residual stresses.

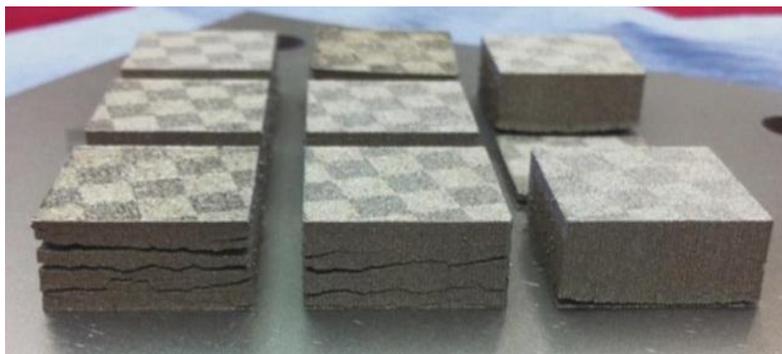


Figure 2.48

The delamination of Ti-6Al-4V samples due to poor inter-layer bonding [219]

Poor inter-layer bonding can also be caused by oxidation. The presence of an oxide film upon the upper surface of each deposited layer can act as a barrier between the metallic atoms of subsequent layers. Should the oxide film be remelted and dissolved as each new track is scanned, increased delamination can still occur, as the presence of oxygen in a melt pool is known to reduce wettability with the underlying layer [220].

As demonstrated by Horn et al. [25], contact with the recoater blade can amplify delamination issues. Building very thin-walled specimens of IN-718, the authors showed that, even for recoaters that use soft carbon fibre brushes, damage from contact was sufficient to dislocate newly-printed layers. The effect was even more pronounced for recoaters that use polymer lips, and more so for steel blades.

2.4.4.4 - Spatter

Spatter refers to the ejection of molten material from the melt pool, which solidifies and falls elsewhere within the build chamber. Though arguably not a defect by itself, spatter can be the cause of significant defects in each category explored in this section. Four related mechanisms explaining this high-velocity material expulsion are found in the literature: Each concerns a different melting regime or set of regimes. The first mechanism occurs as the melt pool transitions from conduction to keyhole regime; the second occurs at all points in the keyhole regime; the third occurs while the melt pool is in keyhole instability; and the fourth occurs at all melting regimes. Some authors have proposed multiple types of spatter, classified based on their mechanism, as shown in Figure 2.49.

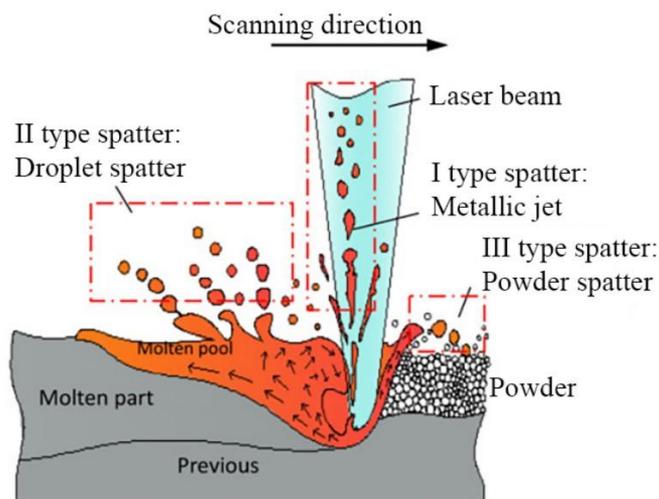


Figure 2.49

A diagram of spatter by Kivirasi, who proposes three classifications of spatter [221]

As demonstrated by King et al. [222] building in 316L steel, there exist certain machine parameters for which the melt pool's regime will transition between conduction and keyhole, stochastically alternating between the two depending on local thermal conditions and powder particle morphology. At the moment that the recoil pressure upon the melt pool overcomes the surface tension, and the vapour depression extends deep below the sides, the molten material that was previously in the pool's centre is displaced. This liquid metal can be ejected out of the pool, and across the powder bed.

During keyhole melting, the ejection of high-velocity vaporised particles from the melt pool results in a high-temperature vapour plume emitting from its surface. The inert gas within the build

chamber can be caught within this flow, creating strong convective gaseous currents surrounding the melt pool [223]. Any material swept into this flow, either powder particles or liquid breaking off the melt pool surface, will be melted and carried high away from the pool, and ejected out across the powder bed.

Huang et al. [224] distinguish between three types of sub-regimes within the keyhole melt mode: quasi-stable (in which the keyhole retains a relatively stable geometry), unstable (in which the keyhole depth dramatically fluctuates due to regular collapse), and a transition region in between these two. It was found that in the latter two of these sub-regimes, the forcing of molten material out of the centre of the melt pool caused an ejection across the powder bed, in a similar manner to that described for conduction-keyhole transition spatter described previously. Huang used X-ray Computed Tomography (XCT) to analyse the sub-layer formation of porosity. These results are shown in Figure 2.50.

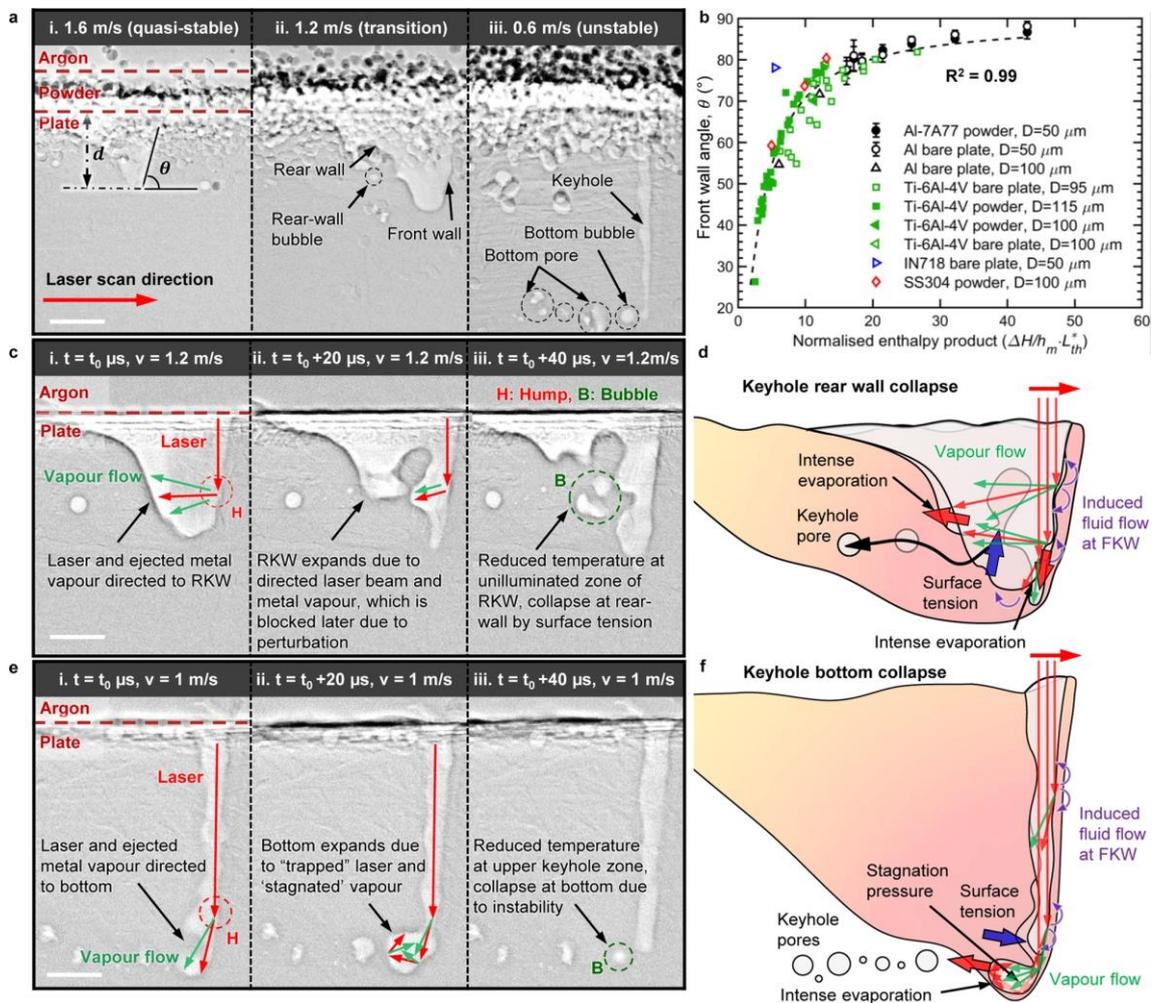


Figure 2.50

From Huang: (a) XCT images of quasi-stable, transition, and unstable keyholes; (b) the relationship between normalised enthalpy (will be discussed later in this chapter) and the keyhole front wall angle; (c) XCT images of a transition keyhole, on a powder-free aluminium plate, undergoing rear keyhole wall (RKW) collapse; (d) illustration of (c); (e) XCT images of an unstable keyhole, on a powder-free aluminium plate, undergoing keyhole bottom collapse; (f) illustration of (e) [224]

Spatter particles, after being ejected from the melt pool, rapidly solidify and are transported away by both their own inertia, and by the flow of inert gas. Though less significant than the three mechanisms listed previously, spatter can also be produced via the mechanism of the flow of this inert gas [225]. Molten material can be carried away from the melt pool, or solid powder particles can be carried into the path of the laser and vapour plume, by the inert gas flow.

No matter which mechanism is responsible for the creation of spatter, its effects can be detrimental to build quality. Some spatter particles may land upon the powder bed, and lie on top of the powder in the path of vectors yet to be scanned in the current layer. Others may land upon solidified tracks, get covered in powder during recoating, and lie beneath and amongst the powder in the path of vectors to be scanned in the next layer. In either case, the inclusion of these particles presents a radical departure from the usual laser-powder interaction for this the machine has been optimised. The energy required to remelt spatter particles, and incorporate them into the melt pool, is significantly larger than the energy required to melt and incorporate standard powder particles. Therefore, spatter upon the powder bed is a noteworthy source of lack of fusion porosity.

Experimenting with Hastelloy X, Esmailizadeh et al. [226] demonstrated a contrast between the size of their powder particles (generally smaller than $45\mu\text{m}$), and the resultant spatter particles (some exceeding $100\mu\text{m}$). It was found that the adhering of this spatter to outer surfaces greatly impart roughness, with R_a increasing from $14.4\mu\text{m}$ to $28\mu\text{m}$ on the part's upper surface.

Spatter particles that protrude significantly from the powder bed before recoating will cause one or both of the following detrimental effects: if the particle is strongly bonded to the part's upper surface, having been adhered to a melt pool, it may damage the recoater blade; if it is either not adhered, or has been broken free by the blade, it will be swept across the bed leaving a deep trough in its wake [227]. In either case, the effect on the next or on all subsequent layers may be significant. These troughs are demonstrated in Figure 2.51.

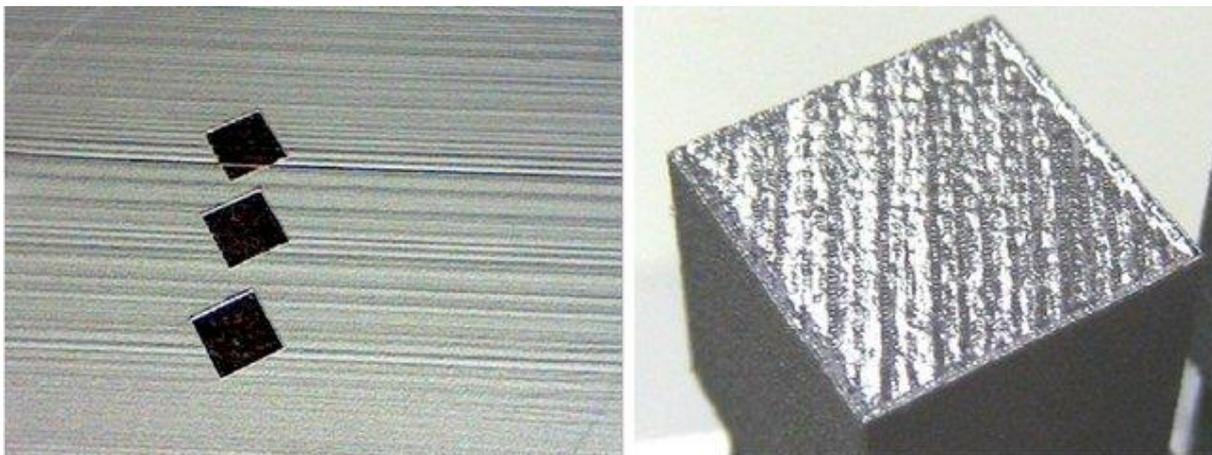


Figure 2.51

Mounds, in-line with parts, due to spatter adhered to part surface causing blade damage; and trenches, away from parts, due to loose spatter on bed being swept by blade. Note that upon visual inspection, any damage to the part's final layer is not immediately obvious. [228]

As discussed previously, the mechanisms of spatter formation largely revolve around the regime in which melting occurs. If parameters can be selected that ensure as many vectors as possible remain in either conduction or stable keyhole regime, then spatter formation can be avoided.

2.4.4.5 - Porosity

Porosity refers to the inclusion of voids within a component, filled with gas or powder. These voids can be divided into three distinct categories, based on their formation mechanisms. These are Lack of Fusion, Keyhole, and Gas pores. These categories are shown in Figure 2.52.

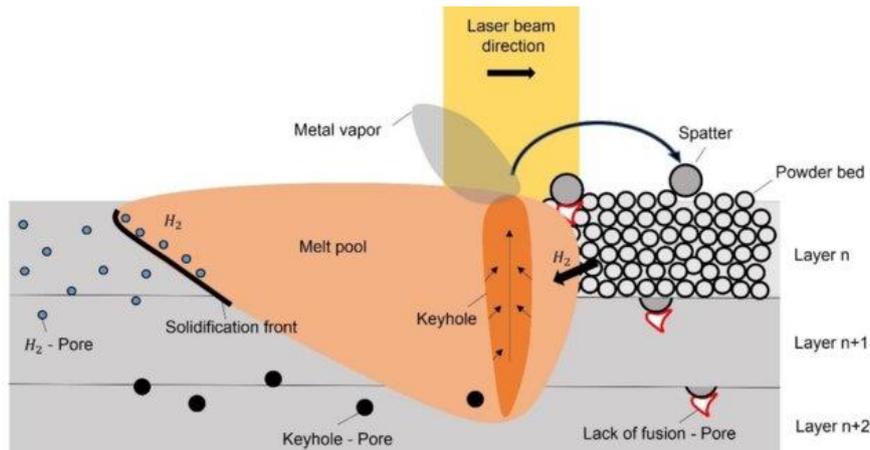


Figure 2.52

The three major types of porosity: keyhole, lack of fusion, and gas (here labelled H_2) [229]

2.4.4.5.1 - Lack of Fusion Porosity

Lack of Fusion (LoF) pores are formed when the laser energy is insufficient to melt all of the powder particles in a scan layer, as demonstrated in Figure 2.53. This error will manifest as insufficient overlap between melt pools, resulting in long, irregular-shaped volumes between vectors, filled with powder particles, partially melted powder particles, and the gas that existed amongst these particles after they were spread across the build chamber.

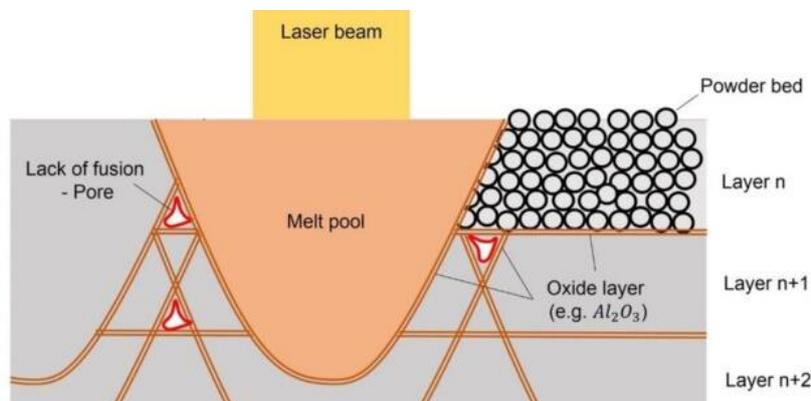


Figure 2.53

A schematic showing how some regions of the powder bed (red), can be missed by all the melt tracks (outlines shown in orange) when hatch spacing is too large [229]

LoF can be quite accurately predicted by basic machine parameters alone. Tang et al. [230] developed a model that predicted the width and depth of a melt pool with the Rosenthal equation, which calculates the temperature at a given point around a moving heat source, based on simple parameters like the source's power and speed, and the substrate's absorptivity and thermal conductivity. Combining these predicted dimensions with layer heights and hatch spacings, it is possible to estimate the extent of melt track overlap (remelting) and/or insufficient melting (LoF) a build will experience. The authors searched existing literature for reports of experiments in which all necessary parameters were given, and found that their model was able to predict the LoF porosity in these experiments with high accuracy.

Whether or not a given scan vector will undergo sufficient melting to prevent LoF porosity can also depend on local powder particle conditions. Wang et al. [231] incorporated these more advanced properties in their model, taking into account the powder particle size distribution, the particle packing pattern (i.e. the amount of gaps amongst the particles), and the difference in conduction conditions between solid and powder. The authors validated their model against builds of Ti-6Al-4V, and found that their predicted values of part porosity fell within 6% of the experimentally measured values.

As mentioned previously, LoF porosity can also be caused by the presence of spatter particles upon the powder bed. These particles are usually significantly larger than the powder, and have developed surface oxides which alter their laser-material interaction properties. Building in IN-718, Snow et al. [232] tracked the locations that spatter particles fell using in-situ monitoring, and later determined the locations of LoF pores using X-ray Computed Tomography. They found that a causal relationship between the two could be determined, with both the likelihood of, and the size of, a LoF pore increased as the size of the spatter particle increased. For spatter particles with an equivalent spherical diameter of above approximately 770 μm , the probability of it causing a LoF pore reached 100%.

Due to their irregular shape and sharp corners, LoF pores can act as stress concentrators and crack initiation sites, significantly reducing both the tensile and fatigue strengths of components. They are arguably the most detrimental of all possible defects upon the mechanical properties of a L-PBF part. Working in 316L steel, Zhu et al. [233] found that an increase in porosity from 0.87% to 2.81%, led to a decrease in yield strength and fatigue strength of 20.14% and 38.35% respectively. These results are shown in Figure 2.54.

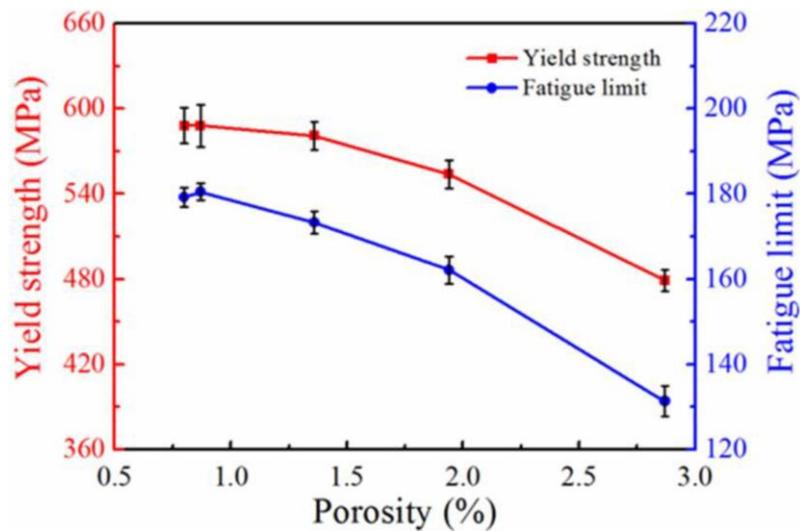


Figure 2.54

How porosity affected mechanical properties from Zhu [233]

As noted by Aliyu et al. [234] it is the LoF pores that make contact with the outer surface that are the main initiation points for ruptures that cause part failure. Therefore, the effects of LoF porosity could be significantly mitigated by scanning outer contours around each layer, to promote a defect-free surface.

Since LoF pores tend to align with the laser tracks, they act as a notable source of mechanical anisotropy. The effects of this anisotropy are not entirely intuitive, however. Ronneberg et al. [235], experimenting with 316L steel, showed that both the ductility and yield strength were lower in the Z (build direction) than the X or Y direction in components with LoF porosity. However, the authors found that once all microstructural anisotropy was removed via heat treatment, the yield anisotropy disappeared too, while the ductility anisotropy remained. It was concluded, therefore, that while LoF defects are the cause of ductility anisotropy, it is microstructural anisotropy alone that is responsible for yield anisotropy.

2.4.4.5.2 - Keyhole Porosity

As discussed in Section 2.2.2, keyhole pores are formed when the walls of a deep keyhole touch and rejoin, thereby sealing off the bottom section as a bubble of metal vapour within the liquid melt pool. Due to the rapid solidification conditions, this bubble will not have time to rise to the surface before the material around it freezes, thus trapping it within the component. As determined by Yu et al. [236], there are five interconnected factors that dictate the stability and fluctuations of keyholes:

- Recoil Pressure - The force generated by the vaporisation of liquid at the melt pool surface. As the keyhole shape changes, and the amount of laser energy reaching each area of the surface changes, this force can vary rapidly.
- Surface Tension - The force that attempts to minimise the melt pool's liquid surface area. It shows an inverse relationship with recoil pressure.

- Vapour Pressure - The force exerted on the walls of keyholes and keyhole pores by the vaporised material. Since the pressure of this gas depends on temperature, any fluctuations in the amount of heat absorbed within the keyhole can lead to large changes in the forces applies on the keyhole walls.
- Liquid Vortex - It is found that the liquid around a keyhole pore spirals as a vortex. The lower part of this vortex plays a role in pinching the keyhole tip, encouraging a collapse.
- Keyhole Morphology - The shape and features of the keyhole, particularly protrusions along its walls, amplify fluctuations. The quantity of laser energy that's absorbed elsewhere inside the keyhole after initially being reflected, also depends largely on the vapour depression's morphology.

The likelihood of a keyhole becoming unstable and collapsing into a pore can largely be predicted by machine parameters and material properties. Laser power is found to be the most significant determinant, with the other parameters affecting volumetric energy density (i.e. laser speed, hatch spacing, layer height) also contributing [237]. A CFD model demonstration of keyhole porosity formation is shown in Figure 2.55.

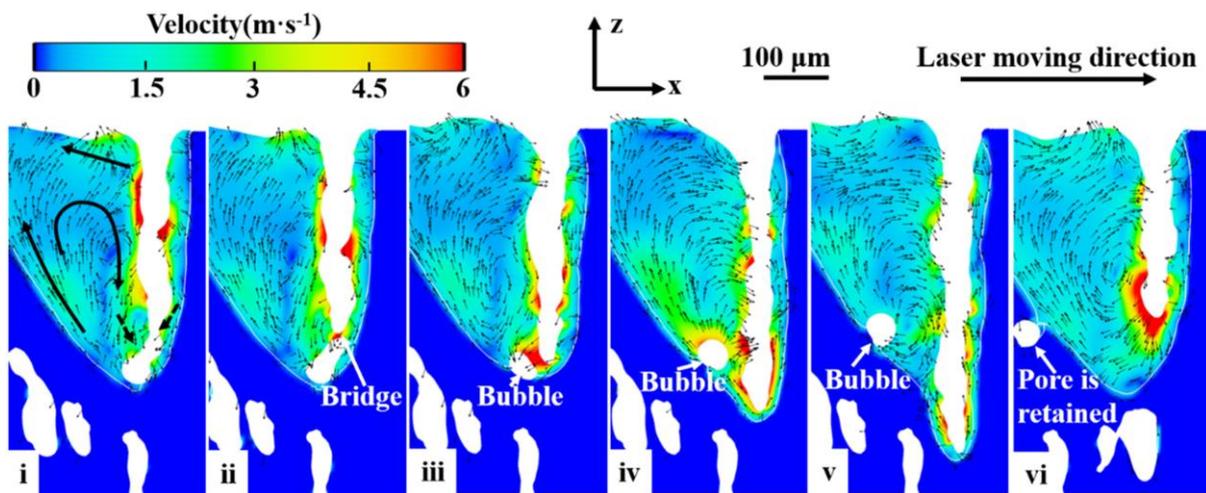


Figure 2.55

A CFD model, showing how when keyhole walls touch (bridge), a bubble gets separated from the main, and can be retained in the solid [238]

Gan et al. [239] developed the universally-applicable dimensionless Keyhole Number, which describes the aspect ratio and stability of a keyhole.

$$Ke = \frac{AP}{(T_l - T_o)\pi\rho c\sqrt{\alpha Vr^3}}$$

where A is absorptivity, P is laser power [J/s], T_l is liquidus temperature [K], T_o is substrate temperature [K], ρ is density [kg/m³], c is specific heat capacity [J/kg/K], α is thermal diffusivity [m²/s], V is laser speed [m/s], and r is laser radius[m].

Building in 6061 Aluminium and 316 steel, it was found that melt pools existed in conduction mode where $Ke < 1.4$, in transition mode where $1.4 \leq Ke \leq 6.0$, and in keyhole mode where $6.0 < Ke$. As the Kehole Number further increased, the stability of the keyhole deteriorates. Where $6.0 < Ke < 16$, there was very little fluctuation and almost no porosity or spattering. Where $30 < Ke$, the keyholes were extremely unstable, with large porosity and spattering occurring.

Due to their spherical nature, keyhole pores are far less damaging to a component's mechanical properties than LoF pores. Montalbano et al. [240], experimenting with Ti-6Al-4V, measured the elastic modulus and tensile strength of samples with varying fractions and types of porosity. They found that the decreases in these mechanical properties in the samples with keyhole pores, could be fully explained by the decreases in cross-sectional area due to the porosity levels: i.e. the presence of keyhole pores does not negatively affect the properties of the solid material surrounding it. In the case of the samples with LoF porosity, the decreases in elastic modulus and tensile strength was more extreme, and could not be accounted for by cross-sectional area reduction.

Reviewing the literature on four materials commonly produced via L-PBF (Ti-6Al-4V, IN-718, 316L steel, and AlSi10Mg), Kan et al. [241] found that keyhole porosity also tends to not negatively impact fracture toughness. When it comes to fatigue life, however, keyhole pores were identified as critical sub-surface defects responsible for early failure. The authors found that in terms of creep and wear behaviour, while porosity is known to be detrimental to both properties, further research is required to differentiate the effects of keyhole and LoF pores.

2.4.4.5.3 - Gas Porosity

It is inevitable that some amount of gas will be churned into the liquid melt pool. Some of this gas may originate from within the powder particles, entrapped during its production process. More will originate from the inert gas found between powder particles in the powder layer. The melting and evaporation of moisture or oxide layers upon the powder particles can also contribute hydrogen and other elements that exist as gas within the liquid [242]. Gas porosity, as well as several other types of pores and defects, is demonstrated in Figure 2.56.

While much of this gas will be ejected from the melt pool due to buoyancy forces, some will remain trapped, due to the chaotic convective and Marangoni forces within the pool, and the rapid rate of solidification. The gas that gets frozen into the component will exist as very small spherical pores, which do not significantly affect to the mechanical properties. In fact, it is found that gas pores are only even noticeable in near-fully dense samples, due to them being masked by the predominance of LoF and keyhole pores in more porous samples [243].

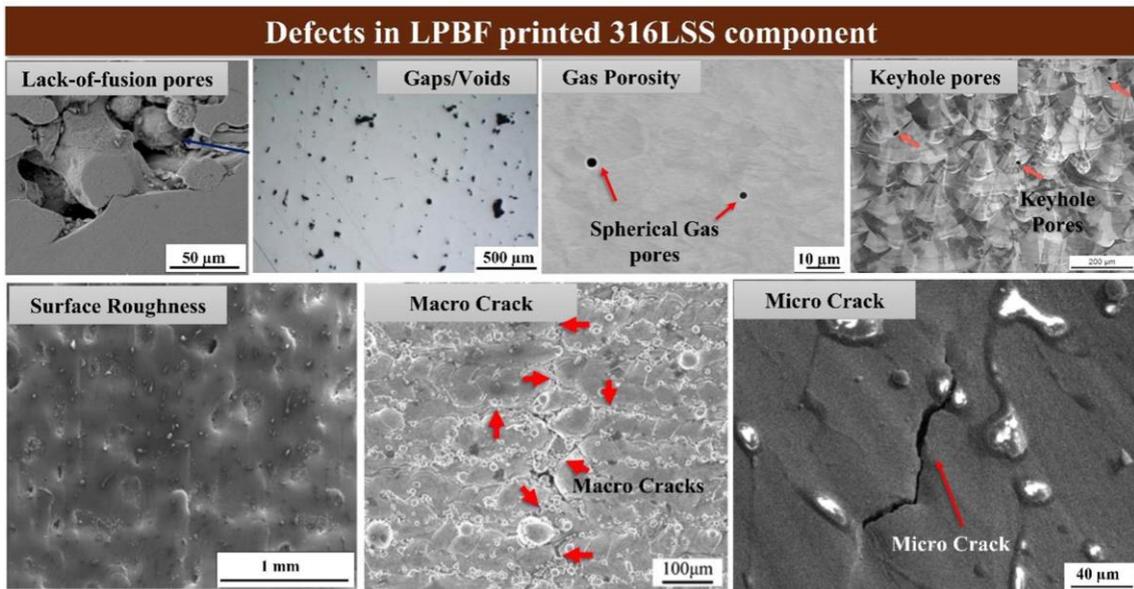


Figure 2.56

SEM images showing a range of pores, cracks, and other defects in 316L steel [244]

2.4.4.6 - Post Processing Mechanical Defects

As mentioned previously, HIP is a very effective post processing technique. Furnaces which incorporate high pressure inert gas allow a heated component to experience uniform pressure in all directions, normal to all surfaces (i.e. isostatically). This method, Hot Isostatic Pressing (HIP), typically exposes L-PBF components to temperatures in the range of 1200°C, and pressures in the range of 100 MPa, for several hours. Subjected to these external forces, the volumes of pores, cracks, and other such internal defects can be significantly reduced as the component is very-slightly compacted, undergoing plastic deformation. It must be noted, however, that any defect which makes contact with an external surface will be unaffected. Coating or encapsulating components before HIP can effectively transform external-touching defects into fully-internal defects, counteracting this issue [245].

Kluczyński et al. [246], working with 316L steel, applied 1150°C and 100 MPa for 3 hours, and achieved porosity reductions of 65% and 78%, when sectioned samples were viewed in the YZ and XY planes respectively, as compared to the as-built samples. It was found that the mechanical properties of samples that were processed with low enough energy density to cause LoF porosity, then underwent HIP, were similar to the properties of those that were built fully dense. XCT images showing the before-and-after are shown in Figure 2.57.

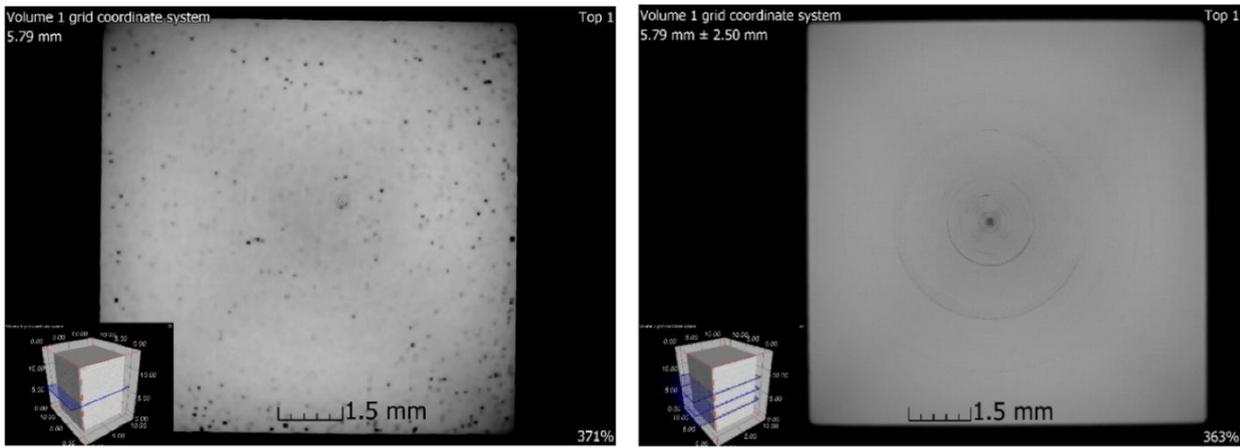


Figure 2.57

XCT images of the XY planes from Kluczyński's 316L steel samples, before (top) and after (bottom) heat treatment [246]

Experimenting with IN-718, Tillmann et al. [247] identified 1150°C, 100 MPa and 4 hours as an ideal parameter set for HIP. As-built density was measured as 99.87% - 99.91%, while after HIP the density rose as high as 99.998%. The authors noted that, in cases where inert gas has entered the pores of the component, achieving 100% density is effectively impossible, due to the increase of pressure exerted internally by this high-temperature gas.

2.5 - Factors Affecting the Build

This section provides a comprehensive overview of the decisions and considerations that a designer or technician must make, when executing an L-PBF build. They will be grouped into the following categories: the component geometry; the laser parameters used during building; the supports and anchors built alongside the component; the strategies to manage thermal and residual stresses during the build; the type of powder used; and the setup of the machine and build chamber.

Due to the interdependent nature of this analysis, where every factor affects every other factor, constraints are required to keep the size of this section manageable. Firstly, only the alloys being experimented upon in this project will be considered: IN-718 and ABD-900AM. Secondly, since the experimental sections of this project are focused on interpreting data from in-situ monitoring, and therefore will not involve any post processing methods, these will not be included in this review section.

2.5.1 - Part Geometry

Components fabricated via L-PBF have the potential to showcase detailed intricacy that would be difficult to achieve with traditional manufacturing methods. At the same time, careful consideration must be paid to the geometries that a machine is being asked to produce. Attempting to produce fine or unusual features without proper planning will likely result in either defects or an outright failed build.

2.5.1.1 - Thin Walls

Herzog et al. [248] found that the minimum thickness with which a thin wall section can successfully build depends upon the wall's length, and whether the wall was supported at its ends. As the length of unsupported sections increased from 10mm to 25mm, the minimum thickness with which they could be fabricated fell from 300 μm to 200 μm . In all cases where walls were too thin to be fabricated, the first layer failed to bond to the baseplate, and all subsequent layers had nothing to bond to.

Wu et al. [249] managed to achieve even thinner results, by building walls that were each just one scan vector thick. It was found that the minimum thickness limit is determined by the width of the minimum melt pool that does not result in severe LoF porosity or delamination between layers. For IN-718, this minimum was found to be $107 \pm 10\mu\text{m}$. These results are shown in Figure 2.58.

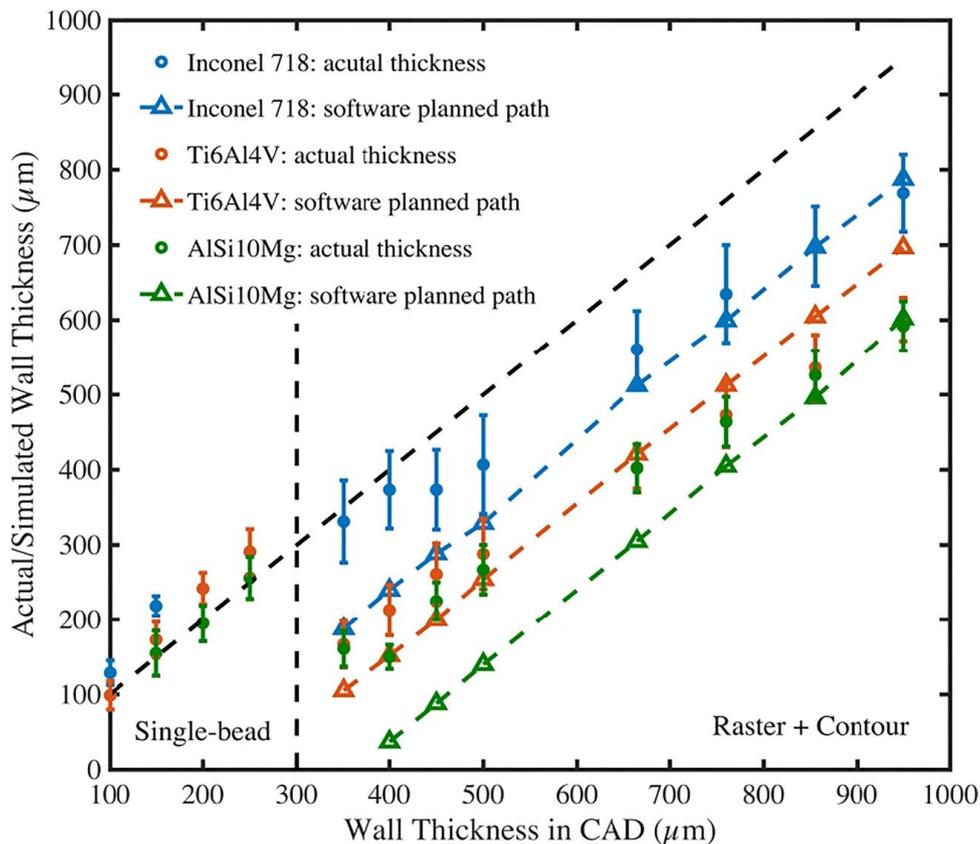


Figure 2.58

Wu's comparison between the CAD wall thickness and: the built thickness (circles), and the AM software's intended thickness (triangles) [249]

Pérez et al. [250] conducted a more application-minded investigation into the question of minimum wall thickness of IN-718. They aimed to find the minimum thickness of walls that a honeycomb structure could be fabricated with, without degrading its mechanical properties. It was found that, by using a contouring scan strategy, walls as thin as 150 μm could be manufactured without the compressive strength of the honeycomb being significantly compromised.

2.5.1.2 - Incline Walls

Fabricating inclined walls with a thickness of 3mm, a width of 20mm, and an intended height of 30mm, Herzog et al. found that overhang angles (i.e. between horizontal and the wall downskin) as low as 25° were achievable. However, as the angle decreases the roughness of the downskin surface increases exponentially, rising from an R_a of 75 μm at 80°, to an R_a of 400 μm at 25°. The upskin R_a consistently remained between 50 μm to 70 μm regardless of angle.

Zhou et al. [251] produced SEM images of downskin surfaces of IN-718 walls produced at various energy densities and overhang angles, to better understand the reasons for this inverse relationship between angle and roughness. They found that at lower angles, more of the downskin melt pool is unsupported by solid material below, and therefore more of it seeps into the powder under gravity and capillary action. This will in turn cause loose powder particles adhere to the downskin surface, with some particles agglomerating to form balls up to 600 μm in diameter, greatly increasing the surface's roughness. Zhou's 45° results are shown in Figure 2.59.

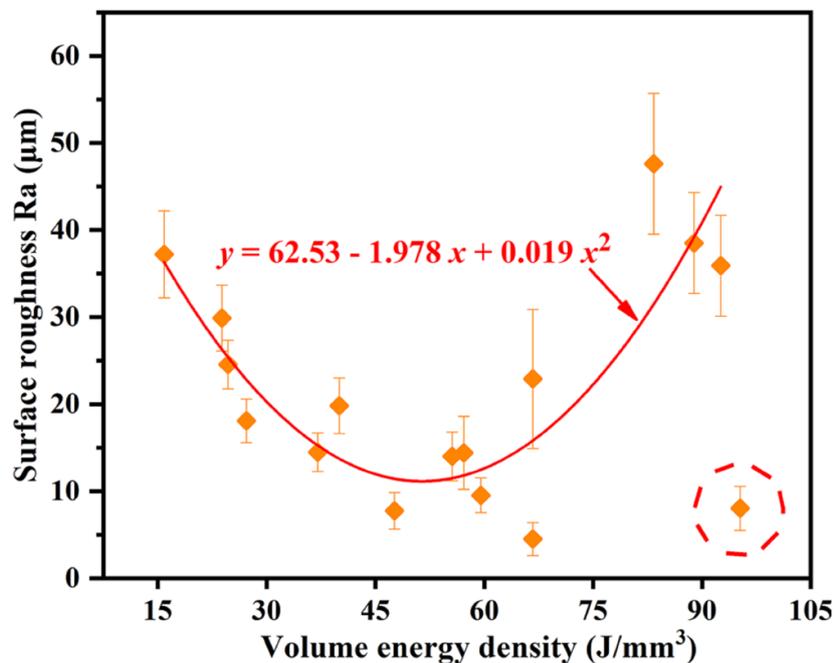


Figure 2.59

Zhou found that surface roughness is lowest at mid-range energy density. Results are here plotted for the 45° angle downskin. [251]

2.5.1.3 - Cantilevers

Fabricating a cantilever (i.e. a horizontal beam that is supported only on one side) is a difficult challenge in layer-wise manufacturing. Herzog et al. attempted to build 5mm wide cantilevers with various overhangs lengths. It was found that at lengths up to 1mm, the cantilever built as a horizontal beam as intended. At lengths above 1mm, however, the angle of the cantilever's underside began increasing from the intended 0°, towards a notable slope. It was concluded, therefore, that lengths greater than 1mm require additional supports.

It should be noted that this 1mm limit refers only to simple cantilever designs: a rectangular beam protruding into space. Different limits may be possible for different forms of horizontal features that are only supported at one point. For example, Sit et al. [252] designed a horizontal IN-718 disk, built upon a narrow rod, as a novel test piece for comparing horizontal downskin roughness of different laser parameters. They found that overhang lengths of 12mm were achievable without deformation or porosity, and lengths of up to 18mm were achievable before failure. This disk is shown in Figure 2.60.

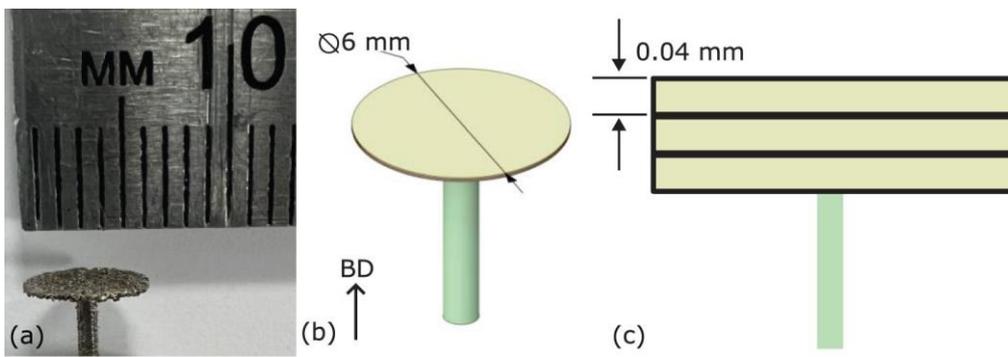


Figure 2.60

Sit's three-layer thick, disk cantilever, shown here at 6mm diameter [252]

2.5.1.4 - Horizontal Bridges

Horizontal beams supported at both ends are easier to build unsupported. Herzog et al. attempted to build bridges of 5mm width, and various lengths up to 5mm, between supporting columns. All lengths could be built, although above 3mm significant amounts of downskin material became lost, via the same mechanism that caused downskin angles to develop in the cantilevers. It was concluded that when a horizontal bridge spans more than 3mm, it should be supported.

This finding is in similar to that of Chen [253], who recommends 4mm as the maximum span of an unsupported bridge in IN-718, but lower than that of Barale [254], who recommends a maximum span of 6mm, as demonstrated in Figure 2.61.

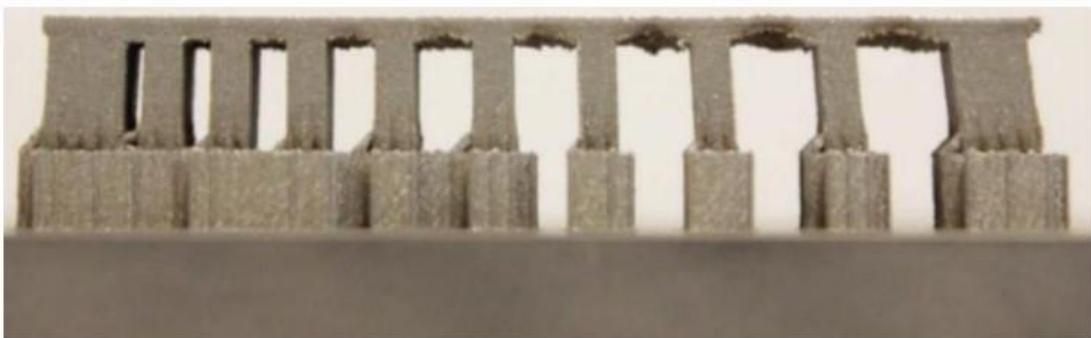


Figure 2.61

Barale's horizontal bridges. The fourth from the left spans 6mm. [254]

2.5.1.5 - Narrow Columns

Attempting to build vertical columns with varying diameters, Herzog et al. found that it was possible to fabricate columns of both 15mm and 30mm height, for diameters of 400 μm and greater. Below this diameter, no layers of the column managed to bond to the baseplate. It was noted that, due to their extreme fragility, all the columns with diameters between 400 μm and 800 μm suffered from bending due to the motion of the powder spreader. It was recommended, therefore, that vertical column should be built with diameters greater than 800 μm .

Znemah [255] also concluded that columns with diameters below 400 μm were impractical, due to the fact that scanning a circular path in which the outer edge of the melt track fits into a smaller space, becomes difficult for most laser parameters. Experimenting with very fine columns of IN-718, Du et al. [256] managed to build diameters of 150 μm , but found that the microstructural heterogeneity of the column's cross-sections were very significant, suggesting that as-built thin columns may be structurally weaker than expected due to anisotropy, and the mechanical properties of narrow sections cannot be extrapolated from larger sections. These results are shown in Figure 2.62.

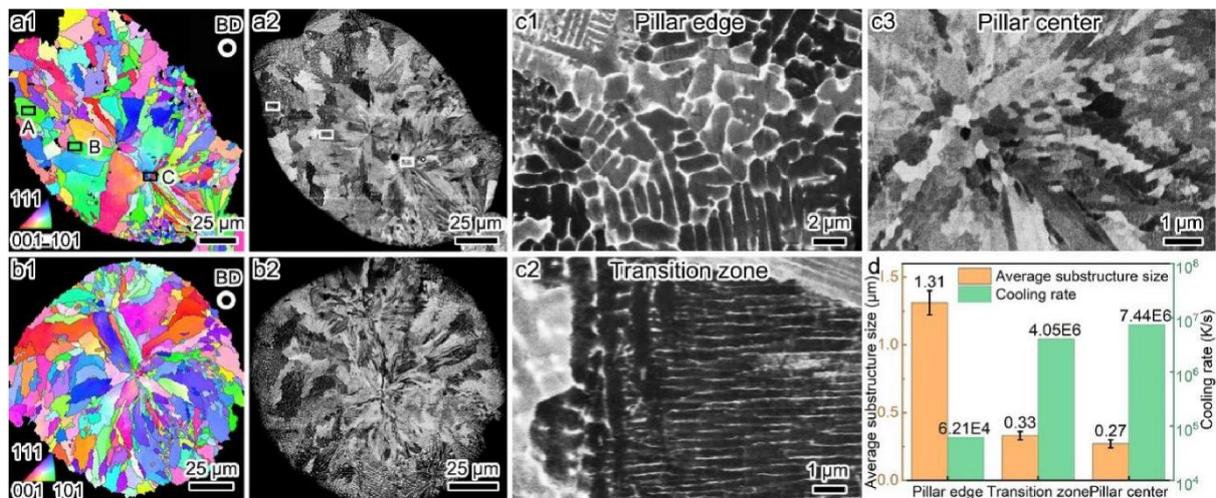


Figure 2.62

From Du: (a1, a2) EBSD and SEM images of a pillar built at 60W and 2500 mm/s; (b1, b2) EBSD and SEM images of a pillar built at 140W and 2500 mm/s; (c1, c2, c3) SEM images of points A, B, and C from (a1); (d) cooling rate and substructure size at points A, B, and C [256]

2.5.1.6 - Horizontal Holes

Fabricating spherical holes can be difficult or impossible, depending on the hole's diameter. At larger diameters, the unsupported downskin at the top of the hole will experience high roughness and material loss, as in the case of horizontal bridges. At very small diameters, the adhesion of powder particles to downskin melt pools may significantly block the hole.

Herzog et al. experimented with three hole geometries: perfectly circular, teardrop-shaped with downskin angles of 30° , and teardrop-shaped with 45° angle. It was found that, for the circular holes, there exists no diameter which is both narrow enough to not require supports, and wide enough to not be significantly affected by powder adhesion. The authors recommended that holes with a diameter of over 4mm could be built, but that supports were always required. These results are shown in Figure 2.63.

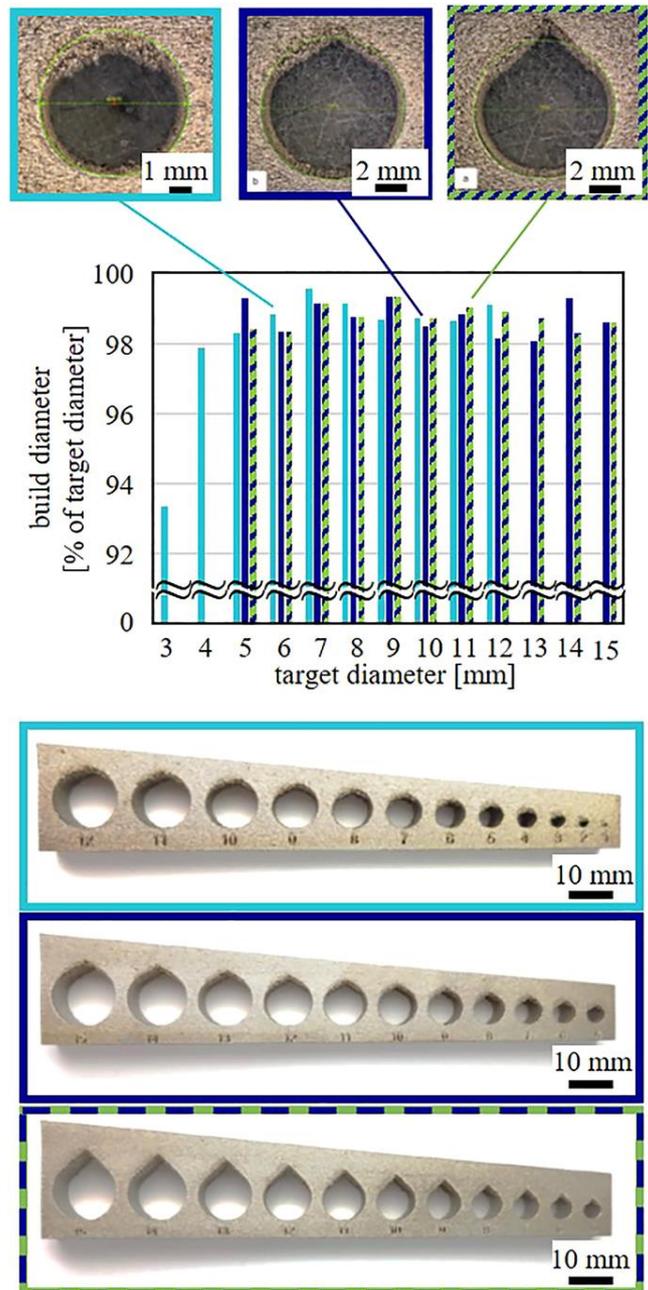


Figure 2.63

Herzog's three shapes of horizontal holes, and how well each matched its target diameter [248]

For the 30° holes, diameters from 1mm to 10mm were built without material loss at the apex. For diameters above this size, supports were recommended. For the 45° hole, all diameters above the size of 1mm built successfully, without any supports required.

2.5.1.7 - Non-Horizontal Channels

The ability of AM to produce components containing internal cooling channels is one of its main advantages. In the case of channels that are curved or at angles that would be difficult to drill, such components would be completely impossible to produce using conventional manufacturing methods.

In order to test this fundamental selling point of L-PBF, Kasperovich et al. [257] built a range of channel diameters (500 μm , 700 μm , and 1000 μm), at different geometries (circular and teardrop), across a range of angles (0°, 22.5°, 45°, and 90°), in IN-718. This experimental setup is shown in Figure 2.64.

The results of these tests matched the expectations that could be drawn from Herzog's tests discussed previously. As the intended diameter of the channels increased from 500 μm to 1000 μm , the deviation between the intended channel geometry and the built channel geometry decreased, as powder particle adhesion and blockage became less of an issue. Similarly, as the angle increased from 0° (horizontal) to 90° (vertical), the deviation decreased, as the channel roof became more supported by underlying solids. Also, for the same reason, teardrop channels outperformed circular channels.

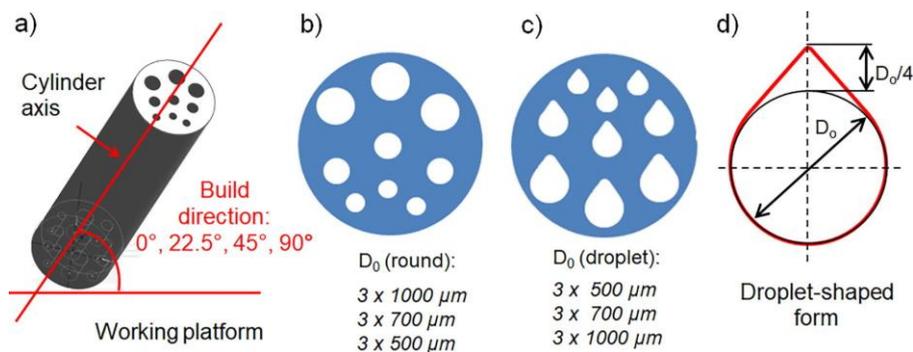


Figure 2.64

Kasperovich's two channel shapes, three channel diameters, and four channel angles [257]

For example, it was found that for both the 500 μm and 700 μm diameters, all circular channels built at 0° and 22.5° suffered from blockages. In the case of the 1000 μm diameter circular channel, as the angles increased from 0° to 22.5°, 45°, and 90°, the mean equivalent diameter of the built channels increased from 805 μm to 820 μm , 952 μm , and 988 μm . The authors noted that at 90°, the channels were devoid of obstructions.

2.5.1.8 - Build Orientation

The previous discussion of the ways in which channel angle can significantly affect build performance, naturally leads to the topic of ideal build orientation. When positioning a component upon a base plate, trade-offs must be made between minimisations of: downskin surfaces, build height, required supports, thermal stresses, staircase effects, and other factors.

Researchers have developed automated methods to suggest ideal orientations for given geometries. Qin et al. [258], for example, designed a system that estimated the results of a component being built at a certain angle, for five factors: support volume, volumetric error, surface roughness, build time, and build cost. Brika et al. [259] designed a similar system with five factors of consideration, but substituted volumetric error out, and mechanical properties in, specifically yield strength, ultimate tensile strength, elongation at failure, and Vickers hardness.

The ability of Birka's system to estimate the mechanical properties of various components at various orientations, relies upon the work of authors like Watring et al. [260], who studied the microstructural and mechanical effects of building IN-718 tensile test pieces at both 0° (i.e. longest and narrowest face of the piece is flush with the base plate) and 60° (i.e. smallest face of the piece is contacting the base plate via one edge). It was found that the 0° samples contained more porosity than the 60° samples (0.23% compared to 0.07%), and demonstrated a greater yield strength (798 MPa compared to 772 MPa).

The authors reasoned that this reduction in yield strength is due to the fact that grains in L-BPF of IN-718 tend to grow as elongated columns aligned with the build direction. Therefore, when a component is being loaded along an axis perpendicular to its build direction (as in the case of the horizontally-aligned tensile piece), the effective width of the component's grains is less than if the component were being loaded along its build direction.

2.5.1.9 - Build Height

Each new layer in a L-PBF build introduces another opportunity for delamination to occur, for swollen parts to damage or halt the powder spreader, or for the machine to fail due to a lack of powder or some other issue. Also, the time and therefore cost that a build requires, is generally proportional to its number of layers, as powder recoating in between layers often takes longer than each layer's scanning time. Therefore, components should be designed and oriented in a way to minimise the number of layers required.

It should also be understood, that the microstructural and therefore mechanical properties of components can differ along their build direction, due to the fact that as a build goes on, and the build chamber heats up, melt pools experience smaller temperature gradients and lower cooling rates during solidification.

Wang et al. [261] built dog-bone creep test pieces of IN-718, located 15mm, 52mm, and 87mm from the base plate. They found that as the build height increased, and therefore the accumulated heat within the chamber and powder rose, grains became smaller, more columnar, and with more low-angle grain boundaries. For these reasons, the creep rate of the test pieces increased as they were built further from the base plate. This progression of grain structure and build height is shown in Figure 2.65.

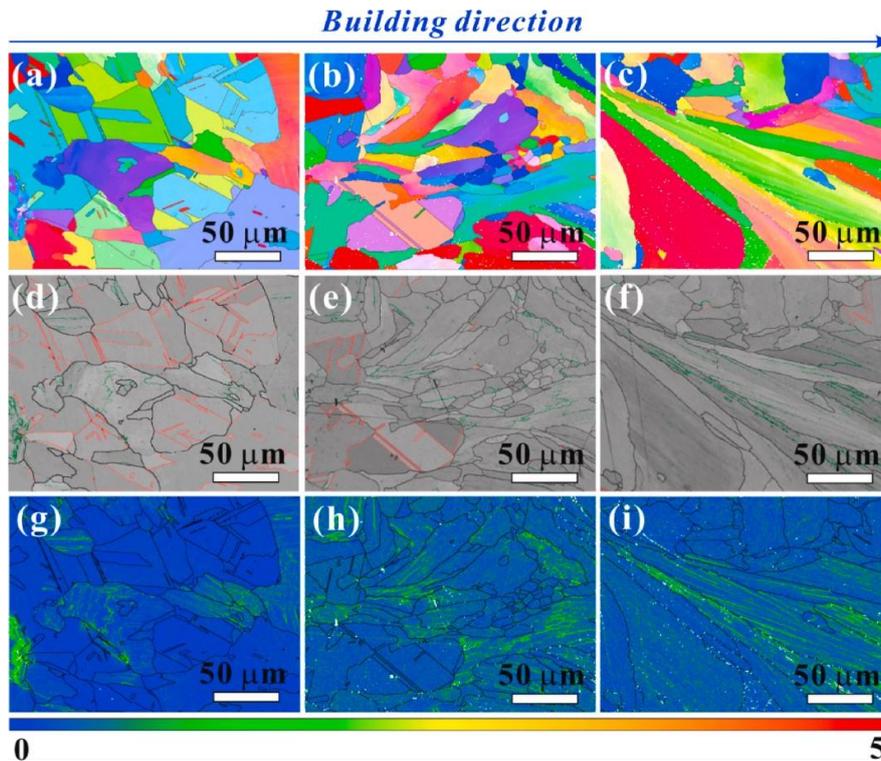


Figure 2.65

IPF maps (top row), band contrast maps (middle row), and misorientation maps (bottom row) of Wang's test specimens extracted from the lower (left column), middle (middle column), and upper (right column) sections of the build [261]

2.5.1.10 - Supports and Anchors

In their review paper on supports in L-PBF, Javidrad et al. [262] identify three scenarios in which components may require to be supported. Firstly, as discussed previously, supports are necessary to prevent material loss on overhang structures. The authors suggest, as a rule of thumb, that any downskin with an angle of less than 45° should be supported. Secondly, supports may be required to provide structural stability and weight-bearing, especially in components that may experience large thermal stresses. In extreme cases, both vertical and horizontal support structures may be required, and must be capable of withstanding both tensile and compressive forces, to restrain all component movement. Thirdly, support structures may be used to conduct heat out of the component, and into the base plate, limiting defects due to overheating.

The authors describe a variety of support geometries. The most straightforward support type is the simple block support, often built as a hollow structure to save time and material (albeit at the cost of ability to withstand strain). These blocks are positioned close to the component, and connected to them with points, webs, lines, pins, or cones. Another support type is the lattice, which can provide superior resistance to strain, with reductions of up to 60% in material usage. More advanced geometries, such as tree branch or voxel-based systems, are shown to provide benefits in specific circumstances. Some of these support types are shown in Figure 2.66.

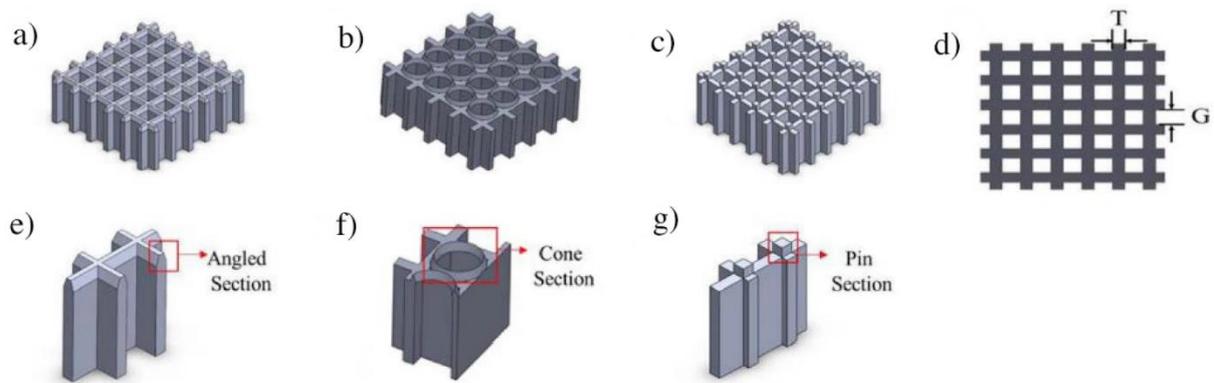


Figure 2.66

Javidrad demonstrates three methods of joining a lattice support to a component: the line (a, e), the cone (b, f), and the pin (c, g). The lattice thickness (T) and gap (G) are shown in (d). [262]

2.5.2 - Parameter Selection

Selecting the optimal laser power, laser speed, hatch spacing, and layer height for a certain geometry can be a very difficult task. Parameters that are shown to be effective for simple cubes may be poor for more complex geometries, or for more advanced scanning strategies. Many researchers have compiled datasets that were appropriate for their components, on their machines, but devising any universally-applicable system for any future build remains challenging.

Marques et al. [263] compiled a review paper of the interdependent effects of power, speed, hatch spacing, and layer height, upon the microstructural and mechanical properties of IN-718 components. Because of the non-linear relationship between the parameters, strict trends were difficult to identify in the data. Some general observations that were made, however, include:

- Higher laser power can result in larger, deeper melt pools, that lower the risk of LoF defects, but raise the risk of keyhole and evaporation defects. High power can promote coarse and columnar grains. Lower power can result in LoF, the effects of which are detrimental on mechanical properties.
- Higher laser speed can result in fine, equiaxed grains, though at the risk of increased LoF. Lower speed can favour coarse, columnar grains, and higher residual stresses, but at the risk of increased keyhole porosity.
- Smaller hatch spacing ensures full melt pool overlap, thereby reducing LoF porosity, but at the cost of increased build time and risks of overheating and defects caused by repeated remelting. Larger hatch spacing reduces build time, at the risk of LoF.
- Smaller layer heights reduce LoF risks, and allow for better feature resolution (due to smaller staircase effect) and microstructural homogeneity. Greater layer heights reduce build time, but increase the risk of delamination.

- Lower linear energy density tends to produce finer equiaxed grains, with reduced chemical segregation. Higher linear energy density tends to produce coarser columnar grains, with more anisotropy.
- A linear energy density of 0.3 to 0.4 J/mm appears to result in a fine microstructure and strong mechanical properties.
- Many authors use the same linear parameters (960mm/s speed, 285W power, and 40μm layer height) to achieve full density in their components.

In order to better understand how combinations of parameters may affect a build, it is easier to group them together rather than attempt individual analyses. Two straightforward examples, which have been introduced previously in this work, are linear energy density (J/mm) and volumetric energy density (J/mm³). A more advanced example, the Keyhole Number, was presented in Section 2.4.4 on keyhole porosity. Several other advanced examples, which often present in dimensionless format, will now be discussed.

2.5.2.1 - Heat Input

Mukherjee et al. [264] defines non-dimensional Heat Input Number as follows:

$$Q^* = \frac{(P/V)}{(P_R/V_R)}$$

where P is laser power [W], V is laser speed [m/s], P_R and V_R are reference values, corresponding to the lowest energy density used in the build.

An increase in Q* is shown to increase both the peak temperature and dimensions of a melt pool, and decrease the cooling rate. In IN-718, lowering Q* and therefore raising the cooling rate resulted in suppression of Laves phases in the as-built component.

2.5.2.2 - Peclet Number

In the same paper Mukherjee et al. also states the Peclet Number as follows:

$$Pe = \frac{UL}{\alpha}$$

where U is the characteristic velocity (laser speed can be used) [m/s], L is the characteristic length (melt pool length can be used) [m], α is thermal diffusivity [m²/s].

Pe quantifies the relative importance of convection versus conduction in heat transfer inside the melt pool. At Pe > 1, convective flow is the primary mechanism in transporting energy away from the point of laser incidence. At Pe < 1, conduction is the primary mechanism. In L-PBF processes, Pe is usually greater than 10. Pe influences melt pool morphology, which in turns influences porosity and the solidification process.

2.5.2.3 - Marangoni Number

Mukherjee et al. also explains the Marangoni Number as follows:

$$Ma = - \frac{d\gamma}{dT} \frac{L\Delta T}{\mu\alpha}$$

where γ is surface tension [N/m], T is temperature [°C], L is the characteristic length (melt pool width can be used) [m], ΔT is the difference in temperature between the melt pool's peak and its edge (liquidus temperature) [°C], μ is the dynamic viscosity [Ns/m²], α is thermal diffusivity [m²/s].

Ma describes the magnitude of convective forces within the melt pool due to Marangoni forces on its surface. At higher Ma there is a vigorous circulation of liquid, whereas at lower Ma heat is transferred primarily by conduction. As with Pe , Ma strongly influences melt pool morphology, which in turns influences porosity and the solidification process. High Ma indicates improved mixing and a deep pool, but a too high Ma might suggest melt pool instability or spatter. Figure 2.67 shows the relationship between Pe , Ma , and heat input.

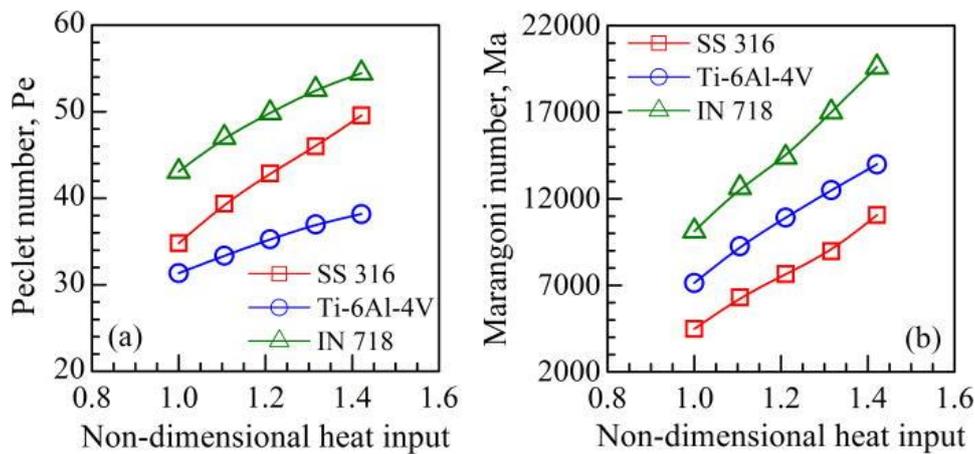


Figure 2.67

The relationship between heat input, and the Pelet and Marangoni numbers, from Mukherjee [264]

2.5.2.4 - Fourier Number

The final dimensionless number discussed by Mukherjee et al. is the Fourier Number, defined as follows:

$$Fo = \frac{\alpha}{VL}$$

where α is thermal diffusivity [m²/s], V is the laser speed, L is the characteristic length (melt pool length can be used) [m].

Fo describes the ratio between the quantity of heat that diffuses through a material, versus the quantity of heat that is absorbed by the material. A high Fo indicates rapid heat dissipation, which suggests faster cooling and smaller melt pools. A small Fo indicates that, despite large temperature

gradients, the solidification velocity is slow. Therefore, grain size and morphology can be implied from Fo. Figure 2.68 shows the relationship between GxR, G/R, and Fo.

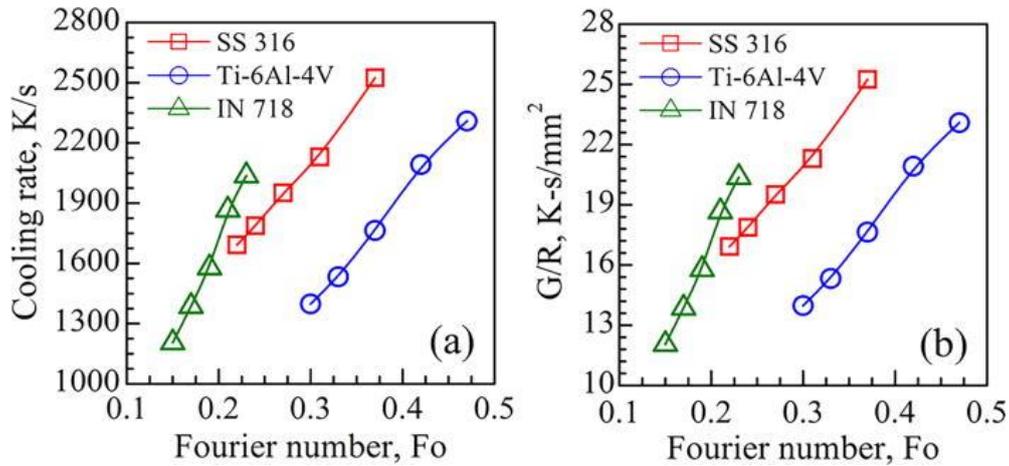


Figure 2.68

The relationship between Fourier number, and GxR and G/R [264]

2.5.2.5 - Normalised Enthalpy

King et al. [222] defines Normalised Enthalpy as follows:

$$H_n = \frac{AP}{h_v \sqrt{\pi \alpha V \sigma^3}}$$

where A is the absorptivity, P is the laser power [W], α is thermal diffusivity [m^2/s], h_v is the volumetric enthalpy of heating [J/m^3], V is the laser speed [m^2/s], σ is the beam diameter [m].

The volumetric enthalpy of heating is:

$$h_v = \rho c (T_l - T_0)$$

ρ is density [kg/m^3], c is the powder's specific heat capacity [$\text{J}/\text{kg}/\text{K}$], T_l is the liquidus temperature [K], T_0 is the powder's ambient temperature [K]

H_n is a comparison between the quantity of heat per unit volume that is input into the powder, and the quantity of heat per unit volume that is required to raise the powder from ambient to liquidus temperature. Where $H_n < 1$, melting would not occur. At $H_n = 1$, perfect conduction mode melting would occur, in theory. At very large values of H_n , keyhole defects and thermal strain become more likely. The effects of H_n on melt pool depth is shown in Figure 2.69.

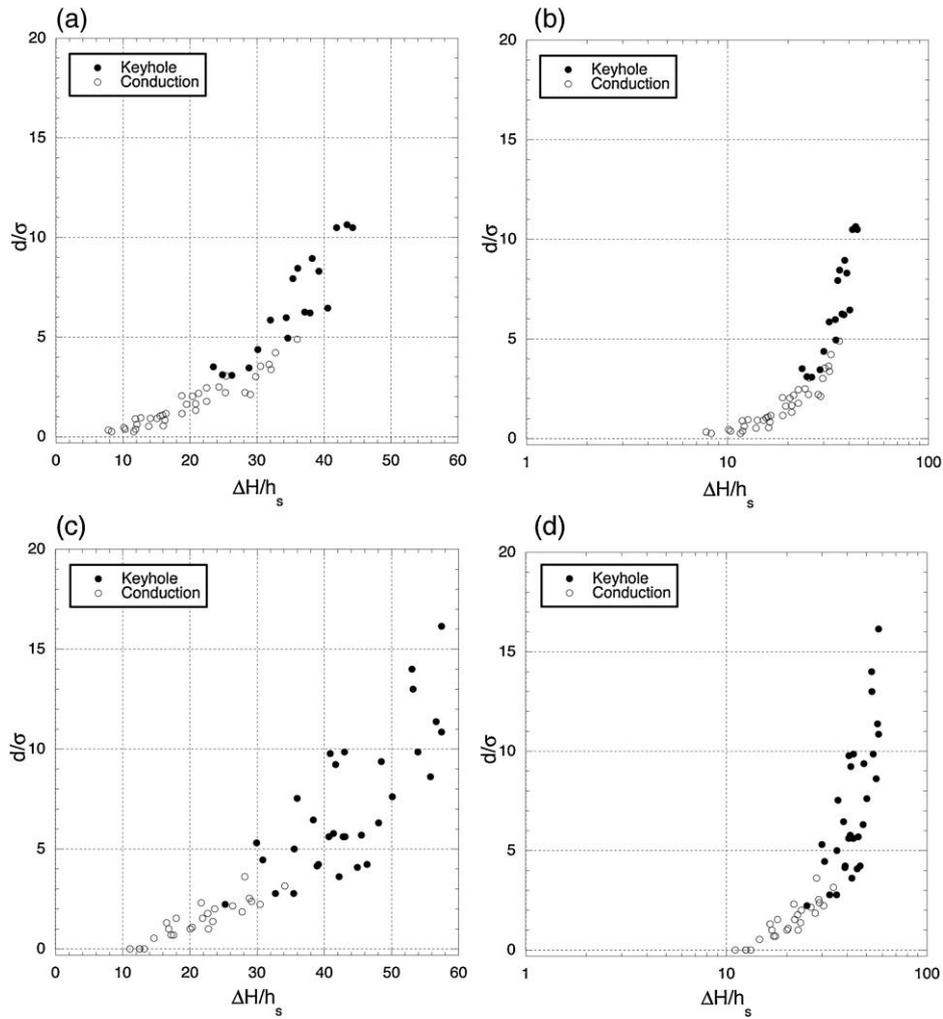


Figure 2.69

Normalised enthalpy (here labelled $\Delta H/h_s$) vs normalised melt pool depth (here defined as depth divided by beam diameter), for beam diameters of $130\mu\text{m}$ (a, b) and $52\mu\text{m}$ (c, d), plotted on linear (a, c) and log (b, d) scales, from King [222]

2.5.2.6 - Normalised Melt Depth

Rubenchik et al. [265] defines the Normalised Melt Depth as follows:

$$d_n = \frac{d}{\sqrt{\frac{\alpha\sigma}{V}}}$$

where d is the melt pool depth [m], α is thermal diffusivity [m^2/s], σ is the beam diameter [m], V is laser velocity [m/s].

The denominator in the equation is the thermal diffusion depth, a measure of the vertical distance that heat diffuses while the laser spot traverses above it. By scaling a melt pool's depth relative to this thermal diffusion depth, it is possible to compare with d_n the geometry of melt pools across a range of processing parameters and material properties, as shown in Figure 2.70.

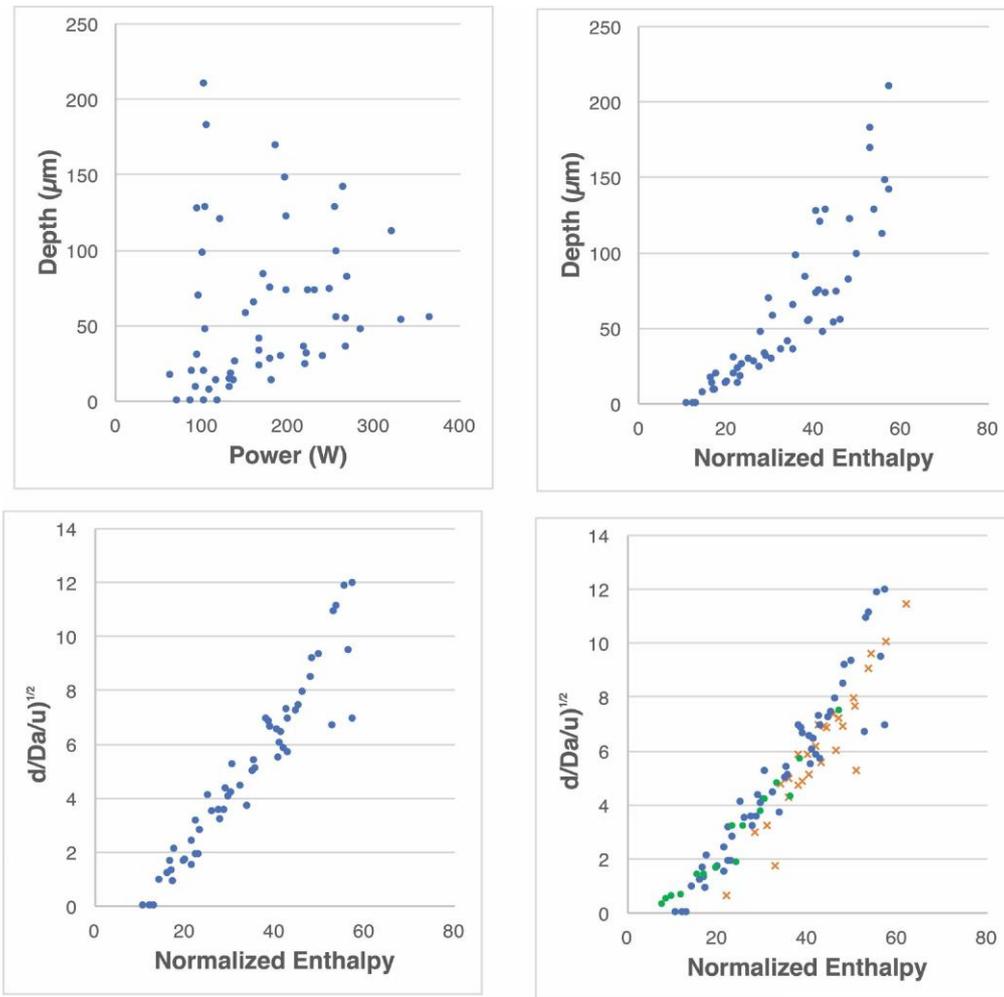


Figure 2.70

Rubenchik used data from King, but plotted normalised enthalpy against the updated metric of normalised melt pool depth, for a better fit [265]

2.5.2.7 - Dimensionless Energy

Ion et al. [266] discussed several dimensionless quantities, intended to remain applicable across all forms of metal laser processing. Dimensionless Energy is defined as follows:

$$q^* = \frac{AP}{r\lambda(T_l - T_0)}$$

where A is absorptivity, r is beam radius [m], λ is thermal conductivity [J/s/m/K], T_l is the liquidus temperature [K], T_0 is the powder's ambient temperature [K].

q^* describes the quantity of laser energy input to the powder, scaled to the melt pool's ability to conduct away that energy. At higher q^* , heat is being imparted far quicker than it can dissipate into the component or powder, resulting in a larger melt pool with a higher peak temperature.

2.5.2.8 - Dimensionless Velocity

Ion et al. also described a Dimensionless Velocity as follows:

$$v^* = \frac{Vr}{\alpha}$$

where V is laser speed [m/s], r is beam radius [m], α is thermal diffusivity [m²/s].

v^* is a measure of laser speed, scaled to the rate at which heat is conducted through the material. At lower v^* values, material in front of the laser spot will experience notable heating before the beam crosses it. At higher v^* values, conduction will not have heated this material significantly. Although v^* does not impart much information about a L-PBF process by itself, this number can be used alongside q^* for cross-build comparison, as demonstrated in Figure 2.71.

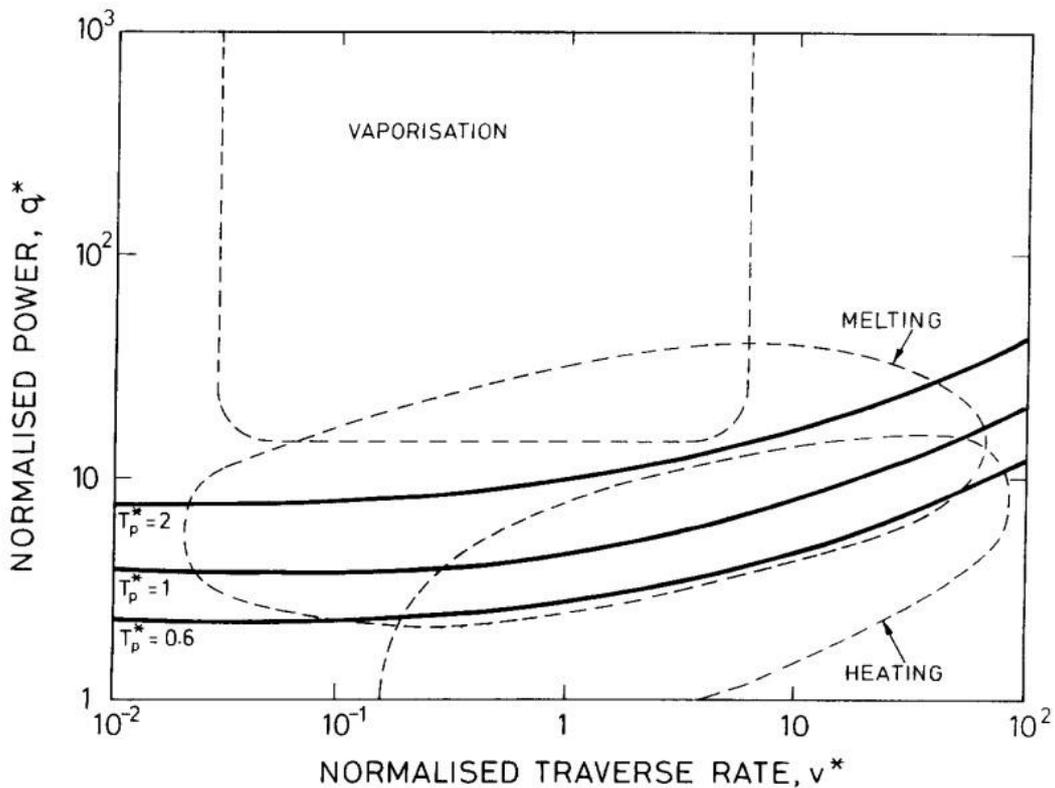


Figure 2.71

Ion collected experimental data for melt pools that exhibited heating, melting, and vaporisation, and plotted this on a diagram of normalised energy vs normalised velocity. The regions into which each data category fit are shown with dashed lines. T_p^* is a dimensionless measure of the melt pool temperature, with T_p^* of 0.6, 1, and 2 approximately signifying the onset of recrystallisation, melting, and vaporisation, respectively. [266]

2.5.2.9 - Dimensionless Hatch Spacing

Thomas et al. [267] introduced a new set of dimensionless numbers, building upon the work of Ion et al. discussed previously. One such number, which is also intended to be used alongside more meaningful numbers for cross-build comparison, is Dimensionless Hatch Spacing, defined as follows:

$$h^* = \frac{h}{r}$$

where h is hatch spacing [m], r is beam radius [m].

h^* is a better measure for understanding the overlap (or lack thereof) between adjacent scan tracks than hatch spacing alone, since the width of the laser spot also significantly impacts this factor.

2.5.2.10 - Dimensionless Layer Height

Another similar number defined by Thomas et al. is the Dimensionless Layer height:

$$l^* = \frac{2l}{r}$$

where l is layer height [m], r is beam radius [m].

l^* approximates a melt pool cross section as a rectangle of $2lr$, then normalises layer height against radius and adjusts the constant to better estimate actual melt pool depth.

2.5.2.11 - Dimensionless Volumetric Energy Density

The final and most consequential value introduced by Thomas et al. is the Dimensionless Volumetric Energy Density, defined as follows:

$$E^* = \frac{q^*}{v^* l^*} = \frac{AP}{2Vlr} \frac{1}{\rho c (T_l - T_0)}$$

where A is absorptivity, P is laser power [J/s], V is laser speed [m/s], l is layer height [m], r is beam radius [m], ρ is density [kg/m³], c is specific heat capacity [J/kg/K], T_l is the liquidus temperature [K], T_0 is the powder's ambient temperature [K].

In practical terms, E^* is the ratio of the quantity of laser energy imparted into the powder each second, versus the amount of energy required to melt the powder that the laser spot passes over each second. When E^* is plotted against the reciprocal of h^* , as shown in Figure 2.72, the graph accounts for enough process and material data to allow cross-experimental comparison.

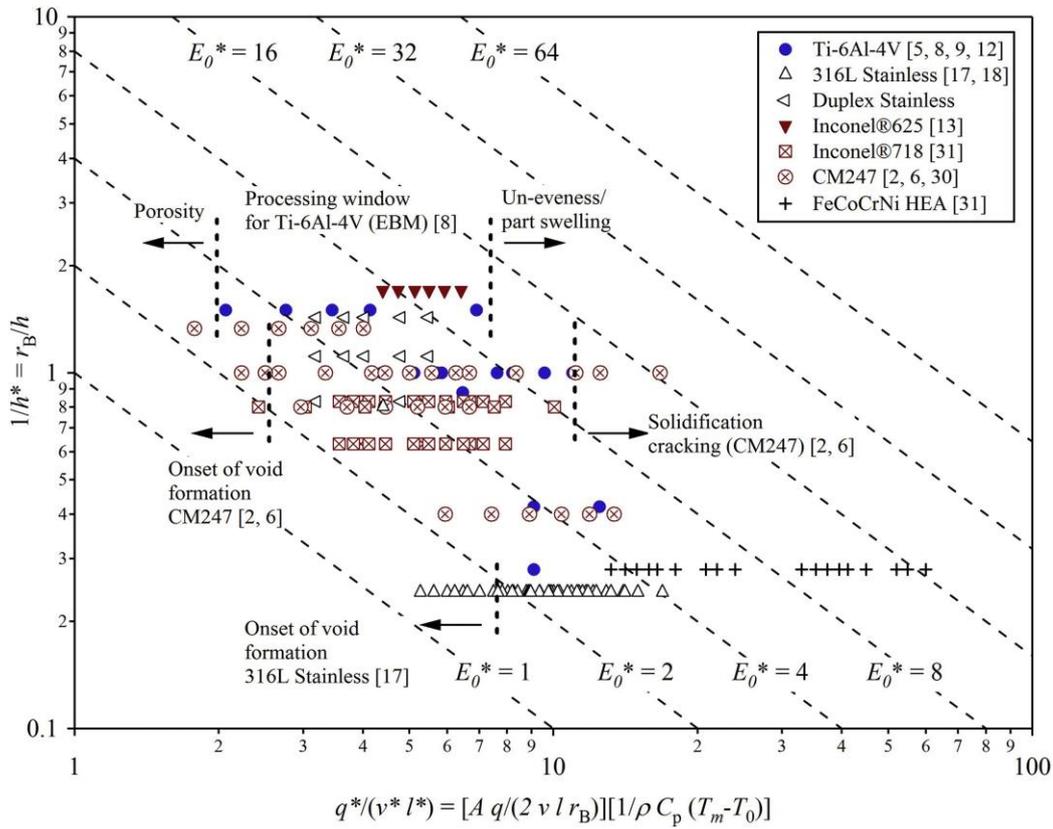


Figure 2.72

Thomas compiled experimental data, and plotted it on a diagram of $1/h^*$ against E^* . The regions of the diagram in which various defects occur are shown. E_o^* is the product of the two terms (i.e. $E_o^* = E^*/h^*$), and serves as a benchmark for the amount of energy the melt pool is receiving. [267]

2.5.2.12 - Ambient Temperature

In many of the previous equations, values are required for T_o , the estimated temperature of the powder in the moment before the laser spot crosses it. Ambient workshop temperatures can be used for this value, although will likely be wildly inaccurate, especially after the L-PBF machine has been operating for several hours. Phan et al. [268] proposed a more rigorous way to estimate T_o as follows:

$$T_0 = T_{pre} + \frac{0.3P}{2\pi\lambda\sqrt{(Vt_r)^2 + h^2}} \times \exp\left(-V\frac{\sqrt{(Vt_r)^2 + h^2} + Vt_r}{2\alpha}\right)$$

where T_{pre} is powder preheat temperature [K], P is laser power [J/s], λ is thermal conductivity [J/s/m/K], V is laser speed [m/s], t_r is laser return time [s], h is hatch spacing [m], α is thermal diffusivity [m²/s].

$$t_r = \frac{L}{2V}$$

where L is scan length [m].

This equation is based upon the Rosenthal Equation, which predicts the temperature field around a moving heat source. It calculates the midpoint temperature of the next scan, as the laser completes the current scan.

The ability to obtain an accurate T_0 value can be of significant value in parameter selection. Phan et al. describe how using this method of estimating T_0 can, for example, bring mathematical clarity to the observation that shorter scan lengths in crack-susceptible builds tend to reduce cracking risk, due to less extreme thermal cycling.

v^* is plotted against q^* , using T_0 in the calculation, in Figure 2.73.

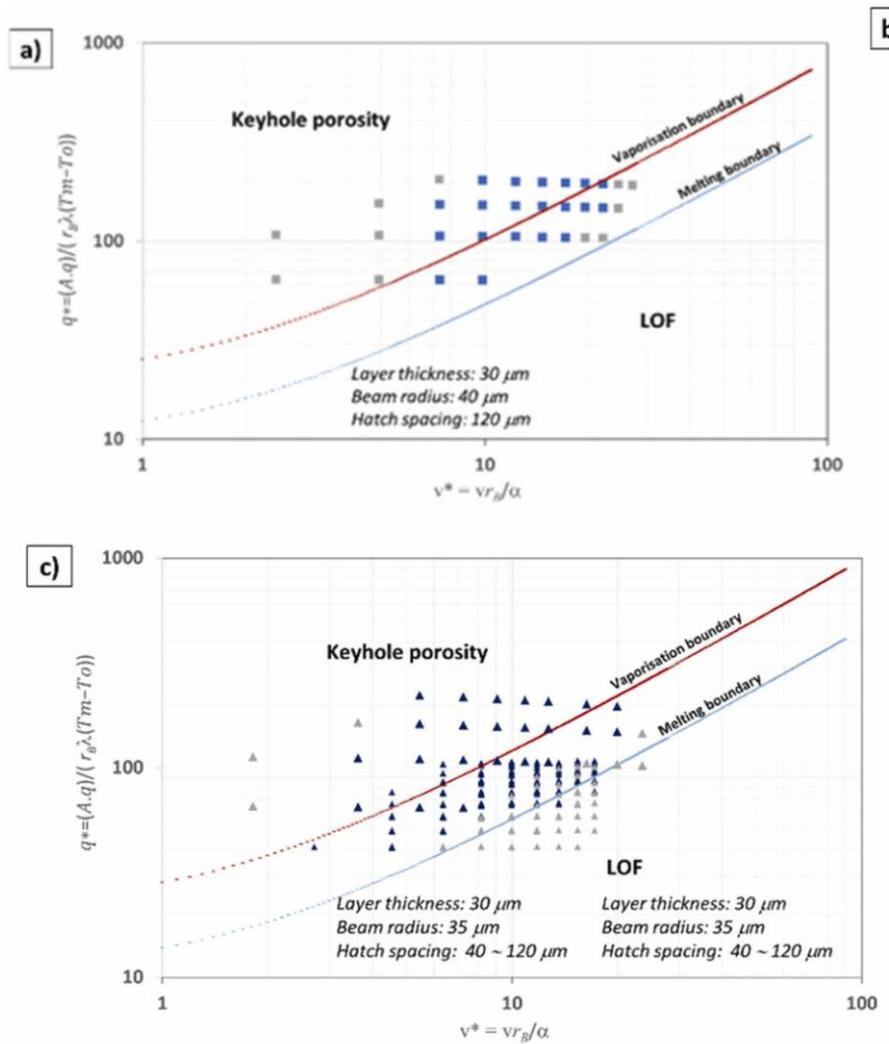


Figure 2.73

Phan updated the work of Thomas, using a more accurate measure of T_0 . Here experimental data for melt pools in IN-718 (squares) and SS316 (triangles) is shown, on both linear and log scales, with boundaries drawn at the expected melting and vaporisation points. Dark shapes show samples with >99% as-built densities, while grey shapes show <99%. [268]

2.5.3 - Thermal and Residual Stress Management

Even when building the same component geometry, from the same machine parameters, from the same powder, differing scan strategies can lead to vastly different outcomes, in terms of surface roughness, microstructure, residual stress, and even porosity and thermal distortion. Many scan patterns have been tested by various researchers, with each reporting its own advantages and disadvantages. As demonstrated in Figure 2.74, Malekipour et al. [269] defines five fundamental patterns:

1. Unidirectional rastering, in which the laser follows a path like words on a page: unidirectional vectors in X, stacked to give a direction in Y. Due to the requirement of the laser to turn off at the end of each vector, and move to the start of each next vector, this method is very time inefficient.
2. Bi-directional rastering, in which the laser reverses direction each vector, 'snaking' its way across the part. This method is much more time efficient, although there is a large heat concentration at the ends of the vectors.
3. Spiralling outwards, starting from the centre and reaching the perimeter. The outer contour allows for better surface roughness, albeit with very significant thermal conditions between the spiral edge and centre.
4. Spiralling inwards, starting from the perimeter and reaching the centre. This method exhibits the same advantages and disadvantages as listed previously.
5. Fractal scanning, in which small patterns are repeated, summing to larger versions of the same pattern, which are repeated to sum to larger versions, etc., until the entire area is scanned. This method has the potential to offer careful control of heat accumulation, at the expense of high complexity and build time.

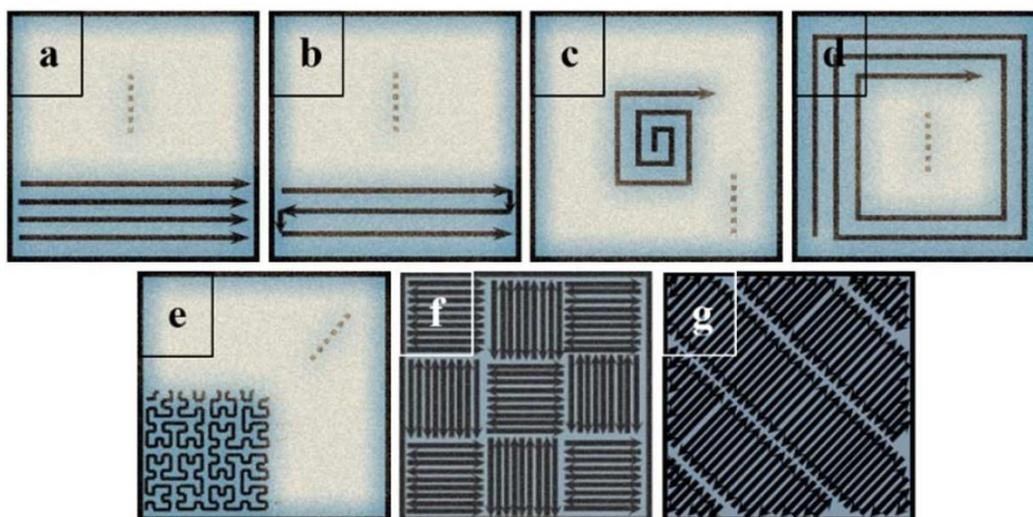


Figure 2.74

The five fundamental scanning patterns as defined by Malekipour (a-e), which can be combined into chessboard sections (f), or stripe sections (g) [269]

Each of these five fundamental patterns either cover the entire layer, or can be combined into a chessboard (squares of equal dimensions) or stipes (rectangles of with longer sides) strategy to cover the layer. The pattern (whether single or combined) can be rotated on successive layers, and the point about which it is rotated can alternate between layers, to prevent the same points of heat accumulation existing on each layer in a build. A 67° rotation is often used between layers, as this number is not a factor of 360° , and results in very little coherence between the vectors in layer N, and those in layer N+1, N+2, N+3, etc.

Ravichander et al. [270] and Liu et al. [271] each tested a ranges of scanning strategies in IN-718 test samples, to find the effects on melt pool dimensions and mechanical properties. Some of the patterns tested were: uni-directional, non-rotating vectors; bi-directional, 90° rotating vectors; bi-directional, 67° rotating vectors; chessboard squares; parallel stripes.

Ravichander found significant differences in melt pool dimensions between the striped and chessboard strategy ($45\mu\text{m}$ depth and $120\mu\text{m}$ width, and $79\mu\text{m}$ depth and $140\mu\text{m}$ width, respectively). Similarly, the smallest average deviation in cube measurements from the intended CAD dimensions were observed in the striped strategy (0.7%), while the largest were observed in the chessboard strategy (1.5%). The chessboard strategy showed the greatest Vickers hardness value (421.4 HV), however, while the striped strategy showed the lowest (374.1 HV). These results highlight both the macrostructural and microstructural differences that can accompany differing scanning strategies, even in geometry as simple as a cube.

Liu found that vectors that spanned the entire layer (whether rotated or not) all showed similar results in terms of yield strength, ultimate tensile strength, and ductility, while the chessboard method showed poorer results in each category.

The reason for the chessboard's poor performance was reasoned to be due to the rapid laser on/off cycles, and the overlapping of melt pools at square perimeters, both of which led to non-uniform melt pool morphologies and thermal conditions, which in turn led to high microstructural anisotropy and micro-segregation.

These findings from Liu et al. highlight the fact that straightforward-sounding ideas (such as using a chessboard strategy to more evenly distribute heat accumulation) can have unintended and counterintuitive results. Scan strategies must be experimentally verified, for the specific component's geometry and purpose, in order to optimise outcomes.

2.5.4 - Powder

There are three major factors that determine the quality of a powder: Particle Size Distribution (PSD), powder morphology, and chemical composition. Due to the stochastic nature of GA and WA powder production, and due to the imperfect conditions in which powders may be transported, handled, and stored, powders in workshops and laboratories always deviate from the perfect ideal:

spheres of uniform diameter, free of surface oxides or moisture, and containing the exact atomic percentages on their data sheets

The extent to which PSD affects mechanical, microstructural, and surface properties, is the focus of much research. Stegman et al. [272] built components from two IN-718 powders from the same supplier, one of which had a typical, unimodal PSD (mean particle diameter 29.1 μm , median diameter 26.2 μm), and the other had a wider, bimodal PSD, featuring peaks of both smaller and larger particles (mean diameter 34.4 μm , median diameter 33.9 μm). These PSDs are shown in Figure 2.75.

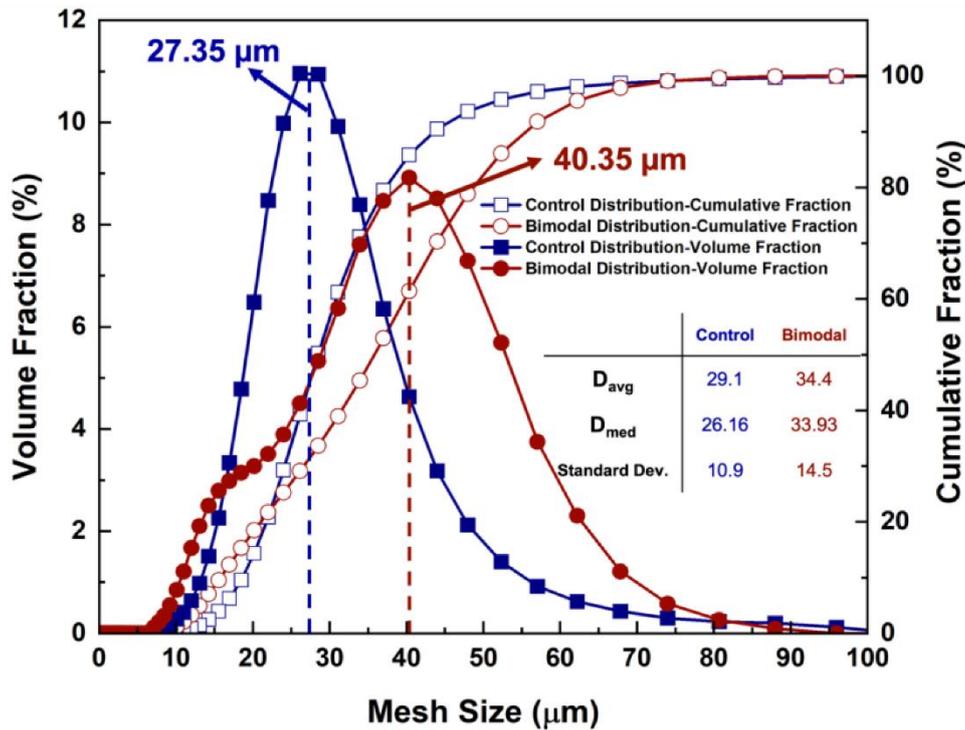


Figure 2.75

The cumulative fraction (hollow points) and volume fraction (solid points), of the unimodal powder (blue) and bimodal powder (red) from Stegman [272]

It was found, across almost all categories tested, that the differences in properties of the components were negligible. In most cases, the unimodal powder resulted in densities up to 0.2% greater than the bimodal, while in one sample the bimodal density was 0.25% higher. The length of columnar grains along the build direction was, on average, 9.2% larger in the bimodal than the unimodal samples (141.2 μm compared to 128.2 μm). The differences in grain widths was not statistically significant. The bimodal samples were also found to have greater fractions of High Angle Grain Boundaries. Despite these slight differences in microstructural properties, the differences in hardness and tensile strength between the two PSDs were not found to be statistically significant.

Attempting to find a linkage between PSD and surface roughness, Sendino et al. [273] experimented with thirteen different batches of IN-718 powders, all from different suppliers and each with a

different PSD. The average surface roughness (R_a) of the as-built components ranged from $6\mu\text{m}$ to $9\mu\text{m}$, and the maximum roughness (R_z) ranged from $40\mu\text{m}$ to $60\mu\text{m}$. A linear relationship was observed between the R_a value, and the mean diameter of particles below D_{50} (i.e. the diameter which is larger than that of 50% of the particles by volume). It was found that the smaller the average sub- D_{50} diameter, the lower the surface roughness, due to the tendency of smaller particles to adhere to outer surfaces, thereby smoothing them. No relationship was found between R_a and the mean diameter of particles above D_{50} .

The results of these two suggest that, in order to achieve optimal results in terms of microstructural, mechanical, and surface roughness properties, using a powder with a wider PSD, skewed towards the smaller end, may be preferable for builds in IN-718.

As well as PSD, the morphology of powder particles can play a role in determining quality, primarily by influencing flowability. Ruan et al. [274] demonstrated this effect of morphology by producing components from two batches of IN-718. The first batch was Gas Atomised (GA), and the second was created with Plasma Rotating Electrode Process (PREP), a process in which the tip of a rotating bar of solid IN-718 is melted with plasma, and centrifugal forces fling off molten droplets which solidify in an inert atmosphere. Images of these powders are shown in Figure 2.76.

The GA powder was found to have a wider PSD, with larger fractions of both smaller and bigger particles, while the PREP powder had a more uniform PSD. The GA powder was significantly more porous than the PREP (99.969% dense, compared to 99.999%). The PREP powder demonstrated better flowability than the GA powder (12.4s for 50g to flow through a standard Hall funnel, compared to 16s), and had a higher density (4.84 g/cm^3 compared to 4.22 g/cm^3). The fact that the GA powder contained more small particles to fill spaces between larger particles, yet still exhibited a lower density, was reasoned to be due to the fact that its irregular morphology prohibited optimal packing. The chemical composition of the two powders, and the absorptivity of 1080nm laser light, was determined to be near-identical.

Many of the differences in the microstructural and mechanical properties of components produced from the two powders, were found to not be statistically significant. The hardness values, ultimate tensile strengths, yield strengths, elongations at fracture, grain sizes, grain morphologies, and precipitate contents, were all found to be very similar. One significant difference between the two powders, was the larger processing window of the PREP. It was found that fully dense samples could be produced from PREP with a range of 100W to 300W, and 400 mm/s to 1500 mm/s. Full density was only achieved for GA powder in the range 150W to 300W, and 600 mm/s to 1500 mm/s. The researchers suggest that the higher flowability and density, and more uniform PSD of the PREP powder, allows for a more homogenous powder bed, which in turn leads to higher thermal conductivity and reduced melt pool size and shape fluctuations, and therefore a wider processing window.

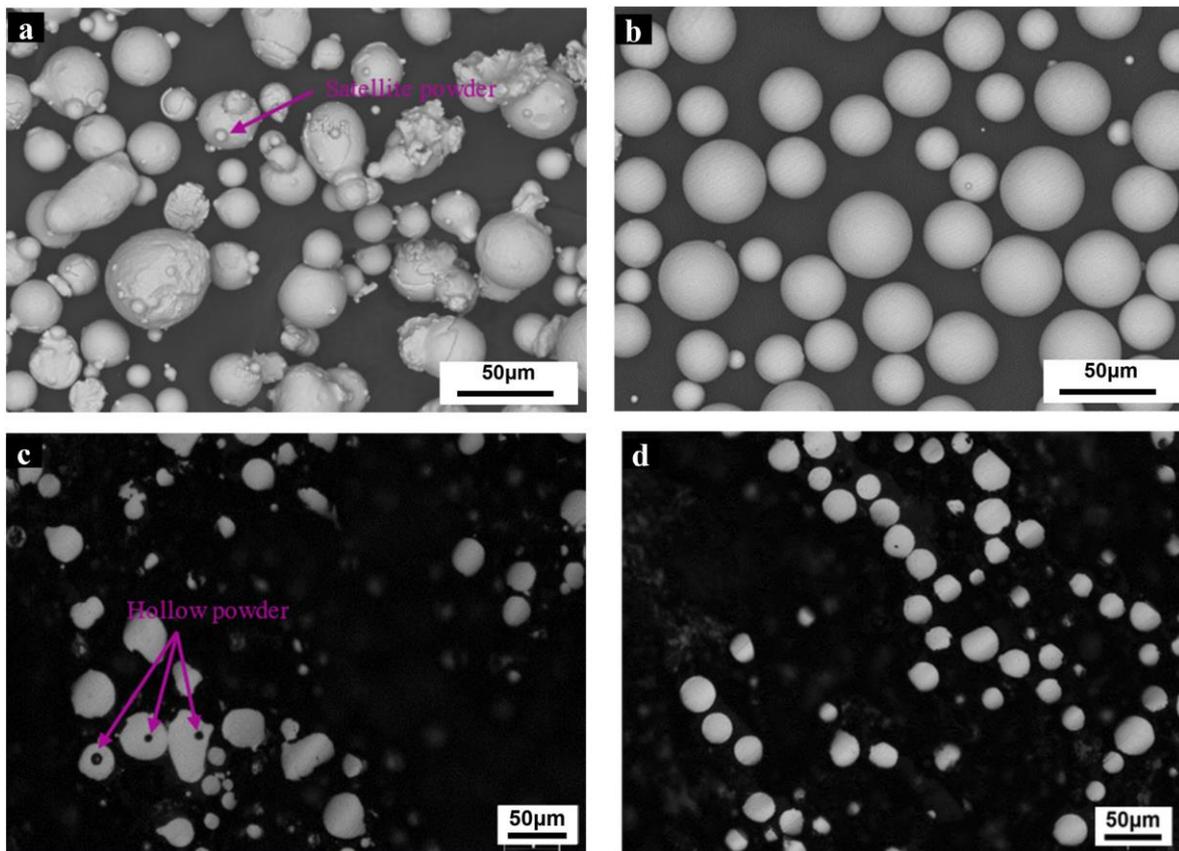


Figure 2.76

SEM (top row) and optical (bottom row) images of the gas atomised (left column) and PREP (right column) powders [274]

These results suggest that, as long as powder is spread evenly across a build plate, and optimal laser parameters are selected for a component, the effects of powder morphology on microstructural and mechanical properties may be slight. However, using more uniform, more spherical particles from a more advanced atomisation process (such as PREP), may make the build more forgiving of non-optimal parameters selections, or poorer powder spreading mechanisms.

The effects of powder recycling on powder PSD, morphology, and chemical composition, and on component microstructure and mechanical properties, has been the basis of much research. Castillo et al. [275] ran virgin IN-718 powder through 21 build-sieve cycles, testing a range of properties of the powder and components at each step. A 53µ sieve was used between builds, to remove spatter and agglomerated particles. In every tested aspect, the effects of repeated build-sieve cycles were found to be minimal.

The D_{10} , D_{50} , and D_{90} of the virgin powder were 12.0µm, 18.6µm, and 30.7µm, and of the powder after 21 cycles were 14.3µm, 18.69µm, and 31.32µm. The fact that smaller particles are being removed from the powder as the cycles progress, yet the effects on the PSD of larger particles is negligible, was attributed to the fact that agglomerates remove small particles from the powder, and are then removed themselves by the sieve. The effects of repeated build-sieve cycles on both the particle morphology (as measured by Scanning Electron Microscopy) and composition (as measured

by Energy Dispersive X-ray Spectroscopy), was found to be negligible. These results are shown in Figure 2.77.

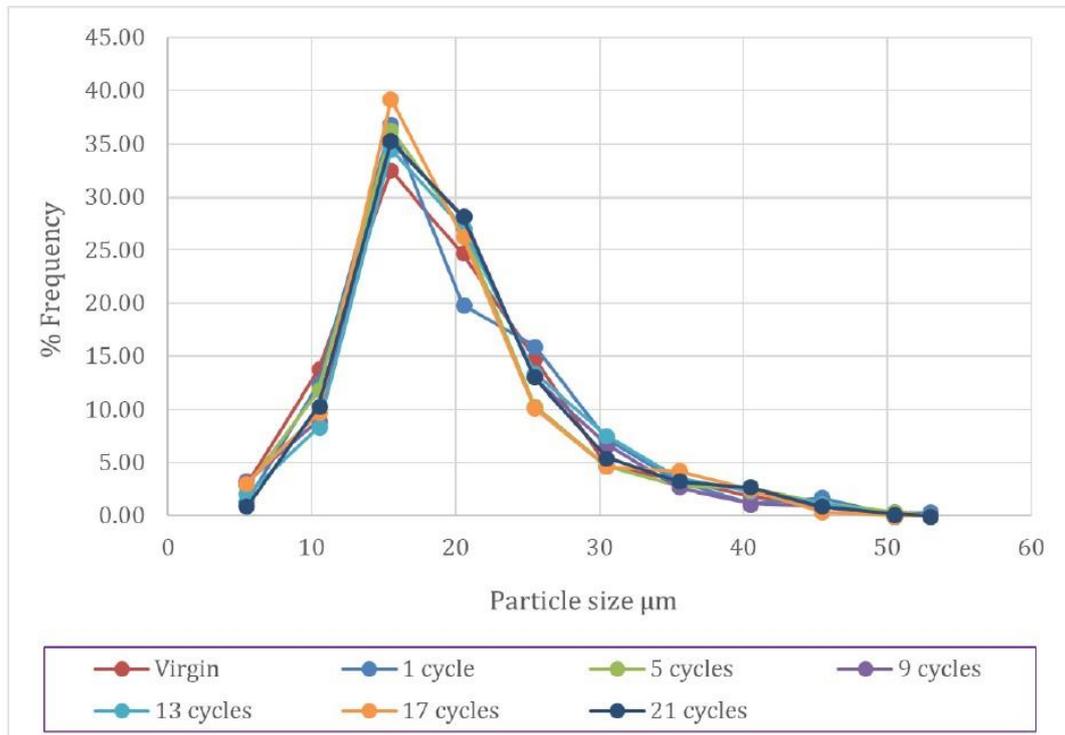


Figure 2.77

Changes to the PSD of Castillo's powder is barely noticeable, even after 21 cycles [275]

Although the average porosity, ultimate tensile strength, and yield strength of the components worsened as the cycles progressed, these changes all fell within the margin of error (i.e. there was more variation between samples at each step, than there was between the averages of all steps). It was therefore concluded by the authors that, so long as powders are carefully sieved between builds, repeated recycling will cause no significant changes to build quality.

It must be noted that repeated recycling can have detrimental effects on powder and build quality, if sieving is either not performed, or an inappropriate sieve size is used. Rock et al. [276] performed a similar experiment to that of Castillo et al. described previously, but used a 140μm sieve on the IN-718 powder between builds, to test a worst-case scenario. In this experiment, after 10 build-sieve cycles, the D_{90} increased from 49μm to 85μm, the oxygen content of the powder increased from 130ppm to 230ppm, and the flowability decreased from 11.3 s/50g to 13.5 s/50g.

2.6 - In-Situ Process Monitoring

This section discusses the various signals that are produced during the L-PBF process, and how these signals can be collected and interpreted to provide 'in-situ' (real time) information about the state of the build. Each subsection will focus on a different category of in-situ signal, breaking down both the specific signals within that category, and the ways in which those signals can be analysed. The

final subsection will explore closed-loop control, the ability of systems to auto-adjust their parameters based on live signals. The categorisation of signal types used in this section is built upon the work of Grasso et al. [277].

2.6.1 - Embedded Sensor Signals

All L-PBF machines incorporate some form of sensors, which are necessary for their basic functionality. Some examples of sensors without which the machines could not operate are: the chamber pressure sensor (which ensures a positive pressure of inert gas within the build chamber); the oxygen content sensor (which ensures the laser only activates when the chamber atmosphere contains below a specified parts per million of O_2); and the inert gas flow sensor (which ensures a steady volumetric flow rate throughout building).

The extent to which real-time data can be extracted from these sensors is limited, and very little mention of their signals is made in literature. One embedded sensor that has received attention, however, is the torque sensor on the recoater's electric motor.

Gallina et al. [278], building in Ti-6Al-4V, demonstrated that by monitoring the current that the recoater blade draws as it pushes powder across a new layer, geometric distortions in the part could be detected. Specifically, spikes in the current drawn were shown to indicate that the blade had made contact with a swollen section of the component. These spikes were detected several layers before the swelling became significant enough to trigger the motor's torque overload warning, that would shut down the machine. Figure 2.78 shows these spikes that were related to defects.

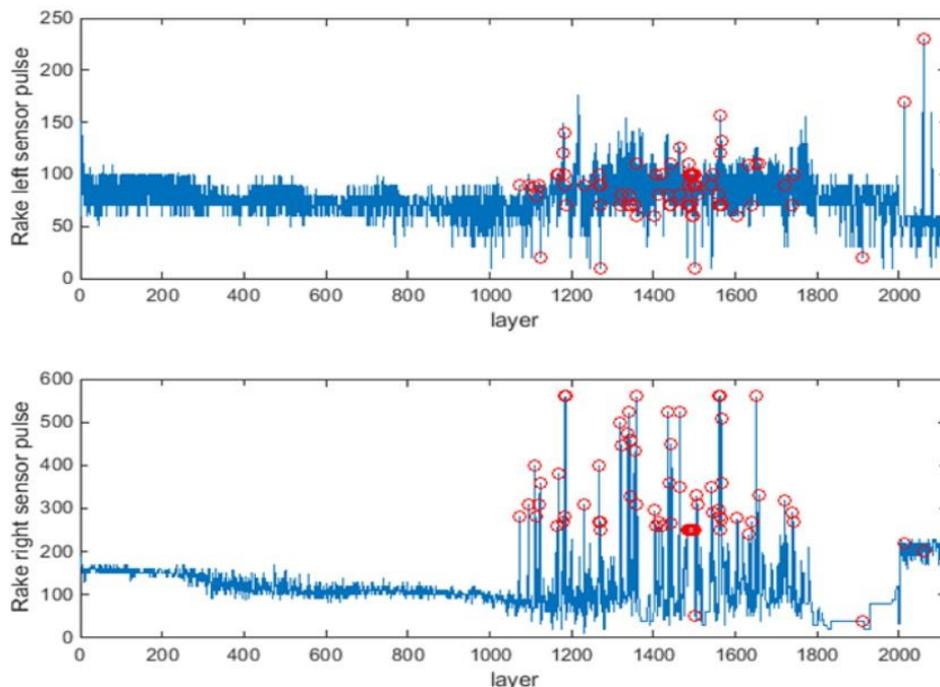


Figure 2.78

The currents being drawn by the left and right motors during Gallina's build. The red circles indicate layers in which defects were identified. [278]

2.6.2 - Powder Bed Signals

Powder bed signals are those that provide data on the state of the entire bed, usually by imaging it immediately before, during, or after powder recoating, when the laser is inactive. The following subsection breaks down the most common techniques used for such imaging.

2.6.2.1 - Photographing The Powder Bed

Scime et al. [279] used off-axis imaging, in the visible spectrum, to search for a range of anomalies both before and after recoating. Experimenting with three different L-PBF machines from different manufacturers, the authors utilised side-lighting and low camera angles to produce high contrast across the bed. Images taken before recoating were used primarily to identify regions of spatter and porosity, while images taken after were used primarily to identify swelling, incomplete powder spreading, and recoater streaking. Depending on the machine and the defect being tested for, the authors found that their off-axis imaging systems detected between 60% to 99% of anomalies successfully. These results are shown in Figure 2.79 and Figure 2.80.

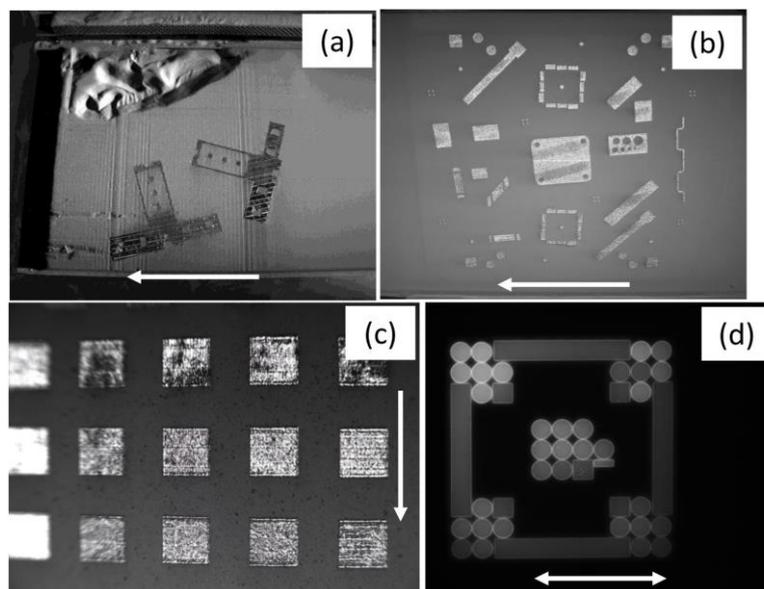


Figure 2.79

Example images taken of the powder beds for three L-PBF build (a-c) and a EB-PBF build (d) by Scime [279]

Similarly, Lu et al. [280] positioned a Nikon DSLR outside the viewing window of an L-PBF device building in 316L steel. The build chamber was illuminated with white LED lights, as a variety of samples were produced at different energy densities. Images were recorded on the camera immediately before recoating, and converted into greyscale format with each pixel representing a trapezoid with sides $10\mu\text{m}$ to $13\mu\text{m}$ upon the part surface. It was found that pixel intensity could later be related to pores, as detected by micro computed tomography, with a Pearson correlation coefficient (a measure of the extent to which two variables are linearly related to each other) of up to 0.78, indicating a very strong correlation. The XCT results are shown in Figure 2.81.

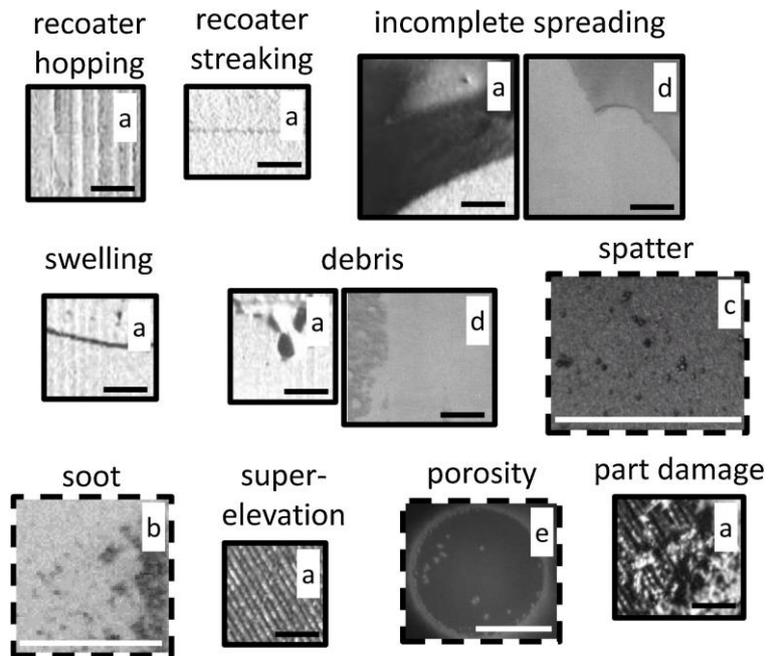


Figure 2.80

Examples of defects automatically detected on the powder bed by Scime. Note, those labelled (d) and (e) are not from L-PBF and should be ignored. [279]

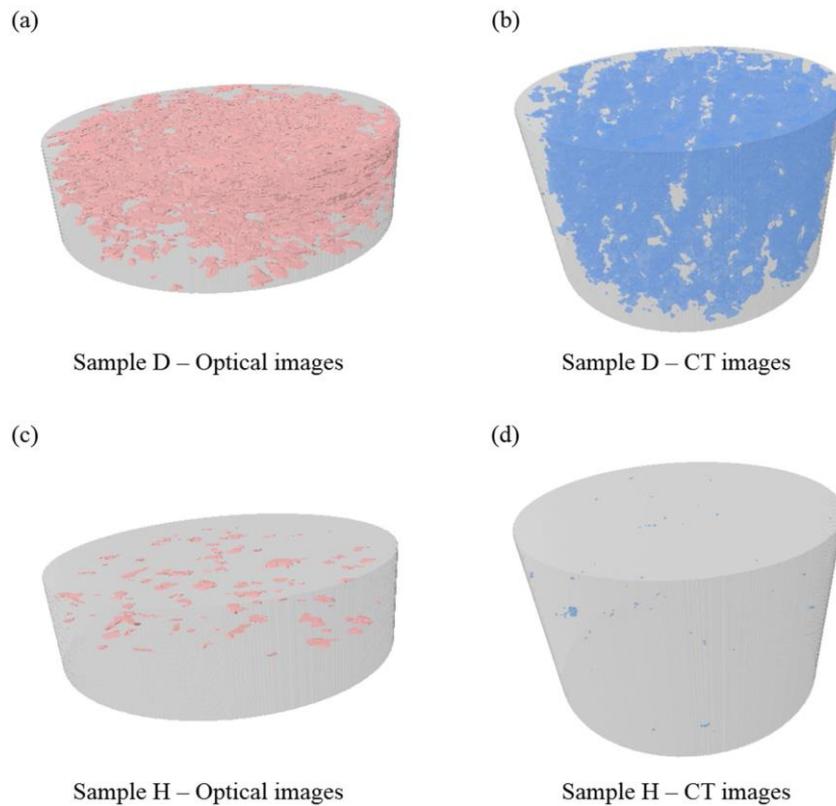


Figure 2.81

Two of Lu's samples, displaying strong correlations between defects identified via the DSLR (a, c), and defects identified via XCT (b, d) [280]

2.6.2.2 - Fringe Projection

Fringe projection is a topography measurement technique involving a projector, which casts light patterns onto a surface, and a camera at a different location, which detects the way the patterns appear on the surface. By analysing the apparent distortion of the patterns from the perspective of the camera, it is possible to determine the topography of the surface.

Zhang et al. [281] used fringe projection to monitor the building of a rectangular test piece in IN-625. By analysing the topography of the build chamber, both before and after recoating, the authors were successfully able to detect regions of spatter, recoater streaking, part roughness, and porosity. These defects could then be correlated to machine parameters. For example, the number of pixels indicating porous surfaces increased from 1.85% to 25.39%, when laser power was reduced from 350W to 230W. A height map of the powder bed is shown in Figure 2.82.

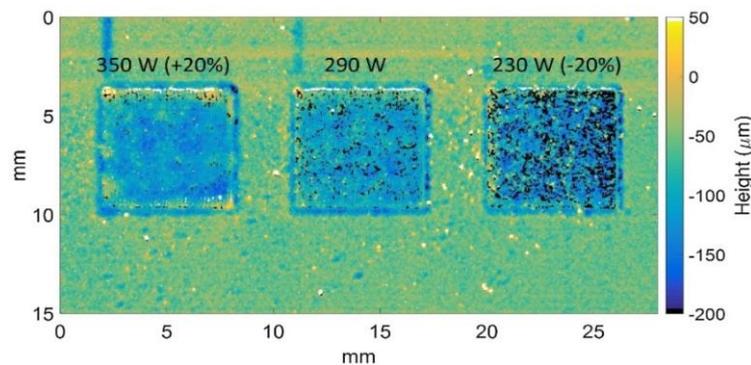


Figure 2.82

A height map of three samples from Zhang, relative to the mean height of the powder bed. Note how LoF porosity increases as laser power is decreased. [281]

Kalms et al. [282] also used fringe projection techniques in a IN-718 build, to analyse the powder bed after recoating with the aim of determining the shape of each printed layer. The detection system was shown to be capable of measuring the edges of the part to within 100 μ m laterally, and all surfaced to within 10 μ m vertically. The authors concluded that the system had the ability of serving as an early warning system against geometric distortions due to thermal strain or residual stress.

2.6.2.3 - Recoater-Mounted Sensors

Sensors mounted directly onto the recoater blade can very efficiently capture images of the powder bed, immediately after recoating occurs. Tan et al. [283] retrofitted a consumer-grade sensor from a Canon document scanner, onto the recoating mechanism of a custom-built L-PBF device. The system was shown to be highly adept at detecting the topography of the layer, largely owing to the extremely narrow depth of field of the photodetectors. Building with 316L steel, the system was shown to be capable of detecting powder bed height changes, representative of swelling, trenches, and partial recoating, to within 50 μ m. The experimental setup, and an example set of results, are shown in Figure 2.83 and Figure 2.84, respectively.

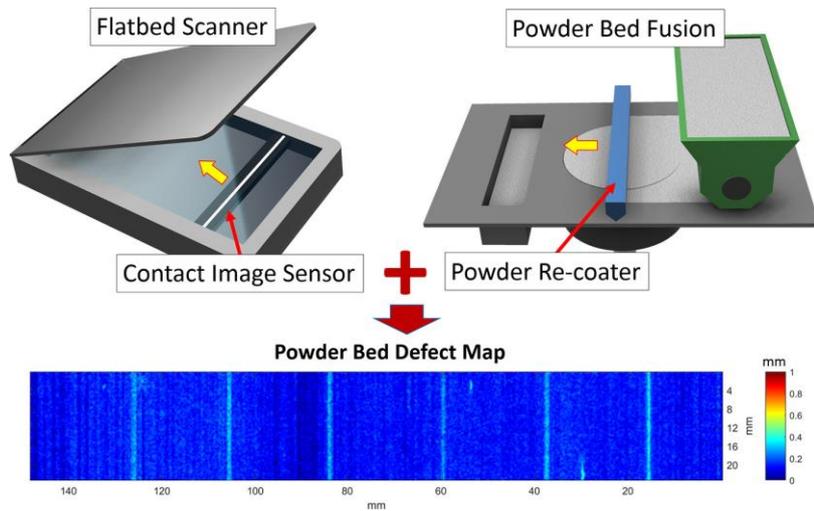


Figure 2.83

A diagram of the document scanner that Tan took the photodetector from, the recoater that it was fastened to, and an example of the height maps that it produced [283]

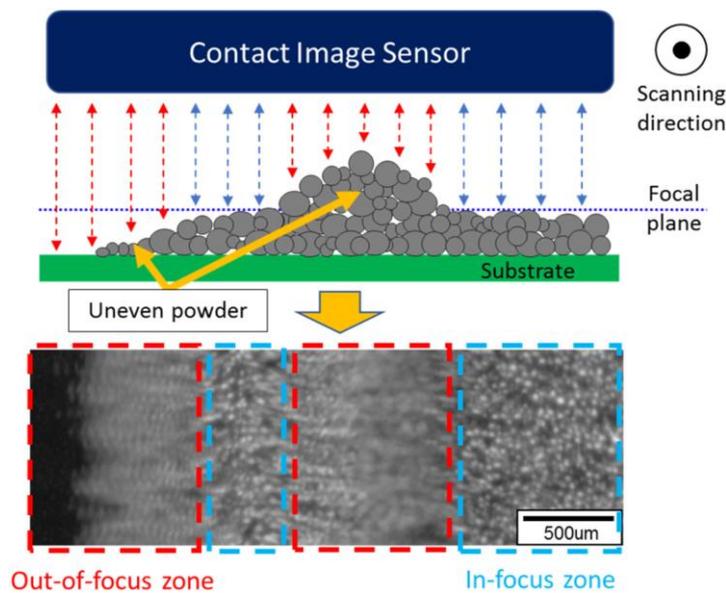


Figure 2.84

Tan's photodetector identified defects by scanning for regions of powder that were raised or fallen [283]

Barrett et al. [284] also produced a height-map of a powder bed, by fitting a Keyence profilometer onto a recoating device. This profilometer used a triangulation system, in which multiple sensors detected the reflections of laser pulses upon a surface, to determine the surface's topography. Building with IN-718, the authors positioned the profilometer so that it scanned the surface immediately before powder was spread. The primary aim of the study was to search for elevated heights across the regions that had just been lasered, which it was reasoned would correspond to non-melted powder particles (since melting is associated with a height drop). The system was capable of detecting all of the regions in samples where the authors intentionally caused LoF porosity due to low energy density.

2.6.2.4 - Low Coherence Interferometry

Low Coherence Interferometry (LCI) is a distance measurement technique that involves a light source that emits with a wide range of wavelengths, known as a broadband. When broadbands overlap, interference will only occur when the waves are near-perfectly synchronised. This fact can be taken advantage of, by splitting a broadband beam in two, and allowing one beam to reflect off a mirror, and the other to reflect off a surface. Only when the mirror and surface are at exactly the same distance (within $<10\mu\text{m}$ for many broadbands), will strong interference occur. When the mirror and surface are not at the same distance, either the mirror distance can be altered until interference occurs, or the returning spectrum can be analysed to determine the phase difference embedded within, from which the distance difference can be obtained.

DePond et al. [285] utilised the spectral analysis method to create a height map of a $4.4 \times 4.4 \text{ cm}^2$ section of powder bed, while building test parts from 316L steel. The spectrometer sampled at 100kHz, with datapoints spaced $100\mu\text{m}$ laterally, and with a $25\mu\text{m}$ resolution vertically. The LCI system was mounted coaxially with the laser beam, and height mapping occurred immediately after the laser deactivated each layer, before powder recoating. The authors primarily used the system to find the roughness of each printed layer. The LCI system successfully determined that, for cubes built with identical scanning strategies each layer, roughness increased by around $1.5\mu\text{m}$ per layer, due to compounding denudation and LoF porosity. For cubes built with 67° rotation between layers, however, a steady-state R_a of $20\mu\text{m}$ was settled upon after only three layers. The experimental setup and an example set of results are shown in Figure 2.85 and Figure 2.86, respectively.

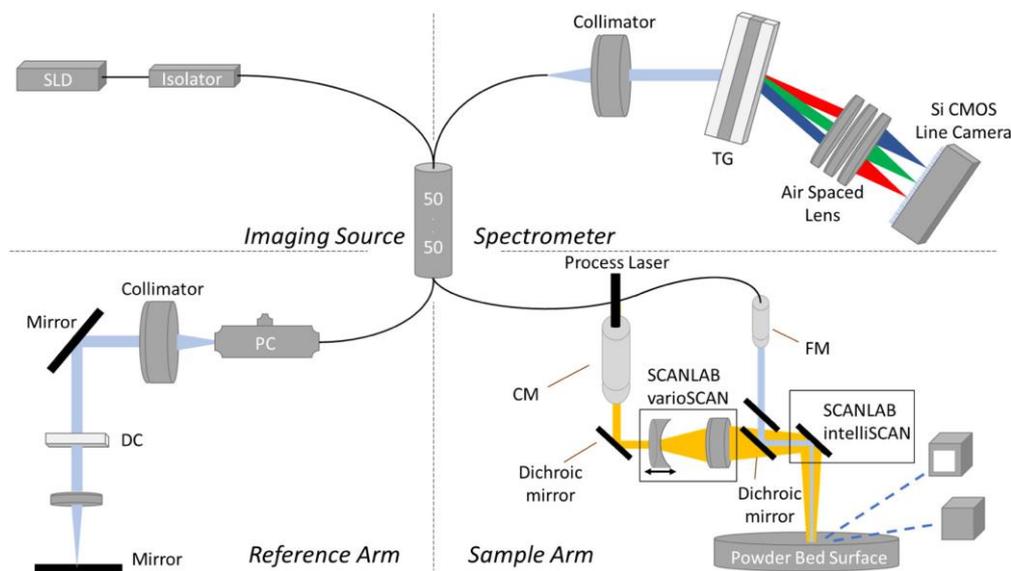


Figure 2.85

DePond's LCI setup, showing: a superluminescent diode (SLD) as the beam source; the beam in the reference arm passing through a polarisation controller (PC) and dispersion compensator (DC); the beam in the sample arm being combined with the processing laser (CM) via a focussing mount (FM); and light returned from both the sample and reference arms being split by dispersive elements (TG) and collected by a CMOS camera [285]

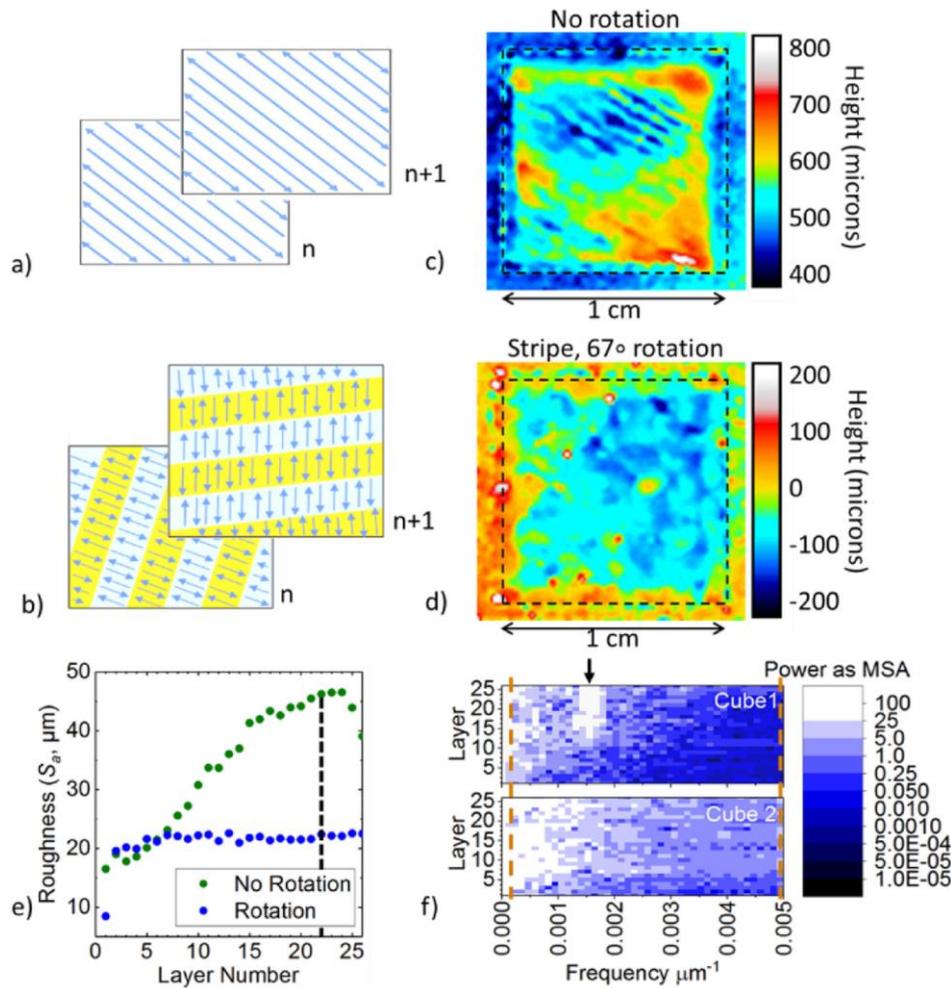


Figure 2.86

Results from DePond's setup: (a) the bi-directional, non-rotating scan strategy; (b) the bi-directional, striped, 67° rotating scan strategy; (c) the height map of layer 22 from scan strategy (a); (d) the height map of layer 22 from scan strategy (b); (e) a graph showing how surface roughness increases with build layers for the non-rotated strategy; (f) a Fourier analysis of the surface heights of all 26 layers from Cube 1 (scan strategy a), and Cube 2 (scan strategy b), showing the more uniform spread across the latter [285]

2.6.3 - Scan Track Signals

Scan track signals are those that are emitted on the scale of vector, while the laser is active. These can be categorised into temperature profiles, process by-products, and air-borne acoustics.

2.6.3.1 - Temperature Profiles

High-speed imaging of the powder bed, during laser scanning, can be used to create a superimposed heat map, in which hot and cool spots can be linked to defects and other regions of interest.

Laguzza et al. [286] mounted an Olympus DSLR outside the viewport of a L-PBF machine printing 316L steel parts, which recorded the powder bed at a frame rate of 300fps, at a resolution of 1280×1024 pixels. All images were recorded in the visible range, using standard, consumer-grade camera settings. Using k-mean clustering, the authors managed to categorise all pixels from the entire image

stack into one of four groups: background pixels (not currently being scanned), normal melting, overheating, and single-point anomalies (such as spatter). A strong correlation was found between locations within the build that corresponded to overheated pixels, and geometric distortions due to thermal stress. A sample of these images are shown in Figure 2.87.

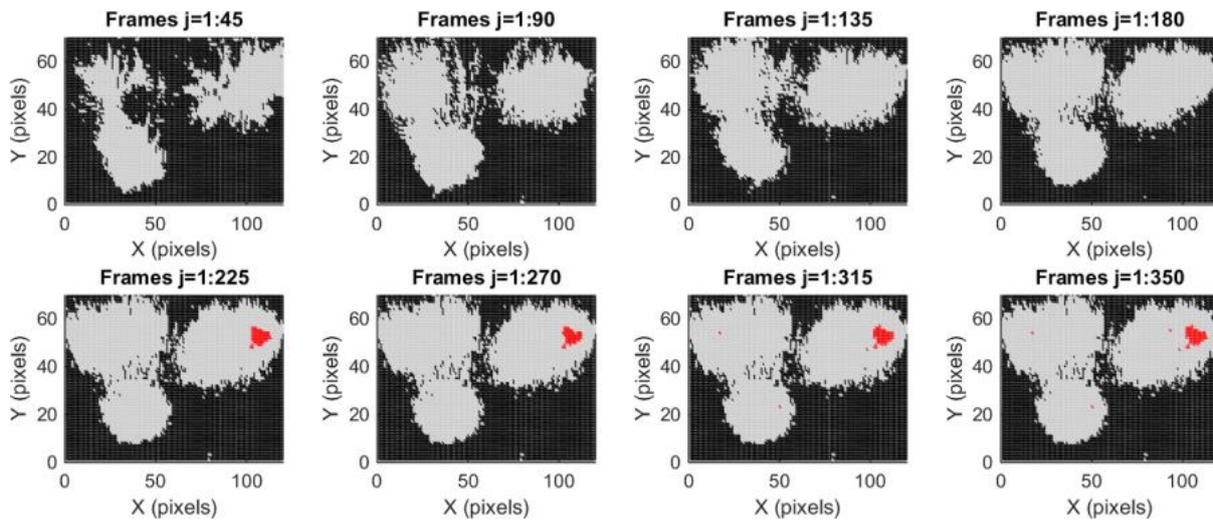


Figure 2.87

A series of cropped images from Laguzza's experiment. All images are from the same layer, but more data becomes available for k-means clustering with each added frame. The part being built is a three-legged strut system. Pixels with temperatures indicating background (no melting) are shown in black, those with normal melting are shown in grey, and those deemed to be overheated are shown in red. [286]

When recording in the visible spectrum, and if the sample set of images is sufficiently large, linkages can be made between the colour and intensity of light coming from melt pools, and the quantity of heat within them. Recording at a single wavelength in the infrared (IR) spectrum allows much more explicit conversions to be made between light intensity and temperature. Additionally, IR light can transmit better through smoke and plume, allowing for a clearer image.

Lough et al. [287] used an off-axis IR camera, filtered to 1450nm wavelengths, to record the production of 304L steel parts through an off-axis viewport. The camera recorded an 80×80 pixel section of the powder bed, corresponding to around $130 \mu\text{m}/\text{px}$, at a frame rate of 2500fps. It was found that, by tracking both the time that pixels spent above the estimated melting temperature, and the rate of decrease from their maximum temperature, it was possible to predict whether a location within a part would experience keyhole porosity, LoF porosity, or dense printing. The authors concluded that the system had the potential to warn of porous builds in real time, allowing for either parameter adjustment, or for builds to be cancelled. An example of these results is shown in Figure 2.88.

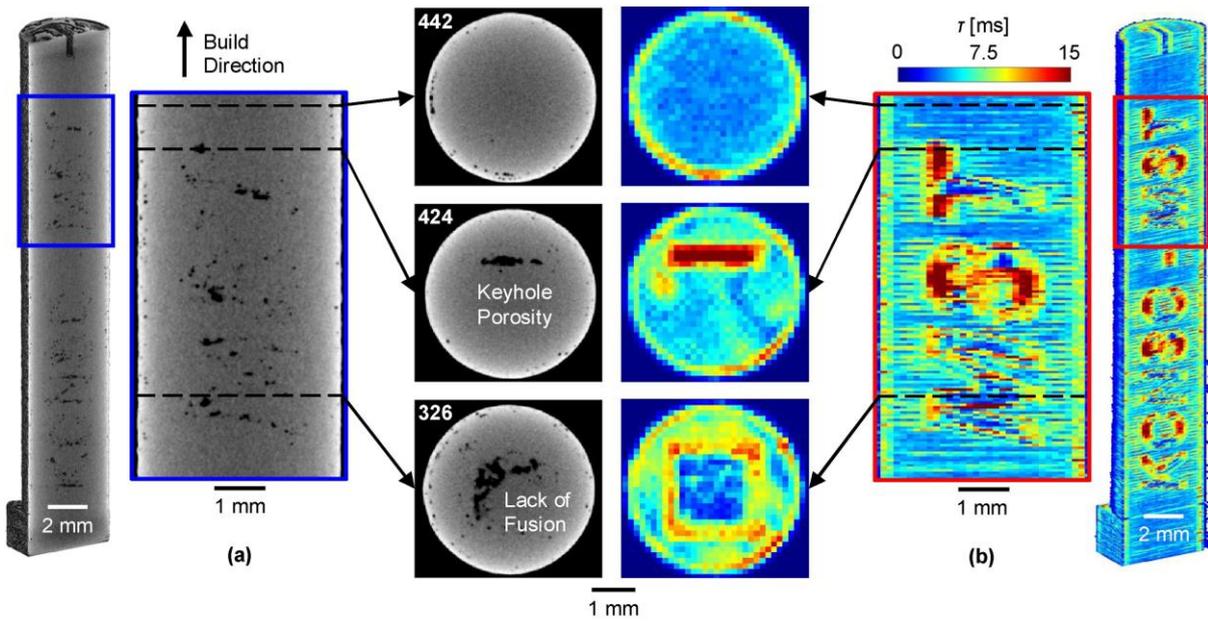


Figure 2.88

Results from Lough's setup: (a) pores identified via ex-situ XCT; (b) features identified by the time their pixels spent above the threshold temperature [287]

In order to determine the exact temperature of the body emitting IR radiation, the emissivity value of the body must be known. However, emissivity values are temperature dependent, shifting significantly for alloys used in L-PBF, as they're heated from solidus to liquidus and beyond. Therefore, in order to determine the exact emissivity value of a body, its temperature must be known. This circular dependency creates a fundamental obstacle for accurate temperature measurement in AM, forcing many researchers to use trends and comparisons in their work, rather than exact values. One method of overcoming this issue is to record radiation at multiple wavelengths simultaneously, which allows for the determination of temperature independent of emissivity, instead using the intensity ratios of the wavelengths.

Mitchell et al. [288] used two off-axis, high-speed cameras, filtered to 750nm and 900nm respectively, to record a 65×80 pixel slice ($21 \mu\text{m}/\text{px}$) of powder bed during the building of 316L steel parts. Images were recorded at 6500 fps. The temperature of each pixel in each image was estimated, over a range of 1100-2800°C. The size, shape, and cooling profile of each melt pool was then calculated, and compiled into a 3D volume image of the part's thermal history. It was found that pores between $30\mu\text{m}$ and $70\mu\text{m}$ in equivalent diameter could be detected with 54% to 74% accuracy, and pores with a greater equivalent diameter that $70\mu\text{m}$ could be detected with 100% accuracy. Images of an example melt track, and of the identified pores, are shown in Figure 2.89 and Figure 2.90, respectively.

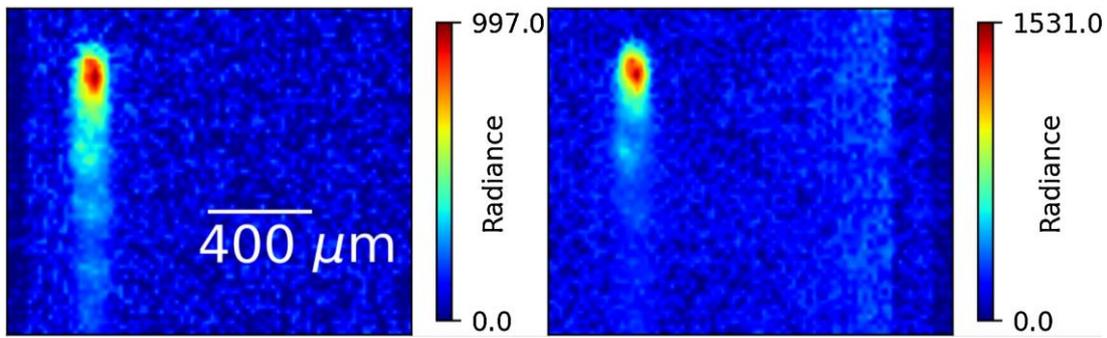


Figure 2.89

Mitchell's two cameras recording the same melt pool: left at 900nm, right at 750nm. Note how more detail from the cooling wake is included at the longer wavelength. [288]

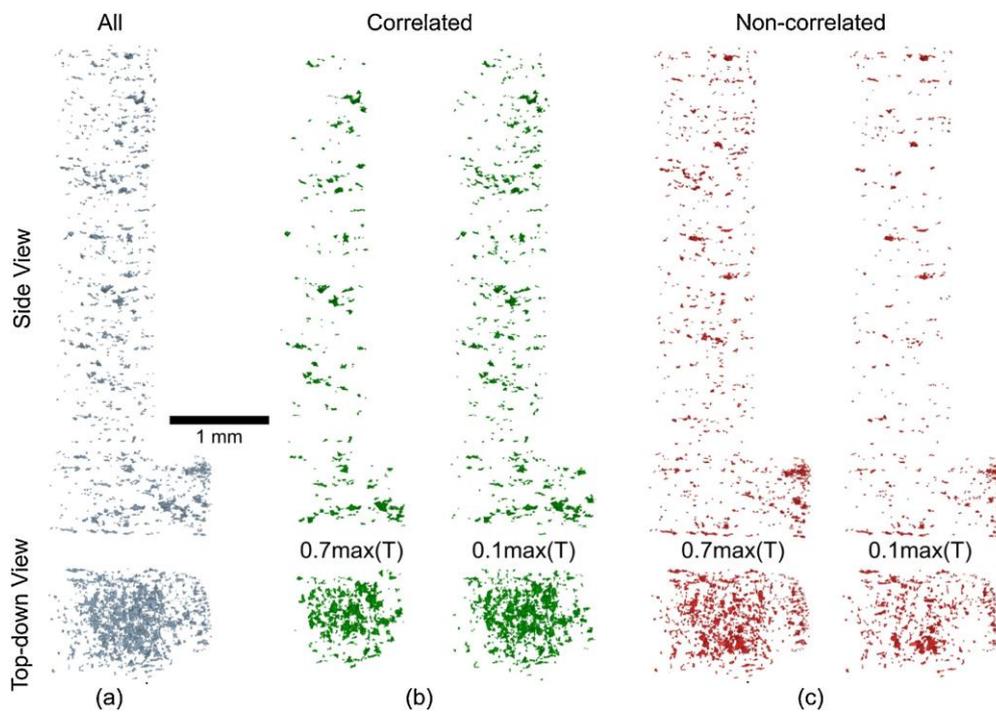


Figure 2.90

Reconstructed volumes, showing pores identified in Mitchell's parts: (a) all pores found via XCT; (b) pores found via XCT that can be linked to anomalies found via the CMOS cameras; (c) pores found via XCT that can't. The size of the anomalies shown in (b) and (c) are determined as the contour lines around the peak temperature (T_p), at both $0.7T_p$, and $0.1T_p$. [288]

2.6.3.2 - Process By-Products

Valuable information about the state of the building process can be gained, by analysing the spatter and the plume that is emitted from the melt pool. These by-products can deflect or absorb laser energy, can result in contamination on the powder bed, or can indicate process instability or improper process parameters.

Yang et al. [289] monitored a build of IN-625 parts with two off-axis cameras, one filtered to a wavelength of 1.3 μ m to 1.6 μ m, and the other recording in the visible range. By combining these two imaging systems, it was possible to both identify molten projectiles that were ejected from the melt pool, and continue tracking these projectiles even if they cooled below solidus temperature. The authors were able to therefore quantify the amount of spatter landing on the powder bed for each part, and correlate these measurements to the degradation of mechanical properties, the presence of pores, and the delamination of layers further along in the build process. Examples of melt track images, and melt pool analysis, are shown in Figure 2.91 and Figure 2.92, respectively.

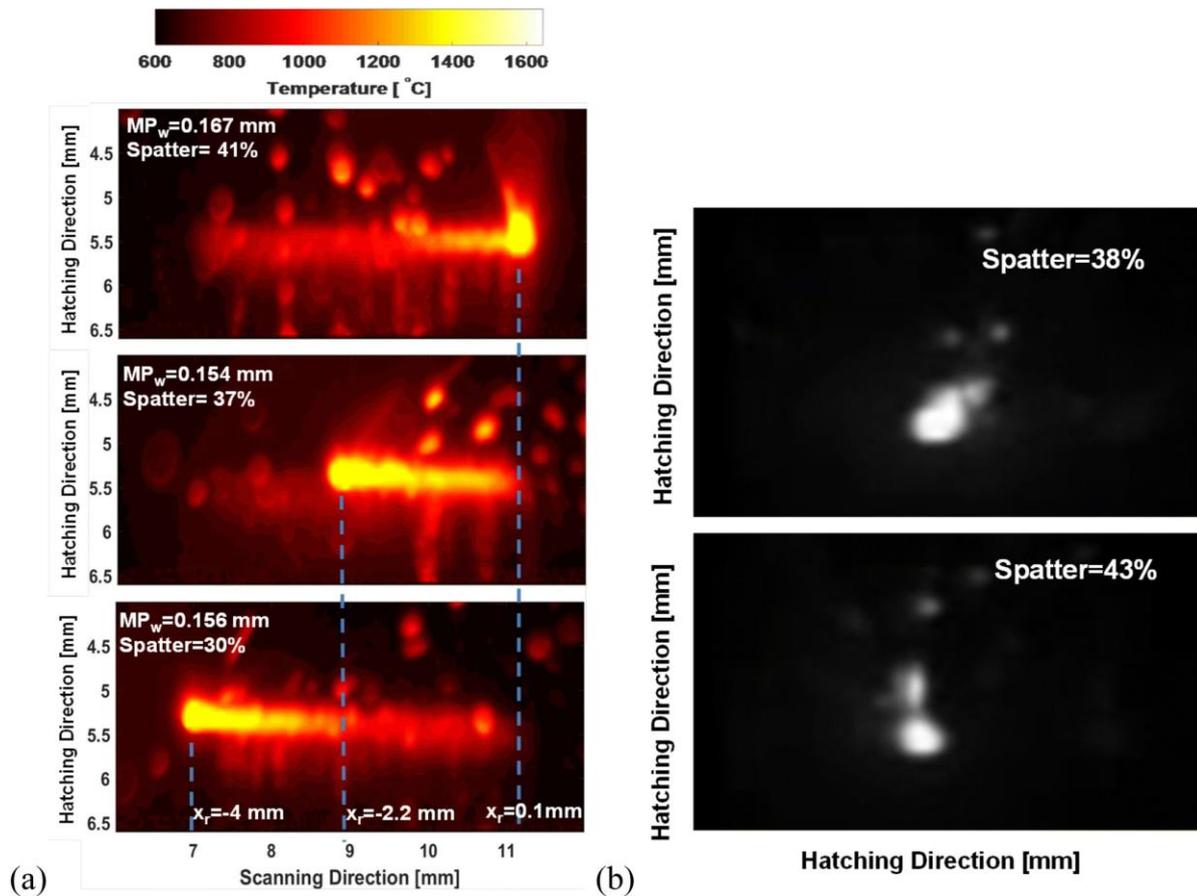


Figure 2.91

A comparison of the data recorded by the filtered (a) and visible (b) cameras in Yang's experiment. Note how the filtered images contain much richer detail regarding temperature profiles, spatter, and the melt pool's wake. [289]

A similar technique to that used by Yang et al., which involves the use of two cameras to monitor the powder bed, is stereoscopic imaging. This method involves dual cameras, positioned at a small distance from each other, recording the L-PBF process from slightly different angles. By analysing the disparity between images taken simultaneously, depth information can be obtained, in the same manner as occurs for human eyes.

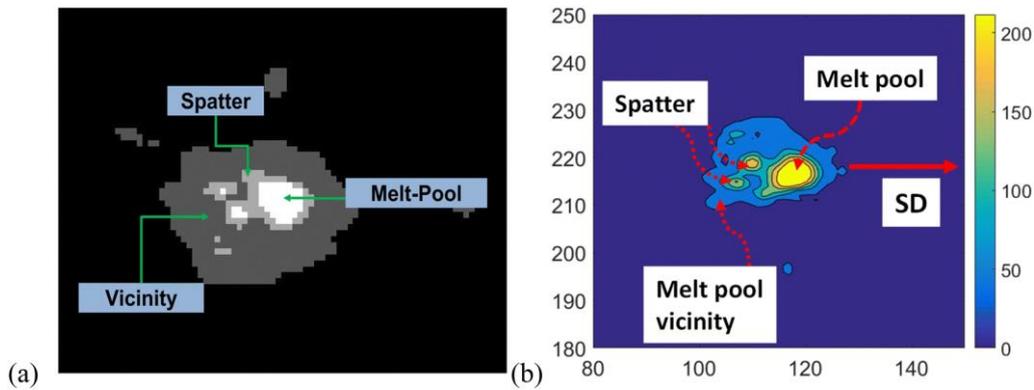


Figure 2.92

Yang identified spatter particles using a *k*-mean clustering method on images from the visible range camera (result on left), and tracked their cooling with the filtered camera (right) [289]

Eschner et al. [290] utilised high-speed stereoscopic imaging in the visible range, to track spatter particles in during a 316L steel build, as shown in Figure 2.93. The system identified glowing particles, and determined their velocities and trajectories. It was found that an increase in spatter velocity can be a valuable predictor of the onset of evaporation within the melt pool, and also used to estimate the depth of the vapour depression. Therefore, when a certain melting mode is desired (e.g. steady-state keyholing, with little evaporation and no collapse or height fluctuation), an analysis of spatter velocity can be a useful tool.

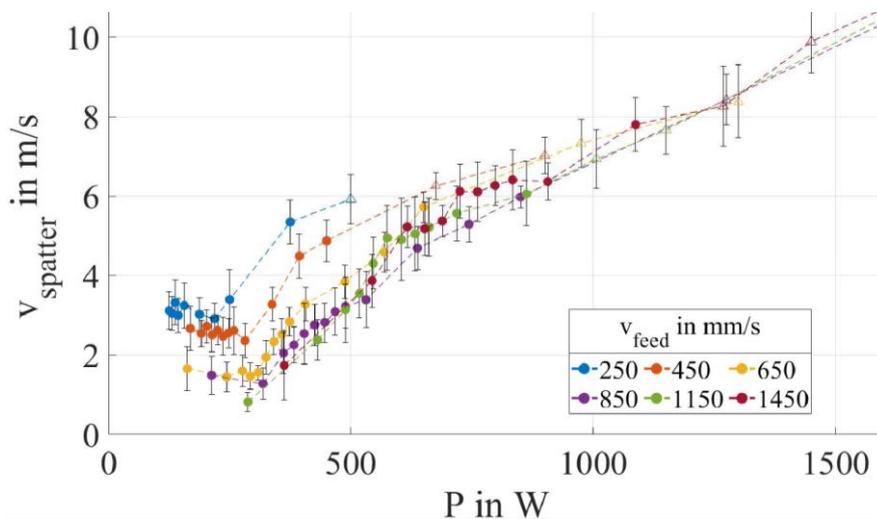


Figure 2.93

Eschner tracked the velocity of spatter particles ($v_{spatter}$) for various laser powers (P) and laser speeds (v_{feed}) [290]

The previous discussion has dealt solely with spatter. Detecting the vapour plume is also possible with IR camera technology, as this cloud of ionised gas, vaporised metal, and condensed nanoparticles also emits IR radiation. Radiation is emitted in the $1\mu\text{m}$ to $10\mu\text{m}$ wavelength band, far longer than that which is usually filtered for with IR monitors of L-PBF. Also, due to the complex and non-uniform emissivity of the fast-moving material, an accurate temperature estimation is extremely difficult.

Grasso et al. [291] used an off-axis FLIR camera, filtered for $8\mu\text{m}$ to $9\mu\text{m}$, to analyse the melt pool and plume of a build of pure zinc powder. The camera recorded at 50 fps, at a resolution of 320×240 pixels. Images were analysed to separate out three distinct regions: the laser heated zone, the spatter, and the plume. The area of the plume in each image was determined, as well as a rough approximation of the plume's temperature, based on pixel intensities. These two metrics were shown to serve as valuable proxies for the extent of evaporation and heat accumulation within the melt pool. Sudden increases in plume area were found to correlate with keyhole collapse, and could even serve as an early warning sign of part failure associated with extreme overheating. This experimental setup and workflow are shown in Figure 2.94 and Figure 2.95, respectively.

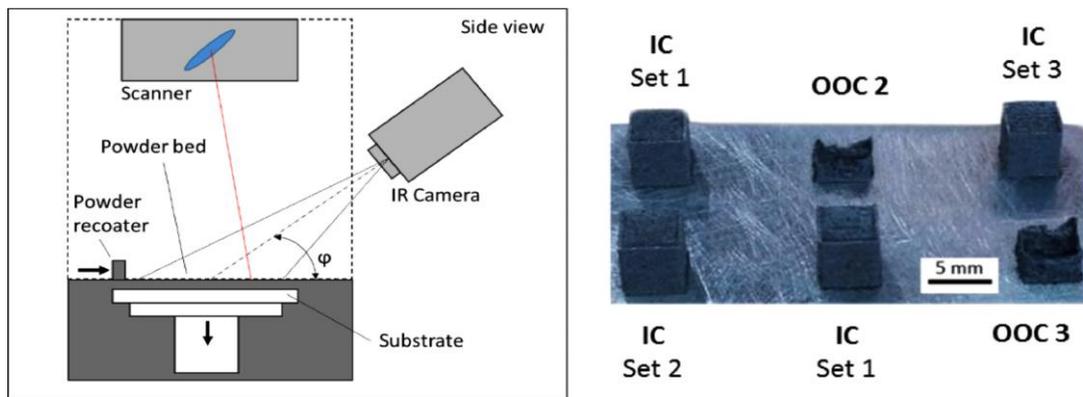


Figure 2.94

Grasso's camera setup, and the 4 in-control (IC) and 2 out-of-control (OOC) parts that were built [291]

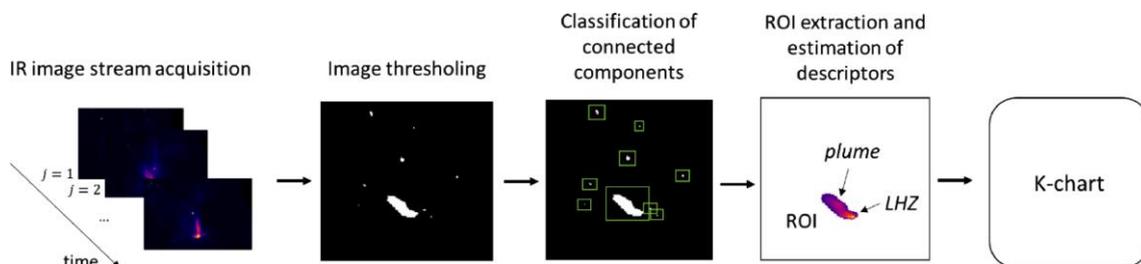


Figure 2.95

Grasso's image workflow thresholded the images; extracted and separated the plume, the melt pool, and the spatter; and sent the result to a k-means clustering method [291]

Another method that can be used to detect the plume, is Schlieren imaging. This technique is used to visualise transparent fluid flows by detecting small changes in refractive indexes, caused by temperature, pressure, or density fluctuations.

Bidare et al. [223] used Schlieren imaging at analyse the vapour plume of 316L steel parts, being built in an argon atmosphere. The authors projected a parallel, high-intensity, rectangular beam from a tungsten lamp, across the powder bed. The beam was reflected into a high-speed camera, which was fitted with filters and polariser to block out all laser light, IR radiation from the melt pool, and any glare or plasma incandescence. By monitoring the plume with this setup, the authors

managed to track the extent to which the laser was absorbed or deflected by the plume, as well as the extent of spatter and denudation occurring in and around the melt pool below. These metrics could then be correlated with porosity and surface roughness in the finished parts. Images of a scan track are shown in Figure 2.96.

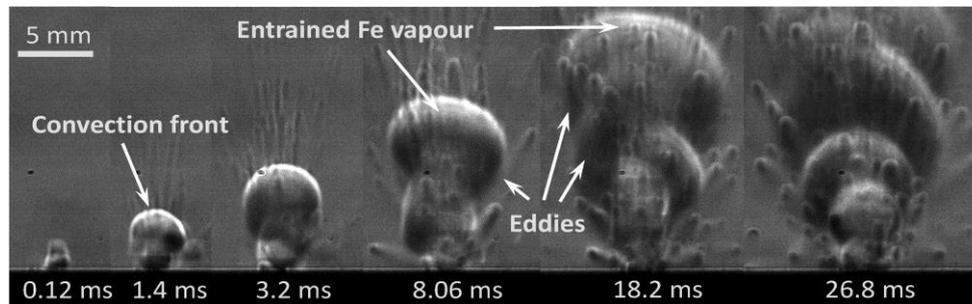


Figure 2.96

Images from Bidare: heated gas with vaporised iron, rising from a scan track that's coming towards the viewing direction, starting at $t=0$ ms, at 100W and 500 mm/s [223]

2.6.3.3 - Air-Borne Acoustics

Phenomena that occurs at or below the powder bed surface, such as micro-cracking, spatter impacts, or keyhole collapse, can result in the creation of pressure waves that transmit through the build chamber as sound. The sensitivity required to detect these minute waves is beyond the capabilities of many standard microphones, and so a Fibre Bragg Grating (FBG) is often used instead. This device consists of an optical fibre, that allows most wavelengths of a broadband source to transmit through it, while reflecting a very specific wavelength. As mechanical pressure waves impact the fibre, the strain causes tiny variations in its refractive index, and therefore in the wavelength of light that is reflected. By tracking these variations in which exact wavelength of the broadband is reflected, the frequency and intensity of sound waves can be determined.

Shevchik et al. [292] placed an FBG inside the build chamber, while printing 316L steel parts. The FBG was continuously illuminated by a narrowband laser (1547 ± 0.01 nm), which failed to reflect while the fibre was undergoing strain due to acoustic waves. The reflected signal was sampled at 1MHz, and fed into a machine learning system. Parts were built at one of three scanning speeds, that resulted in high (1.42%), medium (0.30%), and low (0.07%) porosity. The machine learning system managed to correctly predict which of the three speeds the acoustic signals were generated from, with an accuracy ranging from 78% to 91%. The authors concluded that the physical processes associated with pore formation correspond with distinct acoustic signals, and that these signals, if correctly interpreted, can provide real-time feedback on build quality.

2.6.4 - Melt Pool Signals

This subsection explores signals which are emitted at the scale of the individual melt pool, categorised by the metric being extracted from the data.

2.6.4.1 - Melt Pool Radiation Intensity

A pyrometer is analogous to a single-pixel camera, that is tuned to detect thermal radiation emitted from hot surfaces. Even basic pyrometers can offer far greater sampling rates and measurement resolution than most high-end cameras. Pyrometers can be mounted off-axis, as with all the camera setups described previously, or co-axially, whereby they share the same optical path as the laser, and are therefore always focussed upon the melt pool, whenever the laser is active. In order to allow co-axial sensing, the optical path must incorporate a dichroic mirror: i.e. a mirror which exhibits very different reflectivity and transmission properties, depending on the wavelength of the light.

Forien et al. [293] built samples of 316L steel, with a pyrometer behind a dichroic mirror that reflected the 1070nm laser wavelength, while transmitting radiation between 1500 to 1700nm. The pyrometer operated at a 100 kHz sampling rate, and picked up radiation emitted from a spot of around 635µm diameter, centred at the laser's centre. This spot size ensured that radiation from the entirety of the melt pool was recorded. The beam position coordinates taken from the galvanometers controlling the mirror, were recorded simultaneously, allowing for the XY location of each pyrometer datapoint to be known to within ±10µm.

Samples were produced with a wide range of process parameters, and then internally imaged with XCT to search for porosity. By cross-referencing the pyrometer measurements with the locations of identified pores, it was found that strong correlations existed. In locations where the normalised pyrometer value was between 0 and 0.5, there was a 5% chance of a pore being found. In locations with values between 0.5 and 0.75, there was a sharp increase in the probability of a pore being present. In locations with values between 0.75 and 1, the chances of a pore being present was 95%. It was reasoned that the pyrometer system has successfully managed to differentiate between conduction, transition, and keyhole modes, based on these three categories. These pyrometry results are shown in Figure 2.97.

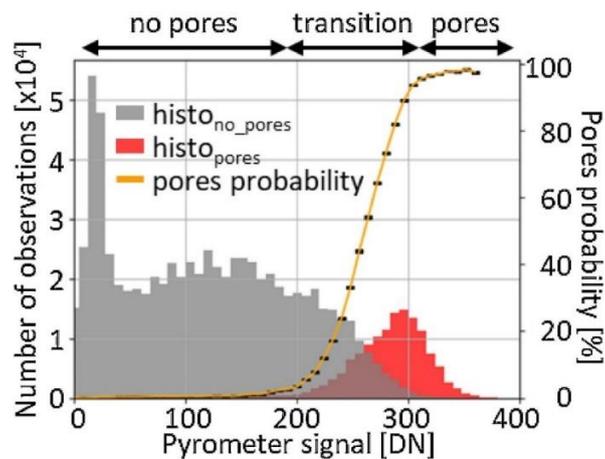


Figure 2.97

Forien's pyrometry results: as the signal increases from 0 to 0.5 (here labelled in the non-normalised form 0 to 200), very few pores are found. Between 0.5 to 0.75, a sharp transition in the probability of porosity occurs. From 0.75 to 1, the probability of finding a pore is 95%. [293]

As discussed previously, conversion from intensity to temperature is complicated by the circular dependency of temperature and emissivity, and this complication can be overcome by using multi-wavelength detectors. Samaei et al. [294] used a dual-pyrometer system to coaxially monitor the mean temperature of melt pools during the building of overhangs in AlSi10Mg parts, with a sampling rate of 200 kHz, and a field of view with a diameter of 5mm. The system correctly detected the melting mode of the pool, as with that of Forien et al. [293] mentioned previously. However, scanning shallow overhang sections, where the lack of solid material for heat to conduct into should result in higher melt pool temperatures, the dual-pyrometer system consistently reported lower temperatures. The authors concluded that these counterintuitive measurements resulted from the fact that all information about the melt pool's size, shape, and temperature profile is lost to the pyrometer. In the case of shallow overhangs, where melt pools are hotter but also narrower and deeper than elsewhere, pyrometry data can therefore be misinterpreted by the analysis workflow as cooler melt pools. Some pyrometer results from the overhang section, and melt pool approximations, are shown in Figure 2.98 and Figure 2.99, respectively.

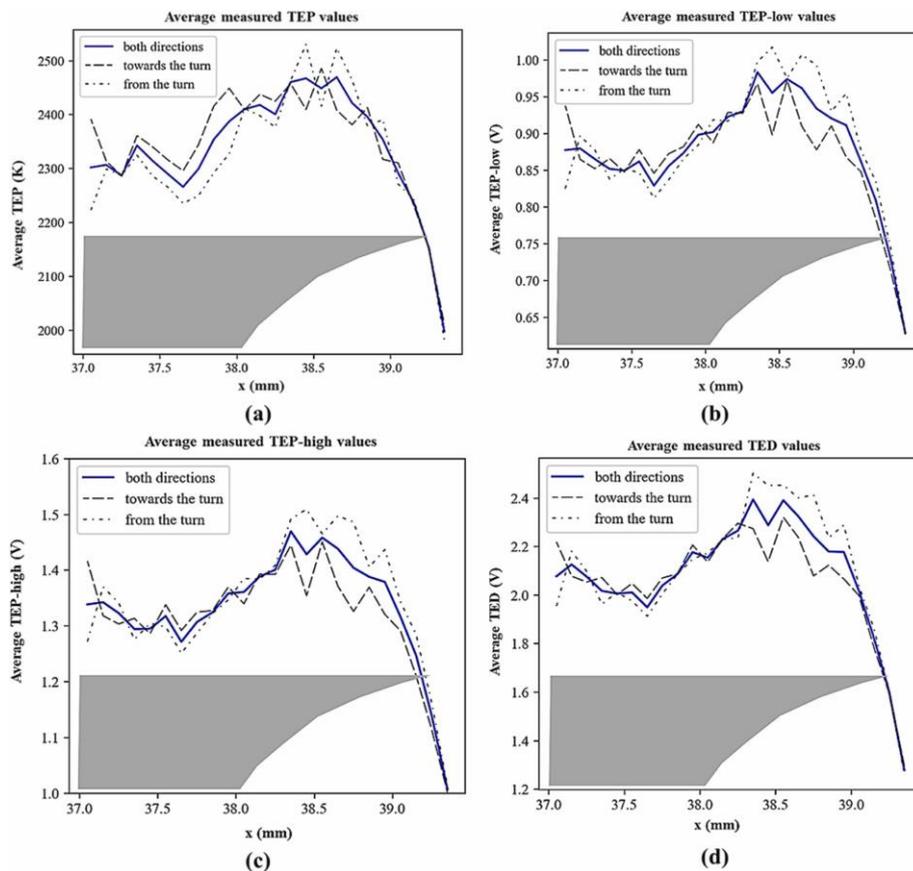


Figure 2.98

Samaei measured the intensity of thermal emission as a melt pool approached (heavy dashed line) and retreated from (light dashed line) an overhang edge. Measurements were recorded in terms of: temperature, as calculated from the dual-pyrometer intensity ratio (TEP); intensity of the 680nm pyrometer (TEP-low); intensity of the 700nm pyrometer (TEP-high); and a separate pyrometer that measured in broadband (TED). All results indicated that temperature dropped as the melt pool approached the overhang edge, which is a counter-intuitive finding. [294]

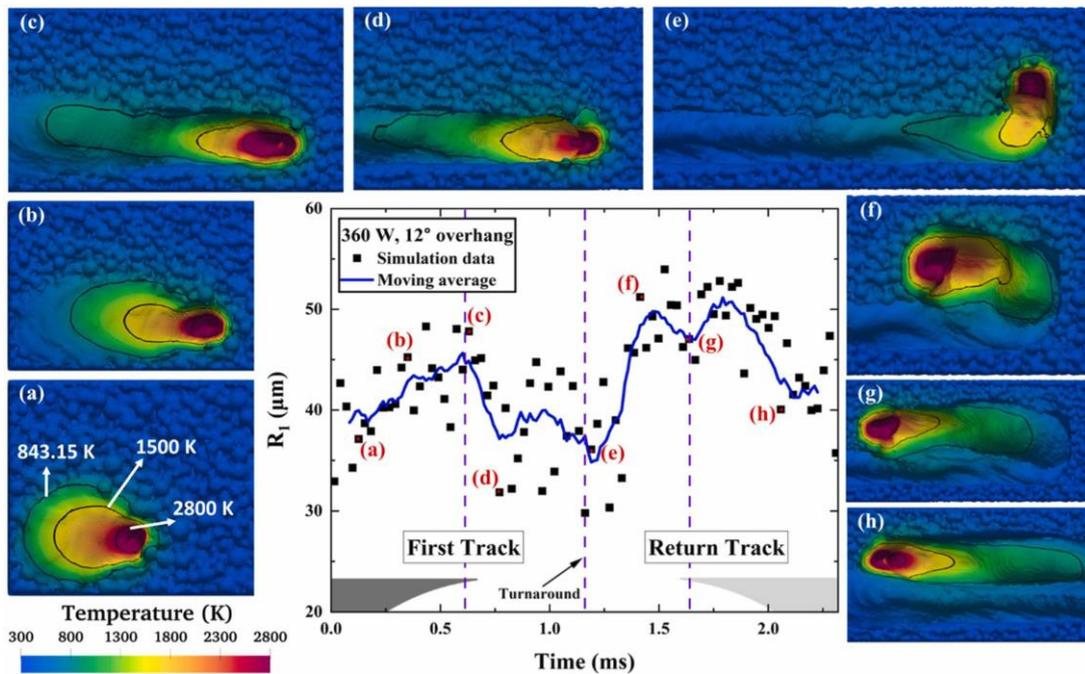


Figure 2.99

A CFD model from Samaei shows how the melt pool morphology changes as it approaches and retreats from the overhang edge, and how these changes result in the counter-intuitive temperature measurements [294]

Although less commonly implemented, it is possible to mount a pyrometer off-axis within the build chamber. Bisht et al. [295] produced Ti-6Al-4V parts, while monitoring the process with a single-wavelength, off-axis pyrometer. The system recorded wavelengths between 1150 to 1800nm, with a field of view effectively covering the entire build platform. The laser XY position was recorded alongside, which was cross-referenced with the pyrometer readings to produce the coordinates of points in the build that experienced sudden fluctuations in thermal emissions. These fluctuations were believed to be the result of pore creation, and parts that incorporated more of them were found to exhibit lower elongation before failure. Images of the thermal emission fluctuations, and the resulting defect structure, are shown in Figure 2.100 and Figure 2.101, respectively.

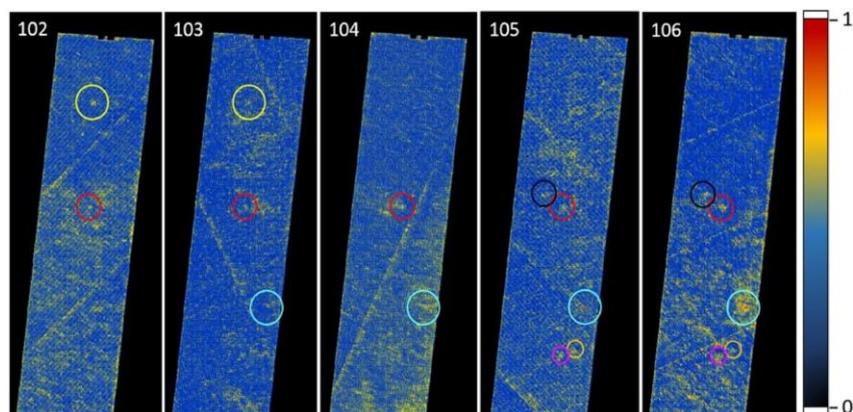


Figure 2.100

Bisht identified several 'events' in the pyrometer data, here shown in XY for layers 102 to 106 [295]

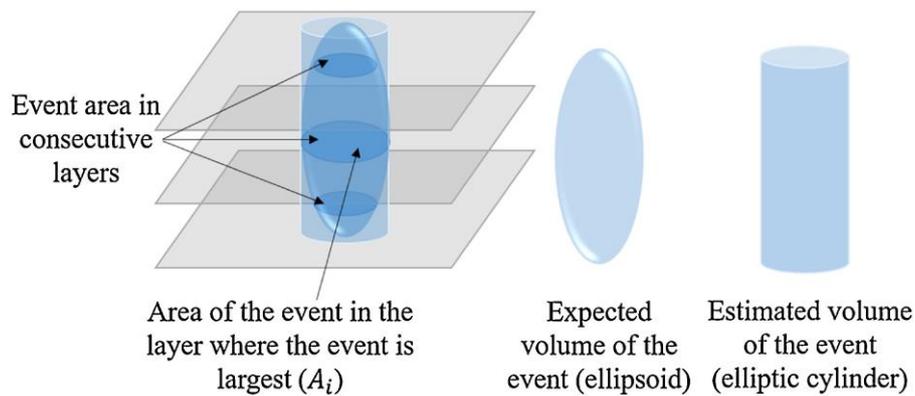


Figure 2.101

Bisht assumed that events would be ellipsoid in nature, but for simplicity estimated them as a cylinder, with a width determined by the maximum event area recorded [295]

2.6.4.2 - Melt Pool Size and Shape

As well as pyrometers, researchers have also successfully mounted full cameras coaxially. This experimental setup allows for full images of the melt pool to be recorded, with data on the size, shape, and temperature profile being possible, depending on the camera type.

Kwon et al. [296] installed a high-speed, visible light camera behind a dichroic mirror that transmitted wavelengths between 450 to 900nm, while a 1075nm laser produced parts from 316L steel. The camera recorded at 2500 fps, with a frame size of 512×512 pixels, corresponding to a $10\text{mm} \times 10\text{mm}$ area around the laser's centre. Each pixel recorded a greyscale intensity, between 0 and 255, which was then binarised at a threshold value of 125. The images were used as training and testing data inside a machine learning system, with corresponding microscopy images that revealed part density. It was found that, once the system was trained, porosity could be predicted from these simple binary images depicting melt pool size and shape with an accuracy of between 78% and 93%. Sample melt pool images, and histograms of intensity results, are shown in Figure 2.102 and Figure 2.103, respectively.

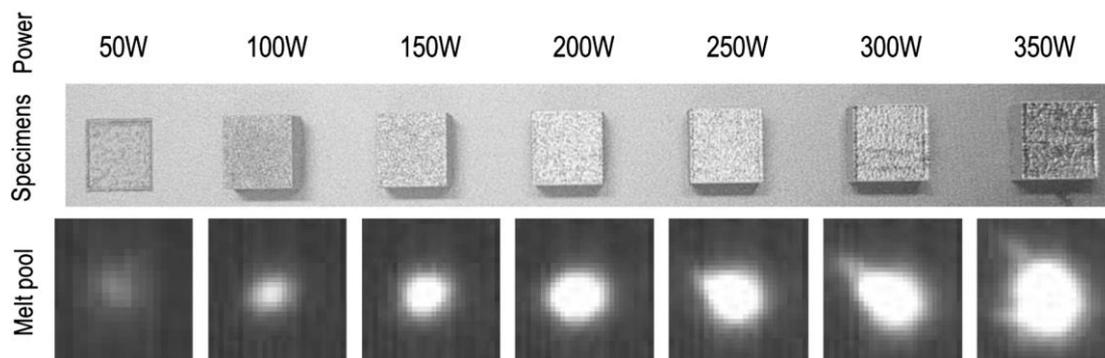


Figure 2.102

A comparison of Kwon's melt pool images, recorded in visible light for a range of powers, and the surfaces that result [296]

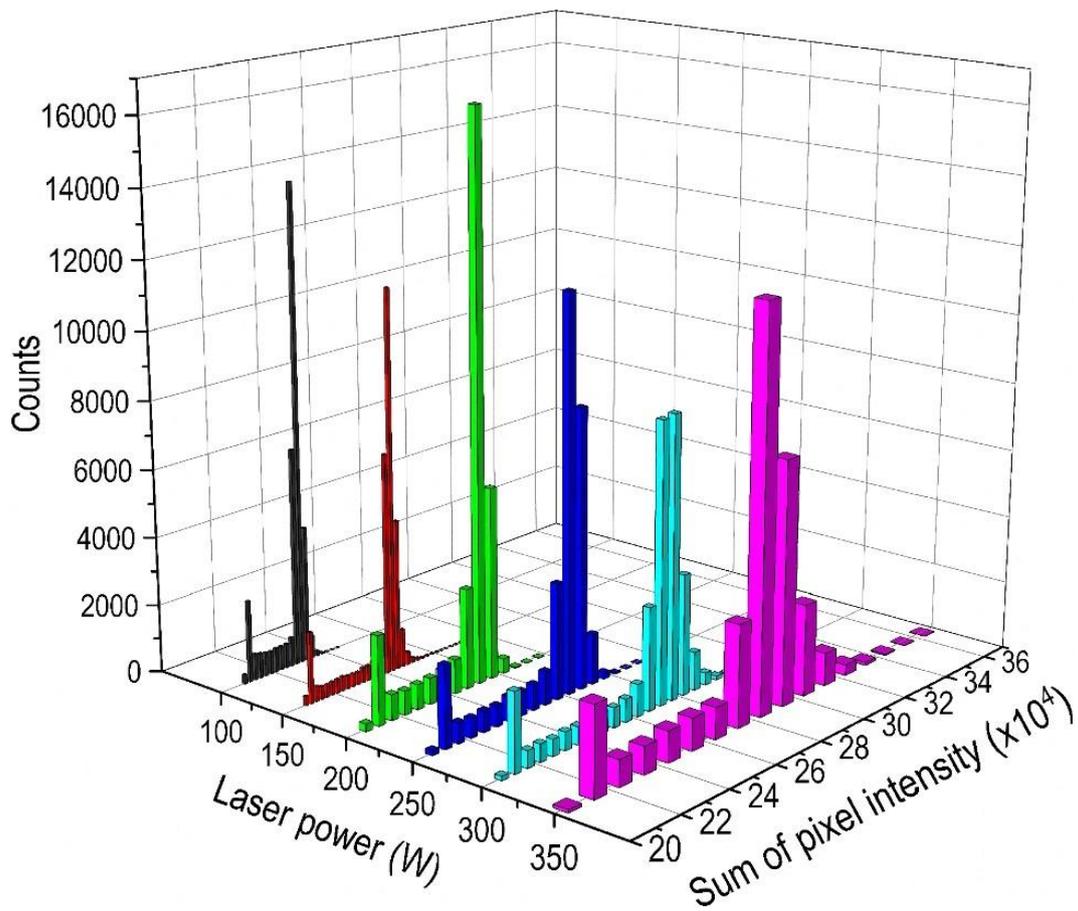


Figure 2.103

Plotting a histogram of the sum of pixel intensity for each image, in each power, shows a clear differentiation in Kwon's visible light results [296]

As with the off-axis camera setups discussed previously, in order to convert pixel intensity values to temperatures, IR cameras, or ideally dual-wavelength cameras, are required. Vasileska et al. [297] mounted a high-speed camera, filtered to record between 850 and 1000nm, behind a dichroic mirror that transmitted between 400 and 1000nm, while building in 316L steel. The camera recorded at 1200 fps, with resolution of 1280×2004 pixels, and a frame size of 4.3×4.3 mm². By monitoring the size, shape, and temperature profile of the melt pool, the authors managed to successfully correlate melt pools that experienced overheating, with regions in the build that exhibited swelling defects. Since the locations of overheating and swelling were systematic (at sharp part edges and the last scanned vector), it was theorised that, with more adaptive control systems, process parameters could be adjusted to correct or mitigate these defects in the layer after they're detected. The experimental setup and results are shown in Figure 2.104 and Figure 2.105, respectively.

Hooper [71] installed two dichroic mirrors in a monitoring system, when building in Ti-6Al-4V. The first lay along the laser's optical path, and allowed for transmission of the 1070nm laser, while reflecting off IR radiation emitted from the melt pool, between 450 and 1000nm. The second split this IR radiation at 800nm, allowed for the transmission of longer wavelengths (which were collected by a high-speed camera filtered to 950nm), and reflecting shorter wavelengths (which were collected

by a high-speed camera filtered to 700nm). The cameras recorded at 200,000 fps, with a resolution of 129×128 pixels, corresponding to a field of view of $2.5\text{mm} \times 2.5\text{mm}$. Melt pool sizes, shapes, temperature profiles, and cooling rates could be detected with high accuracy. Thermal gradients in the range of 5 to 20 K/ μm were recorded, corresponding to cooling rates in the range of 1 to 40 K/ μs . Hooper demonstrated with this system that the beginning and end of scan vectors, as well as turns in hatching, were more thermally volatile. Cooling rates at overhang section were shown to be significantly slower than bulk section, due to a lack of conduction. It was concluded that this system, which exhibited extreme accuracy in both the spatial and temporal domain, held the potential to precisely detect anomalous melt pool characteristics, and warn against defect creation before irreversible damage was done to the part. Some example results are shown in Figure 2.106.

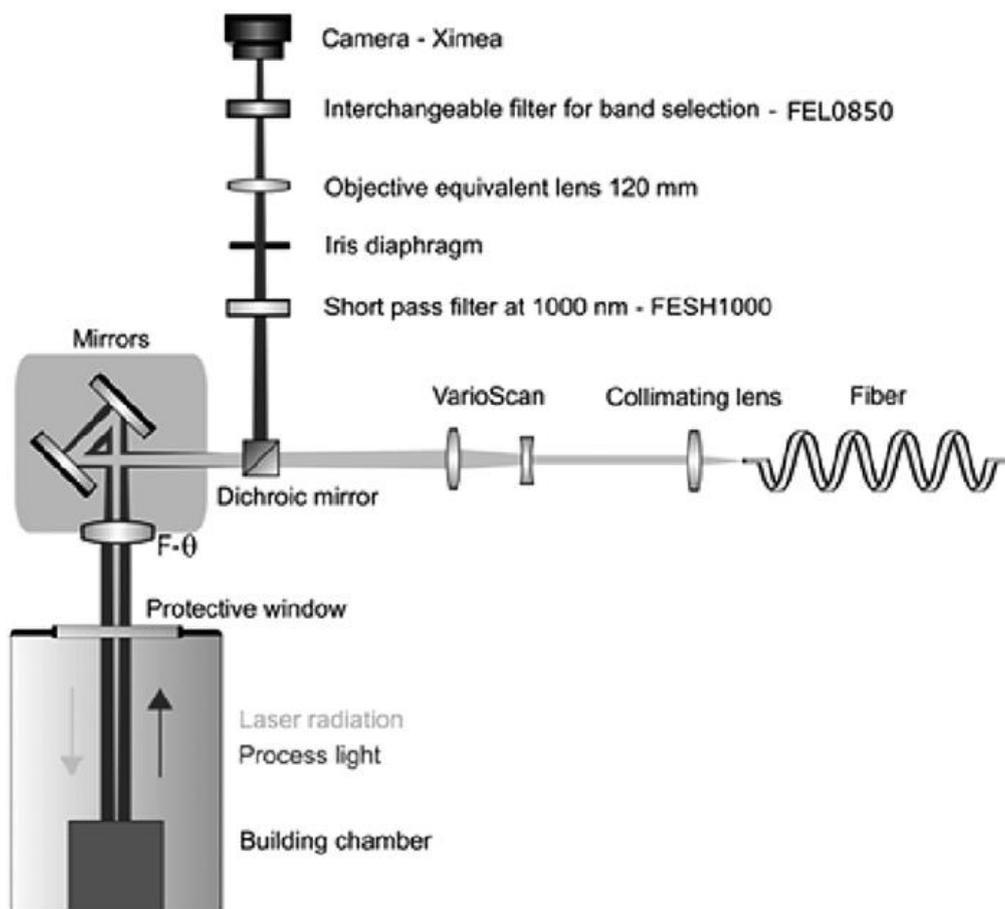


Figure 2.104
Vasileska's coaxial IR setup [297]

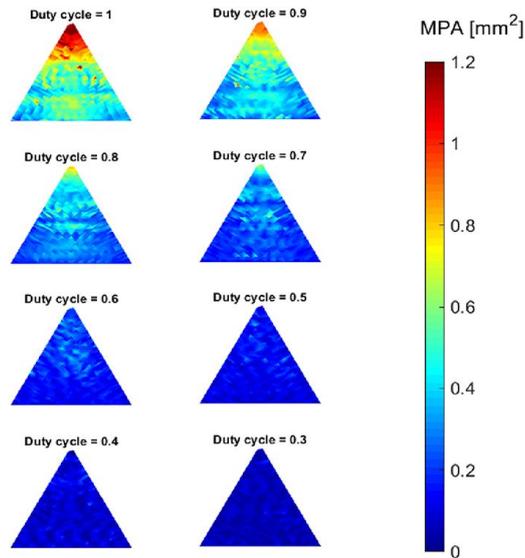


Figure 2.105

Melt pool areas were calculated, from images recorded during the building of a triangular part, at various duty cycles (fraction of time that a pulsed laser is active) by Vasileska [297]

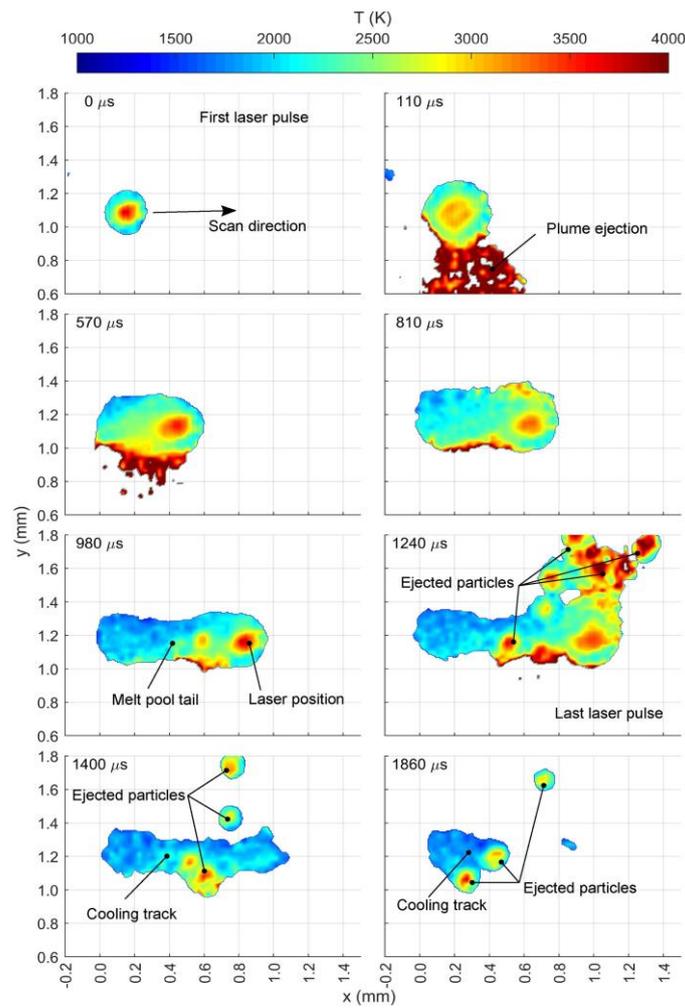


Figure 2.106

Results from Hooper, showing a laser pulse activating, approaching an overhang feature, deactivating, and its pool cooling. The plume, spatter, and melt pool cooling profile can all be discerned. [71]

2.6.4.3 - Melt Pool Emission Spectrum

A similar piece of technology to those discussed previously, but one that is used for a distinct purpose, is the spectrometer. A spectrometer consists of a diffraction grating, which deflects different wavelength of an incoming broadband radiation off at different angles, and a sensor, in which each pixel is aligned to detect a specific wavelength. By monitoring the intensity readings of all these pixels, it is possible to reconstruct the intensity-wavelength spectrum of the radiation being emitted from the melt pool. Using techniques discussed previously, it is possible to decipher a mean melt pool temperature from the wide continuum of the spectrum. It is also possible to search for sharp spikes within the spectrum, which refer to excited electrons falling from one energy level to another in certain atoms, and to determine melt pool chemical composition from these. Additionally, by tracking changes in the spectrum over time, it is possible to monitor for melt pool temperature and morphology fluctuations.

Lough et al. [298] used a coaxial spectrometer to analyse the melt pool emission spectrum, while building in 304L steel. The dichroic mirror reflected wavelengths below 900nm into the diffraction grating, which deflected wavelengths between 400 to 700nm onto the sensor. It was found that the strongest signal within the spectrum corresponded with that of chromium (at 520.6nm), and that monitoring this signal served as a useful proxy for melt pool temperature and geometry. The results are shown in Figure 2.107.

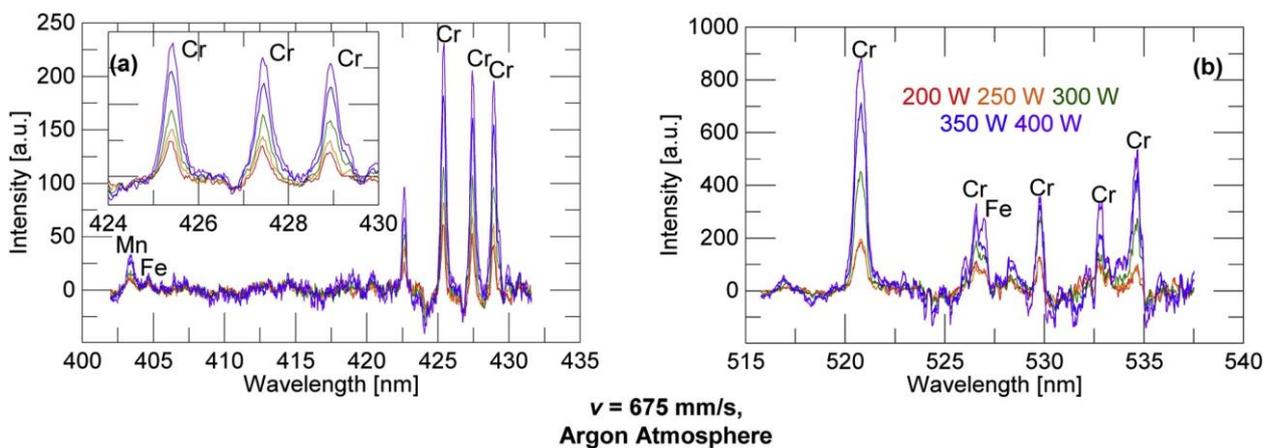


Figure 2.107

The emission spectra of Lough's 316L parts, plotted at various laser powers, and separated into two wavelength charts, with elemental peaks highlighted [298]

Ziefuss et al. [299] used an off-axis spectrometer setup, to monitor emissions while building in a neodymium-iron-boron-based alloy. The spectrum was analysed in the 200 to 550nm range. The authors performed a composition analysis on the powder before building, and compared this to the spectral analysis obtained during building. It was found that, due to the varying tendencies of elements to vaporise, the melt pool composition differed from the feedstock composition. Iron composition fell from 75.4 atomic % in feedstock to 74.0 at. % in the melt pool, for example, while

neodymium composition rose from 7.5 at. % to 8.7 at. %, as shown in Figure 2.108. These values were validated by a final compositional analysis of the finished part.

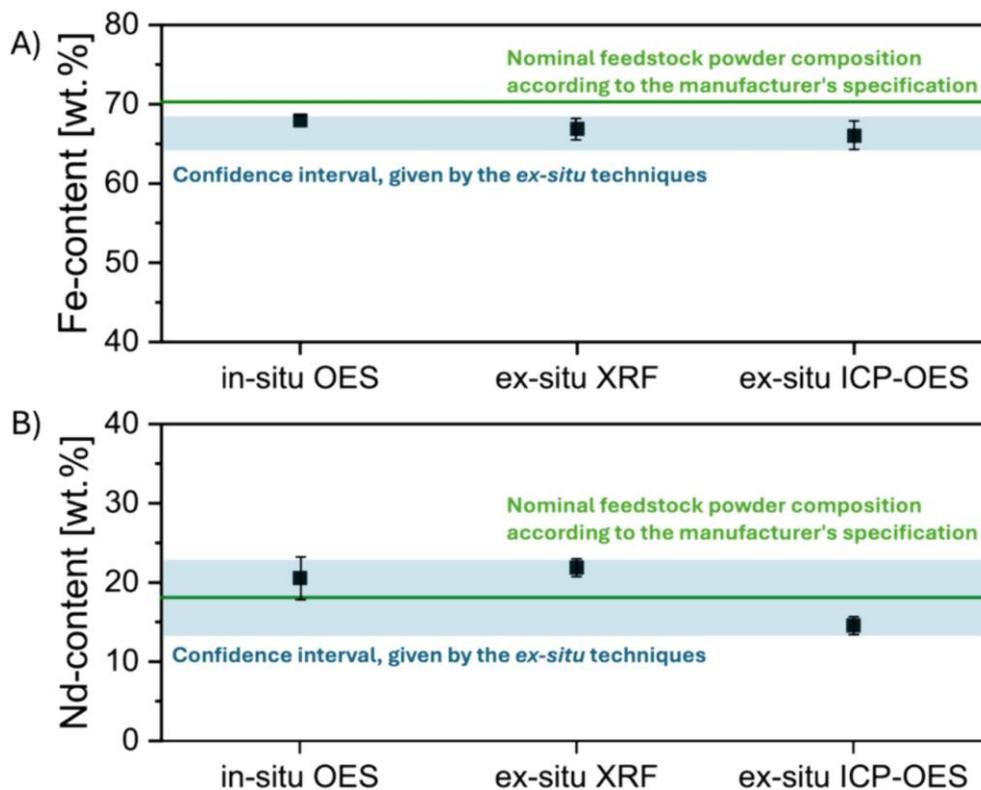


Figure 2.108

Comparing the iron content (top) and neodymium content (bottom), as measured by Ziefuss' in-situ spectroscopy, ex-situ X-ray Fluorescence (XRF), and ex-situ spectroscopy. The green line is the manufacturer's specifications, and the blue bar is the experimental confidence interval. The deviations from the specifications are attributed to the varying tendencies of different elements to vaporise. [299]

2.6.5 - Sensor and Signal Discussion

Over the previous few subsections, various in-situ data-gathering techniques have been presented. While some of these (such as acoustic monitoring or LCI) stand as separate categories, the majority can be classified simply as:

- Either cameras or pyrometers
- Mounted on-axis or off-axis
- Imaging in a single spectrum, in dual spectra
- In the visible spectrum, the infra-red spectrum, or both

Considering that these techniques account for the majority of in-situ research, and also the experimental work in this project, a comparative analysis of their strengths and limitations, and the information that they can and cannot reveal, is warranted.

2.6.5.1 - Cameras vs Pyrometers

In this context, a camera refers to a 2D array of pixels, capable of transforming incoming light into electrical signals, to produce a series of images that provides both spatially-resolved and temporally-resolved data on the L-PBF process. A pyrometer refers to a single-pixel camera, which provides temporally-resolved data only.

2.6.5.1.1 - Camera Strengths

Cameras can provide 2D (or 3D in the case of stereoscopic imaging) maps of the melt pool, the scan track, or large section of the powder bed. Melt pool dimensions and temperature profiles can be calculated from these maps, from which cooling curves can be extracted. They can capture a diverse range of process dynamics, such as spatter trajectories, plume shape, and even the motion of powder particles. Depending on user requirements, consumer-grade, off-the-self camera setups can provide lots of valuable data, without any complex designing or installation.

2.6.5.1.2 - Camera Limitations

The major limitations of cameras, is the processing time and storage requirements for the large data quantity that is output from the pixel array. Frame rates are often limited to a few kilohertz, at most, which may be too slow to capture rapid fluctuations in melt pool temperature, and may lead to motion blur as spatter particles are ejected at high speed, or keyholes fluctuate rapidly. Precise focusing is required, to tune the array of pixels to capture data from a specific location on the powder bed only, and plumes or spatter can often pass across the frame as out-of-focus blur. Calculating exact pixel temperatures can be difficult or impossible, due to the circular dependency of temperature and emissivity. Additionally, cameras can be very expensive pieces of equipment to purchase.

2.6.5.1.3 - Pyrometer Strengths

The major advantage of a pyrometer lies in the device's simplicity. Sampling rates of tens or hundreds of kilohertz are possible. This very fine temporal resolution can provide a measure of melt pool temperature that accounts for rapid keyhole oscillations and melt pool instabilities. Pyrometers are also physically small and simple devices, that are relatively inexpensive, and require only minimal processing power to function. Should the pyrometer be mounted to maintain focus on the laser incident spot, the single-pixel nature of the data allows for a robust tracking of the melt pool's peak temperature.

2.6.5.1.4 - Pyrometer Limitations

The single-pixel nature of a pyrometer removes all context from the melt pool temperature measurement. Pool dimensions, spatter, plume, track overlap, and the temperature profile of the molten alloy, are all lost. As with cameras, the circular dependency of emissivity and temperature makes the calculation of an exact temperature value difficult or impossible. The alignment of a pyrometer is also challenging: very small deviations from perfect can cause the device to be focused

upon the melt pool's periphery, or the melt pool's entirety, rather than its centre. Such a deviation will greatly distort the intensity readings obtained.

2.6.5.1.5 - Overall Comparison

Which type of device is more appropriate for in-situ monitoring will depend largely on the type of data that is being sought, and whether any financial or process-complexity restrictions apply. Should an operator require a spatially-resolved understanding of temperature profiles, melt pool dimensions, or spatter, a camera is necessary. If measuring rapid temperature fluctuations inside the melt pool is desired, a pyrometer may be the preferred option. In cases where equipment budgets are limited, but technicians can devote time to proper focussing, a pyrometer could suffice. If more money can be spent, and either a viewing port pre-exists in the build chamber, or resources can be devoted to a more complex installation, then a camera may yield better results.

2.6.5.2 - On-Axis vs Off-Axis

On-axis and off-axis monitoring setups each possess inherent strengths and limitations, with the preferred option depending on the types of signals desired for study. The following discussion deals only with cameras, since it almost goes without saying that a pyrometer is best mounted on-axis.

2.6.5.2.1 - On-Axis Strengths

The ability of a sensor to track the exact position of the laser incidence spot, allows for much greater spatial resolution. The Field of View (FoV) for an on-axis camera is typically less than 5mm × 5mm, centred around the melt pool, which can provide image clarity on the order of 1 µm/pixel. This clarity can result in highly accurate measurements in terms of melt pool dimensions, and temperature profiles. These are very useful in terms of monitoring the kinds of melt pool fluctuations that accompany keyhole formation, or lack of fusion porosity.

2.6.5.2.2 - On-Axis Limitations

The major limitation of on-axis sensors is that they can't provide data on anything outside of the melt pool's immediate vicinity. Recoater defects, swelling, track overlap, and track cooling profiles are all lost outside the narrow FoV. Also, due to the fact that the camera records from directly above the melt pool, the molten alloy can often be obscured by plume or spatter. Complex optics and precise calibration are required, in order to split the laser wavelength from the wavelengths of monitoring interest, and to focus the camera directly upon the powder bed.

2.6.5.2.3 - Off-Axis Strengths

The FoV of an off-axis camera setup is typically much larger than that of an on-axis, with frame sides usually on the order of 10 to 50mm. This allows for the context of entire tracks or parts to be recorded, including spatter trajectories, recoater defects, swelling, track overlap, and track cooling profiles. The fact that off-axis cameras can be simply placed into the build chamber, or even outside a viewing port, without much complex optics, results in far easier setup and installation.

2.6.5.2.4 - Off-Axis Limitations

The major limitation of off-axis camera setups is the lower spatial resolution. Since the FoV includes large sections of the powder bed, resolution is typically in the order of 10 to 100 $\mu\text{m}/\text{pixel}$. The melt pool may only be recorded in low detail, while the majority of the pixels in each saved image may contain background static, wasting a large amount of processing time and storage. Also, due to the non-vertical angle that the camera records with, there will inevitably be some perspective issues in the images, with closer melt pools appearing larger, and with only melt pools within certain regions if the image being perfectly in focus. Therefore, the restrictions of the camera must be considered from early on in a build's design process, if off-axis monitoring is to yield any valuable data.

2.6.5.2.5 - Overall Comparison

Whether on-axis or off-axis monitoring is preferable, depends on what the user would like to observe. If detailed melt pool dimensions or temperature profiles are required, on-axis is the better approach. If defects and details that lie outside the narrow window around the melt pool are needed, or if the complex optical engineering that is needed to mount an on-axis camera must be avoided, then off-axis might be the more appropriate choice.

2.6.5.3 - Single Spectrum vs Dual Spectra

In this context, single spectrum refers to light that has been filtered to a single band of the spectrum, whether narrow (in the case of cameras fitted with a bandpass filter), or broad (in the case of off-the-shelf cameras that record in the visible range). Dual spectra refers to light that has been split into two different bands, to be recorded on different sensors.

2.6.5.3.1 - Single Spectrum Strengths

The major advantage of a single spectrum sensor (whether a camera or pyrometer), is that the optical requirements of the system are significantly reduced, and only one sensor is needed. In contexts where money, time, or a technician's time is limited, a single spectrum setup may be the more economical and robust solution. Additionally, data processing and storage costs are significantly lower when only one sensor is recording.

2.6.5.3.2 - Single Spectrum Limitations

As discussed in detail in a previous subsection, a single spectrum cannot accurately calculate temperature, due to the circular dependency of temperature and emissivity. When every pixel in a camera assumes a constant emissivity, regardless of whether it is pointing at liquid or solid, temperature estimations be off by several hundred degrees. Also, should a sensor be tuned only to one single spectral band, it can be difficult to decide upon an appropriate band: pixels tuned to spectra that can detect the lower range of liquid temperatures, may become saturated near vaporisation temperature; and pixels tuned to spectra that can detect near-vaporisation temperatures without saturation, may not be able to detect lower liquid temperatures as well.

2.6.5.3.3 - Dual Spectra Strengths

The two-colour pyrometry method, in which a point's temperature can be calculated by comparing the intensity ratios between two pixels recording it at two different wavelengths, bypasses the need for emissivity values. This allows for accurate temperature measurements, that can be compared across experimental setups. The danger of saturation is also lower, as wavelength bands are often chosen to be removed from each other on the spectrum.

2.6.5.3.4 - Dual Spectra Limitations

The major drawback to dual spectra setups is the cost, complexity, and requirements for precise calibration. Beam splitters and multiple filters are needed, and, considering that each pixel on a camera sensor must be mapped to the same pixel on the corresponding sensor, positioning must be exact. The processing and data storage complexity also doubles, when builds are recorded with two sensors instead of one.

2.6.5.3.5 - Overall Comparison

Unlike the other categories being discussed, where different approaches are more suitable depending on circumstances and desired output, in the case of single versus dual spectra, there is a very clear superior. Being able to monitor the build (whether via camera or via pyrometer, whether on or off-axis) in dual spectra allows for more accurate temperature calculations, and reduces the risk of pixel saturation. It is only the increased cost and complexity associated with a dual spectra setup, that may make it the less appealing option.

2.6.5.4 - Infrared vs Visible

The types of phenomena that are detectable in the visible range of the spectrum (400 to 700 nm), versus those detectable in the infrared range (700 to 1400 nm), can be drastically different. Careful consideration must be made, when choosing a spectral band. This discussion is again tailored towards cameras, as pyrometers can be assumed to be filtered to infrared.

2.6.5.4.1 - Visible Strengths

A wide variety of processes inside the L-PBF chamber are visible to the human eye, and therefore visible to an off-the-shelf camera sensor, with no spectral filtering. Melt pool dimensions, hot spatter, plume intensity, and the movement of powder particles, are all detectable within the visible spectrum, whether via incandescence, reflection of light, or distortion of light due to refractive differences in gases. With side-illumination of the powder bed, visible light cameras can simultaneously detect melt pool dimensions, and recoater defects.

2.6.5.4.2 - Visible Limitations

The greatest limitation with visible range camera setups is the lack of temperature data. While melt pool dimensions can be implied from the sharp incandescence difference between solid and molten metal, the temperature profile of liquid within the melt pool is largely unknown. Pixels tuned to

visible light are also highly susceptible to saturation, with all those within the melt pool boundaries jumping to maximum intensity.

2.6.5.4.3 - Infrared Strengths

The major advantage of a narrowly-tuned infrared spectrum camera is that, even with a single spectrum setup, temperature values can be estimated from intensity quite straightforwardly (or exactly calculated, for a dual spectra setup). Thermal gradients and cooling profiles can then be extracted from the images.

2.6.5.4.4 - Infrared Limitations

Only objects that emit a significant quantity of thermal radiation, within the infrared spectrum, are observable. The majority of the powder bed appears to the sensor as a void, as does the plume, and spatter that has cooled. The optical requirements of an infrared setup are more complex, with specific bandpass filters being required. Selecting a spectral band that can detect near-liquidus temperatures, while also not becoming saturated at near-vaporisation temperatures, is not a trivial task.

2.6.5.4.5 - Overall Comparison

Which spectral band is more appropriate for in-situ monitoring depends largely on what features are required to be observed. Should a height map or defect distribution of the powder bed be desired then visible light is required. Should a measure of melt pool dimensions be needed, then visible light would suffice. If any kind of reliable temperature data is wanted, however, then infrared light, tuned to a carefully considered wavelength, is necessary.

2.6.6 - Sub-Layer Signals

In the following subsection, the few signals that can reveal the state of sub-surface processes will be discussed.

2.6.6.1 - X-Ray Computed Tomography (XCT)

Due to their ability to penetrate solid metal, X-rays can be used to probe in-situ, sub-surface dynamics in the L-PBF process. The mechanisms through which X-ray production, focussing, absorption, deflection, and detection can be used to monitor sub-surface features, will be discussed in more detail in the following chapter, where X-ray techniques used in this current work will be explored.

Leung et al. [300] transmitted a 55KeV X-ray beam through a chamber building Invar 36 powder, recording images at 5100 fps, to capture a 2D video of melt pool development, with a pixel size of $6.6\mu\text{m} \times 6.6\mu\text{m}$. The aim of the experiment was to determine the effects of powder oxidation on melt pool dynamics. It was found that iron and nickel oxides altered the Marangoni convection, reversing it from flowing outwards across the pool surface (in virgin powder), to inwards (in heavily oxidised

powder). This inwards flow results in deeper melt pools, and more entrained gas pores. The study highlights the need for careful manufacturing, handling, and storage of powder, to minimise oxygen exposure. Example images of melt pools, and the effects of oxides on Marangoni flow, are shown in Figure 2.109 and Figure 2.110, respectively.

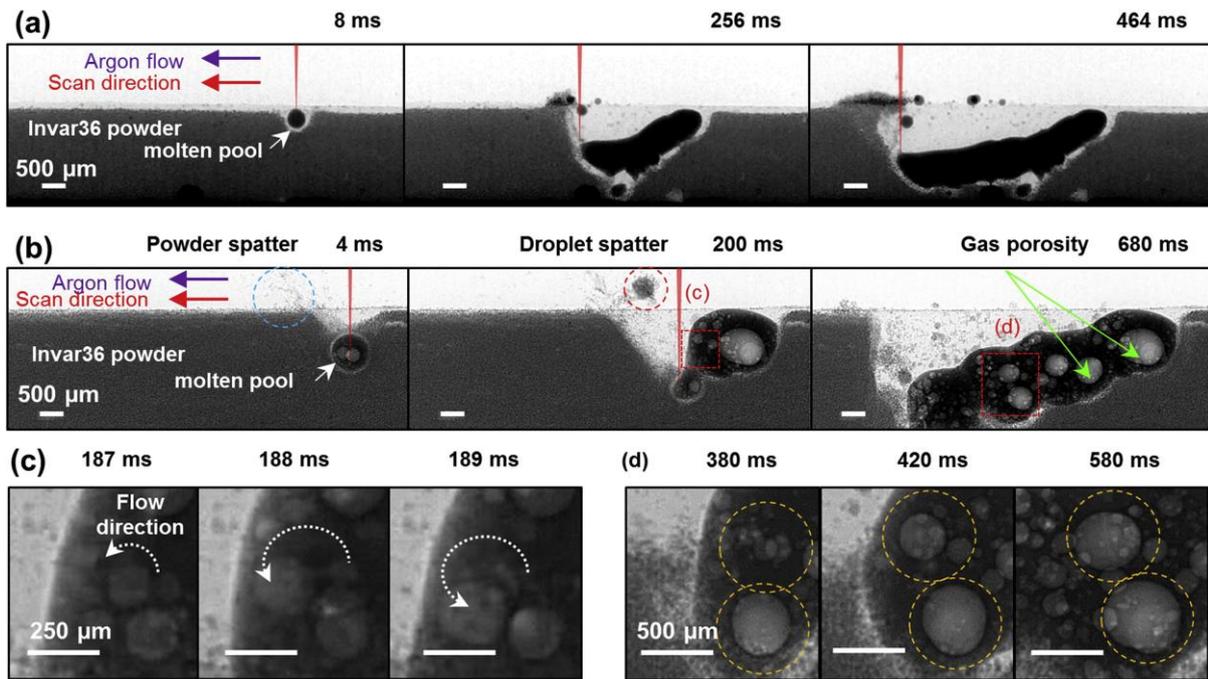


Figure 2.109

XCT results from Leung: (a) scan track with virgin powder; (b) scan track with oxidised powder; (c) zoomed region from (b), showing pore coalescence and migration by the centripetal Marangoni convection (white arrows); (d) another zoomed region, showing pore growth promoted by oxide films (yellow circles) [300]

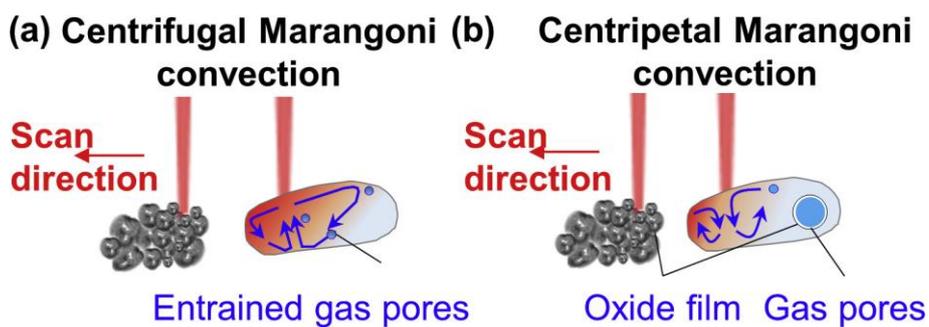


Figure 2.110

Leung's illustration of how an oxide film can change Marangoni convection from outwards (centrifugal) in virgin powder, to inwards (centripetal) in oxidised powder [300]

Calta et al. [301] used a 24keV pulsed X-ray source, transmitting bursts of 22ps, to monitor melt pool dynamics as lasers scanned solid plates of 316L steel, Nickle 400, 6061 aluminium, and Ti-6Al-4V. The authors tested the plates under various atmospheric conditions, including near-total vacuum, ambient-pressure argon with extremely low oxygen, and ambient-pressure argon with 13,000 ppm

oxygen. The relationship between laser power and vapour depression was found to be significantly altered for each alloy, depending on the processing atmosphere. It was reasoned that at reduced pressures, as boiling temperature and therefore melt pool surface temperature is reduced, surface tension increases, leading to lower melt pool aspect ratios. Also, the presence of oxygen in the atmosphere was found to significantly reduce surface tension, therefore leading to higher aspect ratios. These results are shown in Figure 2.111 and Figure 2.112.

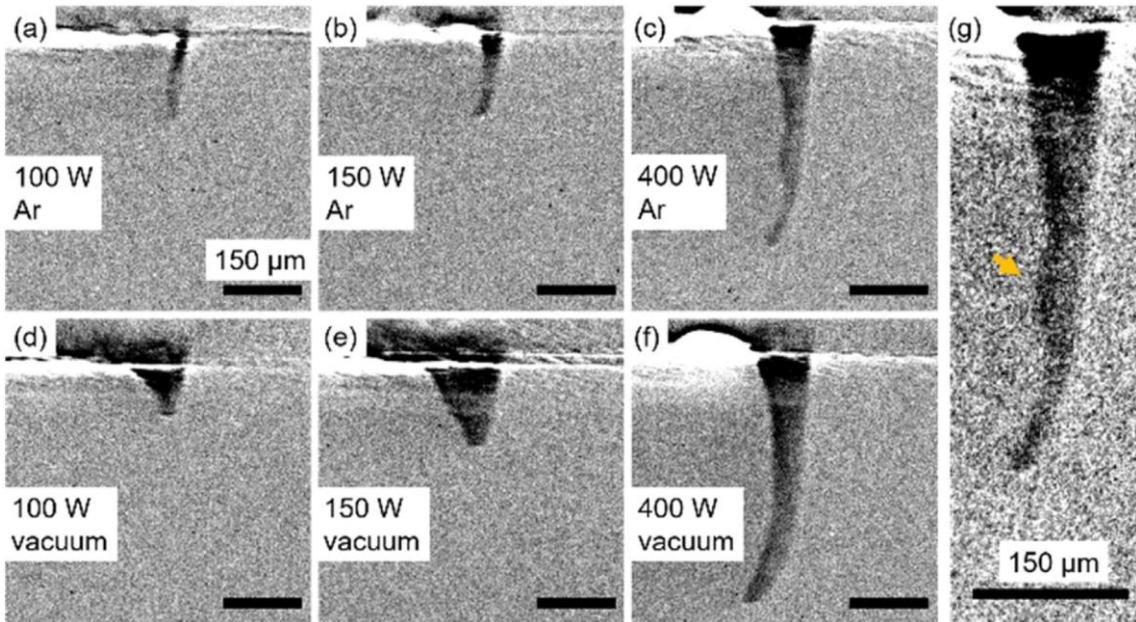


Figure 2.111

XCT results from Calta, showing the effects of both laser power and atmosphere on keyhole morphologies. (g) is a zoomed version of (c). [301]

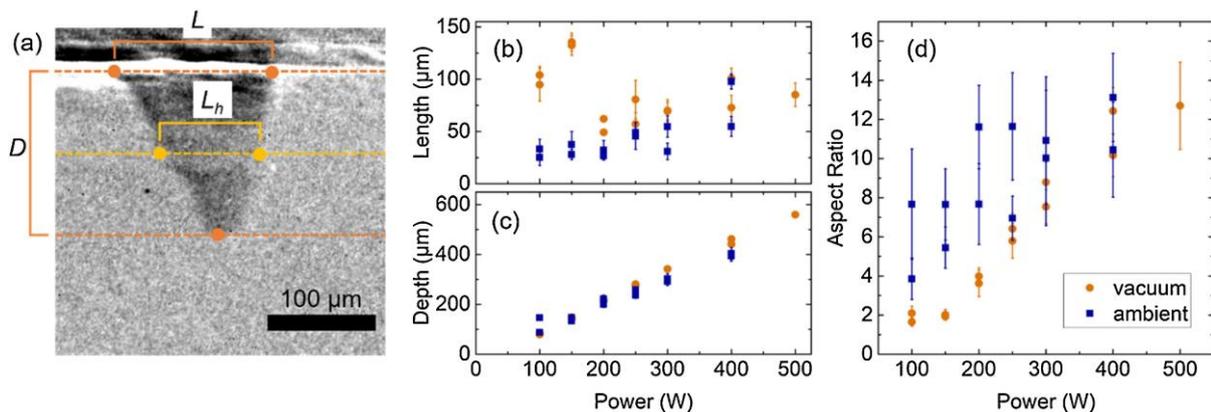


Figure 2.112

Results from Calta: (b) melt pool lengths across various laser powers at two different atmospheres; (c) the pools' depths; (d) the aspect ratios; (a) the method used to measure these dimensions. [301]

Paulson et al. [302] transmitted a 25keV beam through a chamber building in Ti-6Al-4V, to a detector recording at 30,000 fps, for a pixel size of $1.97\mu\text{m} \times 1.97\mu\text{m}$. Simultaneously, an IR camera was used to record average melt pool temperature. The researchers aimed to link the thermal history

of a scan vector, to its likelihood of forming keyhole porosity. It was found that temperature instability, especially when temperatures start high, dip, and then increase, is very likely to indicate keyhole porosity. The authors conclude that this kind of knowledge about the implications of IR thermal data, can be used to gather implications about sub-surface processes. The experimental setup and results are shown in Figure 2.113 and Figure 2.114.

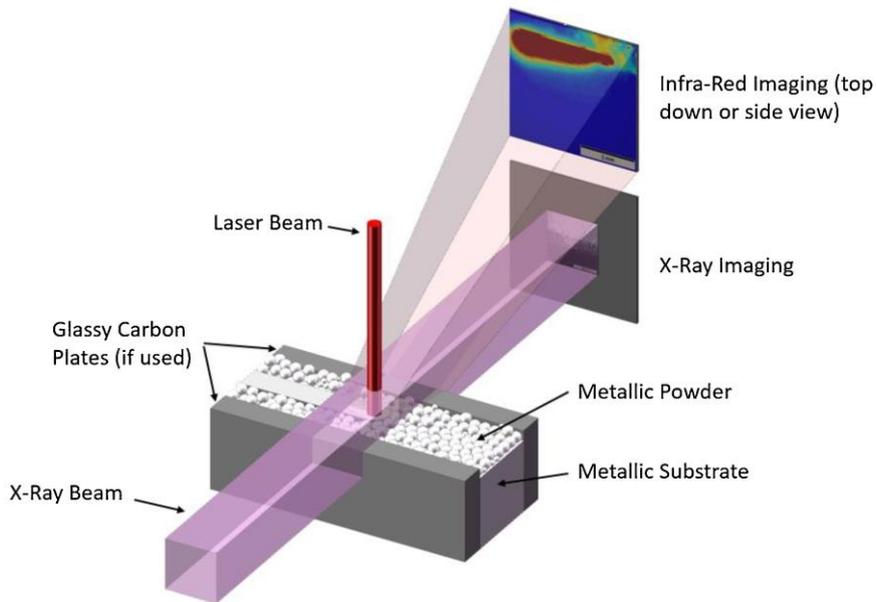


Figure 2.113

Paulson's experimental setup, with IR and XCT recording simultaneously [302]

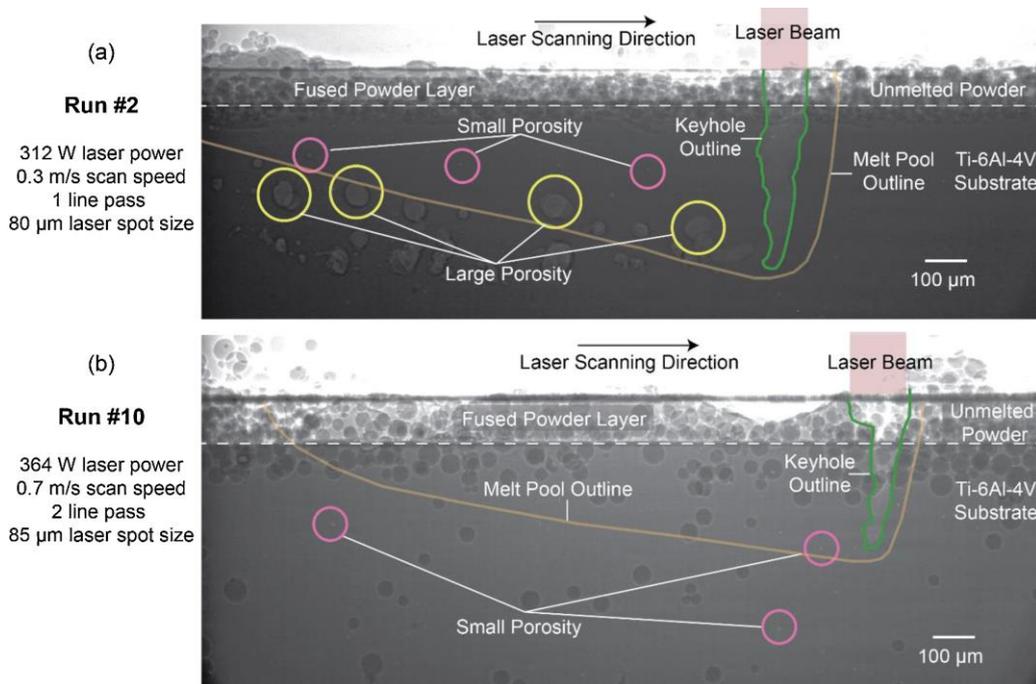


Figure 2.114

Example XCT images from Paulson, showing the melt pool (orange), the keyhole (green), small pores (pink), and large pores (yellow) [302]

2.6.6.2 - X-Ray Diffraction (XRD)

XRD refers to the process of determining crystal structure, phase, strain, and orientation of a crystal lattice, by measuring the manner in which x-ray beams passing through it is diffracted.

Schmeiser et al. [303] used XRD to monitor the lattice spacing of IN-625 during building. By comparing spacing in the X and Y directions to a stress-free control lattice, the authors managed to track both transverse strain and building direction strain, over the course of a build. It was learned that strain states change continuously during manufacture, with complex stresses, alternating between compressive and tensile, forming around the melt pool. It was found that the interplay of thermal contraction, and the fact that yield strength increased as material cooled, caused the greatest strain to exist around 300µm below the build surface. Stress and strain measurements from this experiment are shown in Figure 2.115 and Figure 2.116, respectively.

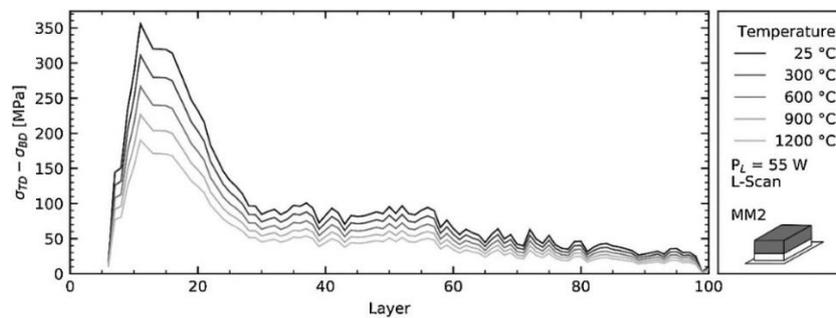


Figure 2.115

Stress difference, calculated as transverse direction (TD) minus build direction (BD), plotted as a function of layers below the top, for various temperatures, from Schmeiser 's results [303]

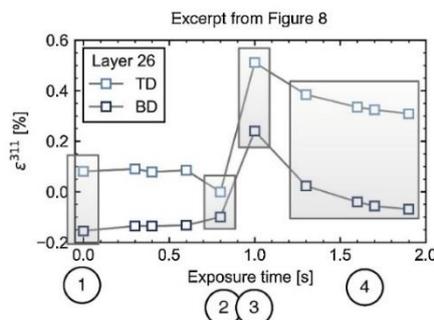
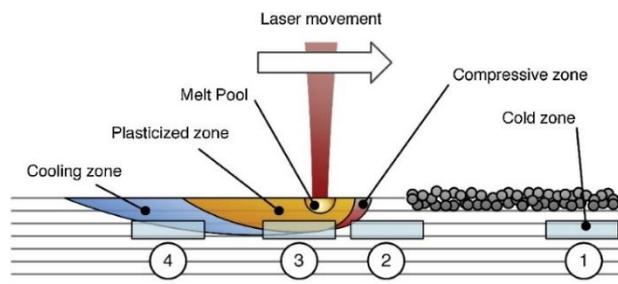


Figure 2.116

Strain, 26 layers (approximately 300µm) below the top surface, in the transverse and build directions, in four different locations, as the melt pool progresses across them. From Schmeiser. [303]

2.6.6.3 - Solid-Borne Acoustics

In Section 2.6.3, the creation of mechanical vibrations was discussed, in terms of these vibrations being transmitted through the inert atmosphere and detected by sensors within the build chamber. These vibrations are also transmitted through the underlying solid material, and can be detected by sensors embedded underneath the baseplate.

Eschner et al. [304] embedded a piezoceramic sensor (a highly sensitive strain gauge that converts pressure waves into electrical current) underneath the build plate during manufacture of 316L steel. Data was sampled at 4 MHz, and fed into a machine learning algorithm in an attempt to interpret the acoustic signatures of melting, vaporisation, solidification, and keyhole pore formation. The model was shown to be 85 to 88% effective at determining part density when tested on the same part geometry that it had been trained on, although this accuracy fell to 56 to 61% when a novel geometry was presented. The authors note the risk that, rather than learning the acoustic signatures associated with thermal defects, the machine may have actually learned the signatures associated with certain process parameters. They caution that this is a common risk in using machine learning with L-PBF data. The experimental setup and results are shown in Figure 2.117 and Figure 2.118, respectively.

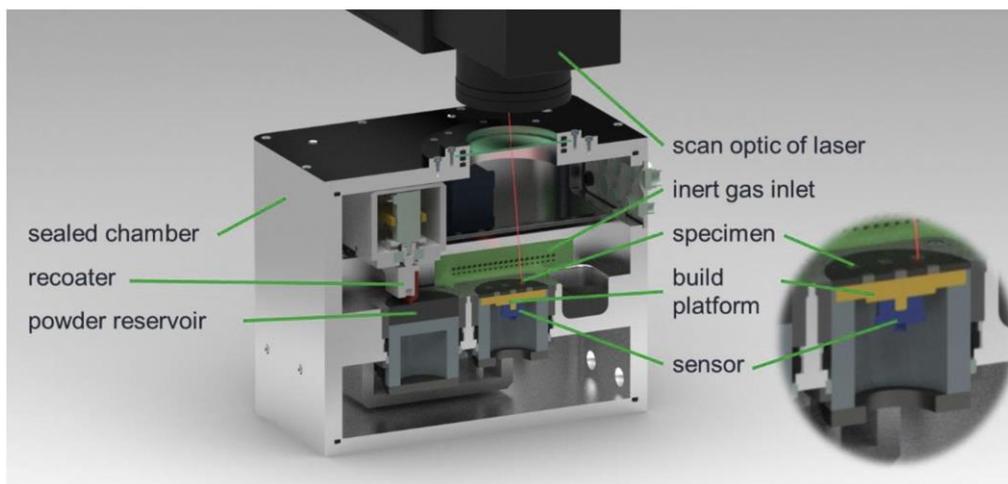


Figure 2.117

Eschner's experimental setup, with the piezoceramic sensor embedded below the build chamber [304]

2.6.6.4 - Baseplate Deformations

As well as monitoring mechanical vibrations permeating through the build plate, it is also possible to monitor strain and deformation in the plate, which can be related back to thermal processes in manufactured component. Hehr et al. [305] developed a 'Smart Build-Plate', with embedded fibre optic sensors, which can determine the extent of strain exerted upon it, by monitoring light transmitted through it. This device was tested in the building of 6061 aluminium components, which were deliberately designed to undergo delamination cracking during the build. By monitoring the strain exerted on the base plate, the authors managed to identify delamination cracking as it occurred, and even successfully specified which regions of the base plate were hosting the parts

undergoing cracking. The ‘Smart Build-Plate’ and strain results are shown in Figure 2.119 and Figure 2.120, respectively.

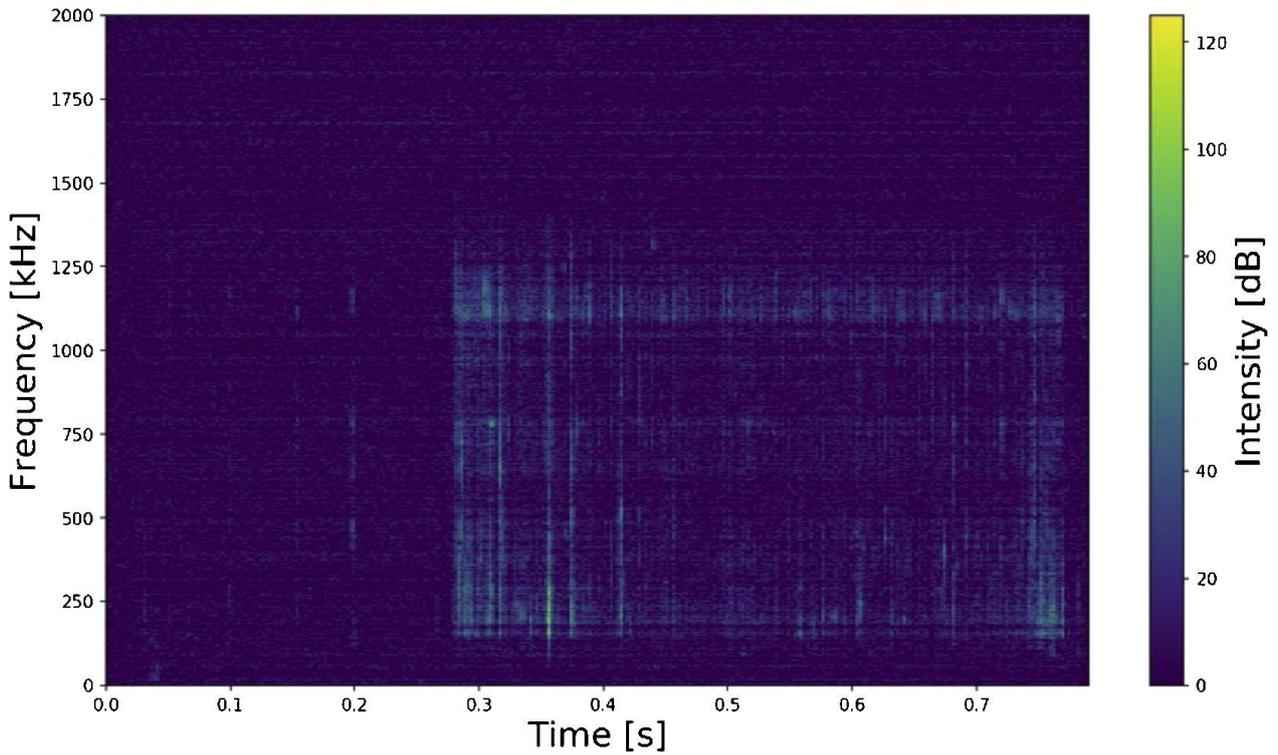


Figure 2.118

The piezoceramic output current of one layer from Eschner, after a Fourier transform has been applied, to identify the intensity of each frequency component. Nine components were built in the layer, with acoustic differences between them clearly visible in the diagram. [304]

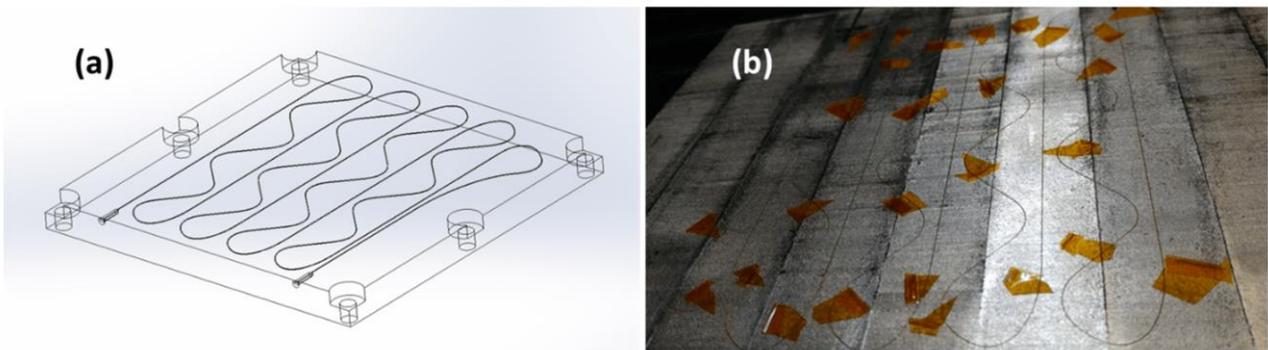


Figure 2.119

A schematic of Hehr's ‘Smart Build-Plate’ (a), and the fibre optic sensors that are embedded into it (b) [305]

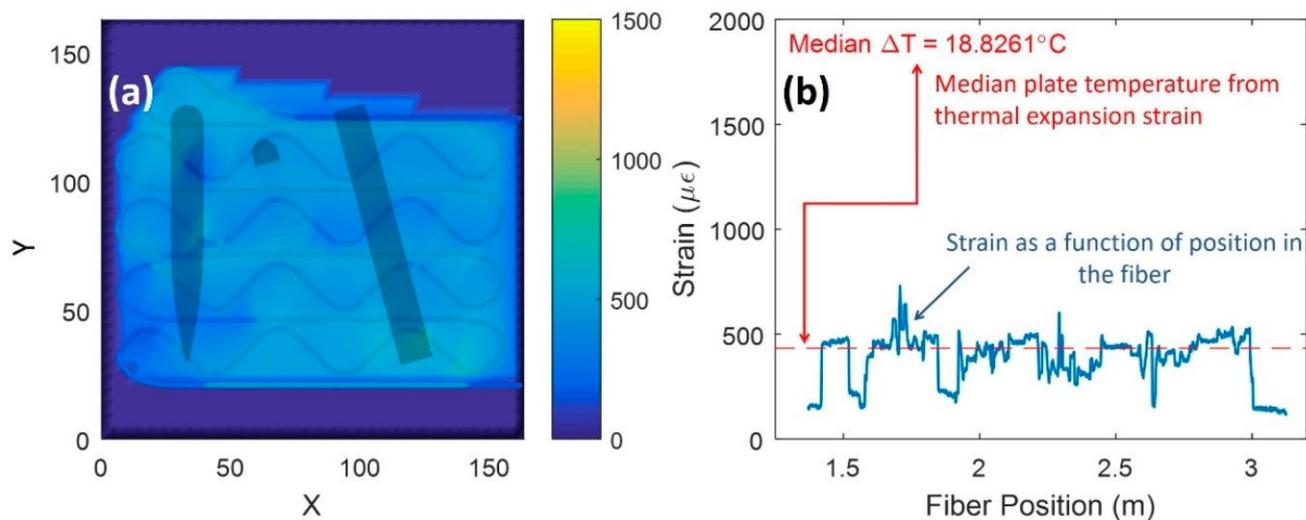


Figure 2.120

Hehr's results: (a) a heatmap showing strain across the build plate; (b) the strain vs fibre position graph that the heatmap is derived from, with a line to indicate background strain as a result of thermal expansion [305]

2.6.7 - Closed Loop Control

Closed loop control, in which defects are detected in-situ and process parameters are automatically adjusted to either correct, mitigate, or prevent propagation of these defects, is the current end-goal of a lot of L-PBF research. The topic is extremely broad and intricate, and only a brief overview will be presented here. The two most significant obstacles to the development and implementation of closed loop control are the processing time required for defects to be interpreted from in-situ data, and the high complexity and non-linearity of the parameter-result relationship inherent to L-PBF. As a result, many researchers attempting close loop control currently only attempt for automatic adjustment of one or two process parameters [306].

Wang et al. [307] used a high-speed, coaxial pyrometer to detect melt pool thermal emissions, while building with Ti-6Al-4V. The data was fed into a machine learning system, designed to interpret signals from over-heating, LoF porosity, swelling, and balling. The system responded to these signals by scaling the laser power, between 100 to 500W, updating the power at a rate of 1 kHz. The authors managed to eliminate each of the aforementioned defects in the controlled parts, and observed significantly less geometrical deviation in these than in the uncontrolled parts. It was noted that the PID controller and pyrometer used in this study cost less than \$50 to purchase, and that the system could be installed as a self-contained unit, without the need for a computer. Images of the uncontrolled and controlled parts, and a height map comparison of these, are shown in Figure 2.121, Figure 2.122, and Figure 2.123, respectively.

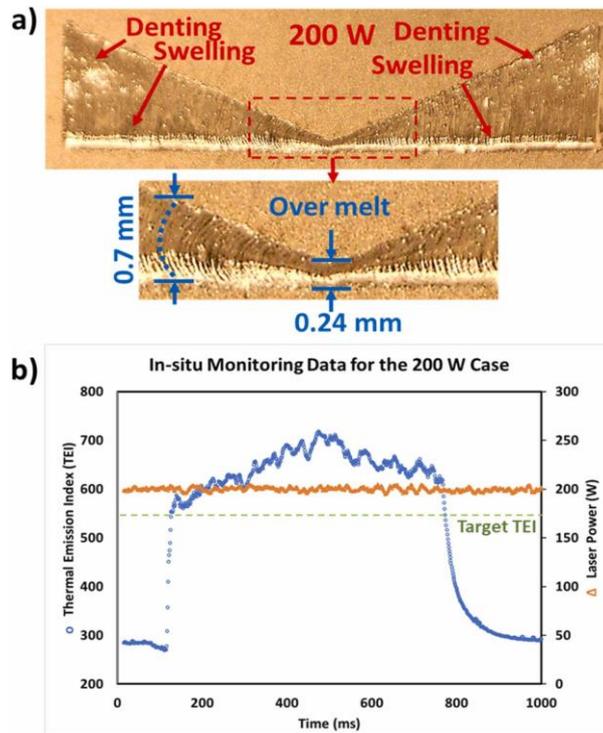


Figure 2.121

Weng's uncontrolled results: (a) swelling and denting are evident across the entire part, and as the vertical scan vectors become shorter, significant over-melt occurs; (b) a constant laser powder (orange line) results in large deviations of Thermal Emission Index (TEI) [307]

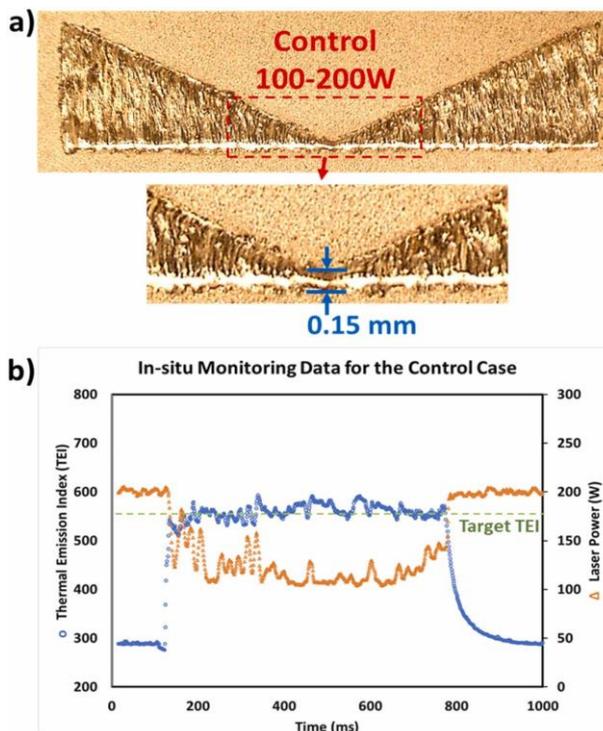


Figure 2.122

Weng's controlled results: (a) the part exhibits far fewer defects; (b) a variable laser power results in a steady TEI [307]

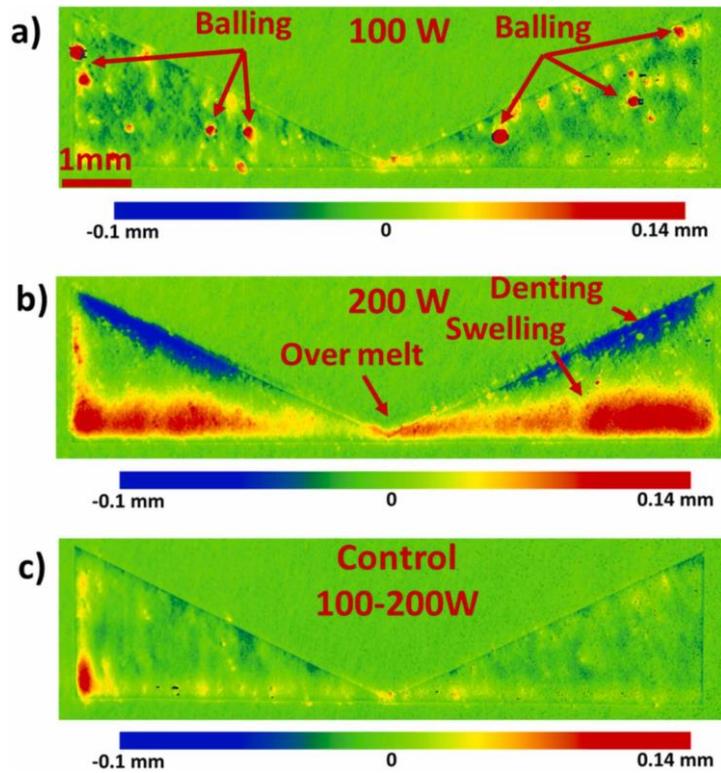


Figure 2.123

Comparing height maps for Weng's results: (a) uncontrolled at 100W, showing significant balling; (b) uncontrolled at 200W, showing significant swelling and over-melting; (c) controlled, with power varying between 100 and 200W, showing very few defects [307]

Nahr et al. [308] monitored the building of Ti-6Al-4V with an off-axis IR camera, recording at 10 fps at a wavelength of 900nm. The system aimed to identify heat accumulation in the parts, and adjust laser power on the subsequent layer to compensate. When control was implemented, a more constant temperature profile was realised across all parts. This was confirmed by the transition of tempered martensite in the uncontrolled parts, to acicular martensite in the controlled parts, which indicated consistently high cooling rates and uniform temperatures. The experimental setup and results are shown in Figure 2.124 and Figure 2.125, respectively.

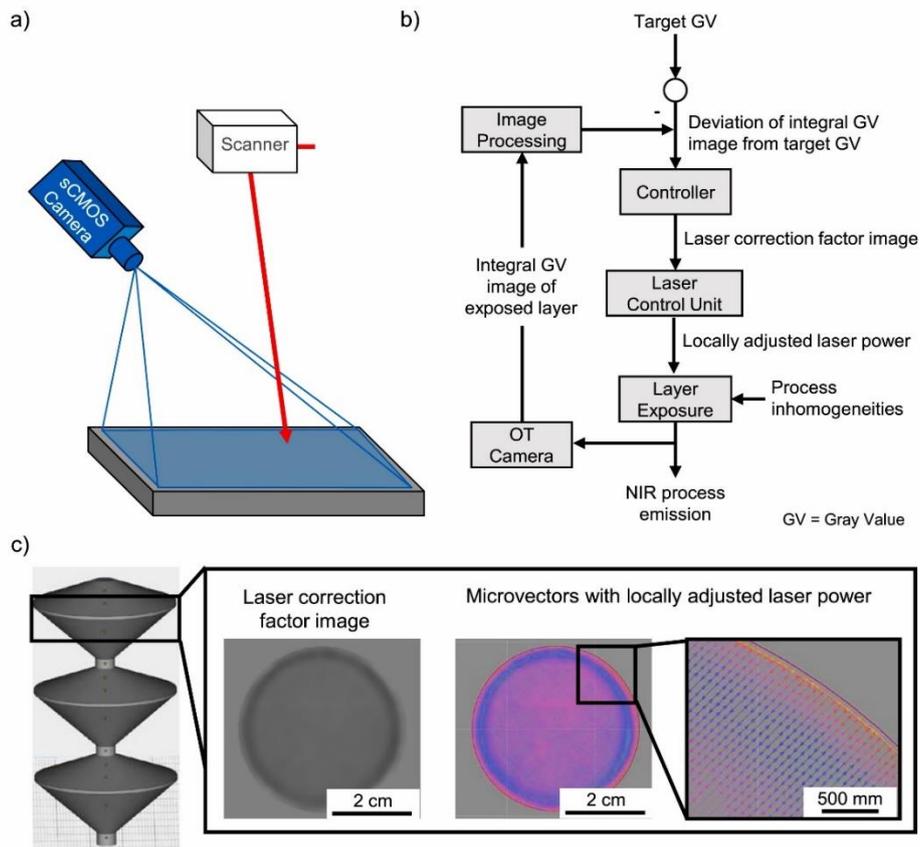


Figure 2.124

Nahr's experimental setup (a); the flow chart outlining the closed-loop that corrects the laser power, in order to achieve a target radiation emission (b); an example layer, showing how the recorded image (greyscale) is transformed into vectors with adjusted power to match the target (c) [308]

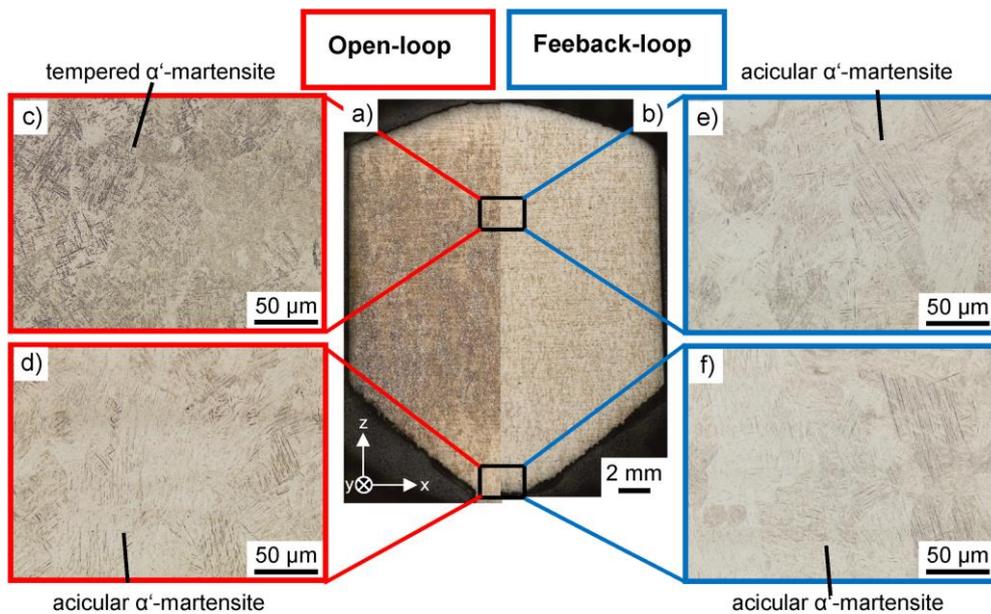


Figure 2.125

A comparison of Nahr's uncontrolled part (red) and controlled part (blue) microstructure [308]

Wenger et al. [309] monitored the building of 316L steel parts, with a coaxial, dual-wavelength pyrometer setup, that tracked mean melt pool temperature. Uncontrolled parts were found to exhibit a wide range of defects, such as cracking, delamination, balling, and porosity. A real-time control system was installed, with the aim of adjusting the laser power in order to keep the melt pool temperature within a target range. The authors found that this approach successfully reduced defects, and allowed for more uniform heat accumulation, when the controlled and uncontrolled parts were compared via off-axis IR video recording. A schematic of the control system's delay is shown in Figure 2.126.

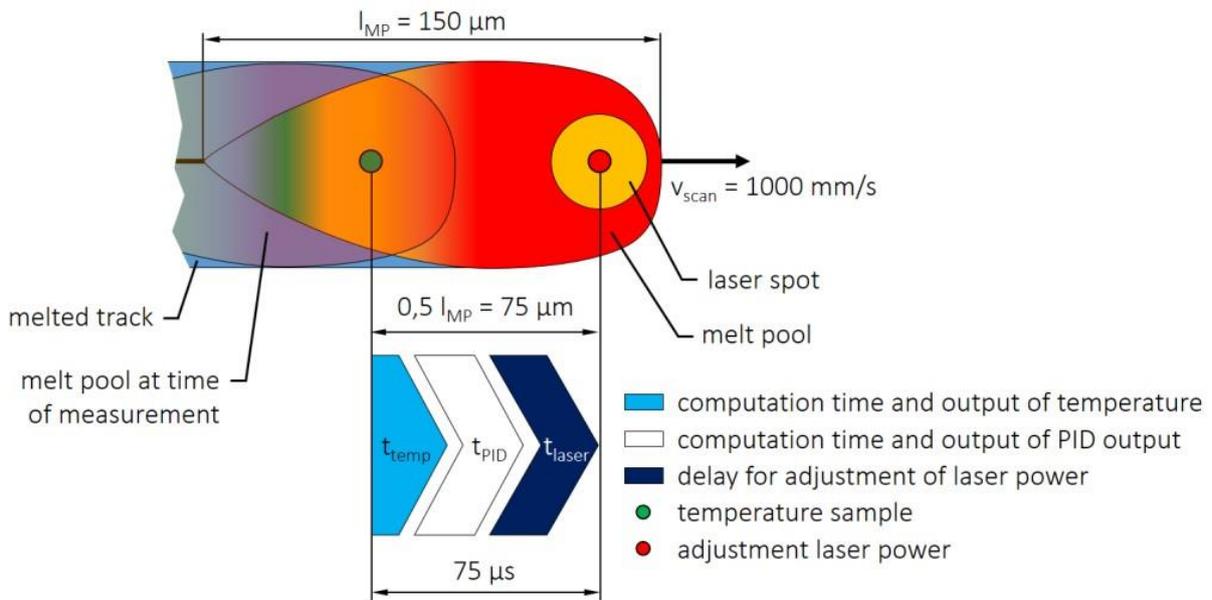


Figure 2.126

A schematic by Wenger, showing the delay between the location at which the temperature is sampled (green dot), and the location at which control is implemented (red dot) [309]

As can be seen from the previous few studies, all of which are very recently published at the time of writing, closed loop control is still in its infancy. Laser power is adjusted in each of the previous, in order to meet the fairly simple benchmark of mean temperature.

It is theorised that someday, much more adaptive and multivariable control will be possible. Lhuissier et al. [310], through the kind of XCT imaging discussed previously, demonstrated that keyhole pores left behind in one track, can be filled in by liquid feeding from adjacent tracks or subsequent layers. The authors found that L-PBF can be an “inherently self-healing process”, should the correct machine parameters be selected. Someday, advanced control systems may be able to detect most or all defects, and adjust not only laser power, but also speed, hatch spacing, scan strategy, and layer height, in order to fix these defects, resulting in near-perfect components, first time, every time.

2.7 - Literature Review Conclusions

Having reviewed literature on the wide range of topics presented previously, some clear questions arise about the predictive value of in-situ monitoring, the challenges of building defect-free, and the limitations of L-PBF.

Despite numerous dimensionless numbers being derived to predict the likelihood that any given set of machine parameters will result in defects, experimental data routinely shows unpredictability and wide margins of error. Much further work is required to better understand the how the fundamental mechanisms inside the melt pool can be controlled by the settings available to technicians and engineers.

While numerous papers describe significant progress in translating in-situ signals into valuable information on build quality, defects, and microstructure, a common problem remains amongst the majority of studies: they require equipment too expensive and specialised, or rely on geometry too simplified, to be relevant to most industrial processes. While researchers may be able to detect pores and phases in cubes, using multiple, state-of-the-art IR cameras, and data fed into advanced machine learning algorithms, these insights would be of little practical value to smaller workshops being commissioned by clients to produce unique and complex components.

Further work is therefore required to bridge this gap. Research is required to better understand how a single wavelength, in-situ monitoring device, that is within the budget of smaller production facilities, can reveal useful information about a complex-geometry build. Such work will involve several subsets of projects:

- The extent to which monitoring device setup affects the results it outputs needs to be evaluated. A camera is of limited use to a workshop if it requires expert-level calibration before each build, in order to make its images accurate.
- The possibility of bypassing emissivity, and being able to distinguish liquid from solid based on pixel-intensity values alone, needs to be investigated.
- Linkages need to be drawn between the data that in-situ monitoring devices output, and a wide range of microstructural and mechanical properties, by cross-correlating camera images with data from modelling, optical microscopy, SEM, EBSD, and hardness testing.

Providing this missing linkage, between that which is possible in high-end research laboratories using cutting-edge monitoring devices, and that which is possible to smaller manufacturing workshops with limited budgets, is going to be the focus of the experimental work in further chapters.

3 - Experimental Methods

This chapter details all of the materials and equipment used in the experiments in this work, and discusses some basic characterisation performed on the powders.

3.1 - Powder Alloys

Two nickel-based superalloys were experimented upon in this work: a generic version of Inconel-718, and ABD-900AM. These materials were chosen for their respective statuses within industry as a proven standard, and as a novel and promising new alloy. The particle size distribution, particle morphology, and flowability of both powders were measured before testing began. The results of these measurements, and all other relevant and available data, will be discussed here.

3.1.1 - Powder Sources

As explained in the Literature Review, *Inconel* is a trademark of the *Special Metals Corporation*, the successor of the company that originally created the 718 alloy. The powder used in this experiment was purchased from *Carpenter Additive*, the AM-focussed business unit of *Carpenter Technology Corporation*. *Carpenter Additive* manufactures and sells their 718 alloy under the trademark *PowderRange*. For the sake of simplicity, the alloy will continue to be referred to as IN-718 throughout this work. All experimental work performed with IN-718 in this project was done with Batch No. 63134, purchased in April 2021.

ABD-900AM is a chemical composition developed by *Alloyed*, a UK company that uses a computational framework (*Alloys By Design*) to formulate new alloys for specific purposes. The powder was produced by French metal manufacturer *Aubert & Duval*. All experimental work performed with ABD-900AM in this project was done with Batch No. MC19155_OSS11001, purchased in 2020.

It is understood that this duration of storage time may have resulted in powder oxidation.

3.1.1.1 - Chemical Composition

Table 3.1
The composition (weight %) of the IN-718 batch

Nickel	52.3 %	Iron	Balance
Chromium	18.9 %	Niobium + Tantalum	4.92 %
Molybdenum	3.10 %	Titanium	1.01 %
Cobalt	0.1 %	Aluminium	0.53 %
Manganese	0.02 %	Silicon	0.04 %
Copper	<0.1 %	Carbon	0.03 %
Nitrogen	0.01 %	Oxygen	0.02 %
Phosphorus	0.015 %	Sulphur	0.001 %
Boron	<0.001 %		

Table 3.1 shows the composition (in weight %) of the IN-718 batch, as per the datasheet that *Carpenter Additive* supplied with the powder.

Niobium and tantalum are listed together, as per industry standards, since they are atomically near-identical, and thus difficult to clearly distinguish via characterisation methods.

Table 3.2 lists the ABD-900AM powder's composition (weight %), as per the datasheets supplied by *Aubert & Duval* alongside the batch.

*Table 3.2
The composition (weight %) of the ABD-900AM batch*

Nickel	49.70 %	Cobalt	20.17 %
Chromium	16.98 %	Tungsten	3.15 %
Titanium	2.31 %	Aluminium	2.09 %
Molybdenum	2.08 %	Tantalum	1.52 %
Niobium	1.82 %	Oxygen	0.0124 %
Nitrogen	0.00934 %	Carbon	0.0482 %
Boron	0.0044 %	Zirconium	<0.005 %

3.1.1.2 - Physical Composition

According to the supplied IN-718 datasheet, the D_{10} , D_{50} , and D_{90} diameters of the batch were 21.7 μm , 35.9 μm , and 57.8 μm , respectively.

According to the supplied ABD-900AM datasheet, 1.88% of the powder (by volume) was measured by sieving to have a diameter of >53 μm , and 2.52% of the powder was measured by laser diffraction to have a diameter of <15 μm .

All of these values were verified in-house via laser diffractometry, using a Mastersizer 3000.

The workflow of the Mastersizer 3000 is as follows:

1. A sample (approximately 5g) of powder is suspended in water inside the machine.
2. A laser beam is passed through the particle cloud
3. The extent of laser scattering depends on the particle sizes (smaller particles scatter at larger angles)
4. Detectors analyse the scattering pattern, and infer the PSD from it

The machine returns results in the form of diameters, ranging from 0.01 μm to 3500 μm , divided into 100 bins, spaced on a logarithmic scale, and the percentage of total powder volume that each bin contributes. The machine assumes spherical powders, with a uniform refractive index.

The test was repeated 5 times with a sample of IN-718 powder, and 5 times with a sample of ABD-900AM powder. No IN-718 particles were found with diameters smaller than 11 μm or larger than 153 μm , and no ABD-900AM particles were found with diameters smaller than 9 μm or larger than 135 μm . The PSD histograms of both powders are plotted in log space (as mean values, with error bars showing standard deviations) in Figure 3.1.

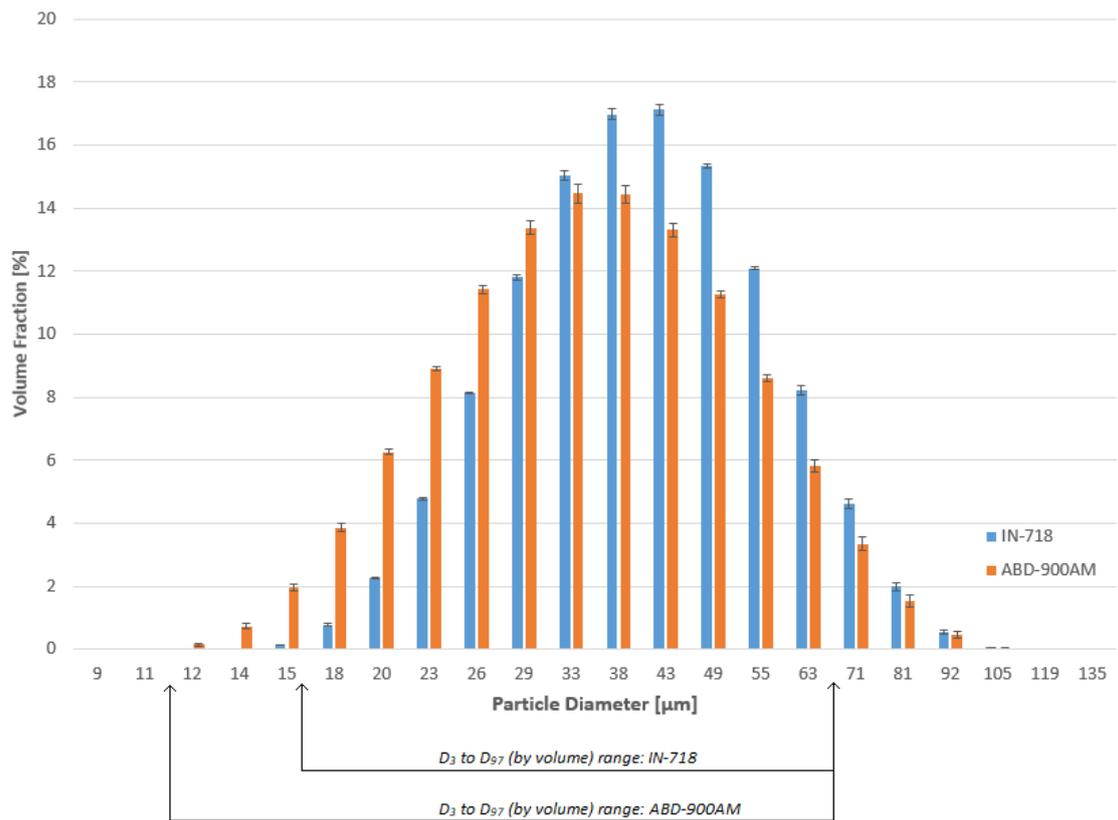


Figure 3.1

The volume fraction distribution of the powders, as tested in-house. The D_3 and D_{97} values lie between the bins that each arrow points between.

The mean particle volume fractions of each powder were plotted as cumulative volumes (in linear space), as shown in Figure 3.2.

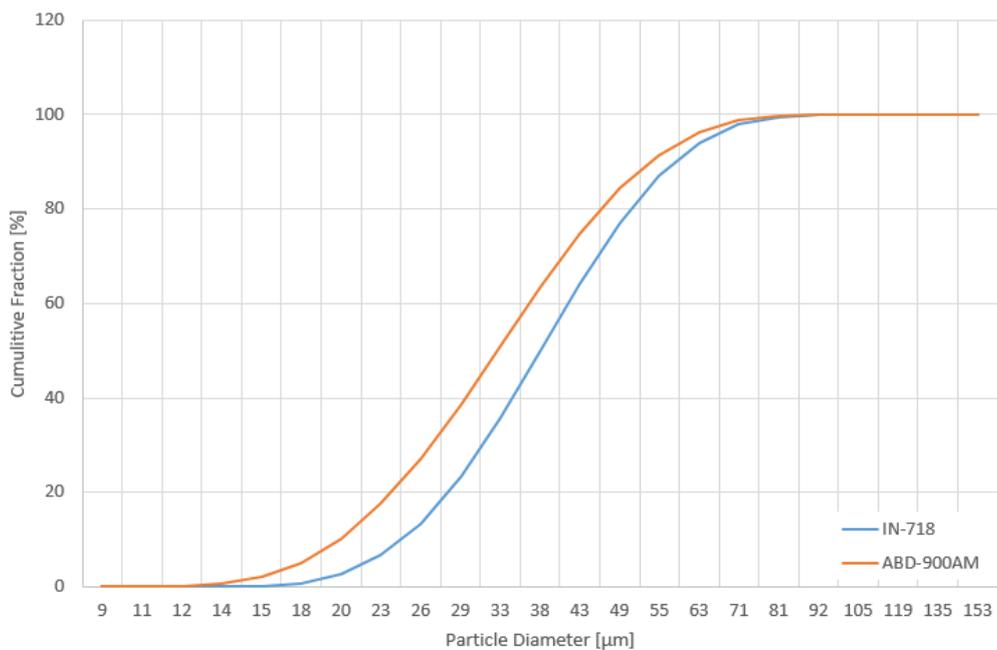


Figure 3.2

The cumulative volume distribution of the powders, as tested in-house

Based on the previous curve, the D_{10} , D_{50} , and D_{90} diameters of the IN-718 were determined to be 23.8 μm , 36.8 μm , and 57.2 μm , respectively. These values vary from those reported on the datasheet by 9.7%, 9.0%, and 1.0%, respectively. This was considered to be an acceptable error, that likely arose from the tendency of particles to segregate by size during a powder's handling. From the curve, the D_{10} , D_{50} , and D_{90} diameters of the ABD-900AM was found to be 20.4 μm , 36.2 μm , and 54.6 μm , respectively.

Neither the IN-718 nor the ABD-900AM datasheets contained information on average particles sphericity. In-house tests were carried out to produce a circularity distribution for both powder batches. Powder was cold-mounted in an epoxy resin puck, inside the Struers CitoVac vacuum curing device. After curing the puck was ground and polished to an OPS finish (explained in Section 3.4.3), and viewed with non-polarised optical microscopy. Three images, each 874 μm \times 874 μm (7950 pixels \times 7950 pixels), were taken of the powder particles.

A MATLAB code was written to analyse the images. The greyscale images were normalized (to set all pixel values between 0 and 1), and a threshold value of 0.5 was selected to isolate particles from background resin. For the small number of images, this visual-inspection method was considered more appropriate than the Otsu method, or other such more-repeatable methods. Using the in-built *regionprops* function, the areas and circularities of all particles were identified. Those with areas below 1000 pixels were ignored, to remove non-particle speckles and anomalies, and histograms were constructed from the circularities of the particles that remained. This workflow is demonstrated in Figure 3.3. It should be noted that in cases where multiple particles were either bonded together, or otherwise appeared to be touching in the optical image, they would have been counted as a single particle, with a larger area and lower circularity than is correct. It was decided that, since the circularity values were being calculated as a sense-check, and would not be used elsewhere in this project, investing the time to write a MATLAB script to rectify this would have been unnecessary. 2,146 IN-718 powder particles, and 1,421 ABD-900AM powder particles, were identified for analysis. Histograms of each powder's circularity distribution are shown in Figure 3.4. The mean circularity was of the IN-718 was 0.73, with a standard deviation of 0.21. The mean circularity and standard deviation of the ABD-900AM was 0.68 and 0.18, respectively. These results suggest that both powder batches were rendered slightly non-spherical by the gas atomization processes. The circularity equation used was:

$$C = \frac{4\pi A}{P} \left(1 - \frac{0.5}{\frac{P}{2\pi} + 0.5} \right)^2$$

where C is circularity, P is perimeter, and A is area.

Although it would have been possible to also measure particle areas from the microscope images, and therefore construct a histogram of sizes, it was decided that an estimate of PSD obtained this way would be of little value compared to those obtained from the laser diffraction method previously,

considering the sectioning bias inherent to imaging powder particles embedded in resin. It was understood that methods exist to compensate for this sectioning bias (e.g. Higginson et al. [311]), however were considering unnecessary to apply given the accuracy of the laser diffraction method.

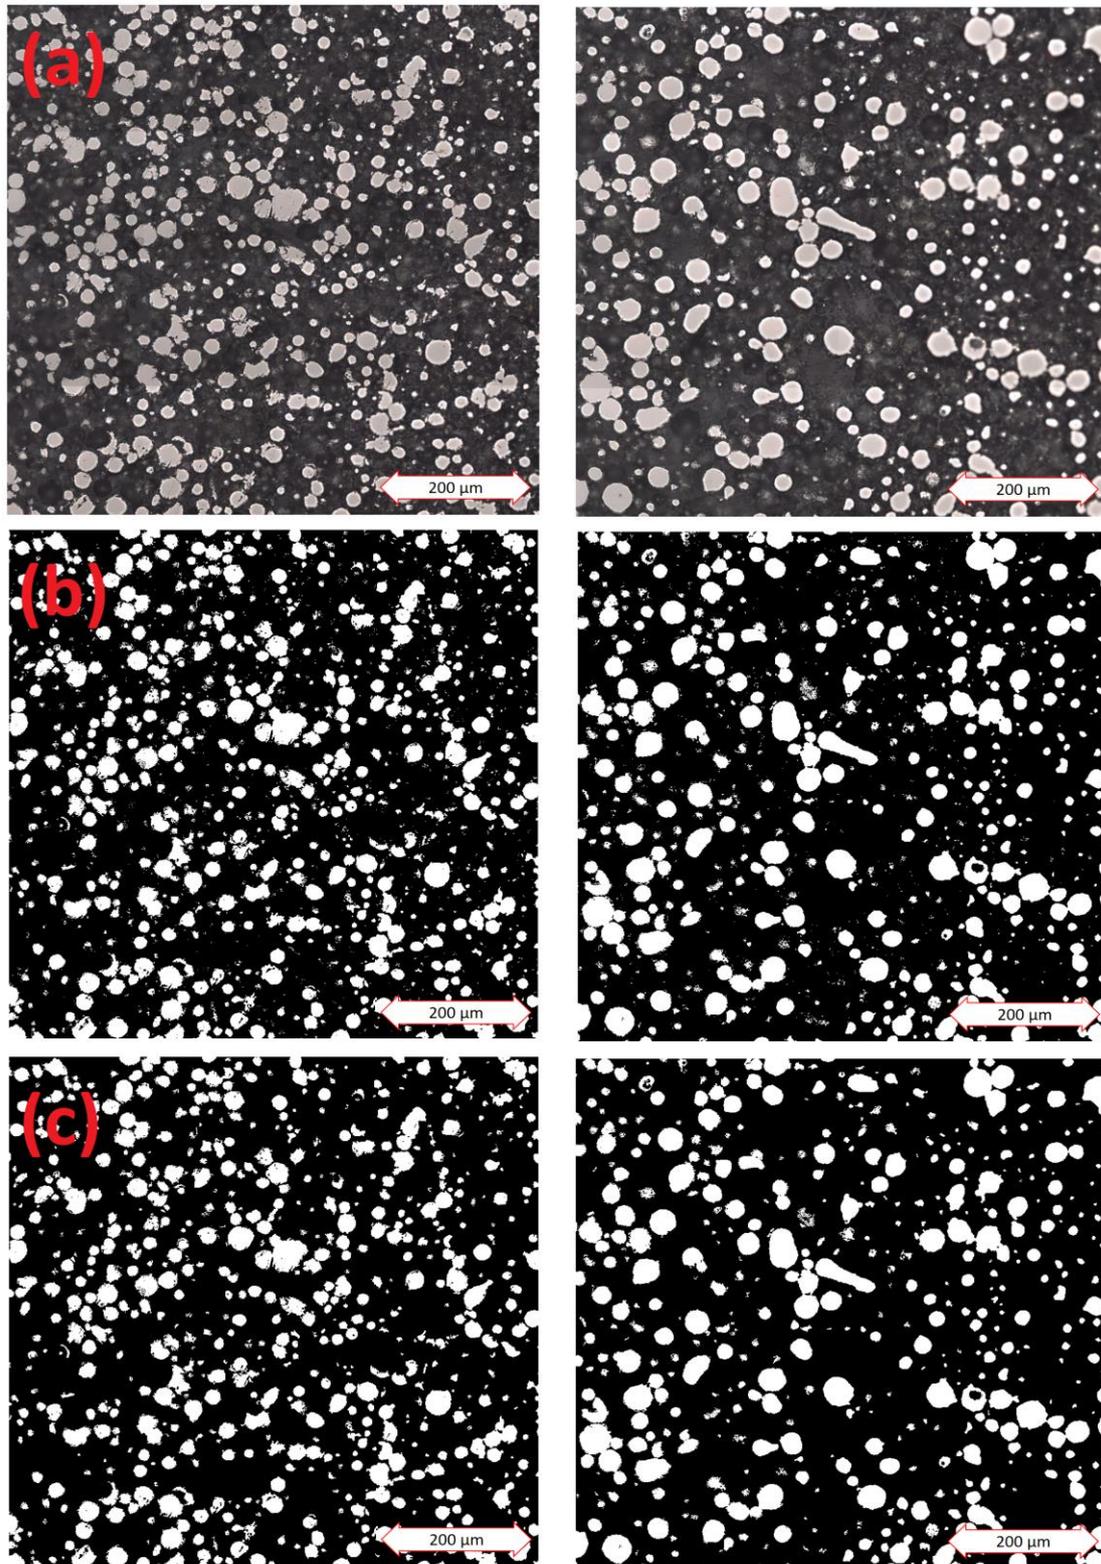


Figure 3.3

Example images from IN-178 (left column) and ABD-900AM (right), showing: raw images from the microscope (row a); the images binarised (row b); and regions of less than 1000px removed (row c)

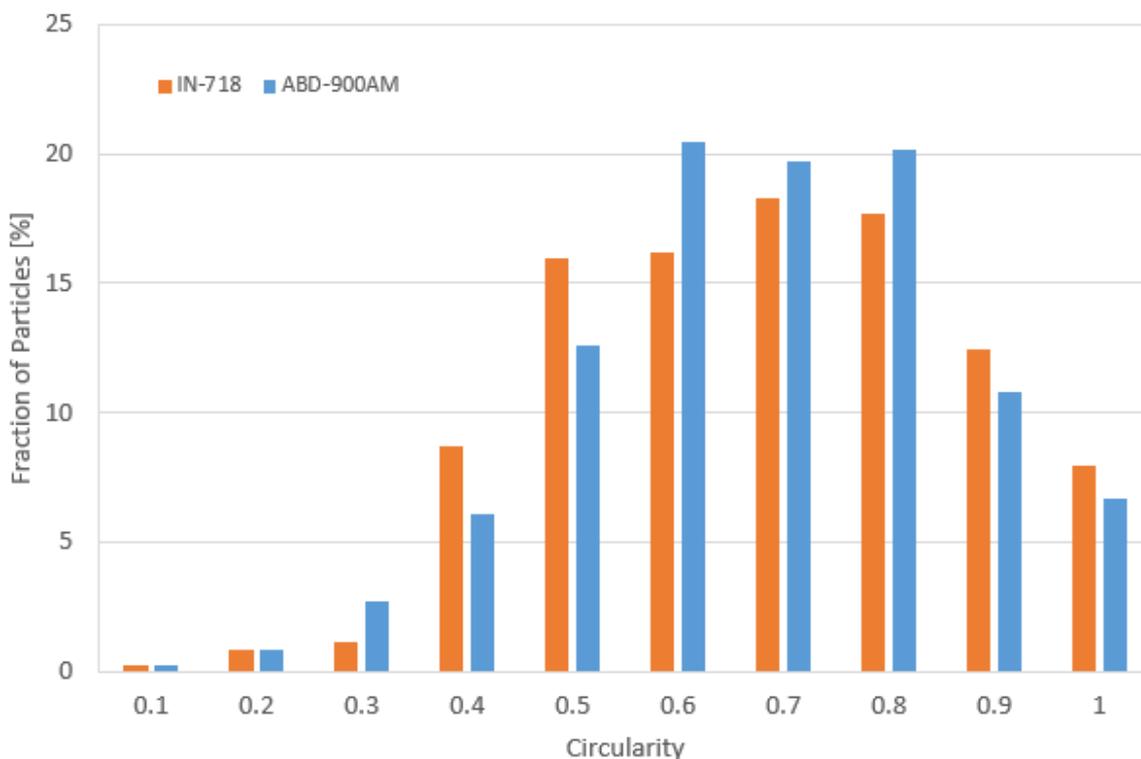


Figure 3.4

A circularity histogram of the powders, showing the percentage of overall particles within each bin

A Hall flow meter is a standardized device, consisting of a funnel with sides 30° off the vertical, and a 2.5mm orifice at the bottom. A powder's flowability is often reported in terms of the time it takes for 50g of it to flow through the funnel. The datasheets supplied with the powders reported the Hall flow of the IN-718 to be 14 s/50g, and of the ABD-900AM to be 12.79 s/50g.

These value was experimentally verified, using an in-house Hall flow meter. 50g of the powders were poured into the funnel, with a latex-gloved finger blocking the orifice. In the case of the IN-718, it was found that after the finger was removed, no powder would flow until the funnel was given a firm tap with a spanner. After this, all powder flowed through, as shown in Figure 3.5.

In the case of the ABD-900AM, 2 spanner taps were required for every test, possibly due to the powder's slightly lower circularity and smaller particle size. The first caused a single column of powder, with roughly the diameter of the orifice and extending up to the powder's surface, to flow through. The second was required to get the rest of the powder to flow. With practice, it became possible to tap the funnel as soon as the initial column stopped flowing, making the disruption to the flow minimal.

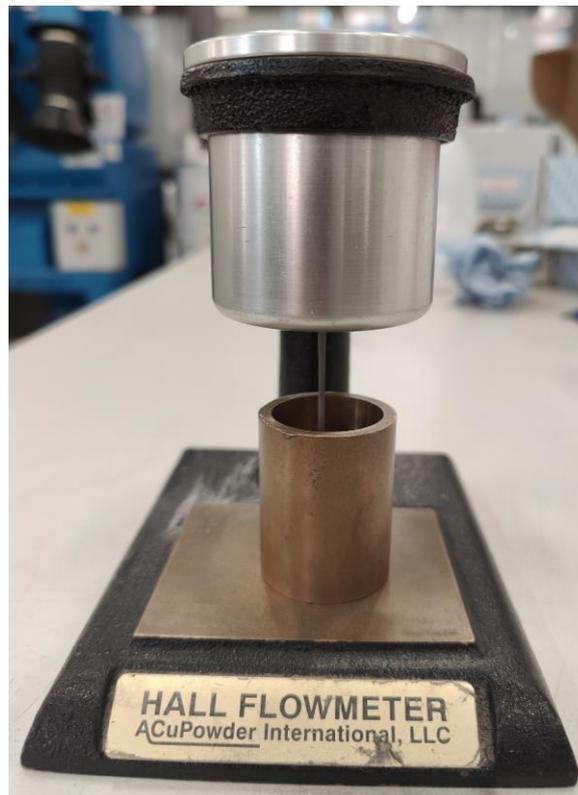


Figure 3.5

IN-718 powder flowing through the Hall Flowmeter

The test was repeated three times for each powder. For the IN-718, the 50g flowed (from the spanner tap to the funnel emptying) in: 13.51s, 13.70s, and 14.09s. This gives an average flow rate of 13.77 s/50g, with a standard deviation of 0.24 s/50g. For the ABD-900AM, the 50g flowed (from first spanner tap to funnel emptying) in: 11.49s, 11.48s, and 11.40s, resulting in an average Hall flow of 11.46 s/50g, and a standard deviation of 0.04 s/50g. These average values were 7.7% larger and 10.4% smaller than the datasheet values, for IN-718 and ABD-900AM, respectively.

The difference between the reported and experimental values is assumed to be an issue of equipment setup. Since the value will not be required further in this project, the discrepancy is not considered important.

3.1.1.3 - Thermo-Physical Properties

Carpenter Additive does not report a solidus and liquidus temperature for IN-718. *Special Metals Corporation*, the current owners of the Inconel trademark, list the alloy's solidus and liquidus temperatures as 1260°C and 1336°C, respectively. *Aubert & Duval* list the solidus and liquidus temperatures of ABD-900AM as 1305°C and 1380°C, respectively. None of these values were verified in-house.

The IN-718 datasheet listed the apparent density (i.e. the powder's density, including the air gaps between powder particles) as 4.12 g/cm³. The ABD-900AM datasheet listed the apparent density as

4.34 g/cm³. Both of these listings followed the ASTM B212 [312] method, which requires that the powder flow freely from the aforementioned Hall flow test funnel, and not be compacted or tapped.

Both of these values were experimentally verified in-house, according to the standard procedure. The cylindrical container used had an internal diameter of 27.7mm, and an internal height of 40.9mm, was filled with powder. Once filled via the Hall funnel, the top of the container was scraped level with a metal ruler, and was placed on a weighing scales. In the case of IN-718, after taking the 137.8g mass of the container into account, the powder inside was found to have a mass of 104.5g, as shown in Figure 3.6.



Figure 3.6

The mass of non-compacted IN-718 powder, inside the cylindrical container, being measured

The mass of the container of powder was measured twice more in this manner (105.3g and 102.8g), to give an average mass of 104.2g:

$$\frac{104.2\text{g}}{\pi \times \left(\frac{27.7\text{mm}}{2}\right)^2 \times 40.9\text{mm}} = 0.00424 \text{ g/mm}^3 = 4.24 \text{ g/cm}^3$$

In the case of the ABD-900AM, the average mass was found to be 108.0g (from 108.7g, 109.6g, and 105.7g):

$$\frac{108.0\text{g}}{\pi \times \left(\frac{27.7\text{mm}}{2}\right)^2 \times 40.9\text{mm}} = 0.00439 \text{ g/mm}^3 = 4.39 \text{ g/cm}^3$$

These flow rates are 2.8% and 1.2% greater than the values reported in the datasheets, for IN-718 and ABD-900AM respectively. The differences between the datasheet and experimentally verified values are assumed to be a matter of powder packing density. The experimental values of 4.24 g/cm³

($\pm 0.06 \text{ g/cm}^3$) for IN-718, and $4.39 \text{ g/cm}^3 (\pm 0.04 \text{ g/cm}^3)$ for ABD-900AM, will be used throughout this work.

3.1.2 - Powder Handling and Storage

Neither of the two batches of powder used in this project were unused when experimentation began. Both tubs had been in use for at least one year. It is uncertain how many times powder was taken from each tub during this time. What is known, is that both before this work began, and during this work, every time powder was taken from either tub, the University of Sheffield's standard procedure for powder use and reuse was followed:

1. The L-PBF machine hopper, build chamber, and overflow container are thoroughly cleaned from previous usage
2. Powder is poured from the tub into the hopper
3. A build is run, and powder gets moved through the chamber and into the overflow
4. After the build, any powder remaining in the hopper is poured back into the tub
5. Powder in the chamber and overflow is collected, and sieved with an appropriate mesh size
6. The sieved powder is returned to the tub
7. The tub is sealed and returned to a low-humidity storage room

For the two powders used in this project, a $53\mu\text{m}$ sieve, designated for use with Ni-based powders only, was used before and during the work. The silica gel sachet that the powder manufacturer's included in the tubs to absorb moisture, remained in the tubs throughout the project. The effectiveness of this pouch five years after its production was undetermined.

3.2 - L-PBF Machine

The L-PBF machine used in this project was designed and manufactured by *Aconity3D*, a German company that specialises in user-configurable devices, allowing workshops and research institutes to add on sensors and similar customisations. The model is the LAB, which was purchased by the University of Sheffield from *Aconity3D* in 2017. The datasheets supplied during purchase list the major specifications of the LAB as follows [313]:

- Build platform diameter - 170mm
- Maximum component height - 200mm
- Maximum laser power - 400W
- Minimum layer thickness - $10\mu\text{m}$
- Maximum laser speed - 4000 mm/s
- Laser wavelength - 1070nm
- Nominal laser spot size - $80\mu\text{m}$

The laser is an Ytterbium-doped (Yb) fibre, manufactured by the American company *IPG Photonics*. The scan head (galvanometers and controls) are manufactured by the German company *Raylase*. The beam profile is Gaussian, with a focal range of -18mm (below base plate) to +8mm (above).

The laser profile datasheets supplied with the system list the widths (defined as the distance at which intensity is $1/e^2$, or approx. 13.5%, of the peak) of the focus levels for a 100W beam as shown in Table 3.3.

Table 3.3
The laser spot size, at various focus levels

Focus Level	Width in X	Width in Y
+ 4mm	126.2 μ m	124.9 μ m
+ 2mm	78.7 μ m	85.2 μ m
0mm	58.8 μ m	68.8 μ m
- 2mm	83.3 μ m	95.6 μ m
- 4mm	134.2 μ m	143.3 μ m

All of the builds in this project will use a focus level of 0mm. 70 μ m will therefore be assumed as the laser diameter, in both X and Y.

The LAB uses a soft blade recoater, with polymer lips. In the workshop at University of Sheffield, argon is used as the shielding gas. The system ensures that the oxygen level is below 2,000ppm at all times when the laser is active. Oxygen levels of around 100ppm are standard during builds. The machine runs on the AconitySTUDIO software.

As input, AconitySTUDIO takes 'Instruction List Table' (.ILT) files, which describe the required laser powers, speeds, positions, and focuses for each vector in each layer in a build. In this project, these files were created using Netfabb, a L-PBF design software created by the American company *Autodesk*. The 2022, 2023, and 2024 versions of the software were used throughout the works, with the differences between these being largely cosmetic.

The LAB is shown in Figure 3.8 and Figure 3.7.

3.3 - Thermal Camera

The LAB is equipped with a thermal camera system, that was designed and fitted in-house. Creating the optical system, calibrating the camera, and developing MATLAB scripts to process the images, required an inter-disciplinary team of researchers from University of Sheffield's School of Chemical, Materials and Biological Engineering (CMBE), School of Mechanical, Aerospace and Civil Engineering (MAC), and School of Electrical and Electronic Engineering (EEE).

3.3.1 - The Camera Setup

The CMOS camera attached to the LAB is designed and manufactured by the Japanese company *Hamamatsu Photonics*. The model is the C11440-42U30, which features a 2048×2048 pixel resolution, at a pixel size of $6.5\mu\text{m} \times 6.5\mu\text{m}$. The camera uses a periscope lens, consisting of six optical elements, incorporated into the lid of the build chamber. The camera observes the build plate from a near-vertical perspective, offset from the build direction by 15° . The sensor's exposure time is set to 1ms. The maximum frame rate that the camera can record at increases, as the frame height decreases, as shown in Table 3.4.



Figure 3.8

An image of the Aconity3D LAB, showing the laser unit, the thermal camera, and the build chamber (with blue viewing port)



Figure 3.7

The Aconity3D LAB from the side

Decreasing the frame height further will not increase the frame rate further, as it is not possible to record more than 1000 frames per second with an exposure time of 1ms.

The rows of pixels that are removed from the frame, as the height shrinks and the frame rate increases, are taken from the top and bottom of the frame equally (i.e., the frame's centrepoint remains unchanged).

Table 3.4
The frame dimensions and maximum recording rates

Frame Width [pixels]	Frame Height [pixels]	Max. Frame Rate [fps]
2048	2048	100
2048	1024	200
2048	512	400
2048	256	800
2048	128	1000

Three filters are placed between the build chamber and the camera:

- A bandpass filter (Thorlabs FL905-10), which blocks all light other than that at $905 \pm 2\text{nm}$ [314]. Acquiring only one specific wavelength is vital to the intensity-temperature conversion, which will be discussed in this section. 905nm sits in the Near-Infrared (NIR) portion of the electromagnetic spectrum, so this type of camera setup may be referred to in literature as a NIR camera. The filter does not perfectly block light outside 905nm, but rather exhibits a narrow transmission window, that peaks at $905 \pm 2\text{nm}$, and drops off on either side. The Full-Width Half-Maximum (FWHM) of this transmission window (i.e. the distance from the peak at which transmission falls to 50% of the peak) is listed as $10 \pm 2\text{nm}$.
- A Neutral Density (ND) filter (Thorlabs NENIR510B), which uniformly reduces the intensity of all incoming light by 90% [315]. The purpose of this filter is to protect the sensor from the very high intensity of light that a melt pool can emit, and to prevent the pixels from becoming oversaturated.
- A short-pass filter (Thorlabs FESH1000), which blocks all light above 1000nm [316]. This filter adds an extra layer of safety, to ensure that none of the extremely high-power laser light (1070nm) reaches the camera sensor.

The camera is connected to a dedicated PC, and is controlled using the HCImageLive software that is included with the purchase of a *Hamamatsu* camera.

The camera can be focused by rotating a gear fitted between it and the build chamber lid. Doing so raises or lowers the height of the camera above the lid, and therefore the height of its focal plane relative to the base plate. Since the camera is offset from the build direction by 15° , rotating the gear also adjusts the focal plane in the Y direction (further into or out of the build chamber). Combined, the result is a line of the base plate being in-focus, lying across the base plate (X direction), that moves forwards or backwards as the gear is rotated. The thermal camera is shown in Figure 3.9.

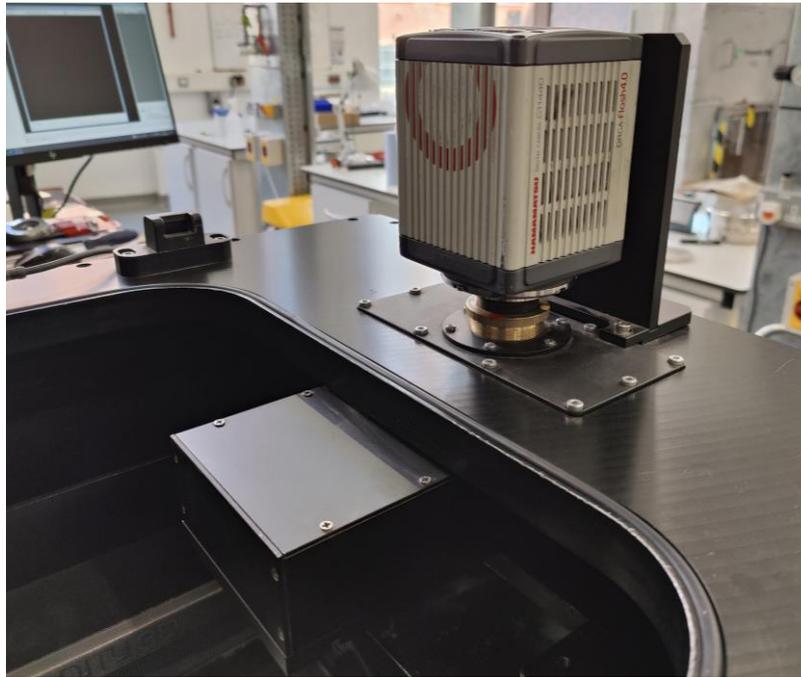


Figure 3.9

A close-up view of the thermal camera, and attached optics inside the build chamber

It was found that, in order to achieve best focus for a build, the following steps should be followed:

1. Ensure that the base plate is level with the build chamber, and the polymer lips of the powder recoater.
2. Scratch very fine lines into the centre of the base plate with a knife, aligned in the Y direction.
3. Turn on the camera and HCLive software, and turn the exposure time to 1 second.
4. Position a desk lamp with an incandescent bulb (which will emit a wide range of visible and infrared light) above the base plate.
5. Turn the focusing gear, and observe the focal line move up and down the length of the scratches.
6. Focus the camera to have the midpoint of the on-screen frame in focus.

3.3.2 - Intensity to Temperature Conversion

The camera and software save images to the PC hard drive in the form TIF images, composed of arrays of 16-bit unsigned integers. Each pixel value in the image refers to the light intensity experienced by its corresponding pixel in the camera sensor. Converting the pixel values from intensity to temperature is a non-trivial process, requiring mathematical relationships between the two, calibrated specifically for IN-718.

The first step in developing these relationships, was to record the 905nm wavelength light emitted from a black body furnace (i.e. no reflection, emissivity value of 1), between the temperatures of 800°C to 1500°C. Planck's Law describes the relationship between the temperature of a black body, and the radiance it will emit at different wavelength:

$$I_{\lambda,T} = \frac{hc^2}{\lambda^5} \frac{1}{e^{\left(\frac{hc}{\lambda k_B T}\right)} - 1}$$

where $I_{\lambda,T}$ is the spectral intensity [W/sr/m³] at wavelength λ [m] being emitted from a blackbody at temperature T [K], h is the Planck constant [J/Hz], c is the speed of light [m/s], and k_B is the Boltzmann constant [J/K].

Plotting $I_{\lambda,T}$ from wavelengths 100 nm to 100 μm , for a few different temperatures, results in the graph shown in Figure 3.10.

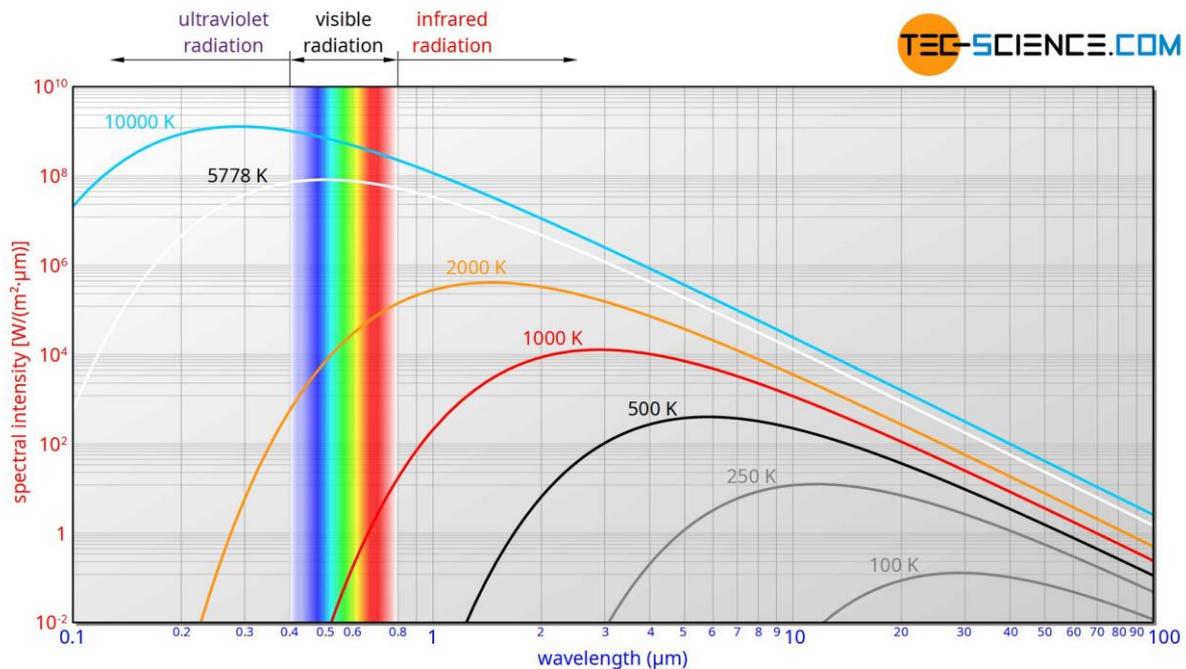


Figure 3.10

The spectral intensities of various black body temperatures, across a range of wavelengths, from Tec-Science [317]

By extracting spectral intensity data for exactly 905nm, it was possible to construct a calibration curve linking intensity values measured by the camera pixels, and temperatures, assuming a black body. Once this calibration curve was applied to pixel values, it was found that the camera could accurately determine temperature, within $\pm 10^\circ\text{C}$, up to 1744°C .

In order to correct for the emissivity values of an IN-718 melt pool, a IN-718 sample was polished to a mirror finish, to simulate the effect. The surface was then heated in an inert gas environment, with radiometers calibrated to 905nm used to determine the surface's emissivity at that wavelength, at temperatures close to those present in the melt pool. An emissivity value of 0.2128 was determined.

By applying this emissivity correction to the black body intensity-temperature calibration curve, it was found that the true temperature at which accuracy falls away is 2125°C , far above the melting range of IN-718.

Once this mathematical relationship was constructed, to convert pixel intensity values to temperatures, specifically calibrated for IN-718, a MATLAB script was written to apply this conversion to every pixel in the raw TIF images. It was assumed that, considering the chemical similarities between the two, the script designed for IN-718 would also be applicable to ABD-900AM.

3.4 - Metallographic Preparation

Much of the ex-situ analysis performed in this project required mirror-flat surfaces before characterisation could begin. Preparing a component from the state it exists in immediately after being removed from the build chamber, to the state required for characterisation, involves some or all of the following steps: removal from the base plate, sectioning to view internal surfaces, mounting in a Bakelite puck, grinding and polishing to a mirror finish, and etching to reveal microstructural detail.

3.4.1 - Removal from Base Plate

Depending on the extent of heat accumulation and residual stress that is expected in a component geometry, it may be possible to build it entirely upon supports that touch the base plate with very minimal contact area. In this case, components may be removed from the base plate, without causing any additional strain or damage, with gentle taps from a hammer and fine chisel. In cases where such supports are not possible, components will need to be removed from the base plate with specialised machinery.

It is possible to remove components with a bandsaw, CNC milling machine, laser cutter, plasma cutter, or waterjet. At the Mechanical Workshop at the University of Sheffield, an Electrical Discharge Machining (EDM) wire cutter is used to remove components. The device is model AQ400L, from the Japanese machining and tooling company *Sodick*.

Cutting on the AQ400L removes about 3mm of material, which must be accounted for in the design of the geometry. The finish left by the EDM wire is far superior to the surface quality of an as-built component, but also far inferior to that required for characterization.

3.4.2 - Sectioning

After a component is removed from the base plate, it often must be cut into smaller sections before being mounted in Bakelite. All sectioning work in this project was carried out on the Secotom-50, a precision, table-top, cutting wheel machine, from *Struers*. Cuts were performed with a 0.76mm thick, stainless steel blade from the *Buehler*. The cutting speed and spindle speed of the machine were altered each time the machine was used, and often throughout individual cuts, in order to maximise cutting speed, while keeping the motor load below 25%.

3.4.3 - Mounting

In order to fit into the grinding and polishing machines, any component that is not deliberately designed as cylinder with a diameter of 32mm and a height of around 20mm, will need to be mounted in Bakelite (or a similar plastic or resin). While cold mounting is possible (as described in Section 3.1.1), hot mounting is preferred for all but the most fragile components, as this method requires a few minutes rather than many hours.

All hot mounting performed in this work took place on the SimpliMet-1000 hot moulder from *Buehler*. After a component is placed (surface of interest down) into the moulder's cylindrical chamber, and around 100g of Bakelite pellets are poured in on top of it, the chamber is sealed, and 29 MPa of vertical pressure is exerted on the contents, at temperatures reaching 180°C. After 2 minutes of heating, and a further 4 minutes of cooling, the puck and encased component is removed from the chamber.

In cases where the components being mounted were intended to be viewed only under the optical microscope, standard, non-conductive Bakelite pellets were used. In cases where conductivity was required, because SEM or EBSD (discussed in detail in this section) was intended, conductive Bakelite pellets, which contain metal powder, were used. In cases where a sample was mounted in non-conductive Bakelite, and it was later decided that conductivity-requiring characterisation was desired, a paint containing silver particles could be applied externally to the puck, to retroactively implement conductivity.

3.4.4 - Grinding and Polishing

In order to transform the as-built, EDM, or sectioned surface into the mirror-like quality required for most characterisation, it must be ground and polished. An AutoMet-250 from *Buehler* was used for both grinding and polishing in this project. This device holds up to six samples, in the form of 32mm cylinders, against a rotating pad of 254mm diameter. Stuck to the pad, either magnetically or by sticky adhesive, is a grinding paper or polishing cloth. The following machine parameters are all controllable:

- The speed of rotation of the pad
- The speed of rotation of the six-sample holder
- The direction of the six-sample holder rotation (with or counter to the pad)
- The force with which each sample is pressed into the pad
- The duration of the operation
- The positioning and flow of a lubricating water tap

Designing an appropriate grinding and polishing route for a given experimental setup involves a lot of trial and error. The steps required will depend on the alloy, and the exact types of grinding papers and polishing liquids being used. Often, it is only once the final polishing steps are performed, that

it will be realised that an error was made earlier, resulting in deep scratches being left in the surface. As a general rule, before a technician can move on to a finer grit, the component should be viewed under an optical microscope, at an appropriate resolution, to ensure that all scratches are uniform and aligned.

After much experimentation, it was found that the grinding route outlined in Table 3.5 worked well for both materials used in this project, with the equipment available at the University of Sheffield.

*Table 3.5
The grinding route used in this project*

Grit	D ₅₀ [μm]	Pad Rot. Speed [RPM]	Holder Rot. Speed [RPM]	Applied Force [N]	Time [min]	Water	Repeat
P60	260	150	60	10	1	Full	X 3
P120	127	150	60	10	1	Full	X 3
P240	65	150	60	10	1	Full	X 3
P400	35	150	60	10	1	Full	X 3
P800	21.8	150	60	10	1	Full	X 3
P1200	15.3	150	60	10	1	Full	X 3
P2500	8.4	150	60	10	1	Full	X 3

Generally, as-built surfaces required the full treatment outlined in Table 3.5, while EDM surfaces could begin from the P120 step, and sectioned surfaces could begin from the P240 step. All grinding papers were purchased from company *MetPrep*. All sample holder rotation direction was counter to the pad rotation direction.

This grinding route was followed by the polishing route, outlined in Table 3.6, using diamond suspensions, Oxide Polishing Solution (OPS), and water.

*Table 3.6
The polishing route used in this project*

Liquid	Pad Rot. Speed [RPM]	Holder Rot. Speed [RPM]	Applied Force [N]	Time [min]	Water	Repeat
3μm	300	60	20	8	Min.	X 1
-	300	60	20	2	Full	X 1
1μm	300	60	20	8	Min.	X 1
-	300	60	20	2	Full	X 1
OPS	300	60	20	10	Min.	X 1
-	300	60	20	10	Full	X 1

Again, all sample holder rotation was counter to the pad rotation direction. Immediately after the final step in this route, samples were submerged in beakers of isopropanol, and placed inside an ultrasonic bath for several minutes, to remove any remaining polishing liquid embedded on their surfaces.

3.4.5 - Etching

Etching refers to the selective corrosion of different features on a sample's surface. Due to segregation during solidification, it is likely that separate phases, grain boundaries, and cell boundaries all have slightly different chemical compositions. By subjecting the surface to chemical solutions which dissolve certain compositions more than others, it is possible to make certain features much more prominent in characterisation.

Determining which chemical etchant is best for revealing which features in which alloys, plus finding a suitable method and duration for applying the etchant, is an extremely complex process, likely to involve a significant amount of trial and error.

After a period of trial and error with the chemicals and equipment available in the University of Sheffield laboratory, it was found that Glyceregia (three-parts hydrochloric acid, two-parts glycerol, and one-part nitric acid) revealed melt pool boundaries, grain boundaries, and sub-grain cell boundaries, in sufficient detail for characterisation.

After the Glyceregia mixture was created in a fume cupboard, the following procedure was enacted:

1. Swab the sample surface with a cotton ball damp with Glyceregia, for X number of seconds
2. Immediately dunk the sample in a beaker of isopropanol, and swirl for a few seconds
3. Remove the sample from the fume cupboard, and view the surface under an optical microscope
4. If the features are not yet visible, repeat from Step One

The duration X of the swab decreased upon repeats, usually from 10, to 5, to 2, then to 1 for all further swabbing.

3.5 – Microscopy

The following microscopy methods were used throughout this project:

3.5.1 - Optical Microscopy

In order to detect porosity, melt pool boundaries, grain boundaries, and sub-grain cell boundaries sized in the range $10^2 \mu\text{m}$ to $10^0 \mu\text{m}$, optical microscopy is often both sufficient, and the most straightforward characterisation method. The optical microscope used in this project was model BX51, designed and manufactured by the *Olympus*. The microscope's motorised stage, and the software used to control the entire system, were produced by *Clemex*. The BX51 is equipped with four interchangeable lenses, allowing for easy switching between 50X, 100X, 200X, and 500X magnification. All images produced on the microscope were analysed with custom-written MATLAB scripts. These scripts, and a variety of others, are available to download on [GitHub](#).

3.5.2 - Electron Microscopy

In situations where optical microscopy cannot provide the detail required, either because features are too small or light cannot detect them, electron microscopy can be a solution. Inside an electron beam device, an electron gun generates a stream of high-energy electrons, which accelerate from a filament (such as a tungsten wire or crystal), past a positively charged anode, and through a series of electromagnetic lens which direct the beam towards the specimen inside a vacuum chamber, striking its surface as a focussed spot with a diameter of several nanometres. As electrons strike this spot they penetrate into the specimen, interacting with the atoms in a tear-shaped volume that extends below the surface, called the interaction volume. Detecting both the backscattering of these electrons, and the emission of secondary electrons from the interaction volume, allows for the imaging of microstructural features through SEM and EBSD.

3.5.3 - Scanning Electron Microscopy (SEM)

The SEM used in this project was the Inspect FEI F50. The working distance, acceleration voltage, magnification, dwell time, spot size, and step size were all varied throughout this project, depending on the exact surfaces and features being imaged. All of these values will be reported each time an SEM image is presented in further sections. All SEM images produced in this project were analysed with custom-written MATLAB scripts.

3.5.4 - Electron Backscatter Diffraction (EBSD)

All EBSD work conducted in this project was performed at the Henry Royce Institute at the University of Manchester, facilitated by Dr Jack Donoghue. The SEM device used was an Apreo C, made by *Thermo Fisher*. The device was equipped with the Symmetry S2 EBSD detector, running on AZtec 6.1 software, from the British company *Oxford Instruments*. The accelerating voltage, working distance, step size, and other such parameters from the images will be laid out in later chapters.

All EBSD image analysis was carried out with a combination of AZtec 6.1, custom-written MATLAB scripts, and the MTEX MATLAB toolbox.

3.6 - X-Ray Computed Tomography

The two methods of microscopy discussed previously involve the bombardment of a surface with light (optical microscopy), and the bombardment of a surface with electrons (SEM/EBSD/EDS). A further methods exist, involving the transmission of electromagnetic radiation through a sample.

X-Ray Computed Tomography (XCT), which involves the transmission of X-rays through a sample, is a very effective, non-destructive characterisation method of detecting pores, cracks, density variations, inclusions, grain structures, and even dislocations inside metal specimens.

The underlying fundamentals behind the kind of XCT facilities used to analyse metal samples, are the same as those behind the X-ray machines used to scan medical patients; and just as rotating an X-ray scanner around a human inside a CT scanner allows for multiple 2D images to be combined into a 3D map of their body, so too does rotating a metal sample inside a XCT facility allows for a 3D map of features.

As an X-ray passes through a metal sample, there are two primary mechanisms which may disrupt its passage to the detector on the far side: absorption and scattering. Absorption occurs when an X-ray's energy is transferred into the sea of free electrons, through excitation of electrons into higher-energy orbitals. Therefore, more absorption occurs in samples with greater electron sea densities, which may arise from higher atomic numbers, tighter-packed crystal lattices, or lower porosities.

Scattered X-rays are deflected from their paths, whether through elastic scattering (whereby an electron-photon interaction alters the X-ray's direction but not frequency) or inelastic scattering (due to the electron absorbing some of the X-ray's energy, both the direction and frequency is altered). In both cases, the deflected X-ray may go on to impact the detector at an unexpected location, thereby causing scatter noise artefacts.

After reconstruction (the process of combining multiple 2D XCT images into a 3D map), each voxel inside the volume is assigned an intensity value, determined by the ease with which the X-rays passed through its space. In the most advanced XCT facilities, voxels of sub-micrometre sizes can be achieved, allowing for the detection of microscopic pores, grain structures, or dislocations, depending on the sample and setup.

All XCT work carried out in this project was performed at the European Synchrotron Radiation Facility (ESRF) in Grenoble, France, on BM18. BM18 uses a polychromatic beam, at 50 to 280 keV. More details on the setup will be provided in further chapters.

Once reconstructed into stacks of TIF files, the data from each XCT scan was several TB in size. Processing this quantity of data required the use of the University of Sheffield's High Performance Computer (HPC). The Stanage cluster was used. This HPC operates on Linux 7, and features 90TB of RAM, split between 12,000 CPU cores capable of operating in parallel.

Reconstructed XCT TIF files, and the custom-written MATLAB scripts that analysed them, were uploaded to the HPC. The HPC performed the analysis, and output data was downloaded. An example XCT image, obtained from the ESRF, is shown in Figure 3.11.

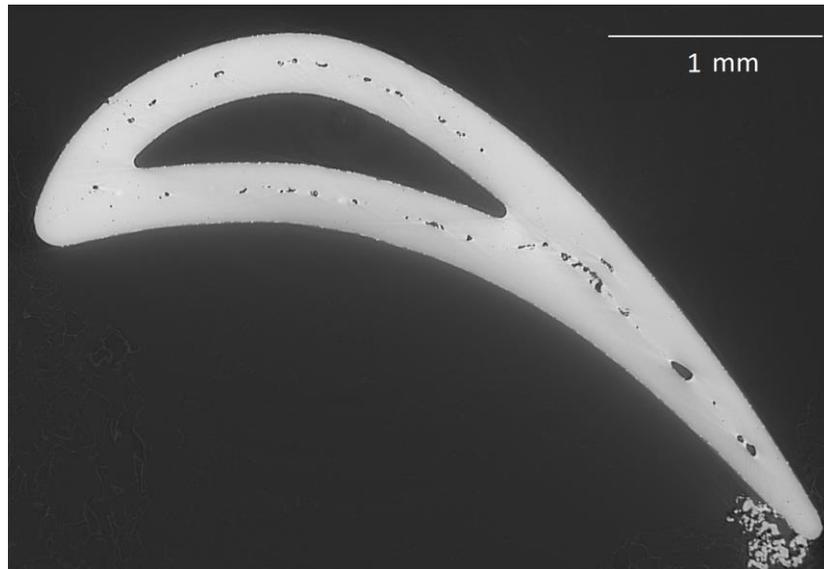


Figure 3.11

An example of a reconstructed TIFF image, of an IN-718 component, as recorded at BM18 at the ESRF

3.7 - Hardness Measurements

All hardness testing in this work was performed on a DuraScan 70. The Vickers methodology was used with a 1kg mass, with at least 5 measurements being taken from each region being tested on samples. All measurement points were selected to ensure that they were at least 1mm away from the nearest visible pore, in order to maximize the chances of the indentation not interacting with porosity.

3.8 - ThermoCalc Modelling

Thermo-Calc 2025a was the CALPHAD (CALculation of Phase Diagrams) software used in this work. The ‘TCNI12: Ni-Alloys v12.1’ database was used for phase properties. Modelling work was only carried out on IN-718, which was defined in the software as having the composition (weight %) shown in Table 3.7.

Table 3.7

The IN-718 composition (weight %) used by ThermoCalc

Nickel	52.194 %	Chromium	19.0
Iron	18.0 %	Niobium	5.0 %
Molybdenum	2.9 %	Titanium	1.0 %
Cobalt	1.0 %	Aluminium	0.6 %
Copper	0.3 %	Boron	0.006 %

Thermo-Calc’s ‘Scheil Calculator’ module predicted the likely phases to exist in this mixture, by running a simulation of the composition cooling from 5,000K, to full solidification using the Scheil method, in temperature steps of 1K, at 1atm pressure.

This data was then passed to Thermo-Calc's 'Additive Manufacture' module, which solves for a moving heat source, using analytical methods. The system predicts steady-state melt pool dimensions, given input parameters such as: ambient temperature; base plate temperature; heat source profile (Gaussian, top hat, etc.); beam radius; laser power; laser speed; powder layer thickness; and several others. The multi-physics approach used by Thermo-Calc accounts for conduction, convection, evaporation, radiation, and fluid flow. An example of a melt pool simulated with this method is shown in Figure 3.12.

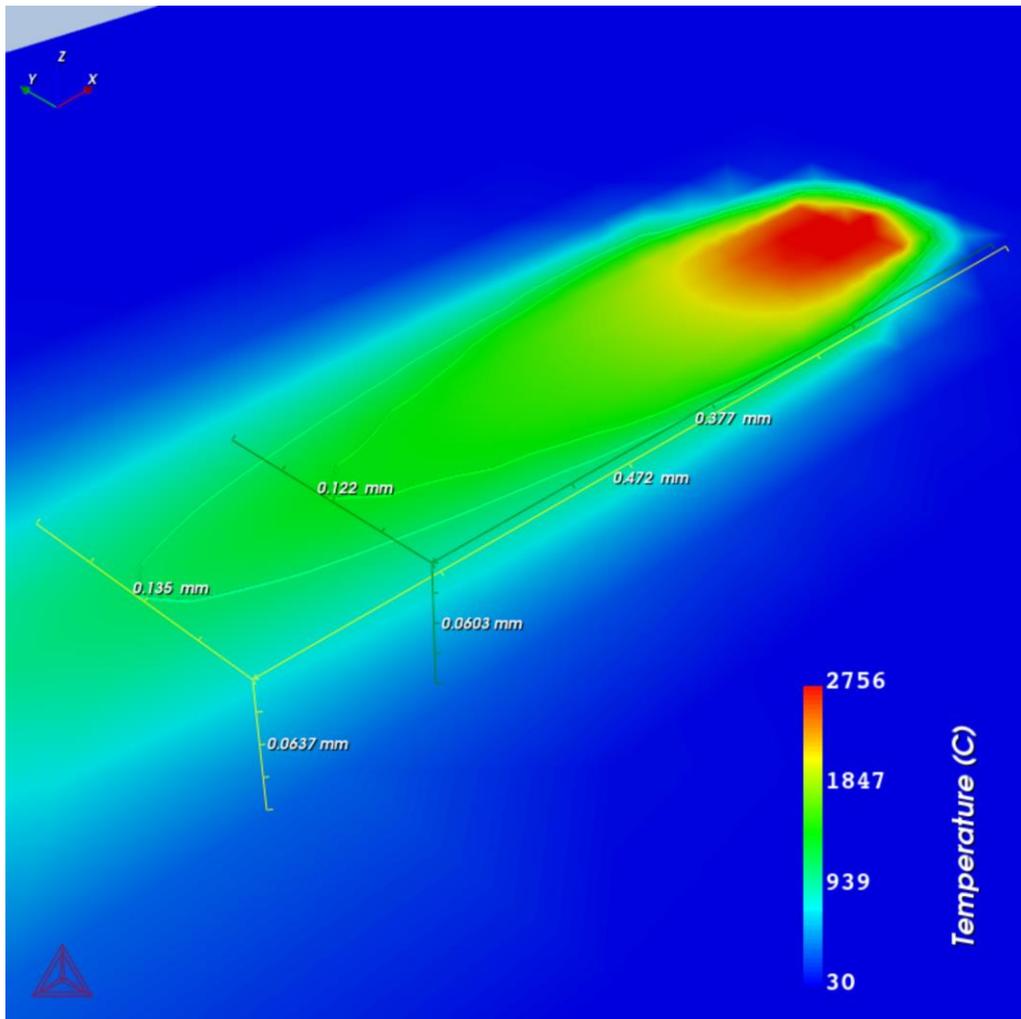


Figure 3.12

An example of the temperature profile and dimensions (liquidus in green, solidus in yellow) of an IN-718 melt pool, as predicted by Thermo-Calc

4 - Camera Characterisation

This experimental chapter details the initial tests undertaken to understand the strengths and limitations of the thermal camera used in following chapters, and to determine the corrections that must be applied to image to account for motion blur, focus levels, and emissivity. Melt pool metrics obtained from the camera are compared to those obtained from the optical microscope and from modelling methods, to explore the fidelity of these methods.

4.1 - Determining Camera Frame Edges

A single-layer geometry was designed, with the intention of being able to track the point at which the melt pool intersected the frame. i.e.:

1. A slow speed (for easier tracking) and low power (to avoid spatter) pattern would be scanned onto a powder-free baseplate, with the camera recording.
2. The image sequence would be carefully analysed afterwards, with the points in the sequence at which the melt pool crossed over the frame being noted.
3. Working backwards, it would be possible to determine the points in the Netfabb design, at which the melt pool crossed out of frame.
4. The dimensions and corners of the frame in the Netfabb software could then be determined.

Designing a geometry that allowed for this process was not straightforward. At the frame edges and corners, where the melt pool crossed out of view, the baseplate was often out of focus. It therefore became difficult to work backwards from the image sequences, and determine which scan vector according to the Netfabb file belonged to which image. Ultimately, a hub and spoke model was used for the test.

A series of 50 concentric circles, with radii 50mm, 49mm, 48mm, etc., were centered upon the approximate midpoint of the camera frame (1mm to the right of the baseplate centre, and 21mm below, as determined by the previous test). A set of 72 lines, extending from the circle centres to the outer circumference, were spaced evenly around the span, at steps of 5°. The Netfabb layout is shown in Figure 4.1.

The circles were scanned first, from the smallest to the largest. By keeping track of the circle numbers in the image sequence, and by noting the circle at which each of the four corners appeared to intersect with melt pools, it was possible to determine the distance of each corner to the centrepoint. Next, the lines were scanned, starting from the topmost and proceeding clockwise. Again, by keeping track of the line number and noting corner intersections, the approximate angle of the corners from the centrepoint could be determined.

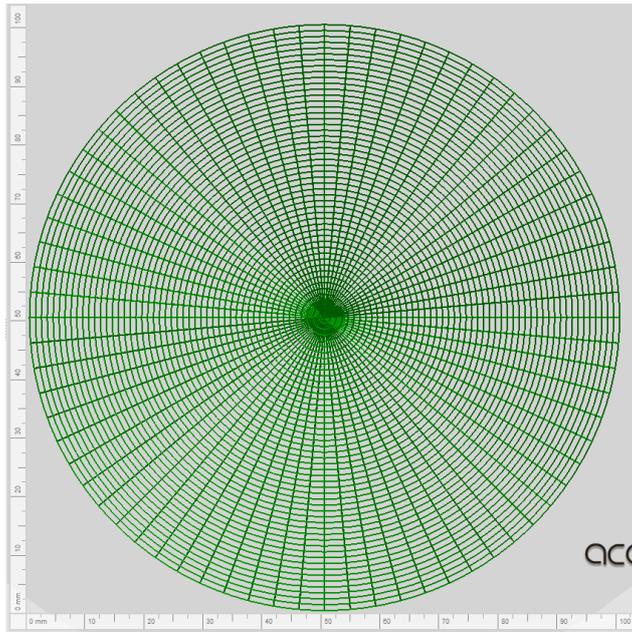


Figure 4.1

The Netfabb layout used to determine frame boundaries with melt pools

Compiling these distances and angles, the coordinates of the corner points, for each of the five available frame sizes, could be determined, as a function of offset from the baseplate centre on the Netfabb centre. Each coordinate is assumed to be accurate to within ± 0.25 mm. These coordinates are shown in Table 4.1.

Table 4.1

The coordinates of the camera frames on the baseplate

Width [pix]	Height [pix]	Corner Coordinates Relative to Baseplate Centre [X, Y]			
		Top Left [mm]	Top Right [mm]	Bot Right [mm]	Bot Left [mm]
2048	2048	[-23.25, -1.5]	[22.75, 0.5]	[17.0, -39.5]	[-18.0, -41.5]
2048	1024	[-22.0, -11.5]	[21.25, -9.5]	[18.25, -29.5]	[-19.25, -31.5]
2048	512	[-21.75, -16.5]	[20.25, -14.75]	[18.75, -24.5]	[-19.5, -26.75]
2048	256	[-21.25, -19.0]	[20.0, -17.5]	[19.0, -22.5]	[-20.5, -24.75]
2048	128	[-20.75, -20.2]	[19.75, -18.0]	[19.5, -20.75]	[-20.5, -22.75]

A sketch drawn in Netfabb is shown in Figure 4.2, which marks the approximate edges of the frames, from largest (red) to smallest (yellow). Three important notes can be made from this sketch:

- All frames share the same lateral edges (shown in red).
- The frame is rotated counterclockwise relative to the baseplate, at an angle of approximately 2.7° .
- Because of the 15° tilt of the camera relative to the vertical, the top edge of each frame is larger than its bottom. This difference decreases from 28%, to 16%, 10%, 5%, and 1%, as the frame height decreases from 2048px, through to 128px. There is no difference (within margin of error) between the left and right edges of each frame.

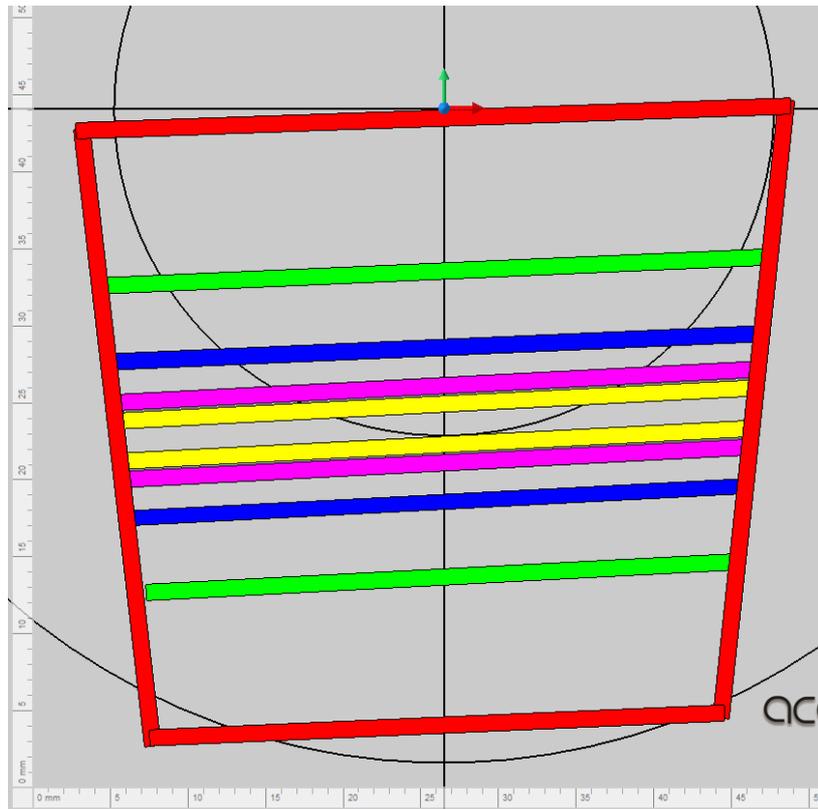


Figure 4.2

A Netfabb model of the frames to be used throughout this project

From this point onwards, each frame will be assumed to be an isosceles trapezium (i.e. top and bottom edges are parallel with unequal lengths, left and right edges are non-parallel with equal lengths), centred at a point 1mm to the right of the baseplate centre and 21mm below, rotated 2.7° counterclockwise. The lengths of the sides are shown in Table 4.2.

Table 4.2

The side lengths of each thermal camera frame

Frame Height [px]	Top Edge [mm]	Bottom Edge [mm]	Left/Right Edges [mm]
2048	45.5	34.5	41.5
1024	43.0	35.0	21.5
512	42.5	36.0	10.0
256	41.5	38.5	6.0
128	40.0	39.5	3.0

When a MATLAB script is converting from pixels to μm , therefore, the conversion factor is not constant, but instead depends on the location of the object being analysed within the frame. Assuming that the camera is focused at the frame midpoint (where 2,048px is 39,750 μm) the conversion factor will decrease as objects appear lower in the frame (to a minimum of 34,500 / 2,048 $\mu\text{m}/\text{px}$ for the bottom of the 2,048 frame), and increase as objects appear higher in the frame (to a maximum of 45,500 / 2,048 $\mu\text{m}/\text{px}$ for the top of the 2,048 frame).

Formalising this, the width conversion factor could be taken as:

$$k_w = 19.41 + 0.00262n$$

where k_w is the conversion factor [$\mu\text{m}/\text{px}$], and n is number of pixel rows that a melt pool's centroid appears above (positive n) or below (negative n) the frame's central row.

However, implementing this complicated conversion will be shown to be unnecessary in Section 4.5, where it will be shown that any melt pool further than ~ 200 pixel rows from the frame's centre will be too defocused for any meaningful analysis. Therefore, for all width conversions, k_w will simply be taken as $19.4 \mu\text{m}/\text{px}$.

The height of the isosceles trapezium defined in Table 4.2 is calculated as $41,130\mu\text{m}$ for the $2,048\text{px}$ frame. Although some complex trigonometry could be used to calculate the extent to which the height conversion factor varies across the span of the frame, for the same reasons as outlined previously, height conversion will simply be taken as k_h is $20 \mu\text{m}/\text{px}$.

4.2 - Determining Spatial Accuracy

Having found that the 15° offset angle of the camera can stretch/shrink the top and bottom edges of the frame by up to 28%, an obvious follow-on question is raised: how well do locations of objects or melt pools on a Netfabb file, match to locations of objects or melt pools on a camera image.

To test this, a simple setup was devised. A grid of dots was lasered onto the baseplate. The grid had gradations of 1mm in both X and Y, and each dot was composed of a 0.1mm tall single-track line, oriented vertically. Two solid single-scan lines were drawn, horizontally and vertically, through the frame centrepoint. The pattern was scanned at 125W and 500mm/s, and recorded in a single long-exposure camera image. The pattern was repeated for each of the five frame sizes, and for different focus levels when larger frames necessitated it:

- 2048 - Repeated 5 times (camera focused at frame's top edge, mid-top, mid, mid-lower, lower edge)
- 1024 - Repeated 3 times (camera focused at frames top edge, mid, lower edge)
- 512 - No repeats necessary (all dots were in focus)
- 256 - No repeats necessary
- 128 - No repeats necessary

Figure 4.3, as an example, shows the long-exposure image of the 1024px frame, when the camera was focused on its middle.

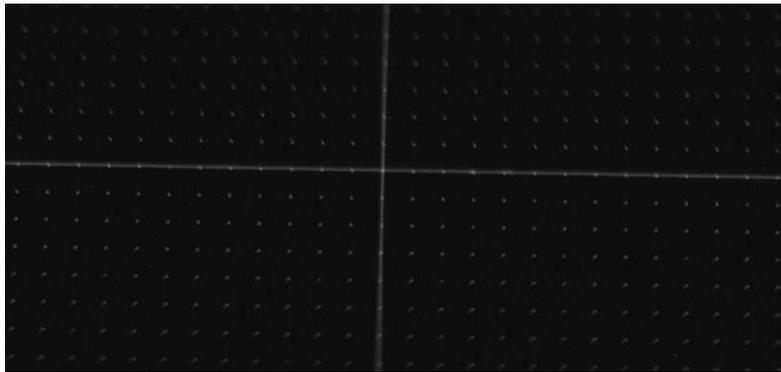


Figure 4.3

A long-exposure thermal image showing all the melt pool dots

Where repeats at different focus levels were used, the images were stacked together. i.e. the uppermost 20% of the stacked 2048px image came from the image with focus on the top edge, the next 20% came from the image with focus on the mid-top, etc. Once a single image was compiled for each frame size, a MATLAB script was written to compare the location of each dot on the images, to the location of the dots on the Netfabb file.

In Figure 4.4, the black dots represent the centroids of the melt pools identified in the composite image of those recorded at 2048px frame size. The red dots represent the locations of the 0.1mm single-track lines according to the Netfabb file. The grid of red dots has been rotated by 2.7° and scaled, so that the centres of both sets of dots are aligned.

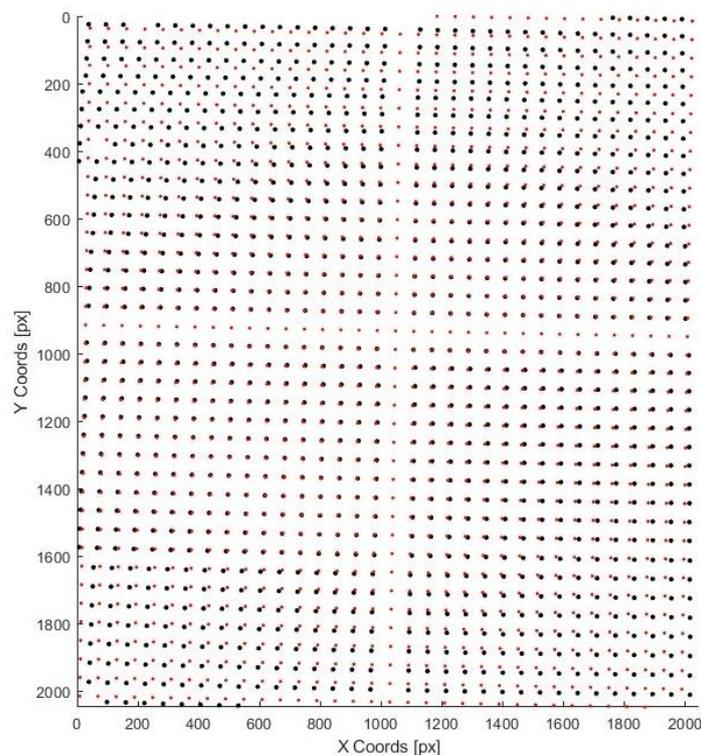


Figure 4.4

The locations of the recorded melt pools (black dots), and the Netfabb lines they were based on (red dots)

As is clear in the image, the further objects are from the centre of it, the greater the misalignment between the Netfabb locations and the recorded melt pool locations. It was found that the 1024px and 512px frame misalignment were visually-identical to that of the 2048px, albeit cropped in the Y direction.

It is noted how, in the lower part of the image, the melt pool locations are further from the centre than their corresponding single-line track locations (i.e. the black dots are ‘moving away’ from the centre faster than the red). In the upper part of the image, where the red dots are moving away from the centre faster than the black, the opposite is true. It appears that both expansive and contractive linear distortion is being applied to the image.

As shown in Figure 4.5, a simple visual was created, to provide a sense of the direction and scale of misalignment between Netfabb object and corresponding melt pool image, at various locations throughout the frame. Each vector shows the extent to which the distance between the Netfabb object (red dot) and centrepoint either shrinks (negative percent), or expands (positive percent) when captured as the distance between the melt pool (black dot) and centrepoint.

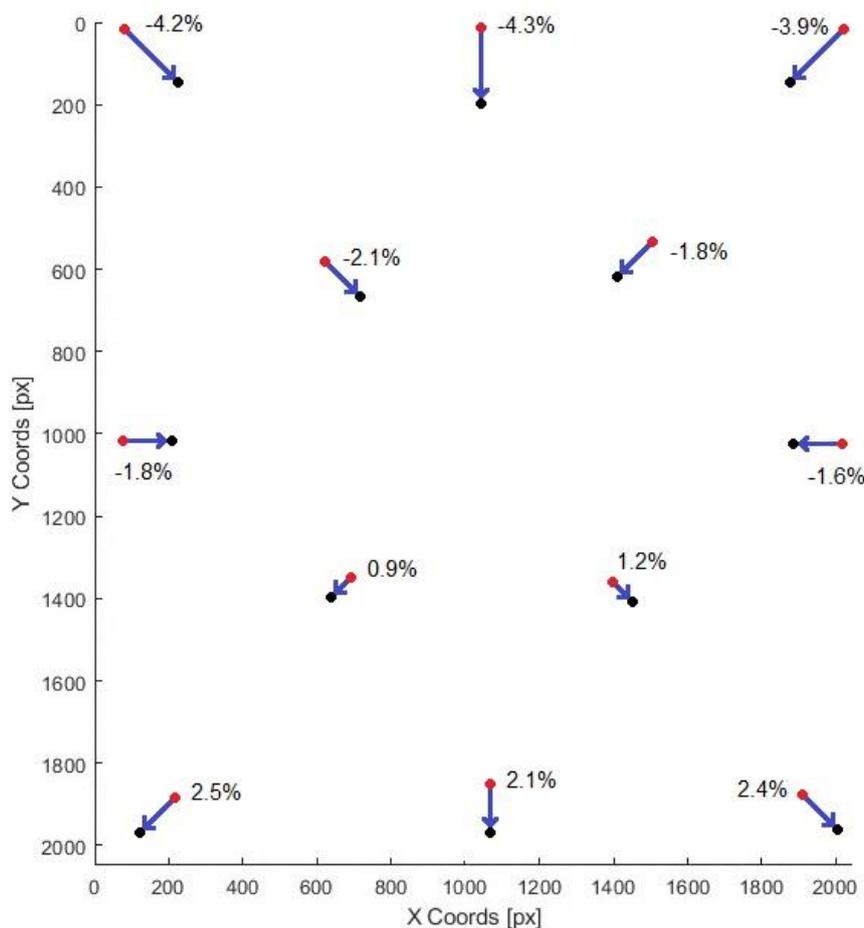


Figure 4.5

A representation of the direction and percentages by which objects in thermal images will scale away (positive percentages) or towards (negative percentages) the centrepoint, as compared to their Netfabb counterpart, based on their location within the frame

It must be noted that the reason that misalignment is largest towards the corners in Figure 4.5, is simply because the centre was chosen as the point to overlay the grids of dots representing the Netfabb and melt pool locations. Should the bottom-right of the image have been chosen instead, for example, then misalignment would have appeared to increase as location became higher and more leftwards. Therefore, only one straightforward conclusion can be drawn from these tests: distance measurements extracted from thermal camera images must be assumed to become distorted as they grow larger, at the rate of approximately 1% per 10mm.

4.3 - Determining Spatial Resolution

The USAF 1951 Resolution Test Chart is a standardized optical pattern, developed in 1951 by the US Air Force. Its purpose is to determine the ability of optical systems to differentiate between distinct objects at a small scale. The pattern is composed of alternating dark and bright lines, in sets of three (i.e. light, dark, light, dark, light), running both vertically and horizontally. There are six elements (line sets), with decreasing thicknesses, in each group, and twelve or more groups of decreasing size in each pattern. The resolution of the line sets follows the formula:

$$R = 2^{\left(G + \frac{E-1}{6}\right)}$$

where R is resolution [line pair per mm], G is group number, and E is element number.

A positive reflective test chart (Thorlabs R3L3S1PR) was used to determine spatial resolution [318]. This chart contained low-reflectivity chrome lines, upon a high-reflectivity chrome background. According to the chart shown in Figure 4.6, at the 905nm wavelength that the camera records, the lines and background are approximately 44% and 61% reflective, respectively.

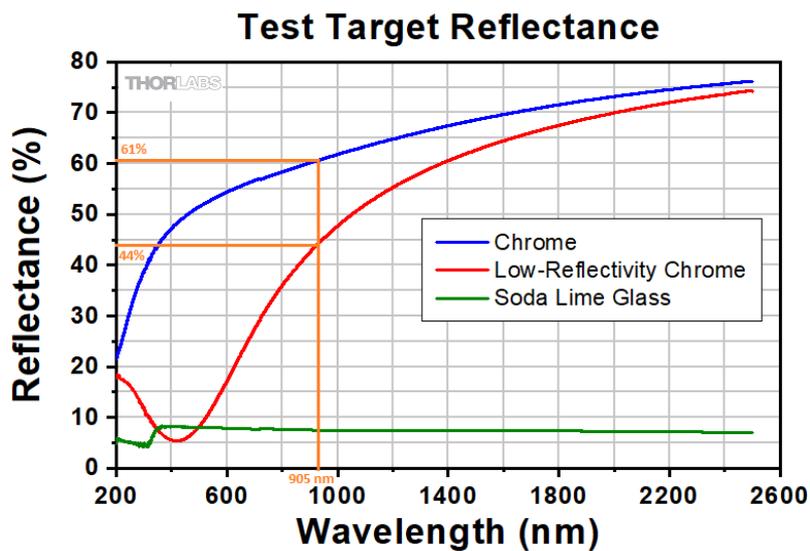


Figure 4.6

The test chart's reflectance, as a function of wavelength [318], with the 905nm reflectances highlighted in orange

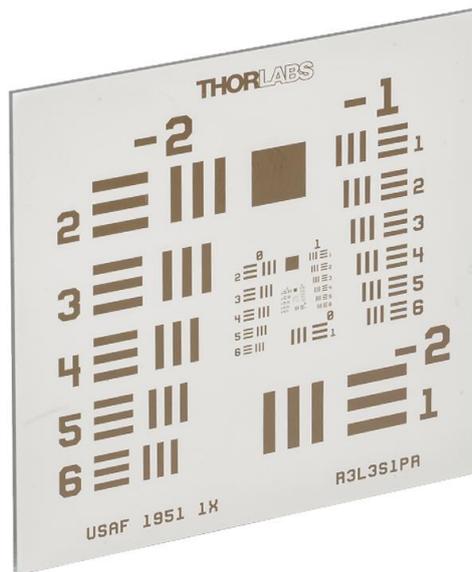


Figure 4.7
The USAF 1951 test chart [318]

The test chart, shown in Figure 4.7, was placed onto the build plate, centred within the thermal camera frame. In a similar method to that described previously, images were recorded at 1s exposure time, while an infrared lamp was shone into the build chamber and moved about, and the camera focus was adjusted, to ensure that each element from each group was both illuminated and in focus, for at least one image. Shown in Figure 4.8, as an example, is an image showing Element 3, from Group 0 (1.26 lp/mm).

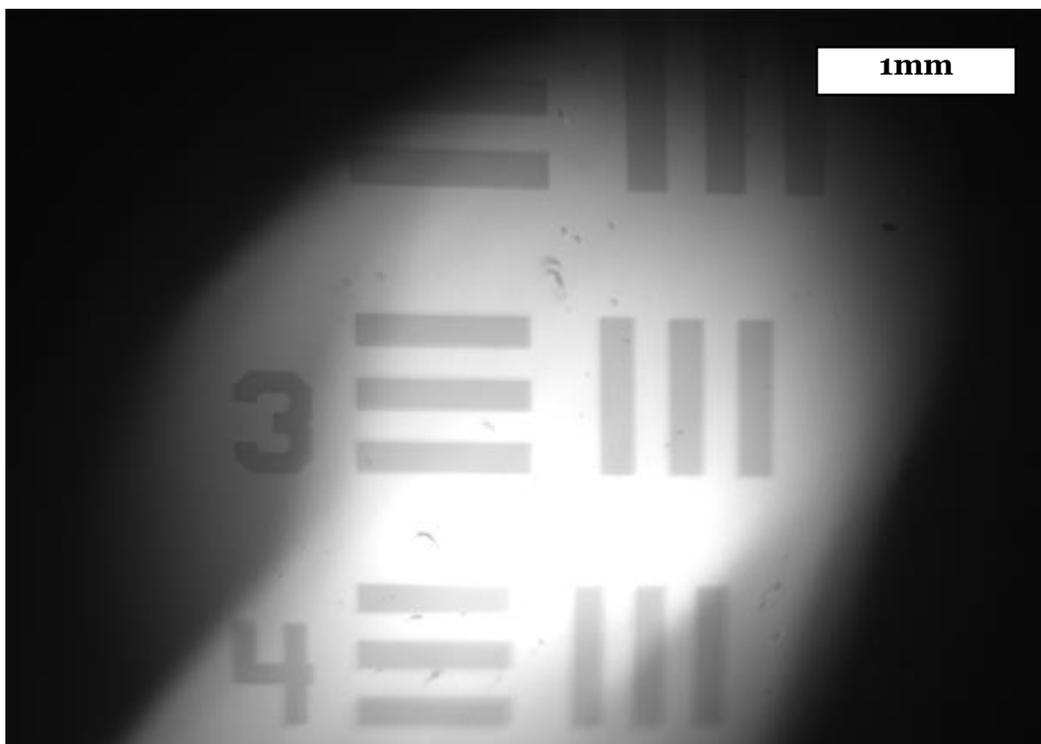


Figure 4.8
A thermal image, showing Group 0, Element 3

In order to quantify how well the camera could distinguish between the lines, a Modulation Transfer Function (MTF) test was performed. An MTF test determines how well the output images of an optical system preserve the contrast of the objects it records, at different spatial lengths. The MTF equation compares the contrast out (i.e. between objects as detected by the camera), to the contrast in (i.e. between these objects in reality):

$$MTF = \frac{C_{out}}{C_{in}}$$

$$C_{out} = \frac{I_{max} - I_{min}}{I_{max} + I_{min}}$$

where I_{max} and I_{min} are the intensity values of the light and dark regions of each element, respectively. The equation for input contrast is the same as that for output contrast, but with the ideal values of intensity being used instead of those picked up by the camera (i.e., for this USAF target, 60 and 42).

For example, in the Figure 4.8 image of element 3 from group 0, the intensities of the light and dark bars were measured as 63,816 and 50,932. Therefore, the MTF for these bars (1.26 lp/mm) is calculated as:

$$MTF_{1.26} = \frac{63816 - 50932}{63816 + 50932} \bigg/ \frac{60 - 42}{60 + 42} = 0.112 / 0.176 = 0.636$$

Calculating the MTF for each element, and plotting as a function of lp/mm, yields the chart in Figure 4.9.

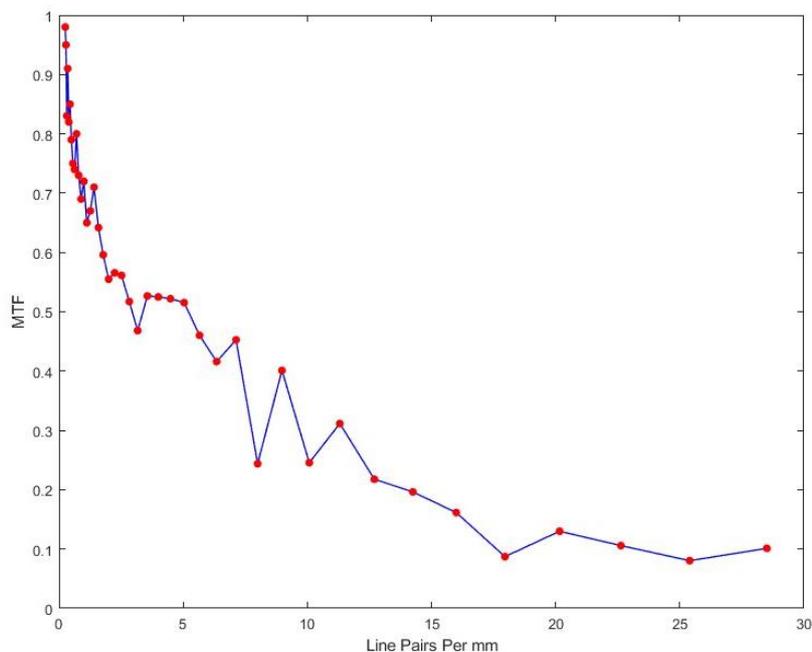


Figure 4.9

The results of the MTF test, in terms of Line Pairs Per mm

As can be seen, the MTF drops rapidly as the lp/mm increases from near-zero. The MTF falls below the industrially-accepted “good enough” value of 0.5, once the lp/mm reaches around 5. Since lp/mm is not a very intuitive metric for humans, it is useful to graph instead by the inverse of twice the lp/mm (i.e. the width of a single line), as shown in Figure 4.10.

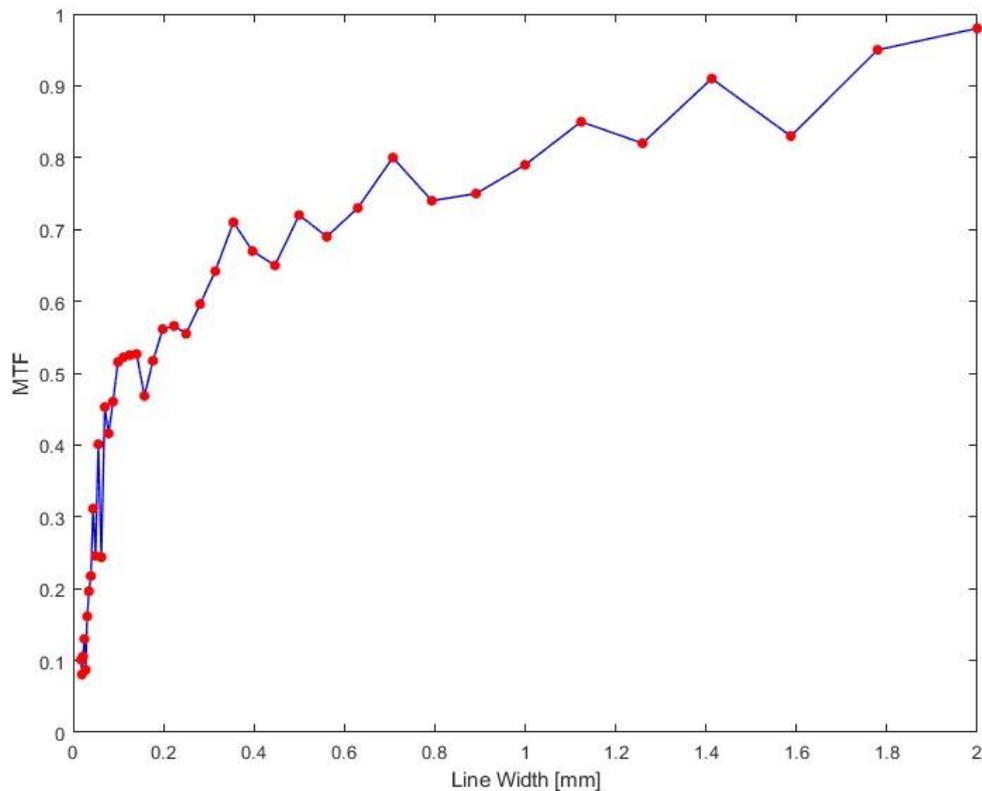


Figure 4.10
The results of the MTF test, in terms of Line Width

From this graph, it can be seen that when the size of a feature is smaller than approximately 200µm, the MTF begins to fall rapidly below the “good enough” target of 0.5.

4.4 - Understanding the Effects of Motion Blur

The exposure time of the thermal camera, outside of characterization tests, was consistently set to 1ms. During this time the melt pool will travel, thereby appearing more extended in the resulting images than it would be in reality (i.e. in a ‘true’ snapshot with an exposure time of 0ms). It stands to reason that the extent to which the image is extended can be simply calculated based on the exposure time and laser velocity (V). For example, if a melt pool is travelling at 1000mm/s, and the exposure time is 1ms, then the extension would be:

$$L_{extend} = 1000\text{mms}^{-1} \times 0.001\text{s} = 1\text{mm}$$

Therefore, it would seem as if every melt pool length measurement extracted from thermal camera images should be subtracted by 0.001V. In order to test this theory, single-line track melt pools were

scanned at various combinations of velocities and exposure times. At least 40 melt pools were recorded, at 125W on a 316L steel baseplate using the 1024px frame size at 200fps, at all combinations of the metrics shown in Table 4.3.

*Table 4.3
The speed-exposure combinations*

Velocities [mm/s]	Exposure Times [ms]
500, 600, 700, 800, 900	1, 2, 3, 4, 5

It was reasoned that, by measuring the melt pool length for each velocity at 5ms, 4ms, 3ms, etc., it would be possible to extrapolate backwards to estimate what the ‘true’ length would be at 0ms. All scanned lines ran horizontal across the frame, at approximately its midpoint in the vertical, where the width of the frame was $40 \pm 0.5\text{mm}$. Therefore, to convert melt pool length from pixels to μm , a conversion factor of $40,000\mu\text{m}$ per $2,048\text{px}$ ($19.53 \mu\text{m}/\text{px}$) was used, with an understanding that this conversion factor was accurate to within $\pm 0.24 \mu\text{m}/\text{px}$.

Images were normalised, and a threshold intensity of 0.4 was set in the MATLAB code to distinguish melt pool from background, and the in-built REGIONPROPS function was used to measure the dimensions of the identified pools. When recording melt pools, the camera occasionally imaged the start or end of a scan track, where the identified pool may be much smaller than it should be. In order to remove these smaller pools, the in-built RMOUTLIERS function was used, with the Median Absolute Deviation (MAD) method.

In this method:

1. The median of the >40 pool lengths was calculated (M_l)
2. The absolute value of the deviation between each length and this median was calculated (D_i)
3. The median of these deviations was calculated (M_d)
4. A multiplier value (k) was chosen to be 3, as per common practice to consider any data further than 3 standard deviations from the mean (i.e. the 0.3% most extreme) as outliers
5. Any length, with deviation greater than 3 times the median deviation, was labelled as an outlier. i.e.

$$D_i > k \times M_d$$

Table 4.4 shows the number of melt pools, and their mean lengths [μm], across all speeds and exposure times, after removing the outliers.

Table 4.4

The number of melt pools, and their mean lengths [μm], for all speed-exposure combinations

	500 mm/s		600 mm/s		700 mm/s		800 mm/s		900 mm/s	
	No.	Length								
1ms	127	960	104	1086	89	1180	80	1271	67	1377
2ms	125	1574	103	1832	82	2073	66	2315	57	2564
3ms	103	2227	84	2617	71	2997	55	3378	48	3763
4ms	76	2901	63	3419	48	3935	44	4456	39	4992
5ms	60	3570	49	4212	42	4886	35	5552	30	6198

In the graph shown in Figure 4.11, the mean length of each data set is plotted, with error bars signifying the standard deviation of the lengths. With the outlier lengths removed, the standard deviation is so small that these error bars are almost invisible.

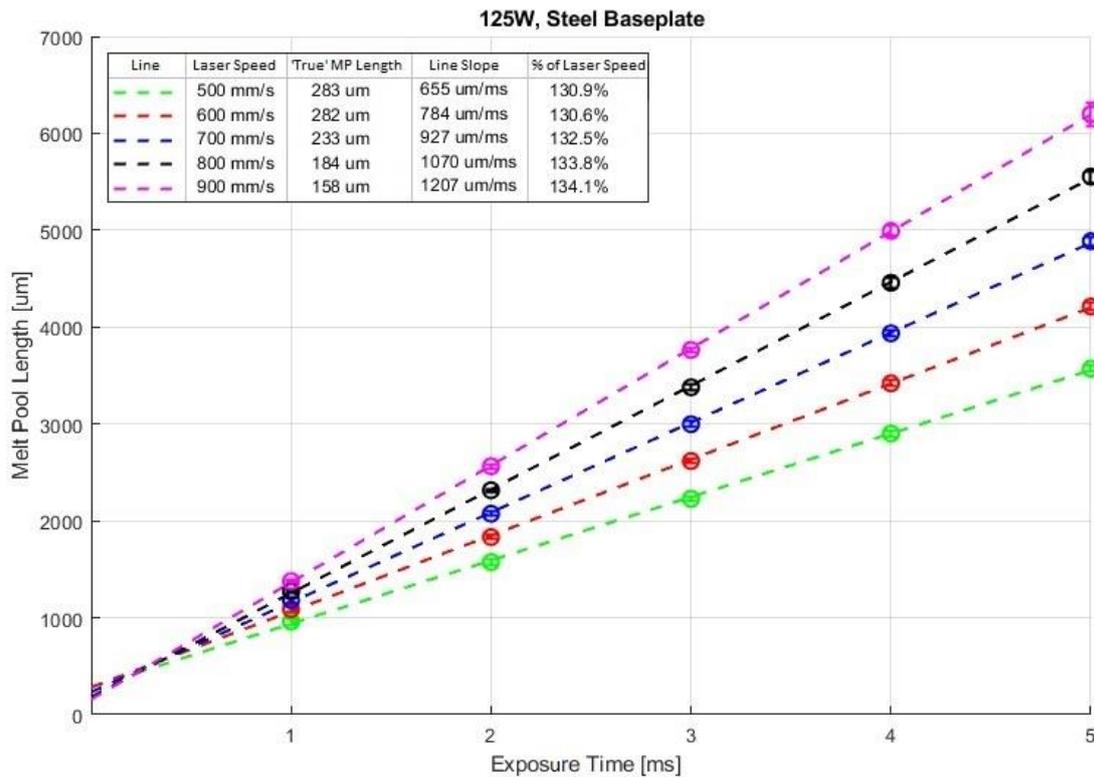


Figure 4.11

Extrapolating backwards, to find what melt pool length would be at 0ms exposure time

A linear line of best fit, calculated using the POLYFIT function, was found for each laser speed. The slope and $x=0$ intercept were extracted. According to the reasoning outlined previously, the intercept should equal the 'true' 0ms exposure time lengths. The line slope (i.e. the rate at which the melt pools extend, as exposure time increases) should equal the laser speed. These numbers are shown both in the chart legend, and in Table 4.5.

Table 4.5
The 'true' length of each speed, and its slope, and speed : slope ratio

Speed [mm/s]	'True' Length [μm]	Slope [$\mu\text{m}/\text{ms}$] (mm/s)	Slope/Speed [%]
500	283	655	130.9
600	282	784	130.6
700	233	927	132.5
800	184	1070	133.8
900	158	1207	134.1

As can be seen, the 'true' melt pool length follows the expected trend, where as the laser speed increases, and the linear energy density therefore decreases, the length decreases. The slope of the graph, however, is consistently around 30% larger than the laser speed. It is believed that this discrepancy can be explained as some combination of the following three sources of error: system error in the exposure time being set accurately; system error in the laser speed; and data analysis error in the removal of outliers.

It was reasoned previously that the extent to which the melt pool extends, while travelling at speed V and imaged at exposure time t , should simply be a product of these factors. i.e., the 'true' length can be determined from measured length by:

$$L_{true} = L_{meas} - V \times t$$

However, should the 130% slope identified be followed (instead of slope = 100%), the 'true' length would be determined as:

$$L_{true} = L_{meas} - 1.3 \times V \times t$$

The graph in Figure 4.12 shows an estimation of 'true' melt pool length, for each of the laser speeds, according to three different methods of calculation:

- According to the $x=0$ intercept, which is only applicable for this experiment, where tests were repeated at various exposure times.
- According to the simple logic with slope = 1, where $V \times t$ is removed from the lengths as measured at 1ms exposure.
- According to the slope identified experimentally of 1.3, , where $1.3 \times V \times t$ is removed from the lengths as measured at 1ms exposure.

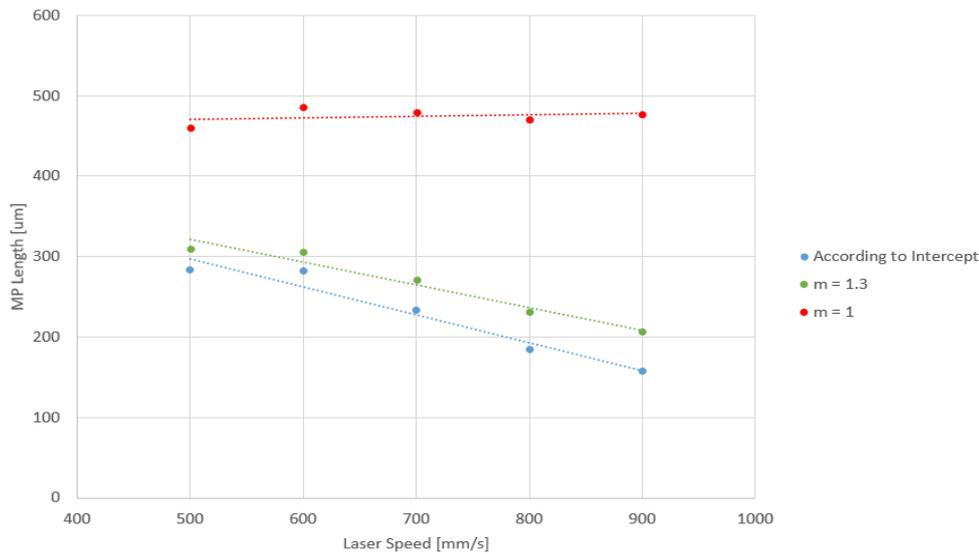


Figure 4.12

Comparing the three methods of estimating melt pool length

As expected, the melt pool lengths as determined by the intercept and $m=1.3$ methods are very similar. The lengths as determined by the $m=1$ method do not follow the expected trend of decreasing in size as linear energy density decreases. Therefore, a slope of 1.3 will be used for all melt pool lengths measured in this project.

In all instances from here onwards where melt pool lengths are discussed, it should be assumed unless otherwise stated that the value has been calculated according to:

$$L = L_{1ms} - 1.3 \times V \times 1ms$$

4.5 - Understanding the Effects of Camera Focus

In order to quantify the extent to which being out-of-focus affects the size, shape, and intensity readings of a melt pool, two data sets were analysed:

- The 0.1mm tall, single-line ‘dots’ that were recorded for the spatial accuracy tests previously, were used to provide a simple measure of the depth of field.
- A rectangle, composed of a set of horizontal lines, was scanned onto a powder-free base plate, for the extraction of more precise data.

4.5.1 - Depth of Field

Deciding which dots should classify as in-focus or out-of-focus, for each of the images, was a very subjective process. The maximum intensity value for each dot was extracted, and plotted for vertical columns. Figure 4.13, for example, shows an extract from the long-exposure image recorded with a 1024px frame, focused at the midpoint, with a column of dots highlighted, and the maximum intensities of the pixel rows plotted.

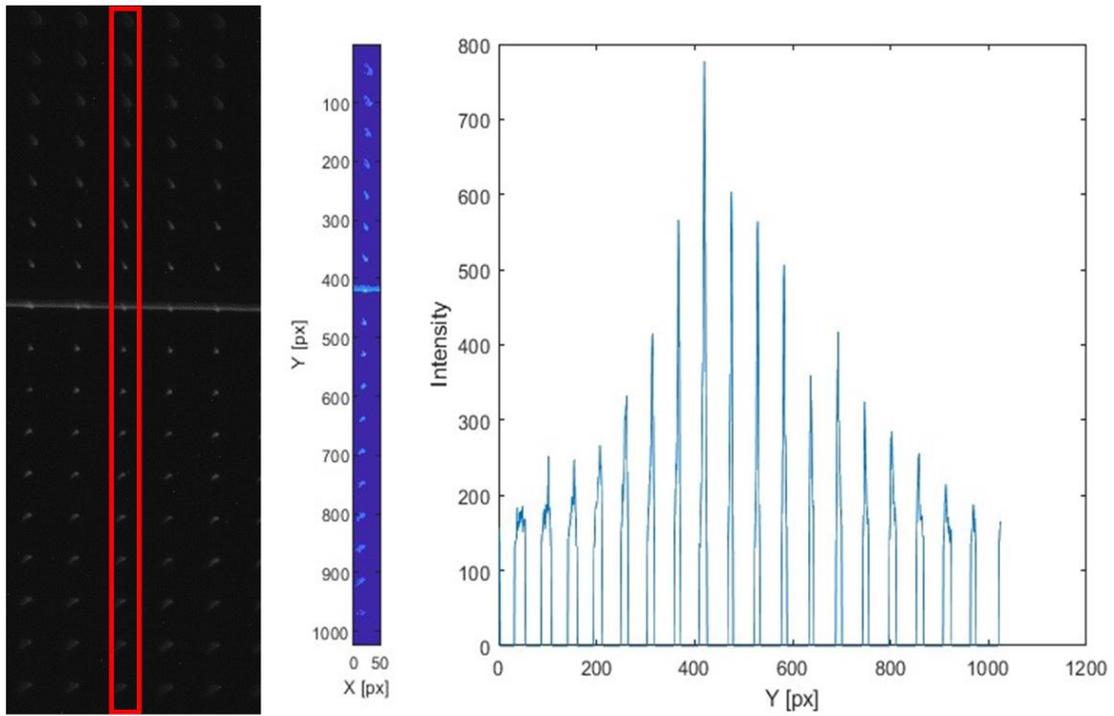


Figure 4.13
Examining the effects of focus on a vertical column of dots

Repeating this method for every column of dots in the image yields 34 such plots. Compiling the plots, and aligning them so that the dot with the greatest intensity is at $X=0$, results in the graph in Figure 4.14.

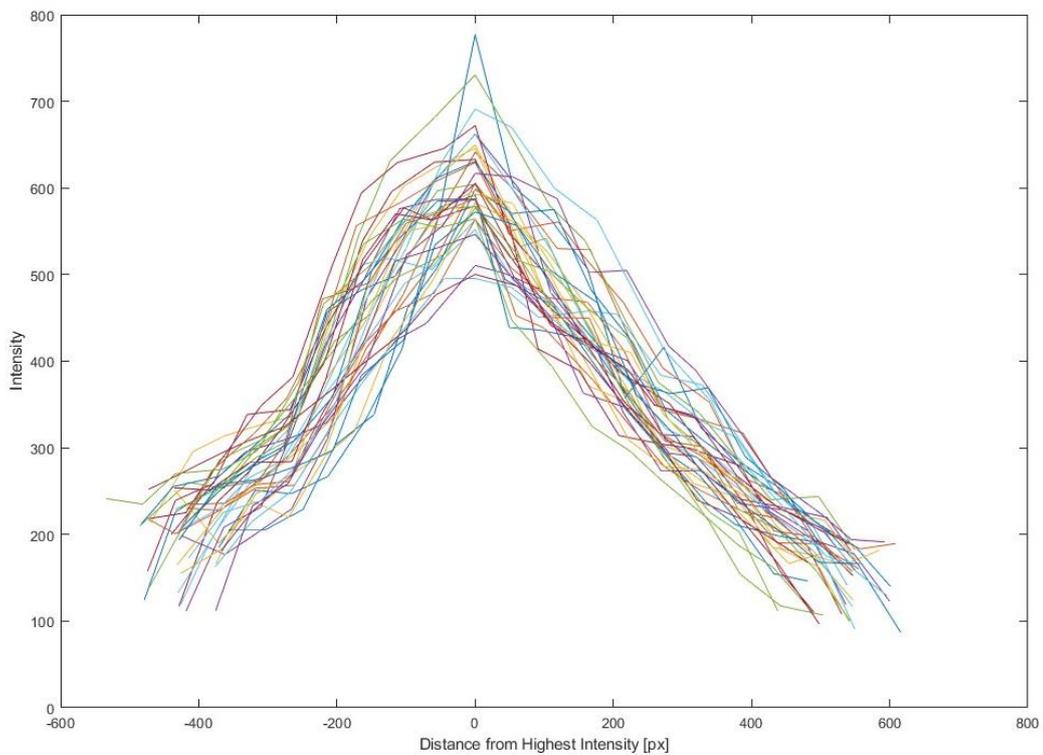


Figure 4.14
Comparing the effects of focus for all 34 columns of vertical dots

It was decided that a dot should be considered out-of-focus, once its peak intensity was less than 2/3 of the peak intensity of the most in-focus dot in its column. This method of determining focus was used rather than a feature sharpness measure, due to the inherent blurriness of all (even in-focus) melt pool edges. By finding the dot which satisfied this metric for each line, taking the pixel distance of this dot from the most in-focus, and calculating the mean of this distance, it was possible to determine a focus range for each frame. In the 1024px frame focused at the midpoint, according to this methodology, the focus range was found to extend 268px above (further into build chamber) the line of focus, and 243px below (closer to viewing window) it. Repeating this method for each image of dots, and converting from pixels to mm (using the previously mentioned conversion factor of 40mm/2048px), yields the ranges shown in Table 4.6.

*Table 4.6
Distance from each line of focus that melt pools can still be considered ‘in focus’*

Frame	Focus	Range Above		Range Below	
		[px]	[mm]	[px]	[mm]
2048	Upper	N/A	N/A	238	4.65
2048	Up-Mid	278	5.43	251	4.9
2048	Mid	263	5.14	249	4.86
2048	Low-Mid	259	5.06	252	4.92
2048	Lower	271	5.29	N/A	N/A
1024	Upper	N/A	N/A	235	4.59
1024	Mid	268	5.23	243	4.75
1024	Lower	276	5.39	N/A	N/A
512	Mid	256	5	243	4.75

Based on these numbers, a general guideline was established: it should not be expected for melt pools that lie further than 5mm above the camera’s line of focus, or further than 4.5mm below the line, to be in-focus.

4.5.2 - Melt Pool Size, Shape, and Temperature

In order to better determine the effect that being out-of-focus has on the size, shape, and temperature profile of a melt pool, horizontal single-line tracks were scanned across a baseplate, and recorded with the 1024px frame, focused on its midpoints. The scans were performed on powder-free 316L steel, at 125W, 500mm/s, with 0.5mm between them. The frame rate was 200fps, and the exposure time was 1ms. Unlike in all the tests discussed previously in this chapter, here the raw camera images were processed using the custom built MATLAB scripts to convert intensity readings to temperature. The method used to perform this conversion was identical to that written about in the previous chapter for IN-718, albeit with the emissivity value of 0.0980 used for 316L steel.

Figure 4.15 shows a superimposed image, composed of all of the melt. All pixels with a temperature reading below the 1,402°C solidus temperature of 316L steel reported in literature [319] were set to have a pixel value of zero. Then, all non-zero pixel regions with an area of < 50px were also set to zero, as any region smaller than this was clearly spatter or an optical anomaly.

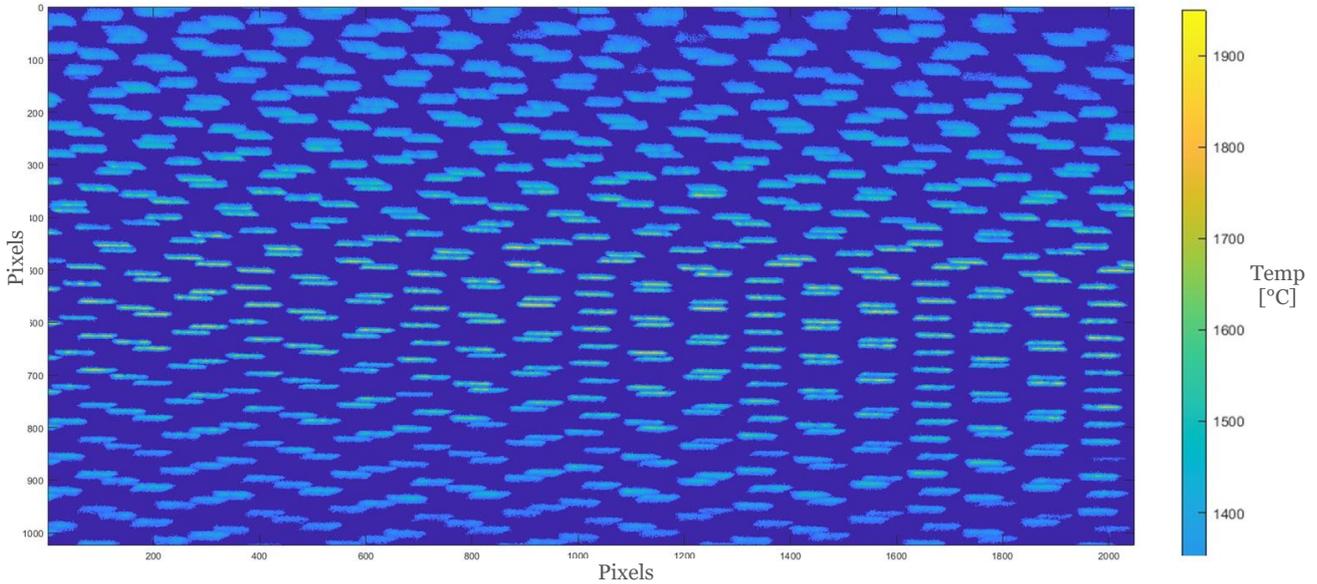


Figure 4.15

The effects of focus level on melt pool dimensions and temperatures

As is clear from Figure 4.15, as melt pools extend further from the line of focus along the frame's centre, their recorded temperatures decrease, and their size begins to swell. Each melt pool had the following data extracted from it:

- Centroid coordinates in X and Y
- Peak temperature
- Mean temperature of all pixels within the pool
- Width and length dimensions

Widths and lengths were found using the fitted ellipse method, which will be explained in more detail in Section 5.3.1.1. The maximum values were found for the peak and mean temperatures, and the minimum values were found for the widths and lengths. These four values were used for comparison for every other melt pool. The percentage difference between each melt pool and these values was then calculated and plotted.

Figure 4.16 shows the percentage difference (D) between the peak temperature of each melt pool (T_p), and the maximum of all these peak temperatures (T_{mp}), according to the formula:

$$D = \left(1 - \frac{T_p}{T_{mp}} \right) \times 100$$

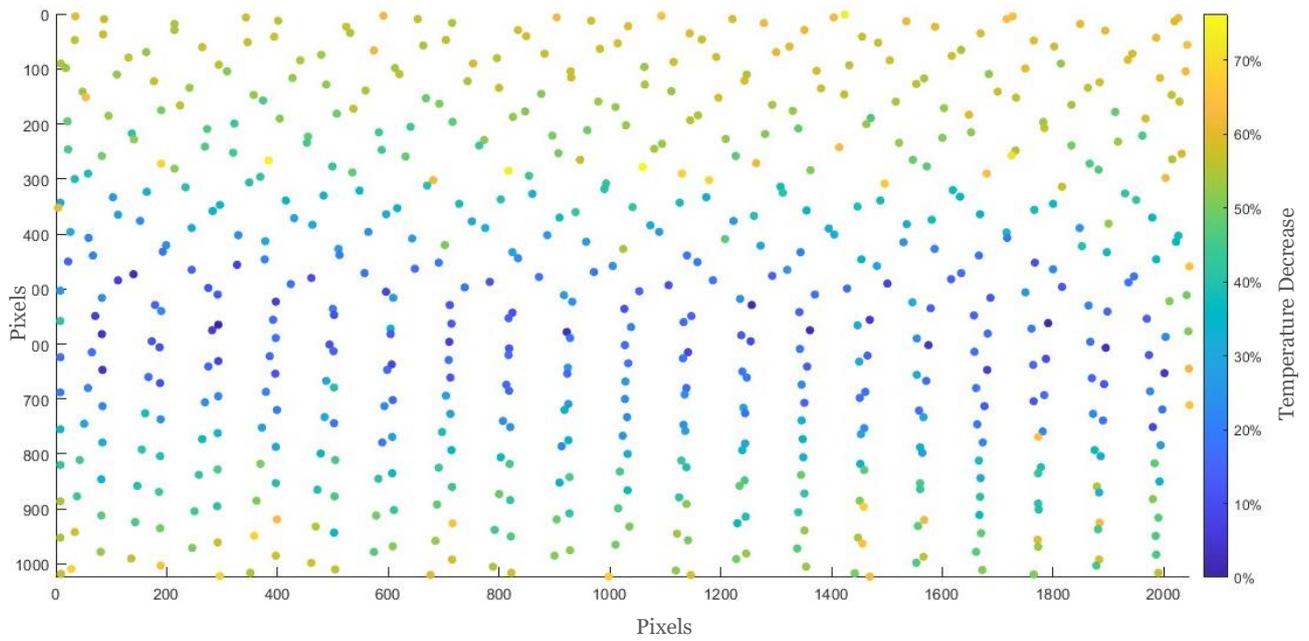


Figure 4.16

The percentage decrease of melt pool peak temperatures, as they move away from the focus line

Likewise, using a similar formula, Figure 4.17 shows the percentage difference between each melt pool's mean temperature, and the maximum of these means.

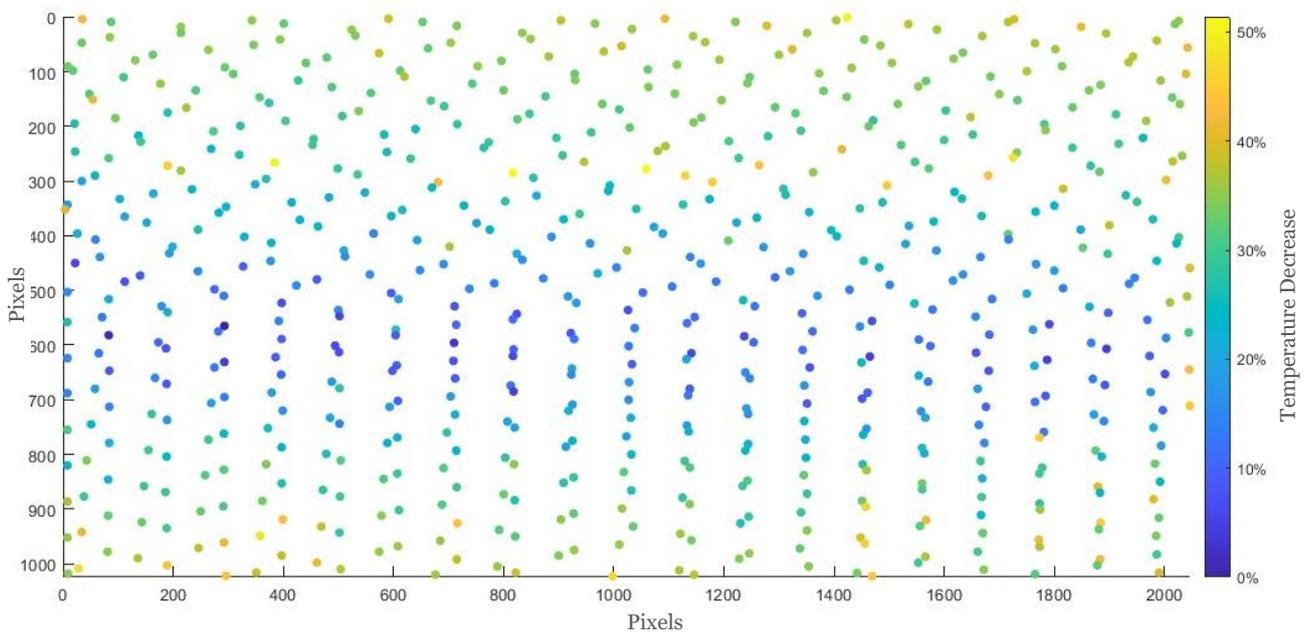


Figure 4.17

The percentage decrease of melt pool mean temperatures, as they move away from the focus line

Figure 4.18 shows the percentage difference (D) between the width of each melt pool (W_i), and the minimum of all these widths (W_m), according to the formula:

$$D = \left(\frac{W_i}{W_m} - 1 \right) \times 100$$

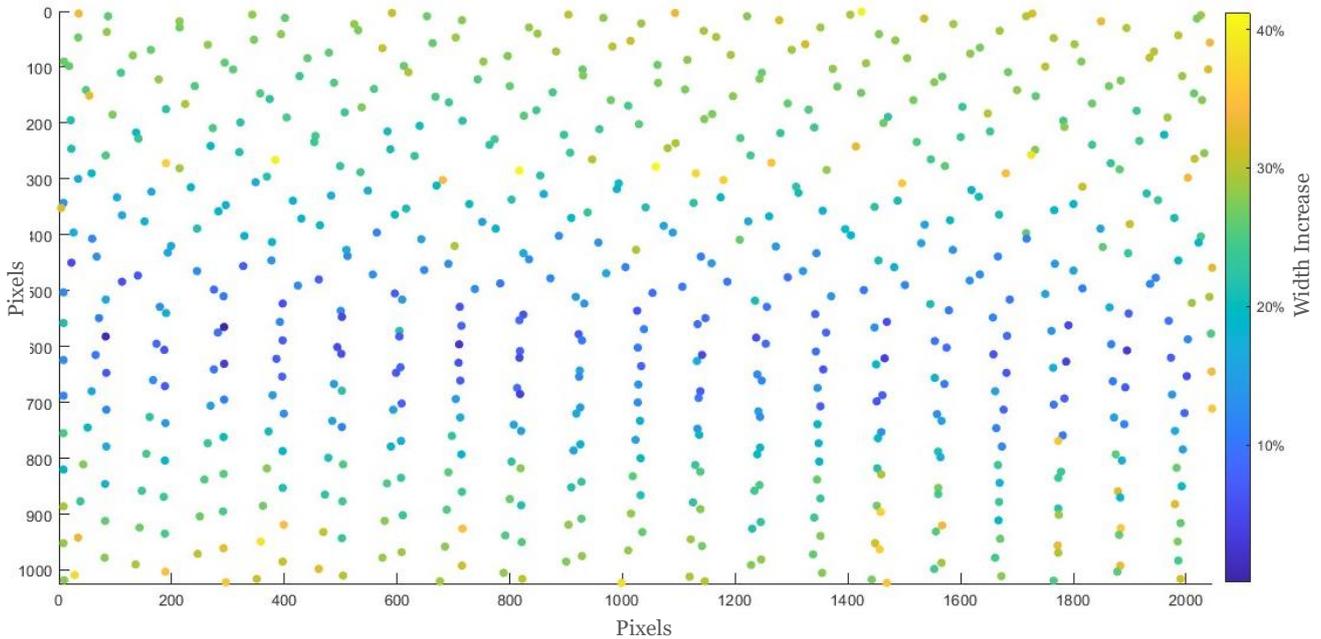


Figure 4.18

The percentage increase of melt pool widths, as they move away from the focus line

No clear trends were found to exist when a similar method was applied to the melt pool lengths. The extent to which the lengths of pools grew as their distance from the line of focus increased, was found to be less significant than the variation of lengths of pools at the line of focus. This follows the logic that as elongated objects blur when out of focus in a camera, their shorter dimension is going to be affected more than their longer dimension.

In these three figures, the line of focus is assumed to lie halfway along the approx. 20mm vertical height of the 1024px frame. The three figures were divided into 20 vertical slices of equal height, each approx. 1mm in height. The average of each percentage change was found in each slice. These averages are displayed in Figure 4.19.

Based on Figure 4.19, general guidelines were estimated, for the rate at which a melt pool's peak temperature, mean temperature, and width, would increase/decrease as the pool strayed from the line of focus. These estimated rates were: -6 % per mm for peak temperature; -5% per mm for mean temperature; and +5% per mm for pool width. From this point onwards, whenever melt pool peak temperatures, mean temperatures, or widths are addressed, this correction has been applied.

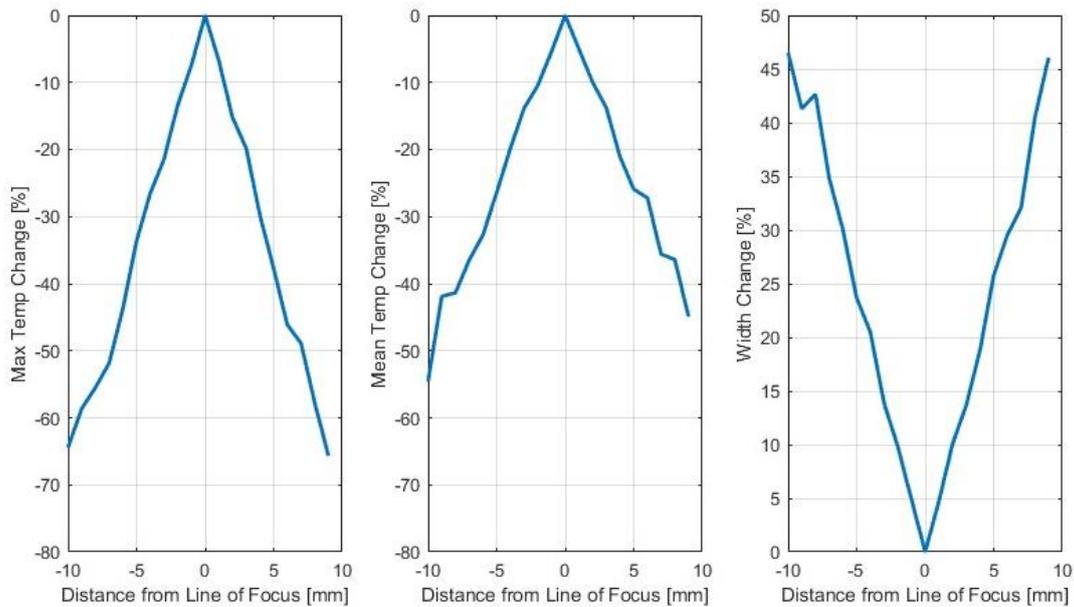


Figure 4.19

The average effects on melt pool peak temperatures, mean temperature, and widths, as they move away from the focus line

4.6 - Preliminary Evaluation of Camera

Four simple experiments were devised, to determine the extent to which melt pool characteristics as recorded by the camera, matched or differed from melt pool characteristics as expected from theory. The four hypotheses to be tested were:

1. The shape of the component being scanned, and the location of the vector within that shape, will affect the thermal characteristics of the melt pool.
2. The volume of adjacent, solid material, for heat to be conducted into, will affect a melt pool's thermal characteristics.
3. The duration of a pause, between one vector and the next, adjacent vector, will affect the thermal characteristics of the melt pool in the adjacent vector.
4. The distance of insulating powder, between a single-line track and its nearest bulk solid, will affect the thermal characteristics of its melt pool.

A build was carried out to test each of these hypotheses. All builds were performed with IN-718 powder, upon a 316L steel baseplate, with a laser power of 125W, a laser speed of 500 mm/s, a hatch spacing of 50 μ m, and a layer height of 30 μ m. The 512px frame was used, with a recording rate of 400fps (the maximum possible at the frame size), and an exposure time of 1ms. The focus was set to the midpoint of the frame, which is where all builds were centered.

The overall aim of these tests was not to determine the exact ways in which melt pool characteristics would be affected, but rather to demonstrate that the thermal camera was capable of detecting the changes to melt pool characteristics, caused by altering component geometry.

4.6.1 - Effect of Component Shape

Four shapes were built:

1. A wide cube, with a side length of 8mm
2. A narrow cube, with a side length of 4mm
3. A wide-based isosceles triangle, with a base of 8mm and a height of 10mm.
4. A narrow-based isosceles triangle, with a base of 4mm and a height of 10mm.

Figure 4.20 shows a top-down view of the four components, surrounded by the estimated perimeter of the 512px frame.

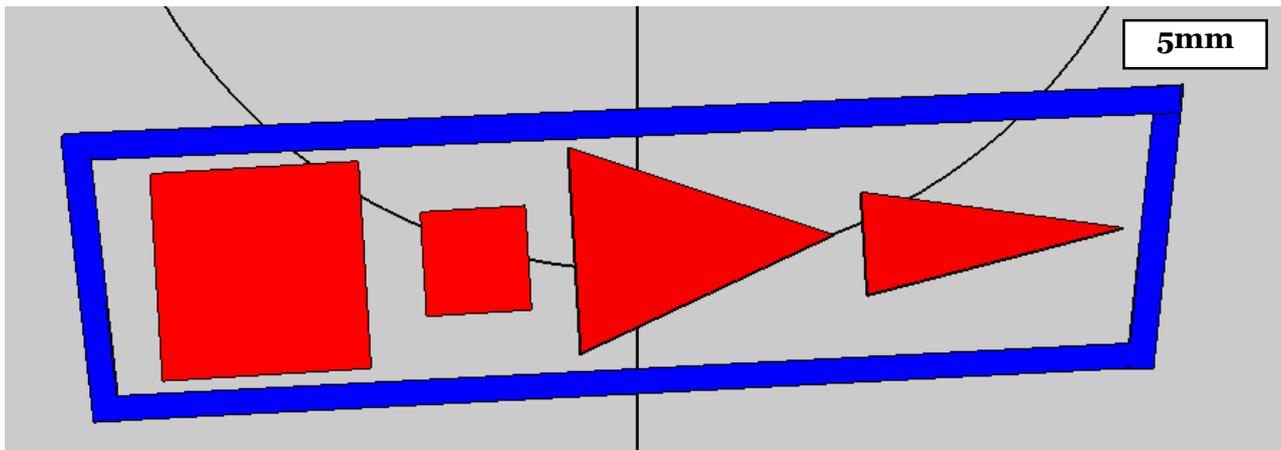


Figure 4.20

The four shapes that were built to test the camera

As can be seen, each component is rotated by 3° counterclockwise in order to align with the frame. The scanning strategy was also rotated by 3° , to align with the components. A simple, parallel scanning strategy was used, with a 90° rotation between layers. The layer that was recorded with the camera was the 100th (build height of 3mm), as it was believed that at this height the thermal conditions would have settled into a steady state. At this layer, the scan vectors were oriented vertically with respect to the components (i.e. 3° off the vertical in Figure 4.20), and progressed from left to right across the layer.

Figure 4.21 shows a superimposed image, with all the melt pools recorded. In each individual image, all pixels with a temperature below 1255°C are set to 0, and each pixel outside the melt pool was set to 0.

As can be seen, both the width and temperature of the melt pools increases as the scan vectors progress from regions with longer vectors and greater return times, to regions with shorter vectors and smaller return times. These findings follow expectations: shorter vectors lead to less time for conduction to dissipate heat energy, which leads to greater heat accumulation, which leads to elevated temperatures, which leads to more material within the melt pool being above the solidus temperature.

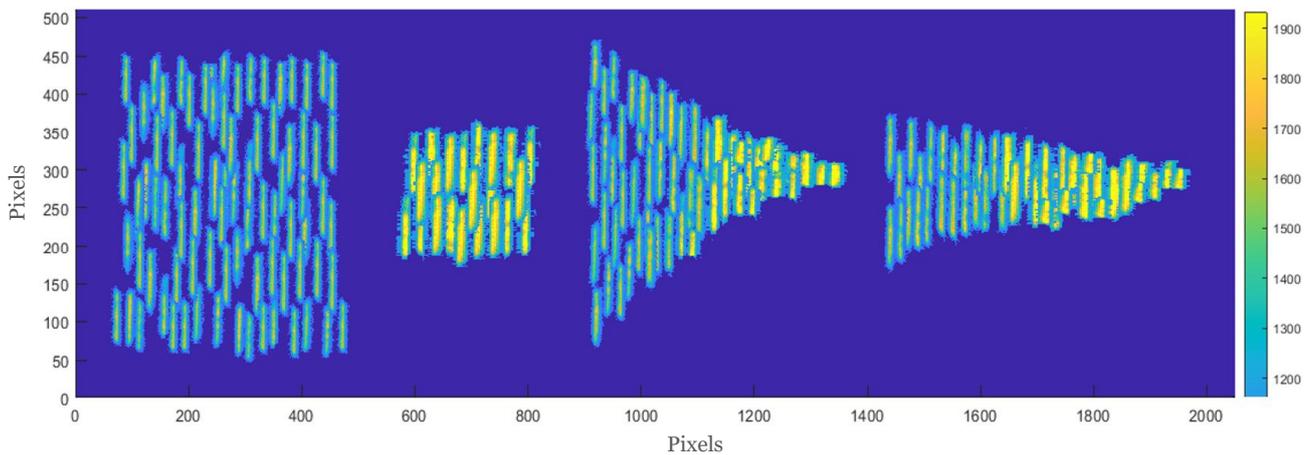


Figure 4.21

A superimposed image of melt pools from the four shapes

The maximum temperature, mean temperature, and width of every melt pool was extracted, and plotted against the length of its vector upon which it is estimated to lie, resulting in the set of graphs shown in Figure 4.23, Figure 4.24, and Figure 4.25.

The fact that the measured mean and maximum melt pool temperatures (and to a lesser extent the melt pool width) appear to show such a strong negative correlation with the vector length, demonstrates the ability of the thermal camera to differentiate the characteristics of melt pools from different geometries.

4.6.2 - Effects of Heat Sink Volume

A single isosceles triangle was built, with a base of 8mm and a height of 35mm. The triangle's centerline was aligned with the frame's line of focus, at a 3° rotation, as shown in Figure 4.22.

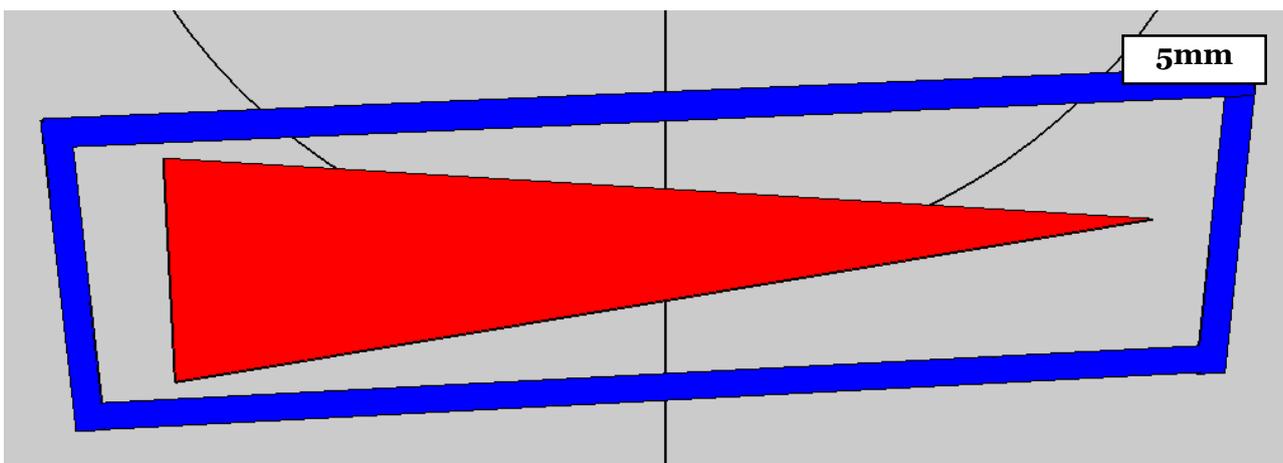


Figure 4.22

The triangle that was built to explore the effects of heat sink volume

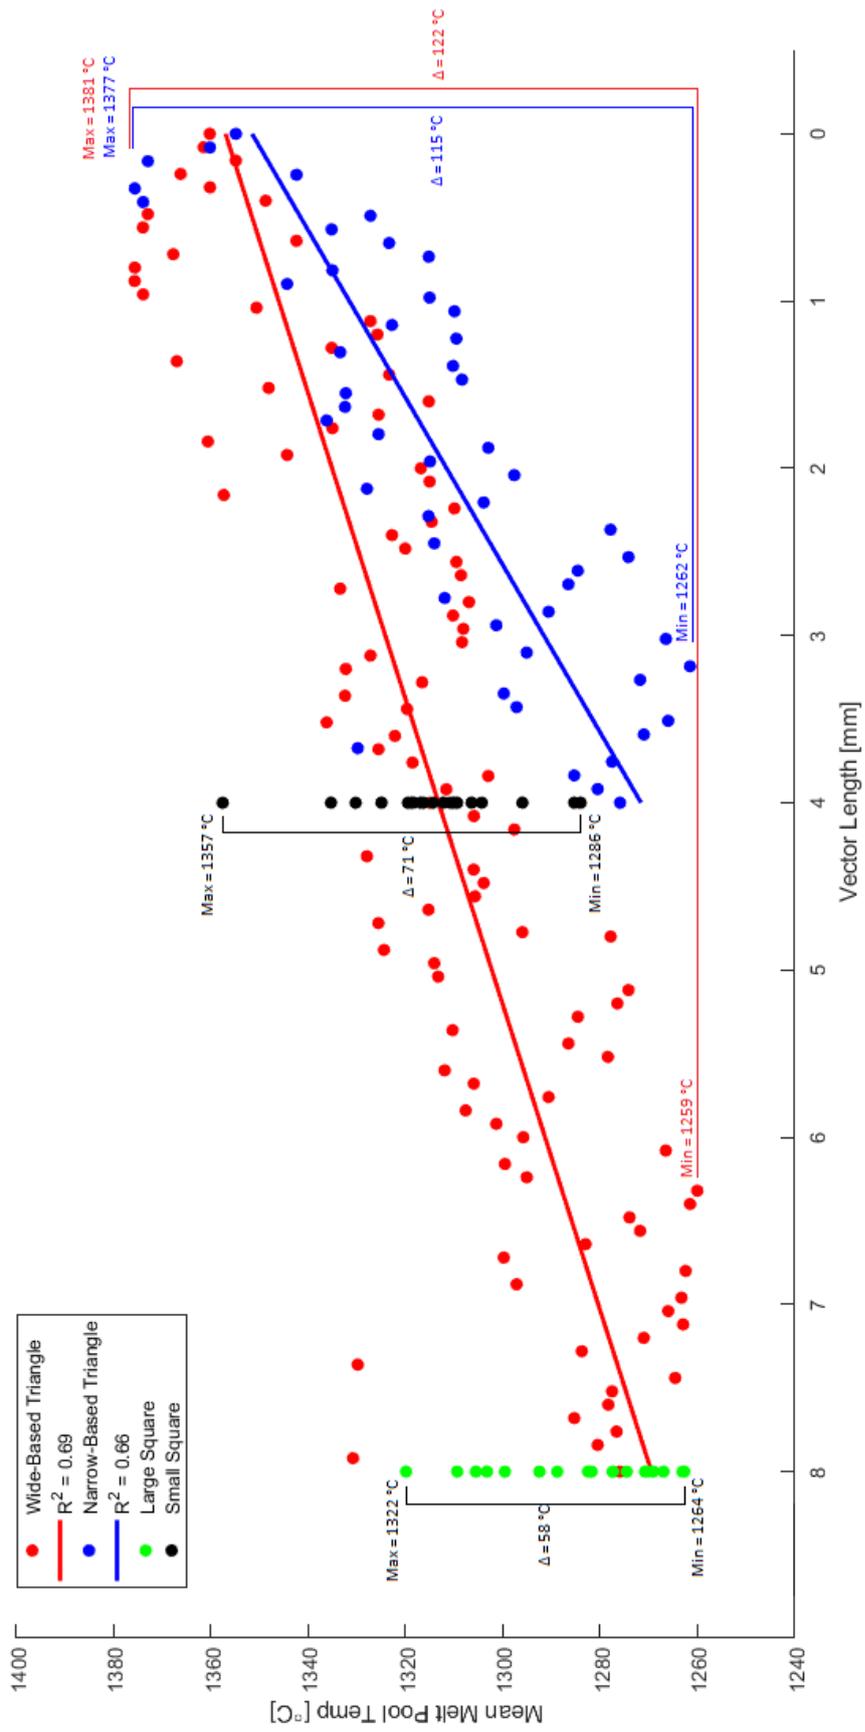


Figure 4.23
The effects of vector length on the mean melt pool temperature

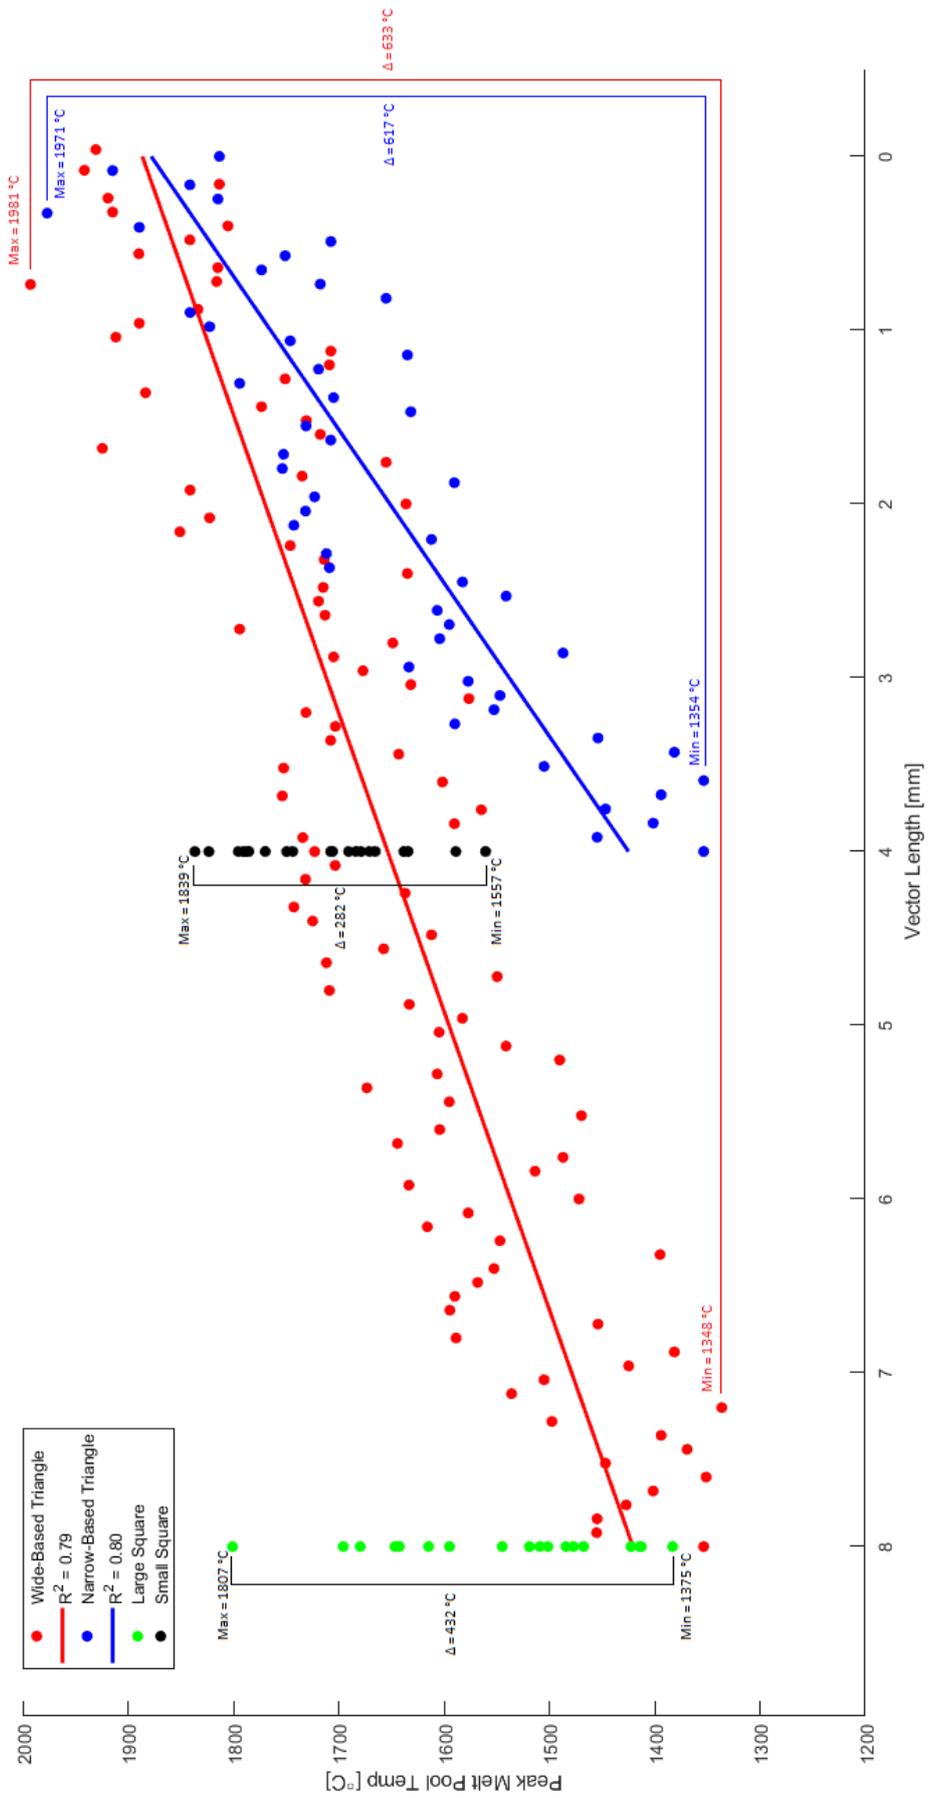


Figure 4.24
The effect of vector length on the peak melt pool temperature

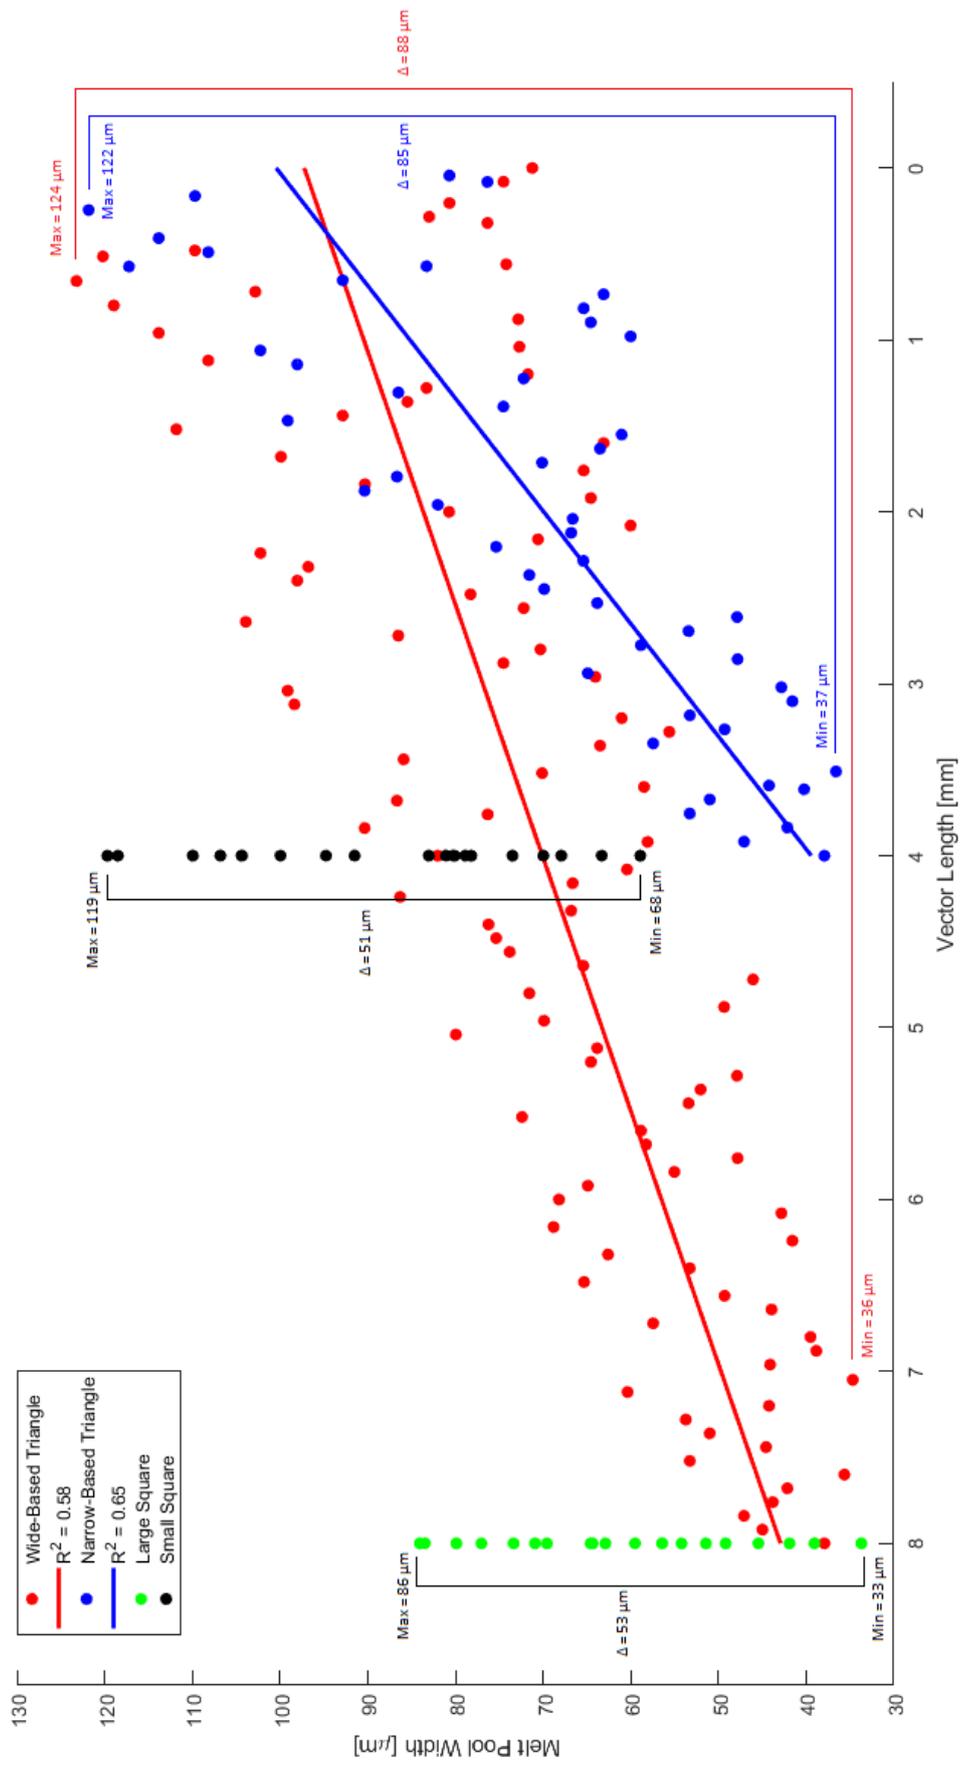


Figure 4.25
The effects of vector length on the melt pool widths

The triangle was built with 30µm layers, and a simple linear scanning strategy, with a 67° rotation between layers. Beginning at layer 100 (3mm build height), and repeating every 10 layers (0.3mm) until layer 190, the centerline of the triangle was lasered from left to right, before the rest of the shape was lasered. It was expected that, due to the larger volume of solid material to conduct away heat, melt pools recorded at the wide end of the triangle would be smaller and cooler than melt pools at the narrower end of the triangle.

Figure 4.26 shows a superimposed image with all the melt pools recorded along the centerline, at layer 150. All pixels either below 1255°C, or not part of a melt pool, have been set to 0.

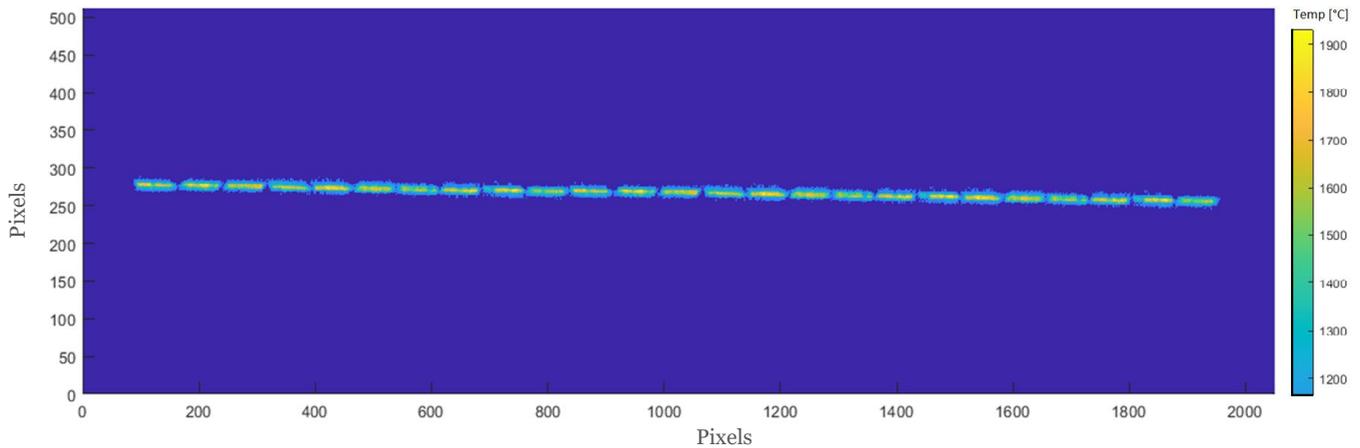


Figure 4.26

A superimposed image of all the melt pools from the triangle's centreline

The mean temperatures, maximum temperatures, and widths of the melt pools were extracted, alongside the width of the triangle (heat sink) at their locations. It was found that, when plotting the melt pool characteristics against the heat sink width, no identifiable trends emerged. Figure 4.27, for example, shows the mean melt pool temperatures are shown, for each of the 10 centreline scans recorded.

As can be seen, the temperatures vary stochastically between just above solidus, and approximately 1440°C. Plotting for melt pool peak temperature, and melt pool width, produce similar stochastic results.

It was reckoned that the cause for this lack of trend was threefold:

- The inherently stochastic nature of melt pool characteristics, especially when the expected trend may only be small in nature.
- The possibility that these differences in triangle width may not be enough to reasonably expect a change in thermal conditions for the melt pool, since the cooling length of the melt pool is sufficiently short to make all triangle geometries appear infinite in size
- The possibility that observing these subtle differences may be beyond the capabilities of the thermal camera.

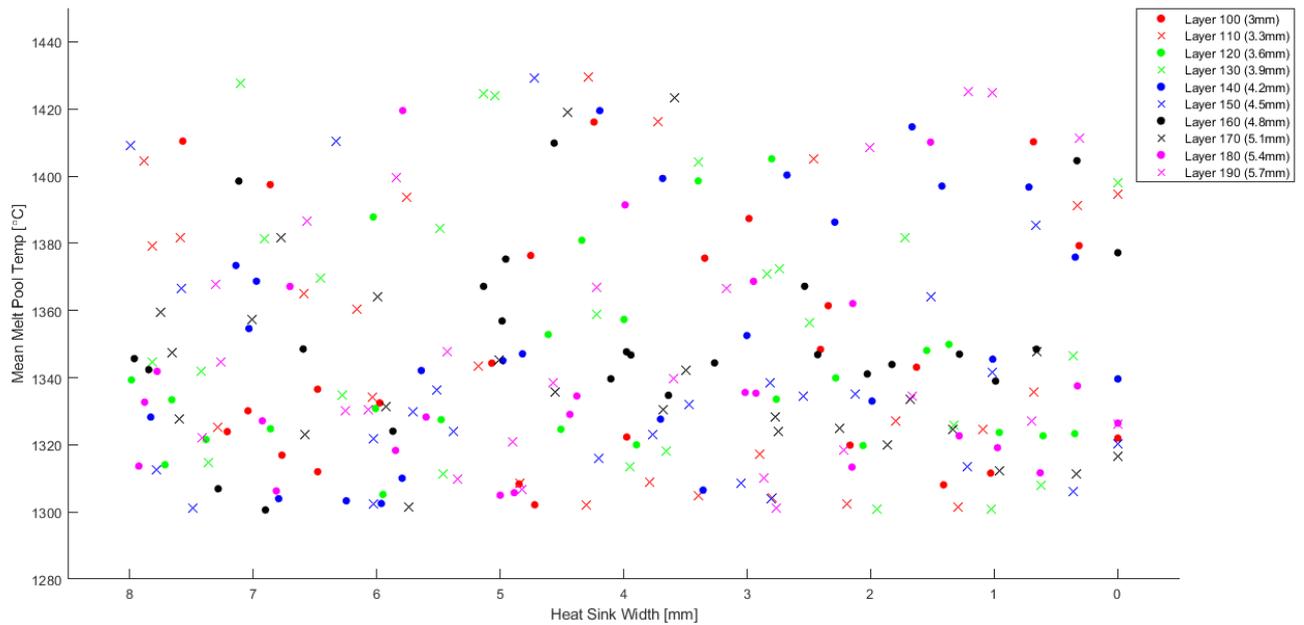


Figure 4.27

All melt pool mean temperatures, plotted against the width of the triangle they were built upon, showing that no trends exist between these parameters

4.6.3 - Effects of Pause Between Vectors

Five cuboids were built, each with widths of 5mm, and lengths of 1mm, 2mm, 4mm, 8mm, and 16mm. The parts were rotated by 3°, and positioned so that their bottom edges lay along the line of focus of the 512px frame (as this is the region of interest, as per the explanation below). Figure 4.28 shows the frame and the five cuboids, as well as a very thin, out-of-frame rectangle, the meaning of which will be explained in this section.

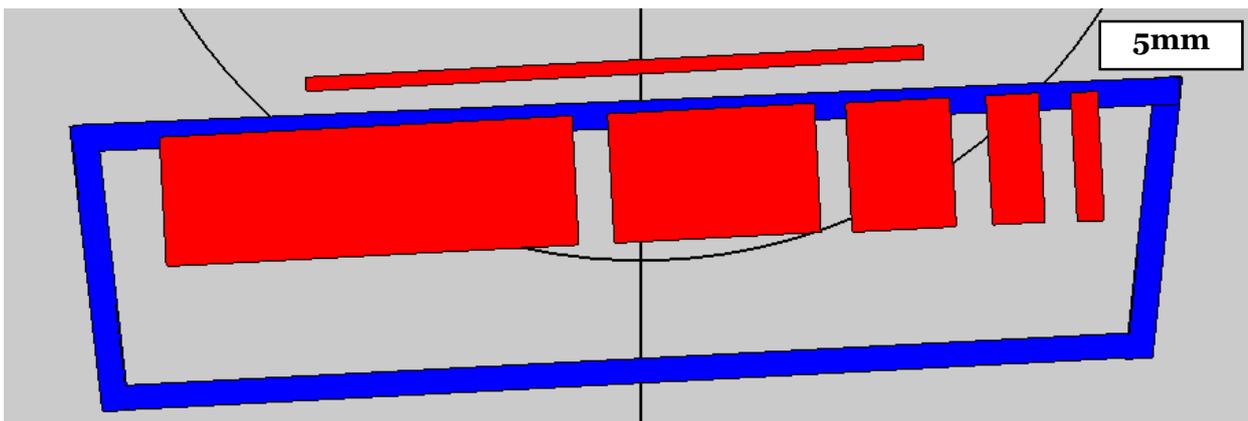


Figure 4.28

The cuboid and pause-rectangle design that was used to test the effects of pauses

A simple hatching scan strategy was used, with vectors running either vertically or horizontally across the parts (i.e. at either 3° or 93° relative to the build chamber), with a rotation of 90° between

layers. Beginning at layer 100 (3mm build height), and repeating every 10 layers (0.3mm) until layer 190, the following occurred for each part, progressing from left to right:

1. Horizontal scanning began, alternating left and right as the laser progressed down the part.
2. All but the last (bottom-most horizontal) vector was scanned.
3. The laser moved out of frame, to the very thin, 20mm long rectangle, which it scanned horizontally at 200 mm/s (i.e. a scanning time of 0.1s)
4. Depending on the layer, the laser may have re-scanned the long rectangle multiple more times, alternating left and right, running over the same vector again and again.
5. The laser returned to finish the final, bottom-most horizontal vector on the cuboid.
6. The laser progressed to the next part to repeat these steps.

At layer 100, the laser scanned the thin rectangle once (0.1s pause), before finishing off each part. At layer 110, the rectangle was scanned twice (0.2s pause). This pattern continues, with the rectangle scanned 10 times (1s pause) at layer 190. These layers that incorporated the pauses were recorded with the camera. Due to the time taken for the galvanometers to position from the parts to the thin rectangle, and back from the thin rectangle to the parts, the pause was slightly longer than prescribed in all cases.

Figure 4.29, for example, shows a superimposed image, composed of all the ‘final-line’ melt pools recorded at layer 150. All melt pools that were not part of the final-lines, have been omitted. In each individual image that makes up the superimposed image, all pixels were set to zero that were either below 1255°C, or outside of the melt pool.

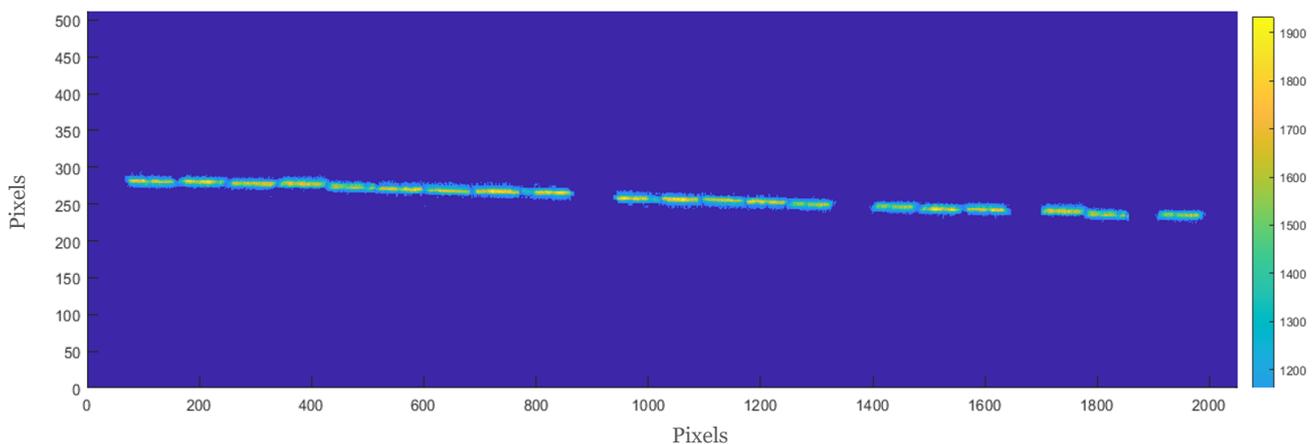


Figure 4.29
The final-line melt pools recorded at layer 150

The mean and maximum melt pool temperatures, and the melt pool widths, were extracted for each melt pool, in each of the five vectors, in each of the ten layers. It was expected that as the pause time increases, the widths and mean and maximum temperatures of the melt pool would decrease, due to the same mechanisms that caused these metrics to decrease with increasing laser return time.

These three metrics are shown in Figure 4.30, Figure 4.31, and Figure 4.32, for each of the five cuboids, and each of the five pause times. It is understood that it would have been beneficial to perform a control with 0s pause time, and that it was an experimental oversight to omit this.

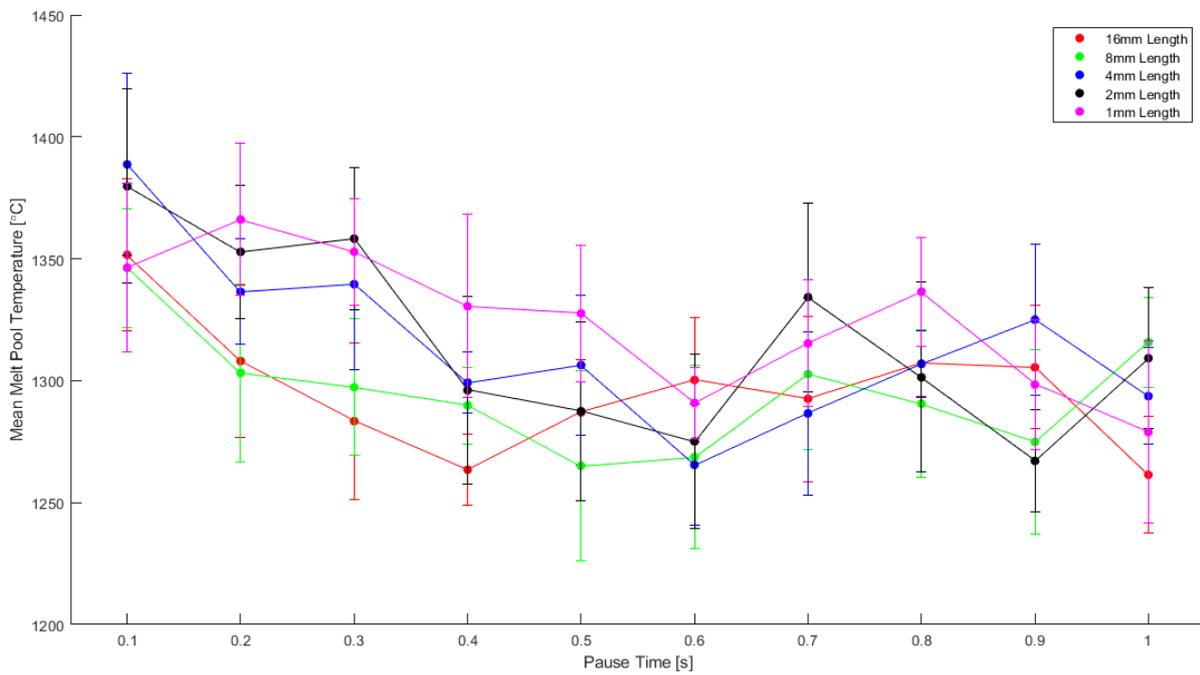


Figure 4.30

The mean melt pool temperature of various final-line lengths, as a function of the pause that came immediately before that final-line was scanned

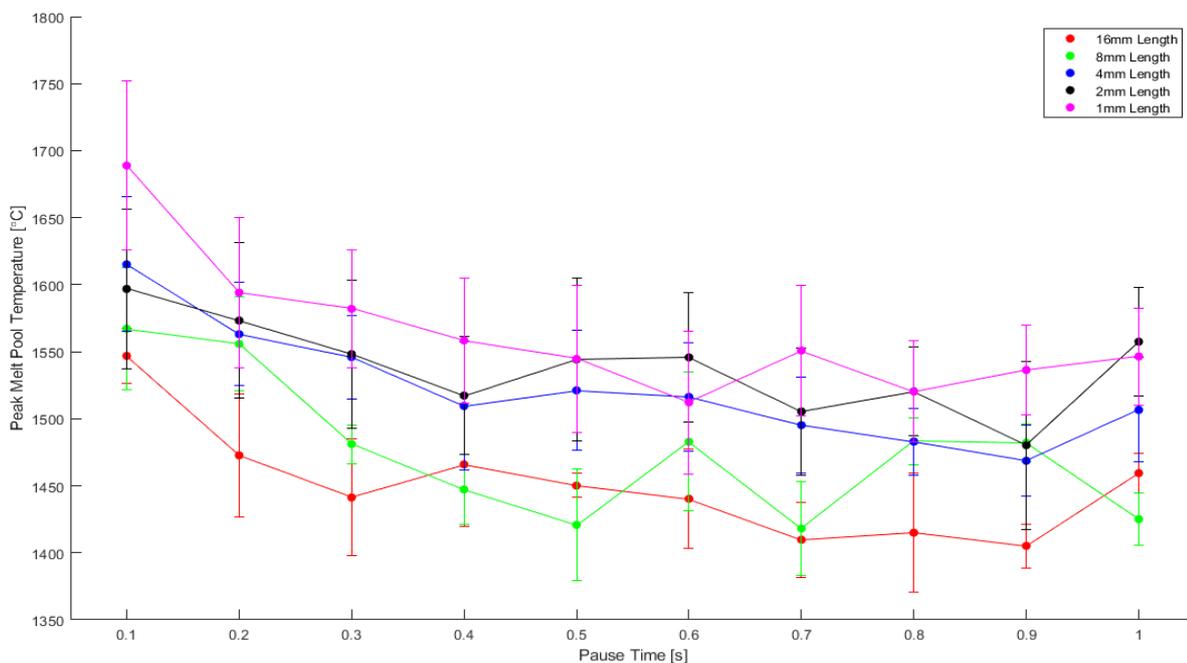


Figure 4.31

The peak melt pool temperature of various final-line lengths, as a function of the pause that came immediately before that final-line was scanned

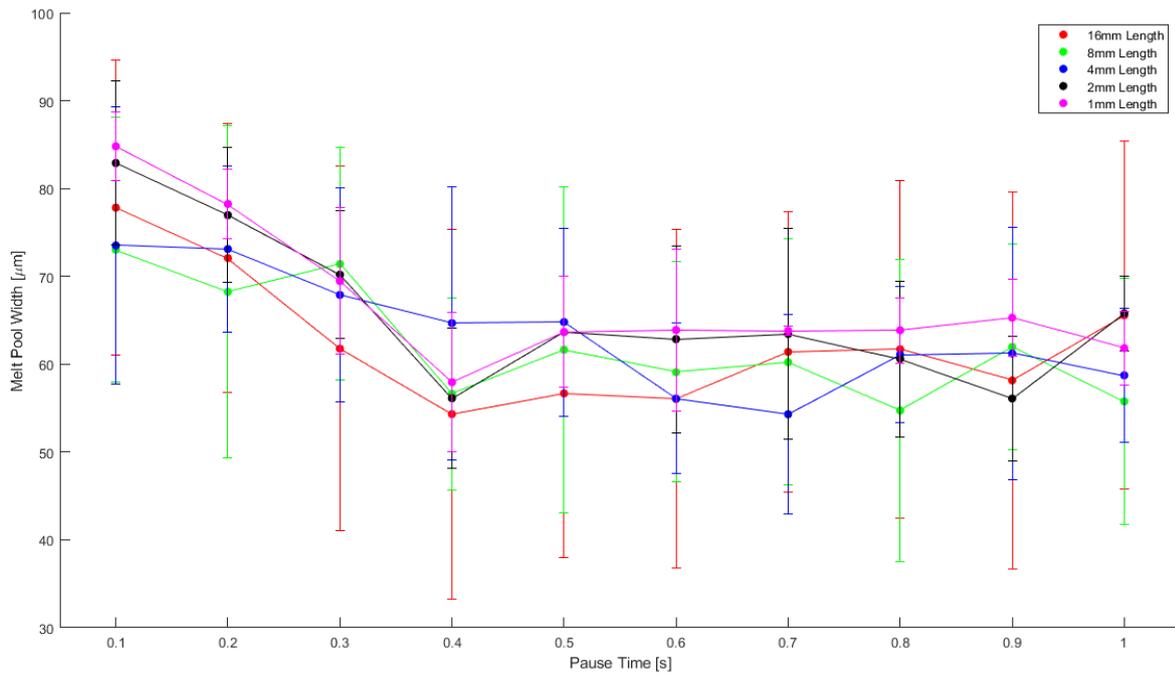


Figure 4.32

The width of various final-line lengths, as a function of the pause that came immediately before that final-line was scanned

In each of these three images, melt pool characteristics are shown to vary coherently as expected, for pause times between 0.1s to 0.4s. After this pause duration, the characteristics settle into a steady state, where any trends that may exist cannot be picked up by the camera.

4.6.4 - Effects of Insulating Powder

To determine if the camera could pick up on subtle differences in melt pool characteristics, caused by varying widths of insulating powder between the scan vector and the nearest solid materials, a simple test was designed. A cuboid with length 35mm and width 3mm was built, just above the line of focus of the frame, as shown in Figure 4.33.

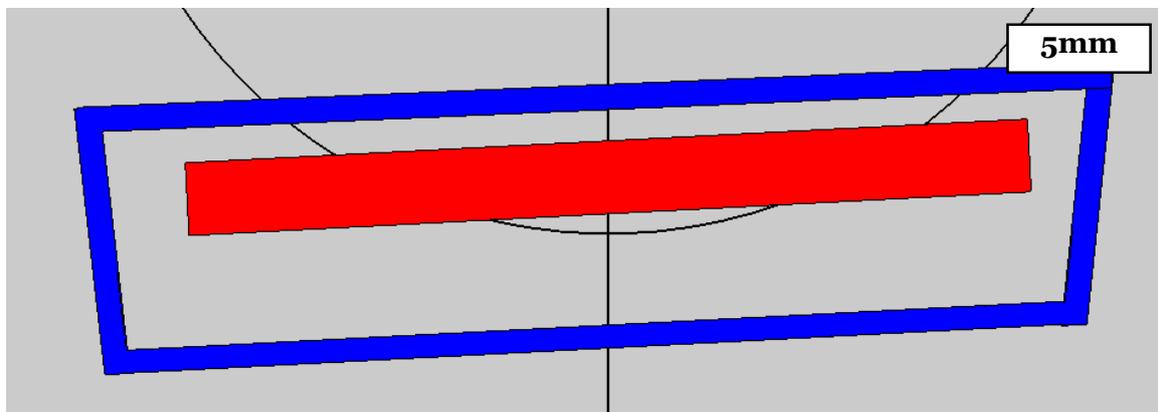


Figure 4.33

The cuboid that was built to test the effects of insulating powder

A simple scanning strategy, with vectors alternating between horizontal and vertical each layer, was used to build the cuboid. The strategy was designed so that on every 10th layer, the final vector was the lower-most horizontal, scanned from left to right. This vector will be referred to as the ‘control’.

On layer 100 (3mm build height), the control was immediately followed by a left-to-right, horizontal ‘test’ vector, scanned directly below it, with a hatch spacing of 50µm being maintained (i.e. the vector touched the bottom of the cuboid laterally, but was built onto loose powder). On layer 110 (3.3mm build height), the control was immediately followed by a left-to-right, horizontal test vector, scanned below it with a hatch spacing of 100µm (i.e. a single hatch was skipped). On layer 120, a hatch spacing of 150µm was used, skipping two hatches. This pattern, skipping zero hatches, then one hatch, then two hatches, was repeated at layers 130, 140, and 150, respectively, and again at layers 160, 170, and 180, respectively. Figure 4.34 shows the first five control-test sets (layers 100 to 140), from the XY perspective.

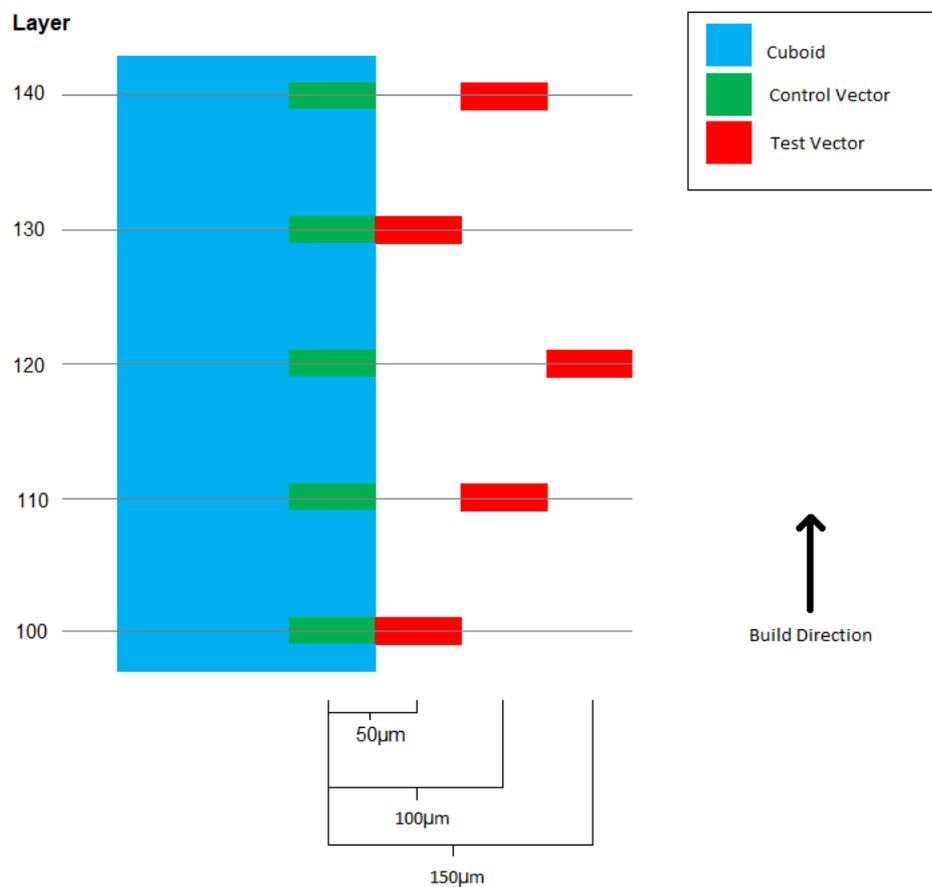


Figure 4.34

An XZ view of the cuboid (blue), showing the control lines (green), the test lines (red), and the powder in between them

The entire layer was recorded with the camera at each of the aforementioned layers, and the control and test vectors were extracted. In the superimposed image in Figure 4.35, the two vectors are shown for layer 120, where there is a 150µm spacing between them. In each individual image composing the superimposed image, all pixels either below 1255°C, or not in the melt pool, are set to zero.

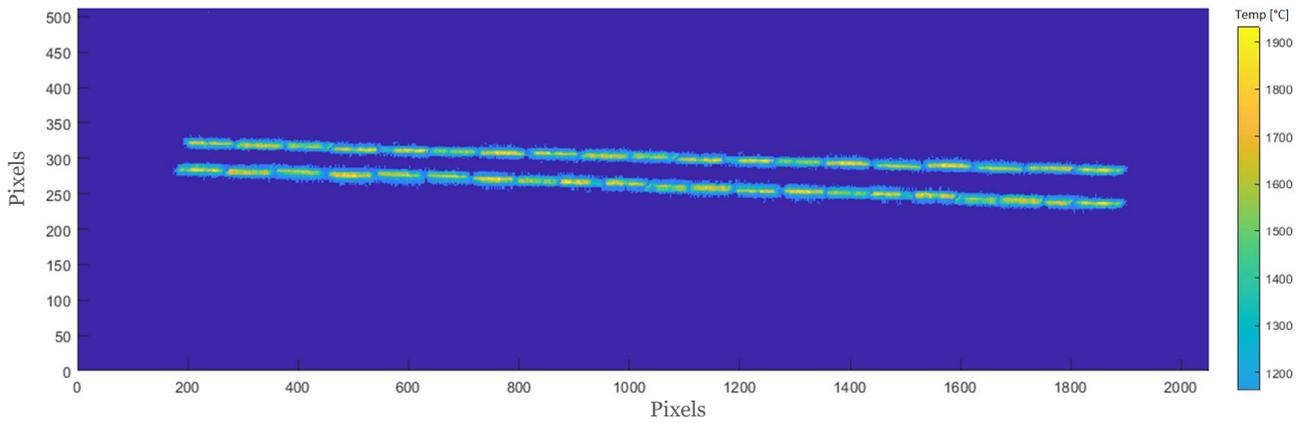


Figure 4.35

The control vector (top) and test vector (bottom) for layer 120

The mean temperature, maximum temperature, and width, was found for each melt pool in both the control and the test vector, for each layer being analysed. It was expected that as the hatch spacing increased, so would the temperatures and widths, due to a reduction of available solid material for heat to be conducted into, despite the decreasing preheat temperature associated with greater hatch spacing. It was also reasoned that the increase would level off after the 100µm vector, as from this vector onwards the melt pool would be completely surrounded by loose powder. When the mean and standard deviation of each melt pool characteristic was determined for each layer, it was found that no trends existed as the layers increased: i.e. all nine of the control vectors could be grouped together, and all three of the 50µm test vectors could be grouped, as could all three of the 100µm test vectors, as could all three of the 150µm test vectors. Therefore, all eighteen vectors can be analysed as only four sets of means and standard deviations, as shown in Figure 4.36.

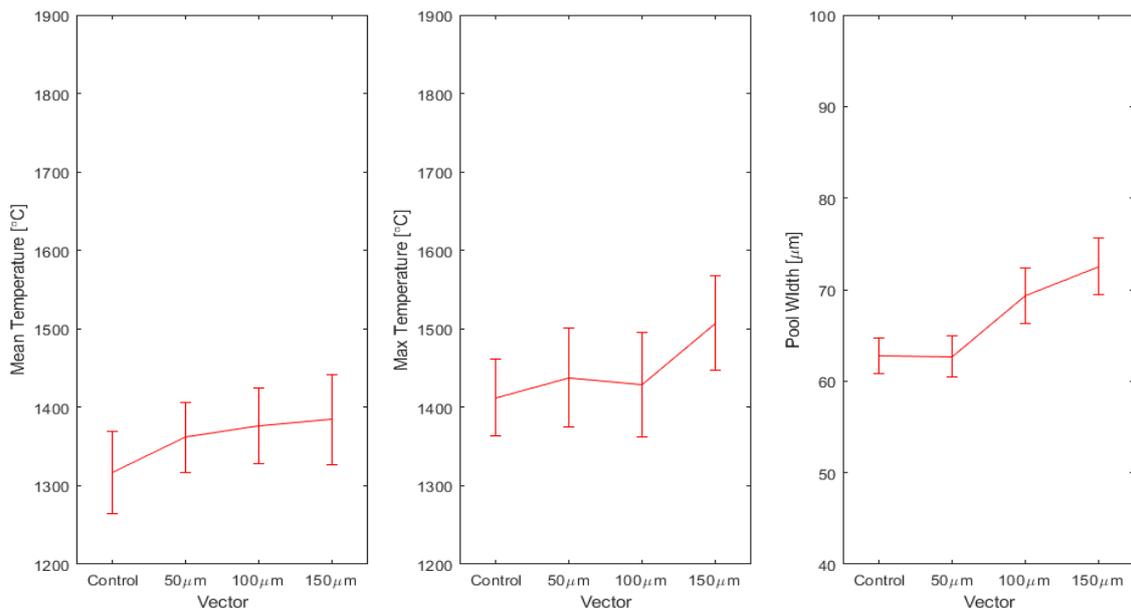


Figure 4.36

The effects of various distances of insulating powder, on mean melt pool temperature, peak melt pool temperature, and melt pool width

As can be seen, very slight increases, mostly within the margin of error, are visible across all three metrics, as was expected. These results further demonstrate both the strengths and limitations of the thermal camera setup: melt pool characteristics can be detected, although large data sets are required, due to the stochastic nature of L-PBF and the large errors associated with any reading.

4.7 - Accounting for Limitations of Emissivity

As discussed in previous chapters, there exists a circular dependency in the use of single-wavelength IR cameras. To convert a pixel's intensity to temperature, one must know the emissivity of the surface the pixel is recording; but also, a material's emissivity changes with temperature. Additionally, viewing angle, surface condition, and laser parameters can affect the value that should be used in equations.

As a result, single-wavelength cameras may be considered accurate in determining temperature trends, especially when images of identical experimental setups and materials are being compared. In terms of determining exact temperature values, however, their accuracy is questionable. This can be problematic in scenarios where an exact measurement of melt pool area is required.

4.7.1 - Work by Prior Researchers

The most common solution to this problem is the use of dual-wavelength camera setups, from which temperature may be calculated from the intensity ratio of the two recordings. A simpler solution that addresses the question of accurate melt pool area (although not the larger problem of accurate temperatures across the entire profile), was proposed by Cheng et al. [320] and Heigel et al. [321].

These researchers obtained the intensity profiles from melt pool images, extracted intensity-distance cooling curves, and analysed these to find the accelerations and decelerations associated with the liquidus and solidus boundaries of the mushy zone. By identifying the intensity readings that corresponded to the solidus and liquidus temperatures, it became possible to draw boundaries around all melt pools, even if the temperature profiles of all that lay inside and outside of these contour lines was unknown.

Heigel et al. first proposed the method of searching for the mushy zone in thermal images of melt pools. The authors recorded a build of IN-625, using an off-axis, high-speed camera, filtered for wavelengths between 1350 and 1600nm. The camera recorded a 13mm × 7mm section of baseplate, at a resolution of 360 pixels × 128 pixels, at 1800 frames per second. Seven power-speed combinations were tested. Figure 4.37 shows some example images of melt pools recorded for each case. Note that an intensity to temperature conversion has already been applied to these images, using a single emissivity value.

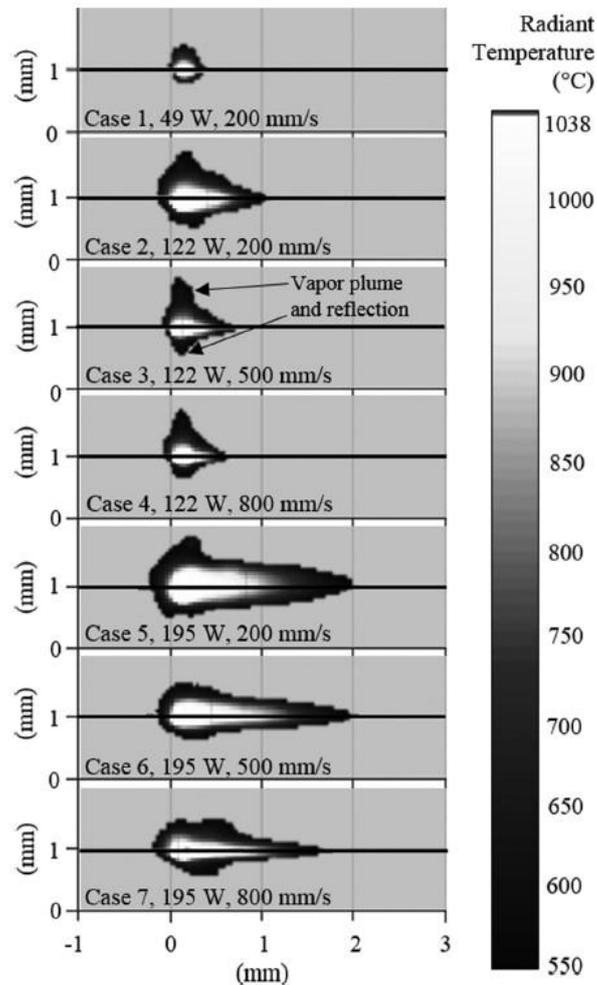


Figure 4.37
Melt pools recorded by Heigel [321]

The temperature profile of the centerline that has been marked in each of these melt pools, was then extracted for analysis. The cooling curve, trailing the peak melt pool temperature, was isolated. By searching for peaks and troughs in the second derivative of this cooling curve (i.e. the points at which the temperature's fall accelerated and decelerated), the mushy zone boundaries were found. Figure 4.38 shows the acceleration (i.e. liquidus temperature), for the melt pool of Case 7 shown in Figure 4.37.

By repeating this procedure for all recorded melt pools, the authors managed to identify the temperatures according to the single-value intensity conversion, that corresponded to the true liquidus temperature of IN-625 (1350°C). Figure 4.39 shows a table showing the laser parameters used in each case, the mean temperature corresponding to 1350°C, and the mean melt pool length that was found using that temperature as the start and end of the liquid pool.

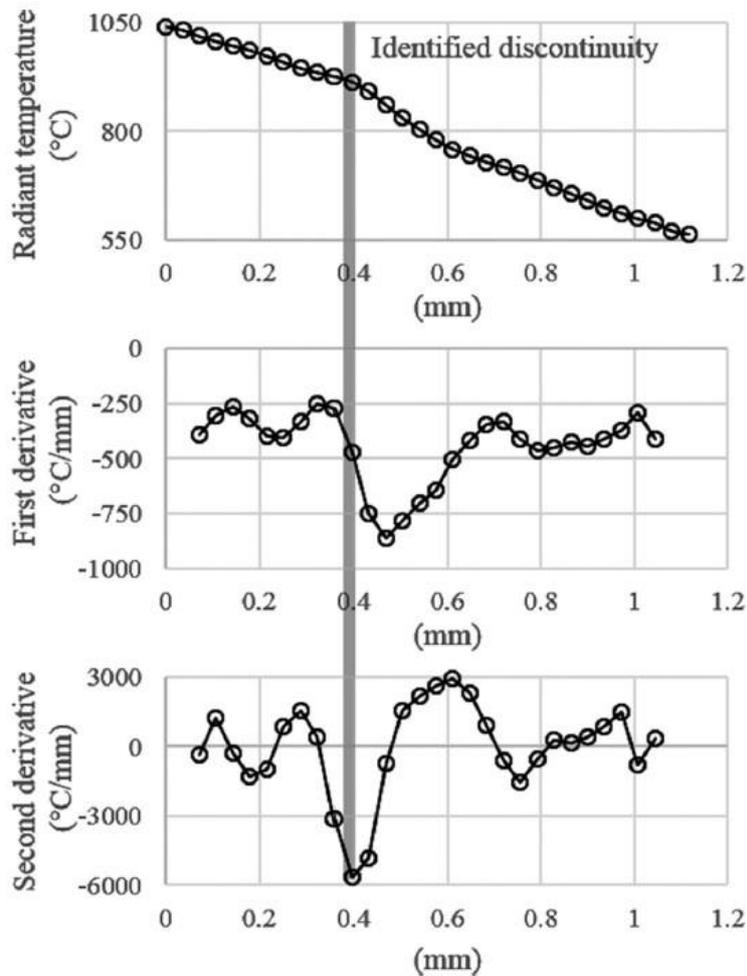


Figure 4.38

Heigel's method of determining liquidus temperature, based on the cooling curve's second-derivative [321]

Case	Power (W)	Speed (mm/s)	Radiant temperature used for detection (°C)	Melt pool length (μm)
1	49	200	942 ± 25	171 ± 16
2	122	200	942 ± 25	519 ± 29
3	122	500	942 ± 25	361 ± 27
4	122	800	942 ± 25	315 ± 27
5	195	200	961 ± 33	824 ± 109
6	195	500	949 ± 26	903 ± 102
7	195	800	936 ± 22	813 ± 79

Figure 4.39

Heigel's results [321]

As is evident from the table, not only is there significant error in the 'radiant temperature', but there is little change in this error, regardless of the laser parameters applied.

Cheng et al. advanced this work, by testing the effects of laser speed on emissivity values in more detail. An off-axis, high-speed camera was used to monitor an IN-718 (liquidus temperature 1336°C) build, recording a $30\text{mm} \times 5\text{mm}$ section of the base plate, with a resolution of $1410 \text{ pixels} \times 275 \text{ pixels}$. The camera was filtered to 670nm , and recorded at 58Hz .

The authors identified more error in their tests than Heigel did. Using the second derivative method, the radiant temperatures of melt pools over a range of build heights. Each set of data points, for each build height in Figure 4.40, is composed of data from tests at 400, 600, and 800 mm/s:

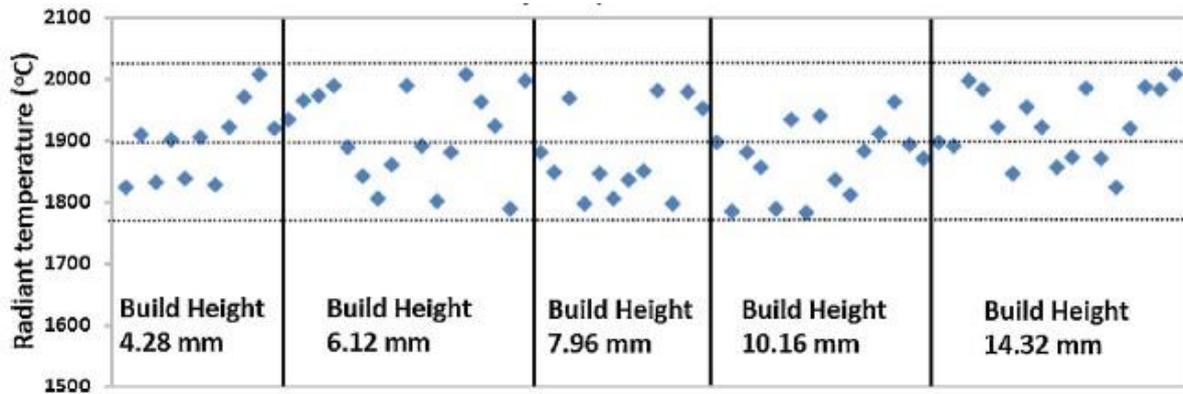


Figure 4.40

Cheng's results for the 'radiant temperatures' of the liquidus temperature [320]

As can be seen, the range of radiant temperatures is quite wide, with all temperatures measured at $1900 \pm 120^\circ\text{C}$. The authors did not show which of the data points arose from which scan speeds, but did state the average temperatures for each speed: 1890°C , 1855°C , and 1912°C for 400 mm/s, 600 mm/s, and 800 mm/s, respectively. The authors did not state any errors on these values.

4.7.2 - Applying to This Project

In order to test the applicability of this second-derivative method to the equipment and materials being used in this project, three experiments were planned:

1. Single-track scans, at a range of laser powers and speeds, were run across a powder-free IN-718 plate, while being recorded with the thermal camera. These scans were a proof-of-concept, to see if the method could be applied at a most fundamental level.
2. Single-track scans, with the same dimensions and laser parameters as previous, were performed on a layer of IN-718 powder, that was spread across rectangles that had been built of the same powder, while being recorded. These scans were considered more applicable to a real-world scenario, where laser-powder interactions could complicate the second-derivative method.
3. In order to test the method in an industrially-relevant setting, the following simple shapes and scanning strategies were built from IN-718 powder, at a range of laser powers and speeds, and recorded:
 - a. Square - Bi-directional vectors, 67° rotation between layers
 - b. Triangle - Bi-directional vectors, horizontal to ensure that the top point is over-heated
 - c. Circle - Uni-directional, inwardly spiraling, concentric circles

4.7.2.1 - Powder-Free Single Tracks

A 60mm × 40mm × 5mm plate of rolled IN-718 was placed onto the powder-free base plate, and levelled with the surrounding build chamber. The thermal camera was set to the 2048 pixel × 512 pixel frame, and the frame's centre was focused on the IN-718 plate. The camera was set to record at 400 frames per second. 27 vertical scans were performed, each spanning the 10mm height of the frame, and with 1mm horizontal spacing, as shown in Figure 4.41.

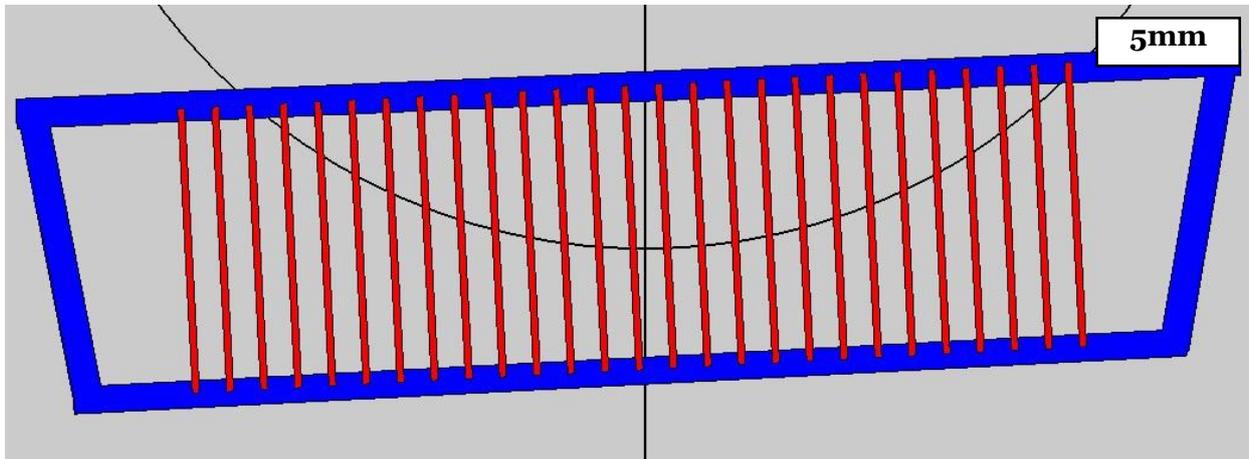


Figure 4.41

The design for the single-line tracks

Each combination of the laser speeds and powers in Table 4.7 were performed, with each combination being repeated three times.

Table 4.7

The power-speed combinations

Velocities [mm/s]	Powers [W]
500, 1000, 1500	100, 150, 200

In each of the recorded thermal images, any pixel that had an intensity value of below 200, or that was not part of the melt pool, was set to 0.

As established in a Section 4.5.2, it was reckoned that any pixels that fell either 5mm above, or 4.5mm below, the centerline of focus, should be discarded. The height of the 512-pixel frame was approximately 10mm. Therefore, all melt pools inside it could have been considered to be in focus.

The locations of each melt pool's cooling curve (i.e. the line from its peak intensity, along the pool's length, to where the normalized intensity falls from 200 to 0) was found. Since multiple melt pools were recorded for each line, even at 1500 mm/s, it was decided that any of those with cooling curves that did not fit entirely within 2.5mm above/below the centerline, could also be excluded.

Figure 4.42 shows the superimposed stack, with all the melt pools with cooling curves outside these limits excluded.

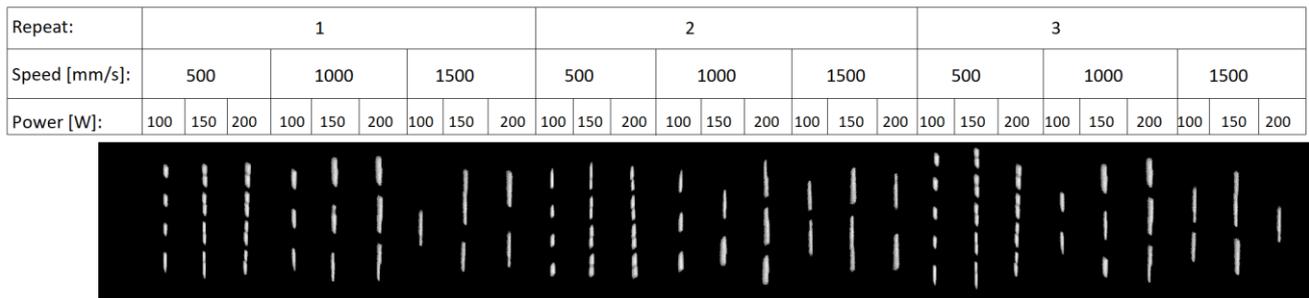


Figure 4.42

The superimposed image of all melt pools that fit within the centreline $\pm 2.5\text{mm}$ boundary

In this image, the melt pools are shown for (from left to right): 500 mm/s (100W, then 150W, then 200W), 1000 mm/s (100W, 150W, 200W), 1500 mm/s (100W, 150W, 200W), then the previous order of nine repeated twice more.

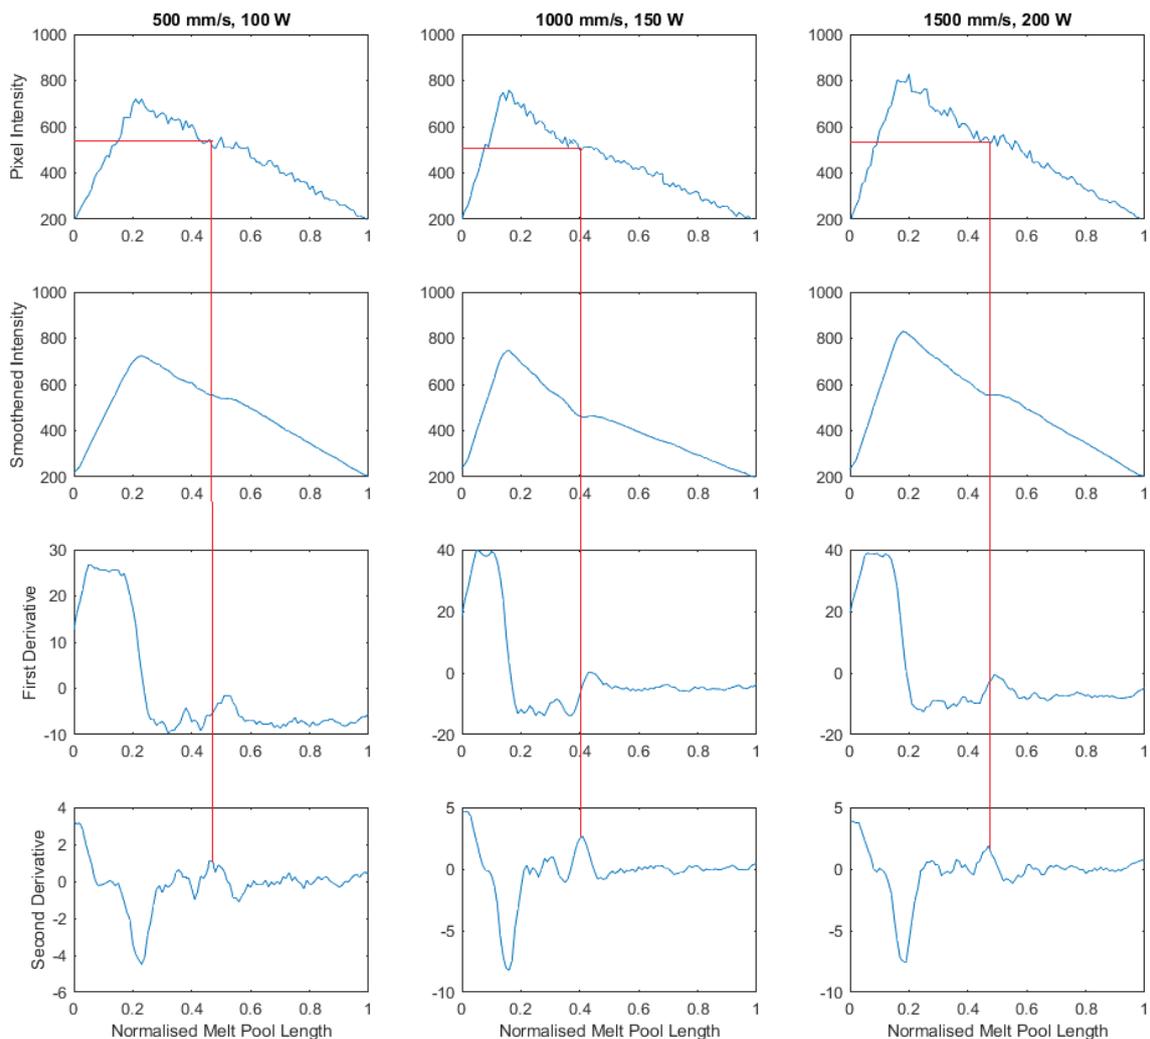


Figure 4.43

Determining the liquidus-intensity from the second-derivative

The cooling curves were extracted from each of the 79 melt pools. Figure 4.43 shows some example images of various cooling curves, along with their smoothed equivalent (rolling mean with window size of 5), and the first and second derivatives of the smoothed curve. It was found that without smoothing, the first and second derivatives varies wildly due to minute temperature fluctuations. The X axis in each of the images has also been normalized, to allow for easier comparison between plots. The reason that the second derivative peak is being extracted here, whereas the trough was extracted by Heigel, is explained simply by Heigel choosing to search for the 'end' of the liquidus arrest (greatest deceleration in temperature change), whereas here the 'start' (greatest acceleration) is being searched for.

As can be seen in Figure 4.43, by tracking the X axis location of the greatest spike in the second derivative of the cooling curve (i.e. after the peak temperature), and matching this location with the pixel intensity, it is possible to estimate the intensity that corresponds to the liquidus temperature of IN-718. These intensity values were 552, 518, and 531, for three melt pool (left to right in Figure 4.43). By identifying the intensity contours around each of the three melt pools, that correspond to these identified liquidus-intensities. It is possible to estimate the length, width, and area of each melt pool. In Figure 4.44, these contours are identified on the greyscale intensity map of the melt pools. Also, using the same scale, the melt pools are shown with the temperature-conversion applied, thresholded at 1336°C, for comparison between the two methods.

In previous sections, the aim of all analysis of melt pool images was to determine what the thermal camera could reveal about melt pool behavior, so therefore melt pool dimensions were presented in micrometers. In this current section, the aim of all image analysis relates to more fundamental experimentation on the camera itself. Therefore, all melt pool dimensions will be presented in pixels.

As can be seen, both the widths and lengths of each of the melt pools are notably larger when determined by the temperature-conversion method (TCM), than when determined by the intensity-contour method (ICM). The exact measurements are shown in Table 4.8.

The large difference between the areas can be attributed to three facts:

1. The intensity threshold that distinguishes liquid from solid will be slightly different between the TCM and the ICM.
2. Due to the small size of the melt pools, the very small difference in threshold can manifest as a large percentage difference. In the case of the 1000 mm/s, 150W melt pool, for example, a width difference of only 3 pixels corresponds to a 37.5% difference.
3. Due to the quadratic scaling nature of area, any increases in length and width will be greatly amplified, hence the 67.3% difference for the area of the same melt pool.

It must be noted, however, that the three melt pools used in the second-derivative methods previously, were specifically chosen for the ease with which the calculations could be applied to them. Many of the recorded melt pools did not feature such clear mushy zones in their cooling profiles.

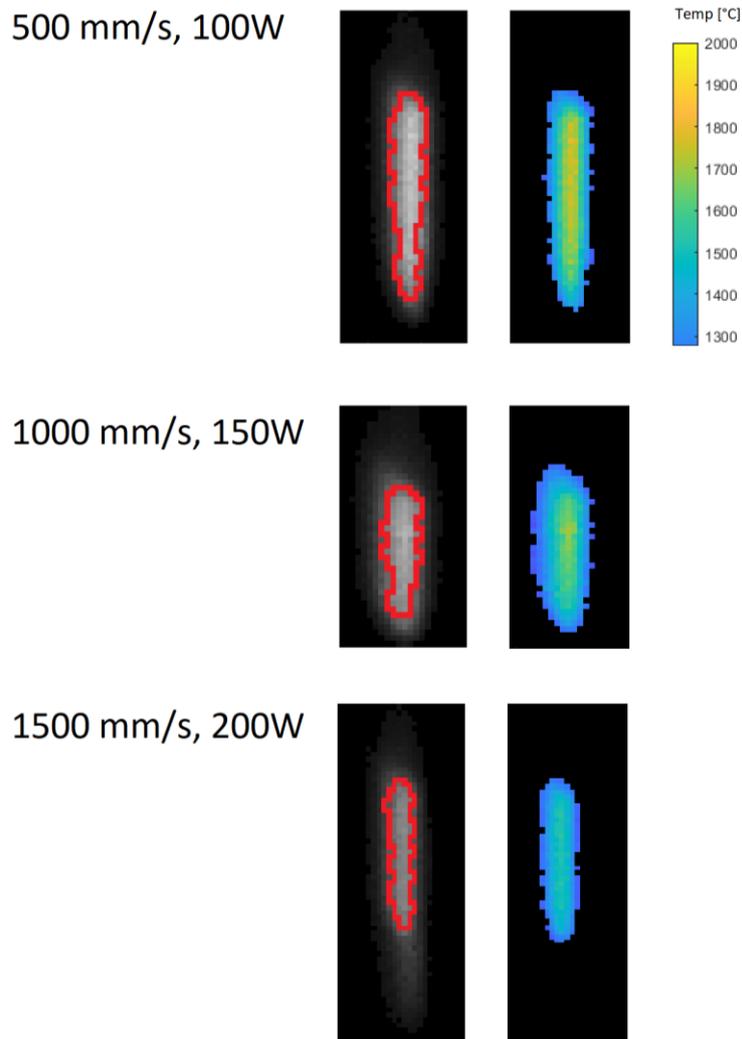


Figure 4.44

Comparing melt pool boundaries as determined by the liquidus-intensity contour (left), and the temperature conversion method (right)

Table 4.8

Comparing the TCM and ICM melt pool dimensions

Melt Pool	500 mm/s, 100W	1000 mm/s, 150W	1500 mm/s, 200W
Width (TCM)	10 pixels	11 pixels	8 pixels
Width (ICM)	8 pixels	8 pixels	7 pixels
Difference	25%	37.5%	12.5%
Length (TCM)	44 pixels	28 pixels	37 pixels
Length (ICM)	40 pixels	23 pixels	31 pixels
Difference	10%	21.7%	19.4%
Area (TCM)	264 pixels ²	184 pixels ²	177 pixels ²
Area (ICM)	192 pixels ²	110 pixels ²	130 pixels ²
Difference	37.5%	67.3%	36.2%

Figure 4.45 features two more examples of cooling curves. The curve on the left has a discernable mushy zone, but only when it is already known that the liquidus intensity appears around 510 to 560. The curve on the right has no discernable mushy zone at all.

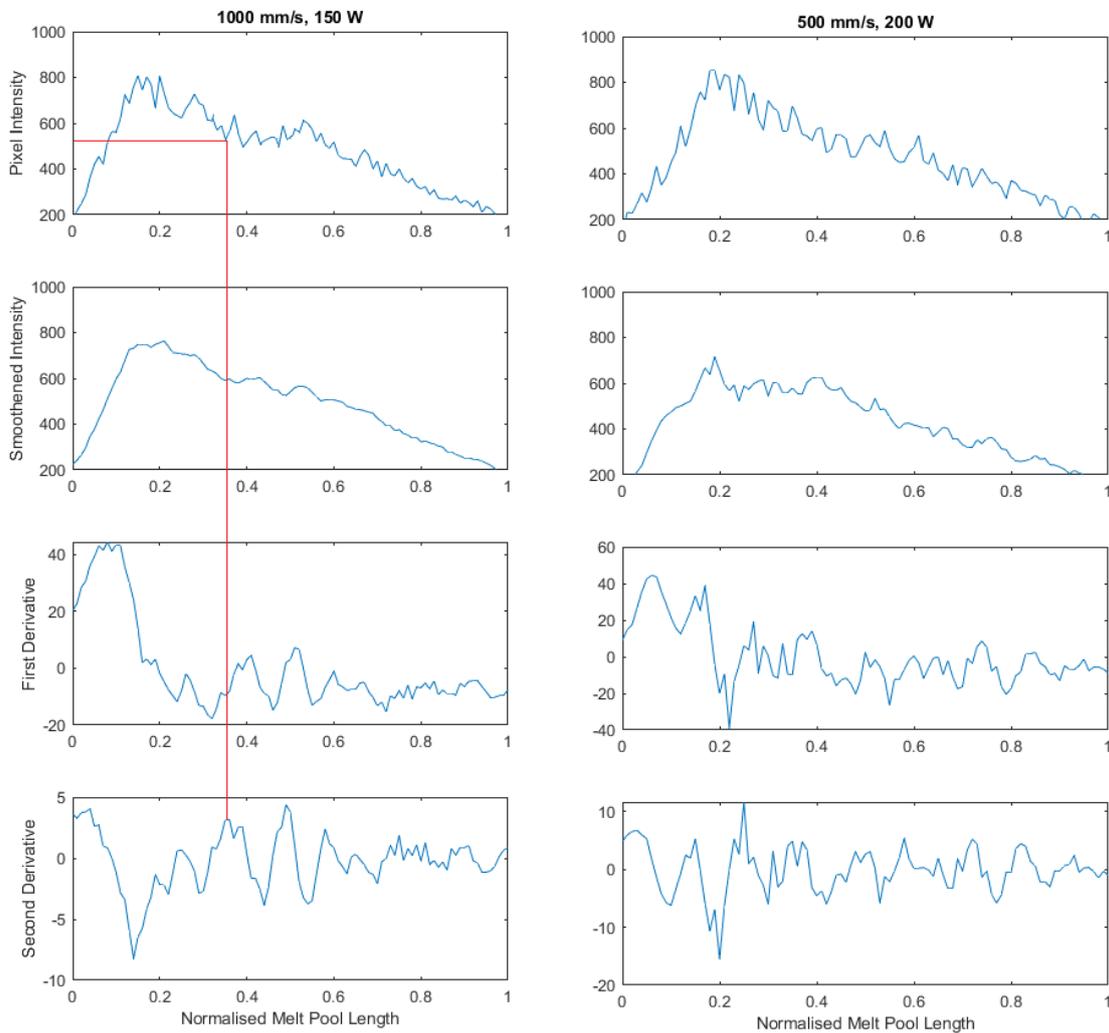


Figure 4.45

An example of a 'Mid' melt pool (left), and a 'Poor' melt pool (right)

Each of the 79 recorded melt pools could be placed into one of three categories:

1. The liquidus intensity could be found by searching for the largest peak in the second-derivative of the cooling curve, as demonstrated with the three 'good' melt pools previously. The value could be found automatically, with a simple MATLAB code.
2. The liquidus intensity could be found, but only when the peak in the second-derivative chart that corresponds to the known intensity range was manually selected, as demonstrated with the 'mid' cooling curve in Figure 4.45. These melt pools required a lot of human effort to find, as it was difficult to write a MATLAB code that could automatically select the correct peaks, without constantly getting confused by anomalies.
3. The liquidus temperature could not be found, either by MATLAB code or human effort, as demonstrated by the 'poor' cooling curve in Figure 4.45.

Importantly, deciding which of the three categories that a given melt pool should fit into, was not a task that could be easily automated. There existed a circular dependency: selecting the correct peak in the second-derivative graph, would reveal the liquidus-intensity to be approximately 530; and the correct peak in the second-derivative graph, was whichever corresponded to a liquidus-intensity of approximately 530. Ultimately, each of the 79 cooling curves and second-derivative graphs had to be checked manually, before they could be divided up into the categories. Table 4.9 shows how many melt pools fit into each of the three categories.

*Table 4.9
The number of 'Good', 'Mid, and 'Poor' melt pools for each power-speed combination*

Speed	500 mm/s			1000 mm/s			1500 mm/s		
Power	100W	150W	200W	100W	150W	200W	100W	150W	200W
Total MPs	13	13	12	8	8	9	5	6	5
'Good' MPs	6	5	5	2	3	3	2	3	1
'Mid' MPs	4	4	2	3	3	4	1	2	1
'Poor' MPs	3	4	5	3	2	2	2	1	3

Table 4.10 shows, for all of the melt pools from which data can be drawn (i.e. those in the 'Good' and 'Mid' categories): the mean and standard deviation liquidus-intensities, and the percentage that the melt pool shrunk, between the TCM and ICM method, expressed as a mean and a standard deviation.

*Table 4.10
The liquidus intensity, and TCM/ICM differences for all 'Good' and 'Mid' melt pools*

Speed	500 mm/s			1000 mm/s			1500 mm/s		
Power	100W	150W	200W	100W	150W	200W	100W	150W	200W
Mean Liq-Int	524.5	529.7	514.1	551.3	552.4	519.6	554.7	546.6	552.0
Std. Liq-Int	25.5	26.2	21.3	27.4	10.0	26.4	17.0	17.3	14.2
Diff. Width [%]	22.3	28.1	14.3	13.1	12.2	14.9	15.6	27.3	19.7
Std. Width [%]	4.1	5.1	2.8	2.6	2.4	3.1	3.2	5.3	3.9
Diff. Length [%]	20.4	15.9	20.0	7.6	22.4	13.8	19.3	8.8	12.2
Std. Length [%]	4.1	2.9	4.0	1.6	4.6	2.6	3.7	1.6	2.2
Diff. Area [%]	39.4	34.3	70.9	73.3	71.0	45.6	60.1	38.6	40.3
Std. Area [%]	7.4	7.5	14.0	15.9	14.7	8.6	13.0	8.4	8.2

As can be seen from the table, there is strong consistency in the liquidus-intensities identified. All intensities fell within the 510 to 560 range, with standard deviations that were typically around 5% of the value. The overall mean liquidus-intensity, calculated from all recorded melt pools, was 538.5, with a standard deviation of 23.4.

The number of melt pools that could be automatically processed, without human input, however, was of concern, especially since this test had been designed to make liquidus-intensity identification as easy as possible, and was not representative of real-world build conditions. The percentages of melt pools which fell into the ‘Good’, ‘Mid’, and ‘Poor’ categories, were 38%, 30%, and 32%, respectively.

4.7.2.2 - Single Tracks with Powder

The same 27 single-line tracks were printed again, following the same 9 power-speed combinations. The only difference between this test and the previous, was that the tracks were here build with powder. In the first five layers of the build, 27 thin rectangles were scanned, for the single-line tracks to be scanned upon in the sixth layer. The rectangles each measured 10mm × 0.5mm, and were positioned so that the single-line track would be at its centerline. A 30µm layer height was used. These rectangles were scanned with a laser power of 125W, a laser speed of 500 mm/s, a hatch spacing of 50µm. The intention of these rectangles was to have the single-line tracks be scanned onto a substrate built from powder, rather than onto a solid baseplate, to better recreate real-world conditions.

Apart from this change, the build, recording, and data processing followed the exact same procedure as previously. Table 4.11 shows the number of melt pools in each category, and the information that was learned about the liquidus-intensity, and the effect on pool dimensions. Due to the fact that all of the 1500 mm/s - 200W melt pools fell into the ‘Poor’ category, no calculations were possible on this combination.

Two major differences are notable between this results table, and those from the powder-free tracks. Firstly, the number of melt pools falling into each of the three categories has shifted. Due to the more chaotic nature of building with powder, with the stochastic nature of particle melting, and the increased tendency for spatter, there are fewer ‘Good’ melt pools, and more ‘Poor’ melt pools. Whereas in the powder-free build the number of ‘Good’, ‘Mid’, and ‘Poor’ pools, were 38%, 30%, and 32%, respectively, here the numbers are 21%, 23%, and 56%, respectively.

Secondly, the liquidus-intensity range has shifted. Whereas the powder-free intensities all fell between 510 and 560, here they fall between 480 and 530. This is likely due to the fact that a material’s emissivity depends on its surface quality, and the liquid of a molten plate is going to differ slightly from the liquid of molten powder in terms of surface shape and stability.

The overall liquidus-intensity from these tracks, calculated from all melt pools, was 501.7, with a standard deviation of 27.5.

Combined, these two shifts point to a noteworthy trend: as experiments to find an alloy’s liquidus-intensity move from the ideal to the real-world, the liquidus-intensity is going to change, and the ability to accurately measure the liquidus-intensity is going to diminish.

Table 4.11

The number of melt pools in each category, their liquidus intensity, and the TCM/ICM difference

Speed	500 mm/s			1000 mm/s			1500 mm/s		
Power	100W	150W	200W	100W	150W	200W	100W	150W	200W
Total MPs	14	12	10	9	7	6	6	7	4
'Good' MPs	3	3	2	2	1	1	2	2	0
'Mid' MPs	4	2	3	1	2	1	2	2	0
'Poor' MPs	7	7	5	6	4	4	2	3	4
Mean Liq-Int	521.5	517.6	488.4	499.0	495.9	510.4	489.5	478.6	-
Std. Liq-Int	29.8	16.3	28.2	27.6	13.4	26.8	28.8	25.0	-
Diff. Width [%]	29.7	26.3	10.1	21.6	12.0	29.6	17.2	13.9	-
Std. Width [%]	6.4	5.2	1.9	3.9	2.5	6.4	3.2	3.0	-
Diff. Length [%]	13.9	7.7	18.6	16.1	21	16.4	20.7	11.6	-
Std. Length [%]	2.8	1.5	3.5	3.3	3.9	3.4	4.4	2.1	-
Diff. Area [%]	74.4	27.8	70.8	74.4	57.9	60.0	71.6	35.1	-
Std. Area [%]	14.6	5.5	12.9	15.6	12.2	11.9	14.4	7.3	-

4.7.2.3 - Simple Shapes

The following shapes, as shown in Figure 4.46, were built from IN-718 powder, upon a 316L steel base plate (due to the metallurgical compatibility of the two materials, and a lack of IN-718 plates):

- A 5mm × 5mm square, with bi-directional vectors, and 67° rotation between layers
- A 5mm (base) × 5mm (height) triangle, with bi-directional vectors, that were all aligned horizontally to ensure that the top point is over-heated
- A circle with 5mm diameter, scanned with concentric circles progressing from outer to inner

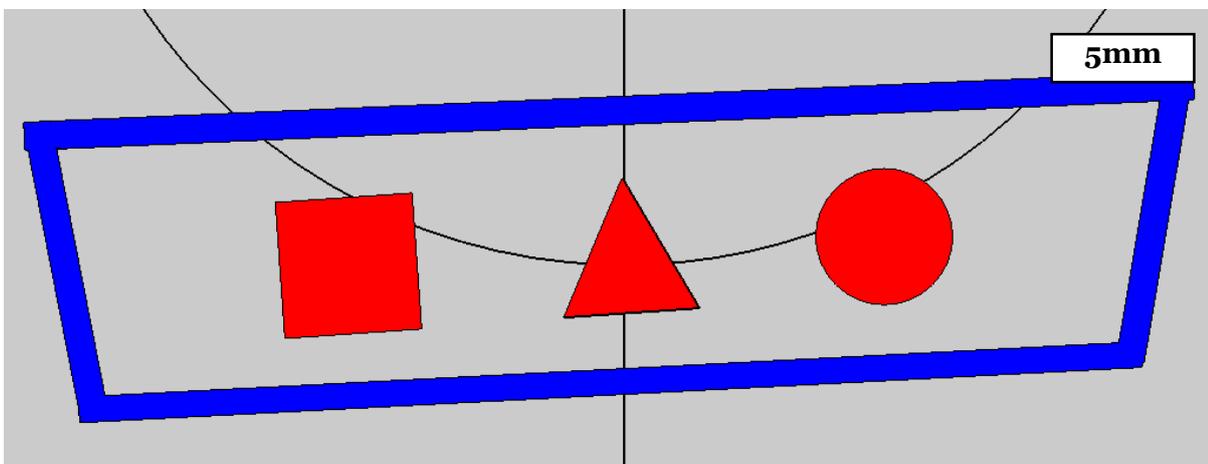


Figure 4.46

The simple shapes used to mimic a 'real-world' application

A layer height of 30 μ m and a hatch spacing of 50 μ m was used. The shapes were built to 45 layers. Each of the nine power-speed combinations that were used for the single-line tracks were used again here, with five layers being built at each combination, and the final of each of the five being recorded. The order of the combinations was: 500 mm/s (100W, then 150W, then 200W), 1000 mm/s (100W, 150W, 200W), 1500 mm/s (100W, 150W, 200W).

Between 200 and 500 melt pools were recorded for each of the nine power-speed combinations: far more than could be checked manually, as had been done for each of the single-line experiments. Therefore, a fully automated method of searching for the intensities that corresponded to the second-derivative peaks was devised.

Firstly, the cooling curve had to be extracted. In the single-line experiments, where all tracks had been vertical and unidirectional, the centerline of each melt pool could easily be found. This was not so easy in the current experiment, where the direction of a melt pool may not have been clear (in the case of the 67° rotation square), or the melt pool may not be longer than it is wide (in the case of the triangle tip), or the melt pool may be curved (in the case of the circle).

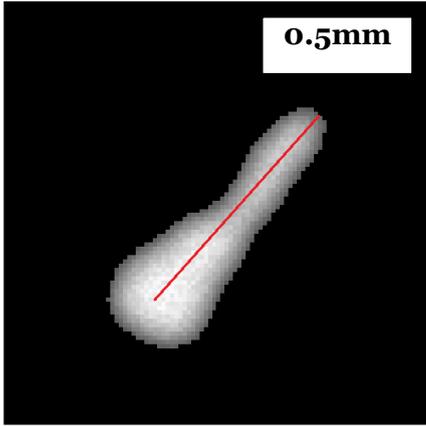
Designing a MATLAB script that could automatically identify the centerline of the melt pool, in its direction of travelling, would have been very difficult. This process could be avoided, however, since it was not necessary to extract 'the' cooling curve from each melt pool, but rather 'a' cooling curve (or a temperature-position curve to be exact): any line of intensities, extending outward from the pool's peak intensity, should in theory capture the liquidus-intensity and associated second-derivative peaks.

In order to find 'a' cooling curve, that should contain the required data, the peak intensity of each melt pool was found, and 360 lines extending outward from it, in 1° steps, were searched. The line that extended the longest, before intensity dropped to below 300, was chosen as the cooling curve. The intensities of all the pixels that lay beneath this line were extracted. Figure 4.47 shows this process for two different melt pools, one from the square (at 1000 mm/s and 150W) and one from the circle (at 1500 mm/s and 200W).

In Figure 4.47, the red line shows the path of the extracted cooling curve, starting at the highest pixel intensity, then proceeding to the melt pool edge.

Based on the single-line experiments, it was assumed that the liquidus-intensity should lie within the 400 to 600 range. This section of the extracted cooling curve was then analysed, starting from the first point that intensity fell below 600, and ending with the first point that intensity fell below 400. This segment was then smoothed, and its second-derivative was calculated, using the same method as described for the powder-free single-line tracks. Each peak in the second derivative that had a prominence (i.e. height difference between it and its nearest trough) of at least five was identified, and cross referenced with its corresponding intensity. For the two melt pools in Figure 4.47, this resulted in the identified intensities demonstrated in Figure 4.48.

Square



Circle

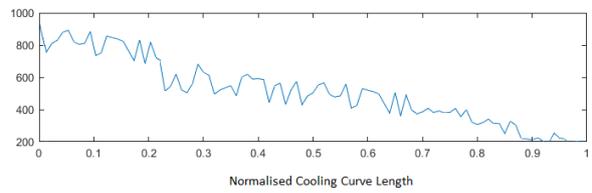
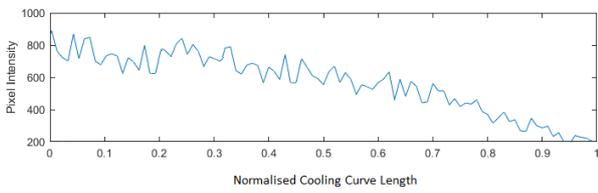
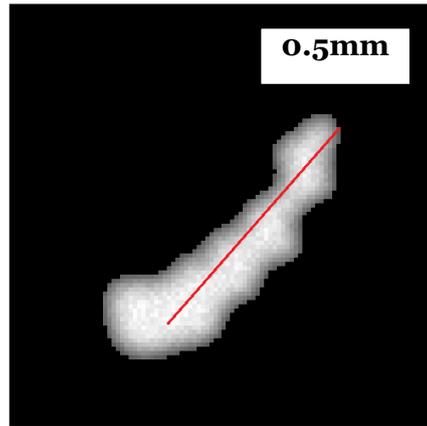
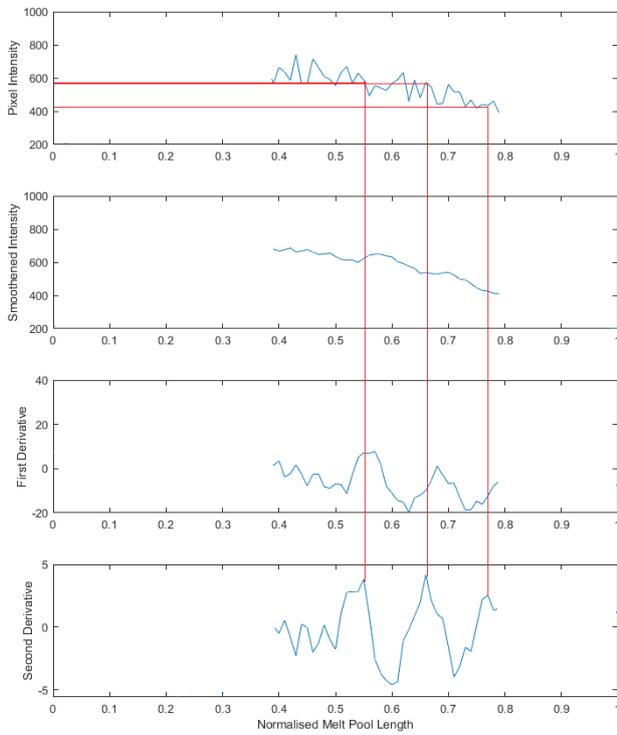


Figure 4.47

Cooling curves, as extracted from a melt pool from the square, and from the circle

Square



Circle

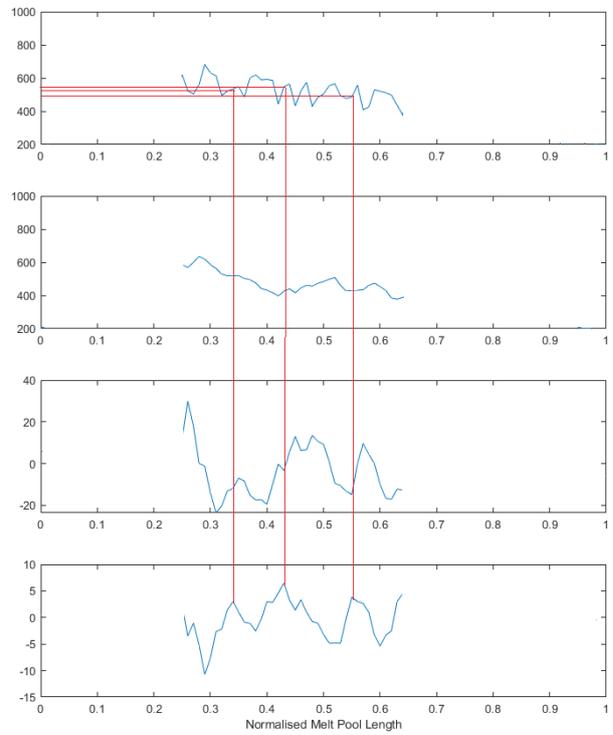


Figure 4.48

The second-derivative method being applied to the square and circle melt pools shown previously

The identified intensities, each which may possibly correspond to the liquidus-intensity, were 429, 575, and 579 for the melt pool from the square, and 497, 542, 568 for the melt pool from the circle. It was hypothesized that, by extracting possible liquidus-intensities like this, from each of the 3146 melt pools recorded, across all nine power-speed combinations, and plotting them on a histogram, the shape of the distribution might reveal which is the ‘real’ liquidus-intensity.

This calculation was performed. An average of 4.3 intensities were extracted from each of the 3146 melt pools, for 12,678 in total. The histogram plotted from these is shown in Figure 4.49.

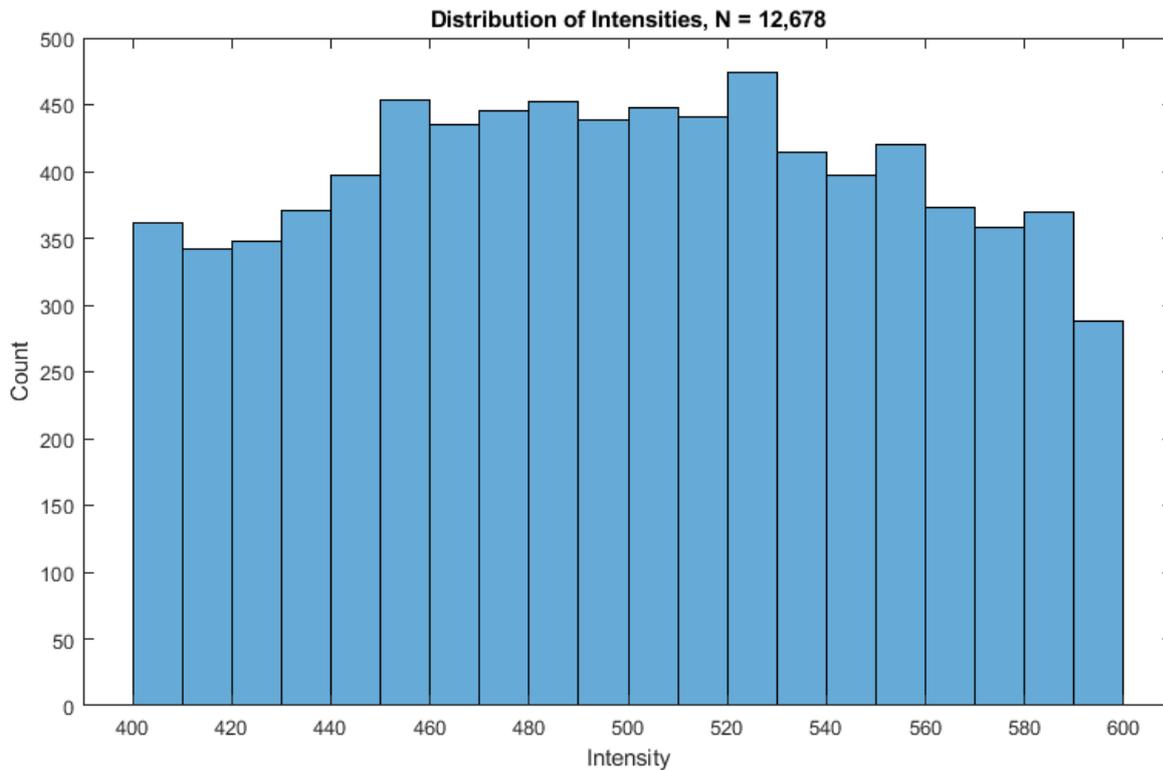


Figure 4.49

A histogram of all the possible liquidus-intensities identified from all power-speed combinations, in all shapes

As can be seen, the distribution is almost uniform, or very-slightly normal. Slight peaks could be argued to exist around the intensities of 450, 480, or 520, although it would be difficult to conclude from the histogram that any single intensity should be considered the ‘correct’ liquidus-intensity.

4.7.3 - Conclusion on Bypassing Emissivity

In the case of the powder-free single-line tracks, it was quite straightforward to identify a mean liquidus-intensity of 538.52, with a standard deviation of 23.48. In the case of the single-line tracks on powder, it was less straightforward, although still possible, to identify a mean liquidus-intensity of 501.75, with a standard deviation of 27.52. In the ‘real-world’ case of building simple shapes however, with more complex scan strategies and geometries, no single liquidus-intensity could be identified.

The fact that a liquidus-intensity can only be obtained when performing single-line tracks, combined with the fact that any liquidus-intensity identified for single-line tracks may not hold for more complex builds, means that the entire second-derivative method should be viewed with skepticism.

Therefore, for the remainder of this work, the temperature-conversion method, using a fixed emissivity value, will be used. It must be understood that, due to this limitation, all melt pool dimensions and extracted temperatures should be considered appropriate for comparison with other dimensions and temperatures from the same experimental setup, but not appropriate as absolute and accurate values.

4.8 - Comparing to Ex-Situ and Modelling Methods

In order to test the fidelity of thermal camera results to ex-situ measurements of melt pool dimensions, made with the optical microscope, a set of single-line, powder-free scans were made on a rolled IN-718 plate. The plate was positioned to be level with the build chamber, and the camera was focussed on it, with the line of focus at the centreline of the frame. The 512px frame was used, at 400fps. Lines of all 25 combinations of the power and speed shown in Table 4.12 were scanned.

Table 4.12

The power-speed combinations

Velocities [mm/s]	Powers [W]
500, 750, 1000, 1250, 1500	100, 125, 150, 175, 200

Three repeats of each combination were scanned, with vectors spaced 0.5mm apart. The Netfabb design for the test is shown in Figure 4.50. The three repeats were spaced evenly across the build (i.e. one set is on the left, one in the middle, one in the right).

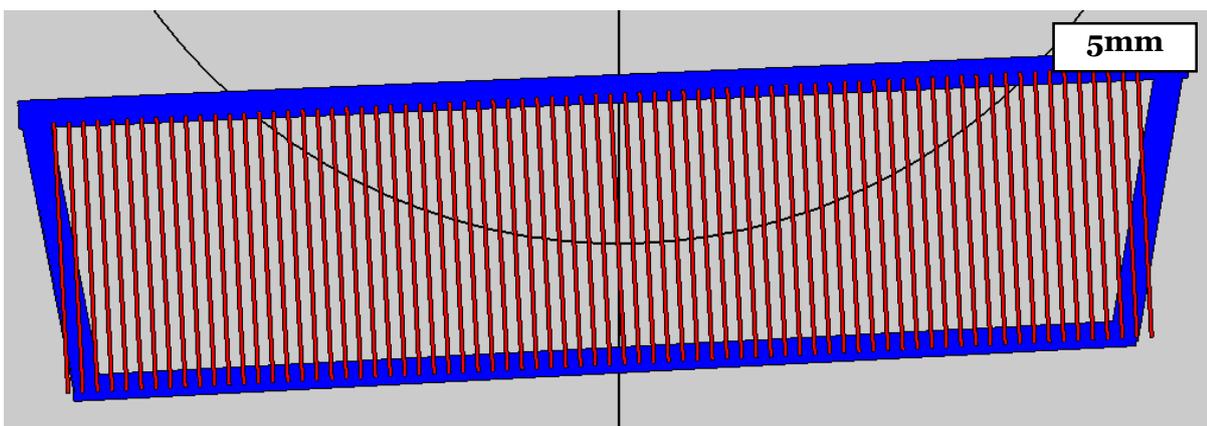


Figure 4.50

The Netfabb design for the comparison test

As can be seen from Figure 4.50, not all vectors lay inside the camera frame. Therefore, simply scanning them left to right would have resulted in uncertainties, when it came to working out which

images corresponded to which vector. To prevent this, all three repeats for each vector were scanned in order (i.e. leftmost 100W - 500 mm/s, centre 100W - 500 mm/s, rightmost 100W - 500 mm/s, leftmost 125W - 500 mm/s, centre 125W - 500 mm/s, etc).

4.8.1 - Thermal Camera Results

After the build was complete, the melt pool images were analysed using the same methods described in previous sections:

1. The images were sorted into stacks of each power-speed combination
2. The intensity-temperature conversion was applied
3. Any pixel below 1336°C or not attached to the melt pool was set to zero
4. Any melt pool that did not entirely fit within 2.5mm above/below the centerline was ignored
5. The dimensions and temperature profiles of the remaining melt pools were recorded

As a reminder, all melt pool lengths have been adjusted to account for motion blur, as described in a Section 4.4. That is why melt pools at higher speeds (which appear longer in thermal images) will often report as having smaller lengths. Table 4.13 shows the number (N) of extracted melt pools for each power-speed combination, plus the means (M) and standard deviations (SD) of their lengths (L), and widths (W). All dimensions are in μm .

*Table 4.13
The thermal camera dimensions for each power-speed combination*

Power	Speed	500 mm/s		750 mm/s		1000 mm/s		1250 mm/s		1500 mm/s	
		M	SD	M	SD	M	SD	M	SD	M	SD
100W	N	10		8		7		6		3	
	L	78.69	28.61	71.10	27.33	80.38	24.25	88.43	19.06	59.29	19.55
	W	25.79	8.01	36.52	13.45	26.81	9.18	29.69	11.08	28.43	9.51
125W	N	11		8		7		5		3	
	L	119.79	36.11	79.65	23.88	100.81	33.58	67.60	19.28	66.12	17.15
	W	44.19	8.87	30.58	9.06	38.59	9.59	27.90	8.90	30.60	10.41
150W	N	10		8		7		6		3	
	L	141.60	38.70	121.36	34.82	104.51	38.27	98.58	39.08	87.67	23.86
	W	55.47	19.77	40.29	8.57	42.93	14.87	29.41	9.60	23.69	6.30
175W	N	11		8		6		6		2	
	L	129.81	41.61	145.11	52.47	116.99	29.98	85.41	30.74	89.49	22.17
	W	49.87	15.49	56.50	13.85	44.33	13.93	40.93	10.75	28.32	7.81
200W	N	11		8		6		6		3	
	L	160.57	62.81	132.24	41.89	133.21	35.50	105.62	22.13	94.32	27.17
	W	63.27	16.22	57.76	16.87	49.14	14.12	44.93	13.28	37.17	14.37

The surface plots in Figure 4.51 show the mean widths and lengths from Table 4.13, for all powers and speeds. It is clear that, despite some fluctuations due to the stochastic nature of L-PBF, the general trend (larger melt pool dimensions at greater energy levels) holds as expected.

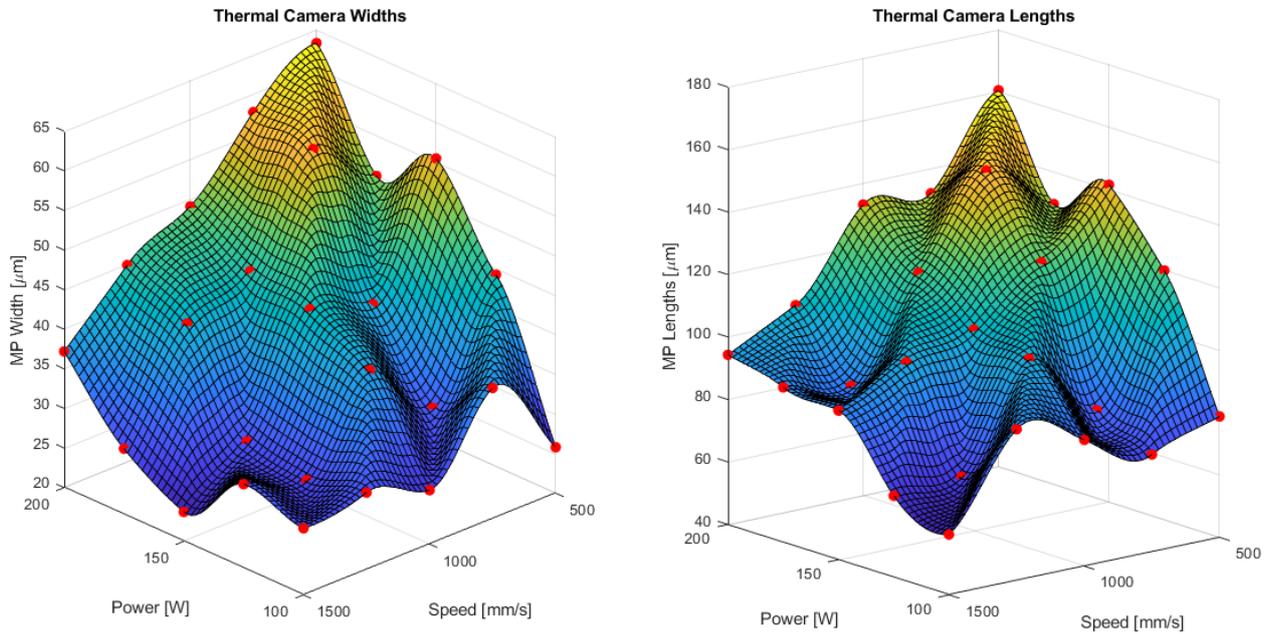


Figure 4.51

The widths and depths according to the thermal camera

4.8.2 - Optical Microscopy Results

After the build, the rolled IN-718 plate was sectioned, mounted, ground, polished, and etched. Sectioning was done in a way to ensure that at least 5 (and in most cases 6) repeats of each power-speed combination was visible upon the polished surfaces, as shown in Figure 4.52.

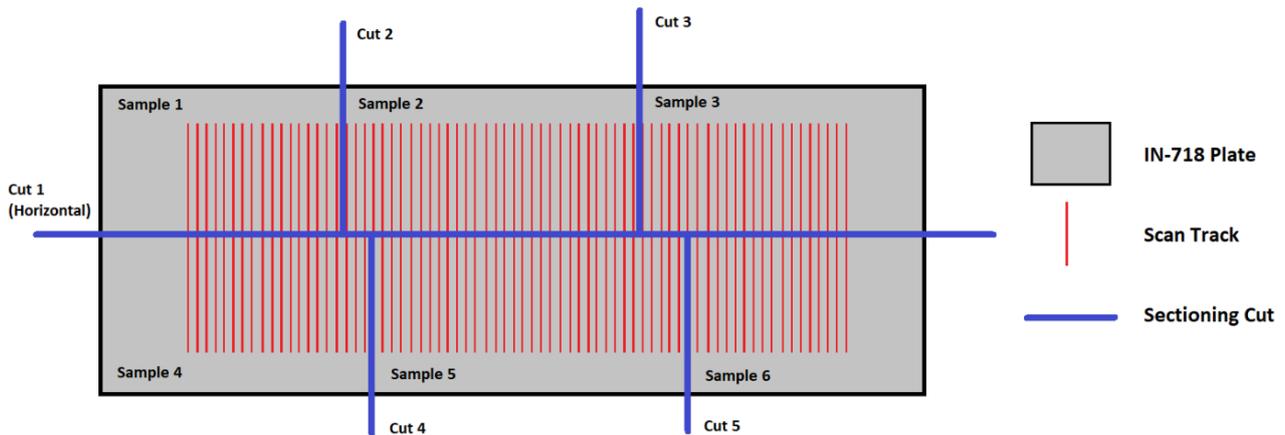


Figure 4.52

A schematic of the way the IN-718 plate was cut

Images were taken of each visible melt pool with the optical microscope. Outlines were drawn around each melt pool boundary, using a Huion Inspiroy H640P graphics tablet and stylus. This was performed manually, rather than with an automated script, due to the nature of the images. Creating a MATLAB function that was robustly capable of determining melt pool boundaries from grain and cell boundaries, or scratches and other such imperfections, proved to be an unnecessarily complex task. Due to the number of repeats, it was assumed that manual segmentation errors would be mitigated.

Once the outlines were extracted, A MATLAB script was written to smooth them, using a Gaussian-weighted average with a window of 100 pixels below the surface, and a window of 400 pixels below the surface. Figure 4.53 shows: one of the 500 mm/s - 200W melt pool images; the outline as drawn by the graphics tablet; the smoothed outline; and the 6 outlines from all repeats, superimposed onto each other.

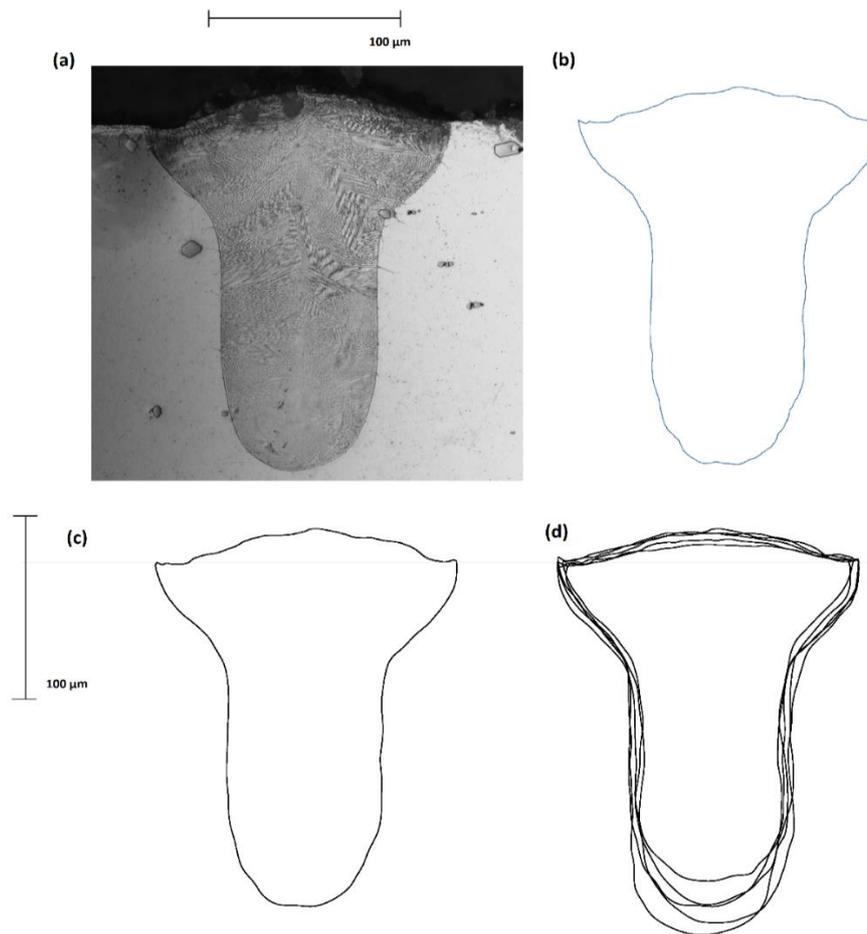


Figure 4.53

An optical image of a 500 mm/s - 200W melt pool (a); the outline of the melt pool (b); the smoothed outline (c); the outlines of the six repeats superimposed onto each other (d)

Table 4.14 shows the number (N) of melt pools for each power-speed combination identified via optical microscopy, plus the means (M) and standard deviations (SD) of their depths (D), and widths (W).

Table 4.14

Optical microscopy dimensions from each power-speed combination

Power	Speed	500 mm/s		750 mm/s		1000 mm/s		1250 mm/s		1500 mm/s	
		M	SD	M	SD	M	SD	M	SD	M	SD
100W	N	6		6		5		6		6	
	W	53.19	17.24	54.07	15.78	58.10	19.86	39.03	8.53	43.29	12.53
	D	94.47	28.12	91.38	24.70	83.88	22.12	77.81	16.11	74.57	22.47
125W	N	6		6		6		5		6	
	W	64.45	17.96	57.57	19.58	60.37	19.20	44.66	14.90	51.72	18.90
	D	139.22	40.51	127.02	36.83	110.21	33.24	100.56	30.15	82.24	30.93
150W	N	6		6		6		6		6	
	W	88.48	18.86	64.65	14.89	73.37	23.87	69.05	19.21	50.39	16.79
	D	138.45	35.90	118.51	27.67	142.04	45.66	87.40	29.01	88.31	26.98
175W	N	5		5		6		6		6	
	W	99.37	32.52	76.57	27.40	77.65	19.17	59.51	20.70	53.12	12.50
	D	221.44	85.51	151.90	40.66	157.95	54.66	110.28	27.55	77.32	16.48
200W	N	6		5		6		6		5	
	W	127.87	31.09	75.95	16.14	68.33	18.84	67.32	13.56	52.08	19.11
	D	201.84	76.42	181.25	65.14	166.03	62.38	136.98	41.07	82.31	24.07

The surface plots in Figure 4.54 show the mean widths and depths from Table 4.14, for all powers and speeds. Again, despite some fluctuations due to the stochastic nature of L-PBF, the general trend holds as expected.

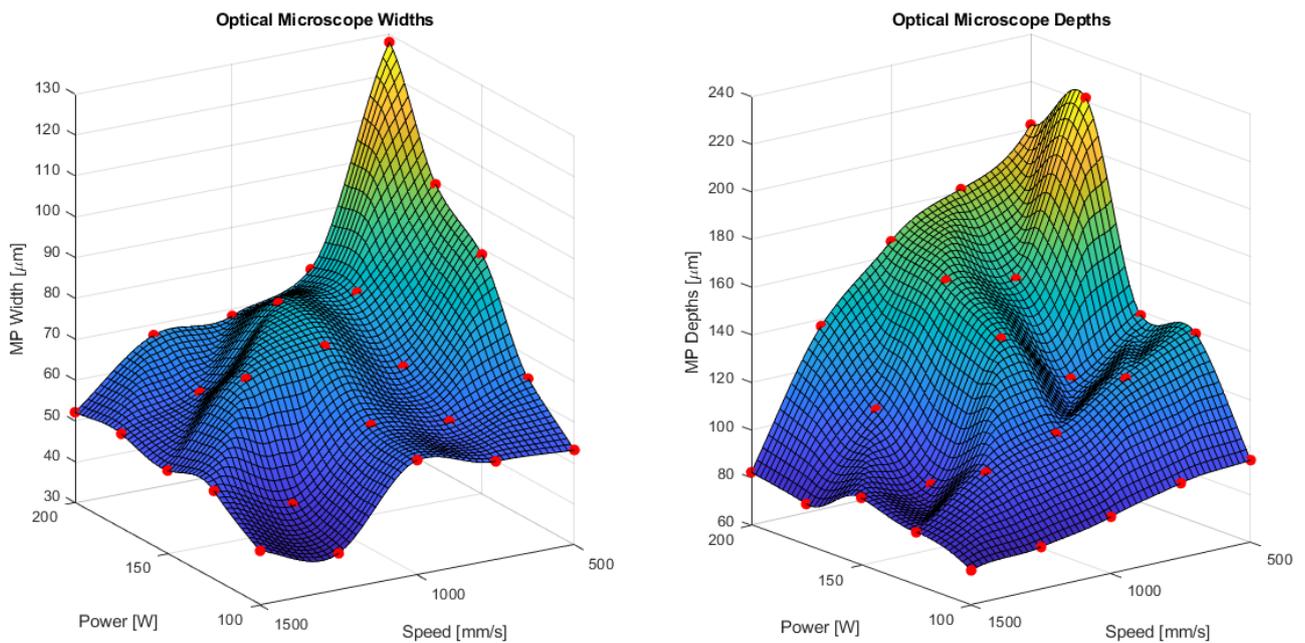


Figure 4.54

The melt pool widths and depths, according to the optical microscope

4.8.3 - Modelling Results

Each of the single-track, powder-free scans were simulated in the ThermoCalc AM Module. The lengths, widths, and depths of the simulated melt pools are shown in Table 4.15.

Table 4.15
ThermoCalc dimensions for each power-speed combination

Power	Speed	500 mm/s	750 mm/s	1000 mm/s	1250 mm/s	1500 mm/s
100W	L	126.6	116.1	125.7	103.9	102.9
	W	48.8	44.5	44.4	41.0	37.4
	D	52.0	49.9	51.7	39.3	43.1
125W	L	145.2	144.2	138.1	114.9	109.8
	W	56.0	57.6	46.8	43.6	39.7
	D	78.4	72.0	58.7	51.9	40.5
150W	L	163.1	154.4	155.4	141.5	119.4
	W	68.6	55.7	49.4	45.7	39.3
	D	96.3	88.2	70.4	55.8	44.6
175W	L	206.8	182.4	166.1	145.9	108.8
	W	72.2	71.9	60.6	54.2	41.1
	D	111.5	103.1	79.1	70.4	48.3
200W	L	202.0	178.8	179.9	134.1	125.7
	W	83.2	73.7	64.4	55.5	48.0
	D	127.3	112.9	99.4	71.1	56.4

The surface plots in Figure 4.55 show the widths, lengths, and depths from Table 4.15, for all powers and speeds. Due to the nature of modelling methods, this surface is far less turbulent than the previous sets.

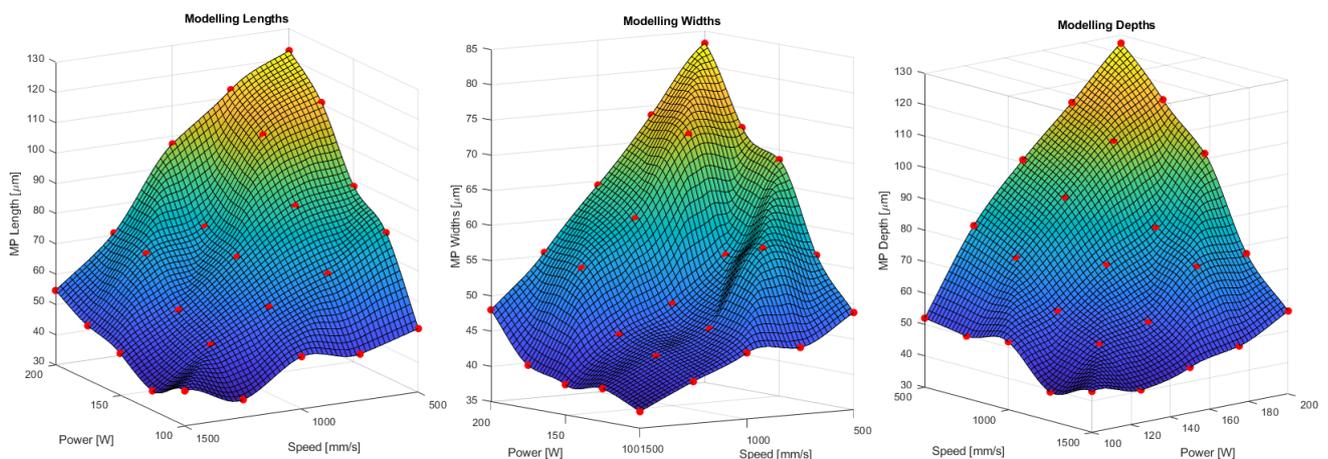


Figure 4.55
The melt pool dimensions, according to ThermoCalc

4.8.4 - Data Comparison

In Figure 4.56, measurements for melt pool lengths, widths, and depths, from each of the three sources, are plotted against each other.

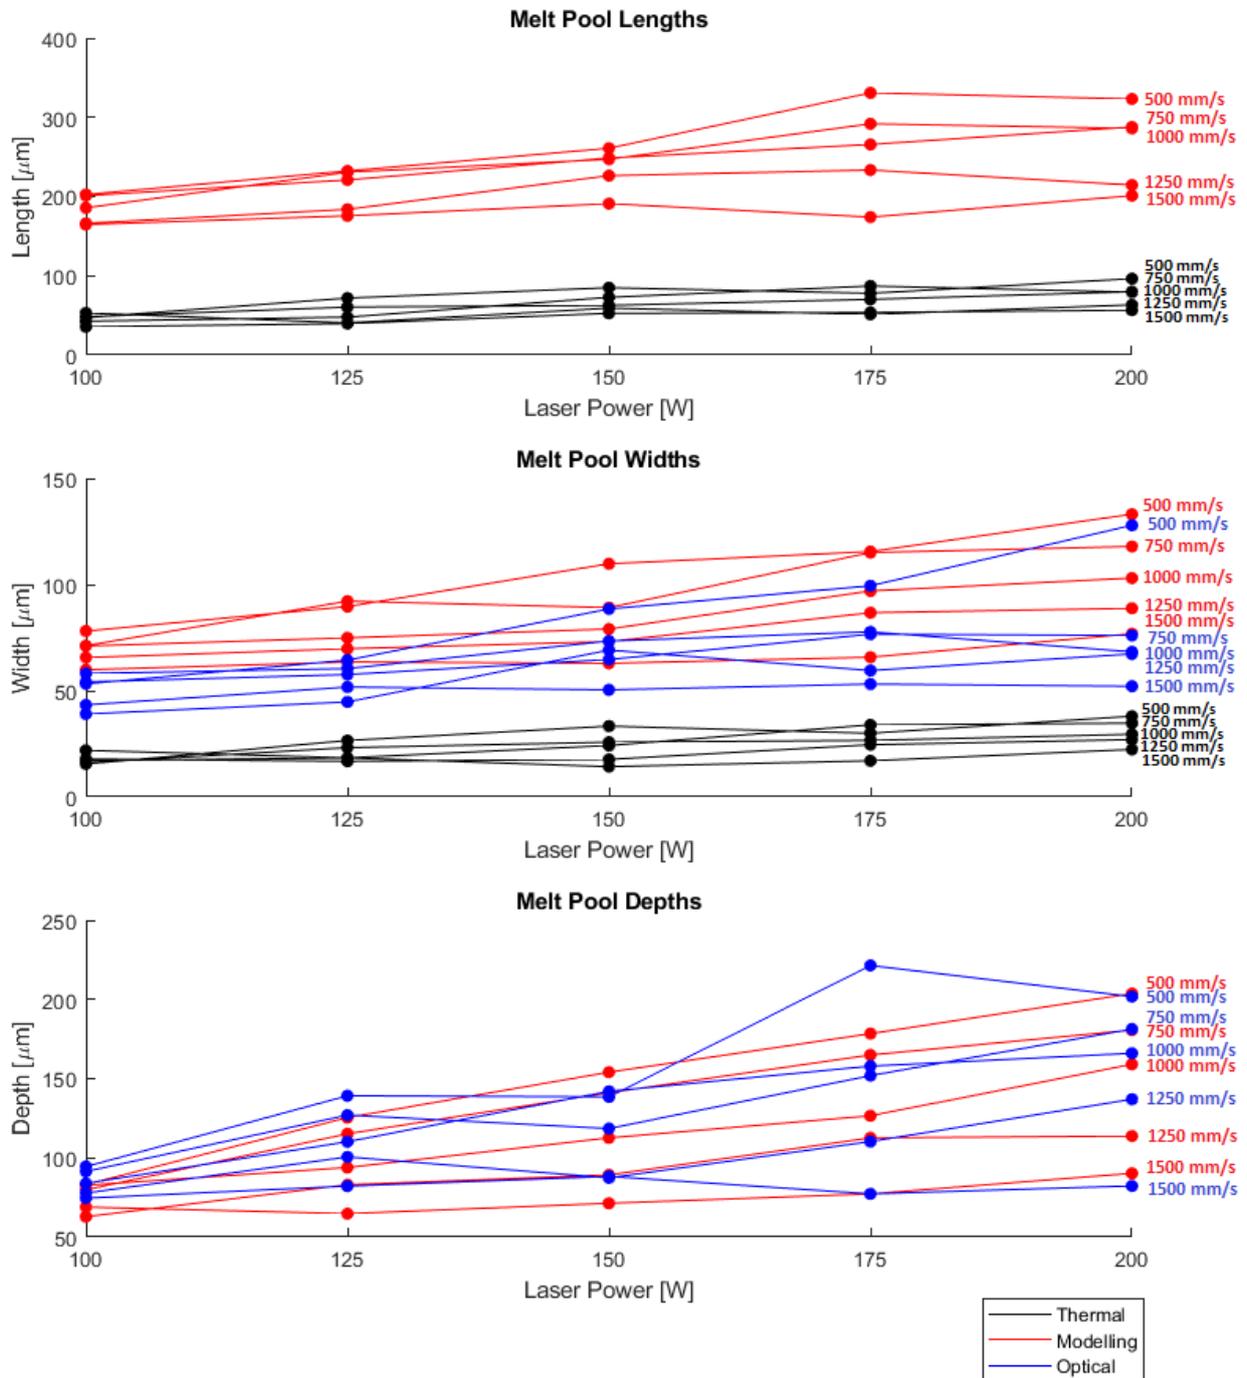


Figure 4.56

A comparison of the melt pool length, width, and depth measurements, according to the three different methods

As is clear from the graphs, there exists strong fidelity between the optical microscopy results, and the modelling results. The thermal camera results, however, significantly under-predict both melt pool length, and melt pool width, in comparison to the other two methods. It is possible that the reason for this is due to the fixed emissivity value making the melt pool area seem smaller than it is in reality. This hypothesis will be discussed in the Future Work chapter.

This finding again suggests that, although the thermal camera can point out trends in data from similar experimental setups, its ability to provide definitive values is limited. The value of such trends will be evaluated in further chapters.

4.9 - Camera Capability Conclusions

The strengths and weaknesses of the thermal camera were assessed, and it was determined that valuable information about the melt pool and solidification process could be learned from thermal images, so long as the necessary compensations, adjustments, and limitations were implemented and respected.

Throughout the remainder of this work, all MATLAB scripts that analyse thermal images should be assumed to make the following adjustments:

1. Whenever a dimension is being converted from pixels to micrometers, the angle of the dimension will be calculated, by finding the angle of the long axis of the melt pool that the dimension belongs too. The conversion factor used will vary between $19.4 \mu\text{m}/\text{px}$ (when the dimension is horizontal relative to the frame) and $20 \mu\text{m}/\text{px}$ (when the dimension is vertical).
2. When melt pool areas are being converted from square pixels to square micrometers, the mean of the two factors will be taken for simplicity (i.e. $k_a = 19.7^2 \mu\text{m}^2/\text{px}^2$).
3. Motion blur is accounted for in all melt pool lengths, by subtracting $1.3V$ from the recorded dimension, where V is the laser speed in mm/s.
4. The peak temperature of a melt pool is increased by 6% for every mm (50 pixel rows, as per the previous conversion) that its centroid appears above/below the frame's central row. Likewise, the mean temperature is increased by 5% / 50 rows, and widths are decreased by 5% / 50 rows. It is understood that this adjustment may introduce errors, due to the possibility that different experimental setups and machine parameters may require slightly different percentage increases or decreases. Since performing this characterisation for each new setup was unfeasible, however, and not applying any adjustments would definitely lead to large errors, the method was used as a best available option.

Based on the evaluation experiments laid out in Section 4.6, the strengths and limitations of the camera have been determined. The camera was found to be highly adept at describing melt pool thermal signatures influenced by vector length, and by pause durations that allowed for material cooling. The camera struggled to detect changes in thermal signatures caused by heat sink volume, or insulating powder. The small size of the data sets from which these determinations were made, however, must be noted.

It was found that, despite the clear theoretical limitations of using a fixed emissivity conversion factor to calculate pixel temperature from intensity, in practice this method is the only viable option for this experimental setup. The stochastic nature of melt pool thermal signatures caused far too much random fluctuation in the second derivative of the cooling curve, for it to be applicable to bypass emissivity.

Comparison with ex-situ and modelling results showed that the thermal camera systematically under-predicted both length and width measurements, likely due to a combination of all the factors outlined previously. Therefore, thermal camera measurements should be considered valuable for highlighting trends in thermal signatures, rather than taken as an absolute measure of any dimension or metric.

5 - In-Situ Monitoring and Microstructure

This chapter details an experiment that was planned with the intention of establishing a link between the thermal signatures of a melt pool, and the microstructure and mechanical properties that result from the pool's solidification. It was theorised that, by predicting from first-principles the ways in which build parameters will affect the thermal gradient and solidification velocity of a set of cubes, it would be possible to predict the as-built grain structure that will develop.

In order to test this hypothesis, ABD-900 AM cubes of various sizes and parameters were built on the Aconity LAB, and monitored with the thermal camera. After the cubes were built, optical microscopy, hardness testing, SEM analysis, and EBSD microscopy was carried out.

5.1 - Experimental Setup

Nine sets of geometrically-identical cubes were built, as shown in Figure 5.1. Each set consisted of six 3mm × 3mm × 5mm (XYZ) cubes, shown in red, four 5mm × 5mm × 5mm, shown in green, and one 10mm × 10mm × 5mm cube, shown in yellow. The numbers superimposed on the squares show the order of scanning.

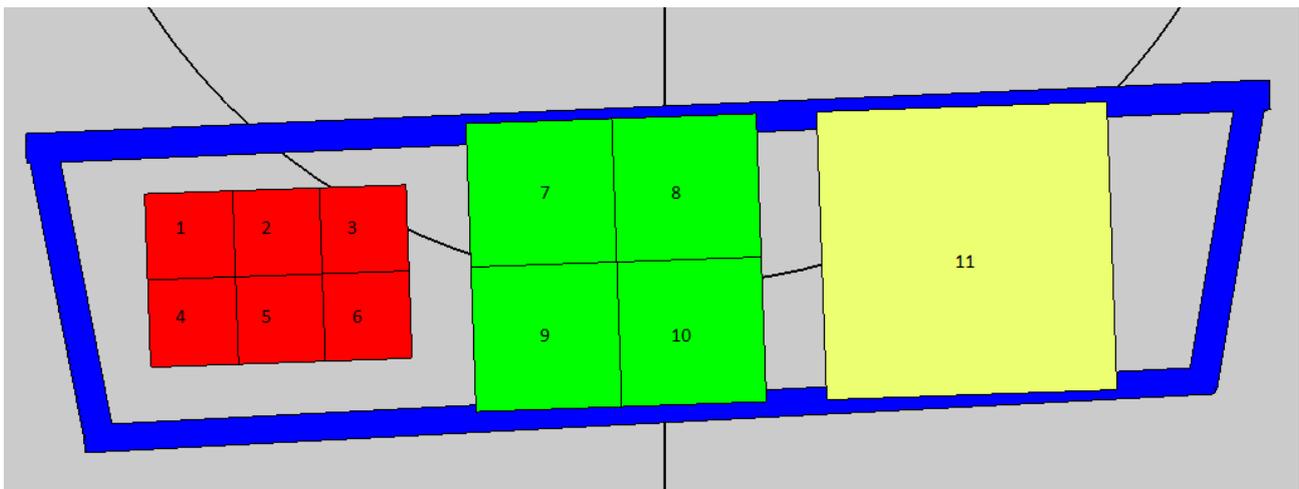


Figure 5.1

A schematic of the layout of the 11 cubes, and their order of scanning

The six 3mm cubes, and the four 5mm cubes, are touching each other but not overlapping (in order to fit all 5mm cubes in the frame, and to ensure consistency with the 5mm cubes, in the case of the 3mm cubes). All builds were performed with a laser power of 360W, and a laser speed of 1050 mm/s. These parameters were selected based on the work of Dr Minh Phan, who ran similar cube tests to those outlined in Section 6.1, and determined that this power-speed combination could result in full density, keyhole porosity, or LoF porosity, depending on other factors.

The two variable parameters were the layer heights and hatch spacings. One build was performed with each combination of the values in Table 5.1.

Table 5.1
The layer height and hatch spacing combinations

Layer Heights [μm]	40, 60, 80
Hatch Spacings [μm]	40, 80, 120

Assuming that, within each of the nine builds, each repeated cube may be considered identical, this provides 27 different data sets (i.e. three layer heights, times three hatch spacings, times three scan lengths). A bi-directional scanning pattern was used, alternating between vertical and horizontal every layer.

After the builds were completed, the cubes were all removed from the base plate. It was found that the weak bonds between the six 3mm cubes, and the four 5mm cubes, often separated during the EDM process. Where they did not separate during EDM, they were split with gentle pressure from two pairs of pliers.

From each build, cube 1 (3mm), cube 7 (5mm), and cube 11 (10mm) were each selected for ex-situ analysis. The cubes were sectioned along their centreline, parallel to the Y axis (i.e. horizontally in the CAD images shown above). The samples were mounted in Bakelite, with the sectioned surface facing downwards. The surfaces were ground and polished according to the previously-described recipe. Once the surfaces had a mirror-finish, they were:

1. Analysed with the optical microscope
2. Analysed with the EBSD
3. Etched, according to the method previously-described
4. Analysed with the optical microscope again
5. Analysed with the SEM

5.2 - Data Gathering

A large quantity of data was extracted from the builds, both in-situ and ex-situ, with the intention of revealing the meso-structural, microstructural, and mechanical properties of each cube. The methods of data gathering will be described here, with the data itself being presented and analysed in a further subsection.

5.2.1 - Thermal Camera

The builds were each recorded with the 512-pixel thermal camera frame, recording at 400fps, and focussed at the frame's centreline.

Four layers were recorded in each of the nine builds:

1. A few layers before the 3mm build height (approximately layer 73 for the 40 μ m layer height builds, approximately layer 48 for the 60 μ m builds, etc.), when scanning was horizontal
2. A few layers after the 3mm build height, when scanning was vertical
3. A few layers before the 5mm build height, when scanning was horizontal
4. The final layer at 5mm (or the layer before), when scanning was vertical

Due to the fact that the orientation of each layer's scanning pattern depended on the layer height, and the amount of time the thermal camera took to save a layer of images depended on hatch spacing, the exact layers at which each of the four recordings occurred was not exact and repeated.

Between 683 and 3172 images were recorded, for each of the 108 cases (9 builds \times 3 cube types \times 4 recorded layers).

Figure 5.2 shows an example melt pool, extracted from the 10mm cube of Build 7, at the 3V recording, with the intensity-temperature conversion and a colour grading applied.

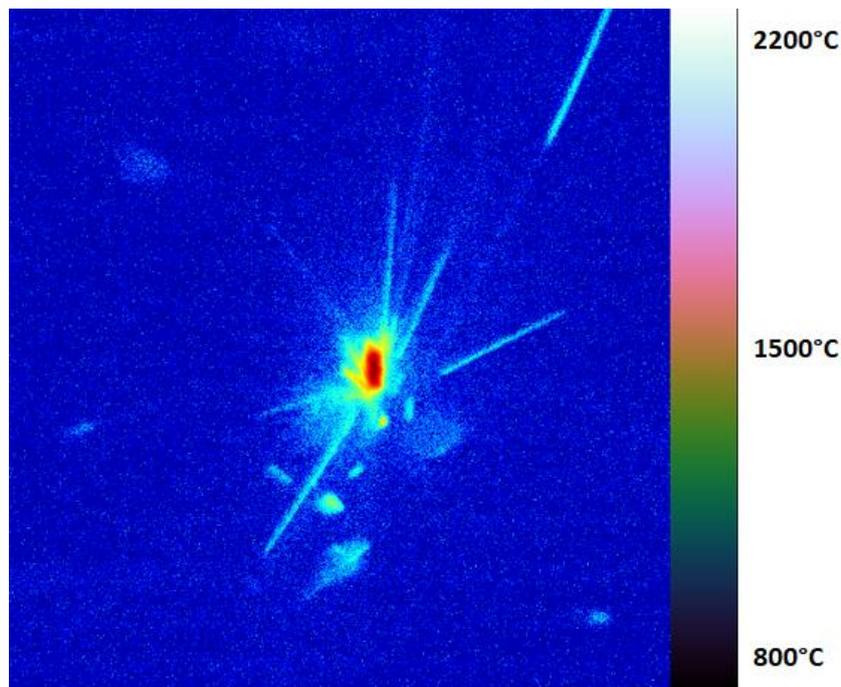


Figure 5.2

An example melt pool, before thresholding was applied

The following types of data can be extracted from each individual melt pool image:

- Maximum temperature
- Width
- Area
- Thermal gradients (front, back, left, right)
- Quantity of spatter
- Mean temperature
- Length
- Aspect ratio
- Cooling rates (implied from laser speed)
- Centroid location

By analysing the sequence of images from each cube, the following can additionally be extracted:

- Means and standard deviations for each of the aforementioned quantities
- Regions of heat accumulation
- Thermal histories of single pixels

5.2.2 - Optical Microscopy

Before etching, each of the 27 cubes were analysed with the optical microscope, in order to search for porosity. Figure 5.3 shows the 10mm cube from Build 9, which experienced huge LoF porosity due to its 80 μ m layer height, and 120 μ m hatch spacing.

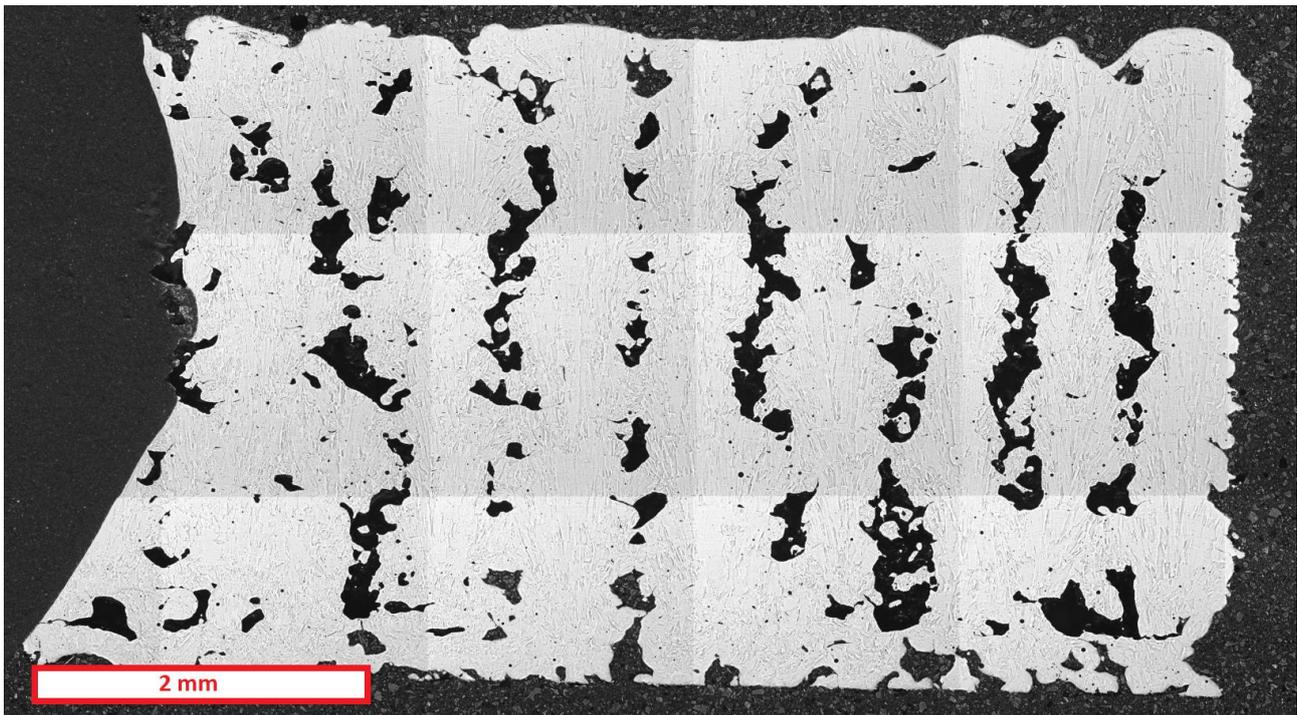


Figure 5.3

An optical image of a surface with much LoF porosity (Build 9, 10mm length, 80 μ m layer, 120 μ m hatch)

From these images, the following data could be extracted:

- Total density
- Mean pore area and standard deviation
- Number of pores
- Mean pore circularity and std. deviation

After etching, the 27 cubes were analysed again with the optical microscope, in order to gain a semi-subjective understanding of the extent of epitaxial grain growth within each. A square of approximately 1.1mm \times 1.1mm was image, located at the centre of each surface. Figure 5.4 shows the image of the 5mm cube from Build 5 (60 μ m layer height, 80 μ m hatch spacing).

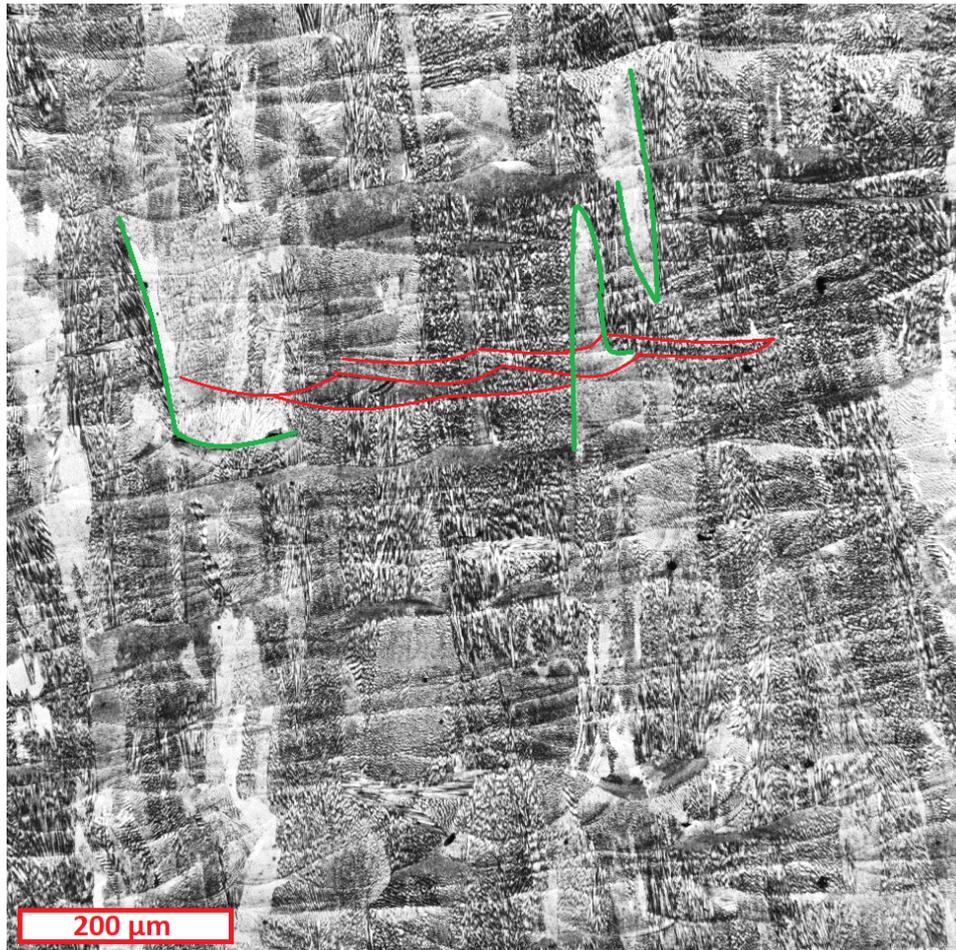


Figure 5.4

An optical image of an etched surface, showing grains stretching across melt pool boundaries (Build 5, 10mm length, 60μm layer, 80μm hatch). Some melt pool boundaries are annotated in red, and some possible grain boundaries are annotated in green.

Although difficult to quantify systematically, it is evident from this image that grains are growing through melt pool boundaries, as demonstrated by the fact that patches of visually-similar crystallographic texture are overlaid with the elongated curves of boundaries.

5.2.3 - SEM

Images were taken from the centre of each of the 27 etched surfaces, at an electron beam accelerating voltage of 10 keV and a working distance of between 10 and 12mm. The intention was to analyse the sizes of sub-grain features (henceforth called cells). Figure 5.5 shows the image recorded of the 10mm cube, from Build 2. A MATLAB script was written, to find the cell sizes. The code first converted the SEM images to binary by normalising all pixels and using a threshold value of 0.5, then searched through every point in a 20-pixel × 20-pixel grid overlaid on the image. At each point that did not lie on a boundary, lines were extended outwards searching for the nearest cell boundary. The length of each line was found, and the minimum of these lengths was chosen as the cell width. The diagram in Figure 5.6, extracted from the same image as Figure 5.5, demonstrates this process.

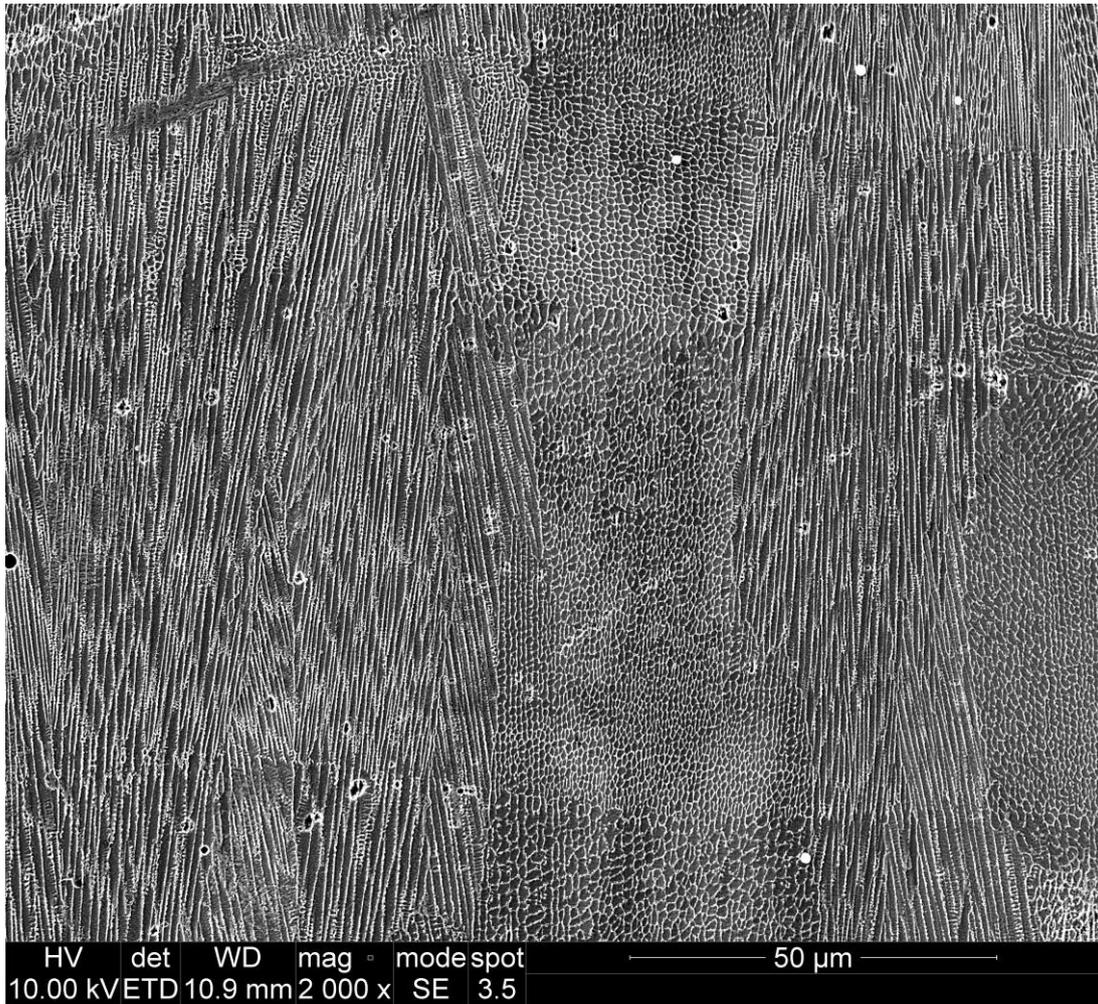


Figure 5.5

An SEM image of the sub-grain microstructure (Build 9, 10mm length, 80μm layer, 120μm hatch)

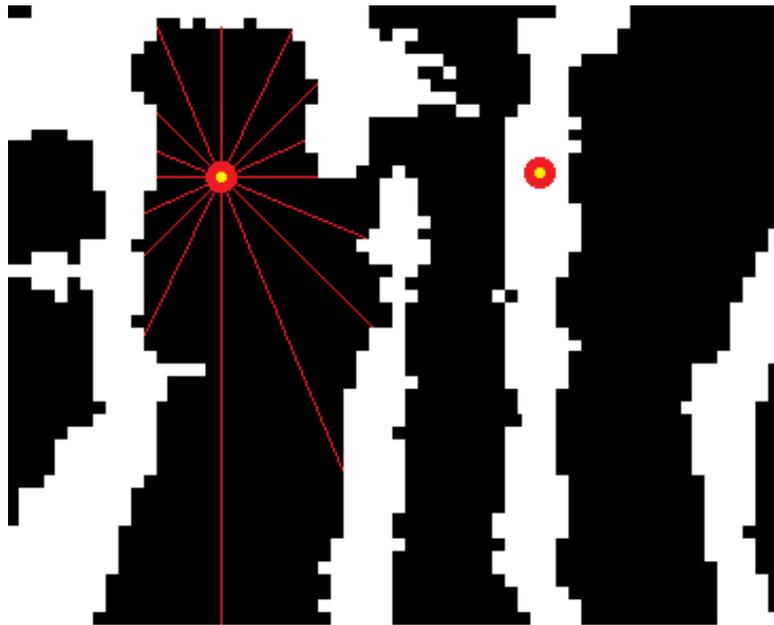


Figure 5.6

A schematic demonstrating the grid and arms search method within each SEM image (Build 9)

In Figure 5.6, the two red and yellow dots are points that have been selected by the 20px × 20px grid. The rightmost point landed on a cell boundary (white), so is ignored. The leftmost point extends pairs of arms outwards (in angle steps of 22.5°), until a boundary is contacted. It is determined that the horizontal pair of arms is the shortest, and therefore their combined length is recorded as the cell width.

It was understood that this semi-random method of selecting points to analyse, would introduce some error. Of the 36,000 points searched in each image, some fell in the corner of cells, where the arms only had to extend a few pixels to have found boundaries. In order to prevent such results from skewing the findings, two methods of data cleaning were applied.

First, outliers were removed from the list of cell widths using the inter-quartile range (IQR) method. Any width was removed that fell outside the bounds of $Q1 - (1.5 \times IQR)$, or $Q3 + (1.5 \times IQR)$. This removal largely affected only outliers that were too big. Secondly, the histogram of remaining widths was analysed. In every case, there existed a significant peak at the very left of the graph, disrupting what would otherwise be a normal distribution, as shown in Figure 5.7.

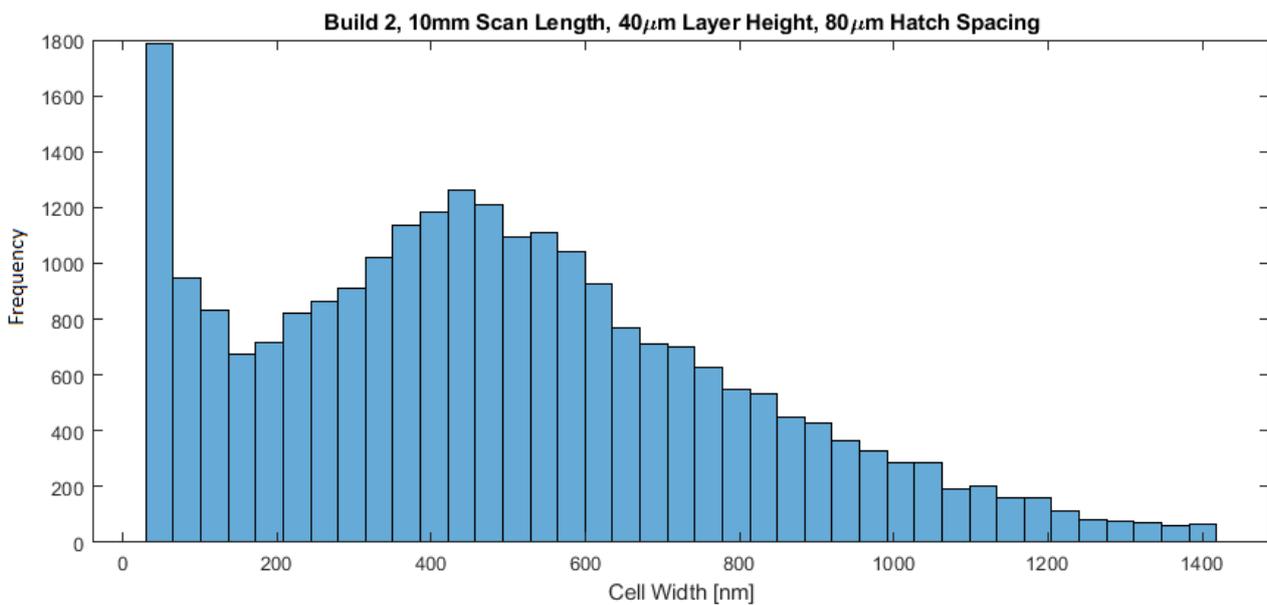


Figure 5.7

The distribution of the cell width measurements, before filtering is applied

The fact that this peak at the lowest widths existed in all 27 distributions, regardless of the mean or variance of the rest of the shape, and regardless of how large or small cells visually appeared in the SEM image, strongly suggested that it was an artefact of these small-corner search points described previously. Based on visual inspections of both the histograms and arm pairs within the images, it was reasoned that any width smaller than 50nm could be assumed to be an artefact, and should be removed. This removal, performed on the same data set as Figure 5.7, yielded the normal distribution shown in Figure 5.8.

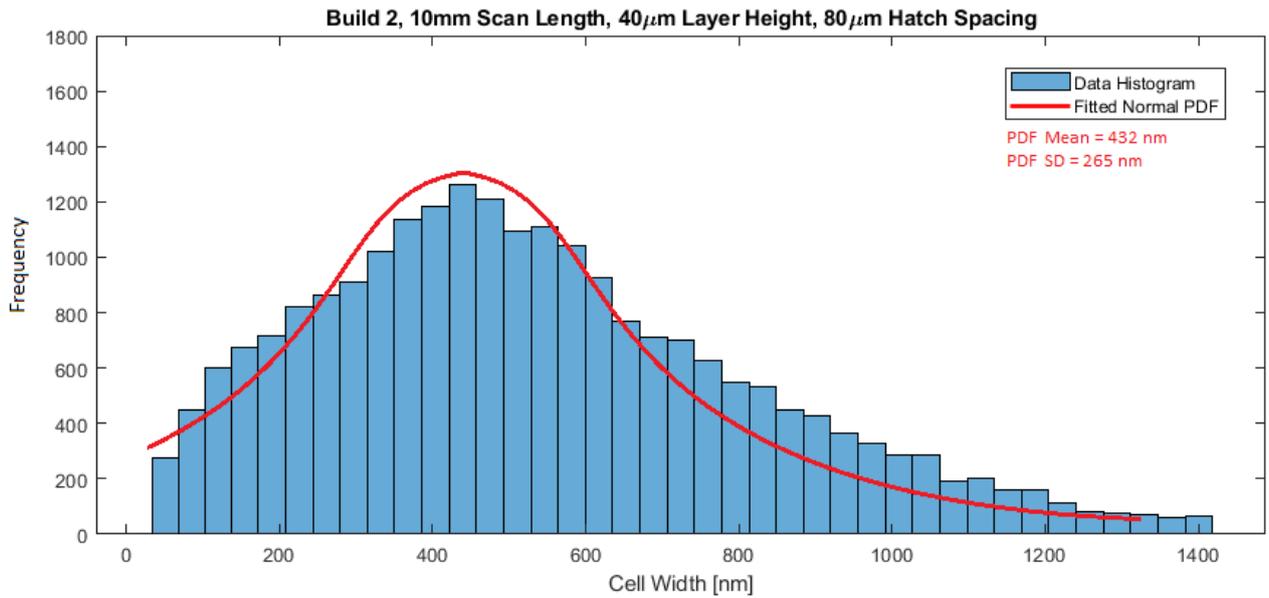


Figure 5.8

The distribution of the cell width measurements, after filtering is applied

Once normal distributions had been found for each cube, the means and standard deviations of the cell sizes could be calculated.

Using the same methodology, a value for cell long-axis was calculated. Instead of recording the shortest value for arm-pair distance, the longest value was recorded. Histogram bins were then reduced to create normal distributions.

There existed extreme variability in these length measurements. Consider the example in Figure 5.9, cropped from the binary image of the Build 4, 5mm cube.

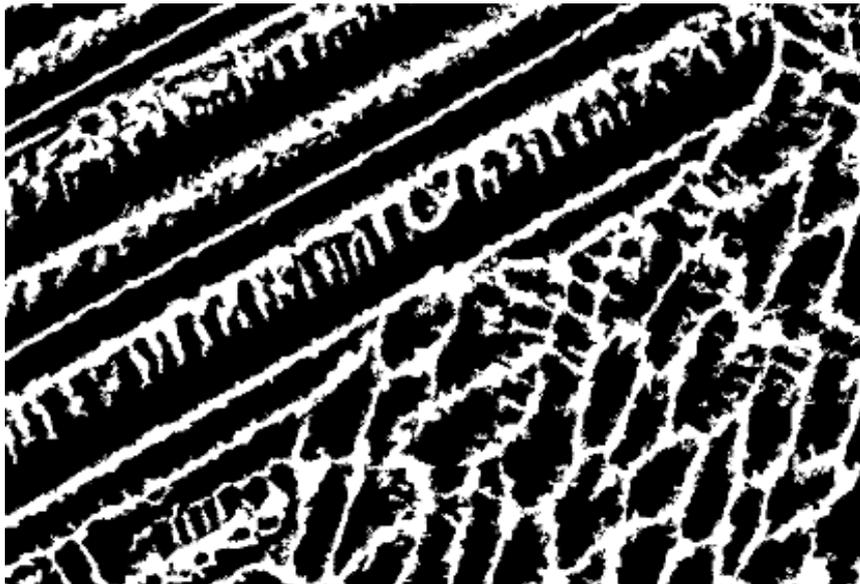


Figure 5.9

An example binarised SEM image, showing extreme variation in cell lengths

It would clearly be incorrect to claim that the cells in the upper-left experienced uninterrupted growth to a much greater extent than the cells in the lower-right. Instead, the plane at which the sample has been cut is parallel to the growth direction of the cells in the upper-left, and perpendicular to that of the cells in the lower-right.

Cell long-axis measurements were therefore not used to estimate actual cell growth, but rather combined with cell widths to estimate cell aspect ratios, which in turn were used to imply growth directions.

After IQR outlier removal and histogram bin reductions were applied to both width measurements and length measurements, the values that remained in each corresponding data set were paired back up (i.e. resulting in only the width-length pairs where neither measurement was filtered out). The aspect ratios were calculated for all the remaining pairs. The standard deviation of these aspect ratios was taken as a useful gauge of the variation of growth directions throughout the imaged surface (i.e. images with high aspect ratio standard deviations likely encompass more growth directions than images with low aspect ratio standard deviation).

5.2.4 - Hardness

Five hardness measurements were taken from each of the 27 surfaces, according to the Vickers method, applying 1kg. The five measurement points were taken from the centre of each sample, with points spread out across a space of approximately $500\mu\text{m} \times 500\mu\text{m}$. Care was taken to make sure that none of the points interacted with porosity or similar mechanical defects.

5.2.5 - EBSD

At the centre of each of the 27 samples, 24 EBSD maps were recorded in a 4×6 grid with 10% overlap area. Each map had a size of $510\mu\text{m} \times 350\mu\text{m}$ ($1025\text{px} \times 705\text{px}$). Stitched together, the total image map had a size of $1.89\text{mm} \times 1.93\text{mm}$ ($3793\text{px} \times 3876\text{px}$). The beam had an accelerating voltage of 30kV, a step size of $0.5\mu\text{m}$, a working distance of 25.8mm, and an average indexing rate (excluding pores) of 99.6%.

Figure 5.10 shows the Y-orientated inverse pole figure, for the same sample as Figure 5.7 (Build 2, 10mm).

The following data was extracted for each grain, from every EBSD image:

- Grain size
- Grain orientation
- Intra-granular misorientation, known as Gran Average Misorientation (GAM)
- Grain aspect ratio
- Inter-granular misorientation

By mapping the grain sizes, aspect ratios, and orientations, it was possible to create heatmaps to determine the global anisotropy within the samples.

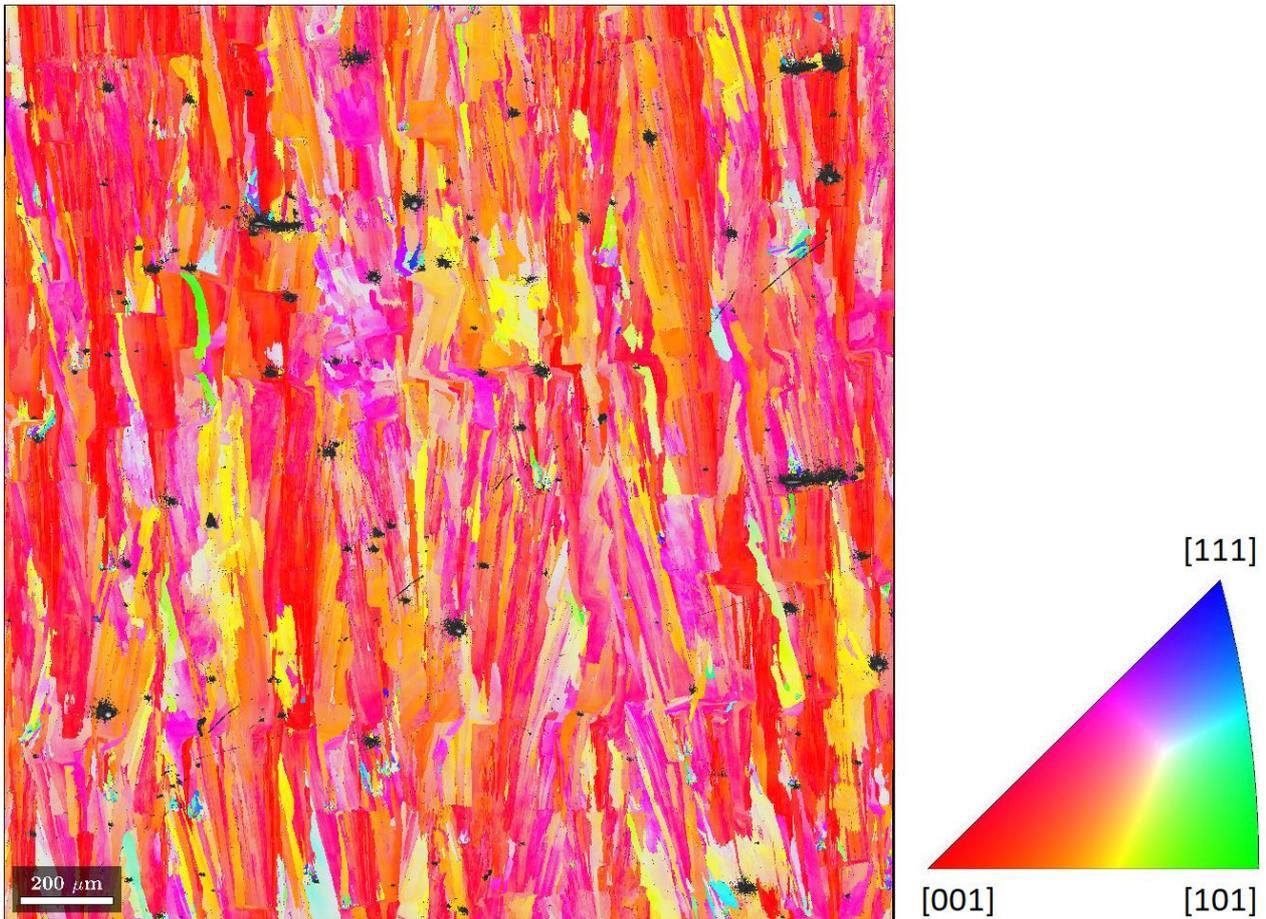


Figure 5.10

An example EBSD image, showing the grains of Build 2 (40 μm layer height, 80 μm hatch spacing), 10mm cube

5.3 - Data Analysis

For each of the 27 samples, an extremely large amount of data existed to be analysed.

Table 5.2 summarises all these different parameters and statistics, and also notes if each is a global data type (applies to the entire sample), or a local type (and therefore has an associated mean, standard deviation, histogram, etc.).

Before data analysis began, a set of hypotheses were drawn up, to be systematically tested for each of the aforementioned data types. A fundamentals-first approach was taken, beginning with the effects of the build parameters on the positioning of the solidification within the GR diagram, as shown in Figure 5.11.

Table 5.2
The parameters extracted for each cube

Data Type	Data	Global	Local
Build Parameters	Scan Length (cube size)	✓	
	Hatch spacing	✓	
	Layer Height	✓	
	Dimensionless Numbers	✓	
Thermal Camera	Peak Temperature		✓
	Mean Temperature		✓
	Melt Pool Length		✓
	Melt Pool Width		✓
	Melt Pool Area		✓
	Melt Pool Aspect Ratio		✓
	Thermal Gradient		✓
	Cooling Rates		✓
	Spatter Quantity	✓	
Optical Microscope	Sample Density	✓	
	Number of Pores	✓	
	Pore Areas		✓
	Pore Circularity		✓
SEM	Cell Size		✓
	Cell Aspect Ratio		✓
Hardness Testing	Hardness		✓
EBSD	Grain size		✓
	Grain Aspect Ratio		✓
	Grain Orientation		✓
	Inter-Granular Misorientation		✓
	Grain Average Misorientation		✓

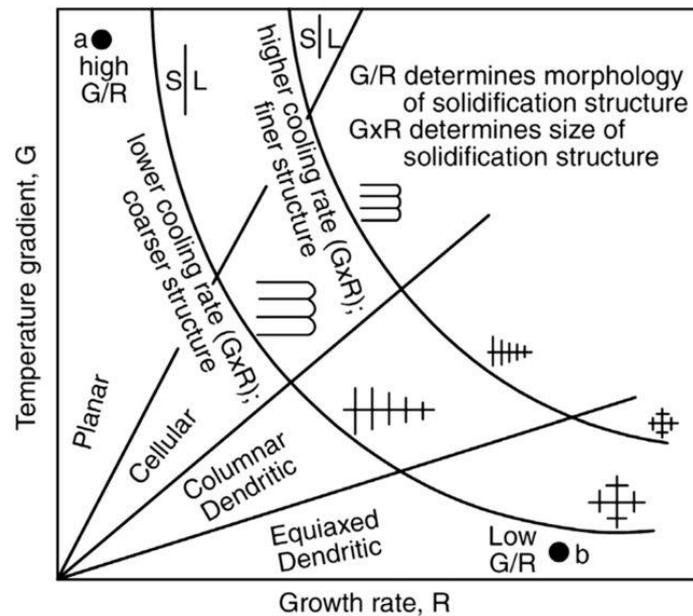


Figure 5.11

The GR solidification diagram, showing how Growth Rate and Thermal Gradient determine microstructure size and morphology [115]

It was reckoned that the effects on Growth Rate and Thermal Gradient would be as follows:

1. As scan length increases, return time will increase. Therefore, the solidified material to the side of the melt pool will be cooler. Therefore, thermal gradient will increase. Growth Rate will be unaffected. On the GR diagram, solidification shifts upwards: from large and equiaxed towards fine and cellular.
2. As hatch spacing increases, the amount of solidified material that is being re-melted on the next hatch will decrease. Therefore, each vector will be encountering less heated solid, and more ambient-temperature powder. Therefore, thermal conduction will decrease, resulting in an increased thermal gradient, and a decreased growth rate. On the GR diagram, solidification shifts upwards and leftwards: from equiaxed towards cellular, with an undetermined effect on grain size.
3. As layer height increases, the amount of ambient-temperature powder that the melt pool contacts will increase. Therefore, thermal conduction will decrease, resulting in an increased thermal gradient, and a decreased growth rate. On the GR diagram, solidification shifts upwards and leftwards: from equiaxed towards cellular, with an undetermined effect on grain size.

Based on these fundamentals, the predicted effects on every data point from each source was considered.

1. Thermal camera predictions:
 - As scan length increases, melt pools will become shorter, narrower, and exhibit lower peak and mean temperatures, due to their cooler surroundings.

- As hatch spacing increases, melt pools will become longer, wider, and exhibit higher peak and mean temperatures, due to lower thermal conduction.
 - As layer height increases, melt pools will become longer, wider, and exhibit higher peak and mean temperatures, due to lower thermal conduction.
2. Optical microscope predictions:
 - At higher volumetric energy densities (small scan length, hatch spacing, and layer height), some keyhole porosity will cause lower density.
 - At lower volumetric energy densities (larger scan length, hatch spacing, and layer height), large amounts of LoF porosity will cause lower density.
 - At medium volumetric energy densities, higher densities will be achieved.
 3. SEM predictions:
 - As thermal gradients increase (i.e. as scan length, hatch spacing, and layer height increase), more rapid solidification will lead to smaller cell sizes, and greater local anisotropy.
 - As thermal gradients decrease (i.e. as scan length, hatch spacing, and layer height decrease), slower solidification will lead to larger cell sizes, and less local anisotropy.
 4. Hardness test predictions:
 - As greater thermal gradients lead to small cell sizes, hardness will increase.
 - As lower thermal gradients lead to larger cell sizes, hardness will decrease.
 5. EBSD predictions:
 - As thermal gradients increase (i.e. as scan length, hatch spacing, and layer height increase), more rapid solidification will lead to smaller grain sizes, and greater global anisotropy, as demonstrated by Hagenlocher [322].
 - As thermal gradients decrease (i.e. as scan length, hatch spacing, and layer height decrease), slower solidification will lead to larger grain sizes, and less global anisotropy.

These predictions will now be systematically tested. In all of the charts presented in this subsection, error bars, where presented, refer to the standard deviations of all measurements.

In cases where the data found in the charts was considered to not add insight, the charts and discussion have been moved to the Appendix.

5.3.1 - Thermal Camera Analysis

In each sequence of images, the melt pools that corresponded to each type of square were identified and separated. The intensity-temperature conversion was applied, and each melt pool image was thresholded at 1380°C (the liquidus temperature of ABD 900-AM). The length, width, peak temperatures, and other statistics of the melt pools were extracted. For each of the 108 recordings (9 builds × 3 cubes × 4 layers), outliers were identified using the inter-quartile range (IQR) method.

Melt pools were removed from the dataset which had a width, length, or peak temperature outside the bounds of $Q1 - (1.5 \times IQR)$, or $Q3 + (1.5 \times IQR)$. After these outliers were removed, between 1750 and 7582 melt pools remained, for each of the 27 cubes.

5.3.1.1 - Melt Pool Length

After each melt pool was identified using thresholding, it was analysed using MATLAB's inbuilt REGIONPROPS function. This function fitted an ellipse to the melt pool, and measured the major axis of the ellipse. These measurements were taken as the melt pools' lengths, and are shown in Figure 5.13. In cases where a long, narrow tail of above-threshold pixels extends from the tail of the melt pool, the fitted-ellipse may not include this details, therefore making the length appear shorter than it actually is. It was reasoned that the large numbers of melt pools being analysed would compensate for these possible errors.

As can be seen from row (1) of the charts, melt pool length very consistently decreases as scan length increases. These findings match the expectations that longer scan lengths will result in greater thermal conduction, and therefore smaller melt pools.

As can be seen from row (2) of the charts, the effects of hatch spacing on scan length are inconsistent. It was expected that larger hatch spacings would lead to lower thermal conduction, and therefore larger melt pools. In many cases, the opposite is true: pool length decreasing with hatch spacing.

One possible reason for this is that, at lower hatch spacings, the material (both powder and solid) in a melt pool's path has been preheated by the previous hatch, to a greater extent. As hatch spacing increases, and therefore preheating decreases, melt pools may therefore encompass less overall heat, leading to lower melt pool lengths.

This alternative theory of the effects of hatch spacing (i.e. greater hatch spacing means lower preheating, which means lower heat), is contrary to the primary theory that has been suggested thus far (i.e. greater hatch spacing means lower thermal conduction, which means higher heat). The validity of this alternative theory will be kept in mind in all further analysis in this section.

As can be seen from row (3) of the charts, melt pool length consistently rises as layer height is increased. This trend matches the expectation that at greater layer heights thermal conduction is reduced, and therefore melt pool length increases.

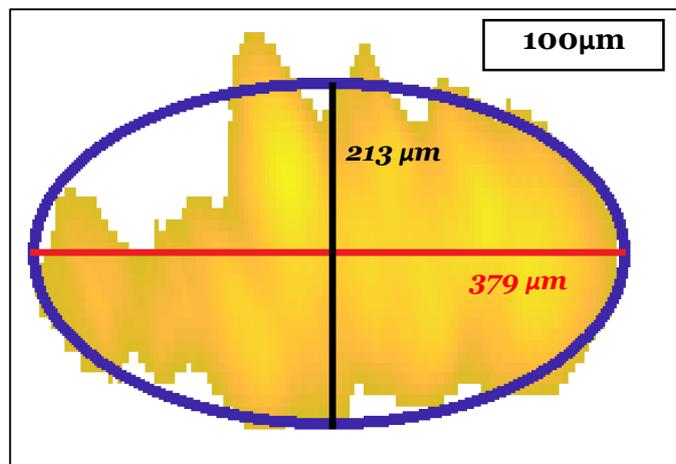


Figure 5.12

A melt pool from Build 9, showing the pixels above threshold temperature (yellow/orange), the fitted ellipse (blue), the major axis (red), and the minor axis (black)

Effect of Machine Parameters on Melt Pool Length

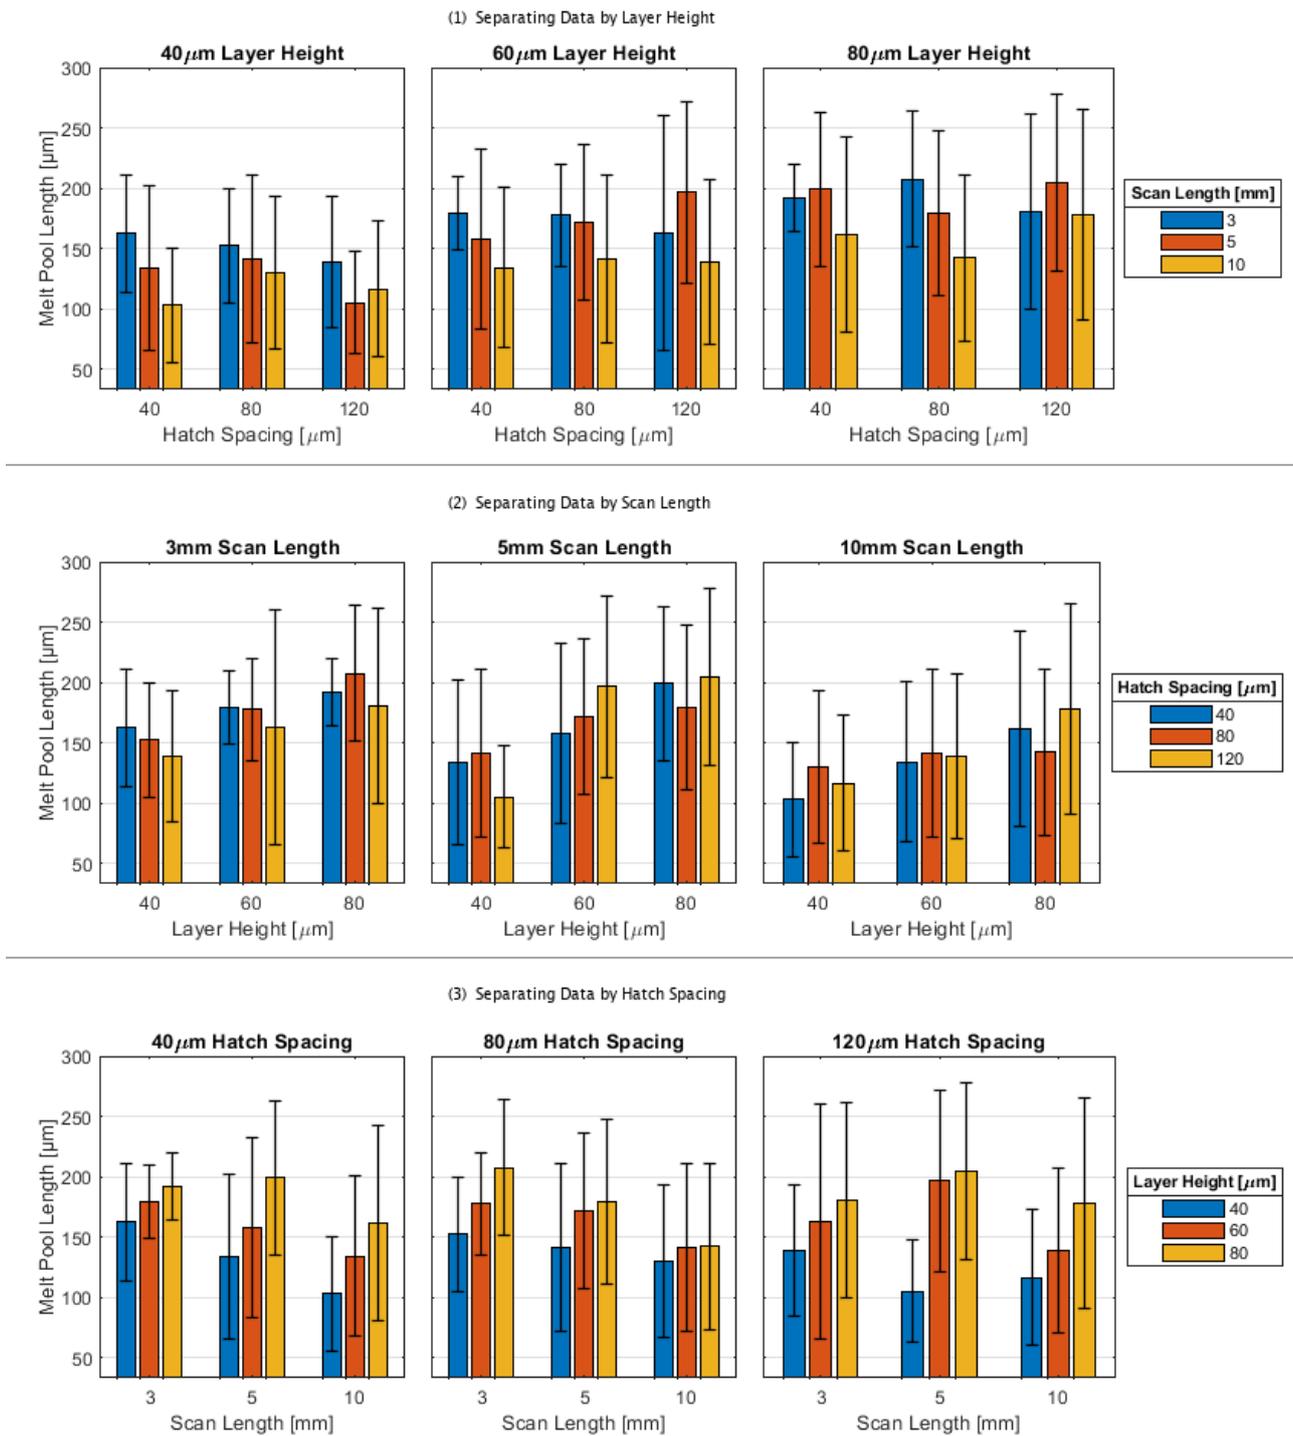


Figure 5.13

The effects of machine parameters on melt pool length

5.3.1.2 - Melt Pool Area

The number of pixels in each melt pool was recorded, and used to calculate the area of the melt pool, as shown in Figure 5.14.

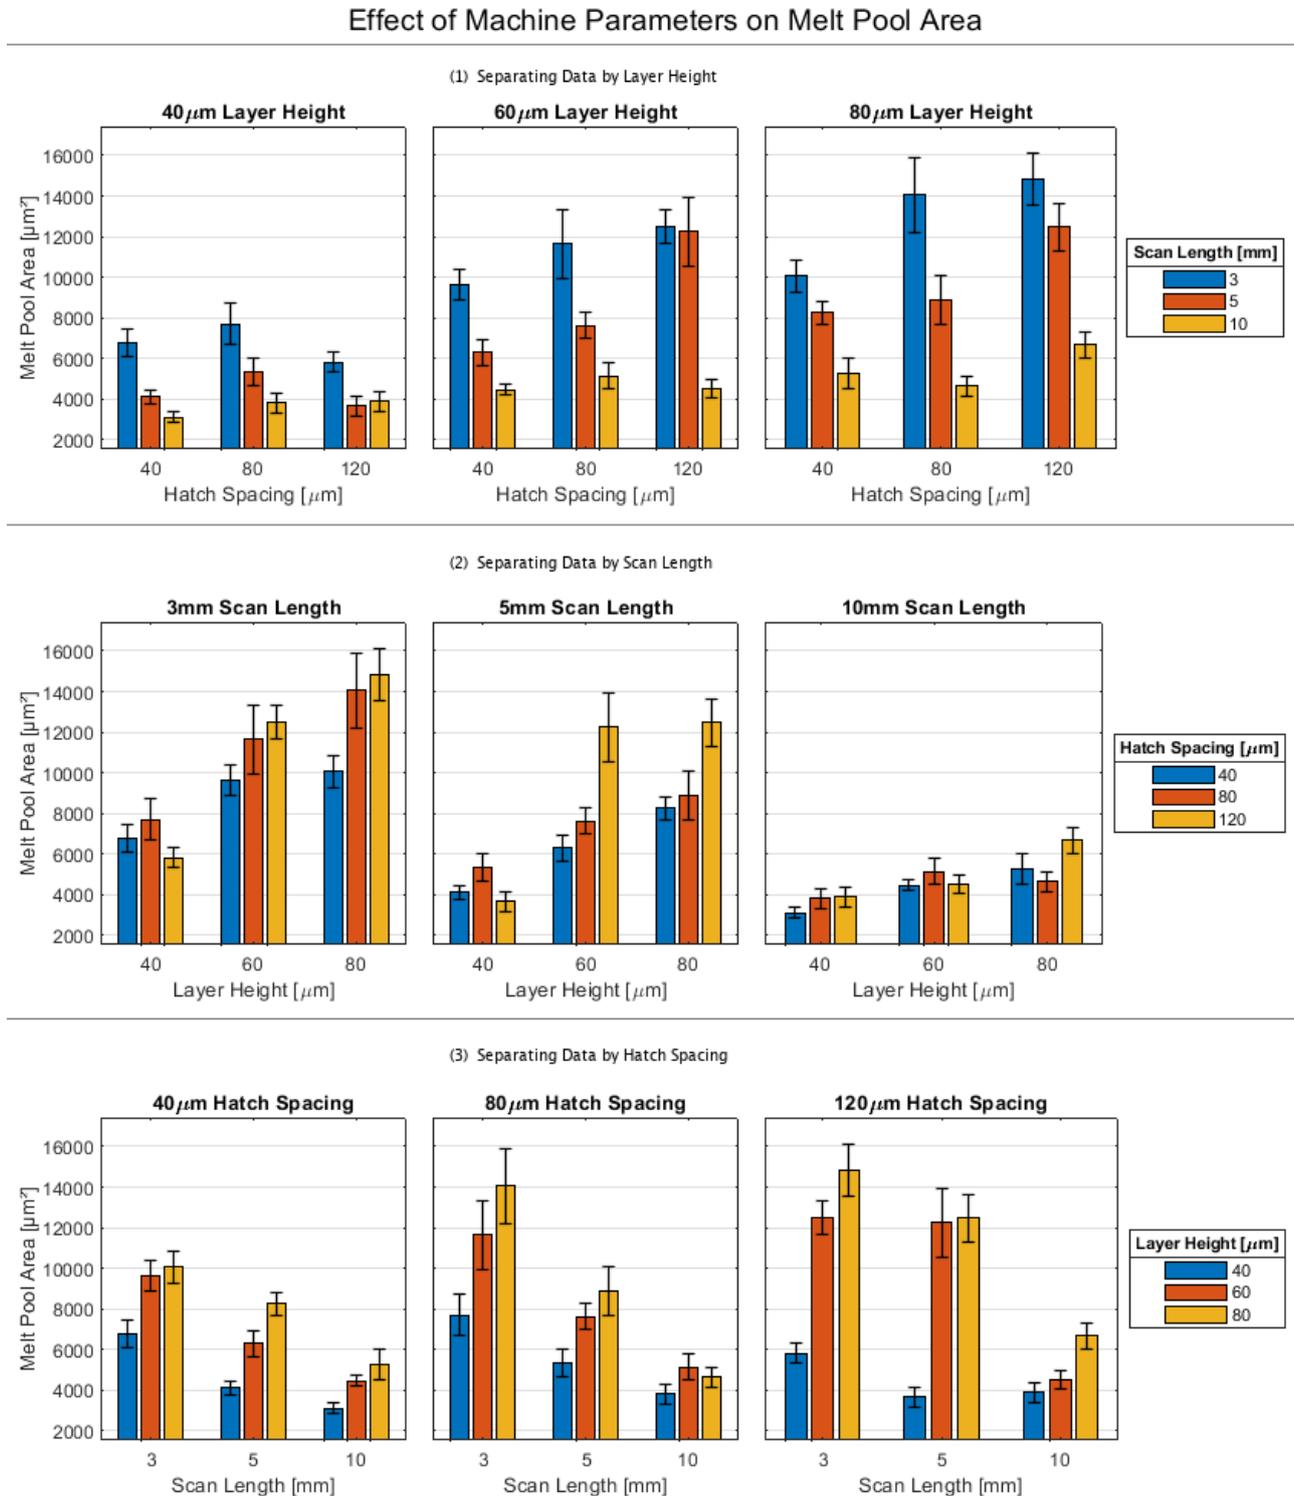


Figure 5.14

The effects of machine parameters on melt pool area

As can be seen in row (1), melt pool area consistently decreases as scan length increases. This trend matches the expectation that longer scan lengths lead to increased thermal conduction, which in turn leads to small melt pools.

As can be seen in row (2), melt pool area generally increases as hatch spacing increases. This trend largely matches the primary hatch spacing theory, which suggests that as hatch spacing increases, thermal conduction decreases, and therefore melt pool area increases.

As can be seen in row (3), melt pool area consistently increases as layer height increases. This trend matches the expectation that as layer height increases, thermal conduction will decrease, and therefore melt pool area will increase.

5.3.1.3 - Melt Pool Aspect Ratio

The aspect ratio of each melt pool was calculated by dividing the length by the width (both of which were determined by the fitted ellipse method). The aspect ratios are shown in Figure 5.15.

As previously discussed, the expected trends for melt pool widths and lengths held quite consistently (increased scan length led to decreased length and width; increased hatch spacing generally led to increased length and width; increased layer height led to increased length and width). What the charts in Figure 5.15 reveal, is that the rates at which these parameters change the widths, and the rates at which they change the lengths, vary significantly.

As can be seen from row (1), when melt pools shrink (in both length and width) as scan length increases, their widths decrease more rapidly, leading to pools with greater aspect ratios. As can be seen from row (2), when melt pools grow as hatch spacing increases, their widths increase more rapidly, leading to pools with lower aspect ratios. As can be seen from row (3), when melt pools grow as layer height increases, their lengths grow more rapidly, leading to pools with greater aspect ratios.

The reasons that increases and decreases in scan lengths and hatch spacings affect width more than length, while increases and decreases in layer height affect length more than width, are reckoned to be due to the differences in the manner of heat conduction out of the melt pool. Increasing or decreasing either the scan length or hatch spacing, will primarily alter the phase and temperature of the material laterally adjacent to the melt pool. Therefore, these parameters affect pool width to a greater extent than they affect length. Increasing or decreasing layer height, however, will primarily alter the phase and temperature of the material below the melt pool, changing the duration required for solidification. Therefore, this parameter affects pool length to a greater extent than it affects width.

Effect of Machine Parameters on Melt Pool Aspect Ratio

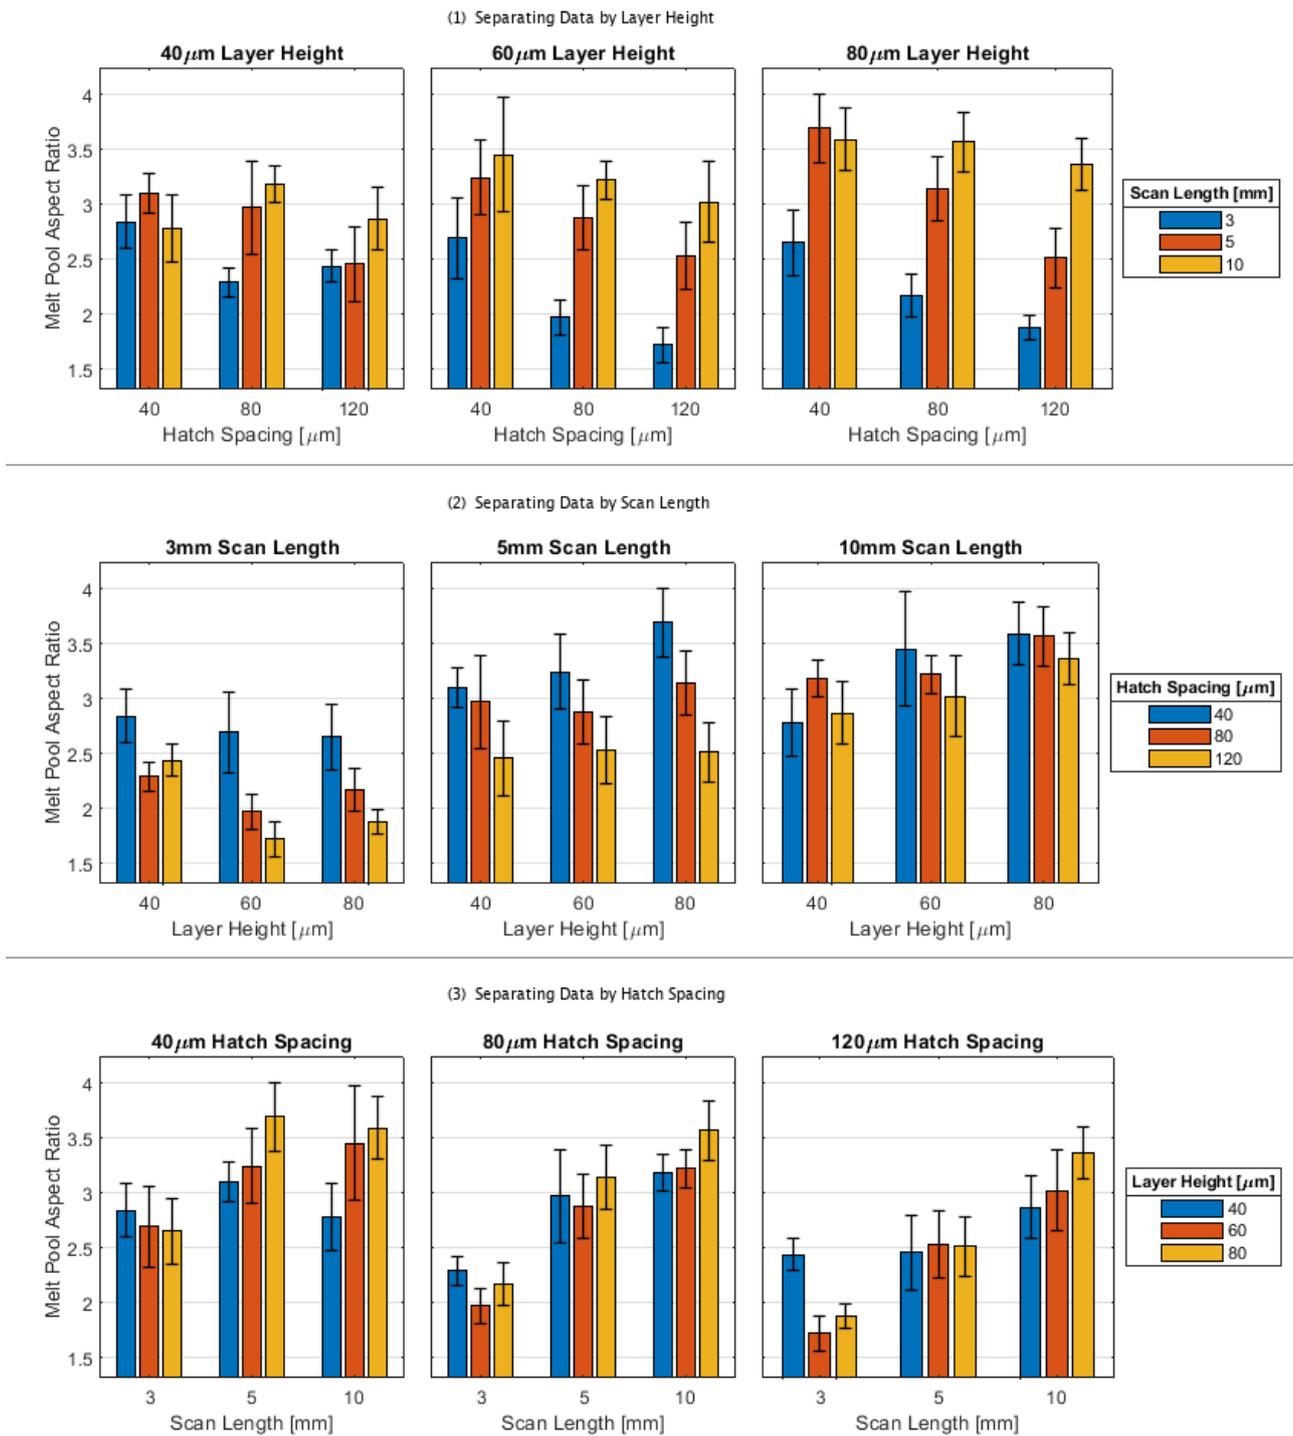


Figure 5.15

The effects of machine parameters on melt pool aspect ratio

5.3.1.4 - Melt Pool Tail Thermal Gradient

The temperature profile of the long axis of the fitted ellipse was extracted for each melt pool. It should be noted that this axis did not necessarily contain the pixel of peak temperature, but it was reckoned that it would still serve as a suitable measure of thermal gradients.

The mean thermal gradient of the tail of the axis (i.e. from its pixel of maximum temperature to its end) was recorded, and plotted against machine parameters in Figure 5.16.

As can be seen in row (1), the thermal gradient quite consistently increases with increasing scan length. This trend matches the expectation that increasing scan length will increase thermal conduction, and therefore increase the rate at which temperatures fall, moving from the centre of the melt pool outwards. The trend does not hold in the case of the 120 μ m hatch cubes: in each of these sets, thermal gradient falls as scan length increases from 3mm to 5mm, then rises again as the scan length is increased to 10mm. No theory was suggested as to why this might be the case.

As can be seen in row (2), the effect of hatch spacing on thermal gradient is inconsistent. For the 3mm cubes, thermal gradient increases as hatch spacing increases. This trend matches the primary theory of hatch spacing, whereby increases in hatch spacing lead to lower thermal conduction, which in turn leads to greater thermal gradients. For the 5mm and 10mm cubes, however, no clear trend exists. It could be the case that, as scan length increases, the alternative theory of hatch spacing (i.e. that increased hatch spacing leads to less preheating, and therefore lower heat accumulation, and lower thermal gradients) may become more dominant.

As can be seen in row (3), thermal gradient very consistently decreases as layer height increases. Initially, this trend appears counterintuitive. As has been demonstrated in previous charts, increasing the layer height leads to lower thermal conduction, due to the lower conductivity of powder vs solid, and therefore greater heat accumulation and peak temperatures. Therefore, greater thermal gradients might be expected. It is the fact that this lower conduction leads to an increased melt pool size, however, that counteracts the greater temperatures. It appears that as layer height increases, both the temperature difference and rear length of the melt pool increase, with the latter increasing more rapidly than the former, in a manner that makes the thermal gradient decrease.

Effect of Machine Parameters on Tail Thermal Gradient

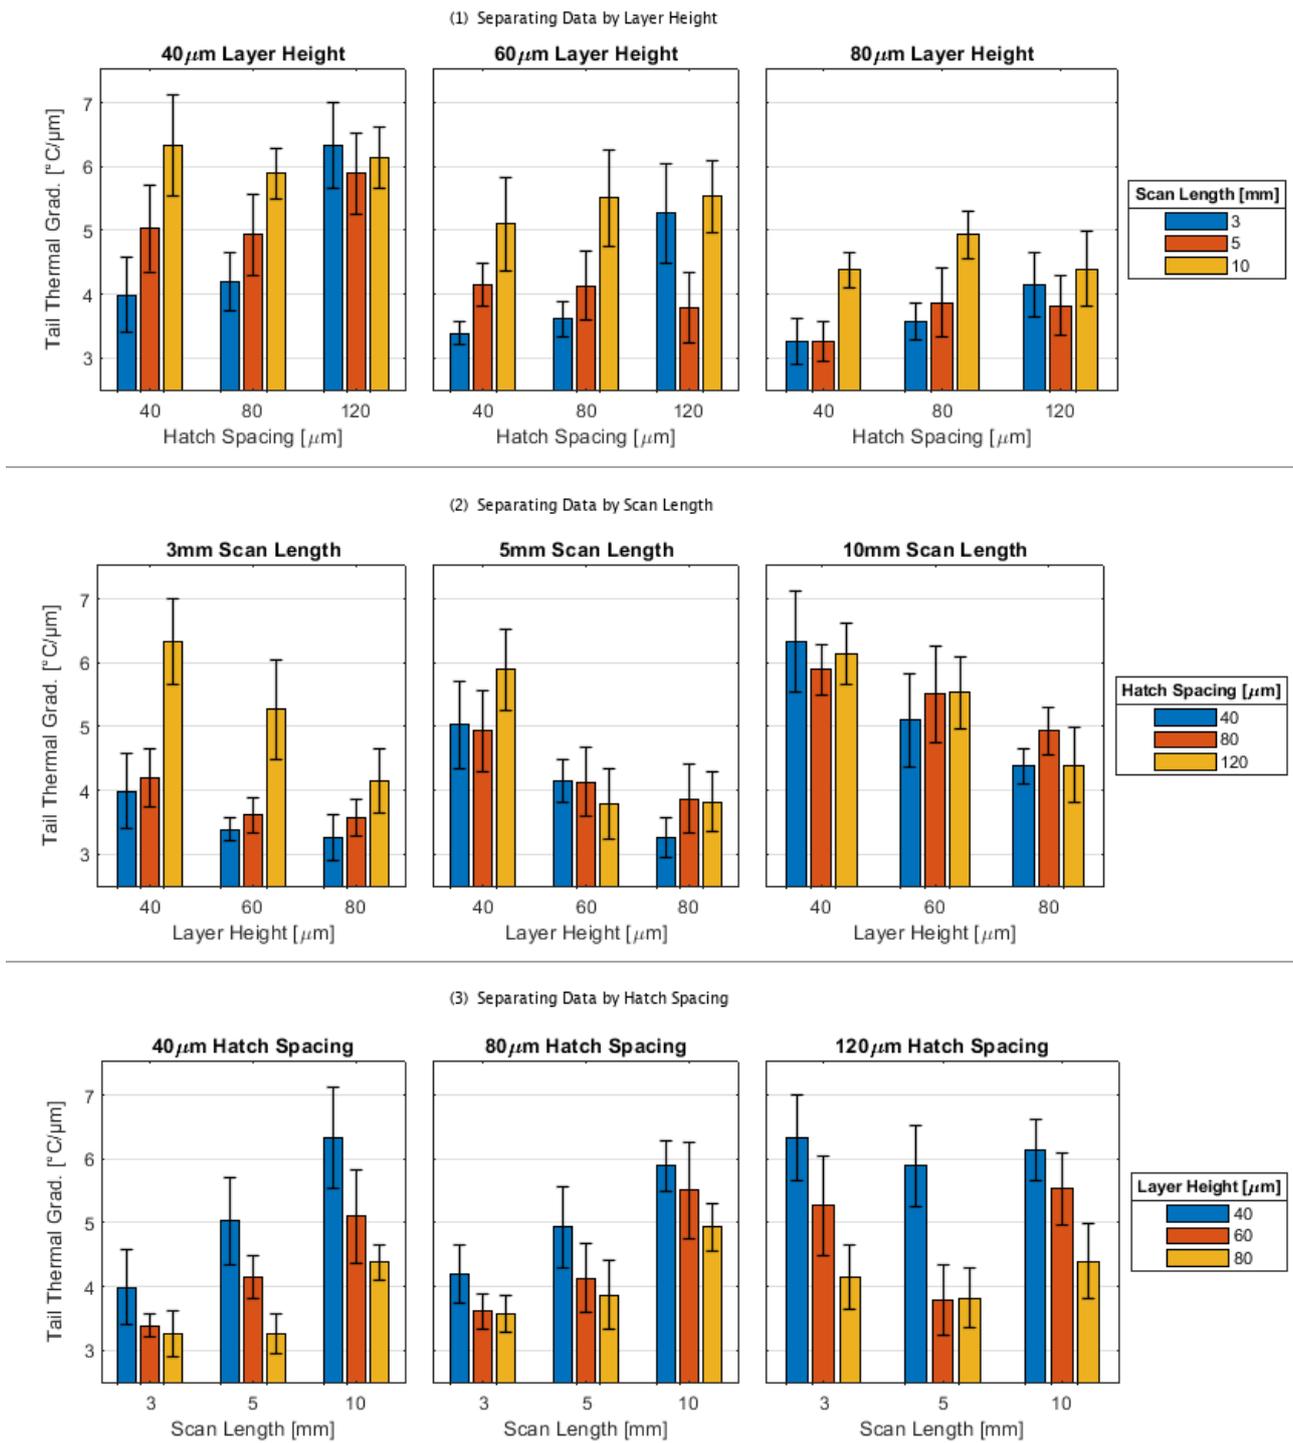


Figure 5.16

The effects of machine parameters on melt pool tail thermal gradient

5.3.2 - Optical Microscope Analysis

Each optical microscopy image was analysed using custom-built MATLAB scripts. For each image, the outer edges were cropped so that only the surface sample was included. Then a histogram of the intensities of all remaining pixels was extracted, and the gap between the lower peak (referring to pores) and the upper peak (referring to solid material) was analysed. The threshold value, used to differentiate between these two, was set as the midpoint between these peaks.

Figure 5.17 shows an example histogram, with the threshold intensity (133) marked in red. The peak at 255 is due to the microscope camera's pixels becoming oversaturated in regions of high reflection.

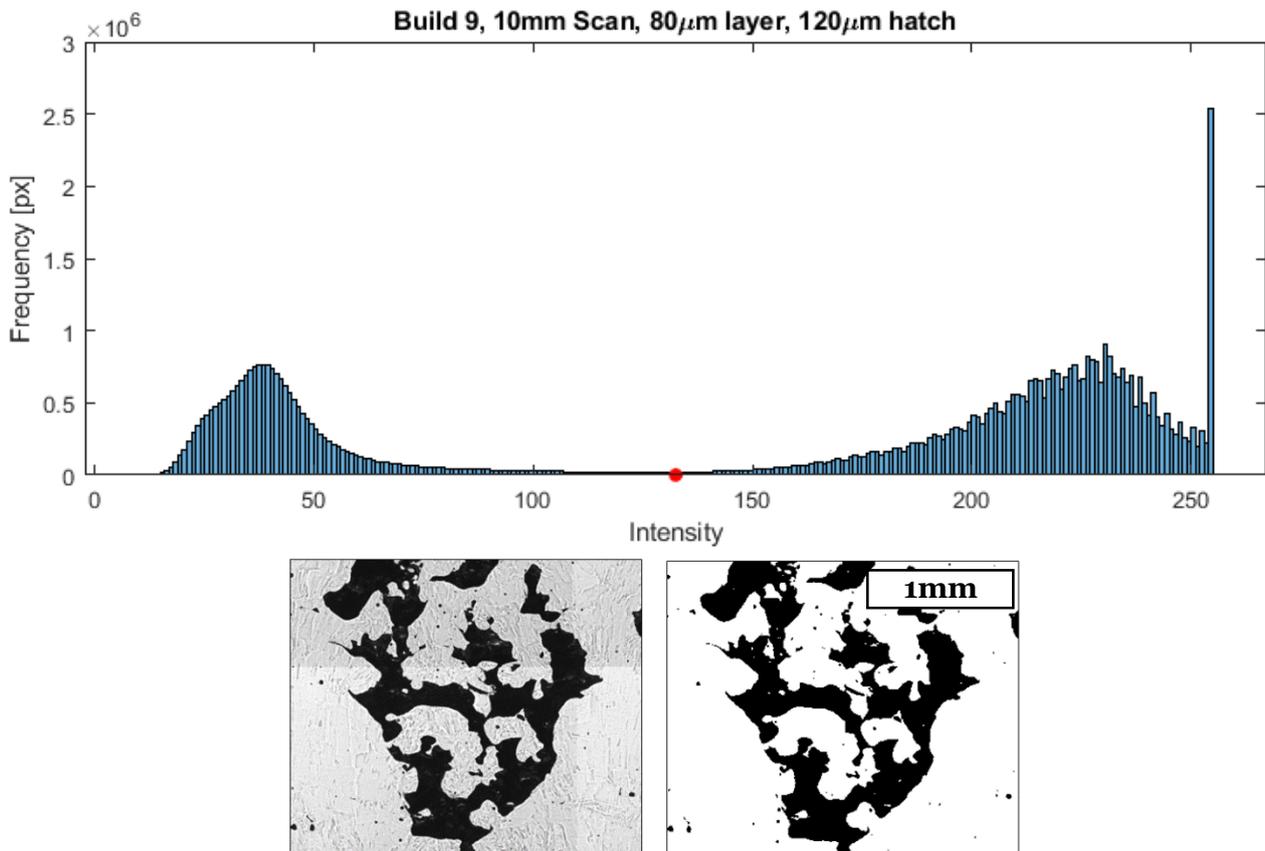


Figure 5.17

A demonstration of the method used to distinguish pore pixels from solid pixels in optical microscope images: the histogram (top), a section of the raw image (bottom left), and the result (bottom right)

5.3.2.1 - Porosity

After thresholding was applied to each image, the number of pixels which corresponded to pores, and the number that corresponded to solid, were counted. Figure 5.18 shows the number of pore pixels of each cube, as a percentage of the total number of pixels. As can be seen from the charts, a strong trend was found to exist between the volumetric energy density of each cube, and its porosity. These findings very strongly suggest that all pores in all samples are due to LoF, rather than keyholing, because if keyholing was present then porosity would be expected to increase at higher energy densities too.

Effect of Machine Parameters on Porosity

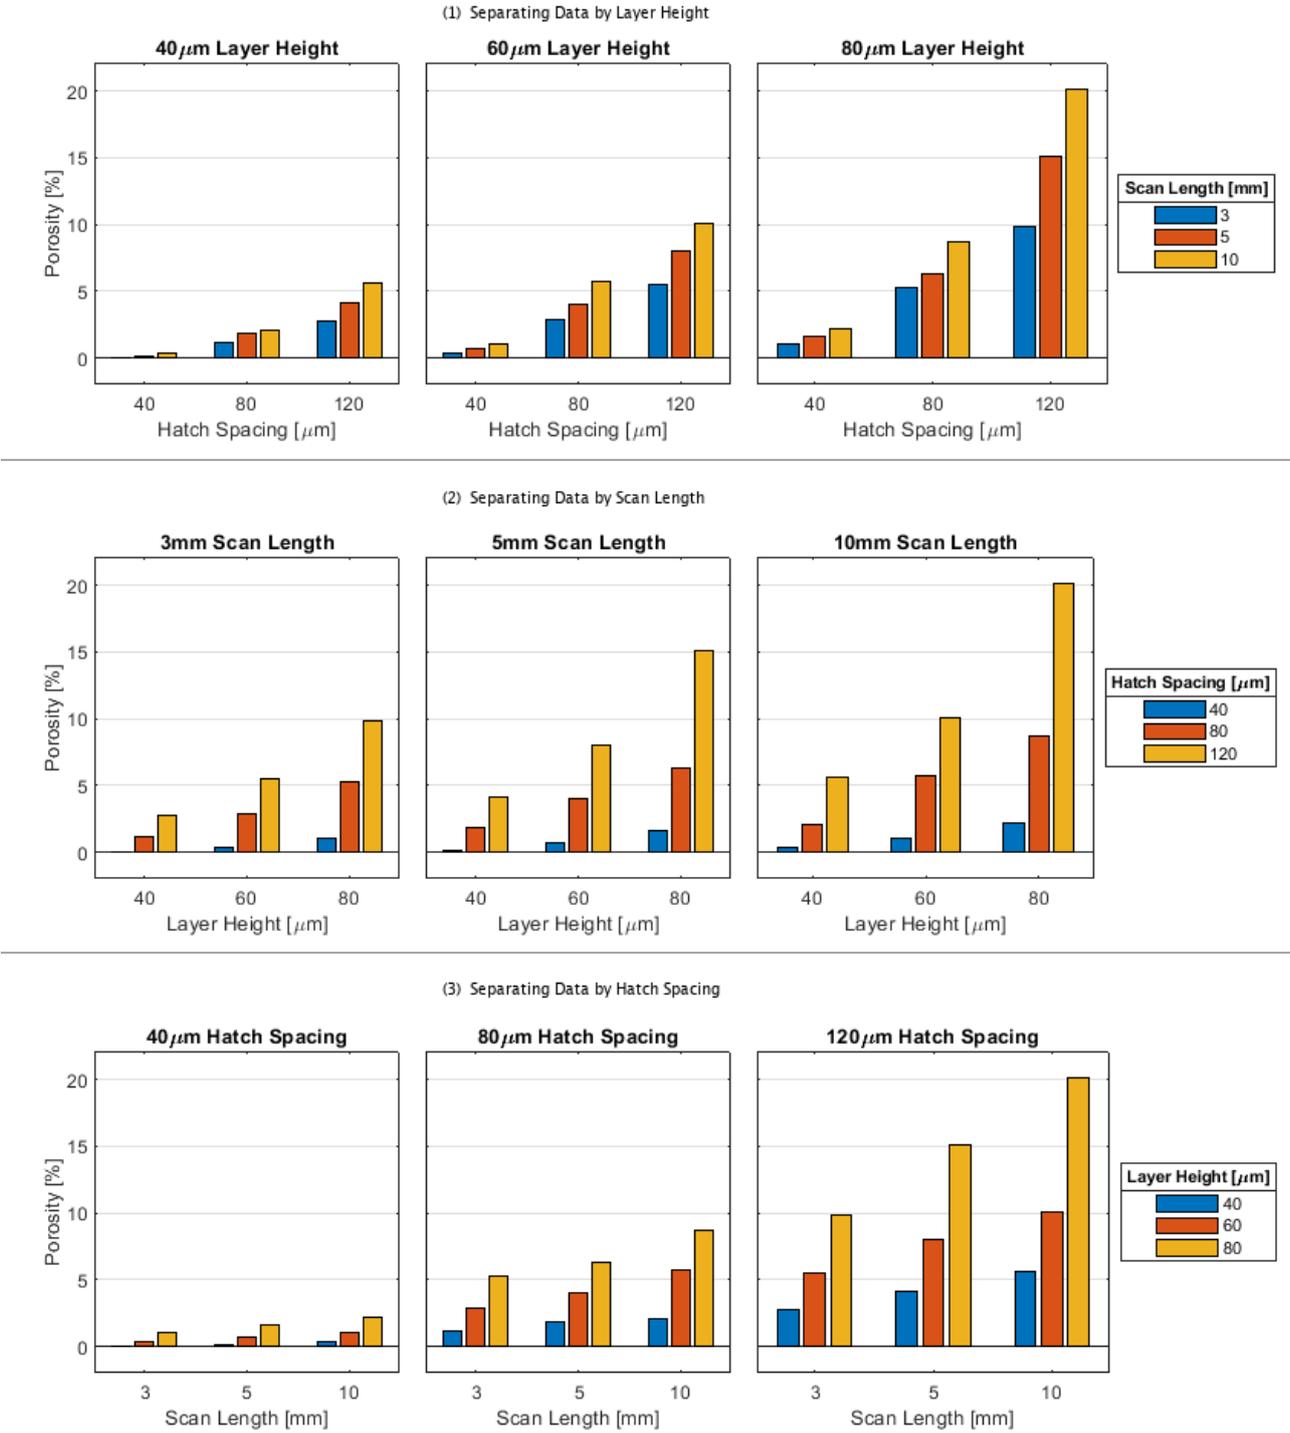


Figure 5.18
The effects of machine parameters on porosity

5.3.2.2 - Pore Area

The area of each pore, in each optical image was counted, and the mean pore area was plotted as shown in Figure 5.19. Again, strong linear trends exist: as volumetric energy density decreases (and porosity increases, as per the previous set of charts), the mean area of the pores increases.

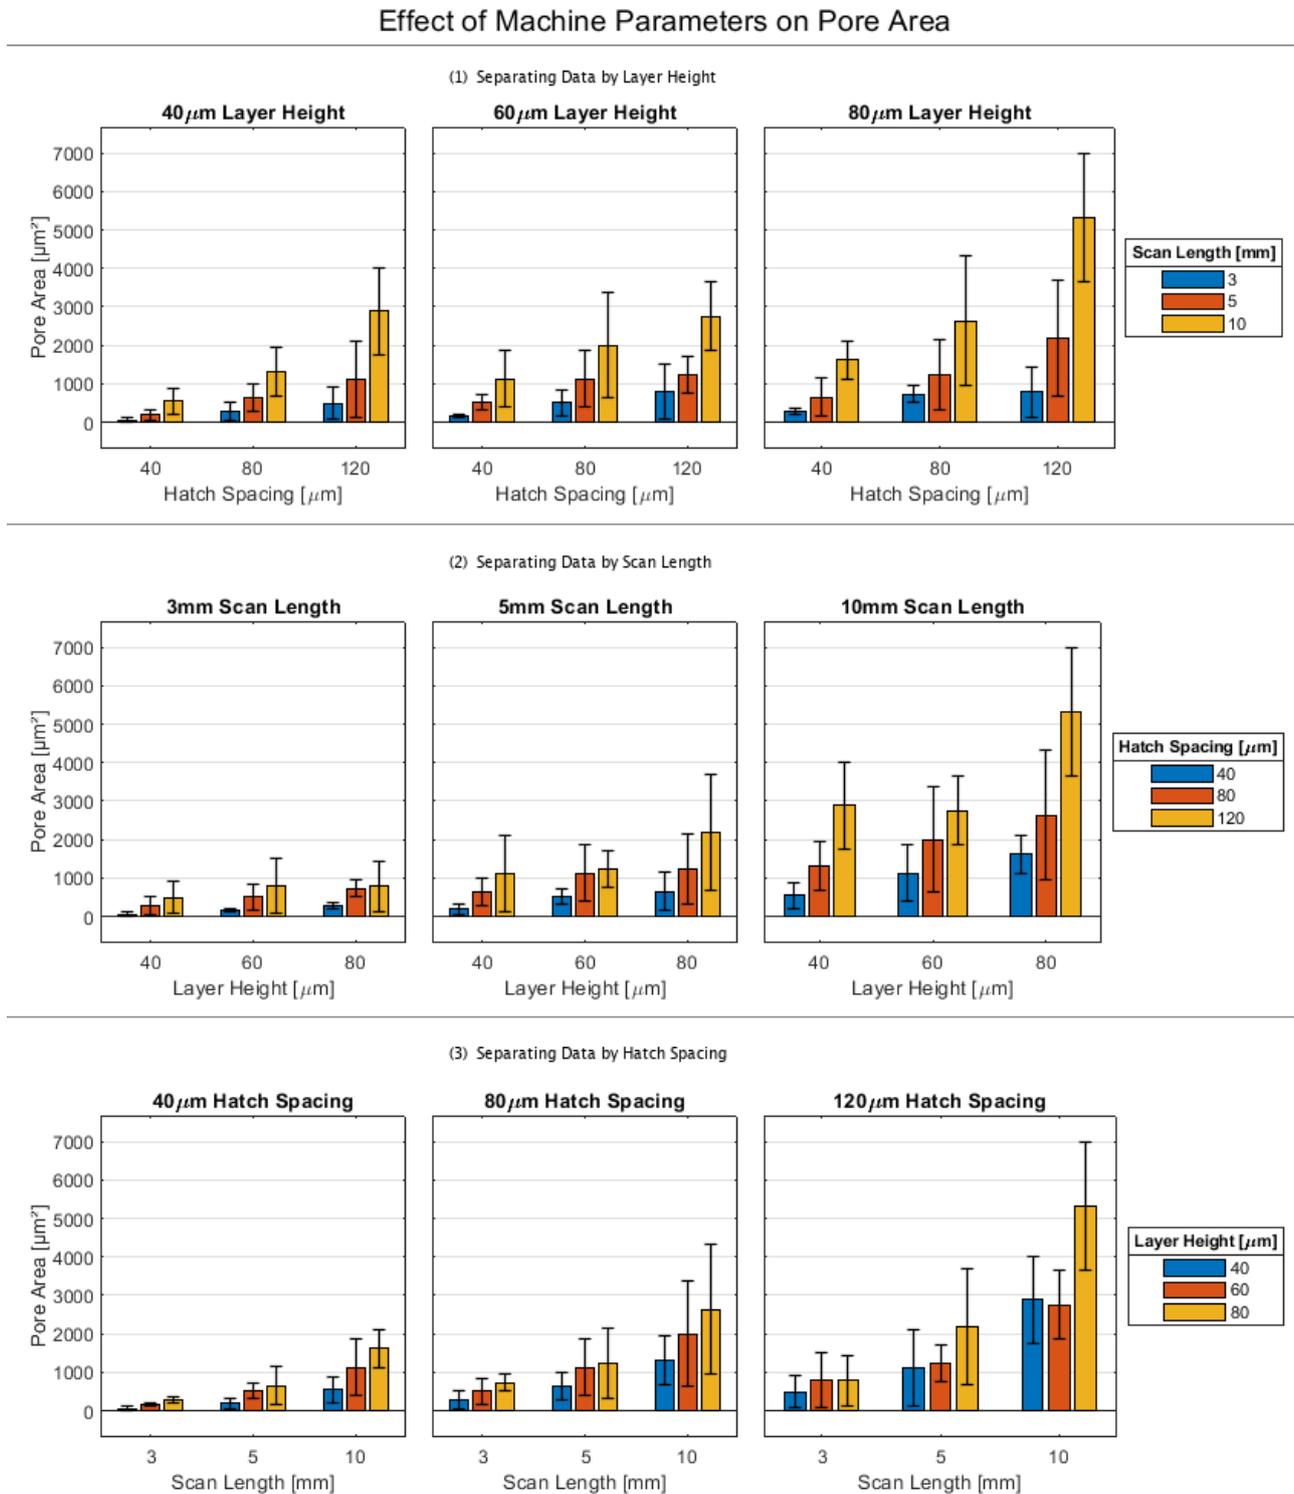


Figure 5.19

The effects of machine parameters on mean pore area

5.3.3 - SEM Analysis

The 27 SEM images were collected and analysed according to the methodology described in a previous subsection. Plotting the mean widths of the cells yielded Figure 5.20.

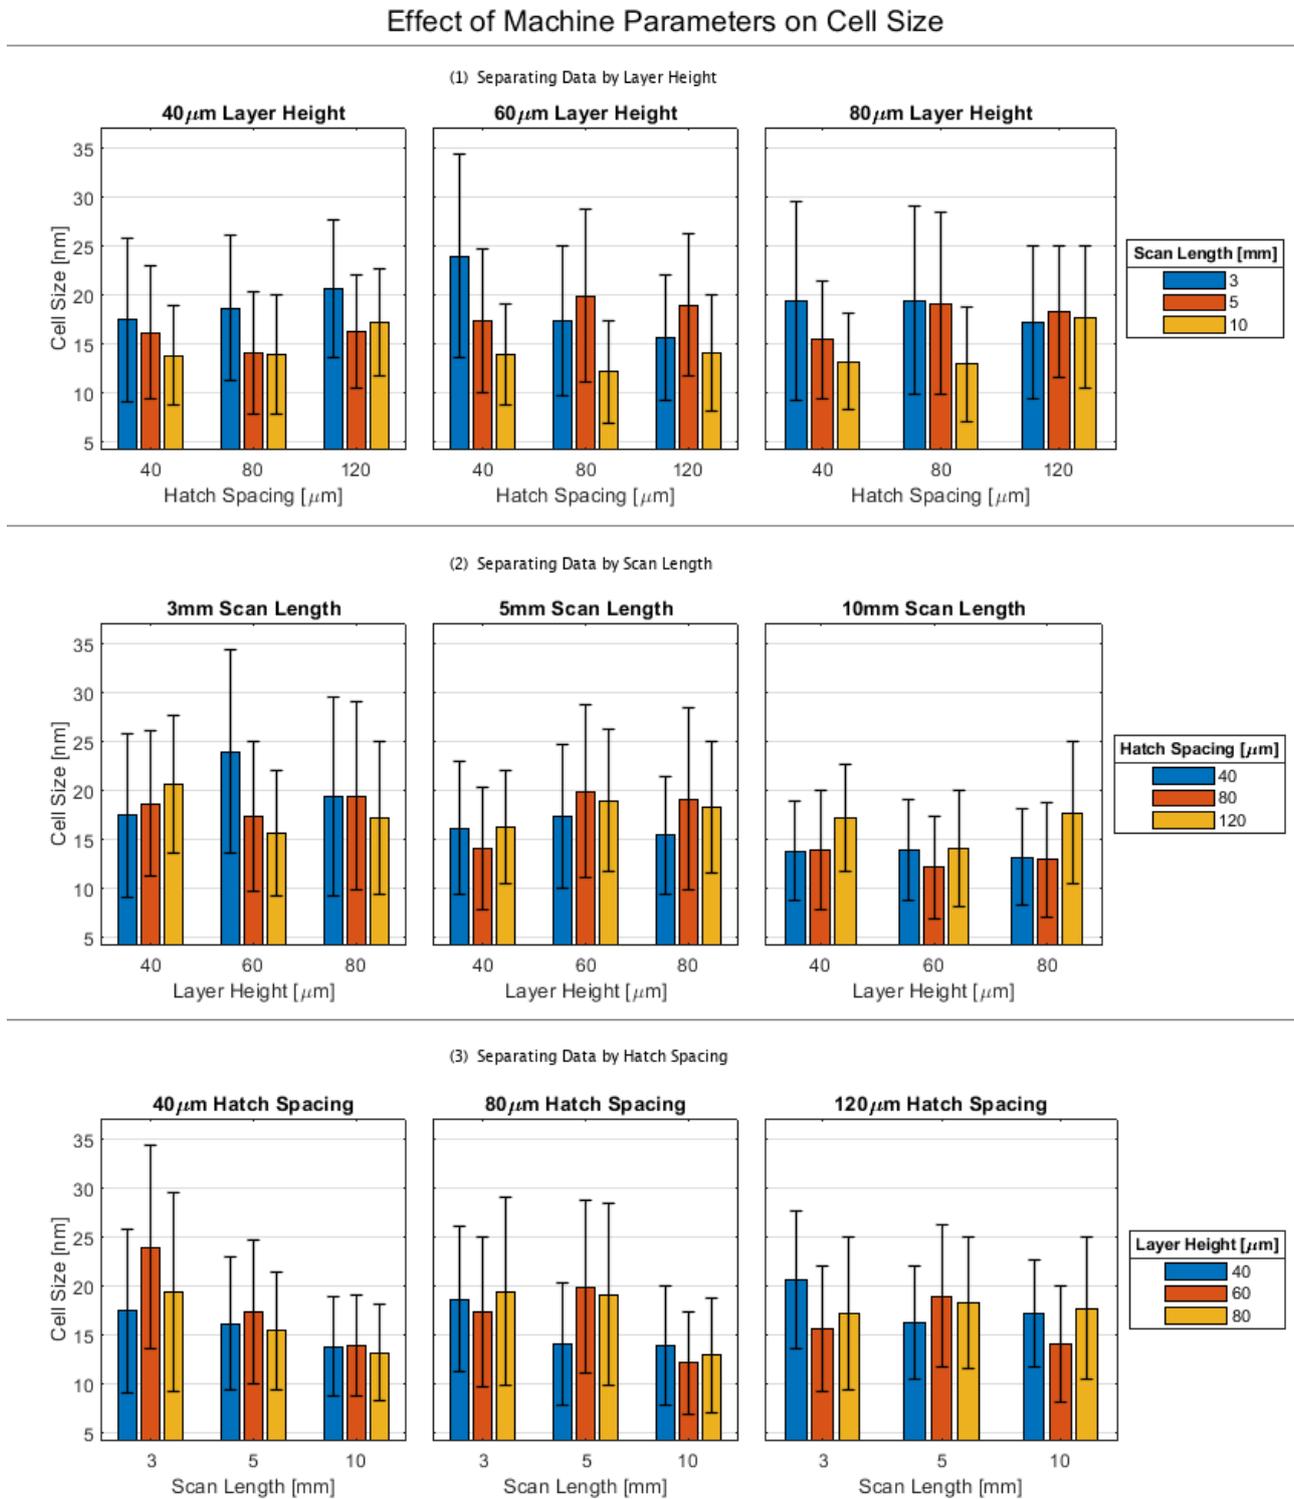


Figure 5.20
The effects of machine parameters on cell size

As can be seen from row (1), cell sizes generally tend to decrease as scan length increases, although the trend is not entirely consistent. This general trend matches the expectation that, as scan length increases and therefore thermal gradients increase, more rapid solidification leads to smaller microstructural features.

As can be seen from row (2), no clear trend exists between hatch spacing and cell size. According to the primary theory of hatch spacing, an increase in hatch spacing will lead to a decrease in thermal conduction, which will lead to a decrease in growth rate, and therefore the larger microstructure associated with slower solidification. According to the alternative theory of hatch spacing, an increase in hatch spacing will lead to less melt pool overlap, and therefore lower heat accumulation, which will lead to lower thermal gradients, with an indeterminate effect on growth rate. Whichever of these cases is more dominant, it is clear in some of these cases cell size appears to significantly increase with increased hatch spacing, in some cases cell size appears to significantly decrease, in many cases the direction of change is unclear, and in all cases the errors are very large, compared to the cell size changes.

As can be seen from row (3), no clear trend exists between hatch spacing and layer height. It was expected that, as layer height increases, lower thermal conduction would lead to slower solidification, which would in turn lead to a larger microstructure. The results do not show this to be true, with the directions of change in cell size being both inconsistent, and small in comparison to the errors.

Plotting cell size against thermal gradient (as measured by the thermal camera), resulted in Figure 5.21.

It was expected that a negative correlation might exist between thermal gradient and cell size, as a larger thermal gradient could be implied to mean more rapid solidification, and therefore a finer microstructure. As is evident from the charts, this correlation cannot be found, even when the data is broken down into separate sections as in rows (2), (3), and (4).

Effect of Thermal Gradient on Cell Size

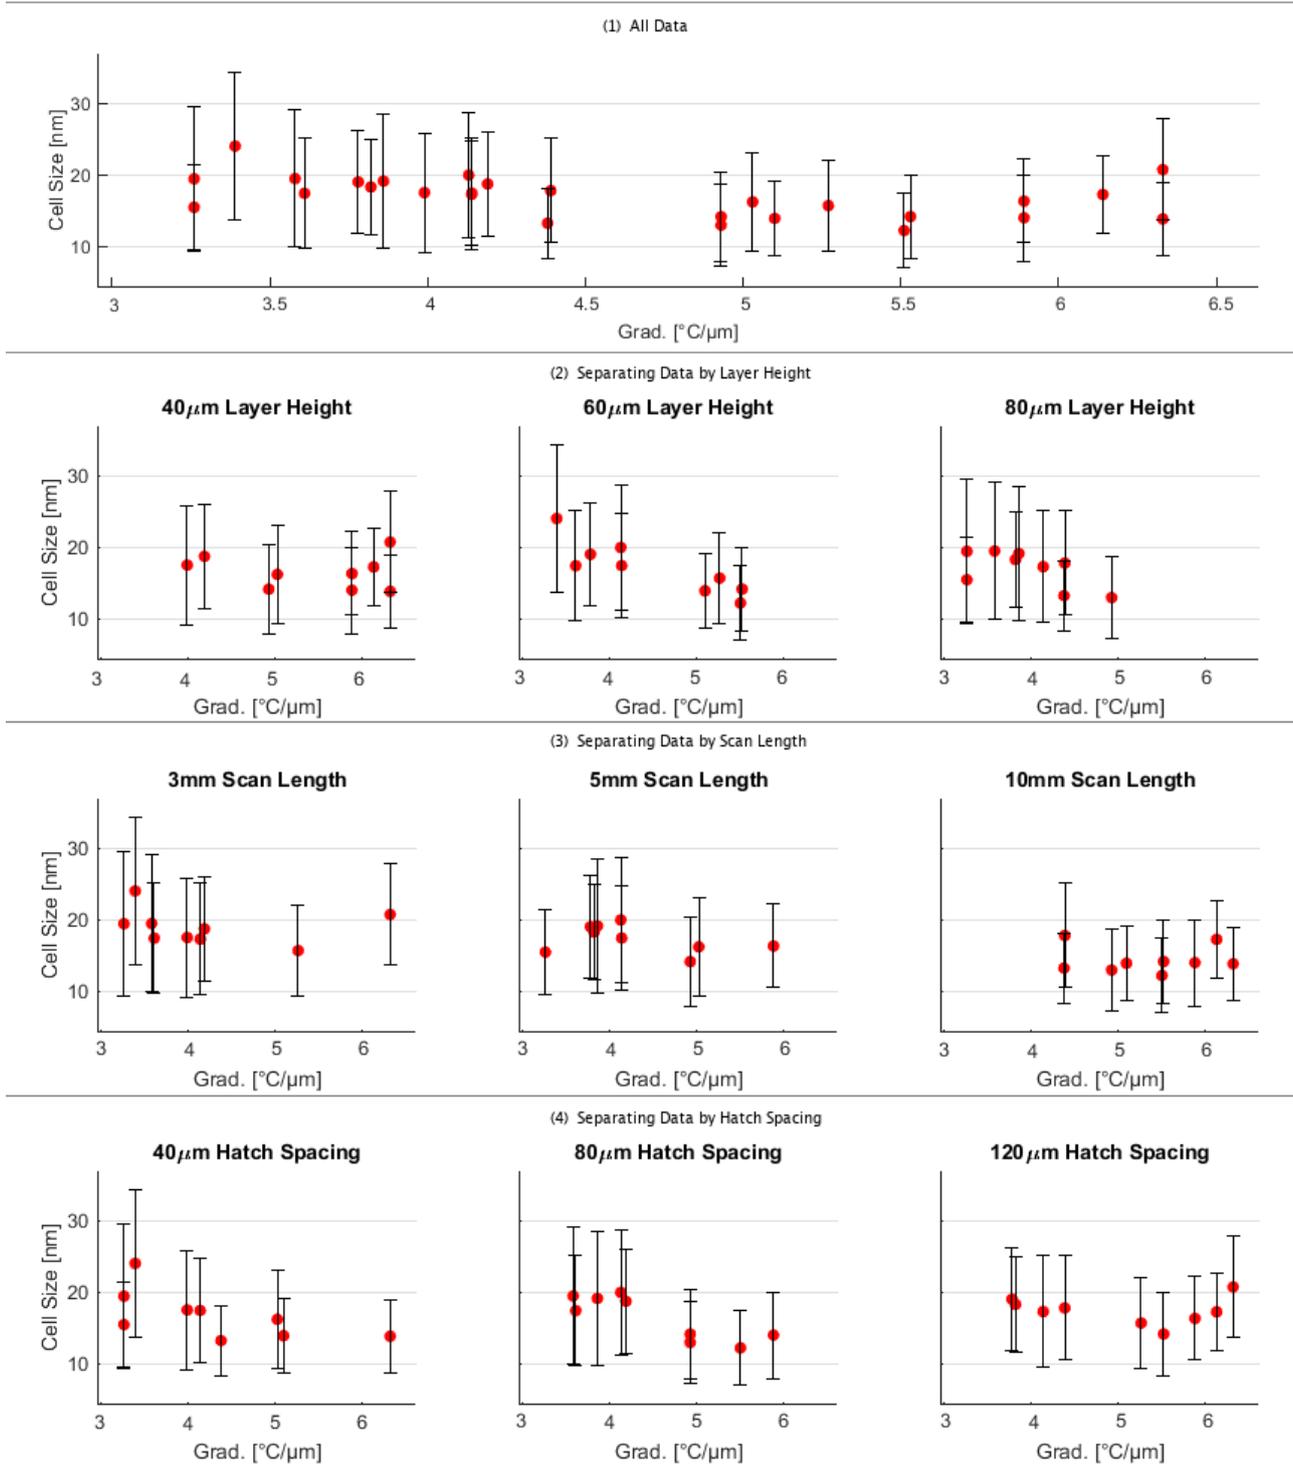


Figure 5.21

The effects of thermal gradient on cell size

5.3.4 - EBSD Analysis

Analysis of EBSD results were performed with Aztec (to extract grain size and shape), and with the MTEX toolbox for MATLAB (to extract inter- and intra- granular misorientation). Each parameter mean in the following results has been area-weighted, according to the formula:

$$\bar{x} = \frac{\sum_i A_i x_i}{\sum_i A_i}$$

where \bar{x} is the area-weighted mean, and A_i and x_i are respectively the area and parameter of each grain in the EBSD image.

The unbiased area-weighted standard deviations of the grains were calculated with the formula:

$$SD_{AW} = \sqrt{\frac{\sum_i w_i (x_i - \bar{x})^2}{\frac{(N - 1)}{N} \sum_i w_i}}$$

where N is the number of grains.

In both the MTEX and Aztec analysis, 10° was used as the threshold to distinguish one grain from the other.

5.3.4.1 - Grain Size

The area-weighted means of the areas of the grains in each map are shown in Figure 5.22. Due to the extreme variability in grain sizes, the area-weighted standard deviations of the grain areas were consistently calculated to be approximately equal to the area-weighted means. In order to make the charts easier to view, these large error bars have been omitted.

As can be seen, no identifiable trends exist in the data. It had been expected that, as scan length, hatch spacing, and layer height increase, the increased thermal gradient would result in more rapid solidification, and therefore a finer microstructure. Since no trends were found, it appears that attempting to predict grain size based on the predicted thermal outcomes of machine parameters may not be a reliable strategy, based on the overwhelming errors involved in such measurements, and the difficulty in linking such a complex phenomenon (solidification), to a simple set of physical variables.

Effect of Machine Parameters on Grain size

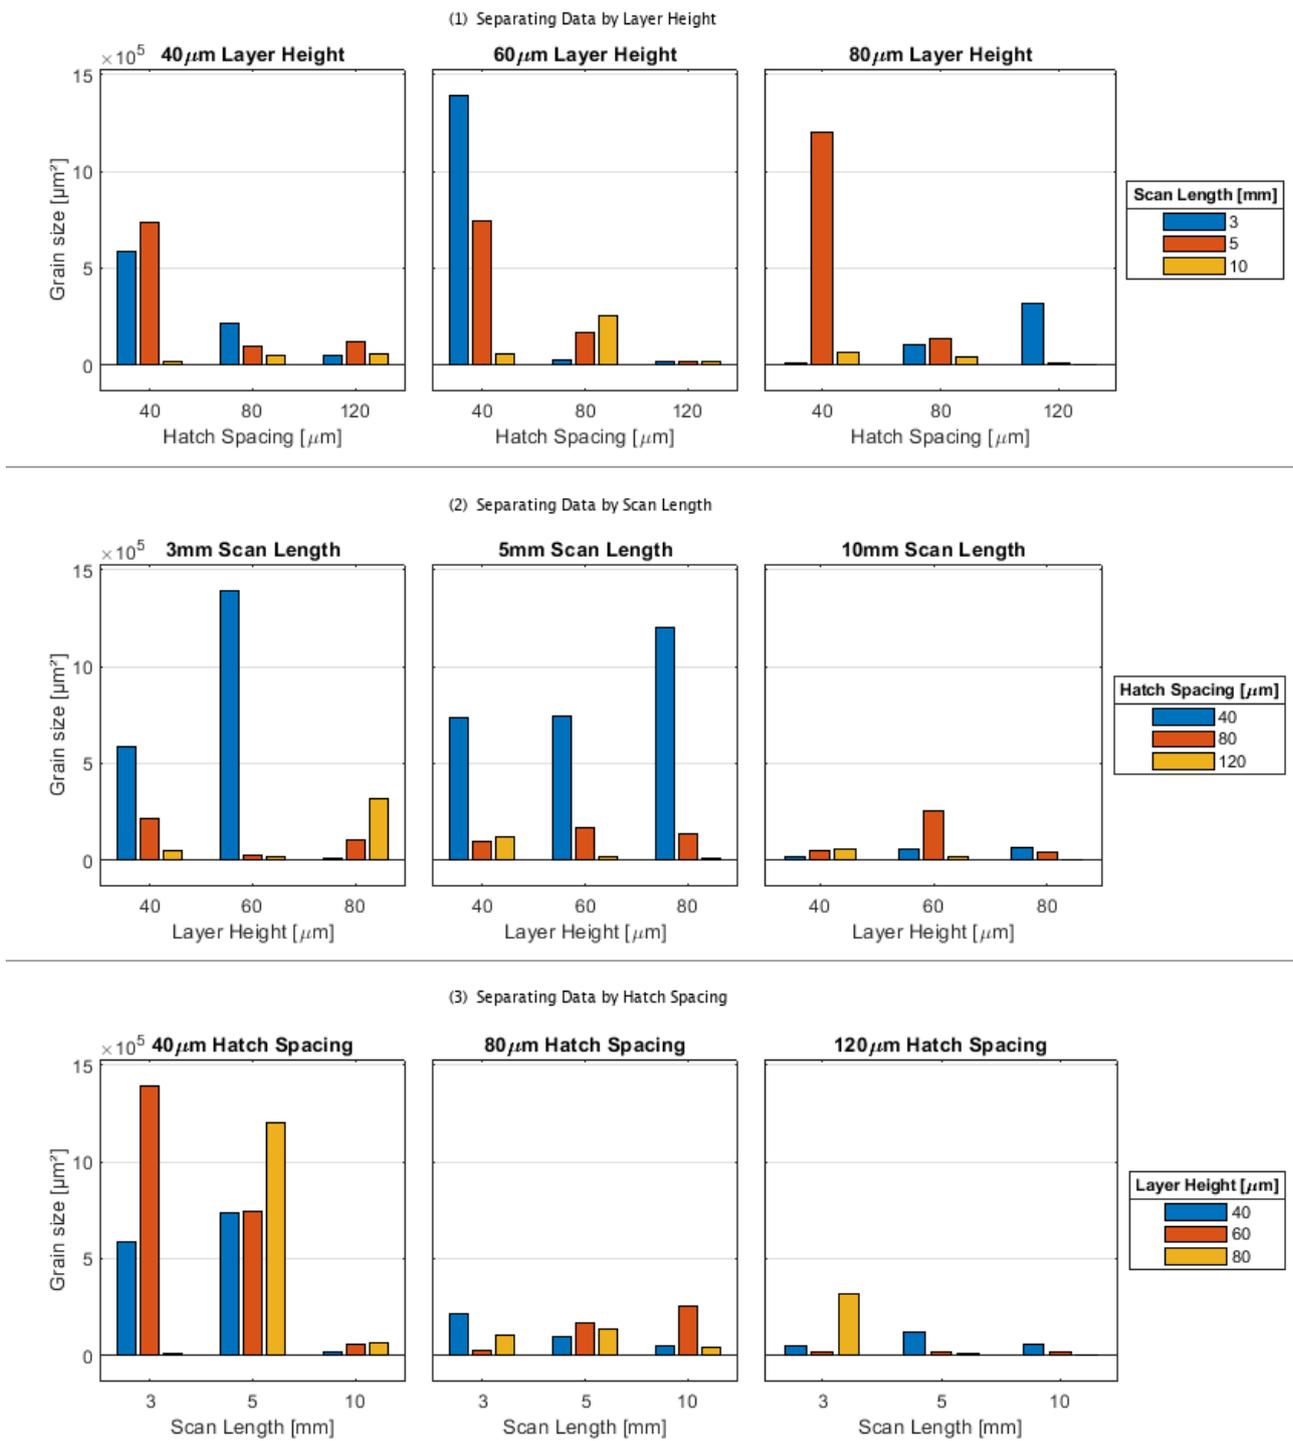


Figure 5.22

The effects of machine parameters on grain size

5.3.4.2 - Grain Misorientation

MTEX codes were written to analyse the absolute value of the misorientation between adjacent grains, along each pixel of all grain boundaries. A mean misorientation (i.e. mean of each pixel's misorientation with its neighbours) was then calculated for each grain, and an area-weighted mean and standard deviation was calculated for each sample. These values are shown in Figure 5.23.

Effect of Machine Parameters on Grain Misorientation

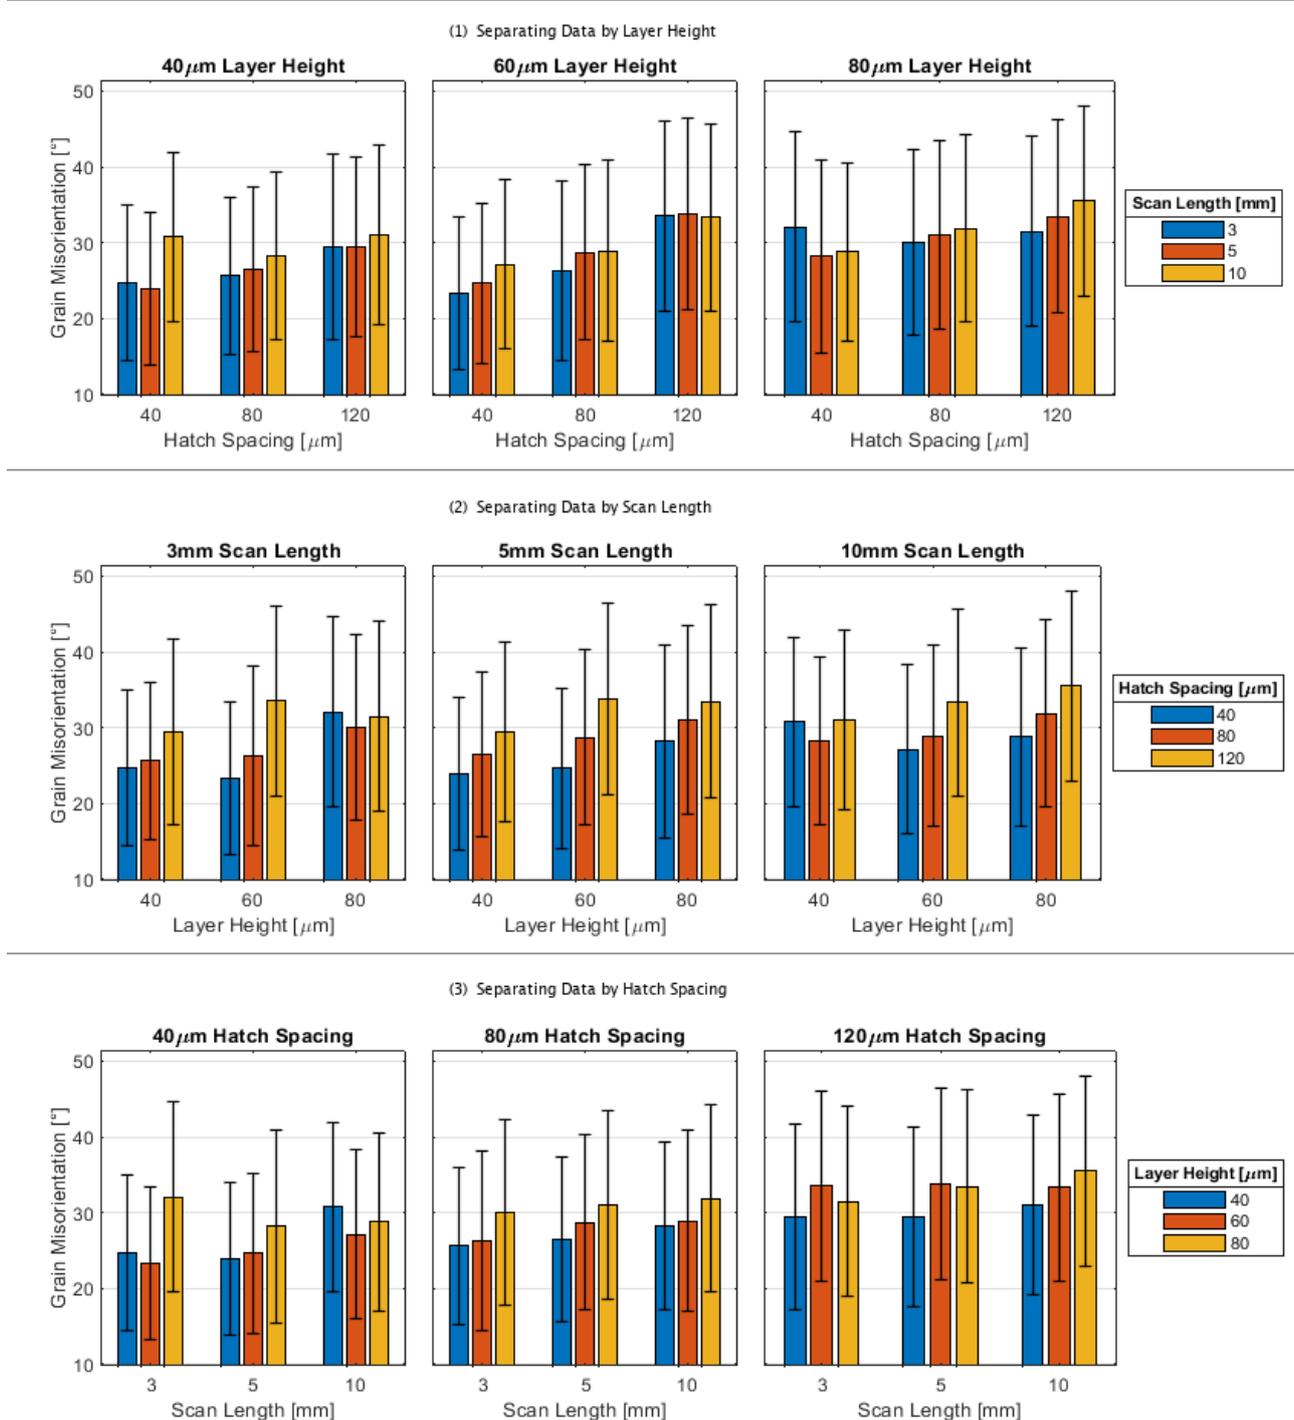


Figure 5.23

The effects of machine parameters on grain misorientation

As can be seen in row (1), misorientation appears to generally increase, as scan length increases. This rough trend opposes the expectation that, as scan length increases and therefore thermal gradient, solidification will shift upwards on the GR diagram, from the equiaxed (high misorientation) region towards the columnar/cellular (low misorientation) region.

This reversal of the expected trend suggests either that increasing scan length does not increase thermal gradient as predicted, or that increasing thermal gradient does not shift solidification from the high-misorientation to the low-misorientation region of the GR diagram as predicted, or both.

As can be seen in row (2), misorientation appears to consistently increase as hatch spacing increases. This trend also opposes the expectation that, as hatch spacing increases and therefore thermal conduction decreases and thermal gradient increases, solidification will shift upwards on the GR diagram. This reversal of the expected trend also calls the value of the solidification predictions into question.

As can be seen in row (3), misorientation appears to roughly increase as layer height increases. This trend also opposes the expectation that, as layer height increases and therefore thermal conduction decreases and thermal gradient increases, solidification will shift upwards in the GR diagram. This reversal likewise questions the value of the predictions.

5.3.5 - Hardness Analysis

The mean and standard deviation of the five Vickers hardness measurements taken, for each of the 27 cubes, are shown in Figure 5.24.

As can be seen in row (1), no clear trend exists between scan length and hardness. Based on the results presented in the previous subsection, it was expected that as scan length increased, and greater thermal gradients led to more rapid solidification and smaller microstructural features, hardness would increase too. The fact that this trend is not found, suggests that the link between cell size and hardness is not so straightforward.

As can be seen in row (2), hardness quite consistently and significantly decreases as hatch spacing increases. This trend runs contrary to the expectation that larger hatch spacings will lead to lower thermal conduction, and therefore greater thermal gradients, more rapid solidification, and a finer microstructure. The trend is also interesting, considering that no trend could be found between hatch spacing and cell size. The fact that a trend exists between hatch spacing and hardness, whereas none did between hatch spacing and cell size, further weakens the assumption that hardness and cell size are related [139, 323].

One possible reason why hardness decreases as hatch spacing increases, is that the remelting of previously-solidified material in an adjacent hatch (which occurs to a greater extent at lower hatch spacings) may cause homogenisation and grain refinement. These processes can result in improved mechanical properties, which would lead to greater hardness.

Effect of Machine Parameters on Hardness

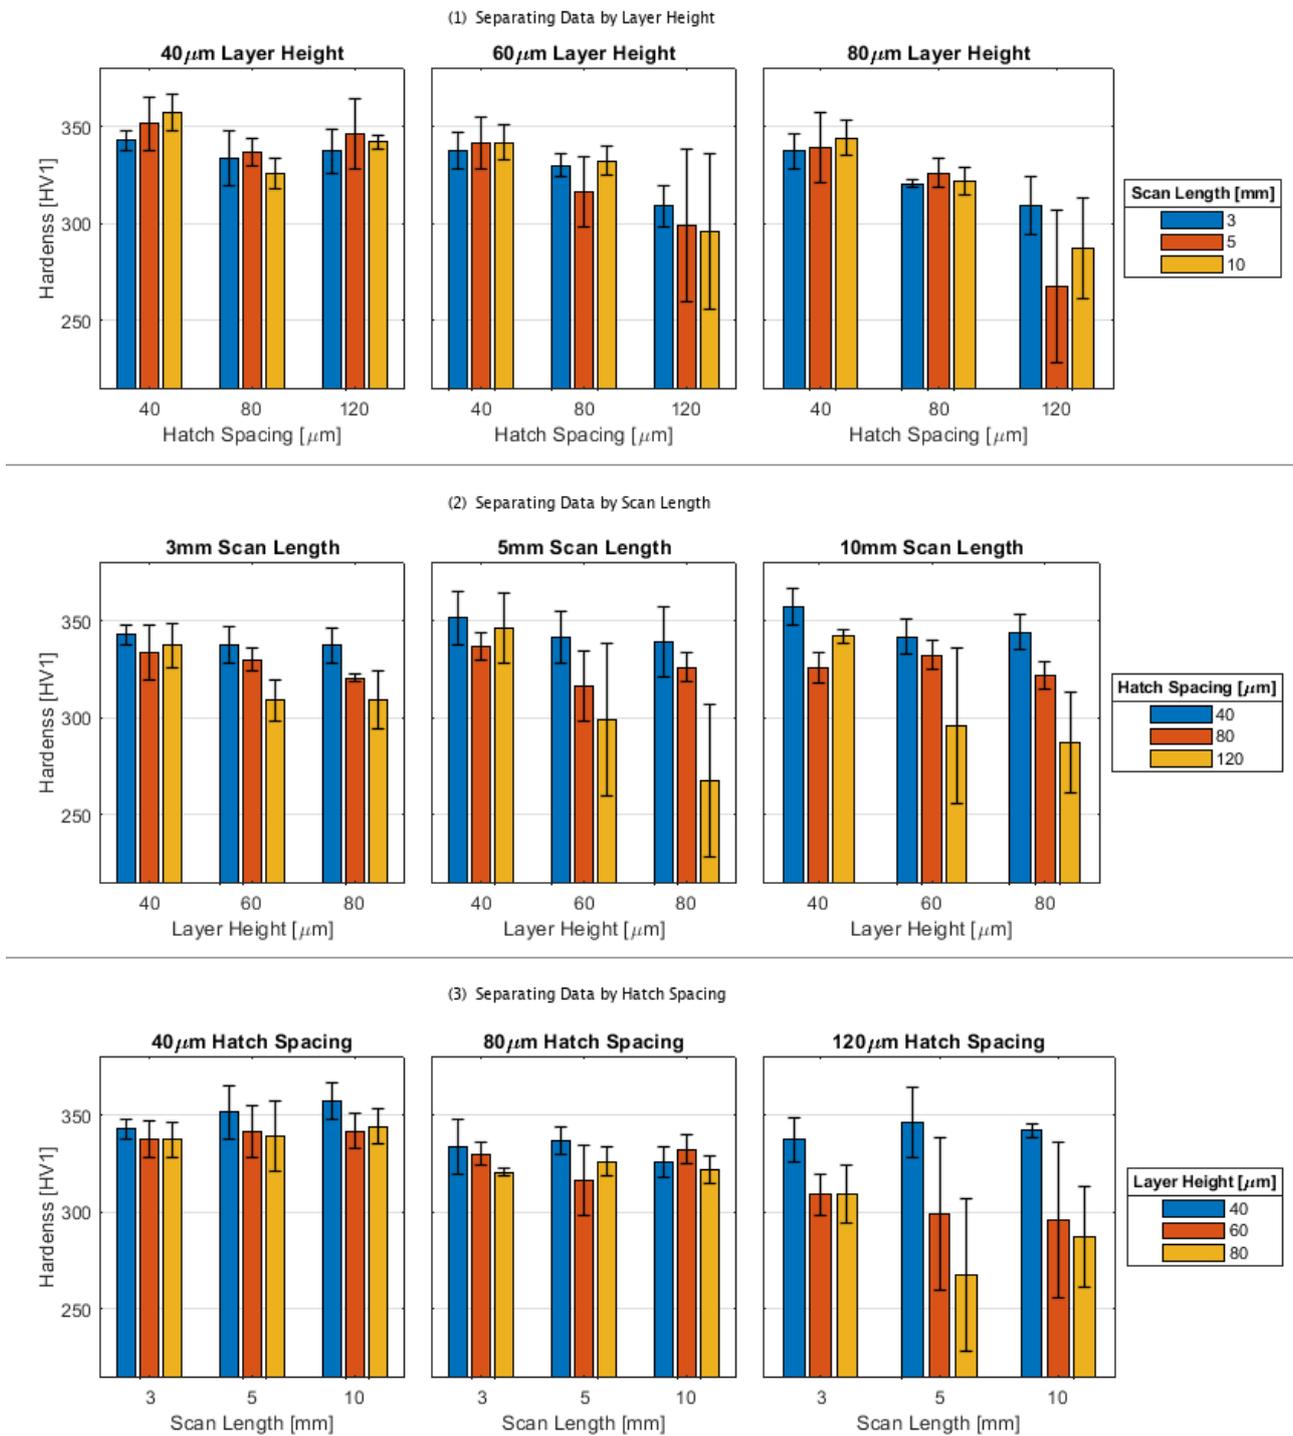


Figure 5.24

The effects of machine parameters on hardness

As can be seen in row (3), there exists a very slight tendency for hardness to decrease as layer height increases. This trend is more pronounced at greater hatch spacings. This slight decrease is contrary to the expectation that greater layer heights will lead to greater thermal gradients, and therefore more rapid solidification, a finer microstructure, and increased hardness. One possible suggestion

for this decrease, is the same as that which was suggested previously for the hatch-hardness decrease: smaller layer heights lead to more remelting, which may lead to more homogenisation and grain refinement.

In order to conclusively demonstrate that the relationship between mean cell size and hardness is either more complex than a simple correlation, or impossible to correlate because of the error associated with the measurement methods, or both, the two were plotted against each other as shown in Figure 5.25.

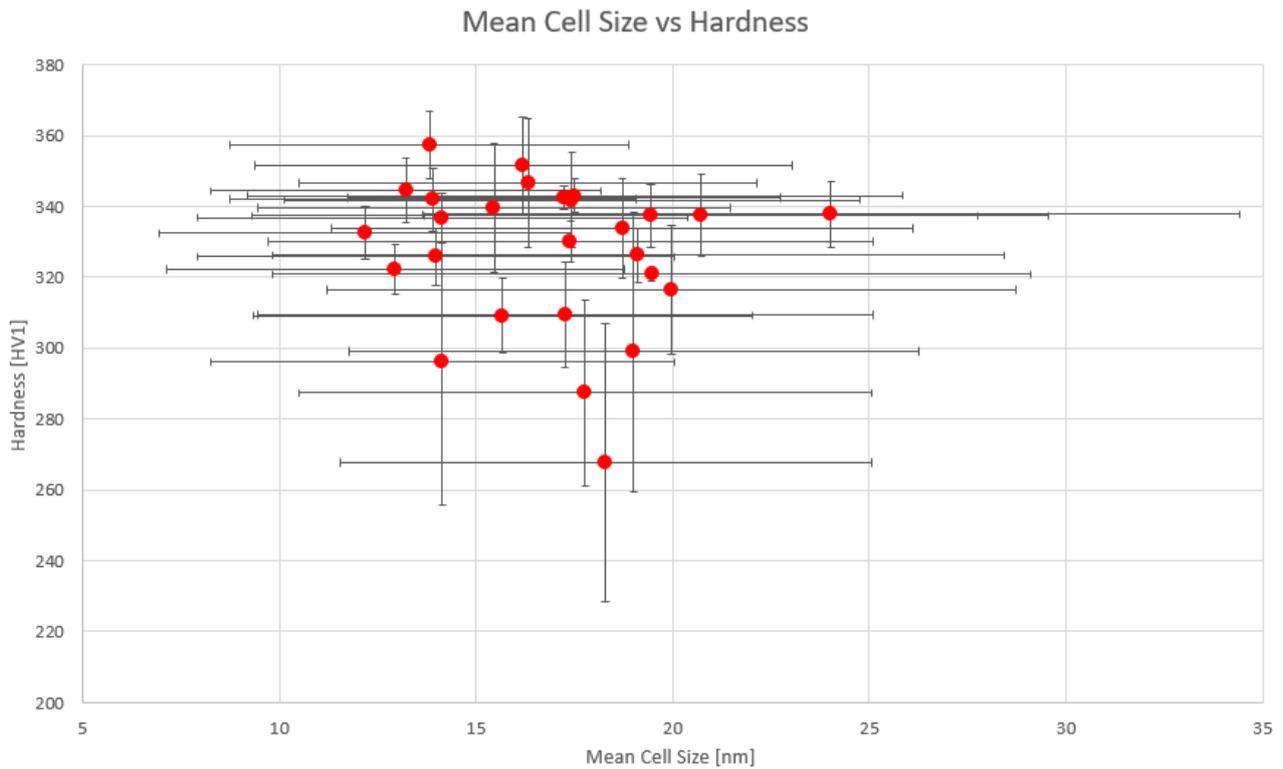


Figure 5.25

The relationship between mean cell size and hardness, and the error in both measurements

As is evident from the this chart, no clear correlation exists between the means, and also the errors (standard deviations) are far too significant for any trend to be considered relevant. Therefore, it appears that no simple trend links the results of this set of experiments.

According to literature, the Hall-Petch relationship between yield strength (σ_y) and grain size (d) states that:

$$\sigma_y = \sigma_o + \frac{k_y}{\sqrt{d}}$$

where σ_o is the base strength, and k_y is a materials constant. In order to see if any relationship could be found between (feature size)^{-1/2} and hardness, the data was plotted for both cell size (SEM) and grain size (EBSD), as shown in Figure 5.26.

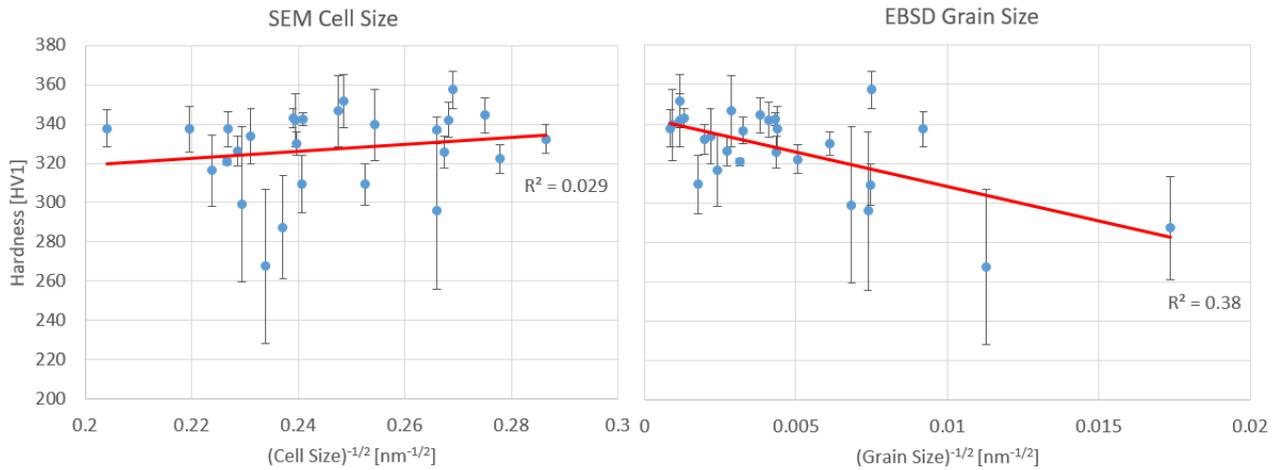


Figure 5.26

The Hall-Petch relationship applied to SEM and EBSD results

As can be seen from the graphs, the Hall-Petch relationship does not seem to hold for either the cell or grain sizes. In the case of the cells, the hardness is increasing as $(\text{cell size})^{-1/2}$ increases, as per the equation, but with a linear line of best fit only fits the data with an R^2 of 0.029. The grain size graph has a linear line with an R^2 of 0.38, but decreases as $(\text{grain size})^{-1/2}$ increases. It must be noted however that the standard deviation errors in hardness are very large, especially for the data points towards the right of the grain size chart, which are significant in dictating the angle of the line.

In order to calculate errors along the X dimension, it is not possible to simply use $(\text{standard deviation of feature size})^{-1/2}$. Instead, the error propagation formula must be used:

$$\Delta f(x) = \left| \frac{df(x)}{dx} \right| \Delta x$$

where $f(x)$ is the function being applied to the data, $\Delta f(x)$ is the error in the function output, and Δx is the error in the data.

In this case:

$$f(x) = x^{-1/2}$$

$$\frac{df(x)}{dx} = \frac{1}{2x^{3/2}}$$

$$\Delta f(x) = \frac{\sigma_x}{2x^{3/2}}$$

where σ_x is the standard deviation.

When these errors are added to the grain size chart, any confidence that the linear relationship is accurate drops even further, as shown in Figure 5.27.

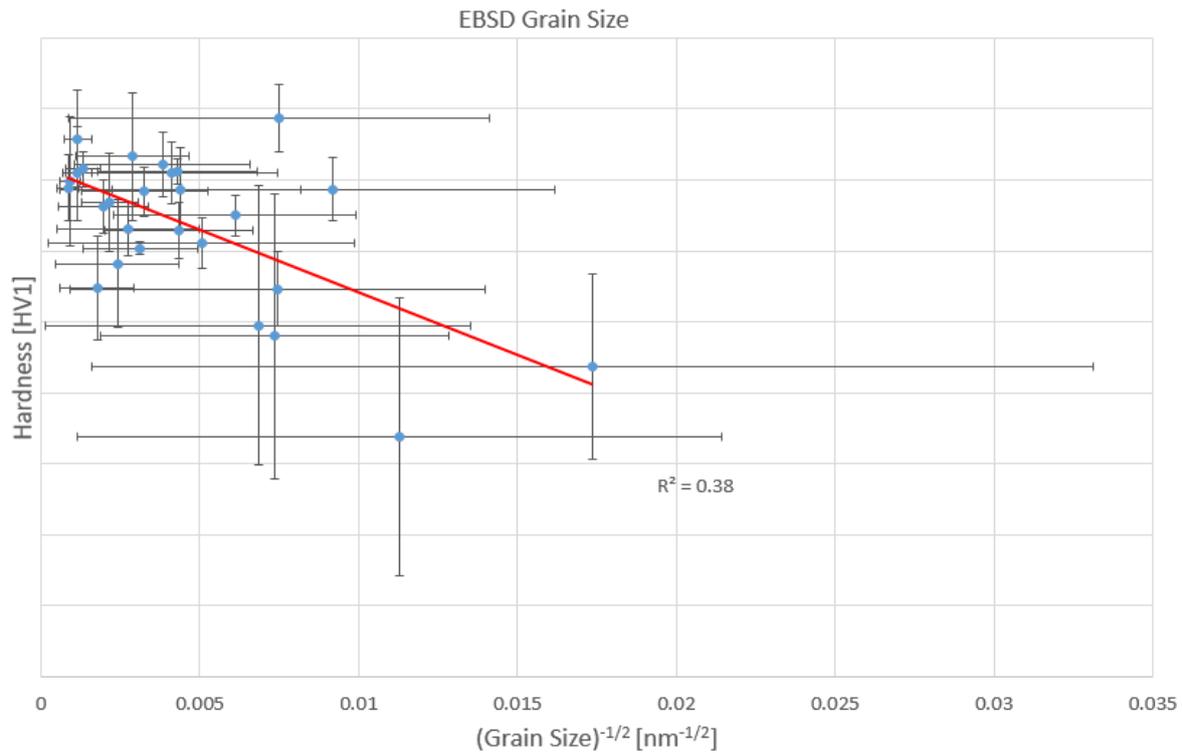


Figure 5.27

The Hall-Petch Relationship applied to the grain size, with the X and Y error bars

Based on these findings, it appears conclusive that expecting a coherent trend to emerge between hardness and microstructural feature size, is an oversimplification for samples produced in L-PBF.

5.4 - Regression Modelling

Table 5.3 shows all derived and measured values for each cube was compiled. Excluding trivial dimensionless numbers (such as v^* , which remained constant for all cubes), and the errors of all measured values, there remained 29 numbers of interest for each cube. Using MATLAB's in-built FIT function, a regression analysis was performed on every possible pair-combination of metrics, using five model types: linear, quadratic, exponential, logarithmic, and power. The R^2 value was calculated for each regression, and those above 0.6 were identified for further consideration. While such a simple analysis method can have drawbacks (low R^2 values can still provide useful trends, while high R^2 is not inherently good), it was reasoned to be useful in searching for realted variables in such a large data set.

Those with trivial explanations for a strong correlation were excluded, for example:

- Scan length is strongly correlated with both Keyhole Number and Normalised Enthalpy, as these numbers both use Return Time in their derivation, which is a function of scan length.
- Many strong correlations exist within the metrics extracted from the thermal camera, all of which can be explained with the truism that hotter melt pools tend to be larger and that an increase in one melt pool dimension usually implies an increase in another.

- The same is true for metrics extracted from the optical microscope: more porous samples contain larger and less circular pores.

Table 5.3
The metrics extracted for each cube

Type	Value	Type	Value	Type	Value
Build Parameters	<i>Scan Length</i>	Thermal Camera	<i>Mean Temp</i>	Optical Microscope	Porosity
	<i>Hatch Spacing</i>		<i>Peak Temp</i>		Pore Area
	<i>Layer Height</i>		<i>MP Length</i>		Circularity
	<i>Peclet</i>		<i>MP Width</i>	Hardness	Hardness
Dimensionless Numbers	<i>Marangoni</i>	Thermal Camera	<i>MP Area</i>	EBSD	Grain Size
	<i>Fourier</i>		<i>Aspect Ratio</i>		Aspect Ratio
	<i>Keyhole</i>		<i>Thermal Grad.</i>		Misorientation
	<i>N. Enthalpy</i>		<i>Spatter</i>		GAM
	<i>E*</i>	SEM	<i>Cell Size</i>		Orientation
	<i>To</i>		<i>Aspect Ratio</i>		

Also excluded were those that have been previously discussed, for example:

- Scan length and melt pool width correlate with a linear R^2 of 0.71
- Layer height and melt pool length correlate with a linear R^2 of 0.73

Once such correlations had been excluded, six remained that were deemed worth discussing further.

5.4.1 - Porosity and Grain Misorientation

It was found that grain misorientation increased with sample porosity, according to a logarithmic function, with an R^2 of 0.65, as shown in Figure 5.28.

It follows logical reasoning that porosity and misorientation must increase together, for the following reasons:

1. Pores interrupt epitaxial growth in L-PBF solidification, disrupting the columnar grains commonly observed, forcing the existence of new orientations and nucleation sites.
2. Porosity alters both thermal gradient and growth rate, shifting solidification's position on the GR diagram, changing the locally preferred grain size and morphology.
3. Pores disrupt liquid flow within the melt pool, affecting pool dimensions, solute segregation, and the extent of remelting, all of which in turn affect grain growth direction.

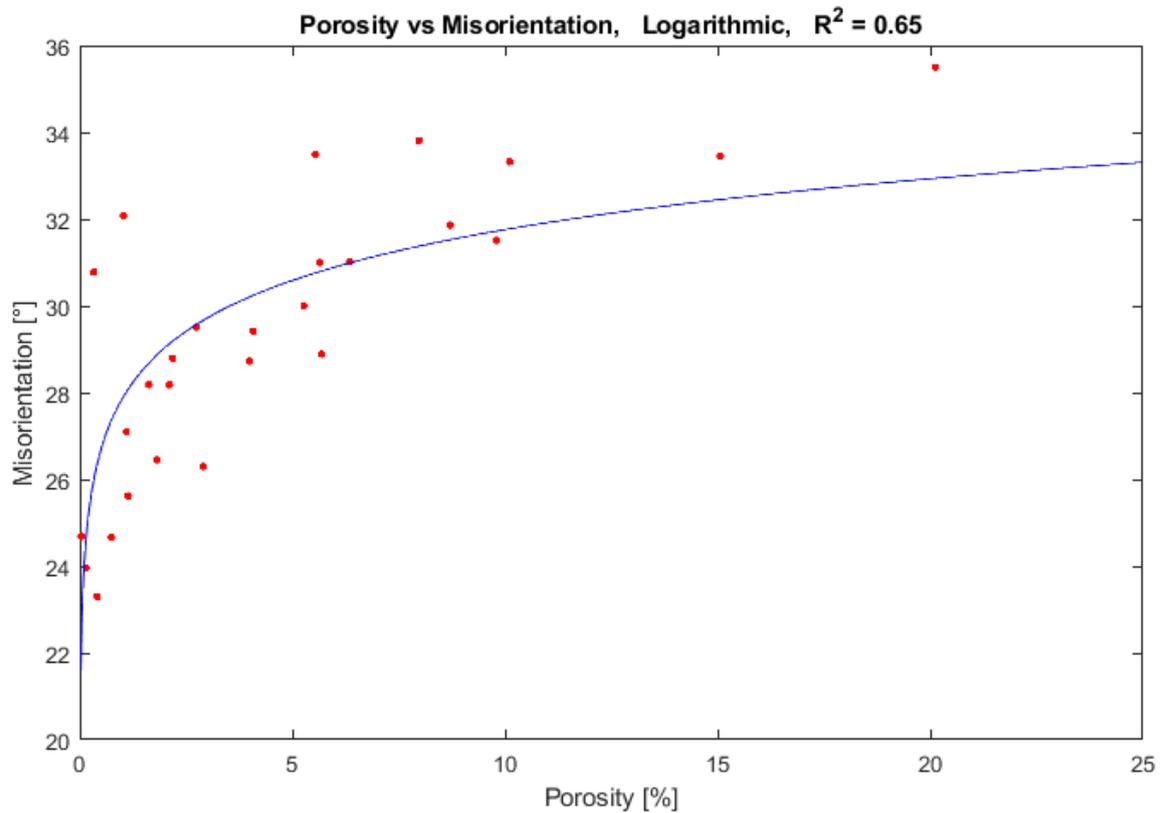


Figure 5.28

The logarithmic relationship between porosity and grain misorientation

It likewise stands to reason that the relationship between porosity and grain misorientation should be non-linear. The disruption to solidification caused by each additional volume of porosity that exists within a sample, is influenced by the disruption caused by previous volumes of porosity, with the disturbances compounding.

The implication of this finding is that a small increase in porosity may reduce anisotropy and promote equiaxed growth. Depending on the component's intended function and post-processing strategy, this trade-off may be desirable. For example, should a component be created with a minimal amount of keyhole pores, which are later closed up by HIP, then the increase in misorientation may be beneficial.

5.4.2 - Porosity and Hardness

Hardness was found to decrease with sample porosity, according to an exponential function, with an R^2 of 0.63, as shown in Figure 5.29.

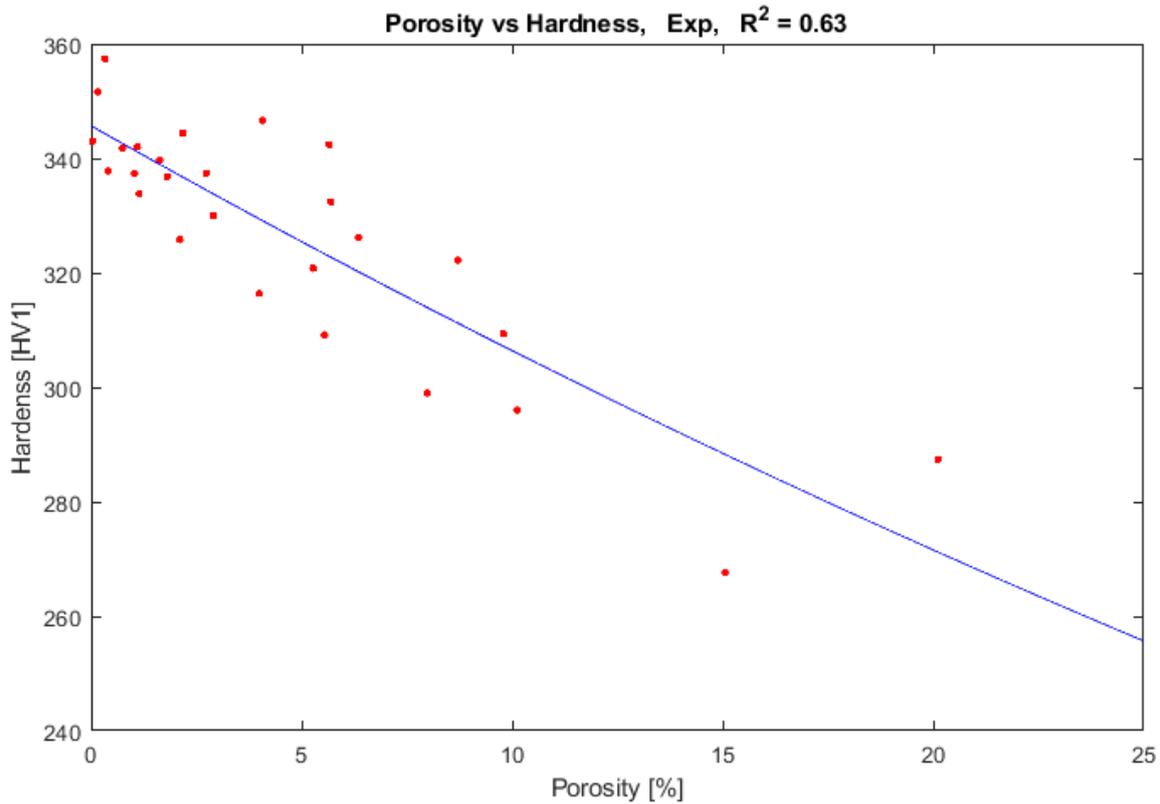


Figure 5.29
The exponential relationship between porosity and hardness

Again, it follows logical reasoning that micro-hardness should decrease with porosity, according to an exponential decay. Starting with a base hardness at zero porosity, as each volume of material is subtracted, mechanical strength will be degraded. Since hardness testing sites were deliberately chosen to not interact with pores, however, this degradation will not be linear. Indentations were kept at least 1mm from any pore, however it was not possible to keep indentations away from subsurface porosity. The rate of degradation will slow, assuming that non-porous sites can still be found, moving towards a new baseline of the hardness of a small fragment of solidified material, surrounded by a sponge of pores.

The implications of this finding are not significant for practical applications. In industrial settings, the micro-hardness of a location within a component, specifically chosen for its non-porosity, is unlikely to be relevant, compared to its macro mechanical properties. Therefore, the fact that the decrease in Figure 5.29 appears to follow an exponential decay, may serve as evidence of the limitations of the experimental method.

5.4.3 - Inter-Granular and Intra-Granular Misorientation

Mean GAM was found to decrease with grain misorientation, according to a linear function, with an R² of 0.63, as shown in Figure 5.30.

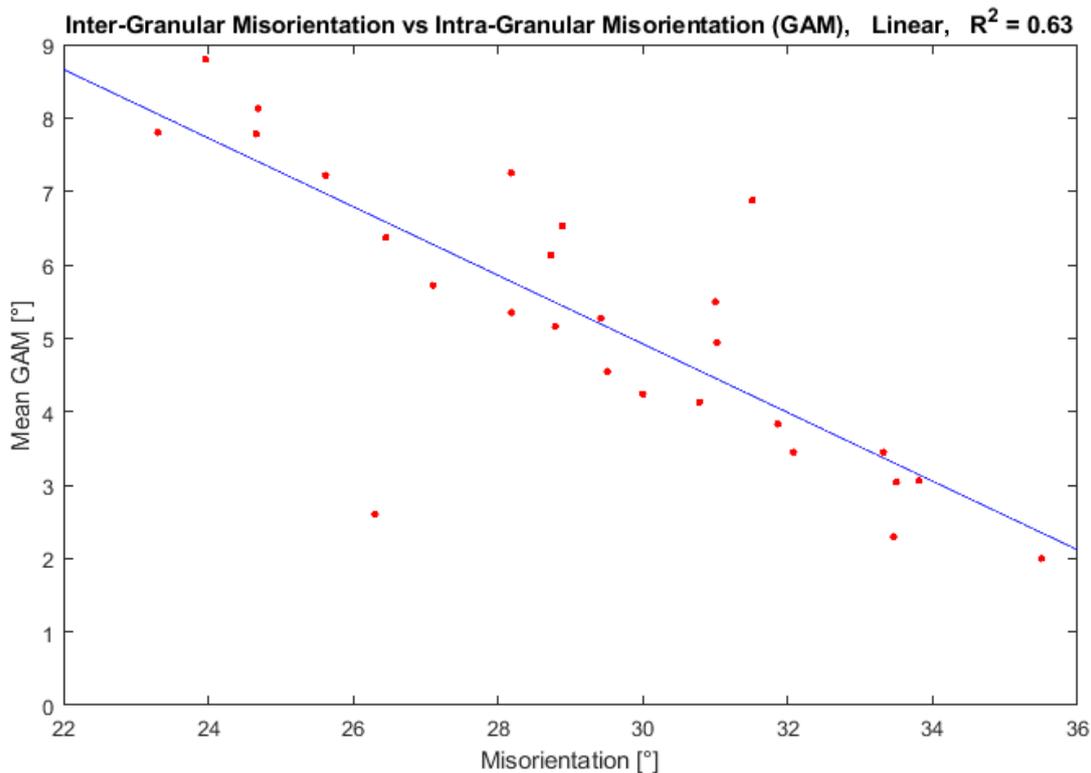


Figure 5.30

The linear relationship between misorientation and mean GAM

Both inter- and intra- granular misorientation can serve as indicators of the rapidity and turbulence of solidification in L-PBF [324]. It was expected that the kind of conditions that resulted in differing angles between grains (quickly advancing solidification fronts, insufficient time for atomic diffusion, abundant perturbations in the solid-liquid boundary, etc.) would also result in differing angles within grains. The fact that the opposite appears true, was surprising.

One reason suggested for this observation, is that grains with lower misorientation are likely the product of columnar epitaxial growth. It may be the case that as these columns grow upwards, contracting under thermal contraction, the accumulation of strain with each grain may be greater than if the grains solidified without epitaxial growth.

Another possibility is that this apparent linear decrease is simply a by-product of grains growing with greater GAM being more likely to separate into different grains, due to internal stress. Therefore, it would not be the case that grains with more misorientation experience less GAM, but rather grains with more GAM divide and share out their differing angle in the form of separate grains with lower GAM but more misorientation.

5.4.4 - Melt Pool Length and Thermal Gradient

Thermal gradient was found to decrease with melt pool length, according to a linear function, with an R^2 of 0.78, as shown in Figure 5.31.

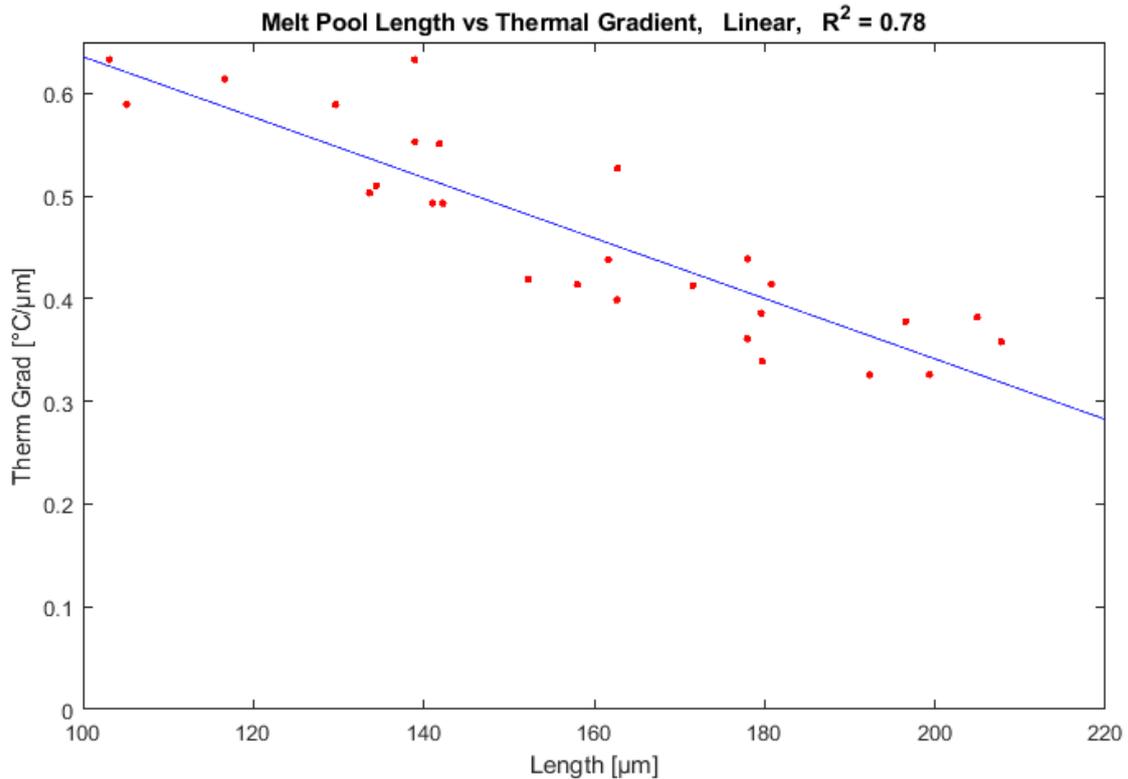


Figure 5.31

The linear relationship between melt pool length and thermal gradient

Previously, results were presented as a function of certain parameters (grain size, misorientation, etc.) plotted against thermal gradient. It had been assumed that thermal gradient, as measured from the peak-temperature-pixel to last-liquidus-pixel in the melt pool by the thermal camera, served as a reliable proxy for the thermal gradient of the solidification front. This chart questions this assumption.

Since thermal gradient reliably and linearly decreases as melt pool length increases, it could be the case that the former is simply a function of how large the latter happens to be. If a hypothetical set of melt pools all had the same peak temperature, but the extent of trailing above-liquidus material varied stochastically, then the relationship between length and measured thermal gradient would be linear and negative, as in this chart. In this scenario, it would be uncertain if the peak-temperature-pixel to last-liquidus-pixel measurement method would be an appropriate proxy for what's happening at the actual solidification front.

This finding begs an obvious question: to measure the thermal gradient, rather than simply divide the peak-to-liquidus temperature difference by the peak-to-tail length, why not analyse the cooling curve closer to the liquidus boundary? As demonstrated in a previous section, this cooling curve is extremely chaotic, and attempting to extract meaningful gradients from it proves unreliable.

The implication of this finding is profound: using the available thermal camera equipment, there may be little value in extracting cooling curves from the melt pool images.

5.4.5 - Grain Size and GAM

The mean GAM was found to increase with grain size, according to a power function, with an R^2 of 0.81, as shown in Figure 5.32. It should be noted that grain size, and not $(\text{grain size})^{-0.5}$ is being used here, since the Hall-Petch analysis used above related to yield stress, and not to misorientation.

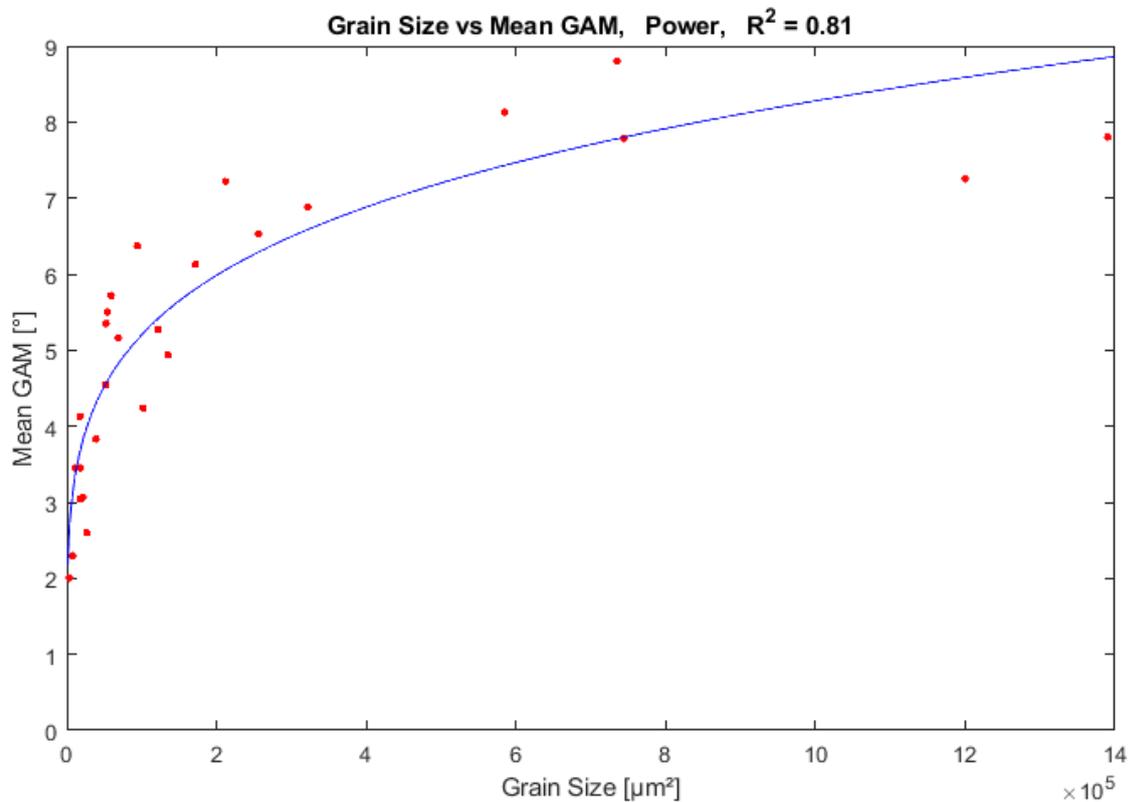


Figure 5.32

The power-law relationship between grain size and mean GAM

As per the previous discussion about misorientation and GAM, it was expected that a negative correlation should exist between grain size and GAM: grains with greater internal strains should split into smaller grains. It is therefore surprising that the exact opposite appears to be true. Larger grains very consistently show increased GAM.

Once again, epitaxial growth is a possible explanation for this result. The largest grains were shown in EBSD images to extend through multiple melt pool boundaries. It could be the case that these boundaries, which exist inside large grains, contribute significantly to GAM, while smaller grains, which do not include such boundaries, experience lower GAM.

This finding serves as an argument for recrystallization in post-processing treatment. Such treatments reduce the anisotropy of the microstructure, transforming large, epitaxial grains that encompass melt pool boundaries (which can serve as crack initiation sites), into smaller isotropic grains with increased strength and fatigue resistance.

5.4.6 - Normalised Enthalpy and Pore Area

Pore area was found to decrease with normalised enthalpy, according to an exponential decay, with an R^2 of 0.69, as shown in Figure 5.33.

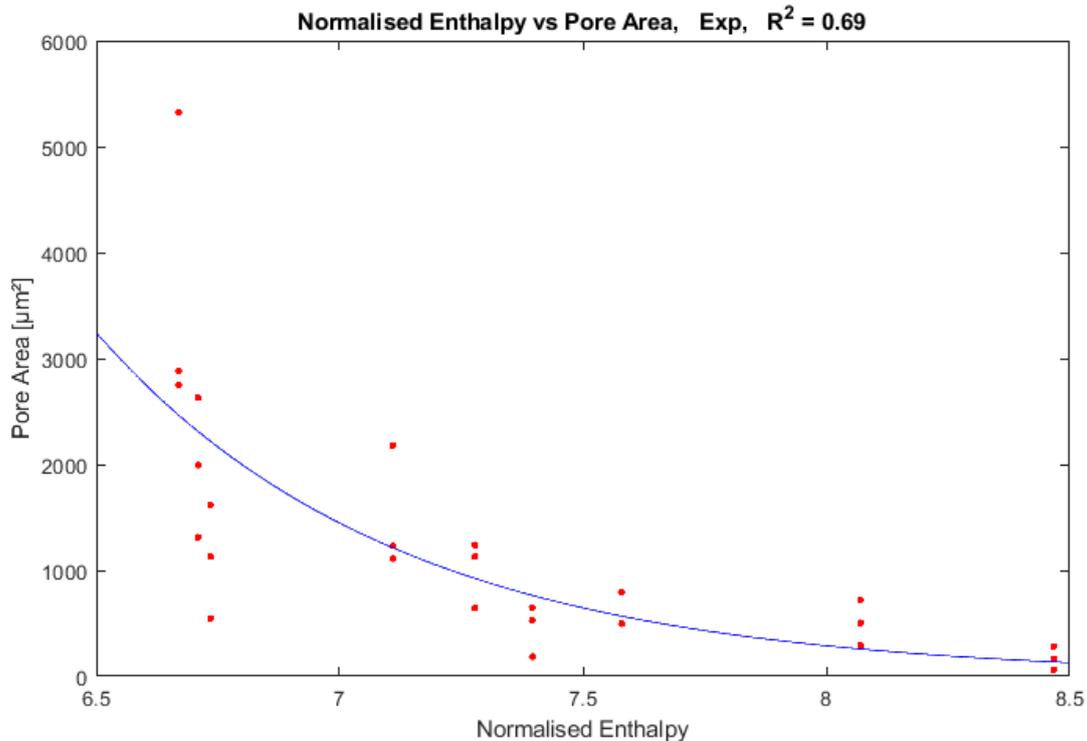


Figure 5.33

The exponential relationship between normalised enthalpy and mean pore area

By itself, this finding is not very impactful. Both normalised enthalpy and pore area have previously been shown to be cause and effect in the extent of LoF porosity in a sample. It is understandable that as the latter decreases, the former begins to rise exponentially, as more pores and larger pores begin appearing.

What is more interesting, is the fact that this relationship is the only identified trend that contains a dimensionless number. Despite seven dimensionless numbers and nineteen physical measurements being included in this correlation search, only this quite trivial finding included a dimensionless number. There were not, for example, trends between the Marangoni number (which describes melt pool turbulence) and porosity, or between volumetric energy density and grain morphology.

This lack of findings hints at the limits of dimensionless numbers in predicting microstructural results. Though dimensionless energy was demonstrated in a previous subsection to be useful in predicting porosity, the fact that no SEM, EBSD, or hardness results are correlated with any dimensionless number, implies that their value may be curtailed to density measurements.

5.5 - In-Situ Microstructure Conclusions

The primary intention of the experiments performed in this chapter was to establish a link between the thermal signatures detected by the thermal camera, and the resulting microstructure and mechanical properties. The underlying hypothesis of the experiments was that variations in build parameters would result in predictable changes to both melt pool thermal signatures, and microstructural characteristics as determined from first principles upon the GR diagram.

Such attempts to predict either melt pool thermal signatures, and microstructural characteristics, proved unreliable. Determining links between any metrics pulled from different data sources (e.g. melt pool length from the thermal camera, grain size from the EBSD, cell size from SEM, etc.) remained challenging.

A significant limitation was discovered, when a strong linear correlation between melt pool length and melt pool tail thermal gradient was found (R^2 of 0.78). This finding suggested that thermal gradient measurements may simply be a straightforward measure of peak temperature divided by pool length, rather than a proxy for the gradient at the actual solidification front, as per experimental assumptions.

It was found that machine parameters largely had the expected effects on melt pool dimensions: increasing scan length consistently decreased melt pool length, width, and area, while increasing layer height consistently increased melt pool size and temperature due to decreased thermal conduction. A perfect trend was also found between volumetric energy density and porosity, allowing for the conclusion that LoF porosity was the primary source of defects in the samples, rather than keyholing.

Regression modelling further demonstrated the lack of correlation between metrics from different sources. Porosity and grain misorientation was found to have $R^2 = 0.65$, likely due to the fact that pores disrupt epitaxial growth, and hardness significantly decreased as hatch spacing increased, possibly due to reduced remelting at higher spacings leading to less homogenisation and refinement. Very few other trends emerged, however, demonstrating the limitations of extracting meaningful data from the thermal camera to predict microstructural or mechanical properties.

6 - In-Situ Monitoring and Mesostructure

This chapter explores the ability of the thermal camera to detect the creation of keyhole and LoF porosity. Test pieces were created, with porosity throughout that was placed deliberately via machine parameter and scan strategy. These pieces were taken to the ESRF to be scanned with XCT, in order to determine correlations between thermal signatures from in-situ monitoring, and porosity embedded within the alloy.

6.1 - Parameter Selection

Before more in-depth experimental work was carried out, it was necessary to test the ability of the Aconity LAB to produce dense parts, under a range of machine parameters. For IN-718, this was performed with a high-throughput build.

6.1.1 - Experimental Setup

In order to efficiently build a large number of samples simultaneously, without the need for individual mounting, grinding, and polishing, a 32mm wide sample holder was designed, inside which multiple cubes could be built, and which could fit directly into the Automet 250. Figure 6.1 shows the Netfabb design for the holder (grey), and the seventeen 4mm × 4mm × 5mm (X, Y, Z) cubes built upon it.

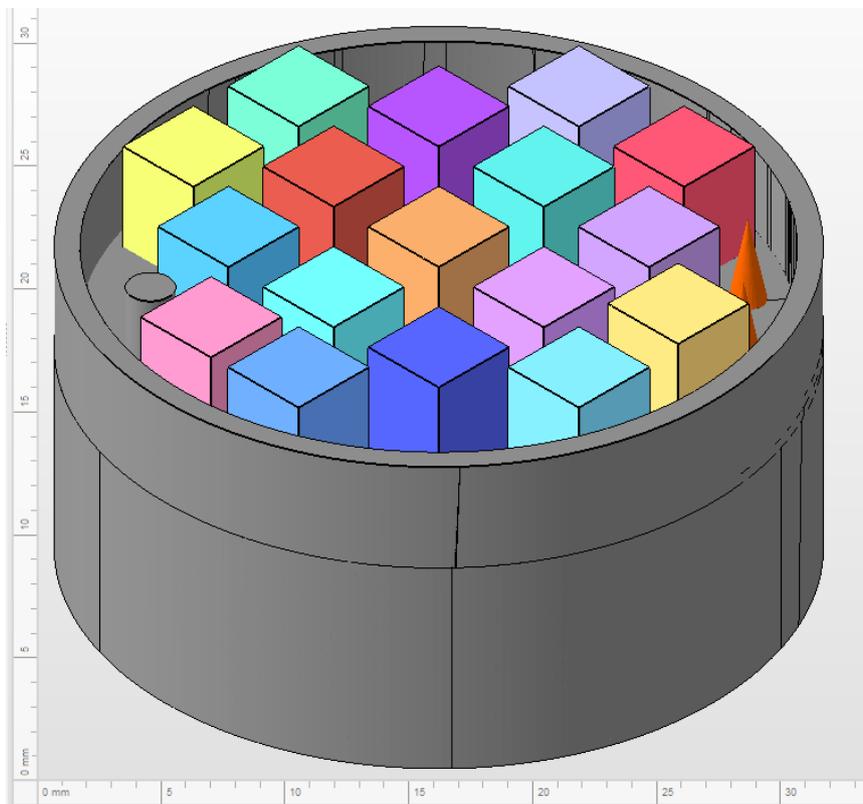
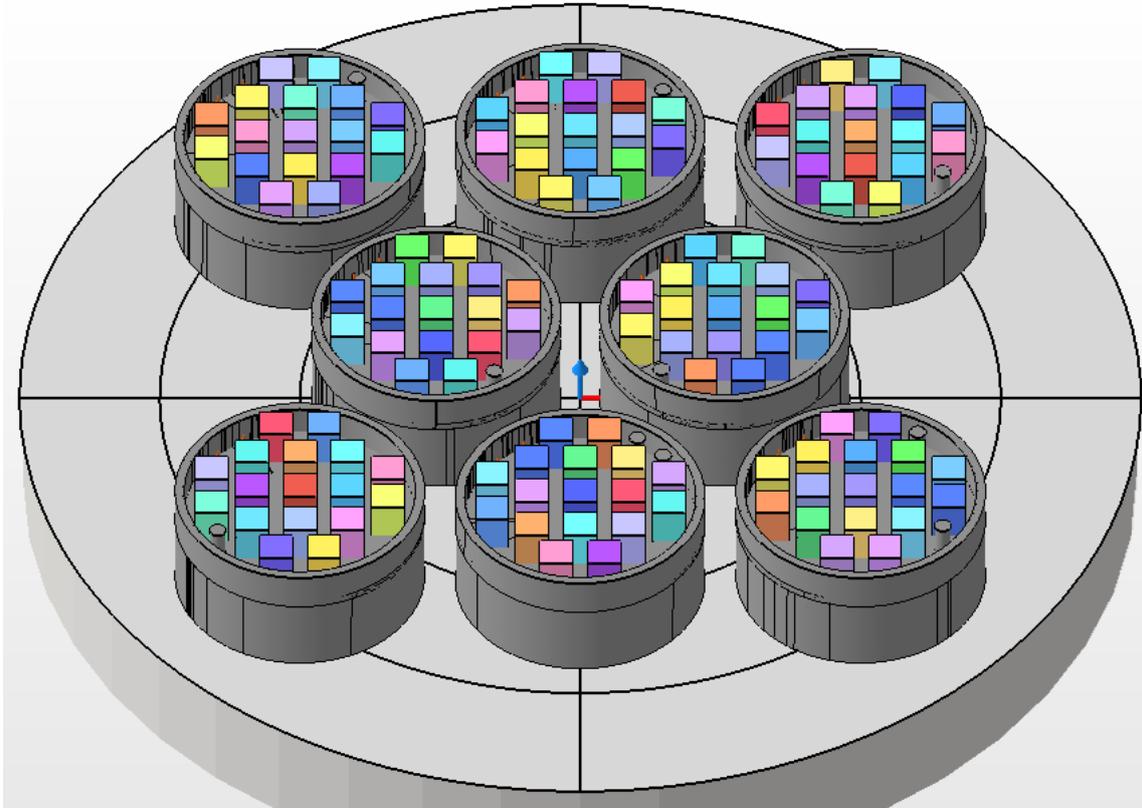


Figure 6.1

A close-up of the 17 cubes in the holder

The cubes stick 1mm above the surrounding cylinder wall, which allows for any shards ground away by the Automet to not become trapped within the sample. In Figure 6.1, the small cylinder and two orange cones that are visible at the left and right edges of the holder, respectively, are used to identify it amongst the set of 8 that were built simultaneously, as shown in Figure 6.2.



*Figure 6.2
The eight cube holders on the base plate*

Each of the 136 cubes were built with bi-directional scanning, with 67° rotation between the layers. A $30\mu\text{m}$ layer height was used. The hatch spacing, laser power, and laser speed were varied between each cube, with most combinations of the ranges in Table 6.1 being tested.

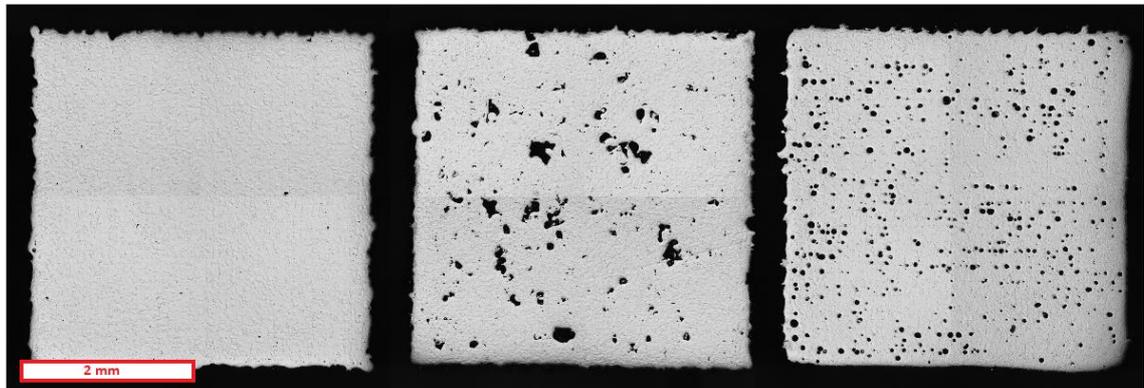
*Table 6.1
The parameters used for the IN-718 cubes*

Laser Powers [W]	85, 100, 115, 130, 145, 160, 175, 190
Laser Speeds [mm/s]	300, 500, 700, 900, 1100, 1300, 1500, 1700, 1900
Hatch Spacings [μm]	30, 50, 70, 90

Of the 288 possible combinations of these parameters, the largest and smallest energy density combinations were omitted, as it was known based on literature that these parameters would result in LoF and keyhole porosity [325-327]. Combinations with volumetric energy densities between 58.48 J/mm^3 and 166.67 J/mm^3 were chosen.

6.1.2 - Experimental Results

Bright-field OM images were taken of each cube surface. These images were then analysed in MATLAB, to extract the total density, the mean pore area, and the mean pore circularity. In Figure 6.3 are three example images, showing a near-fully dense sample, a sample with LoF porosity, and a sample with keyhole porosity.



Power	190W	100W	115W
Speed	1900 mm/s	1700 mm/s	300 mm/s
Hatch	50 μ m	30 μ m	90 μ m
Density	99.96%	94.87%	93.58%
Num. of Pores	27	375	1,083
Mean Pore Area	24 μ m ² (<i>7 μm²</i>)	975 μ m ² (<i>327 μm²</i>)	1,120 μ m ² (<i>874 μm²</i>)
Mean Pore Circularity	0.65 (<i>0.13</i>)	0.45 (<i>0.35</i>)	0.78 (<i>0.28</i>)

Figure 6.3

An example of a fully dense (left), LoF (centre), and keyholed (right) cube. *Italicised values are standard deviations.*

In Figure 6.4, the laser powers and speeds are shown, for each of the four hatch spacings, as was presented in a diagram previously. The size of each data point depends on its porosity (i.e. larger points are more porous, less dense), and its colour depends on the sample's mean circularity.

As was expected, porosity increases at higher and lower energy densities, and circularity is lowest at lower energy densities (indicating LoF), and highest at higher energy densities (indicating keyholing).

Figure 6.5 shows how circularity increases with energy density, making the distinction between LoF and keyhole regimes clear. The size of each data point depends on its mean pore area.

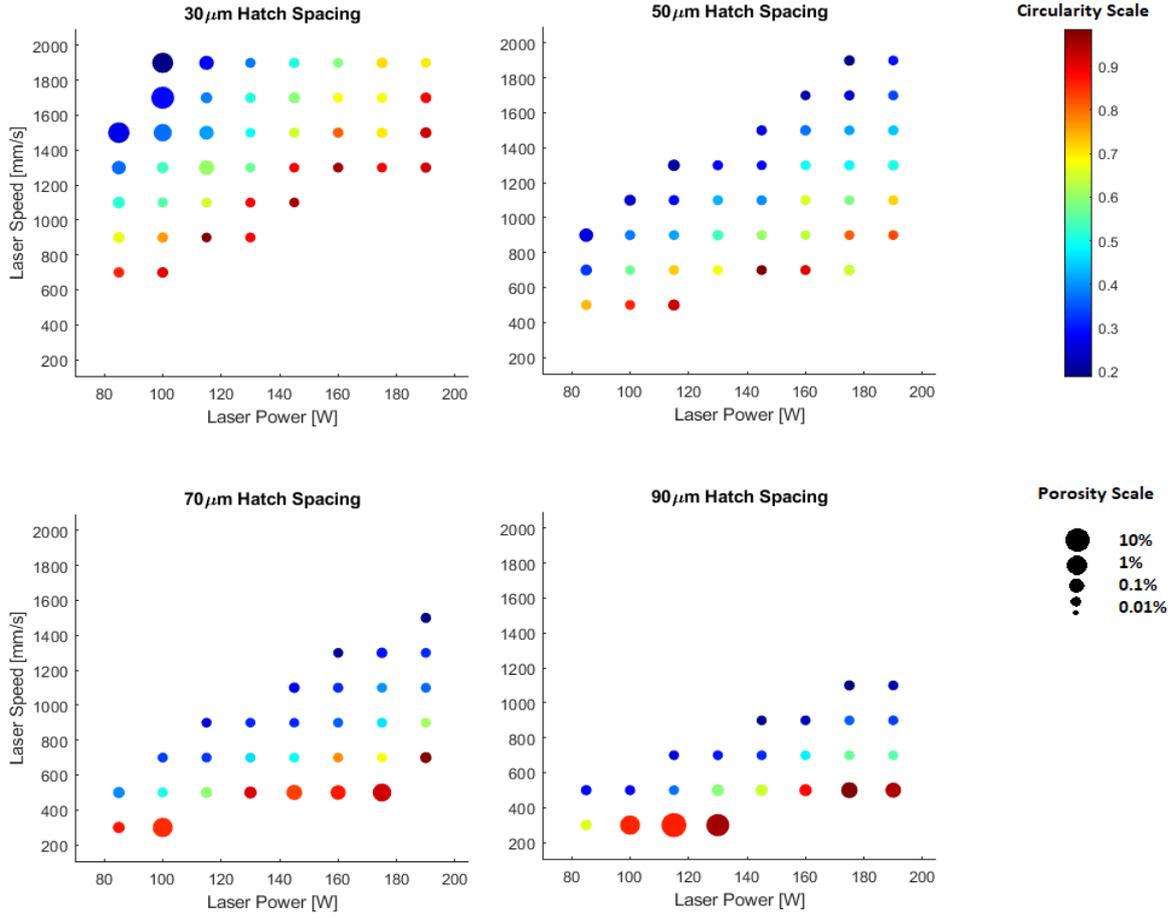


Figure 6.4

The porosity and circularity of each cube, across the entire parameter space

The normalised energy method of dimensionless analysis, as proposed by Thomas et al. [267] and discussed in a previous chapter, was applied to the results from this build. In this method, E^* is plotted against $1/h^*$, where

$$E^* = \frac{q^*}{v^* l^*} = \frac{AP}{2Vlr \rho c (T_l - T_0)}$$

and

$$h^* = \frac{h}{r}$$

In order to apply the method as rigorously as possible, the method proposed by Phan et al. [268] for calculating T_0 was applied:

$$T_0 = T_{pre} + \frac{0.3P}{2\pi\lambda\sqrt{(Vt_r)^2 + h^2}} \times \exp\left(-V\frac{\sqrt{(Vt_r)^2 + h^2} + Vt_r}{2\alpha}\right)$$

where

$$t_r = \frac{L}{2V}$$

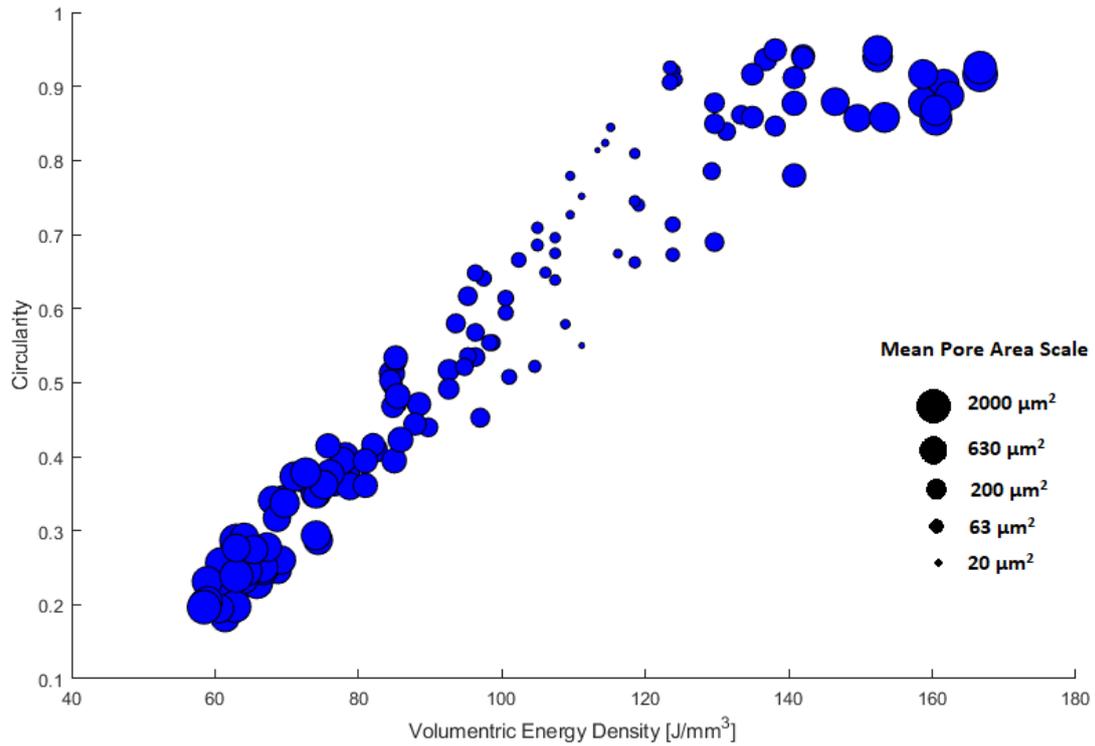


Figure 6.5

The mean pore area and circularity of the cubes, plotted against volumetric energy density

The numbers in Table 6.2 were used in these calculations.

Table 6.2

The values used in the dimensionless calculations

Number	Value	Number	Value
Laser Power (P)	Variable	Laser Speed (V)	Variable
Hatch Spacing (h)	Variable	Absorptivity (A)	0.45 [325]
Beam Radius (r)	35μm	Powder Density (ρ)	4.24 g/cm ³
Spec. Heat Capacity (c)	485 J/kg/K [328]	Liquidus Temp. (T _l)	1336°C
Layer Height (l)	30μm	Thermal Conductivity (λ)	30 W/m/K [145]
Thermal Diffusivity (α)	5 mm ² /s [329]	Scan Length (L)	4mm
Preheat Temp. (T _{pre})	25°C		

This yielded the graph shown in Figure 6.6.

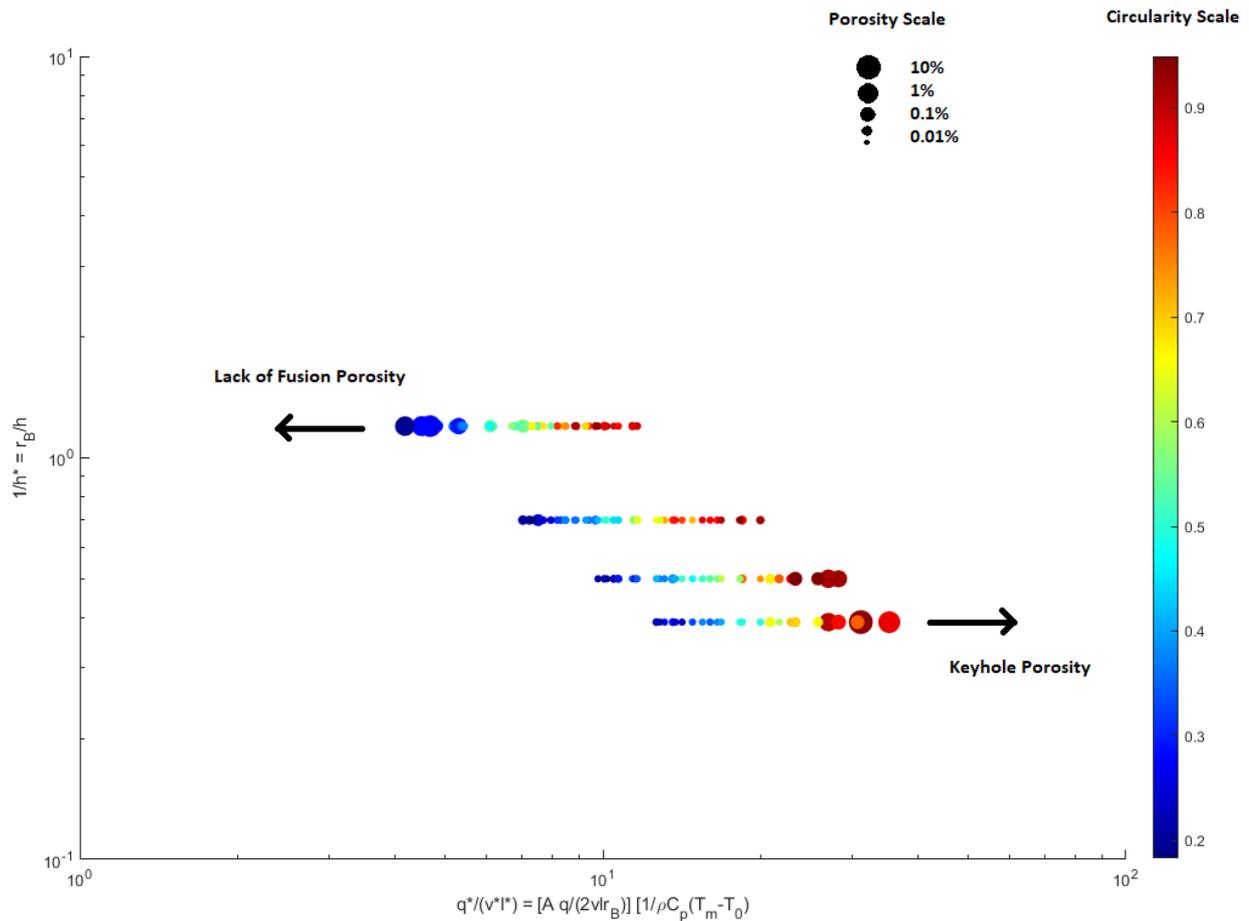


Figure 6.6

The circularity and porosity of the cubes, plotted on the same normalised energy density axes used by Thomas [267]

The regions at which low normalised energy leads to LoF, and the regions at which high normalised energy lead to keyholing, become very obvious in this graph. When this data is overlaid onto the chart produced by Thomas et al., it is clear that these results are in keeping with others from literature, as shown in Figure 6.7.

In Figure 6.6, the red dots are the data points from this experiment, with their sizes based on porosity, as per the previous scale. At both the lower and higher energy densities, where LoF and keyhole pores appear, the energy densities are consistent with similar defects in other alloys and experimental setups.

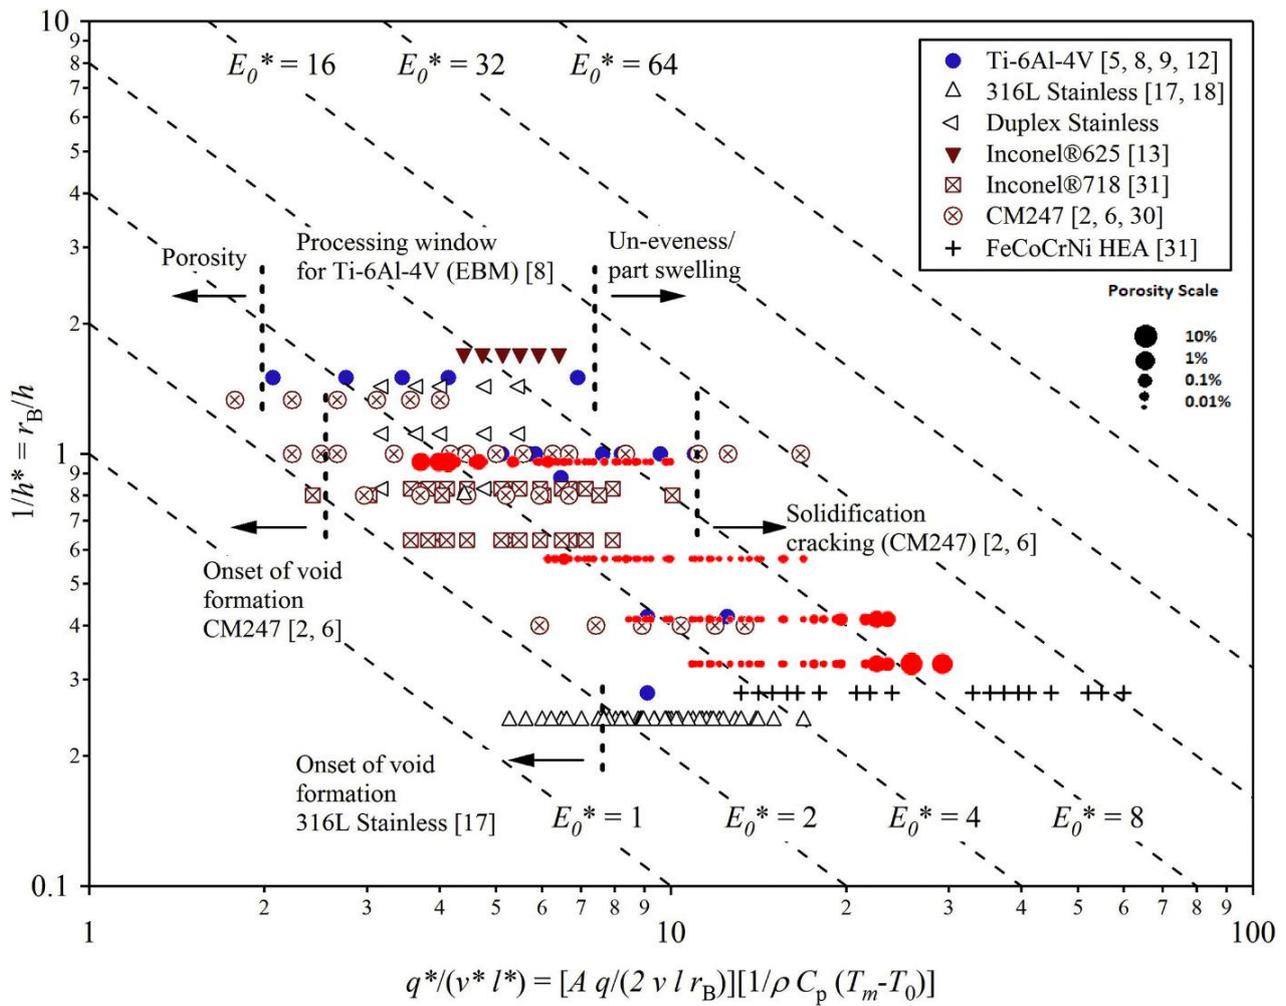


Figure 6.7

The results from this experiment (red dots), superimposed on the normalised energy density graph used by Thomas [267], with marker size denoting porosity.

6.1.3 - Selecting Parameters

When selecting preferred parameters for IN-718 builds to be monitored with the thermal camera, there are several criteria which should be considered, here ranked in order of importance:

1. Porosity should be as low as possible (to allow for full density with proper scan strategies)
2. Laser speed should be low, to allow for more recorded melt pool images
3. Hatch spacing should be small, for the same reason
4. Laser power should be low, to reduce spatter

In order to find the parameter set that best fits these criteria, the 60 that achieved densities of above 99.9% were considered. By sorting these sets first by velocity (lowest to highest), and then by hatch spacing (lowest to highest), the most ideal parameters set was found to be: 100W power, 500 mm/s speed, 50µm hatch spacing.

6.2 - In-Situ Monitoring and Deliberate Porosity

In the previous chapter, the ability of porosity to be predicted based on thermal camera data was explored. However, results were presented in terms of averages: the mean thermal signatures from a cube, and the total porosity of that cube. It was not possible to attempt to match individual pores, with the images of the melt pool that made them.

In this subsection, this attempt will be made. Two test pieces were designed, using machine parameters and geometries that were likely to introduce LoF porosity into specific points in one, and keyhole porosity into specific points in the other.

Each test piece was produced on the Aconity LAB from IN-718, and monitored with the thermal camera recording at a frame height of 512px, and a frame rate of 400 fps. Pieces were then brought to the ESRF in Grenoble, for XCT to be carried out on BM-18. A 195 keV monochromatic X-ray beam was used, with samples at a working distance of 10m. The beam size at the sample was 241mm × 12.1mm, resulting in a voxel definition of 4μm × 4μm × 4μm.

The thermal camera and XCT datasets were then cross-referenced, to match specific pores to specific thermal signatures. This was intended to test the hypothesis that the precise thermal signature of an individual pore could be detected during that pore's creation.

6.2.1 - Creating Lack of Fusion Porosity

A cuboid was built, with dimensions 35mm × 7.25mm × 4.5mm (XYZ). The previously identified machine parameters of 100W, 500 mm/s, and a 50μm hatch spacing were used. A 30μm layer height was used throughout. Inside the bulk cuboid, there were 45 channels running along the X axis, built with different laser powers. The speed and hatch spacing used in these channels was the same as in the bulk.

Above each channel, the power returned to the bulk setting of 100W. After a certain number of layers were built at the bulk parameters, a layer was recorded with the thermal camera. A channel was built, with each combination of the following three variables:

- Power used to build the channel (0W, 20W, 40W, 60W, 80W)
- Width of the channel (3, 5, and 7 hatches)
- Number of layers of 100W power, in between the top of the channel, and the 100W layer that was recorded (0, 1, and 2)

Both the bulk and the channels were built with bi-directional scanning, aligned to the X axis, with no rotation between layers. This ensured identical conditions for the building of channels, and also for the recording of layers. It was understood that linearly varying power with height could introduce a systematic bias due to the heating of the build chamber, but it was reasoned that the large height of the support structures would allow for a steady state to be reached before the channels of interest.

In Figure 6.8, the channels are shown in red, and the recorded layer is shown in blue:

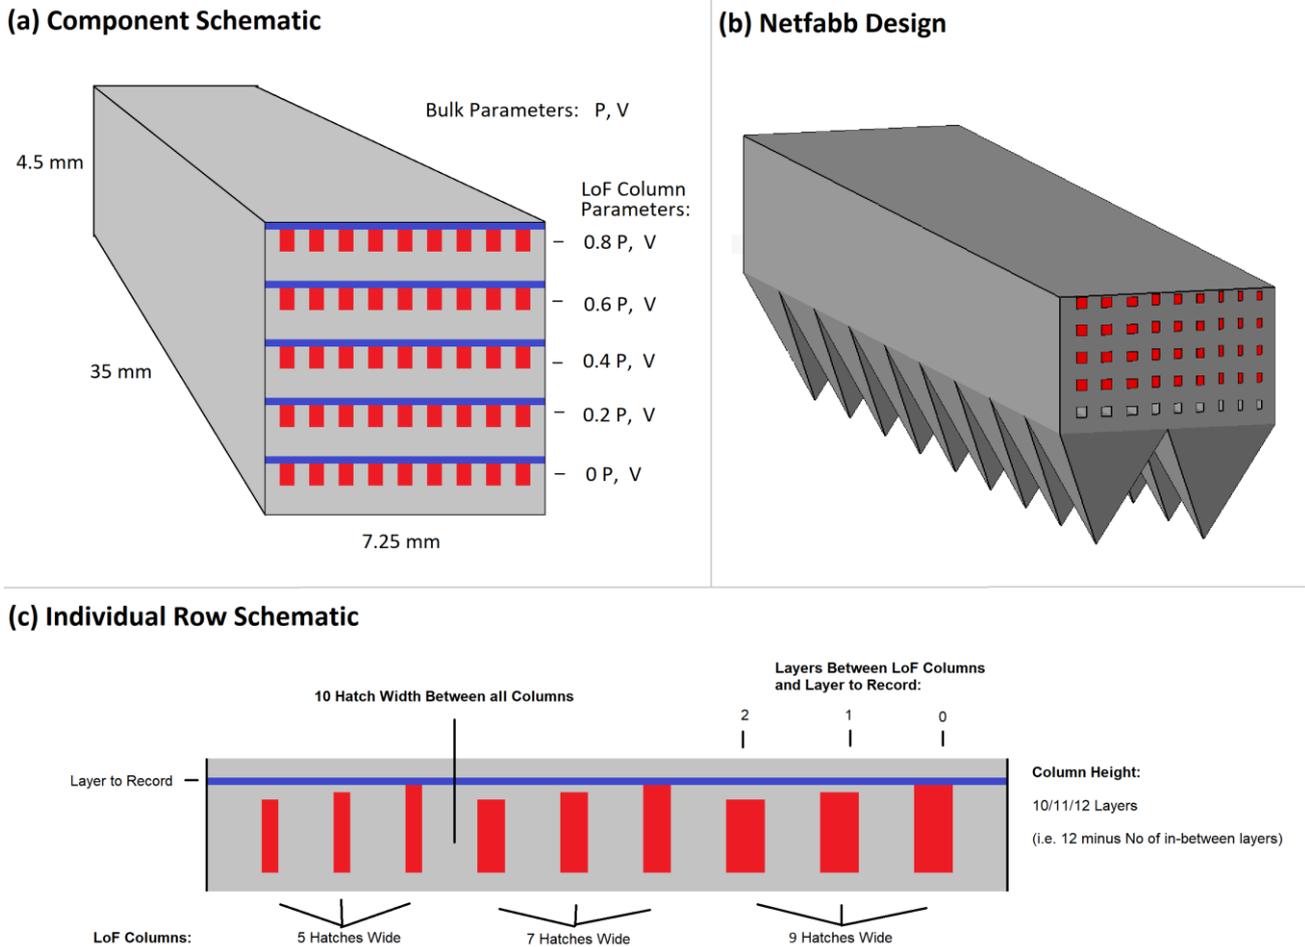


Figure 6.8

The placement of LoF channels within the test piece

6.2.1.1 - LoF in Thermal Images

Each recorded layer consisted of 143 hatches, of which:

- 80 scanned over fully solid material
- Three sets of five scanned over LoF material (with 0, 1, and 2 solid 'in-between' layers)
- Three sets of seven scanned over LoF material (with 0, 1, and 2 solid 'in-between' layers)
- Three sets of nine scanned over LoF material (with 0, 1, and 2 solid 'in-between' layers)

For each of the five recorded layers, melt pools were separated into these ten categories. Assuming that the 80 full-solid melt pools from each layer can be grouped together, this resulted in 46 total groups of melt pools. There were a minimum of 125 images in each group.

Each melt pool image was analysed using the same techniques described in previous chapters:

- The image was thresholded at 1,336°C
- The region with the greatest sum of pixel-temperatures was assumed to be the melt pool

- Any pixels outside this region were set to zero
- The REGIONPROPS function was used to extract the melt pool's stats

It was expected that, due to the lower thermal conduction when a melt pool is building with LoF porosity, the temperature and dimensions of melt pools would both increase, as: the hatch-width of channels increased; the number of in-between solid layers decreased; and the laser power with which the channel was made decreased.

In each of the following charts, the data is presented separated by LoF channel width, by solid in-between layers, and by LoF channel laser power. In each plot, the horizontal red line refers to the mean value of the 10,278 'solid' melt pools.

6.2.1.1.1 - Melt Pool Length Above the LoF Channels

The effect of LoF channels on melt pool length are shown in Figure 6.9.

As can be seen from the charts in row (1), no consistent trend appears to exist between the power with which the LoF channel was built (i.e. the extent of porosity), and the melt pool length. It is assumed that the decrease in length at the 5 hatch width, 2 in-between layers, is simply a coincidence, considering that no similar trend exists elsewhere.

As can be seen from row (2), melt pool length very consistently increased as channel width increases, across all variations of power and in-between layers. This trend matches the expectation that wider LoF channels would result in lower thermal conduction, which would in turn result in longer melt pools.

As can be seen from row (3), no consistent trend appears to exist between the number of in-between layers and the melt pool length, in the case of the 5-hatch-wide and 7-hatch-wide channels. In the case of the 9-hatch-wide channels, however, length appears to decrease as in-between layers increase, as expected. It may be the case that for the narrower channels, sufficient thermal conduction into the nearby bulk exists, to render any further changes in conduction caused by changes in the number of in-between layers, as negligible.

Effect of LoF Channel on Melt Pool Lengths

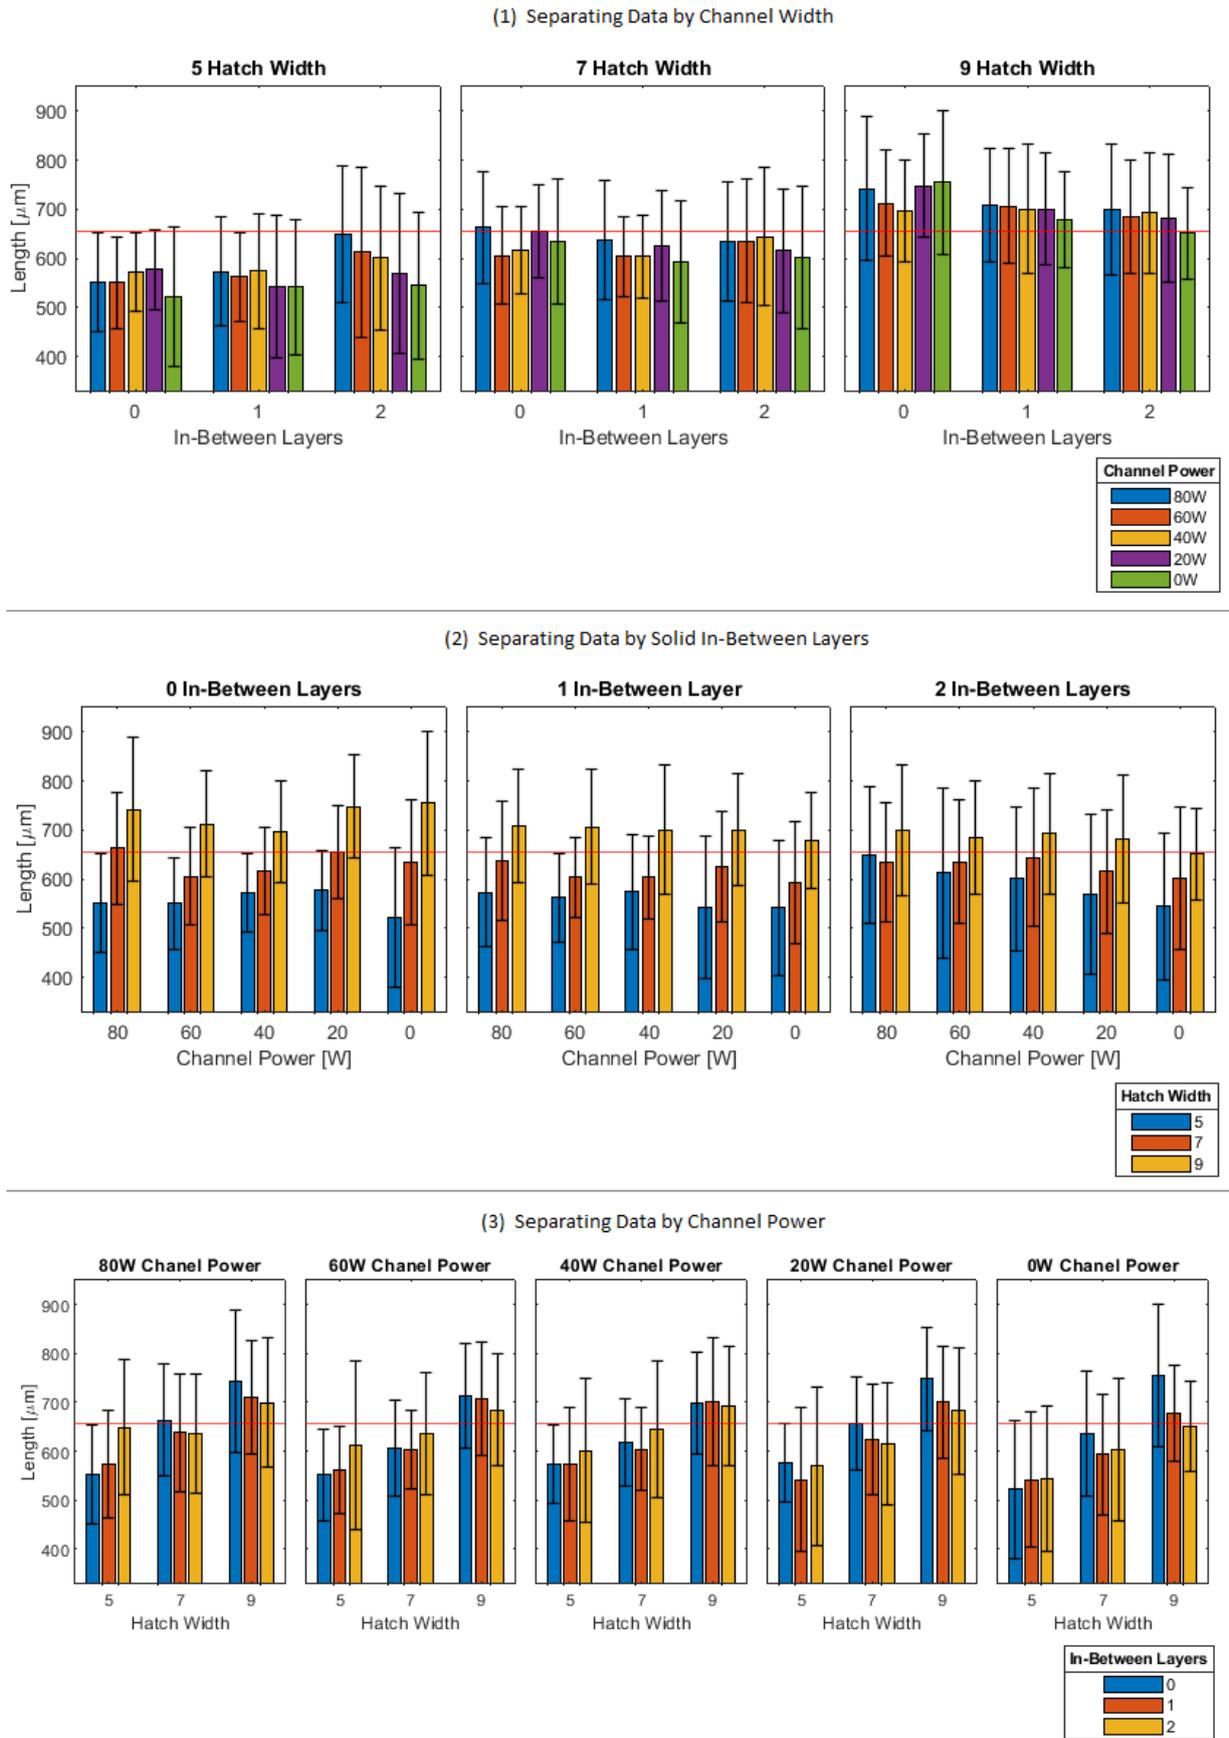


Figure 6.9

The effect of LoF channels on melt pool length

6.2.1.1.2 - Melt Pool Width Above the LoF Channels

The effect of LoF channels on melt pool width are shown in Figure 6.10.

As can be seen in row (1), a slight but consistent trend exists, whereby the width of the melt pools decrease as channel power is decreased. This result is contrary to the expectation that as channel power is decreased (and therefore LoF porosity increases), the decreasing thermal conduction would cause melt pool widths to increase. One possible reason for this may be the melt pool flowing into the voids of the LoF pores beneath it. Especially at the lower powers, where the material beneath the melt pool is largely composed of loose powder, the low density allows a lot of volume for the molten metal to sink into.

Again, there is a very consistent, and expected, trend for melt pool width to increase as LoF channel width increases, as can be seen from row (2).

As can be seen in row (3), melt pool width appears to consistently increase as the number of in-between layers increases, for the 5-hatch-wide channels only. Results from the 7-hatch-wide and 9-hatch-wide channels are inconsistent. It was expected that as the thickness of this fully dense layer between the LoF porosity and the melt pool increased, the increased thermal conduction would cause the pool width to decrease.

The fact that this contrary-to-expectations result exists only for the narrowest channels may be purely coincidental (especially considering how the error of the values far exceeds the differences between them), or may point to flaws in the initial expectations of melt pool behavior. Again, one possible reason for this counterintuitive behavior may be the fact that the molten liquid can sink into the porous substrate when the in-between layers are lower, thereby reducing the pool width.

Effect of LoF Channel on Melt Pool Widths

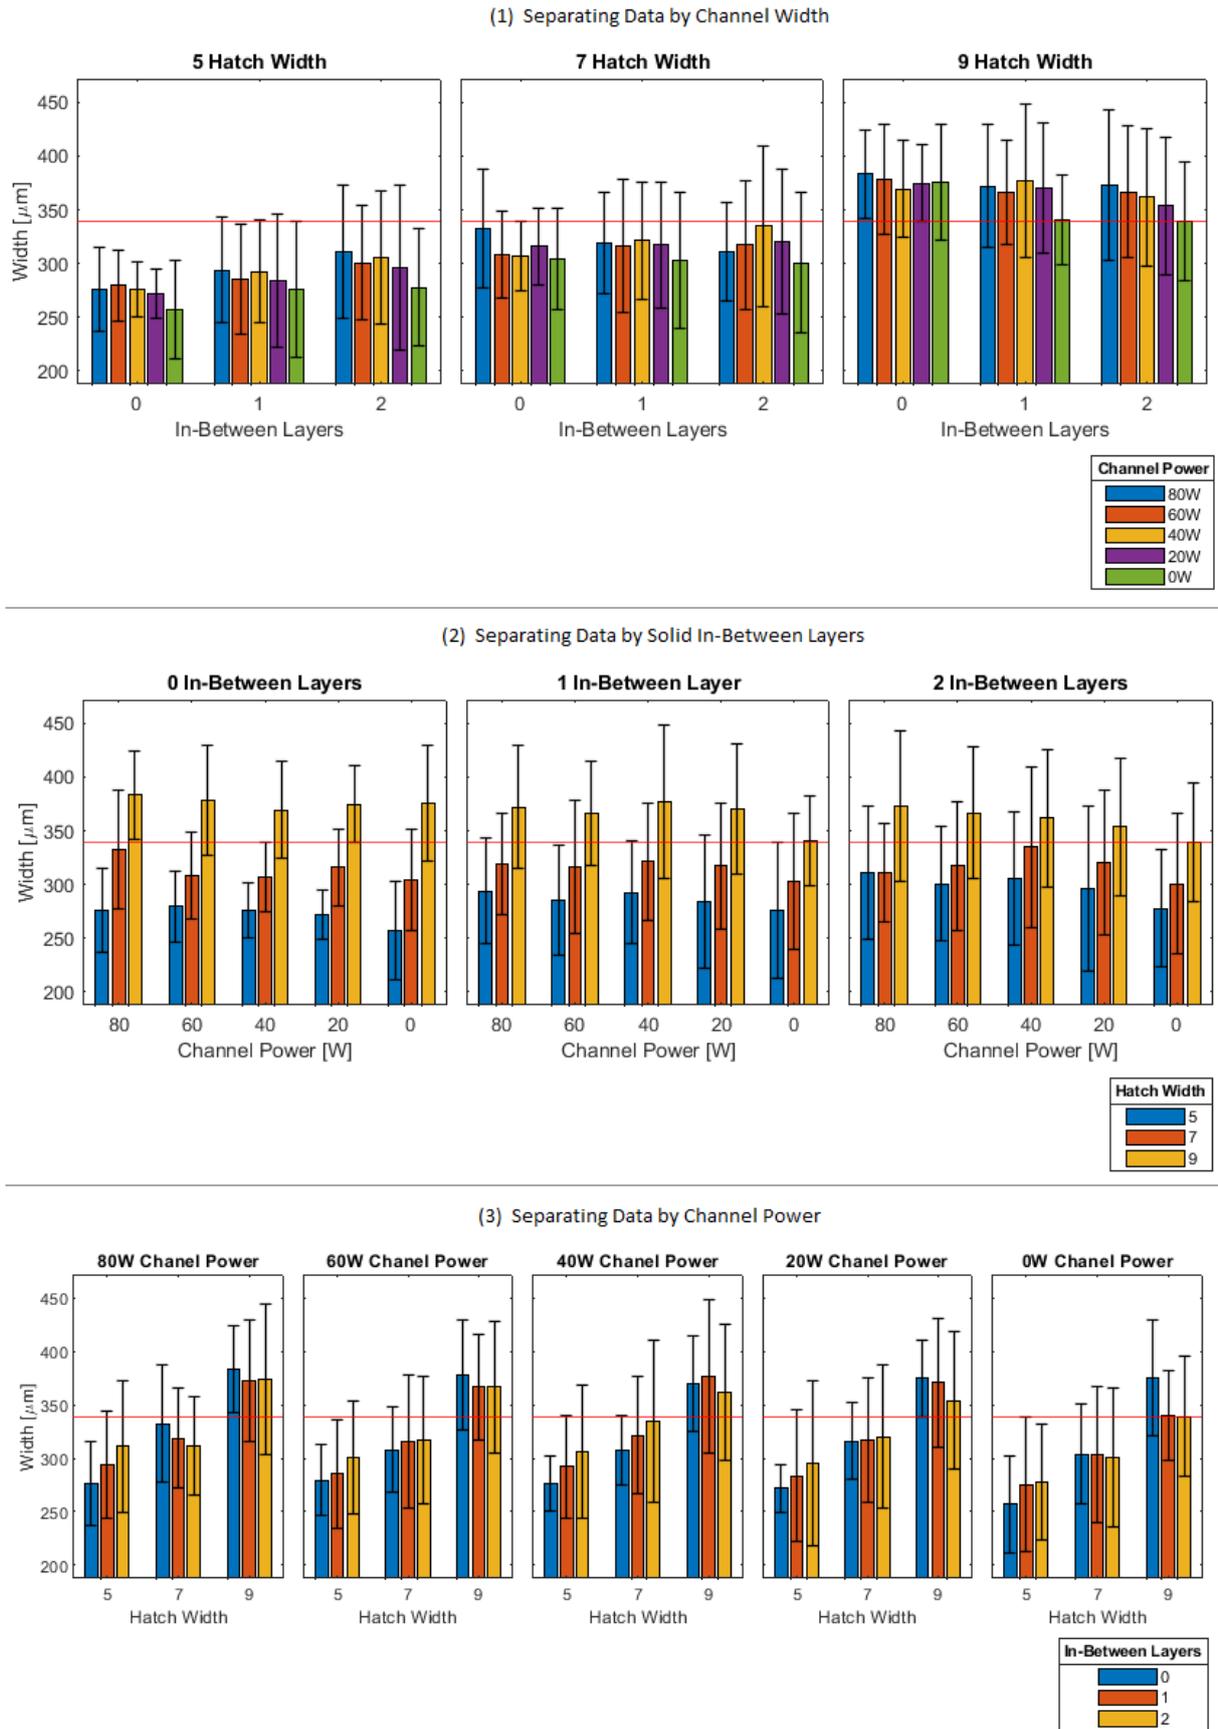


Figure 6.10

The effect of LoF channels on melt pool width

6.2.1.1.3 - Melt Pool Peak Temperature Above the LoF Channels

The effect of LoF channels on melt pool peak temperature is shown in Figure 6.11.

As can be seen in row (1), there exists a very small trend for melt pool peak temperature to increase as channel power decreases (and therefore LoF porosity increases). This trend matches the expectation that as thermal conduction decreases, melt pool temperatures will increase. It must be noted, however, that the trend is only visible in the 7-hatch-wide and 9-hatch-wide cases, and even then only in the 1 and 2 in-between layer cases, and also the increases in peak temperature are significantly smaller than the errors. Therefore, it would be unwise to draw firm conclusions from this set of results.

Considering the peak temperatures of the melt pools at the 9-hatch-wide and 0 in-between layer channels, it is clear that in these cases, where LoF porosity is most extreme, peak temperature is far lower than in the solid bulk (red line). This finding strongly contradicts the initial assumption that melt pools producing LoF porosity, and therefore experiencing lower thermal conduction, will exhibit higher peak temperatures.

As can be seen in row (2), peak temperature consistently decreases as channel width increases. This finding also contradicts the initial assumption that melt pools above channels with lower thermal conduction, will exhibit higher peak temperatures.

As can be seen in row (3), peak temperature appears to consistently increase as the number of in-between layers increases, for the 7-hatch-wide and 9-hatch-wide channels. The results for the 5-hatch-wide channels are inconsistent. This finding appears to again show that the initial assumption (that more LoF would cause higher peak temperatures) was the opposite of what exists in reality.

These three sets of findings all suggest that the initial assumption about porosity and peak temperatures, was incorrect. Two possible conclusions may be drawn from this finding:

1. As thermal conductivity decreases, melt pool peak temperature decreases, rather than increases as per assumptions. This conclusion would contradict much of the knowledge already established in literature [330, 331]. It could be possible that, due to these specific experimental conditions, the reason for such a decrease is due to increased convective or Marangoni flow within the melt pool. This possibility also contradicts established knowledge that it is melt pools in the keyhole regime that experience greater flows [184, 224].
2. Thermal conductivity is not actually decreasing as channel hatch width increases, as intended in the experimental setup, but is in fact increasing. The mechanism by which a melt pool's thermal conductivity could increase, as the width of the pore-filled channel it builds upon increases, is possibly due to the molten material sinking into the porous substrate, as described previously. Methods to test this possibility will be discussed in Section 8.3.1.

It is believed that the latter of these reasons is more likely. Further experimentation is needed to find if this sinking is the reason for the effect. This will be discussed in the Further Work chapter.

Effect of LoF Channel on Melt Pool Peak Temps

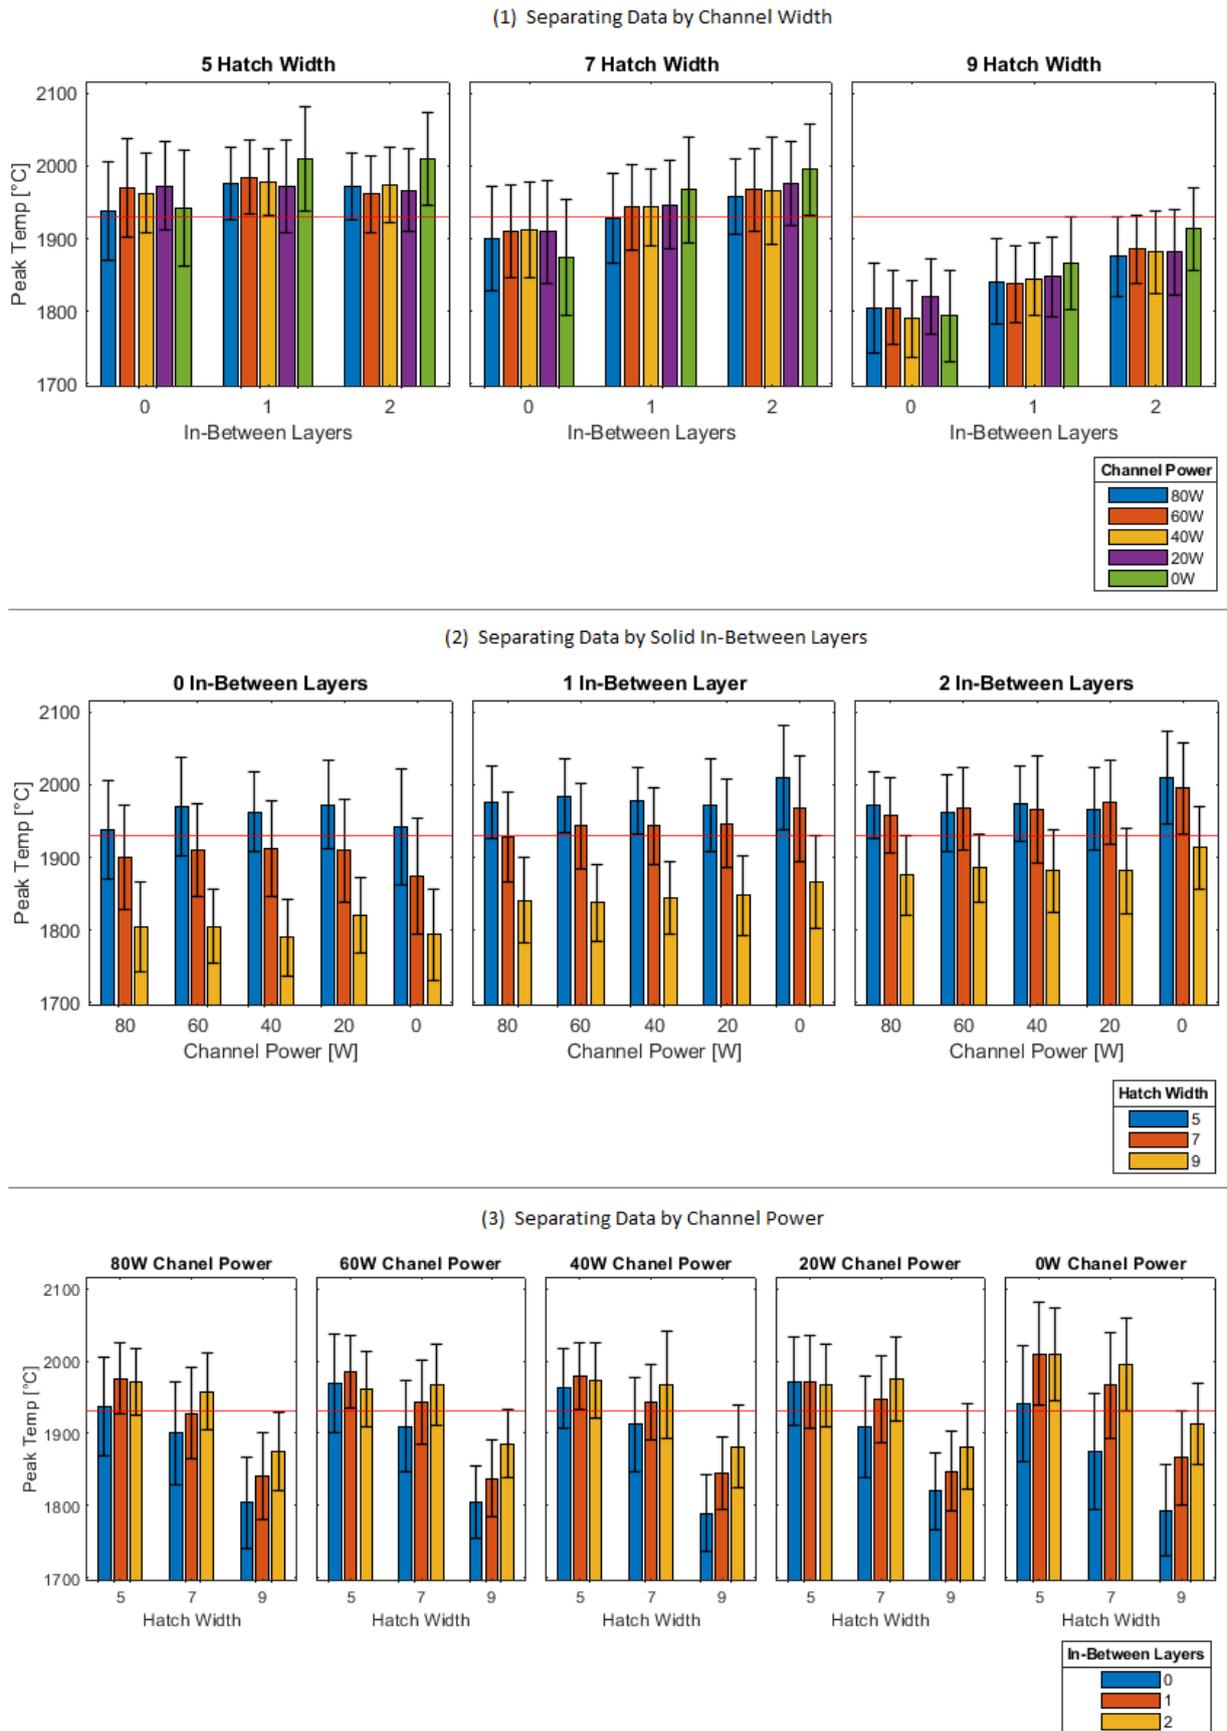


Figure 6.11

The effect of LoF channels on melt pool peak temperature

6.2.1.2 - LoF in XCT Data

The reconstructed data acquired at the ESRF consisted of 11,400 TIF files, each with size 8029×8029 pixels. Figure 6.12 shows an example TIF image, cropped to show only the 2D slice of the printed part.

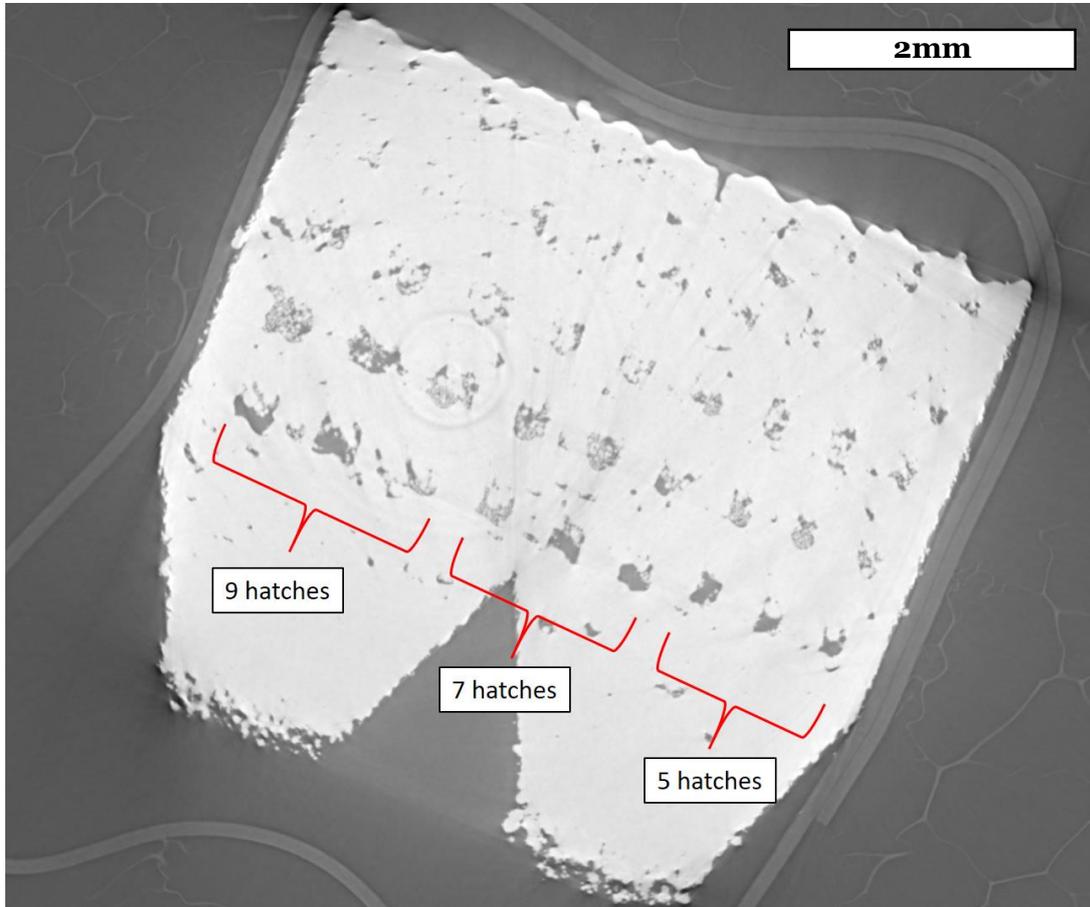


Figure 6.12

An example TIF image of the LoF channels, returned from XCT at the ESRF. Note that the concentric rings and faint lines visible are artefacts of the XCT process ('shadows' of material the X-ray has passed through), rather than actual features of the sample.

Due to the placement of the sample within the XCT setup, the image is reversed compared to the diagrams shown previously. As can be seen in this image, the channels are ordered 9 hatch width, then 7, then 5.

The first step in analyzing the images was to identify the 45 LoF channels. The four corners of the space that spans the channels were manually selected, for the first and last TIFs in the stack. By interpolating over the height of the entire stack, it was then possible to estimate the pixel coordinates of the channel-space for all images. By dividing this space into 5 rows and 45 columns using the channel-width ratios, it was possible to draw a rectangle around the predicted location of each of the LoF channels. Figure 6.13 demonstrates this, using the same TIF as shown in Figure 6.12.

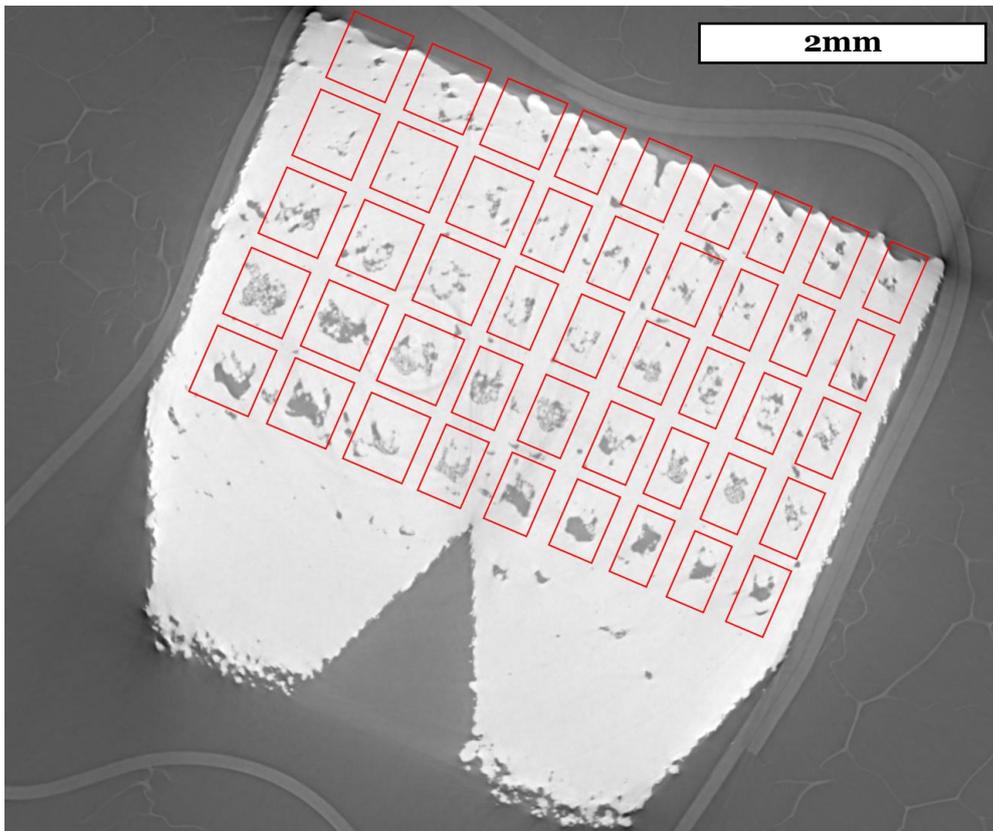


Figure 6.13

The same TIF image as previous, with the channels ready for individual extraction

Each rectangle was then extracted, rotated (via projective transformation using MATLAB's inbuilt `fitgeotrans` and `imwarp` functions), and saved as a separate TIF image. Figure 6.14, for example, shows the 215×185 pixel extract of the lower-left channel.



Figure 6.14

The lower-left channel in the previous image, extracted and ready for analysis

As can be seen in this image, a large section of the loose powder in the channel (9 hatches wide, 12 layers tall) has been melted by the laser scanning the layer on top.

6.2.1.3 - Comparing Thermal Images and XCT

For each of the 45 LoF channels embedded within the part, between 117 and 252 melt pool images were extracted, and 11,400 cropped and rotated XCT images were extracted. By calculating the X coordinates of the leftmost and rightmost pixel within each melt pool, it was possible to estimate which section of the channel's span that the melt pool occupied. It was then possible to map this section to positions within the 11,400 stack of XCT images, and therefore match up melt pools from the thermal camera with the XCT data that lay under them.

A MATLAB script was written to loop through every recorded melt pool, from each of the 45 channels, and extract the XCT images that they corresponded to. These XCT images were then thresholded, with any pixel with an intensity less than 85% of the image's peak intensity assumed to be a pore. This threshold was determined by visual inspection. Any pore that touched the image's top edge was discarded, as it was assumed to be the space above the uppermost layer of channels, rather than an actual LoF pore.

By stacking these binary images together, and using MATLAB's REGIONPROPS3 function, every pore within the 3D stack could be analysed. Using similar equations to those used in previous chapters, the volume-weighted mean volume and volume-weighted mean sphericity of each stack were calculated.

This resulted in the following metrics being available for each of the 8072 melt pools:

- Thermal camera data: melt pool lengths, widths, areas, aspect ratios, peak temperatures, and mean temperatures
- XCT data: volume-weighted (VW) mean pore volumes, VW mean pore sphericities, and number of pores

Each thermal camera metric was plotted against each XCT metric (i.e. 18 plots in total), to search for any trends within the data. Many of the plots did not show any meaningful correlations. For example, in the set of charts in Figure 6.15, it would be difficult to suggest that calculating the peak temperature, area, or mean temperature of a melt pool with the thermal camera, could reveal anything meaningful about the VW mean pore volume, the VW mean pore sphericity, or the number of pores, respectively.

In some cases, however, meaningful patterns could be interpreted from the charts. These patterns came in the form of boundaries, outside of which no data points fell. In the set of charts in Figure 6.16, there are clear regions that the melting process falls into.

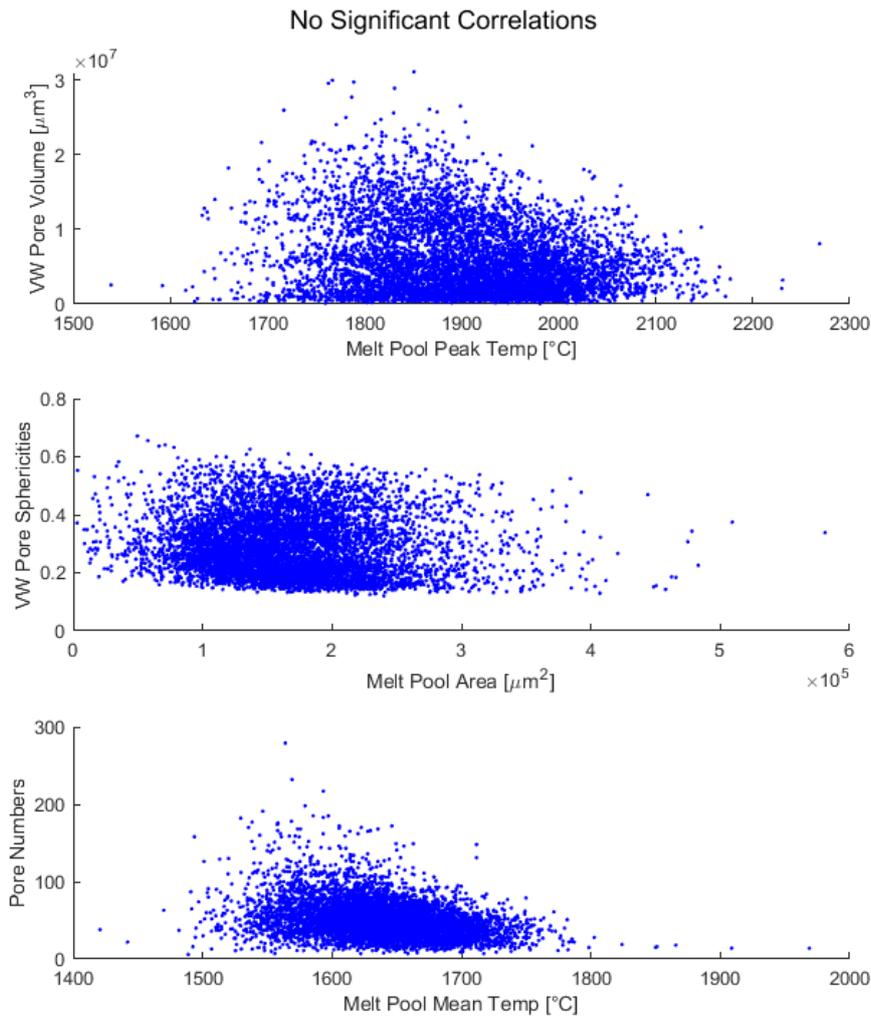


Figure 6.15

Examples of XCT and thermal image analysis, in which no significant trends were found

The upper limit of the number of pores being formed underneath a melt pool, and the upper limit of the VW mean volume of those pores, appears to increase as melt pool length, width, and area increase, and decrease as melt pool mean temperature increases. Therefore, it may be possible to monitor the thermal profile of a build, and predict the likelihood that the process has strayed into an unacceptable pore size/frequency region, based on the pool's dimensions and temperatures.

The fact that high porosity appears to correlate with greater melt pool area, as demonstrated in Figure 6.16, follows the logical reasoning that LoF porosity involves melt pools that cannot conduct their heat into surrounding solid material, and therefore are larger and hotter than melt pools building fully dense components.

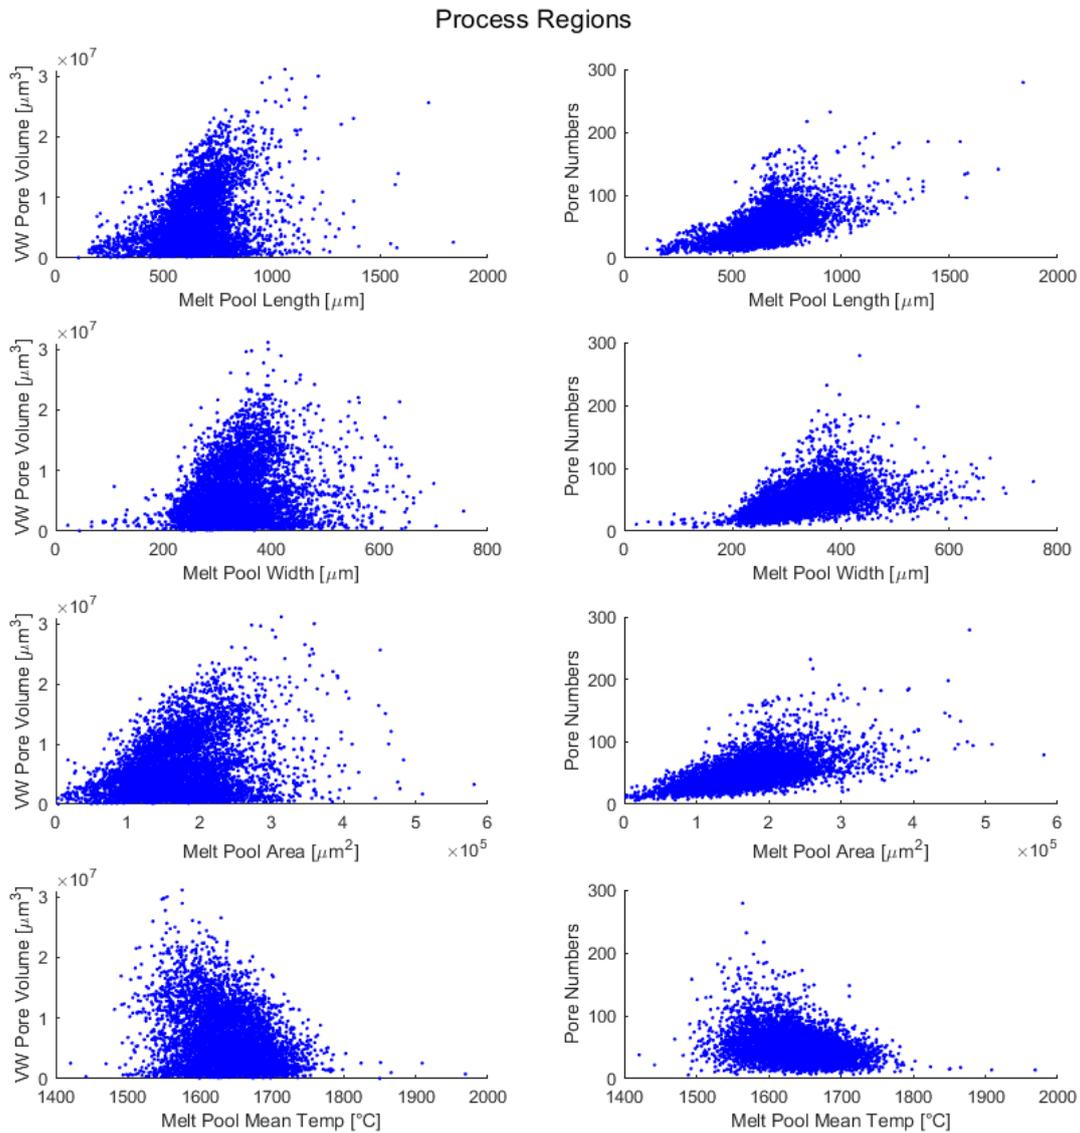


Figure 6.16

Examples of XCT and thermal image analysis, in which processing regions could be identified

Based on this data, it does not, however, appear to be possible to pinpoint with certainty the formation of specific pores, based on melt pool data thermal signatures alone. Several possible reasons exist for this:

1. LoF porosity arises due to insufficient overlap between subsequent layers and vectors. This lack of overlap can manifest in three forms. In increasing order of severity, these are:
 - a. Narrow slivers of non-overlap exist near the bases of the U-shaped pools, in the space where neither of the adjacent pools, nor the previous layer, managed to melt.
 - b. Delamination, where melt pools do not reach the previous layer, due to line energy being too low, and/or layer height being too large.
 - c. Adjacent vectors not meeting, due to line energy being too low for the chosen hatch spacing.

The first two of these forms are subsurface processes. Therefore, despite the effects that they may have on a melt pool's thermal conductivity, which in turn may affect the pool's dimensions and temperature profile, the process might be largely invisible to the thermal camera, which only monitors the powder bed surface.

2. A thermal signature may exist, but not be detectable by the thermal camera used in this experiment. Most melt pools appear in the images as only 10 to 30 pixels, in any dimension. This spatial resolution may be too small to detect the subtle thermal signatures that indicate LoF porosity.
3. Likewise, the temporal or thermal resolution of the camera may be insufficient to detect LoF porosity. The $1\mu\text{s}$ exposure time and the 400 fps rate may be too slow to capture individual pore formation, and the fixed-emissivity method of intensity-temperature conversion may miss the small temperature fluctuations associated.
4. The angle of the off-axis camera, which introduces foreshortening and distortion to the image, may reduce its ability to pick up slight variations in melt pool dimensions that accompany LoF creation.

Some changes in experimental setup, which could offset these limitations, will be discussed in the Future Work chapter.

6.2.2 - Creating Keyhole Porosity

For a given set of machine parameters, it is difficult to create different conditions under which keyholing porosity will or will not happen, via the altering of part geometry alone. The aim of this experiment was to obtain two sets of thermal images: the first of melt pools that were not forming keyhole pores, and the second of melt pools, using the same machine parameters, that were forming keyhole pores. It was predicted that altering the part geometry may allow this non-keyhole/keyhole transition, without altering laser power, laser speed, or hatch spacing.

Five long triangles, with a base (Y) of 2mm, a length (X) of 35mm, and a height (Z) of 4.5mm, were built. As before, these were built with the standard bulk parameters of 100W, 500 mm/s, and $50\mu\text{m}$ hatch spacing. A $30\mu\text{m}$ layer height was used throughout. All vectors were bi-directional, aligned with the X axis, and without rotation between layers.

Inside each triangle, aligned with the X axis, five line tracks were scanned, upon solidified material after their layers were finished. It was expected that as these lines progressed from the base of the triangle, towards the tips where thermal conduction was much lower, they would transition towards keyhole porosity, with the pores left trapped in the material beneath.

In order to preserve these pores, and not have them filled in by the next layer's melting, horizontal channels were built above each line track, that became narrower over the next ten layers. As the final line tracks were built on the final layer of the build, no channels were needed above them.

Twenty-five line tracks were scanned in total, encompassing all combinations of the following variables:

- Width of line track (1, 2, 3, 4, and 5 hatches)
- Power of the line track (100W, 120W, 140W, 160W, and 180W)

The speeds and hatch spacings used in the line tracks was the same as in the bulk. In each of the five layers containing line tracks, the tracks were scanned after the rest of the layer was built. The thermal camera recorded these layers.

In Figure 6.17, the line tracks are shown in red.

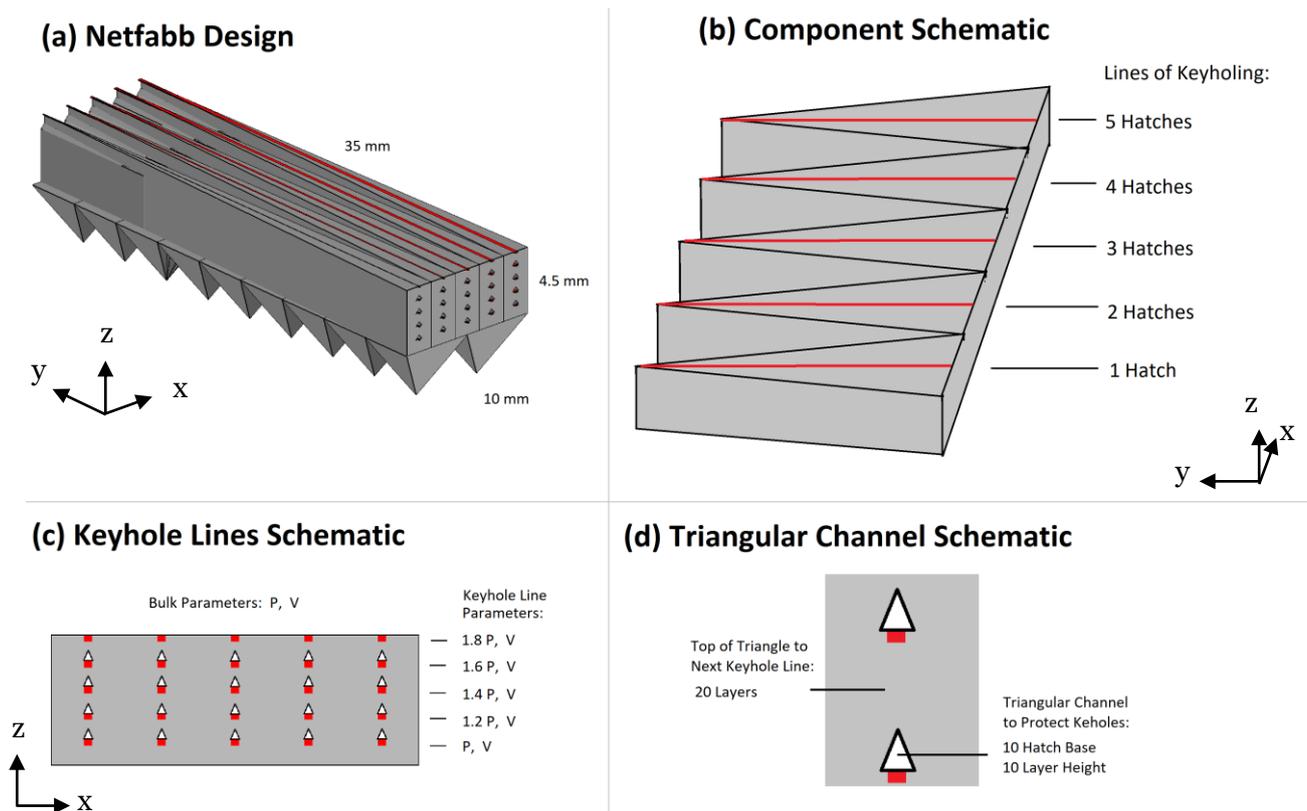


Figure 6.17

The keyholes channels deliberately placed in the part

6.2.2.1 - Keyholes in Thermal Images

In each of the 25 groupings (5 powers \times 5 hatch numbers), there were between 32 and 154 thermal images recorded. Each melt pool image was analysed using the same techniques described previously (thresholded at 1,336°C; highest temperature pixel assumed to be the melt pool; all others set to zero; REGIONPROPS function used to extract statistics).

It was expected that, as laser power increased, and as melt pools progressed along the X axis (from the wide part of the triangle towards the narrow), the pool's increased heat and decreased thermal conduction, would increase the likelihood of keyhole porosity.

The effect of increasing the number of hatches in the line track was difficult to predict. On one hand, it was believed that by increasing the number of hatches, there would then be more solidified material for subsequent hatches to conduct heat into, thus reducing the likelihood of keyhole porosity. On the other hand, it was believed that by increasing the number of hatches, there would then be additional heat in the triangle's tip for subsequent hatches, thus increasing the likelihood of keyhole porosity. Ultimately it was decided that the effect of increased hatches on keyhole likelihood was impossible to predict without experimental data.

It was expected that melt pools in the keyhole regime would exhibit shorter lengths, narrower widths, and higher peak and mean temperatures, than those in the conduction regime.

In order to allow for better analysis of the data, the progress of each melt pool along the X axis was not taken as a continuous value. Instead, the axis was divided into ninths, and separated into 3 bins as shown in the schematic in Figure 6.18.

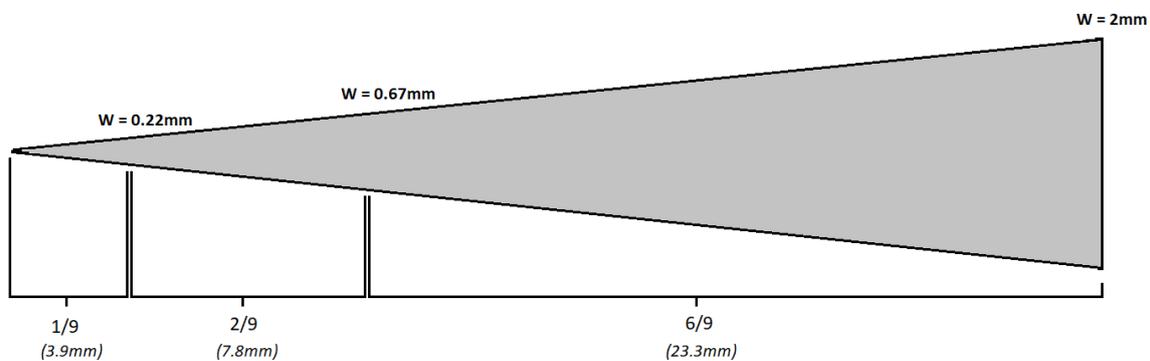


Figure 6.18

The divisions that separate the triangles into three parts

The results of the thermal images will now be presented, in terms of melt pool lengths, widths, peak temperature, and mean temperature.

6.2.2.1.1 - Melt Pool Length Above the Keyhole Line

The effect of parameters on melt pool length are shown in Figure 6.19.

As can be seen in row (1), no consistent trend exists, as the width of the triangle increases. This finding contradicts expectations that melt pools in the narrower section would exist in the keyhole regime, and therefore be longer.

The results from row (2) are both expected and trivial: as laser power decreases, melt pool length decreases too.

As can be seen in row (3), no consistent trend exists in the narrow section of the triangle, as line width increases. In the mid and wide sections of the triangle, where the keyhole regime was less expected, melt pool length appears to slightly increase as the number of hatches in the line increases.

Effect of Parameters on Melt Pool Length

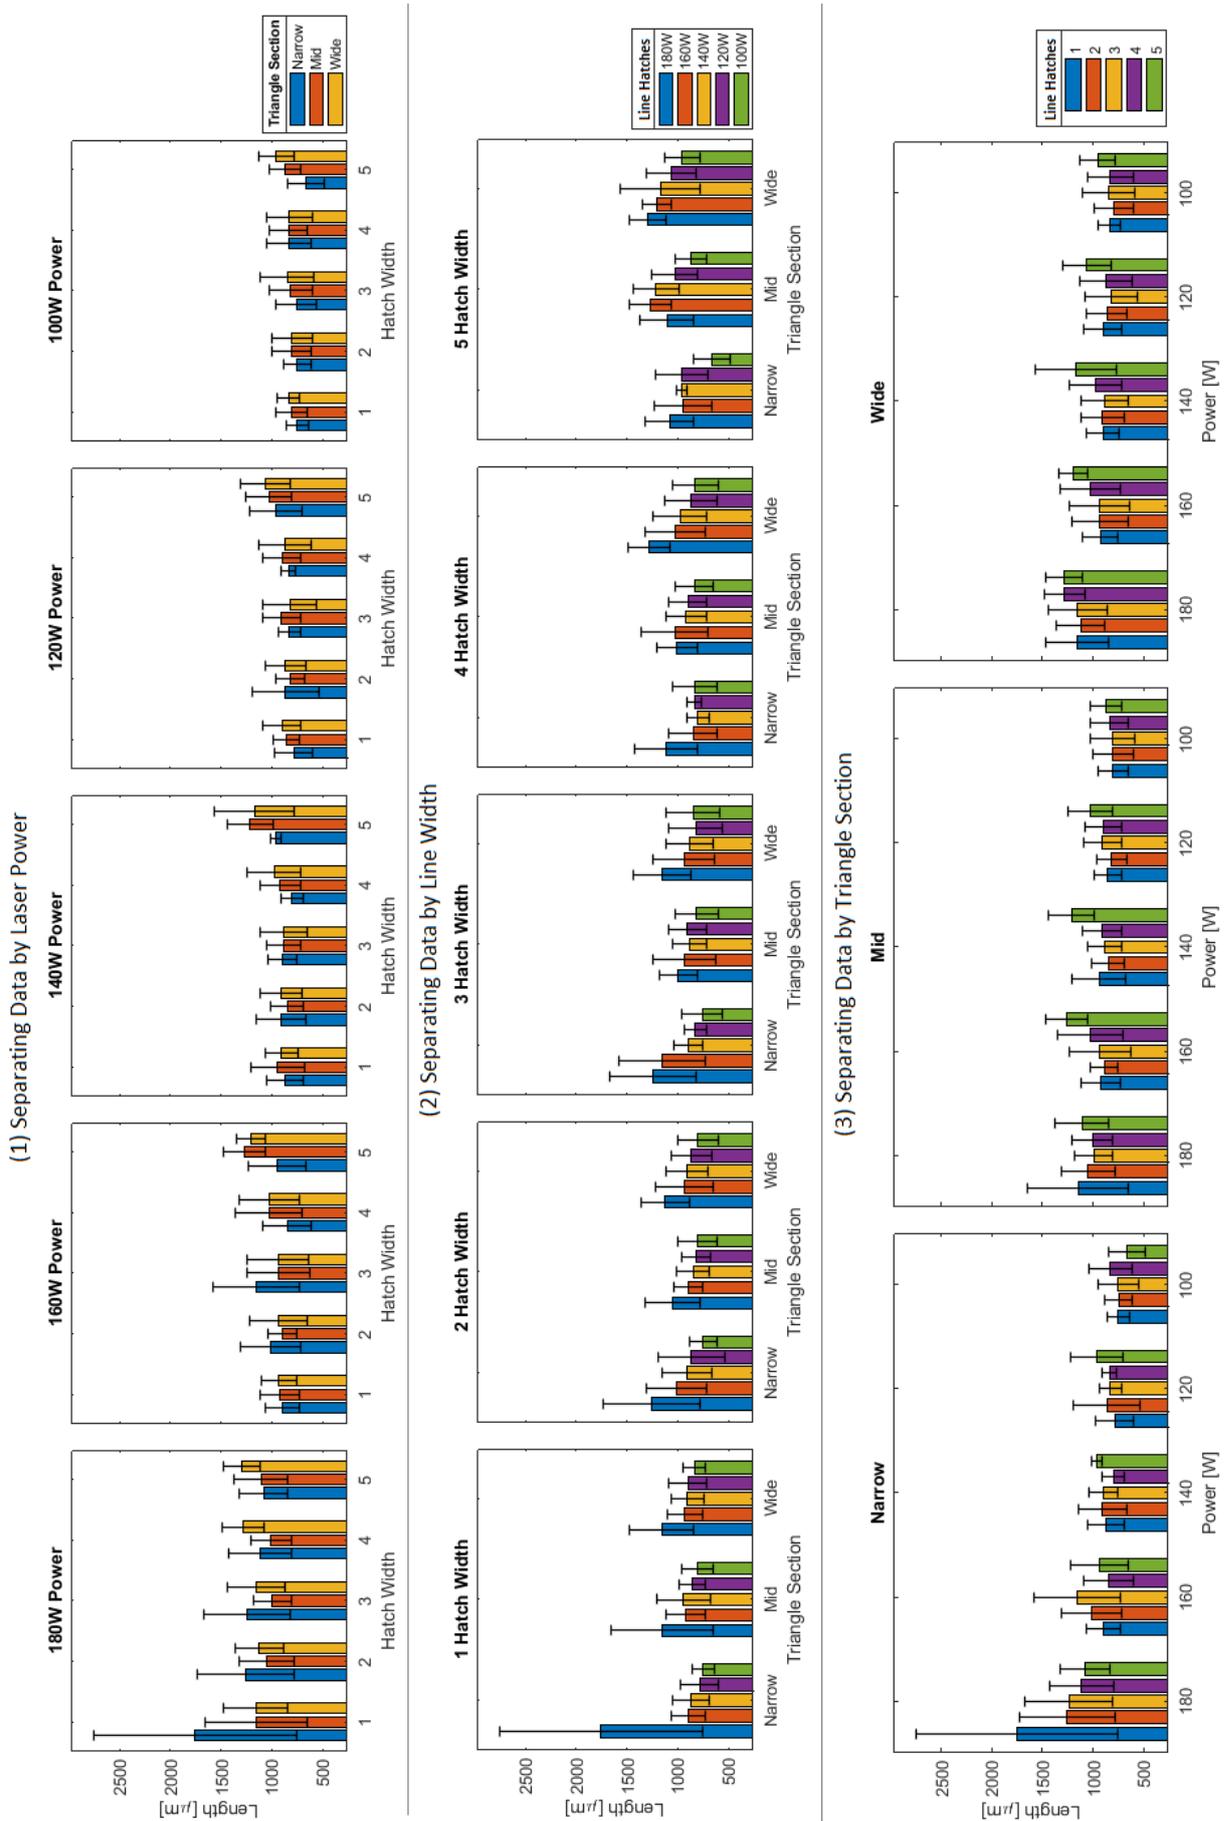


Figure 6.19

The effect of parameters on melt pool length

This latter result contradicts the explanation that (assuming keyholing is not present) increasing line width will increase thermal conduction, and therefore decrease melt pool length.

A possible explanation for this apparent increase in melt pool length, as the heat-conducting ability of the substrate increases, is the fact that changes convective flow might be playing a role. As the triangle section widens, and its ability to draw out heat from the melt pool therefore increases, the temperature gradients within the pool may increase. This, in turn, will create convective flows, churning the melt pool, thereby possibly delaying the amount of time taken for full solidification to occur.

6.2.2.1.2 - Melt Pool Width Above the Keyhole Line

The effect of parameters on melt pool width are shown in Figure 6.20.

As can be seen in row (1), melt pool width consistently increases as the width of the triangle section increases. This finding strongly contradicts expectations that the keyhole regime (and therefore wider melt pools) would be more prevalent in the narrower sections. This finding suggests two possibilities:

1. Melt pools in the keyholing regime are actually narrower than those in the conduction regime. It is uncertain the extent to which this finding would contradict established knowledge. While literature consistently shows keyhole melt pools to be wider, there exist no studies (that the author is aware of) that compare keyhole and conduction regime melt pools that were produced with identical machine parameters, as had been attempted in this experiment.
2. The melt pools at the narrow section of the triangle were actually less likely to exhibit keyhole regime melting than conduction regime. The mechanism by which this could be possible is unknown.

Again the results from the charts on row (2) are expected and trivial: as laser power decreases, melt pool width consistently decreases.

The narrow section results in row (3) appear to consistently show a decrease in melt pool width, as the width of the scan line increases. This finding suggests that the increased width of the line causes increased thermal conductivity, which in turn decreases the width of the melt pool.

Effect of Parameters on Melt Pool Width

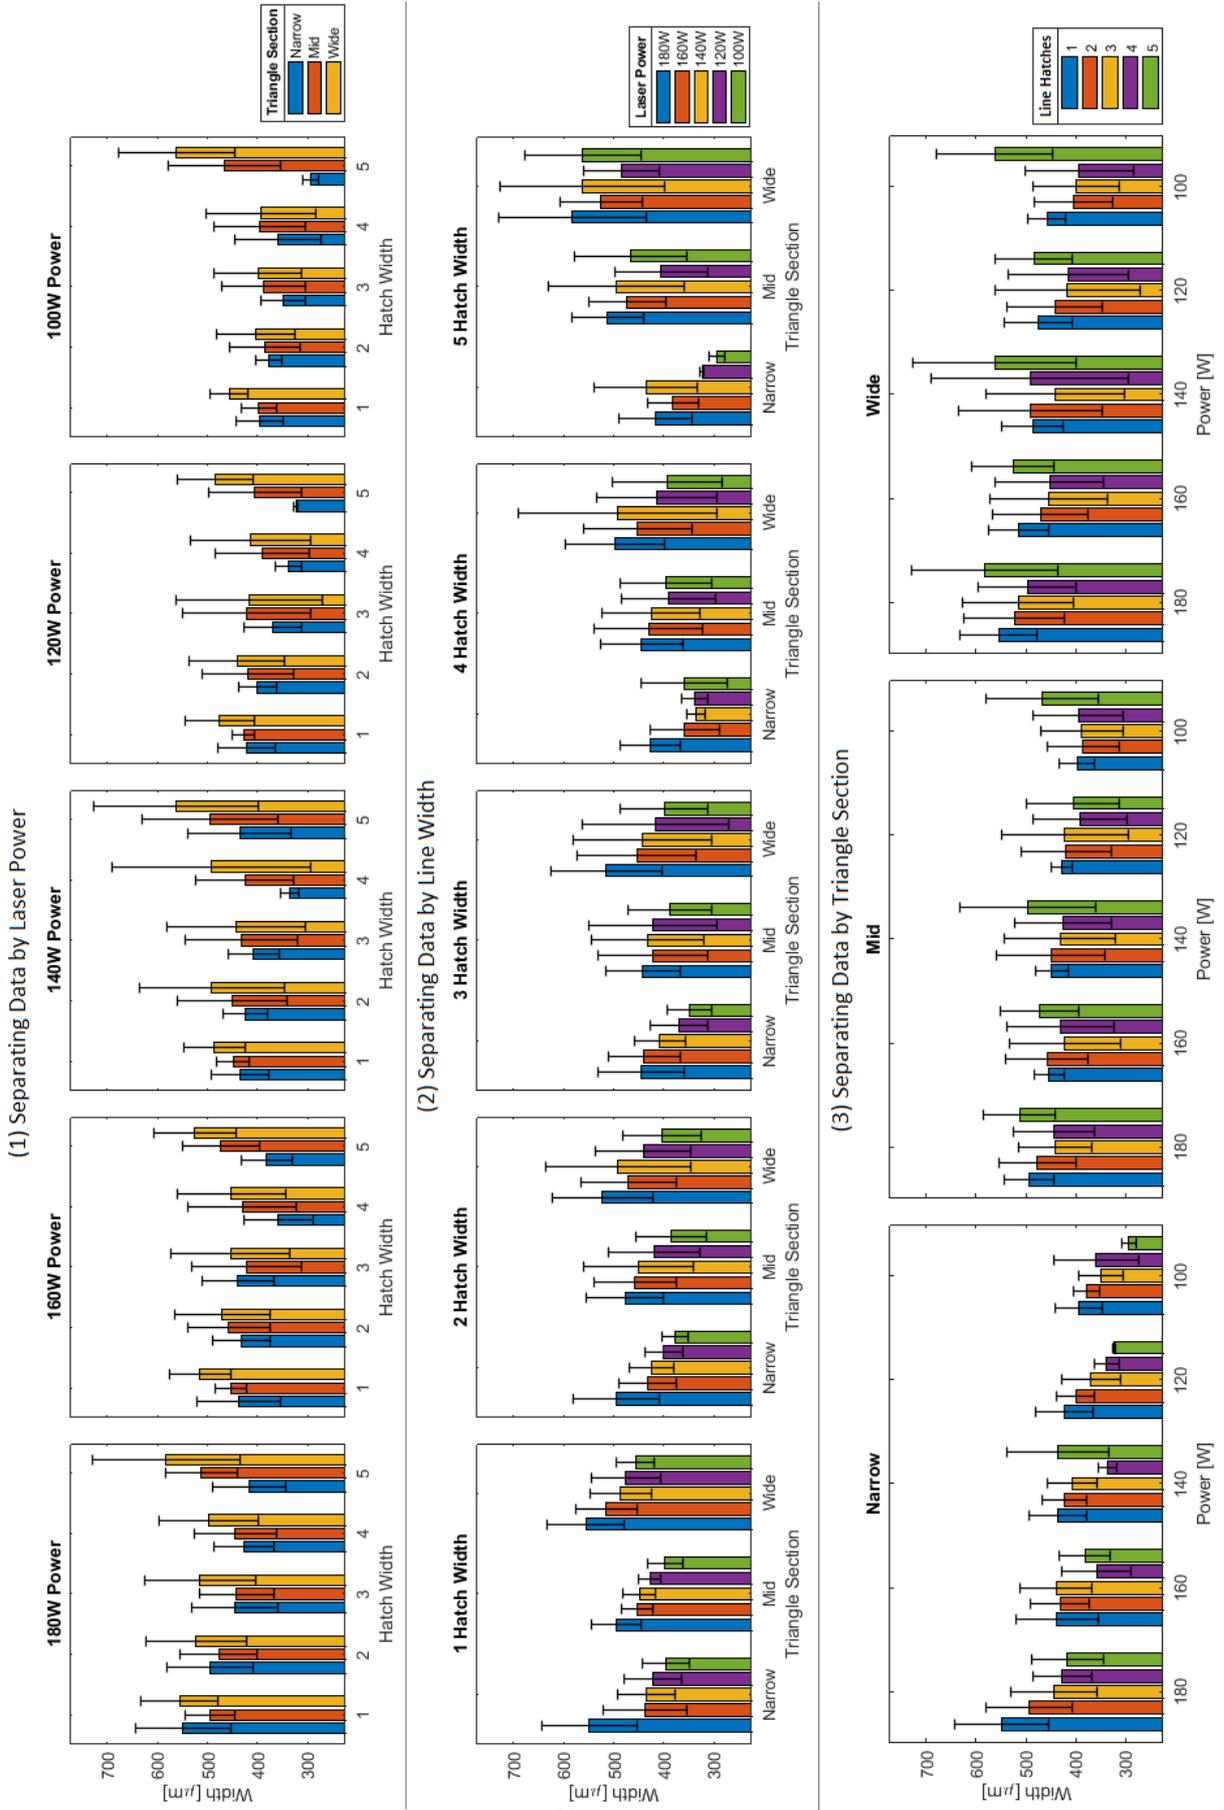


Figure 6.20

The effect of parameters on melt pool width

6.2.2.1.3 - Melt Pool Peak Temperature Above the Keyhole Line

The effect of parameters on melt pool peak temperature are shown in Figure 6.22.

As can be seen in row (1), no clear trend exists in the response of the melt pool's peak temperature, as the width of the triangle section changes. This result is contrary to expectations that as the width of the triangle section increases, and therefore the substrate's heat-conducting ability increases, melt pool peak temperature should decrease. One possible reason for this lack of decrease is that, while the increased heat conduction reduces the temperature of the pool periphery, it has limited effect on the temperature at the laser incidence spot. Another possible reason is that as the increased conduction cools the periphery of the pool's surface, this in turn increases Marangoni flow, which can force hotter molten material upwards from the pool's core, back into the path of the laser.

Again, the results of row (2) are expected and trivial: as laser power decreases, the melt pool's peak temperature decreases too.

As can be seen in row (3), melt pool peak temperature consistently increases as the line width increases, across all triangle sections. The fact that this trend holds true at even the widest triangle section, suggests that it is the result of accumulated heat: that as the number of hatches increases, the temperature of the material in the path of the melt pool increases, thus increasing the melt pool's peak temperature.

6.2.2.2 - Keyholes in XCT Data

As was the case with the LoF XCT images, the reconstructed keyhole XCT data acquired at the ESRF consisted of 11,400 TIF files, each with size 8029×8029 pixels. Figure 6.21 shows an example TIF image, cropped to show only the 2D slice of the printed part.

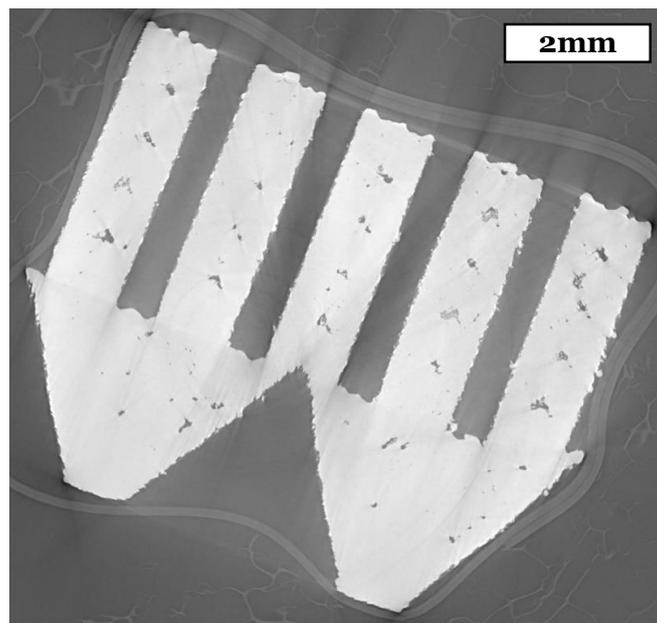


Figure 6.21

A TIF image of the keyhole channels, from XCT at the ESRF

Effect of Parameters on Melt Pool Peak Temperature

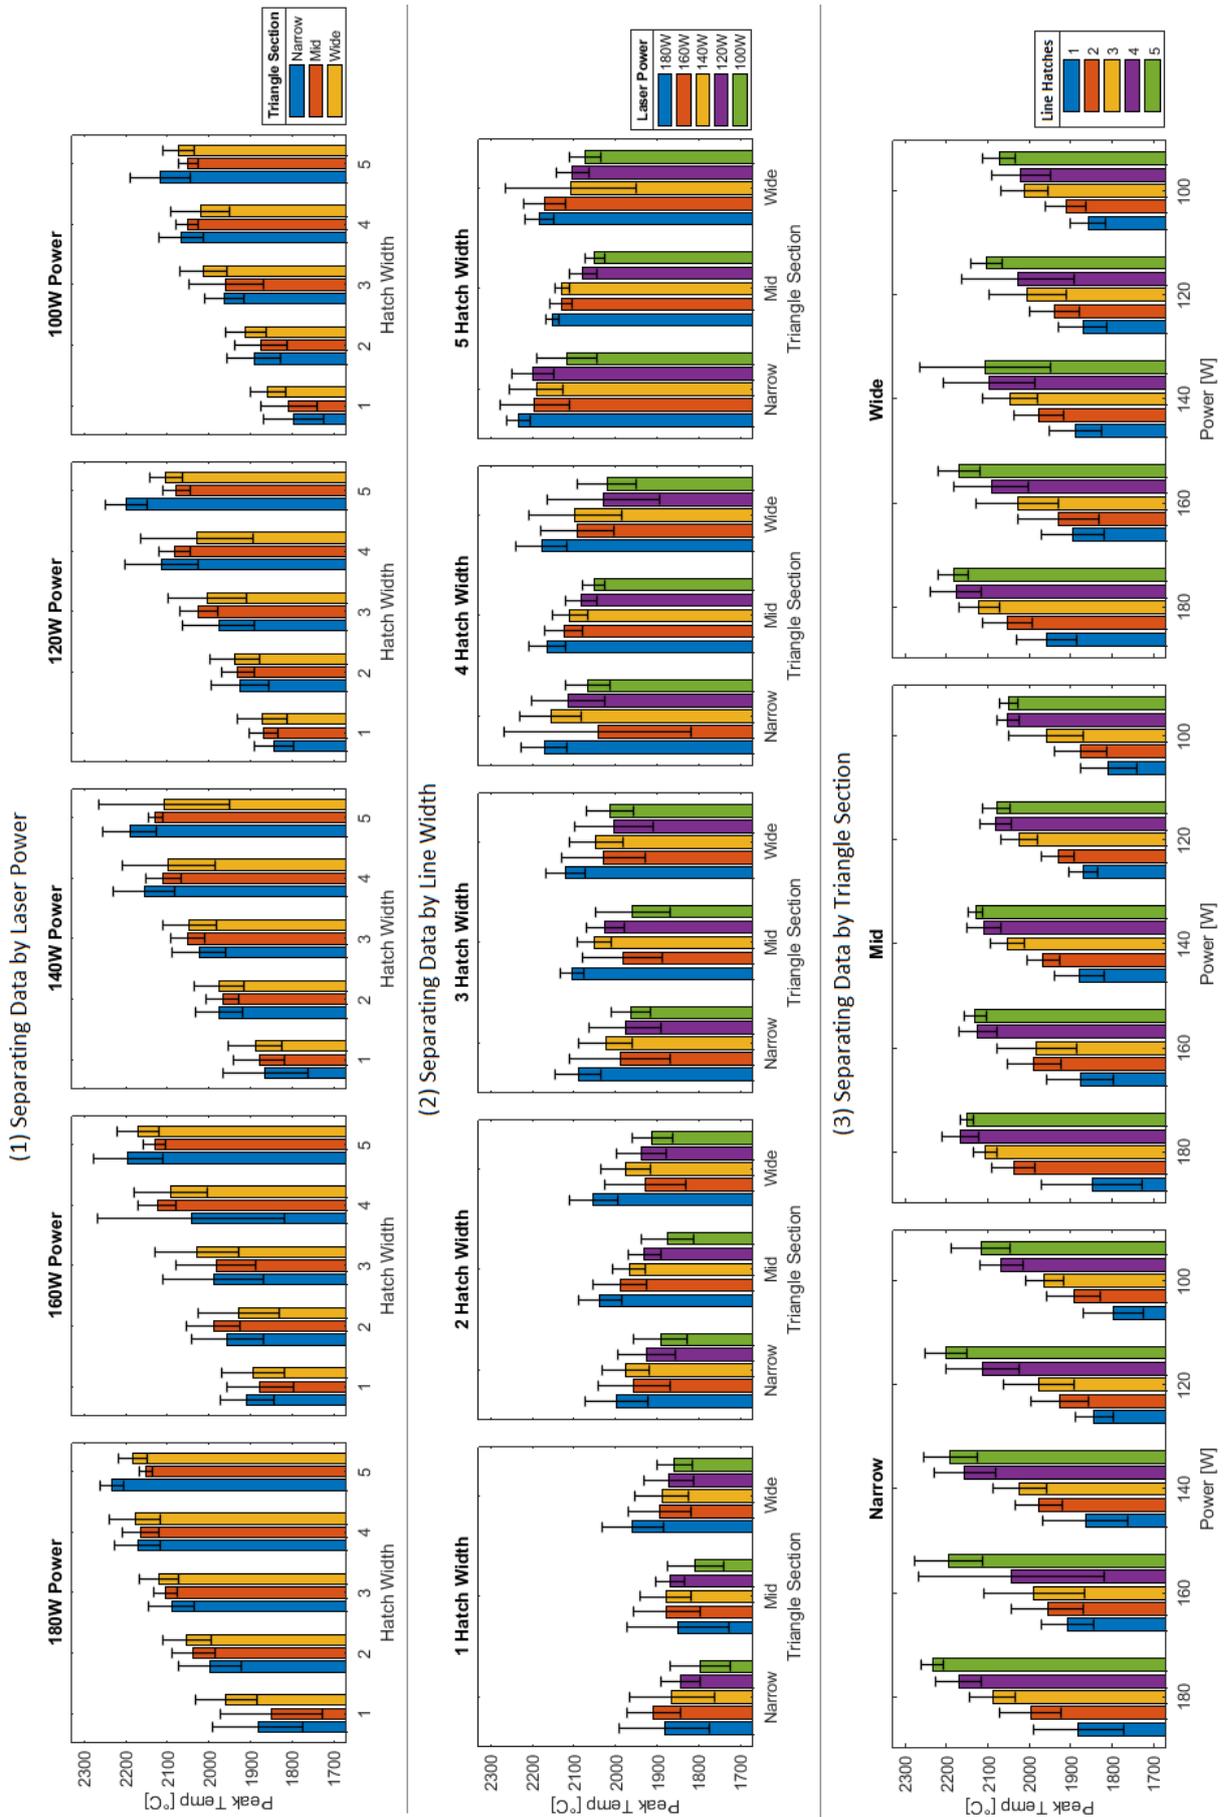


Figure 6.22

The effect of parameters on melt pool peak temperature

The first step in analyzing the images was to identify the 25 channels, beneath the line tracks, into which keyholes would form. This was achieved by manually selecting the four corner line tracks in the first and last image, and interpolating over the entire stack of TIFs, to assign corner points to each image. Then, by dividing the space between these points in each image into rows and columns, it was possible to assign a 100×150 pixel rectangle around the predicted location of every channel, in each image. Figure 6.23 shows this demonstrated for the same 2D slice as shown in the previous image.

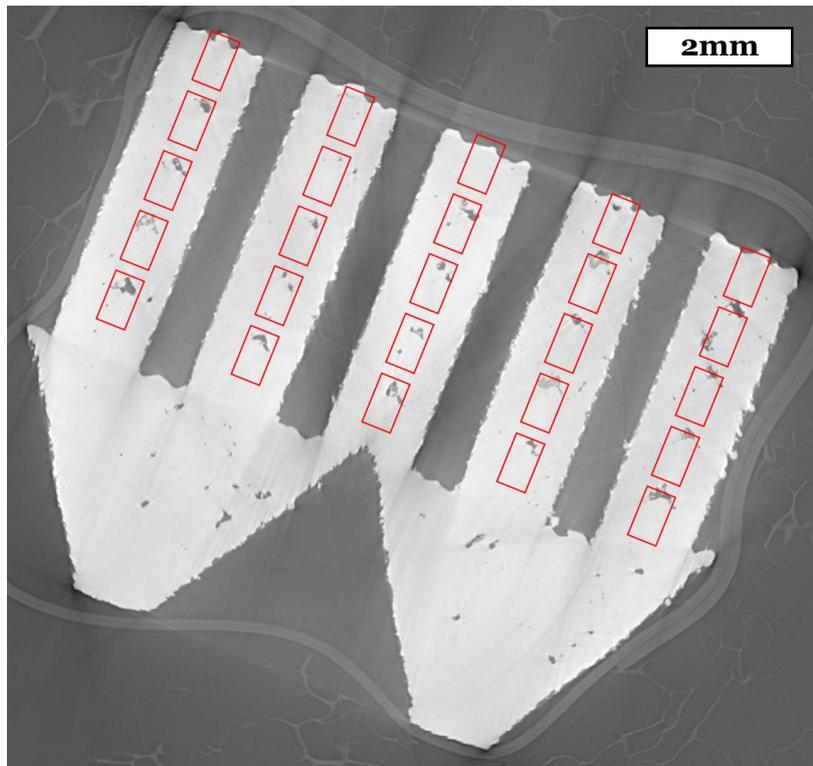


Figure 6.23

The same TIF as in Figure 6.21, but with the keyhole channels ready for extraction

Each of these 25 rectangles were then extracted, rotated, and saved as separate files. Figure 6.24 shows the extracted lower-left channel from Figure 6.23.



Figure 6.24

The lower-left channel in the previous image, extracted and ready for analysis

6.2.2.3 - Comparing Thermal Images and XCT

Using the same methods as described in the LoF section, each melt pool recorded with the thermal camera was matched with the stack of XCT images that lay beneath it. These XCT images were then thresholded, with any pixel with an intensity less than 94% of the image's peak intensity assumed to be a pore. This value was chosen by visual inspection, proving to be the most versatile at separating pores from solid across a wide range of images. Any pore that touched any of the image's edges was discarded, as it was assumed to be either the protective channel above the line scan, or the narrow edges of the triangle, rather than an actual keyhole pore.

The binary images were stacked together, and the REGIONPROPS3 function was used to calculate the VW mean pore volumes, the VW mean pore sphericities, and the number of pores, for each of the 1,946 melt pools recorded by the thermal camera. As was done in the LoF section, these three metrics were plotted against six metrics extracted from the thermal images (melt pool length, width, area, aspect ratio, peak temperature, and mean temperature), for each melt pool.

No discernable trends or processing regions were identified in any of the plots, either when all melt pools were included in the data, or only the 180W melt pools (the most likely to exhibit keyholing) were included. As examples, the same plots that were shown to contain processing regions for the LoF analysis, are shown for the keyhole analysis in Figure 6.25.

As can be seen in Figure 6.25 (which includes all melt pool data), it would be very difficult to extract any meaningful trends. It could be argued that the number of pores beneath a melt pool roughly increases with its length, although the impact of this finding is limited.

Based on this almost-total lack of trends, three possible conclusions could be drawn:

1. The formation of keyhole pores is not associated with any kind of thermal signature that differs from those of non-porous melt pools. Based on various studies highlighted in the literature review section of this work, this possible conclusion is highly unlikely to be correct.
2. The thermal signatures associated with keyhole pore formation could not be detected with the thermal camera that was used in this experiment, due to an insufficient resolution, or frame rate, or other similar camera quality issue.
3. No keyhole pores were formed, even in the 180W line tracks at the narrowest part of the triangle. Visually inspecting some of these narrow-triangle sections, a lack of porosity is noted, making this possible conclusion likely seem correct.

This lack of keyhole porosity is surprising. Figure 6.26 shows a cube built for the parameter selection experiment in the previous chapter. This cube was built at 175W, 500 mm/s, and 50 μ m, and despite the much greater thermal conduction available to the melt pool, abundant keyhole porosity was observed.

Process Regions

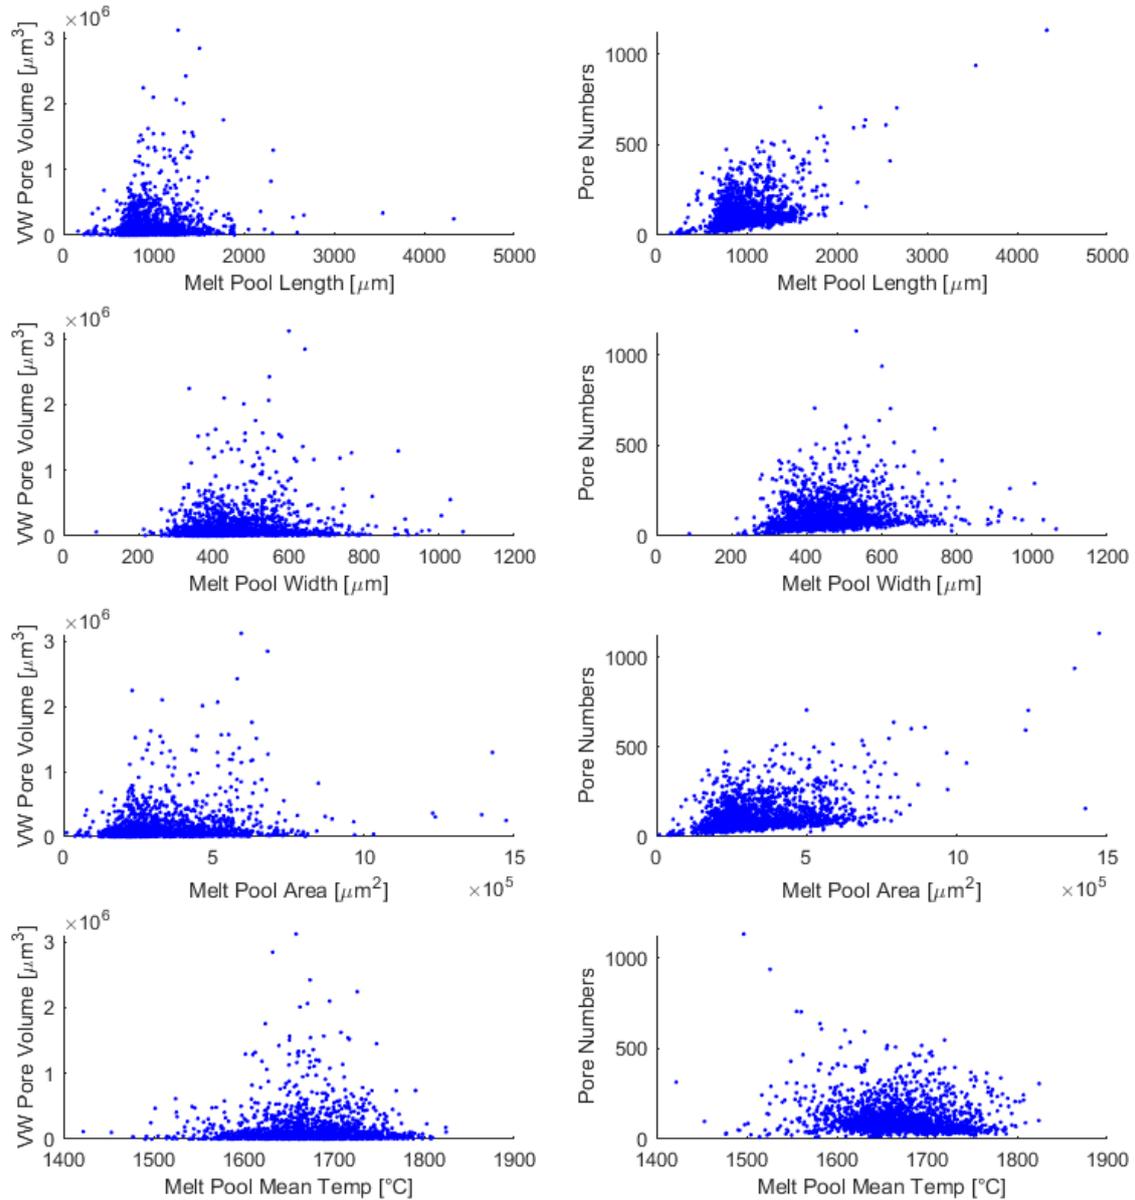


Figure 6.25

Example charts from the keyhole XCT and thermal image analysis, showing the lack of processing regions

It is possible that the fundamental assumption of this experiment (i.e. that the lower levels of thermal conduction inside the narrow triangle should result in increased keyhole porosity) was flawed. One possible explanation for this counterintuitive result is that, since thermal conductivity actually increases with temperature for IN-718, the accumulated heat within the triangle tips made keyholing less likely than in the triangle bases, or the cube. An experiment to address this possibility will be described in the Future Work chapter.

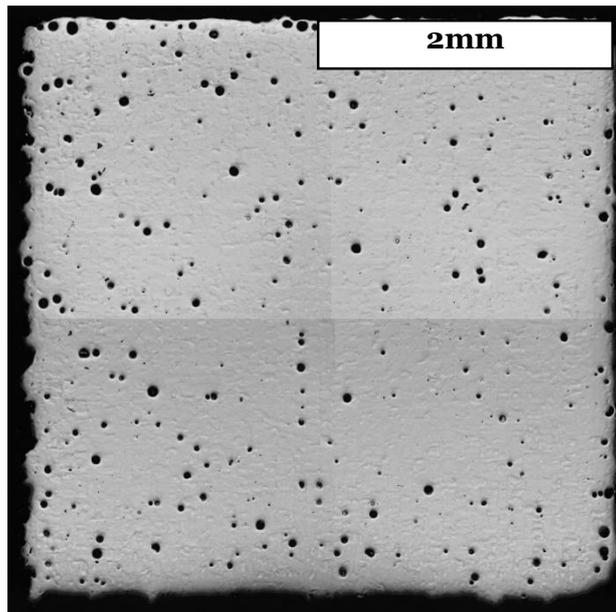


Figure 6.26

A cube built for a previous experiment, showing that keyhole pores were expected for these machine parameters

6.3 - Industrial Applications of L-PBF

All builds in this project have so far consisted of single lines, cubes, or other basic shapes. Attempts were made to link thermal signatures from these builds to characteristics and defects in the micro- and meso-structure. Any findings generated from these experiments automatically included the caveat that their applicability to more complex builds, such as those performed in industrially-relevant settings, may be questionable.

In order to test this applicability, three turbine blades were built, with identical geometries and machine parameters, but different scan strategies. The standard bulk parameters of 100W, 500 mm/s, a 50 μ m hatch spacing, and a 30 μ m layer height were used throughout. Figure 6.27 shows the layout and dimensions of the part geometry.

The five major sections of the turbine blade were designed to include a range of geometrical features. From the bottom upwards these are:

1. Fir Tree - Upskin and downskin slopes, with overhang angles as low as 20°.
2. Flange - A large and bulky section
3. Platform - An extreme overhang section, with unsupported, horizontal spans of up to 3mm
4. Airfoil - Tapering to a narrow point at its thin edge
5. Cap - A complex assortment of upskin, downskin, and thin-wall sections

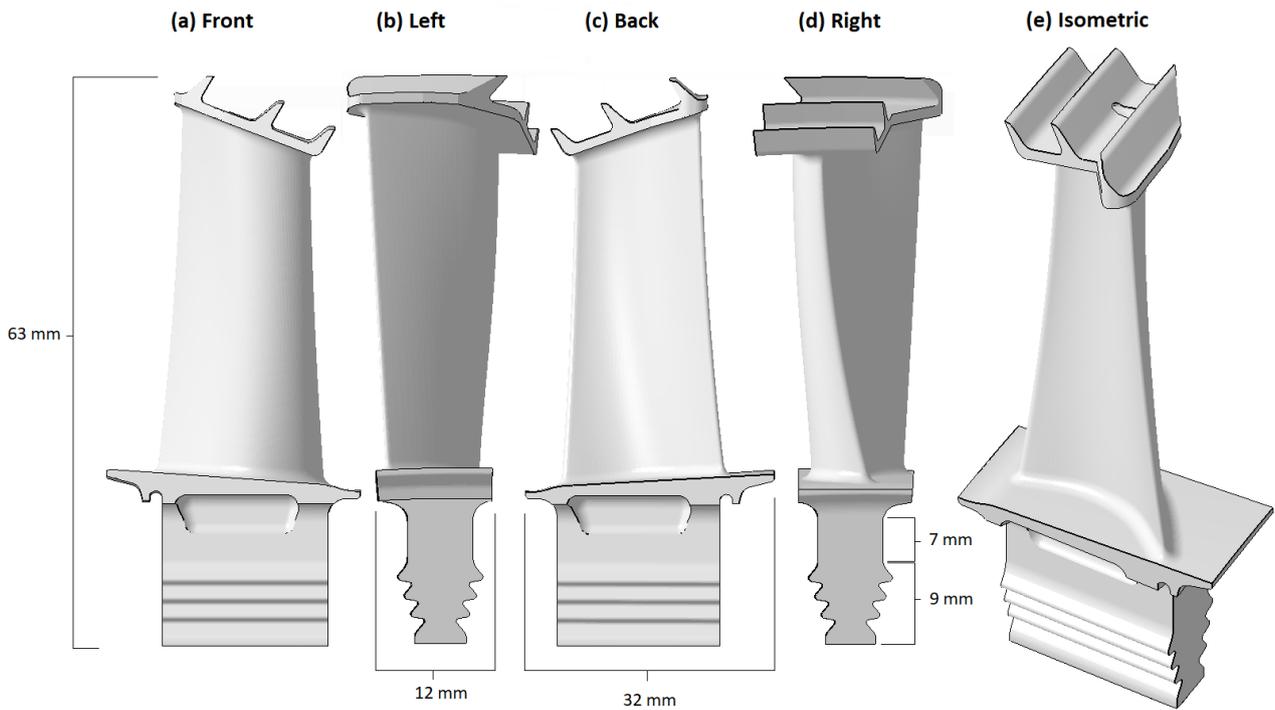


Figure 6.27
The design of the turbine blade

The three scanning patterns used in the builds were as follows:

1. Simple - Bi-directional scans, with 90° rotation between layers, aligned to the X and Y axis
2. Contour - Concentric perimeters, with $50\mu\text{m}$ spacing between them, starting at the outer edges of the part, and progressing inwards to the centre
3. Chessboard - $3\text{mm} \times 3\text{mm}$ squares, scanned in a random pattern, with 67° rotation between layers

6.3.1 - Thermal Camera Data

Each turbine blade was 2,099 layers tall. Since recording all of these layers with the thermal camera would have been neither possible nor useful, certain layers were selected for recording, with an emphasis on overhangs and other such geometrically interesting features. The layers were recorded with the 1024 pixel frame, at 200fps.

The selected layers on each turbine blade that were recorded are shown in Table 6.3. Each melt pool image was analysed using the same techniques described previously (thresholded at $1,336^\circ\text{C}$; highest temperature pixel assumed to be the melt pool; all others set to zero; REGIONPROPS function used to extract statistics).

Each melt pool image was analysed using the same techniques described previously (thresholded at $1,336^\circ\text{C}$; highest temperature pixel assumed to be the melt pool; all others set to zero; REGIONPROPS function used to extract statistics).

Table 6.3. Each melt pool image was analysed using the same techniques described previously (thresholded at 1,336°C; highest temperature pixel assumed to be the melt pool; all others set to zero; REGIONPROPS function used to extract statistics).

Each melt pool image was analysed using the same techniques described previously (thresholded at 1,336°C; highest temperature pixel assumed to be the melt pool; all others set to zero; REGIONPROPS function used to extract statistics).

*Table 6.3
The layers recorded for each turbine blade.*

Feature	Layer	Height [mm]	Feature	Layer	Height [mm]
Foot	20	0.60	Shank	300	9.00
Fir Tree 1 Upskin	40	1.20		350	10.50
	60	1.80		400	12.00
	80	2.40		500	15.00
Fir Tree 2 Downskin	90	2.70		Platform Downskin	550
	95	2.85	552		16.56
	100	3.00	554		16.62
Fir Tree 2 Upskin	110	3.30	556		16.68
	130	3.90	558		16.74
Fir Tree 3 Downskin	150	4.50	Airfoil	560	16.80
	160	4.80		750	22.50
	165	4.95		900	27.00
Fir Tree 3 Upskin	170	5.10		1050	31.50
	190	5.70		1250	37.50
Fir Tree 4 Downskin	210	6.30	Cap Upskin	1500	45.00
	220	6.60		1810	54.30
	225	6.75		1817	54.51
Fir Tree 4 Upskin	230	6.90		1830	54.90
	235	7.05		1850	55.50
Fir Tree 4 Upskin	250	7.50	Cap Downskin	1900	57.00
	270	8.10		2000	60.00

6.3.2 - XCT Data

Reconstructed data was acquired from the ESRF, in the form of approx. 9,000 TIF images for each of the three turbine blades, with each image being 3924 × 1890 pixels in size. Figure 6.28 shows, as an example, a cropped TIF image of the top of the shank of the Contour blade.

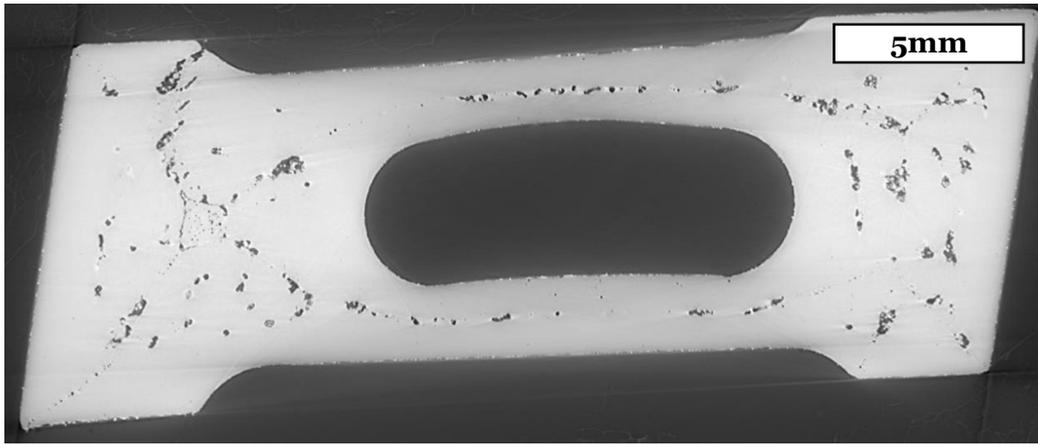


Figure 6.28

A TIF of the shank of the Contour blade, from XCT at the ESRF

It was the intention of the experiment to match each melt pool recorded with the thermal camera, to its 2D ‘footprint’ in the appropriate layer of XCT data. A major complication in achieving this was the fact that, since the turbine blades had been removed from their baseplates by EDM before the XCT took place, they were not placed into the x-ray chamber exactly upright. As a result, each TIF image did not show a perfectly horizontal slice of the turbine blade, and even appeared significantly distorted.

Figure 6.29, for example, shows an upskin section of fir tree two of the Contour blade. It is clear that the shape as portrayed in the image, does not match the true form of a horizontal blade slice. The distortion is so large, in fact, that the subsequent downskin section appears to be free-floating below the part.

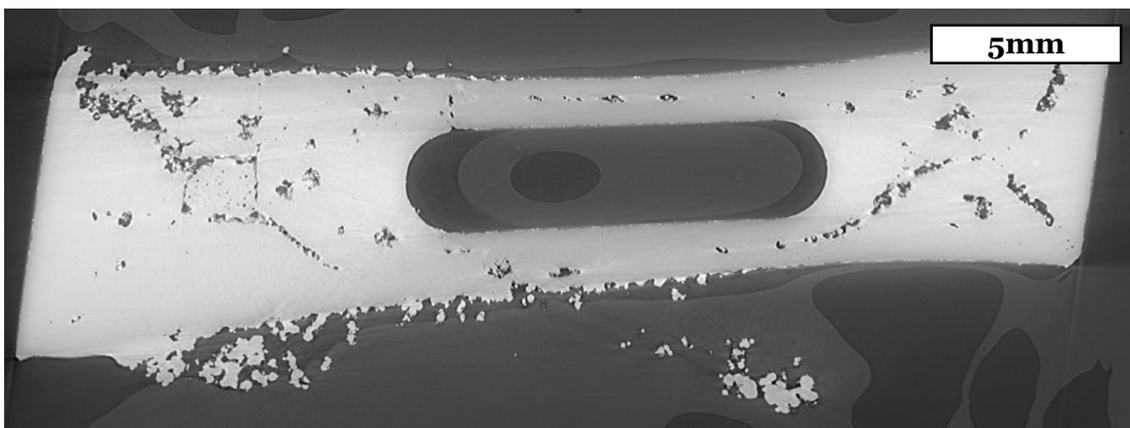


Figure 6.29

Significant distortion in the images of the fir tree in the Contour blade, due to the blade not being perfectly vertical during XCT

In order to compile pixels from the appropriate TIF images, to create horizontal slices that match the layers for which thermal camera data was recorded, it was necessary to calculate required rotations for each turbine blade. To start, at least three coordinates (where Z is the TIF image

number, and XY are the pixel coordinates within that image) were required for each blade, that were known to be horizontally-aligned in reality, even if they were not according to the TIFs. The corners of the EDM cut, at the very bottom of the blade, were chosen for this.

After manually selecting XYZ coordinates of the top left, top right, and bottom right corners of this cut for each blade, a MATLAB script was written to calculate the rotation matrix required to rotate the cut's plane (around the top left corner point), to make it fully horizontal.

Once this rotation matrix had been calculated for each blade, it was possible to apply it to every pixel, across each of the approx. 9,000 TIF images, to construct a full, rotated copy of the blades. This was achieved by:

1. Initialising 9,000 blank images
2. Progressing through each of the raw (non-rotated) TIF images
3. Progressing through every pixel in each raw image
4. Rotating it according to the defined rotation matrix
5. Saving that pixel's value, in its rotated position, in the blank image stack

Due to the extreme computation cost required for this, all processing was done on the HPC.

By measuring each turbine blade with a Vernier callipers, it was possible to estimate the amount of material that had been removed from it, when the blade was cut from the base plate by EDM. The first 75, 80, and 80 layers appeared to have been removed from the Simple, Contour, and Chessboard blade, respectively. Using the knowledge that these numbers of layers were absent in the XCT data, and the fact that each XCT voxel measured $7.7\mu\text{m} \times 7.7\mu\text{m} \times 7.7\mu\text{m}$ while each build layer was $30\mu\text{m}$ in height, it was possible to determine which image in the rotated TIF stacks corresponded to which layer in the thermal camera recordings.

Figure 6.30 shows layer 130, at the upskin of fir tree two, in the Contour blade. This is the same section that was shown in the raw TIF stack previously. The difference between the raw and rotated image is significant.

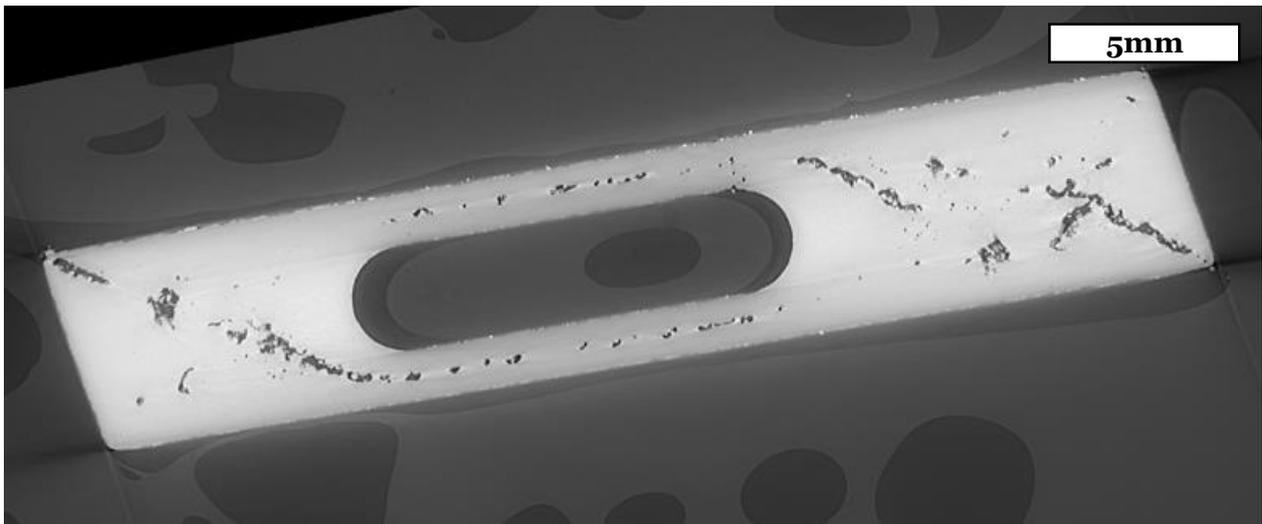


Figure 6.30

The same fir tree image as previous, now rotated in 3D to remove distortion

6.3.3 - Comparing Thermal Images and XCT

Once a rotated TIF had been produced to correspond to every build layer recorded with the thermal camera, the position of each melt pool had to be mapped to a location within the TIF. A single XCT slice was chosen to correspond to each thermal image layer recorded.

This analysis achieved by (for each individual layer-TIF pair) superimposing together all melt pools into one image, comparing this superimposed image to the TIF, manually selecting corner points in each, and applying translations and rotations until the superimposed stack matched up with the TIF. Figure 6.31 shows this process for layer 130 of the Contour blade, with the white and red dots in images (1) and (2) being the points that were manually selected.

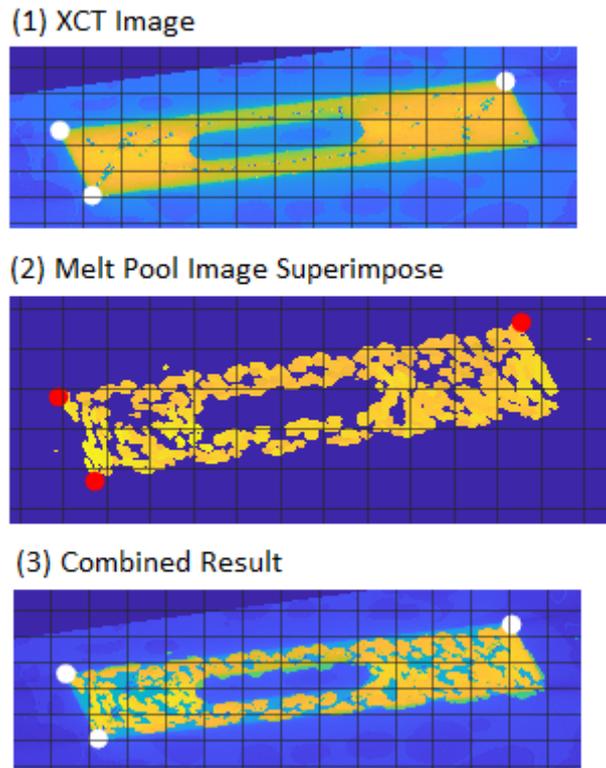


Figure 6.31

Mapping the rotated XCT images to the melt pool images

Once every melt pool, in each layer in each turbine blade, had its footprint in the XCT data defined, these pixels could be extracted and analysed for porosity. A MATLAB script was written to threshold the XCT footprint pixels at 90% of their median pixel intensity, as this value proved upon visual inspection, the most consistent at differentiating pores from solid. Once pores had been identified, those that touched the boundary of the part in the TIF image were excluded, as these were assumed to be products of surface roughness, rather than actual pores. The remaining pores were analysed using the REGIONPROPS function, to provide for each melt pool footprint: the footprint porosity [%], number of pores, the area-weighted mean pore area [μm^2], and the area-weighted mean pore sphericity.

These four metrics were plotted against seven metrics extracted from the melt pool images (length, width, area, aspect ratio, peak temperature, mean temperature, tail thermal gradient).

It had been expected that since keyhole mode melt pools are generally hotter and exhibit a greater thermal gradient than conduction mode melt pools, the following trends would have been noticeable:

1. As melt pool thermal gradient increased, the footprint porosity, number of pores, and mean pore area would increase.
2. As melt pool peak temperature increased, the same effects would be present.

Of the 28 graphs generated, any that plotted for number of pores, area-weighted mean pore area, or area-weighted mean pore sphericity, showed no discernible trends, correlations, or distinct

processing regions. Figure 6.32 shows, as an example, three such plot types in which the placement of the data points appears entirely random within certain bounds.

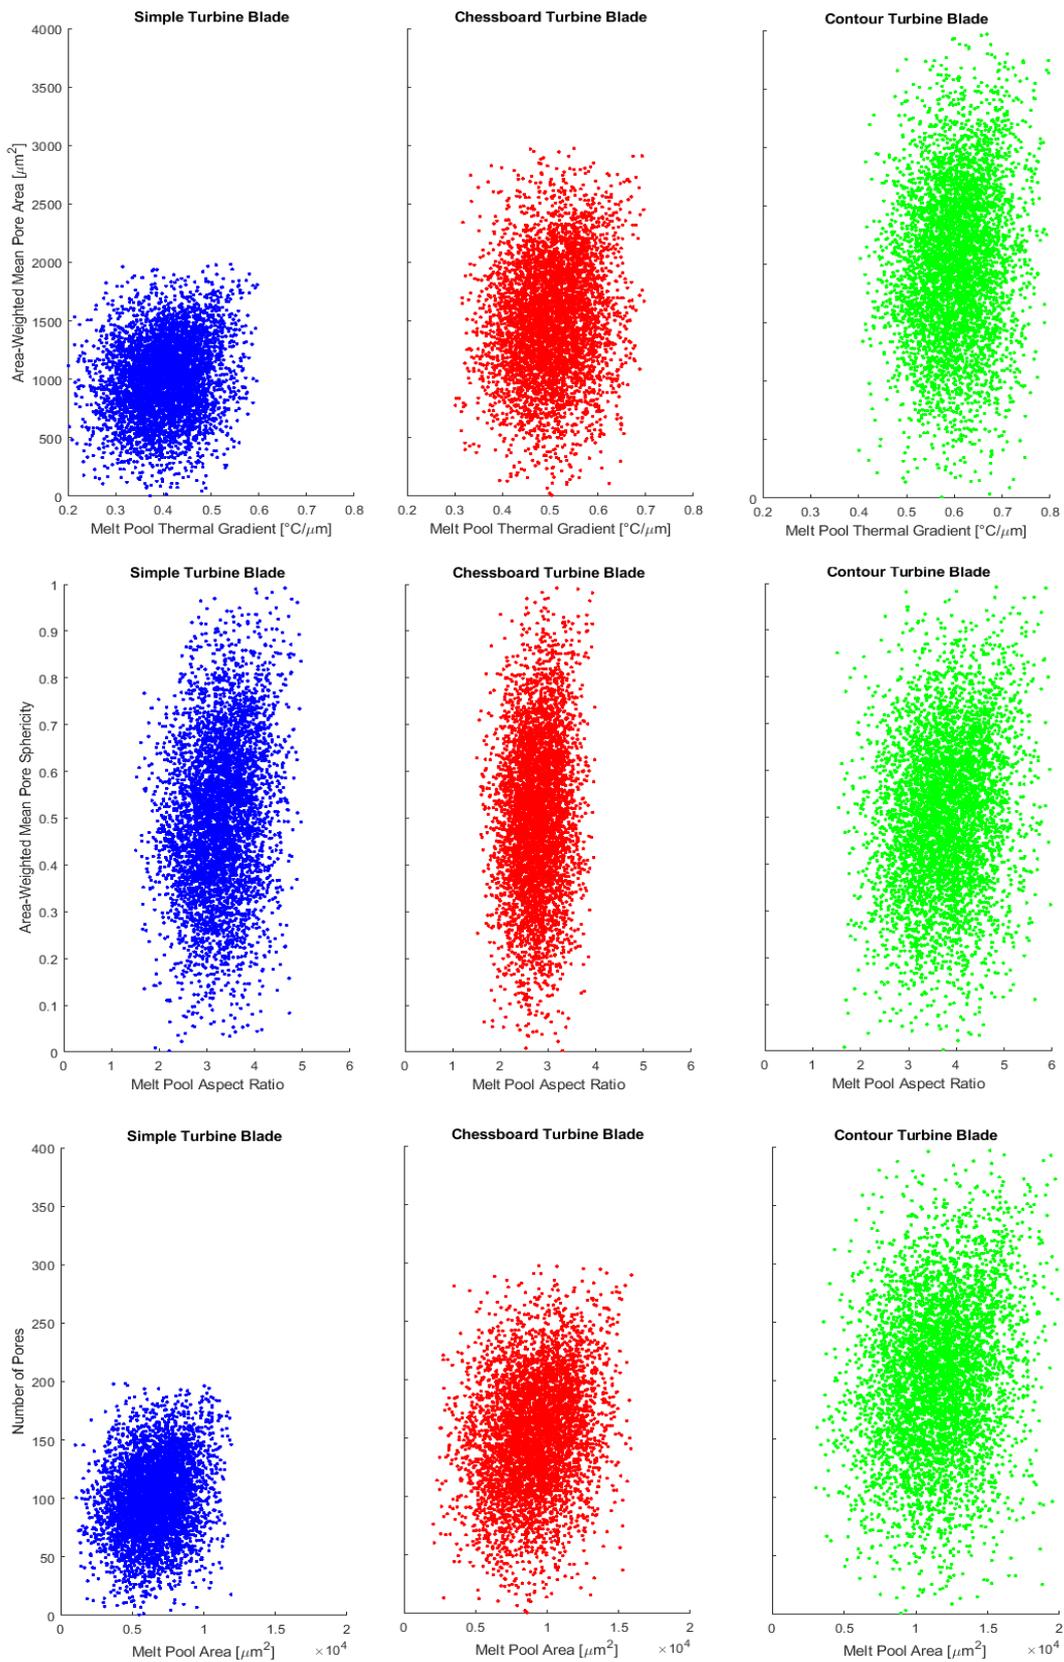


Figure 6.32

Example plots showing no trends or processing regions from XCT and thermal image analysis

When plotted for footprint porosity (i.e. the amount of porosity in the XCT image that falls beneath the footprint of each melt pool), however, distinctive processing regions begin to emerge. In each of the three turbine blades, the minimum porosity that may lie beneath a melt pool rises alongside its mean temperature and area (and similarly its length and width), as shown in Figure 6.33. No such distinctive regions were found when plotting for the melt pool's aspect ratio or tail thermal gradient.

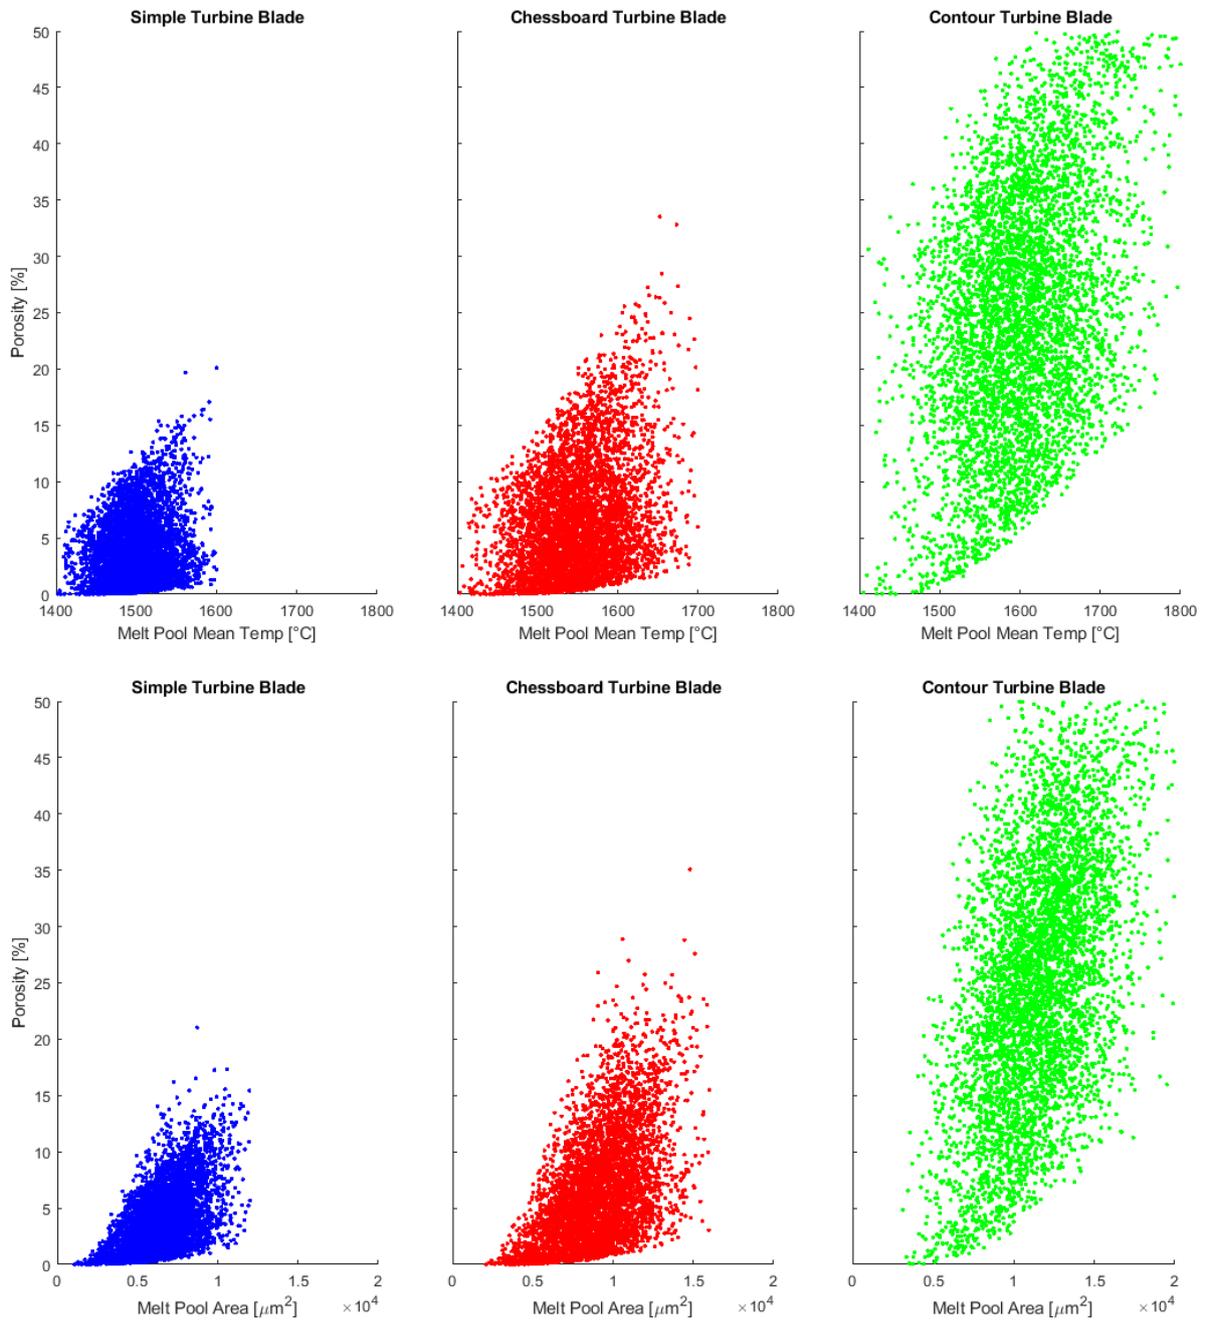


Figure 6.33

Example images showing distinct processing regions, when thermal signatures are mapped against porosity

The clearest processing regions of all, are found when plotting porosity against a melt pool's peak temperature, as shown in Figure 6.34.

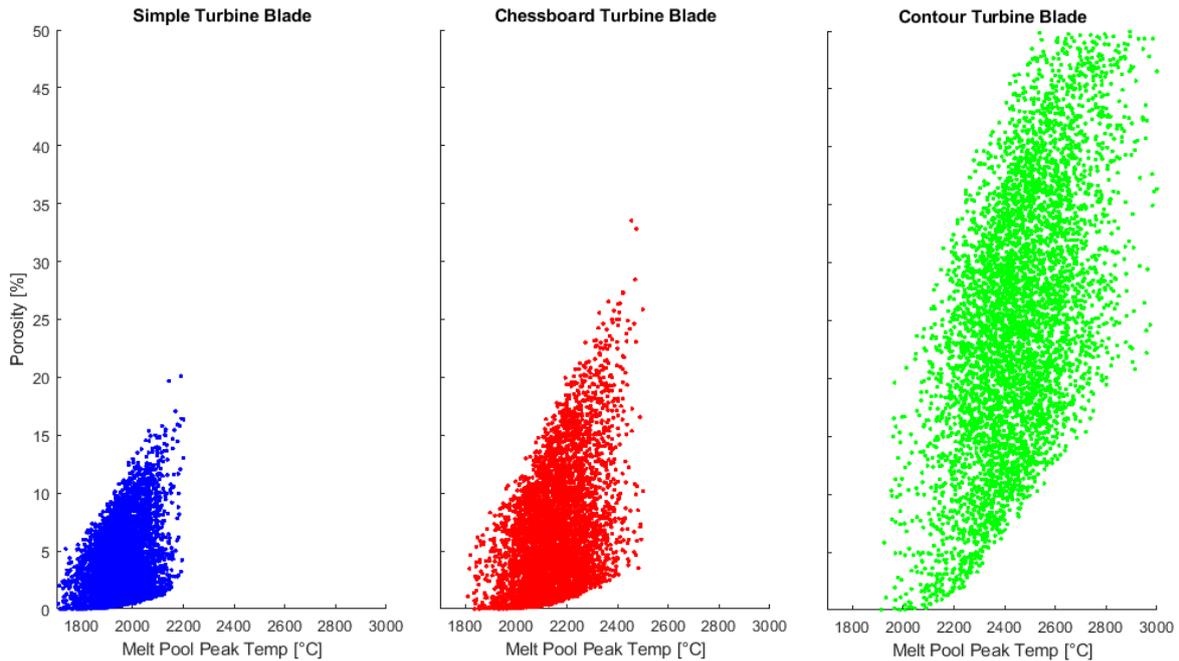


Figure 6.34

Very clear processing regions, obtained when melt pool peak temperature is mapped against porosity

In these charts, a very clear pattern is present: when a melt pool image shows an elevated maximum temperature, the possibility that it is building free of pores becomes diminished. The same is true, albeit to a lesser extent, when a melt pool image shows an elevated mean temperature, or area.

6.3.4 - Analysing XCT Independently

As well as extracting porosity beneath melt pool footprints, the pores and surface roughness of each turbine blade was analysed individually, to see if trends could be drawn between the different scanning strategies and pore frequency, size, and morphology at various sections of the blades.

6.3.4.1 - Pore Analysis

A MATLAB script was written to extract the XYZ coordinates of pore pixels from each 2D slice of the rotated stack, and save them in an array. This array was then searched for adjacent pixels, allowing all pores to be reconstructed into 3D volumes, complete with their centroid's location within the turbine blade, and their volume, sphericity, and aspect ratio. Figure 6.35 some example pores, which intersect with a downskin layer of the Contour blade fir tree.

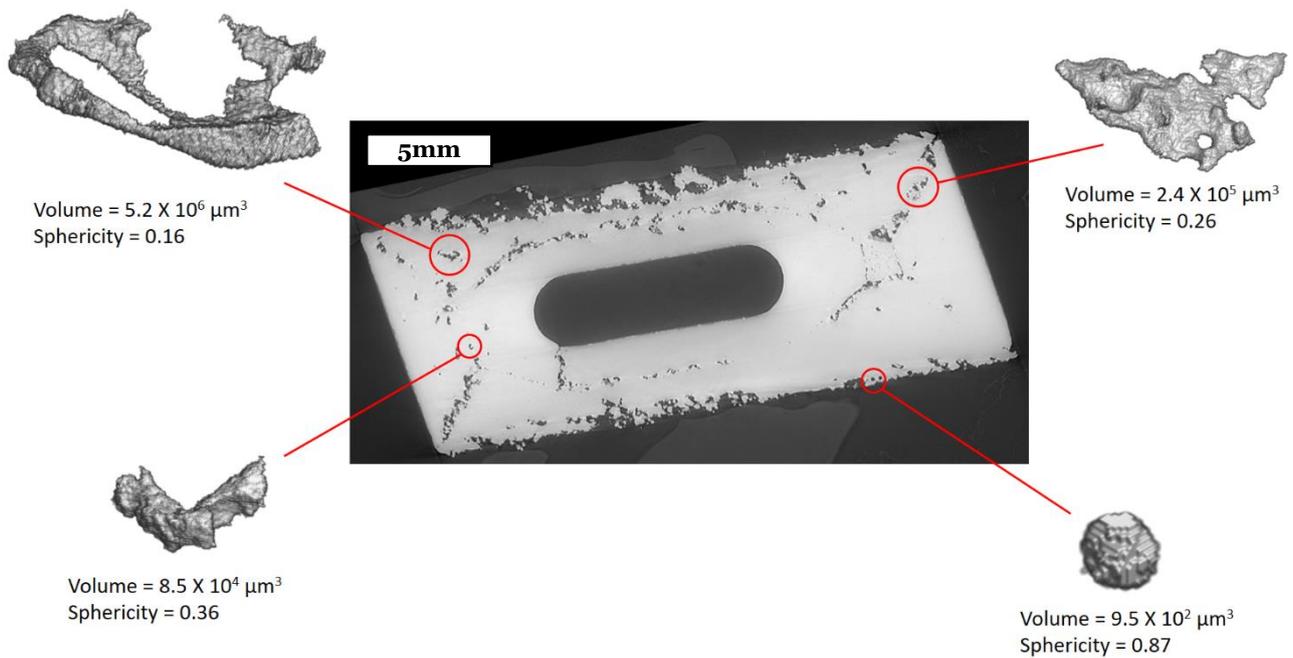


Figure 6.35
Some example pores from the Contour blade fir tree

12,058 pores were identified in the Contour blade, with a combined volume of $34.76 \times 10^{10} \mu\text{m}^3$, thereby accounting for 8.15% of the total blade volume. 6,468 pores were identified in the Chessboard blade, with a combined volume of $8.47 \times 10^{10} \mu\text{m}^3$, thereby accounting for 1.98% of the volume. 5,286 pores were identified in the Simple blade, with a combined volume of $7.21 \times 10^{10} \mu\text{m}^3$, thereby accounting for 1.73% of the volume. Figure 6.36 shows a histogram of the pore volumes in each blade, and the combined volume of each of the bins of pores.

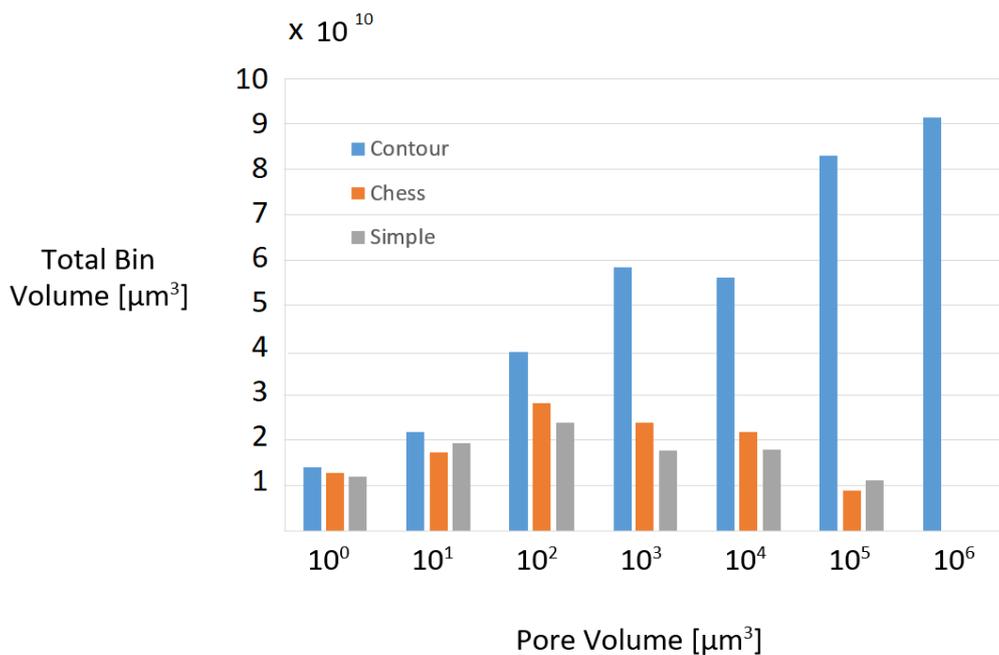


Figure 6.36
The total volume of each pore size, in the three blades

It was found that the pores in the Chessboard and Simple blades were generally smaller and more spherical, and spaced evenly throughout the component, while the pores in the Contour blade were larger, more irregular, and concentrated at points where the spiralling contour vectors met.

Table 6.4 shows a few examples of pore morphologies and volumes. The pores have been selected from the list compiled from each blade, with the pores ranked in order of volume. As can be seen, the pores from the Chessboard and Simple rapidly become spherical as their ranking decreases, whereas the Contour pores remain large and complex.

6.3.4.2 - Surface Roughness Analysis

A MATLAB script was written to extract a surface outline for every rotated TIF image, for each of the three turbine blades. Figure 6.37 shows some example outline, extracted from the Simple airfoil and Contour fir tree upskin.

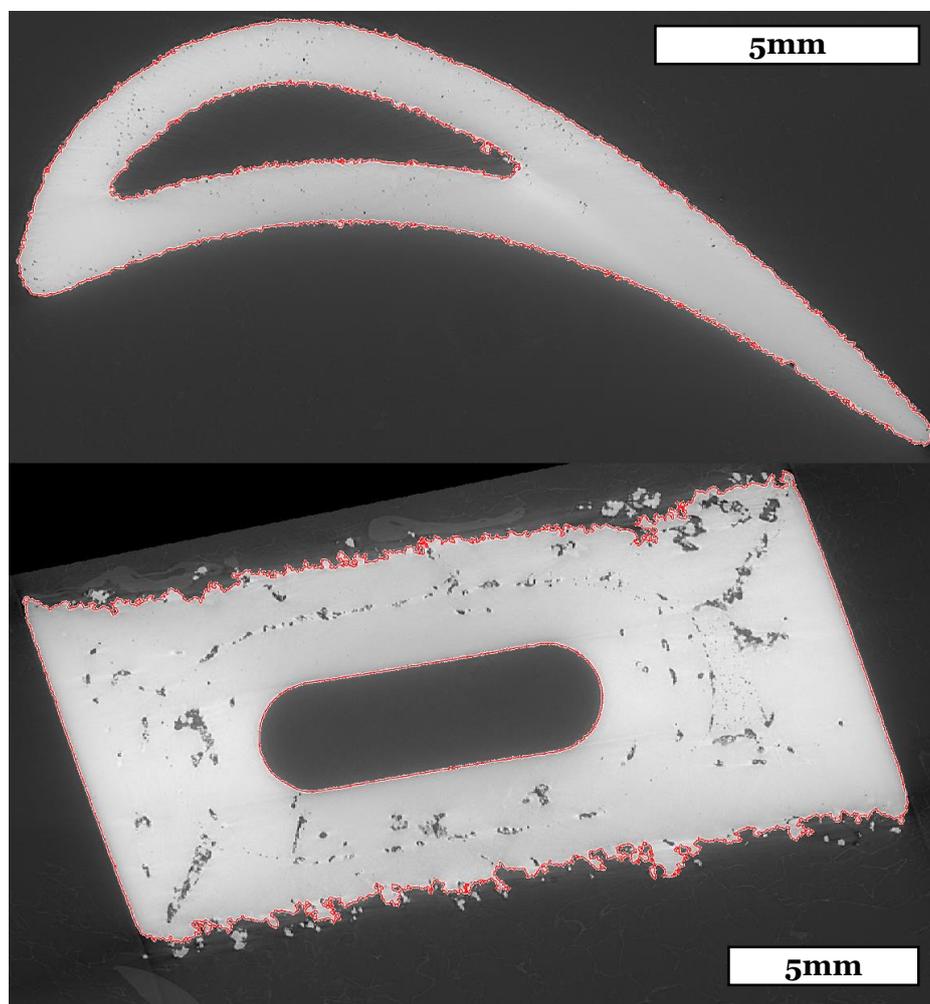
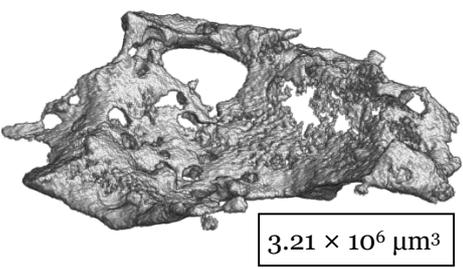
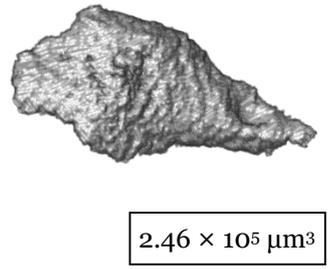
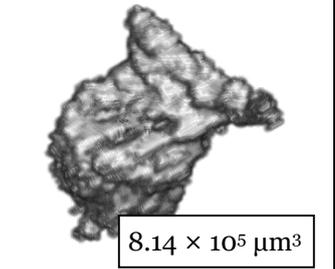
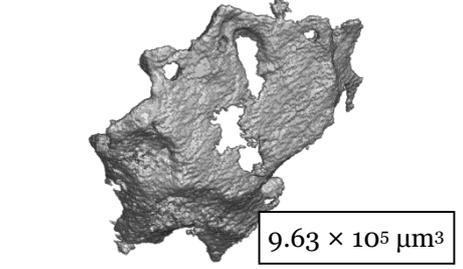
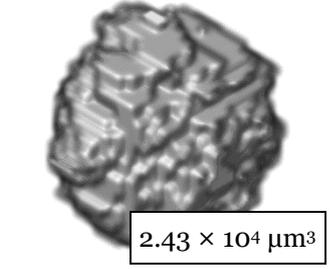
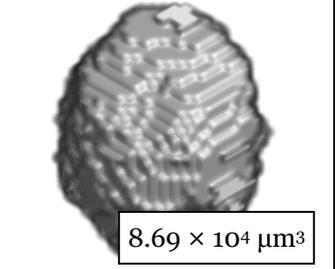
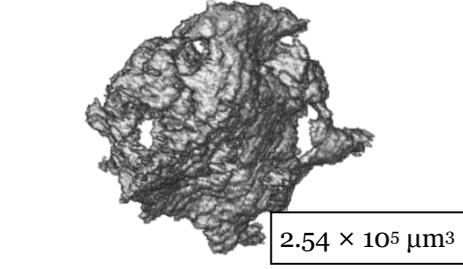
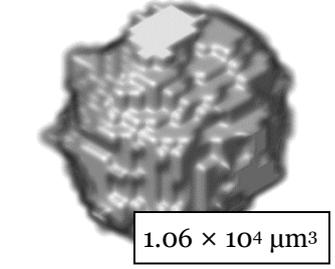
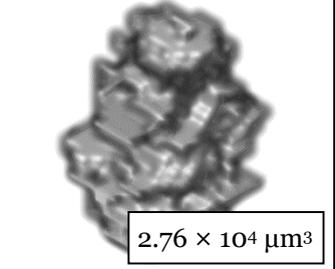
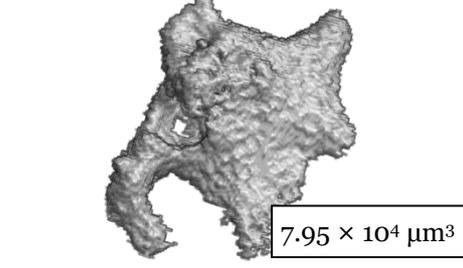
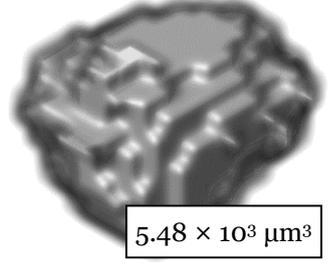
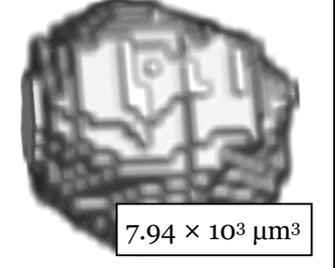
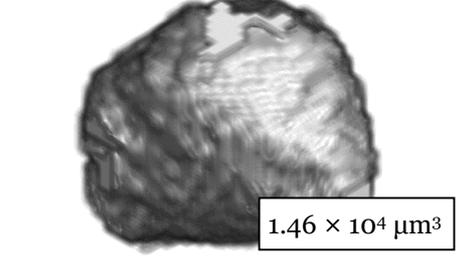
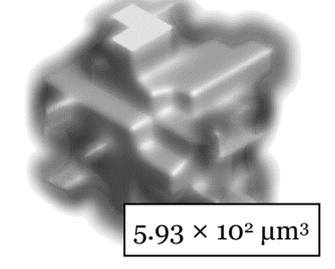
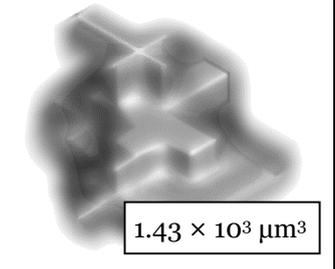
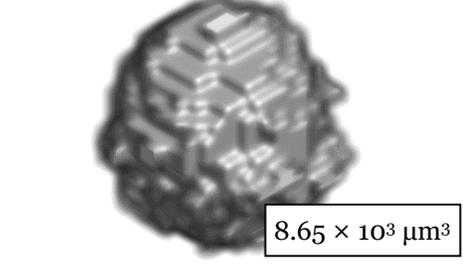
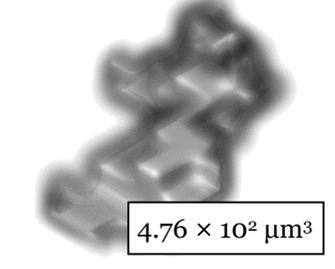
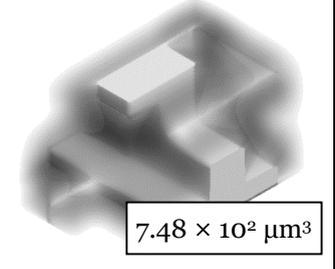


Figure 6.37

The extracted surface outlines (in red), overlaid onto rotated TIFs for the layers in the Simple airfoil (top) and Contour fir tree upskin (bottom)

Table 6.4

A selection of pores from each blade, drawn from a list ranked by volume

Pore	Contour	Chessboard	Simple
1 st Biggest	 $3.21 \times 10^6 \mu\text{m}^3$	 $2.46 \times 10^5 \mu\text{m}^3$	 $8.14 \times 10^5 \mu\text{m}^3$
10 th Biggest	 $9.63 \times 10^5 \mu\text{m}^3$	 $2.43 \times 10^4 \mu\text{m}^3$	 $8.69 \times 10^4 \mu\text{m}^3$
100 th Biggest	 $2.54 \times 10^5 \mu\text{m}^3$	 $1.06 \times 10^4 \mu\text{m}^3$	 $2.76 \times 10^4 \mu\text{m}^3$
500 th Biggest	 $7.95 \times 10^4 \mu\text{m}^3$	 $5.48 \times 10^3 \mu\text{m}^3$	 $7.94 \times 10^3 \mu\text{m}^3$
1,000 th Biggest	 $1.46 \times 10^4 \mu\text{m}^3$	 $5.93 \times 10^2 \mu\text{m}^3$	 $1.43 \times 10^3 \mu\text{m}^3$
5,000 th Biggest	 $8.65 \times 10^3 \mu\text{m}^3$	 $4.76 \times 10^2 \mu\text{m}^3$	 $7.48 \times 10^2 \mu\text{m}^3$

In order to obtain a measure of surface roughness from these outlines, they were first broken into two categories:

- Those that ‘should’ (according to the Netfabb design) be straight lines. All external surfaces in the fir tree, the long sections of the fir tree cooling channel, most of the platform overhang, and some section in the cap fell into this category.
- Those that shouldn’t be straight lines. The short section of the fir tree cooling channel, the air foil, some sections of the platform, and most sections in the cap fell into this category.

Attempts were made to design an automated system that received all of the approx. 9000 outline images, for each of the three blades, and divided each image up into outline sections that should and should not be straight lines. It was soon discovered, however, that the alternative approach would be much more straightforward: significantly downsize the number of images being processed, and divide them up manually.

The same 42 layers that were chosen for thermal camera recording (see Table 6.3) were chosen for each blade. Each melt pool image was analysed using the same techniques described previously (thresholded at 1,336°C; highest temperature pixel assumed to be the melt pool; all others set to zero; REGIONPROPS function used to extract statistics).

A MATLAB script was written to allow user selection to divide up the sections. Figure 6.38 shows this division for a fir tree upskin layer from the Chessboard blade. A non-rotated XCT slice has been chosen for this example (i.e. the top edge shows an upskin section, while the bottom edge shows a vertical section), in order to highlight the difference between the two in terms of roughness.

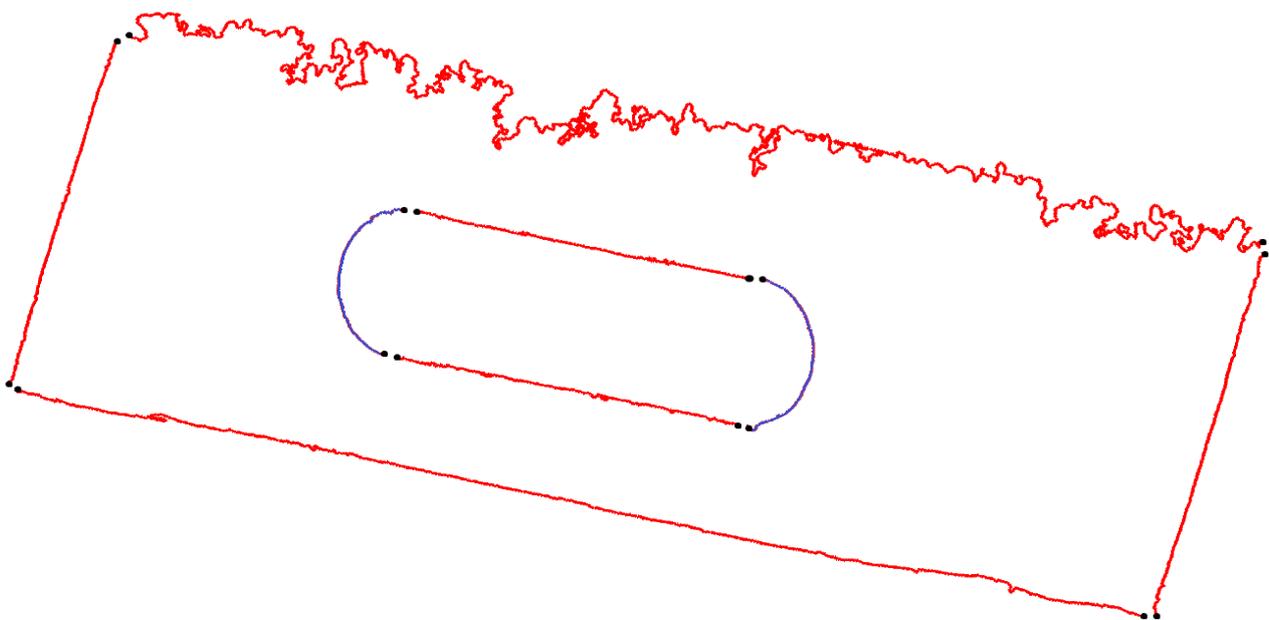


Figure 6.38

A Chessboard fir tree upskin layer divided up into sections that should be straight lines (red), and sections that shouldn't (blue). Each section is ended with a black dot for visual clarity.

In order to calculate the surface roughness of the straight line sections, each section had a line of best fit calculated, and the mean distance of each outline pixel from this line was found. Figure 6.39 demonstrates this, using the same line section as in the fir tree in Figure 6.38. A possible limitation of this method is demonstrated, whereby the long, narrow, downward deviations of the red line cause the entire blue line to shift downwards.

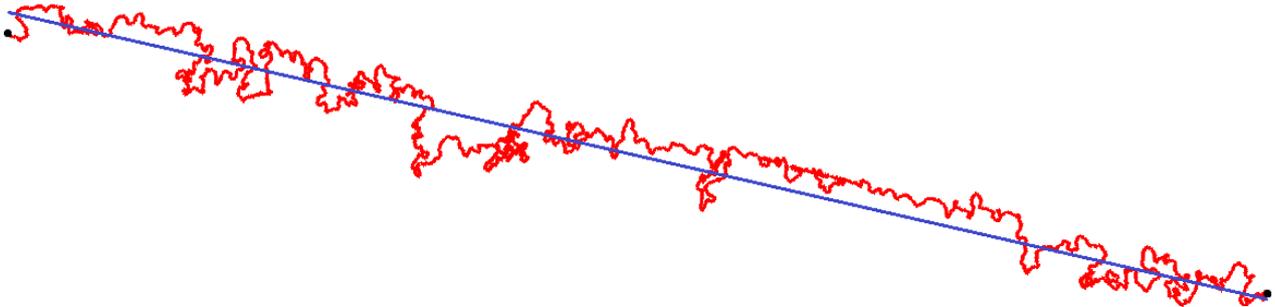


Figure 6.39

The line of best fit (blue) is applied to the fir tree outline section (red), and the absolute perpendicular distance of each outline pixel from the line is calculated.

To calculate the surface roughness of the curved sections, the boundary outlines were smoothed using a moving mean, with a window size of 80 pixels, and the mean absolute perpendicular distance of each boundary pixel from the smoothed curve was calculated. Figure 6.40 demonstrates this, with the same airfoil tail section as shown in Figure 6.37.

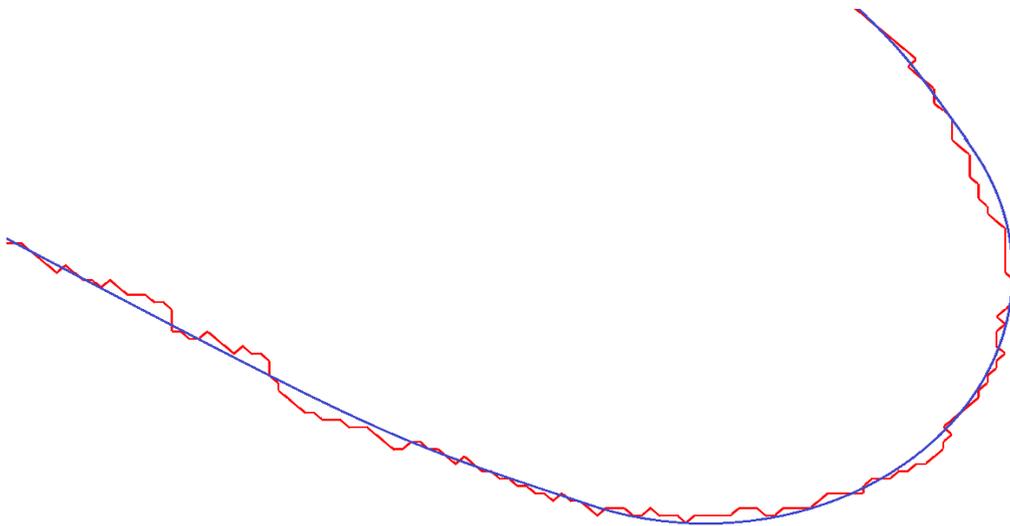


Figure 6.40

The moving mean smoothed outline (blue), and the original outline (red) of a Simple airfoil tail section

Using these two methods, it was possible to produce an R_a for each section, of each of the 42 selected layers, for each of the three turbine blades. These values are summarised in Figure 6.41, which show average R_a for each section, with standard deviations as errors.

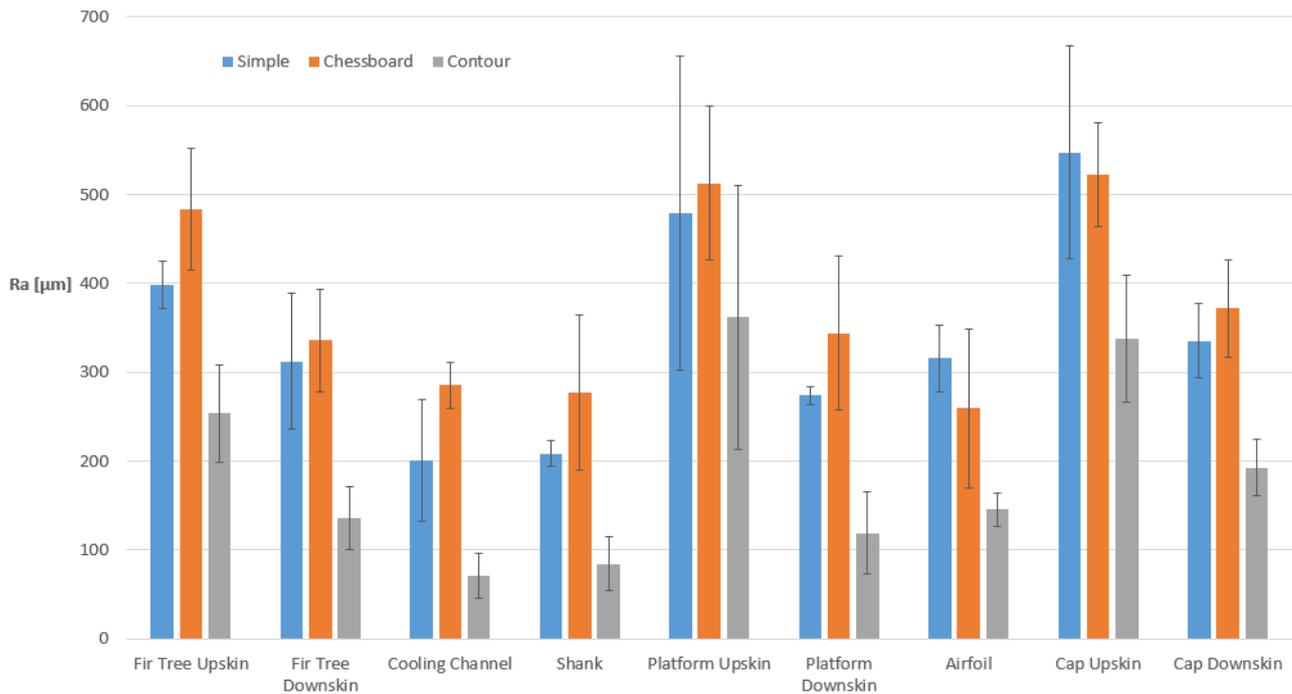


Figure 6.41

The average R_a , for various blade geometries. Shank here includes the vertical sections of the fir tree.

Two notable conclusions can be drawn from these values, each of which align with established literature. Firstly, the fact that R_a is consistently highest in upskin sections (fir tree, platform, and cap upskins), medium in downskin section (fir tree, platform, and cap downskins), and lowest in vertical sections (cooling, shank, and airfoil), matches what other researchers have found about the effects of dross [332, 333] and staircases [152, 334].

Secondly, it is clear that the Chessboard blade has the roughest surfaces overall, with the highest R_a values in almost every category. The Simple blade is slightly smoother, though often the differences between the Chessboard and Simple are within error. The Contour blade has the lowest roughness by far, due to the fact that its perimeters consisted of a single, unbroken vector, as opposed to the other two blades, whose perimeters consisted of many vector ends.

Based on these results, and the pore analysis results presented earlier, a clear recommendation can be drawn, that is also supported by literature [335, 336]. A chessboard scan strategy is likely to reduce internal porosity within the interior of a component. The external perimeter, however, should then be finished off with a contour outline, to reduce surface roughness. A combination of these two strategies is most likely to reduce both forms of these mechanical defects.

6.4 - In-Situ Mesostructure Conclusions

The experiments in this chapter primarily investigated the ability of the thermal camera to locate specific pores, whether placed deliberately or not, within samples. Optimal parameters for full density in IN-718 components were first determined with high-throughput cube testing. The parameter set chosen was found to be consistent with previous findings in literature.

Samples were produced with geometries and powers that were intended to cause LoF and keyholing porosity within parts. These pores were detected with XCT, and their locations were cross-correlated with melt pool images. Trends were searched for between these pores, and the thermal signatures of the melt pools that formed them.

It was found that linking specific thermal signatures with the creation of pores was challenging, due to the stochastic nature of L-PBF. Meaningful processing windows were discovered for the LoF pore samples. These windows showed that the upper limits of both the number of pores, and the volume of pores, beneath a melt pool footprint, correlated predictably with melt pool and temperature. This suggests that the thermal camera can predict the likelihood of the process straying into an unacceptable porous region. No such windows were discovered for keyhole porosity, possibly due to a flawed assumption that accumulated heat increases keyholing likelihood in IN-718, whereas the material's thermal conductivity increase with temperature may actually decrease this likelihood.

When applied to a complex turbine blade geometry, it was found that one significant process window that results is that of peak melt pool temperature. Namely, when a melt pool exhibits an elevated peak temperature, the possibility that the pool is building free of pores becomes diminished. This finding suggests an applicable usage of thermal camera monitoring in industrially relevant processes.

7 - Conclusions

The experimental work carried out in this project provided a comprehensive exploration into the characterisation of the in-situ monitoring equipment available at the University of Sheffield, and the equipment's value in revealing the microstructural and mesostructural properties of IN-718 and ABD-900AM in L-PBF.

It was discovered that a range of corrective measures had to be taken, to account for distortions in thermal camera images arising from the camera's 15° angle offset from vertical, and the defocusing that resulted in turn. Melt pool widths, peak temperatures, and mean temperatures were found to change by -6%, -5%, and +5% respectively, for each mm that their centroids drifted from the frame's line of focus. This result showcased the limitations of off-axis imaging, as only select areas of the baseplate can ever be entirely in focus, and melt pool images from outside of these areas will require corrective measures.

It was found that, while melt pool lengths did not exhibit noteworthy size changes as they strayed from the focus, their values nonetheless required corrections to account for motion blur. Experiments showed that $1.3Vt$ (where V is laser velocity and t is camera exposure time) had to be subtracted from the length measurement that appeared in thermal camera images.

Even with these corrections, it was still found that thermal camera images were under-predicting melt pool dimensions by a factor of between 3 and 5, when compared to optical microscopy and ThermoCalc models. Based on this, it was concluded that thermal camera results should be considered accurate for the discovery of trends, but should not be considered accurate in contexts when absolute dimensions are required, such as when comparing different machine setups, or when attempting cross-experimental dimensionless analysis.

The circular dependency of temperature and emissivity remained problematic, despite attempts to verify methods of removing emissivity from calculations, by instead searching for the liquidus arrest in cooling curves. The liquidus arrest should, in theory, exist as a brief slowing of the falling temperature, as the curve progresses from melt pool interior to exterior.

Due to the random fluctuation of temperature along the melt pool cooling curve, creating an automated system to detect the liquidus arrest proved difficult for single line tracks, and impossible for more complex shapes. A fixed emissivity was therefore used for all temperature conversion calculations.

Despite all these limitations of the thermal camera, many coherent trends were identified across thermal signatures arising from varying build parameters. Melt pool length, width, area, and peak temperature consistently responded to changes in scan length and layer height, in the manner that was expected from the theory of thermal conduction. Increasing layer height, in particular, was found to increase melt pool dimensions as thermal conduction was lowered.

Such coherency was not universal. Some build parameter changes did not result in any trends, while increasing the quantity of LoF porosity beneath a melt pool's path was actually found to decrease its dimensions and temperature, contrary to expectations. A theory was proposed as to how this seemingly counterintuitive observation was possible: the liquid melt pool may sink into the underlying porosity, thereby reducing the area visible to the thermal camera.

In terms of both microstructural and mechanical properties, no consistent trends were identified between thermal signatures, and the size or morphology of microstructural features, or hardness. This lack of a trend strongly suggests that an expectation that machine parameters and thermal signatures can act as an accurate proxy for the thermal gradient and solidification velocity at the solid-liquid interface, is oversimplified. The only significant correlation between metrics from different data sources was that between sample porosity (identified from optical microscopy) and grain misorientation (identified from EBSD). The correlation implied that LoF pores disrupt epitaxial columnar grain growth, promote new nucleation sites, and thereby result in greater misorientation within the alloy.

Additionally, it was found that the ability of the thermal camera to identify the formation of specific LoF or keyhole pores was limited. Several reasons were suggested for why the camera system may not be capable of detecting the subtle thermal signatures associated with pore formation, some of which were related to image quality, while others were related to the signatures themselves.

While individual LoF pore creation detection remained elusive, clear processing boundaries were established that suggested where LoF porosity was likely. It was found that as melt pool dimensions increased, and as mean temperature decreased, the upper limit of the resulting pore volume and frequency below the pool, increased. When these analysis methods were applied to a complex, industrially-relevant build, it was also found that a strong correlation exists between an elevated melt pool peak temperature, and an increased porosity beneath the melt pool.

Taken all together, these findings highlight the strengths and weaknesses of off-axis thermal monitoring. Careful calibration is required before the system can be considered ready for use. The extent to which porosity formation can be determined from any individual melt pool image is very limited. Process maps can be drawn to suggest the likelihood of components building in porous or non-porous regions. Expecting to predict microstructural or mechanical properties from thermal signatures is to over-simplify the complexities of solidification.

8 - Future Work

The results of much of the experimental work carried out in this project lead to further research questions that require investigating. These research questions will be broken down by the research category they belong to.

8.1 - In-Situ Monitoring Equipment

The following research questions arose from the first experimental chapter, which investigated the calibration of the thermal camera.

8.1.1 - Motion Blur

The experiment intended to determine the effects of motion blur on the length of melt pools in thermal images yielded some unexpected results. It stands to reason that the extent to which a melt pool is extended should be simply a product of the laser speed (V), and the camera's exposure time (t). For example, if the laser is travelling at 500 mm/s, and the exposure time is 1ms, then the extension due to motion blur should be:

$$500 \text{ mm/s} \times 1 \text{ ms} = 500 \mu\text{m}$$

However, the experimental results consistently showed that the extension is not Vt , but $1.3Vt$. It is currently unknown how or why the 1.3 constant arises. Possibly the number is an artefact of some imperfection in the camera, the recording software, the L-PBF machine, or some piece of equipment. It could be the case that the laser speed with which Sheffield University's Aconity LAB builds is 130% of the speed it's programmed to. Experimental investigation is required to uncover the meaning behind the number.

In order to test if the Aconity's laser velocity is significantly faster than it is programmed to be, a large rectangle could be printed, with long vectors to minimize the effect of skywriting time. A phone camera could be used to record the printing through the viewport, and the video analysed to calculate the time required to print the shape.

From this, should the Aconity laser velocity be shown to be accurate, a similar rectangle could be printed and recorded with the thermal camera, to see if it is recording at a frame rate that matches expectations.

8.1.2 - Emissivity

The experiments conducted to see if melt pool dimensions could be calculated based on liquidus arrest, rather than on emissivity conversion, proved unreliable. When single line tracks were scanned across a powder-free base plate, the liquidus arrest could only be automatically detected in the cooling curve in 38% of melt pool images, required human input to be found in 30% of images,

and could not be found at all in 32% of images. When single line tracks were scanned across powder, these numbers were 21% (automatic detection), 23% (human required), and 56% (detection impossible). When more complex shapes were scanned in powder, detection was discovered to be entirely unfeasible.

The reason that detection proved so difficult was due to the large temperature fluctuations in any cooling curve that was extracted from a melt pool. The liquidus arrest was hidden amongst other rises and falls, and differentiation was often impossible.

Further work could be undertaken to see if these difficulties could be overcome. It is possible that superior methods of either extracting a non-fluctuating cooling curve, or distinguishing the liquidus arrest from the fluctuations, could result in the melt pool dimensions being determinable from any melt pool image.

Attempting the experiments again, with different powers, speeds, and materials, or possibly using a different camera setup, may prove to alleviate these difficulties.

8.1.3 - Thermal Camera Dimension Under-Prediction

When melt pool dimensions, as determined from thermal camera images, were compared to dimensions that were determined from either optical microscopy images or ThermoCalc models, it was found that the latter two of these methods yielded very similar results, while the thermal camera consistently under-predicted lengths and widths by a factor between 3 and 5.

Due to the fact that the trends within the thermal camera results were self-consistent, this inconsistency with other measurement methods was interpreted as further evidence that the camera should not be relied upon for exact dimensions. Therefore, for the remainder of the work, no comparison was made between thermal camera dimensions, and any other type.

Further work could be undertaken to better understand why this under-prediction occurs. It could be a case of melt pools being under-sized during image analysis, or it could be due to melt pool peripheries not emitting sufficient IR light, or it could be some other factor. A better understanding of the reason for this inconsistency would allow for more confidence in results derived from the thermal camera.

8.2 - Microstructural Development

The following research questions arose from the second experimental chapter, which investigated the ability of the thermal camera to predict microstructural development.

8.2.1 - Effect of Hatch Spacing on Heat Accumulation

The effect of increasing hatch spacing on the thermal properties of a sample remained unclear, after the experimentation in this chapter. There was some evidence to suggest that as hatch spacing increased, there was a decrease in preheating of powder in the laser's path, so therefore lower heat in the melt pool. There was also some evidence to suggest that as hatch spacing increased, and therefore more of the material in the laser's path transitioned from solid to powder, the thermal conduction of the melt pool decreased, so therefore there was greater heat in the melt pool.

It is possible that one of these effects is entirely dominant over the other, or it could be possible that the effect that is more dominant depends on the hatch spacing and other factors. Further work is needed to better understand which theory should be considered more viable, and under which circumstances.

A test could be designed to test a wide range of hatch spacings, while keeping other parameters constant, and repeated for several different powers and speeds. When components built under these conditions are analysed with the thermal camera, more information about the effects of hatch spacing on heat accumulation will be revealed.

8.2.2 - Value of Thermal Gradient Measurements

Due to the strong linear correlation between thermal gradient, as extracted from the tail of melt pool images, and the length of these melt pools, it was suspected that this method used to estimate thermal gradient was not actually a true representation of the gradient at the solidification front, but rather just a function of the pool's length and peak temperature.

Further work is required to determine a better way of extracting a thermal gradient for the solidification front. This work would be closely linked to that described in Section 8.1.3, since an accurate determination of the boundary of the melt pool is needed for both.

Further analysis could be undertaken, with a more exact method of extracting thermal gradient measures, rather than those based on the melt pool length. For example, the few hundred degrees of cooling immediately after the peak temperature pixel, could be analysed instead.

8.2.3 - Relationship Between Hardness and Microstructural Features

According to the Hall-Petch relation, as grain size decreases, strength (here inferred by hardness) increases. Such a straightforward correlation was not found in this experimental work, however, with the microstructural feature sizes according to neither EBSD results, nor SEM results, neatly correlating with sample hardness.

There could be several reasons for this lack of a trend. Measurement issues, with any of the three methods, could have played a role. The Hall-Petch relation may be oversimplified for this kind of an analysis. Hardness and strength may not be perfect proxies for each other. Further work would be required to fully understand whether the initial assumptions that SEM/EBSD feature size should be correlated with hardness is correct. Such work may reveal whether the contributions of feature size to hardness are small in comparison to other effects associated with rapid solidification in L-PBF, such as those of microsegregation, high residual stresses, dislocation densities, and anisotropy.

Another major reason for this lack of a relationship, could revolve around the measurement technique used to test for hardness. Since indents were deliberately selected to stay 1mm away from pores horizontally, but may have been close to pores vertically, the method cannot be said to be unbiased. Using a more rigorous and thorough selection method may result in a trend being present.

8.2.4 - Predictions Based on the GR Diagram

Many of the predictions about the effects of machine parameters on grain size and morphology were based on the Thermal Gradient - Solidification Velocity diagram. The extent to which these predictions were proved correct was highly questionable. There could be several explanations as to why the linkage between machine parameters and grain size/shape output broke down. The effects of machine parameters on thermal gradients and solidification velocities could have been misunderstood. There could have been errors in microstructural measurements.

Additionally, it could be the case that, even were it possible to accurately control thermal gradients and solidification velocities via machine parameters, and also possible to measure grain size and shape with complete accuracy, it may still be an oversimplification to expect the latter to be easily predicted by the former, when it comes to the stochastic and chaotic nature of L-PBF.

Further research on this question would be valuable. Specifically, an experiment could be designed to carefully alter both the thermal gradient and solidification velocity of a melt pool, and to determine the extent to which its microstructure matches what the GR diagram predicts.

8.3 - Mechanical Defect Development

The following research questions arose from the third experimental chapter, which investigated the ability of the thermal camera to predict mesostructural development.

8.3.1 - Melt Pools Sinking into LoF Porosity

Melt pools were often found to decrease in size, as the amount of LoF porosity in the substrate beneath them increased. This finding was entirely contrary to the logical reasoning that, as thermal conduction decreased due to porosity, melt pool size should increase. One possible explanation for this finding is that as porosity increases, the molten alloy in the pool can sink into the voids beneath it, therefore decreasing the area of its surface visible to the thermal camera. Based on the fact that dross at overhangs is known to protrude downwards from the downskin, this theory seems very possible. An experiment focused entirely on this possible phenomenon would be valuable.

Designing an experiment that tests this phenomenon would involve running single line tracks over various types of porosity, ranging from fresh powder to lightly sintered powder, to fully dense solid. By recording such tracks with the thermal camera, it will be possible to determine if the molten metal is sinking into the pores.

8.3.2 - Lack of Keyholes in Narrow Sections

When a 2mm × 2mm cube was built with 175W power, 500 mm/s speed, and 50µm hatch spacing, abundant keyholing was observed. When a narrow triangle section, with far less volume for heat to conduct away from the melt pool into, was built with 180W power, 500 mm/s speed, and 50µm hatch spacing, no obvious keyholing was observed. This result appears completely counter to logic.

One possible explanation that was proposed involves the thermal conductivity of IN-718. Because conductivity increases with temperature, it may be possible that in the cube, where less heat accumulation is possible, and therefore the temperature of the solid alloy surrounding the melt pool is lower, heat was actually conducted away from the melt pool at a lower rate than in the triangle tip, and therefore keyholing was more likely.

An experiment designed to test this hypothesis would be very important in understanding the effects of complex geometries on porosity. It could be the case that in bulk sections, less prone to heat accumulation, risks of keyholing are more pronounced than in narrower sections, and therefore engineers and technicians should adjust machine parameters accordingly.

8.3.3 - Camera Quality and Pore Formation

It was determined that the ability of the thermal camera to detect the formation of any individual pore is limited. A range of reasons for this fact were proposed. It could be the case that the resolution of the camera (either spatially, temporally, or thermally) is insufficient to detect the subtle thermal signatures involved in pore formation. Alternatively, it could be the case that such thermal signatures (or rather their differences from non-pore-creating thermal signatures) are too subtle for detection by any state-of-the-art camera system.

Searching for the formation of pores with an on-axis, higher-speed, two-wavelength camera, which allows for greater resolution entirely dedicated to the melt pool, would be a necessary first step in understanding this question.

8.4 - Future Work Conclusions

The ultimate aim of much current research into L-PBF is the implementation of closed-loop control, in which in-situ monitoring systems detect in real time when a build process is producing mechanical or microstructural defects, and adjusting machine parameters to correct the process.

This current work has shown that there remains much understanding to be gained in the context of relating individual thermal signatures to defects. Should the art form ever advance to the stage where specific pores can be detected at their moment of creation, or where the microstructure beneath a melt pool can be accurately predicted, then much higher resolution monitoring equipment is needed, in combination with XCT and other ex-situ methods, to find if individual thermal signatures can reveal substantive information about their solidification.

9 - Appendix

The following sections have been removed from their respective chapters, and moved to this appendix, as the content contained within them was not considered sufficiently relevant or significant to include in the main thesis.

9.1 - Literature Review

A subsection on Machine Setup was originally included in Section 2.5 - Factors Affecting Build. It will be presented here:

In addition to the most obvious and changeable machine parameters (laser speed and power, hatch spacing, and layer height), there are several background-type settings and configurations that can significantly affect build quality. A few of these will now be explored.

9.1.1 - Powder Coating Type

There are five major categories of powder recoating systems, used to spread powder across the build plate:

- Hard blades, usually composed of steel, and which move across the build plate at high speed [25]
- Soft blades, usually polymer lips, which can elastically compress upon contact with minor swelling from components [25]
- Rolling cylinders, composed of hard or soft materials, often rotating counter to its direction of movement across the build plate [337]
- Brushes or bristles, often made from carbon fibre, that gently push powder particles across the build plate without damaging fragile components [338]
- Contactless systems, usually based on conveyor belts, vibratory devices, or air-blades, which ensure zero physical contact with components on the plate [339]

As determined by Uriati et al. [340] and Reijonen et al. [341], when it comes to hard vs soft blades (the two most commonly used recoater types) there exists a clear trade-off. Hard blades tend to deliver better powder bed uniformity, and increased dimensional accuracy, until a swollen feature causes a collision. At this point either the feature will break free, resulting in significant damage to the component, or the build will halt entirely, requiring a total reset. Soft blades deliver poorer powder uniformity, and lower dimensional accuracy, but are much more forgiving of swollen features, and can continue working even if significantly damaged. A damaged recoater blade will, however, lead to uneven tracks of powder running across the bed, which may also render a component unusable.

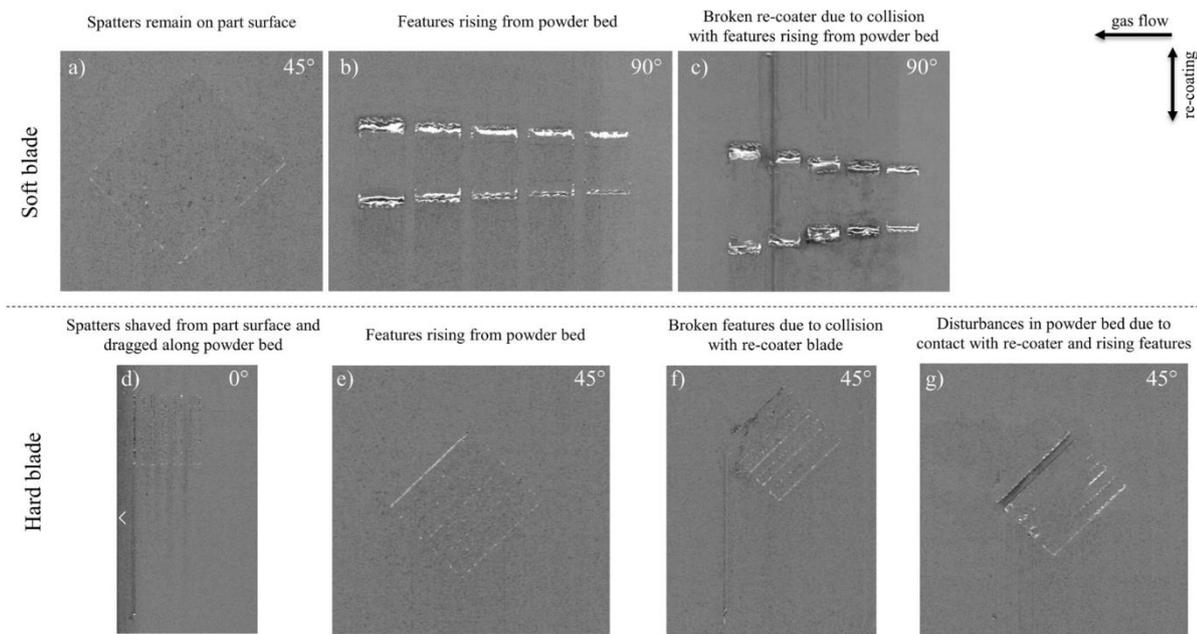


Figure 9.1

Reijonen demonstrated how soft blades are more forgiving of swelling, but may become damaged, leaving tracks in the powder bed; whereas hard blades shave off spatter and swollen features, which may damage the part [341]

Doris [342] attempted to build thin-wall from IN-718 across 16 different L-PBF machines, 5 of which used soft blade recoaters, 7 of which used hard blades, 3 of which used brushes, and one of which used a non-contact system. Only four of the systems successfully built the thinnest walls (0.1mm thick), three of which used soft blades and one of which used non-contact. It was noted that amongst the many reasons for machines to fail at building these thin walls, contact with recoaters was common.

9.1.2 - Inert Gas Flow

The type of inert gas pumped through the build chamber, the positioning of the outlets, the velocity of the flow, and the positioning of the components relative to the outlets, can all have significant impacts upon the melt pool dynamics and component quality.

Pauzon et al. [343] studied the effects of using argon vs nitrogen gas during the build process of IN-718. It was found that very few observable differences resulted from the gas choice. The difference in nitrogen content between the powder and component (due to nitrogen evaporation from the melt pool), was found to be about 400ppm for argon gas, but only 350ppm for nitrogen gas. It was suggested by the authors that this decrease of reduction in nitrogen content may lead to increased nitride and carbonitride phases, which are known to act as both crack initiation sites and grain refiners, depending on their context.

Decker et al. [344] also experimented with variations in inert gas composition for IN-718, namely: 100% argon; 70% argon, 30% helium; 30% argon, 70% helium; and 68% argon, 30% helium, 2%

hydrogen. Spatter generation was found to be reduced by increasing helium content, arguably due to helium's higher thermal conductivity and therefore ability to cool airborne spatter faster. This reduction in spatter in turn led to a reduction in porosity for the samples processed in high helium content. Both the microstructure and mechanical properties of all samples were not affected by the choice of gas in a statistically significant way. The authors suggest that engineers should consider the addition of some helium into the standard argon gas used in L-PBF, for the purpose of increasing process stability.

The velocity and uniformity of gas flow that a melt pool experiences, will differ greatly depending upon where it is relative to the gas inlet and outlet nozzles. Abeyta [345] experimented with the gas flow across an EOS M290 machine, and found that the flow velocity was reduced by 64% as gas moved from the inlet nozzles positioned 35mm above the upper build plate (2.5 m/s) to the outlet nozzles at the lower front (0.9 m/s). At a height just 8mm above the build plate, flow velocity varies between 0 and 2 m/s.

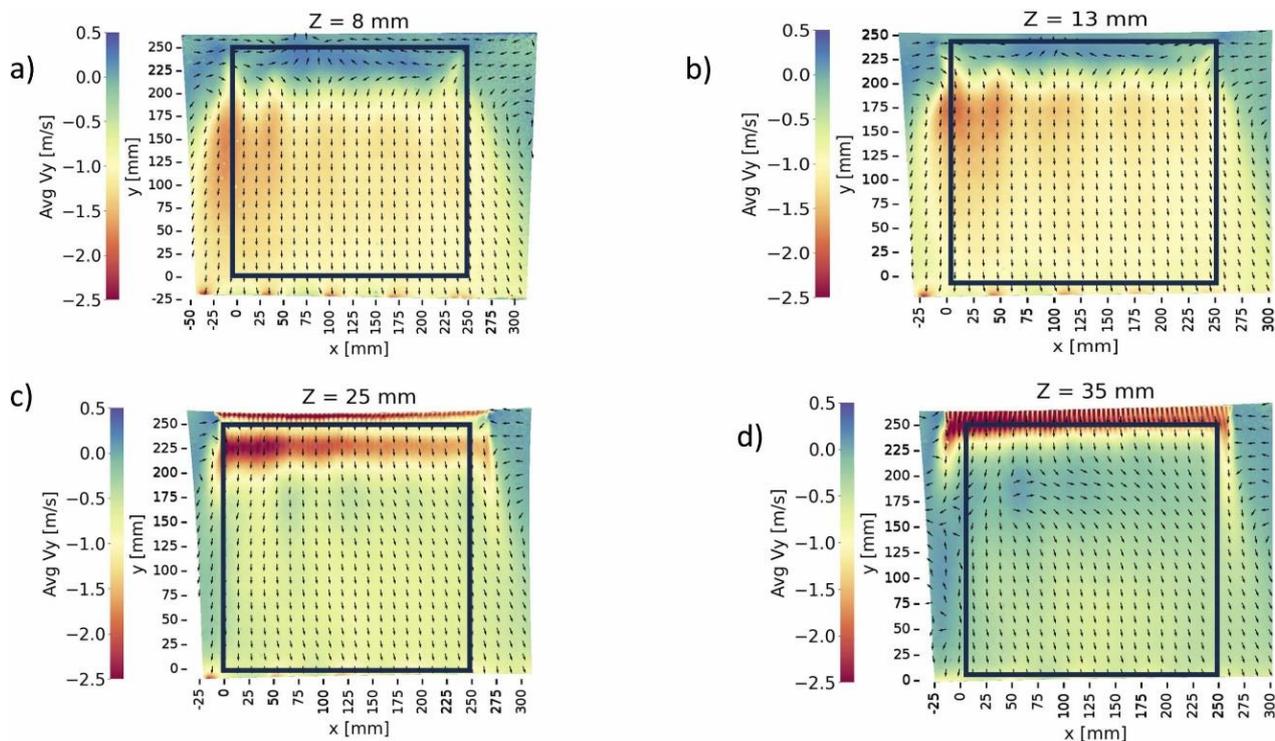


Figure 9.2

Flow velocities at various heights above the build plate, from Abeyta [345]

These results naturally prompt a question about the importance of flow velocity to melt pool properties and component quality. Chien et al. [346], using a coupled Computational Fluid Dynamics (CFD) and Discrete Element Method (DEM) simulation, showed that spatter particles are ejected from melt pools with velocities of 4 to 7 m/s, and that higher inert gas flow velocities, approaching these spatter velocities and ideally directed opposite to the scan direction, are more effective at removing spatter particles from the laser's path. Greater flow velocity was also shown to remove accumulated heat from the gases above the melt pool, thereby significantly increasing

cooling rates. Combined, these effects were suggested to improve part quality notably. It is noted, however, that a trade-off exists: very high gas flow velocities will lead to the removal of finer powder particles from the bed.

9.1.3 - Beam Profile

A beam's profile is a measure of the laser energy that it delivers across the plane normal to its direction. Though all relevant data is contained in the profile, for practical purposes three interdependent parameters are usually discussed: the beam shape (i.e. Gaussian, top hat, ring, etc.); the spot diameter; and the laser's focal point above or below the build plate. It can be very difficult to alter one of these parameters without significantly affecting the others.

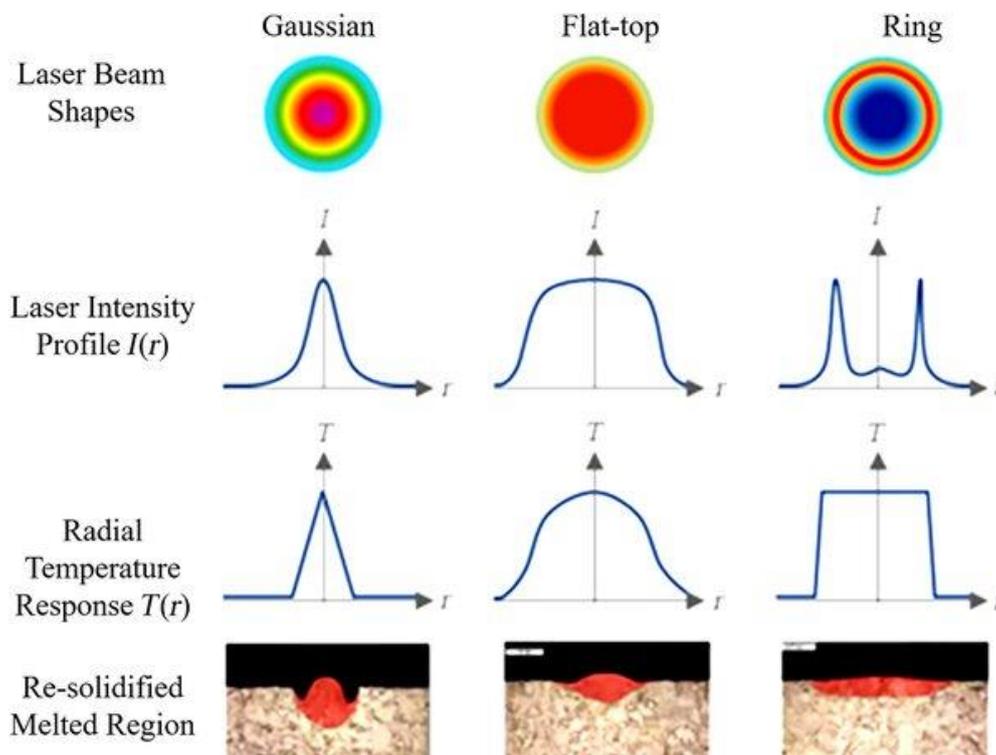


Figure 9.3

The shape and effects of three common beam types [347]

Cozzolino et al. [348] experimented on the effects of using a Gaussian vs ring beam shape in the building of IN-718. The Gaussian beam diameter (defined as the point at which intensity is $1/e^2$ of the maximum) was $80\mu\text{m}$. The ring diameter was around $240\mu\text{m}$, with 90% of the laser energy being concentrated at the edges.

Components produced with the ring profile were found to have wider, shallower melt pools, as opposed to the narrow and deep pool from the Gaussian. The ring components exhibited more directional solidification, and thinner columnar grains. The effects of beam profile on mechanical properties and surface roughness were considered negligible.

The major advantage of the ring profile was a 19% decrease in build time, attributable to the lower number of scan vectors required for the larger spot. Since a VED of 67 J/mm³ was used for both builds, the laser energy consumed was near-identical in each. Due to the decrease in machine time, however, the entire process consumed 10% less energy for the ring than the Gaussian.

Mirzabeigi et al. [349] identified similar trends, experimenting with different diameters of Gaussian beams in IN-718. By using a beam of diameter 400µm, build rates of up to 16 mm³/s were achieved, which is around four times the industrial norm. This increase in productivity came at the cost of wide, shallow melt pools, with fine, columnar grains that may be harmful to a component, depending on intended usage (e.g. if it were loaded in tension along the build direction).

While many commercial L-PBF machines can not alter the diameter of their laser beams, a similar effect can be replicated by adjusting the beam focus. Zhong et al. [350] built IN-718 samples with a Gaussian beam, at five different focus levels ranging from 0mm (focussed on the base plate, with a diameter of 70µm), to +20mm (focussed above the base plate, with an effective diameter of 197µm). It was found, for the power, speed, and hatch parameters tested, that 0mm and +5mm focus led to melt pools too narrow and deep to overlap with each other, that +10mm focus led to full density, and that +15mm and +20mm focus led to insufficient melting. It therefore stands to reason that, by adjusting power, speed, and hatch to provide an appropriate VED regardless of spot size, using defocussed lasers could lead to the productivity gains without microstructural sacrifices that other researchers achieved.

9.1.4 - Base Plate Preheating

Preheating the base plate in a L-PBF build can have two major effects: thermal gradients and cooling rates inside the melt pool are reduced; and the temperatures to which material is cooled while still inside the build chamber, is increased. Combined, these effects can influence residual stresses, cracking behaviour, microstructure, and mechanical properties.

Park et al. [351] demonstrated the effects of preheating the build plate in IN-718, on each of these four categories, with temperatures ranging from 50°C to 150°C. It was found that in the X direction, preheating did not significantly affect residual stress, only reducing the value from 77 MPa to 68 MPa as the temperature increased from 50°C to 150°C. In the Z direction, however, residual stress was reduced by 22%, from 422 MPa to 332 MPa. This reduction of residual stresses in turn led to a reduction of cracks within the components.

As preheat temperature increases, density increased from 99.982% at 50°C to 99.992% 150°C. The size and shape of melt pools were not observed to vary across the temperatures, although the anisotropy of grains was found to decrease as temperature increased. As temperature increased, inter-granular misorientation was found to increase, while intra-granular misorientation decreased.

Interestingly, there was no significant difference in hardness, tensile strength, yield strength, or elongation at failure, as preheat temperature increased. It is therefore concluded that increasing preheat temperature is a practical remedy to many of the defects that arise from residual stress (thermal warpage, cracking, etc.), but will not have a major impact on components that are not at risk of such defects.

9.1.5 - Base Plate Material

Welding two different alloys onto each other is not a trivial task. Significant mismatches between melting ranges and thermal stress responses can make bonding difficult from a mechanical perspective, and an inability of one chemical composition to nucleate onto the other can cause incompatibility from a metallurgical perspective.

There exists in the literature very few experiments that report an explicit comparison of IN-718 build quality on different base plate alloys. Shrivastava et al. [352] tested the bonding between IN-718 powder and wrought IN-718. As expected, these chemically-identical alloys were highly compatible, with excellent bonding between them.

Chen et al. [353] demonstrated that the bonding of IN-718 onto a 316L steel substrate is more achievable via L-PBF than it is through traditional welding, due to the rapid cooling rates, which suppress Laves formation. It was suggested that, with careful parameter control, adequate metallurgical bonding can take place.

9.1.6 - Future Machine Possibilities

Many researchers are beginning to experiment with very advanced L-PBF techniques, in which machines are augmented and enhanced to allow for better printing and higher quality builds. A very quick summary of some such techniques will now be presented.

Within a melt pool, charged metallic atoms circulate due to convective and Marangoni forces. Should a magnetic field be applied to the melt pool, a Lorentz force will be generated, further impacting the molten alloy's movement. Via careful positioning of this field, forces can be exerted on melt pools which cause them to be deep and narrow, or wide and shallow. Yu et al. [354] demonstrated that the application of a static magnetic field to IN-718 melt pools can result in both grain refinement, and an increased fraction of equiaxed grains. Liu et al. [355] experimented with non-static magnetic fields in IN-718, to create an electrostatic stirring effect inside the melt pool. Convective currents in the liquid were increased, which inhibited the segregation of alloying elements like Nb and Ti.

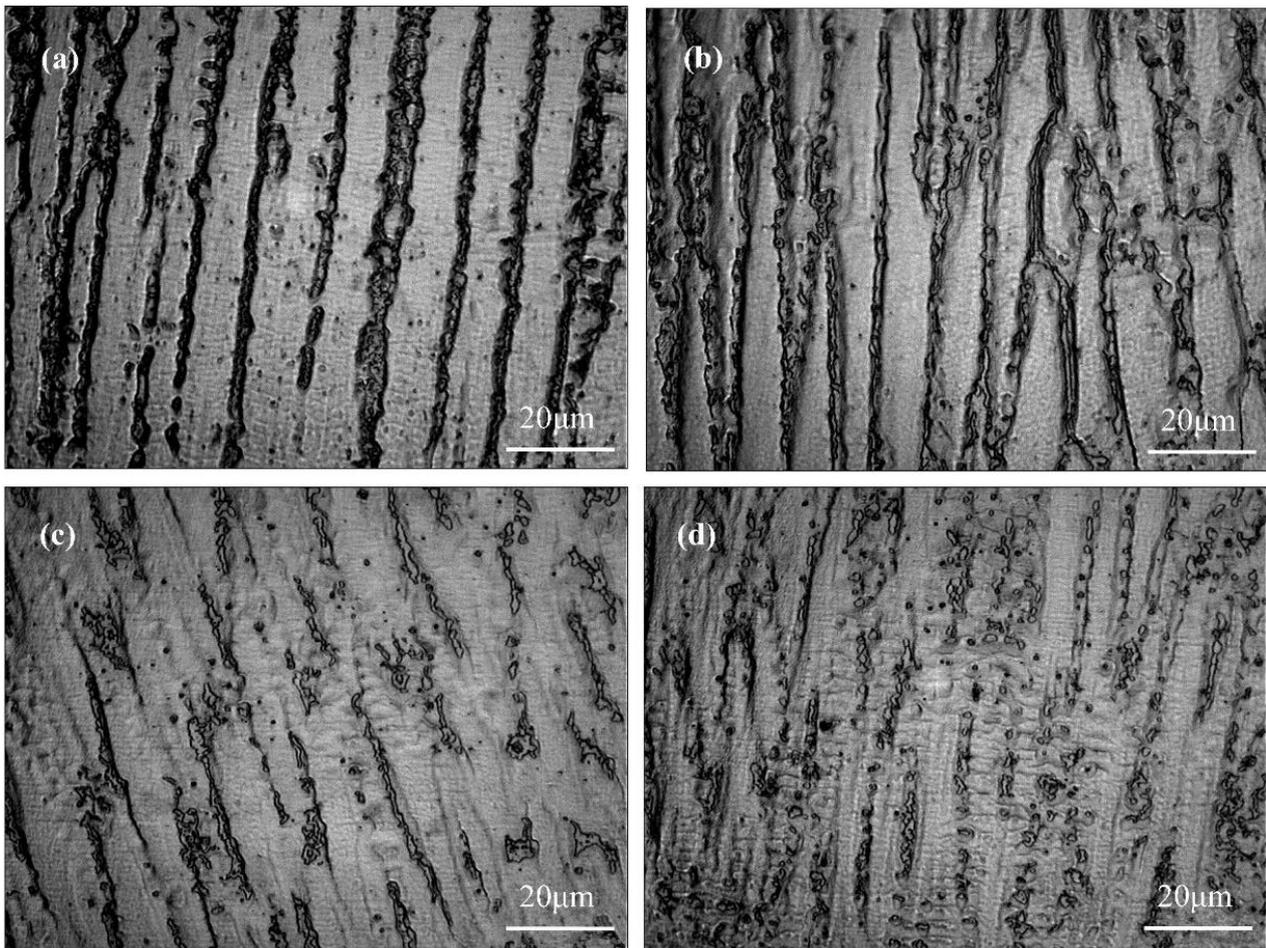


Figure 9.4

Li's microstructures, built under 0mT (a), 30mT (b), 50mT (c), and 80mT (d). Note how segregation lessens as magnetic flux density increases. [355]

The use of ultrasonic vibrations is another method of increasing melt pool stirring. Wang et al. [356] demonstrated that ultrasound can induce cavitation and acoustic streaming (the generation of macroscopic fluid flow from sound waves) in a IN-718 melt pool, which resulted in better mixing, the homogenisation of heat and density, and ultimately more equiaxed grains and fewer Laves phases.

Using a pulsed laser, Wu et al. [357] successfully used plasma to enhance the L-PBF process. After each melting pulse (around 200ns long, and 0.03mJ in energy), an ionising pulse (around 4ns and 0.1mJ) followed to turn some of the melt pool's surface into plasma. The very high pressure exerted on the melt pool by this plasma plume promoted better flow in the liquid, suppressed balling, and reduced porosity.

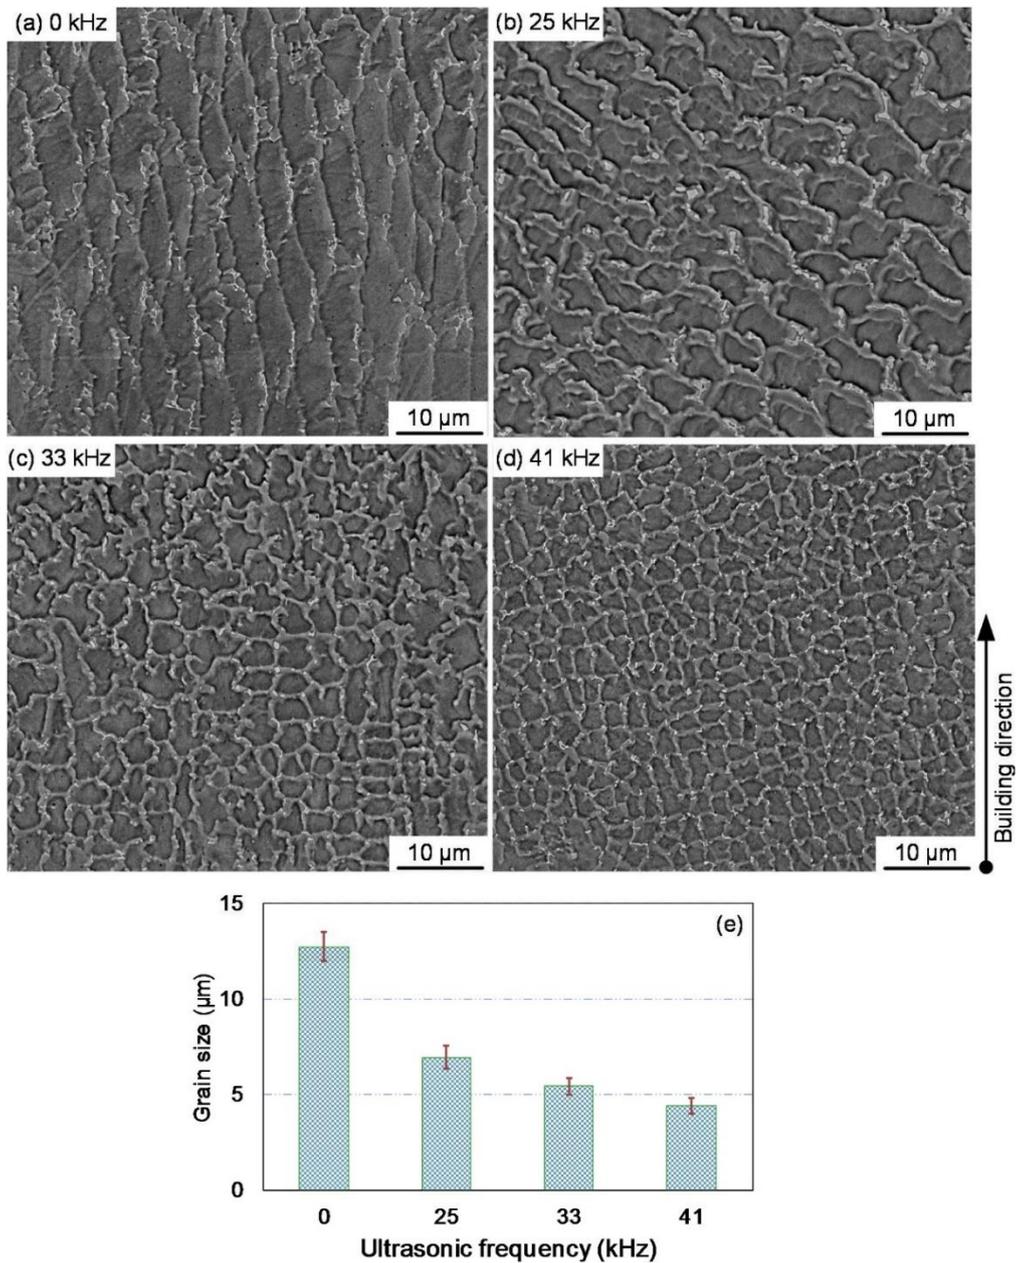


Figure 9.5

Wang's microstructures, built under various ultrasonic frequencies, and a plot showing the effects on mean grain size [356]

9.2 - In-Situ Monitoring and the Microstructure

The following charts and discussion were originally included in Section 5.3 - Data Analysis, but have been moved due to their lacking in significant insight.

9.2.1 - Thermal Camera Data
9.2.1.1 - Mean Melt Pool Temperature

The mean temperature of all the pixels inside the melt pool was calculated, for each melt pool image. In Figure 9.6, the means of these mean temperatures are displayed, for each of the 27 cubes.

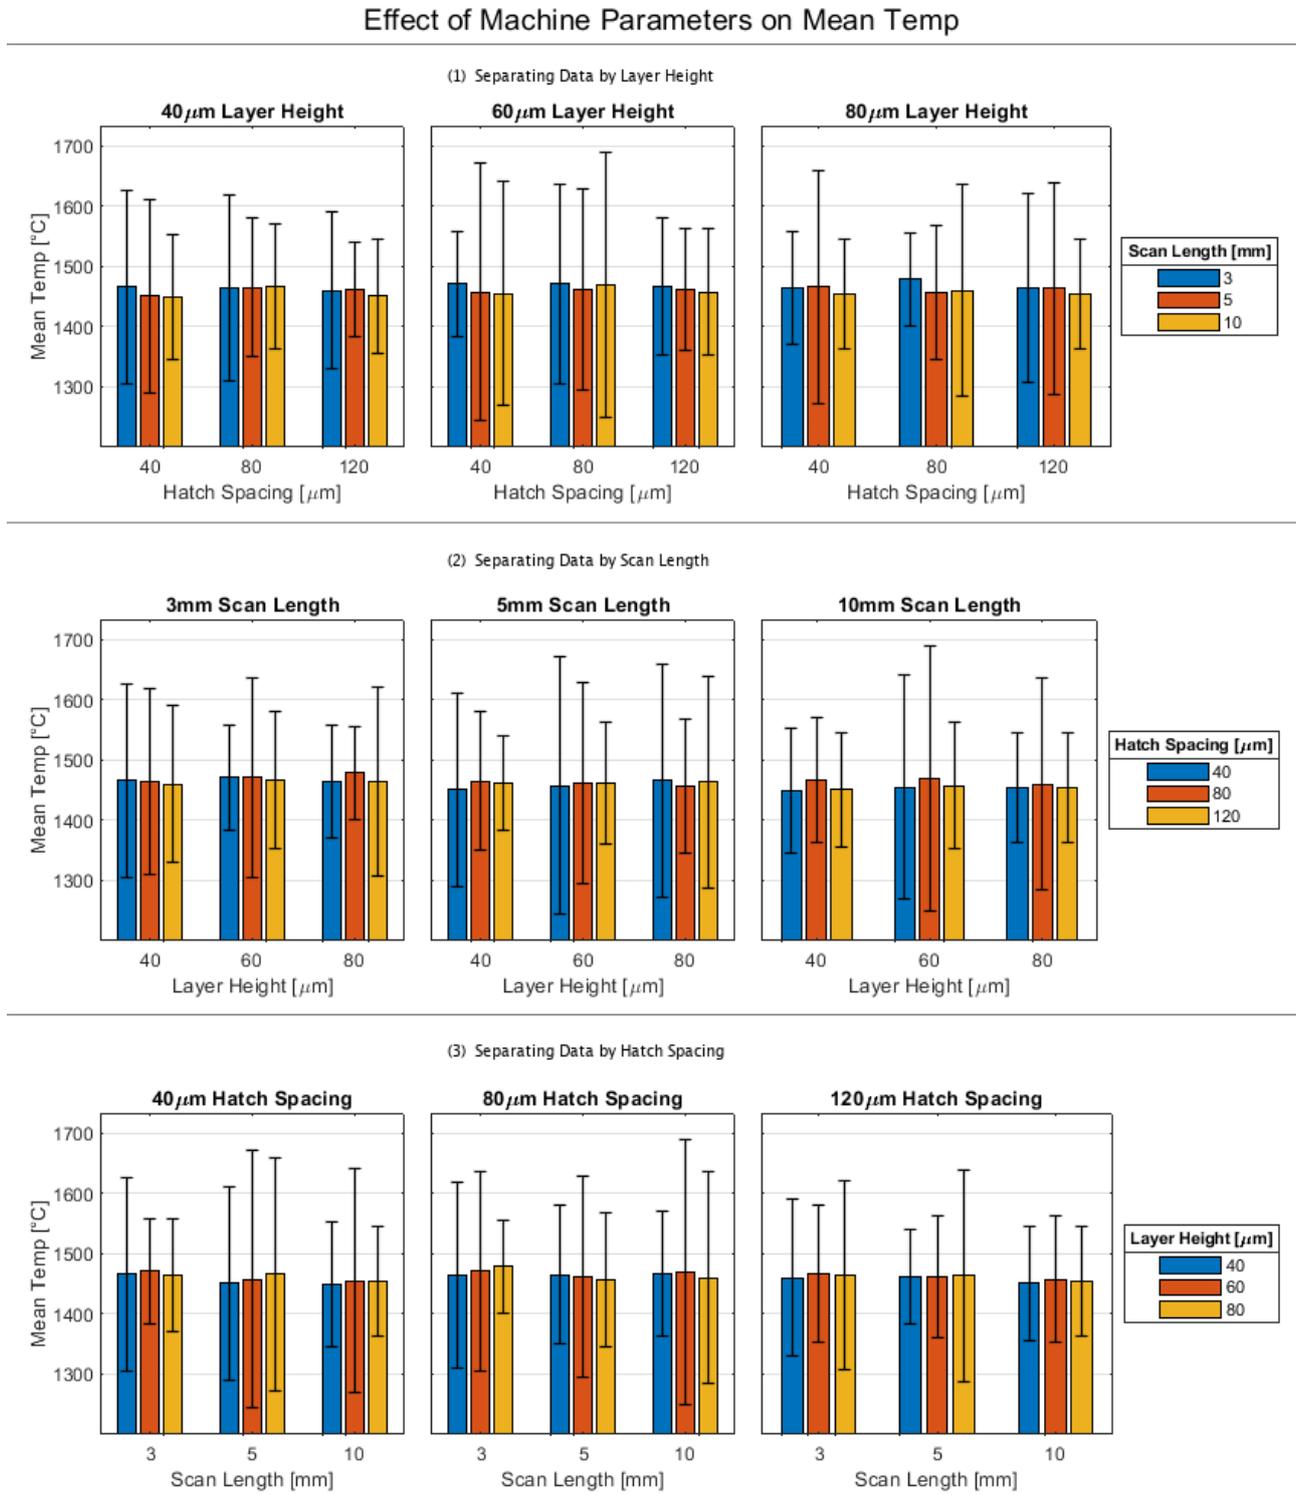


Figure 9.6
The effects of machine parameters on mean melt pool temperature

As can be seen from these charts, little valuable information can be drawn from analysing the mean melt pool temperatures. The errors (standard deviation of the measurements) are far more significant than the differences between the means. This extreme variability in the mean of each pool (which in turn leads to little variability in the mean of means) is explained by the fact that the number of peripheral pixels in each pool, close to the liquidus threshold, can vary significantly.

9.2.1.2 - Peak Melt Pool Temperature

A single value for peak temperature was extracted for each melt pool. Figure 9.7 shows these peaks.

As can be seen from row (1) of Figure 9.7, where the data is separated by layer height, there exists a trend for peak temperature to decrease as scan length increases. This trend matches the expectation that longer scan lengths will result in the temperature of solid material adjacent to melt pools being lower, therefore increasing thermal conduction, and lowering the temperature of the liquid material.

As can be seen from row (2), the effect of hatch spacing on peak temperature is inconsistent. It was expected that at higher hatch spacing, due to lower thermal conduction, peak temperatures would be higher. This is generally true for the cubes with 5mm scan length, but not for the 3mm or 10mm cases. There does not appear to be any alternative trend that is evident in the 3mm or 10mm samples. It was reckoned that the non-existence of a trend is either due to the hatch spacing expectations being incorrect (this will be addressed further later), or due to the stochastic nature of the measurements and the large error in them.

As can be seen from row (3), peak temperature generally tends to increase with layer height, although not always to a significant extent. This finding roughly matches the expectation that increasing layer height will decrease thermal conduction into the solid material below the melt pool, therefore increasing melt pool temperature.

Effect of Machine Parameters on Peak Temp

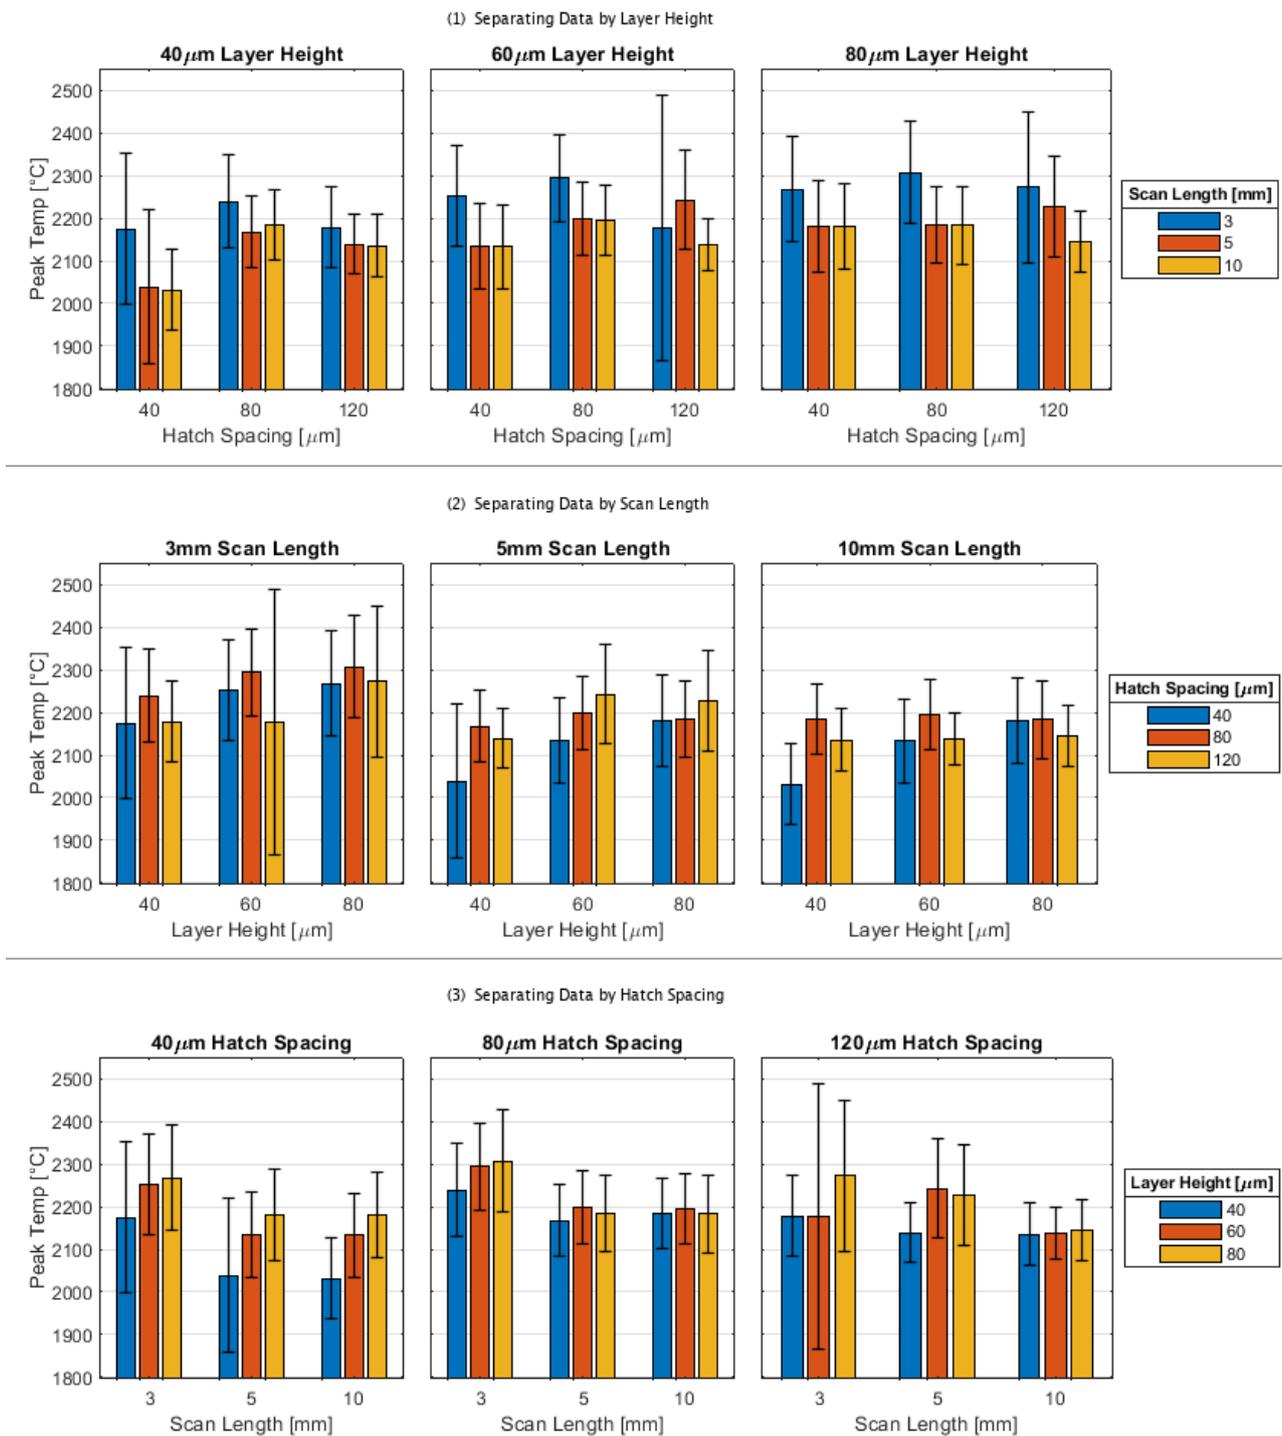


Figure 9.7

The effects of machine parameters on peak melt pool temperature

9.2.1.3 - Melt Pool Width

Using the REGIONPROPS function, the minor axis of each melt pool was measured. These distances were taken as the pool's width, and are shown in Figure 9.8.

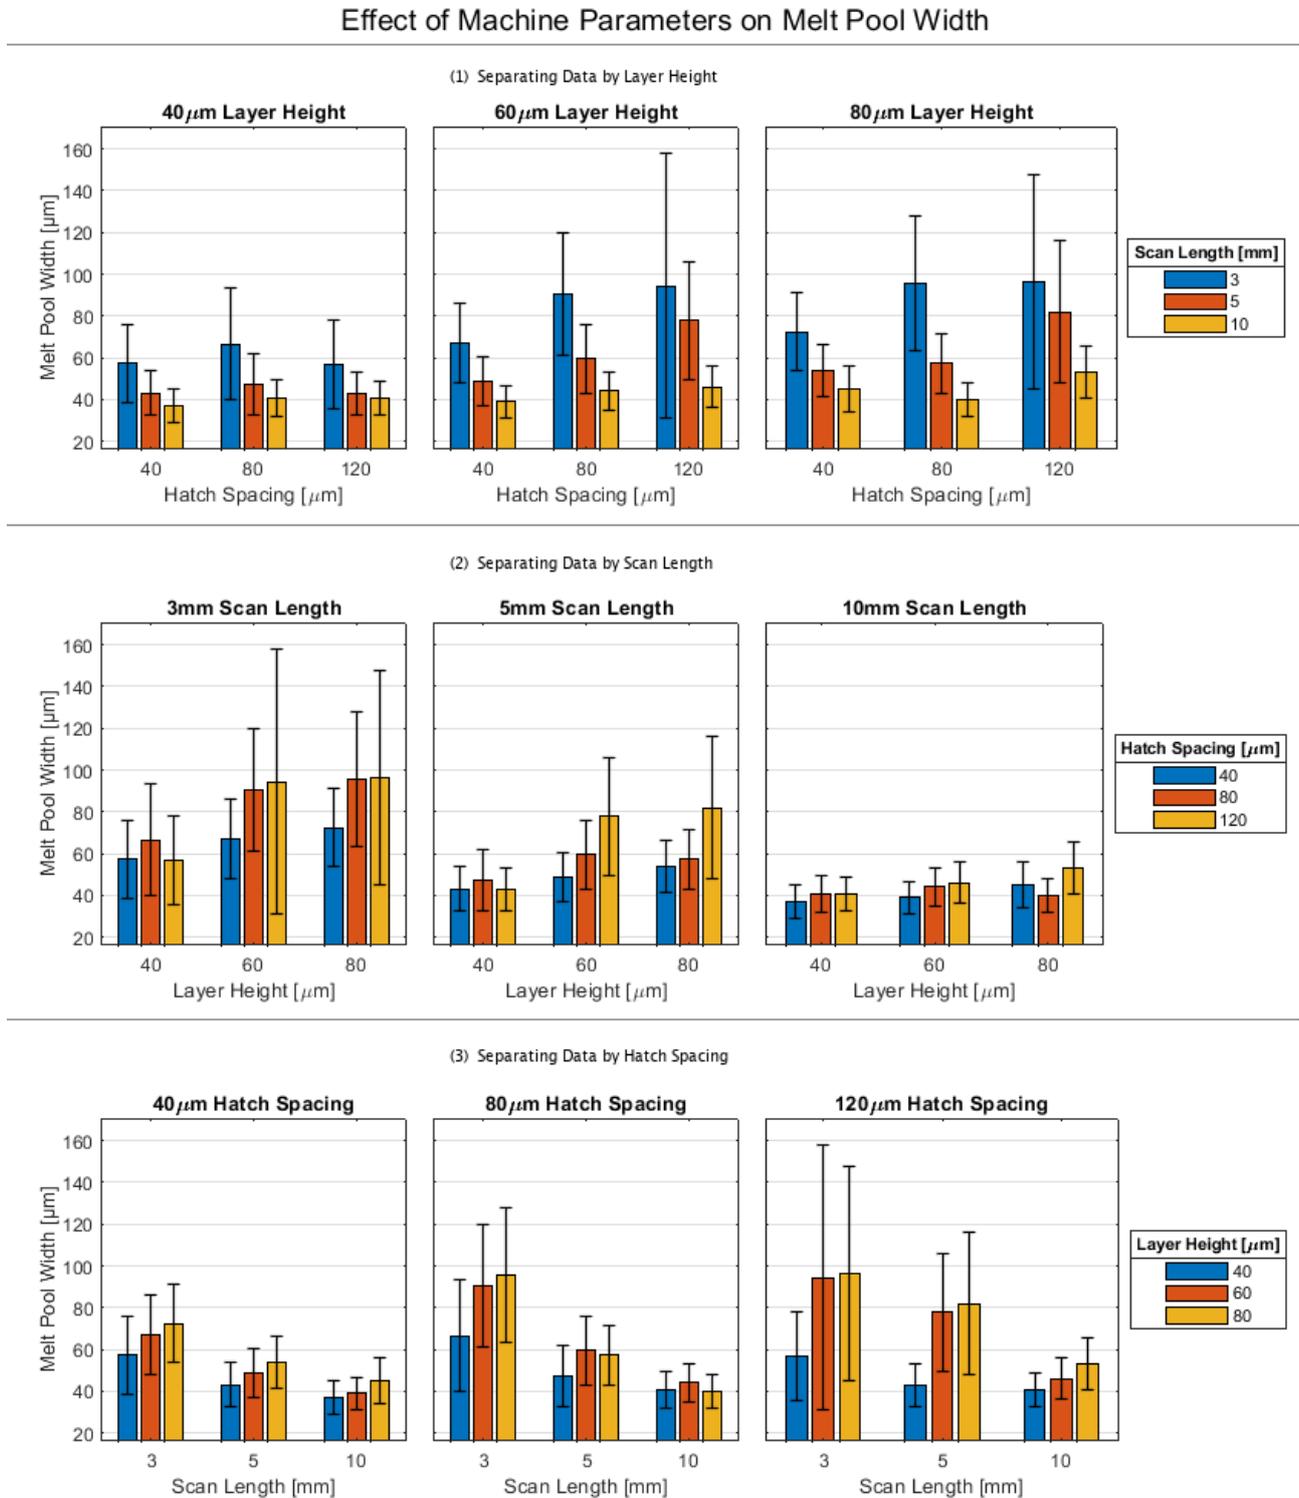


Figure 9.8
The effects of machine parameters on melt pool width

As can be seen from row (1), the widths of the melt pools consistently decrease as scan length increases. This finding matches the expectation that an increased scan length leads to increased thermal conduction, which leads to a lower melt pool width.

As can be seen from row (2), melt pool width generally tends to increase with increased hatch spacing. These increases are not very significant in comparison to the errors. According to the primary theory of hatch spacing, as hatch spacing increases thermal conduction will decrease, and therefore pool widths will be greater.

As can be seen from row (3), melt pool width consistently increased as layer height increased. This trend matches the expectations that larger layer heights lead to lower thermal conduction, which leads to increased melt pool widths.

9.2.1.4 - Melt Pool Tail Cooling Rate

Figure 9.9 presents the same data as that in the Tail Thermal Gradient chart, but with all thermal gradients multiplied by the laser speed, to find the cooling rate.

Though these charts do not provide additional insight, beyond that which was already provided by the thermal gradient charts, they still serve as a useful sense check. The cooling rates all lie between 3 and 7×10^6 °C/s. This range of cooling rates agrees with those found in literature, demonstrating that the method used to extract thermal gradients from the melt pools was sound.

9.2.1.5 - Spatter

An attempt was made to quantify the amount of spatter being ejected from the melt pool, in each cube build. Each image was analysed, after 1380°C thresholding, but before IQR filtering, to search for regions of hot material outside the melt pool. Whenever a region measuring at least 10,000px in area was found, that image was declared to ‘contain spatter’. The percentage of images with spatter was found for each cube, and is shown in Figure 9.10.

As can be seen, no clear trends exist. The percentages vary wildly, from as low as 0.6%, to as high as 98.4%. It would be expected that some correlation exists between quantity of spatter and LoF porosity. As will be demonstrated in the next subsection, however, no such correlation can be found. It was therefore reasoned that either an error existed in the methodology of determining spatter from thermal camera images, or that the thermal camera was simply unable to capture spatter in a meaningful way. It was reckoned that, should the latter be the case, the camera’s extremely shallow depth of field (which would greatly reduce the temperature reading for any airborne projectile) would be to blame.

The images were analysed again, this time using 1,000 pixels and 100,000 pixels as the area thresholds for spatter. Using these new thresholds significantly increased and decreased the spatter percentages, respectively, without providing any kind of clarity on trends that may exist.

Ultimately, it was concluded that the thermal camera is effective at capturing low-flying spatter particles, but not higher-flying particles, and that analysing low-flying particles alone is of little benefit in determining anything useful about the state of the build.

Effect of Machine Parameters on Tail Cooling Rate

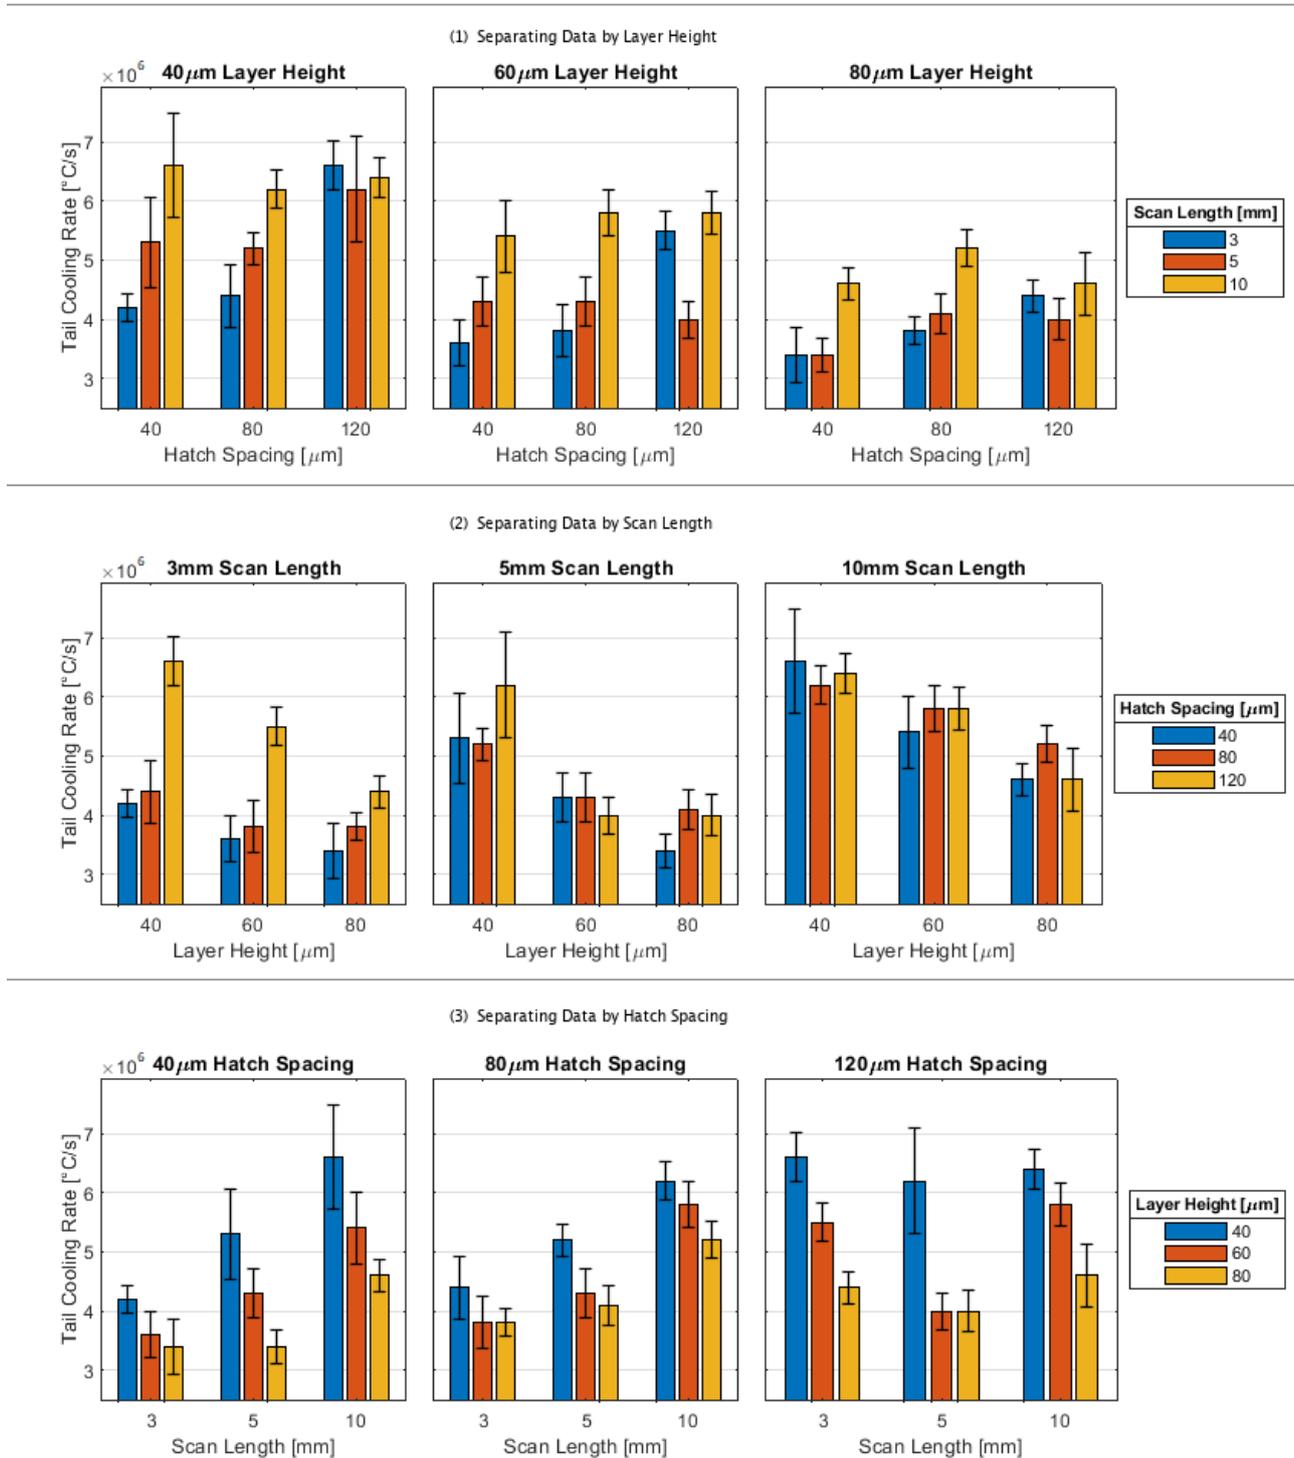


Figure 9.9

The effects of machine parameters on melt pool tail cooling rate

Effect of Machine Parameters on Spatter

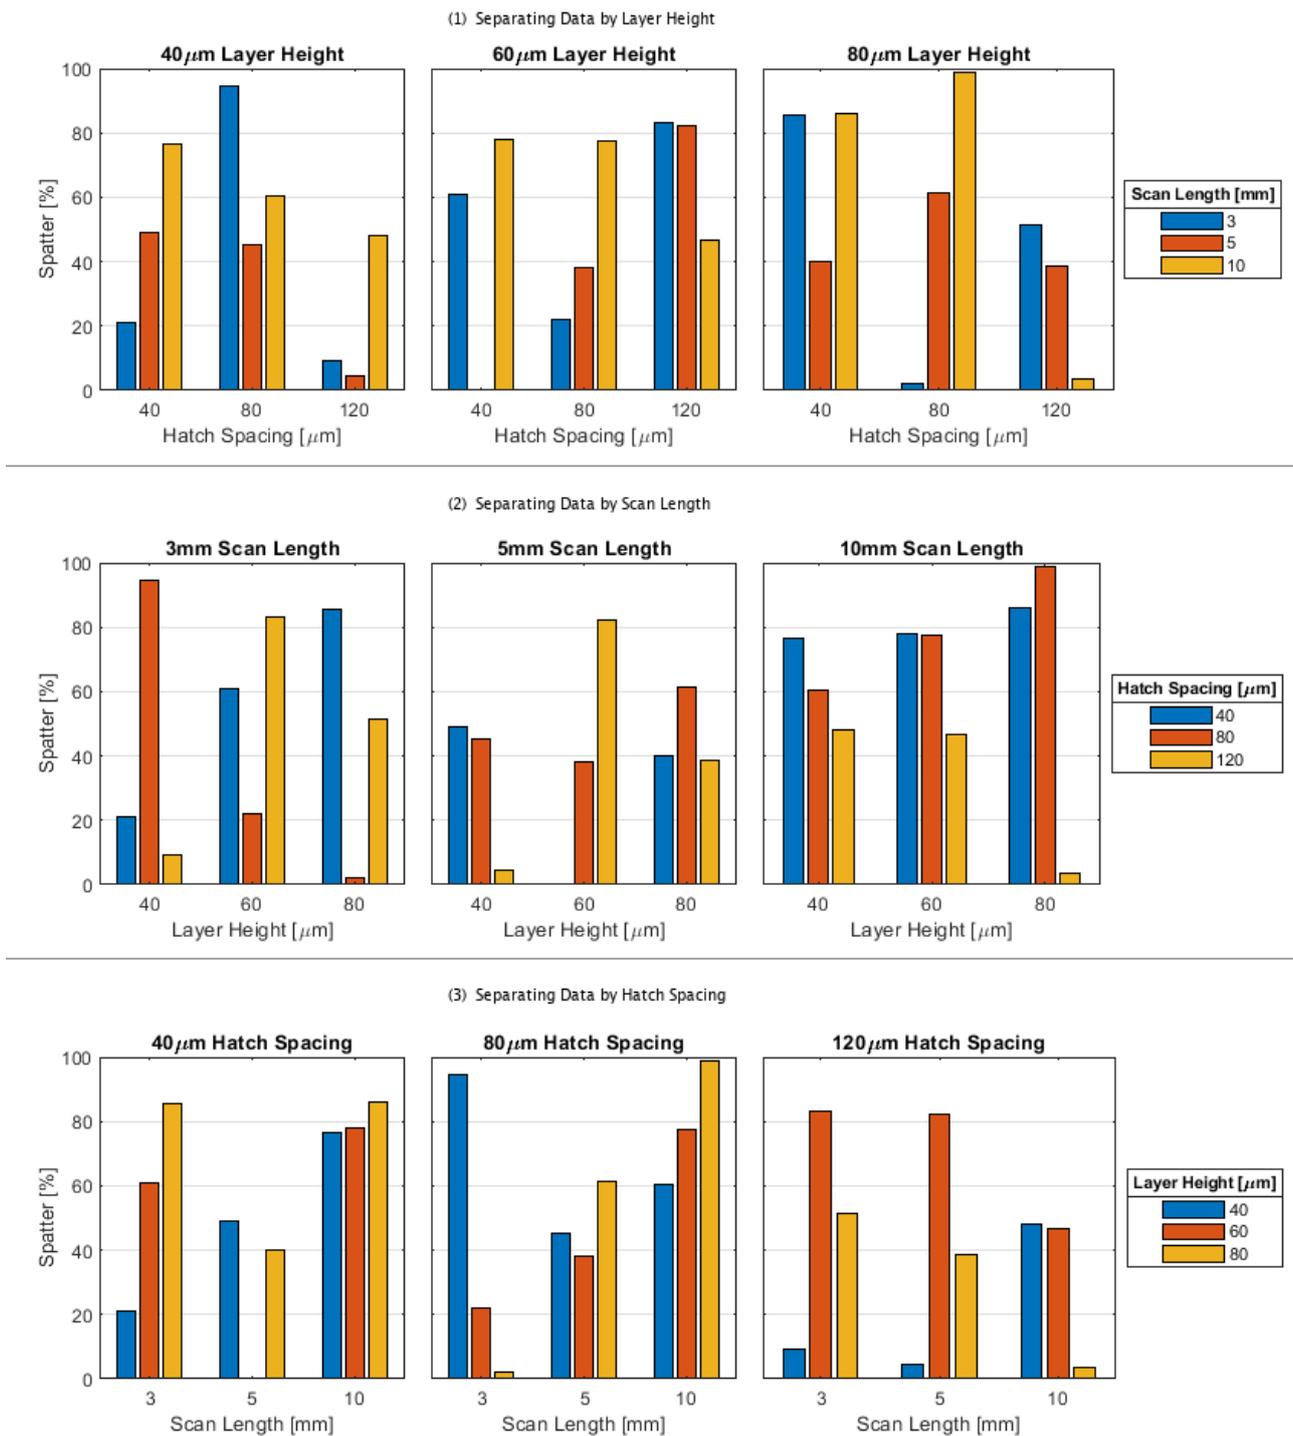


Figure 9.10

The effects of machine parameters on spatter

9.2.2 - Optical Microscope Data

In each optical image, the REGIONPROPS function was used to calculate the circularity of each pore. An area-weighted mean pore circularity, and an area-weighted standard deviation, were calculated, as shown in Figure 9.11. The equations used to find these numbers will be described in the EBSD subsection.

As can be seen from Figure 9.11, the area-weighted mean pore circularity is close to 1 (an almost-perfect circle) for cubes with high energy density and low porosity, and decreases to closer to 0.5 (highly elongated) for cubes with low energy density and high porosity. These results further demonstrate that LoF (which always exhibits irregular shapes) is the primary source of porosity.

9.2.3 - SEM Data

As described previously, it was theorised that measuring the standard deviation of the aspect ratio of the cells inside an SEM image, could serve as a metric of the anisotropy of the cells. It was expected that in cubes with lower thermal gradients, and therefore slower solidification (i.e. smaller scan lengths, hatch spacings, and layer heights), there would be lower anisotropy in the cells. It was expected that in cubes with greater thermal gradients, there would be greater anisotropy. The standard deviation of the cell aspect ratios are plotted Figure 9.12.

As can be seen, no clear trends of any kind appear to exist, strongly suggesting that either the standard deviation of cell aspect ratio is not a good proxy for cell anisotropy, or that cell anisotropy cannot be easily predicted by machine parameters, or both.

The standard deviations of the cell aspect ratios were plotted against thermal gradient, as shown in Figure 9.13.

Again, no clear trends exist, adding weight to the suggestion that no simple correlation exists between the temperature gradient inside/around a melt pool, and the anisotropy of the sub-grain microstructure that exists after solidification.

Effect of Machine Parameters on Pore Circularity

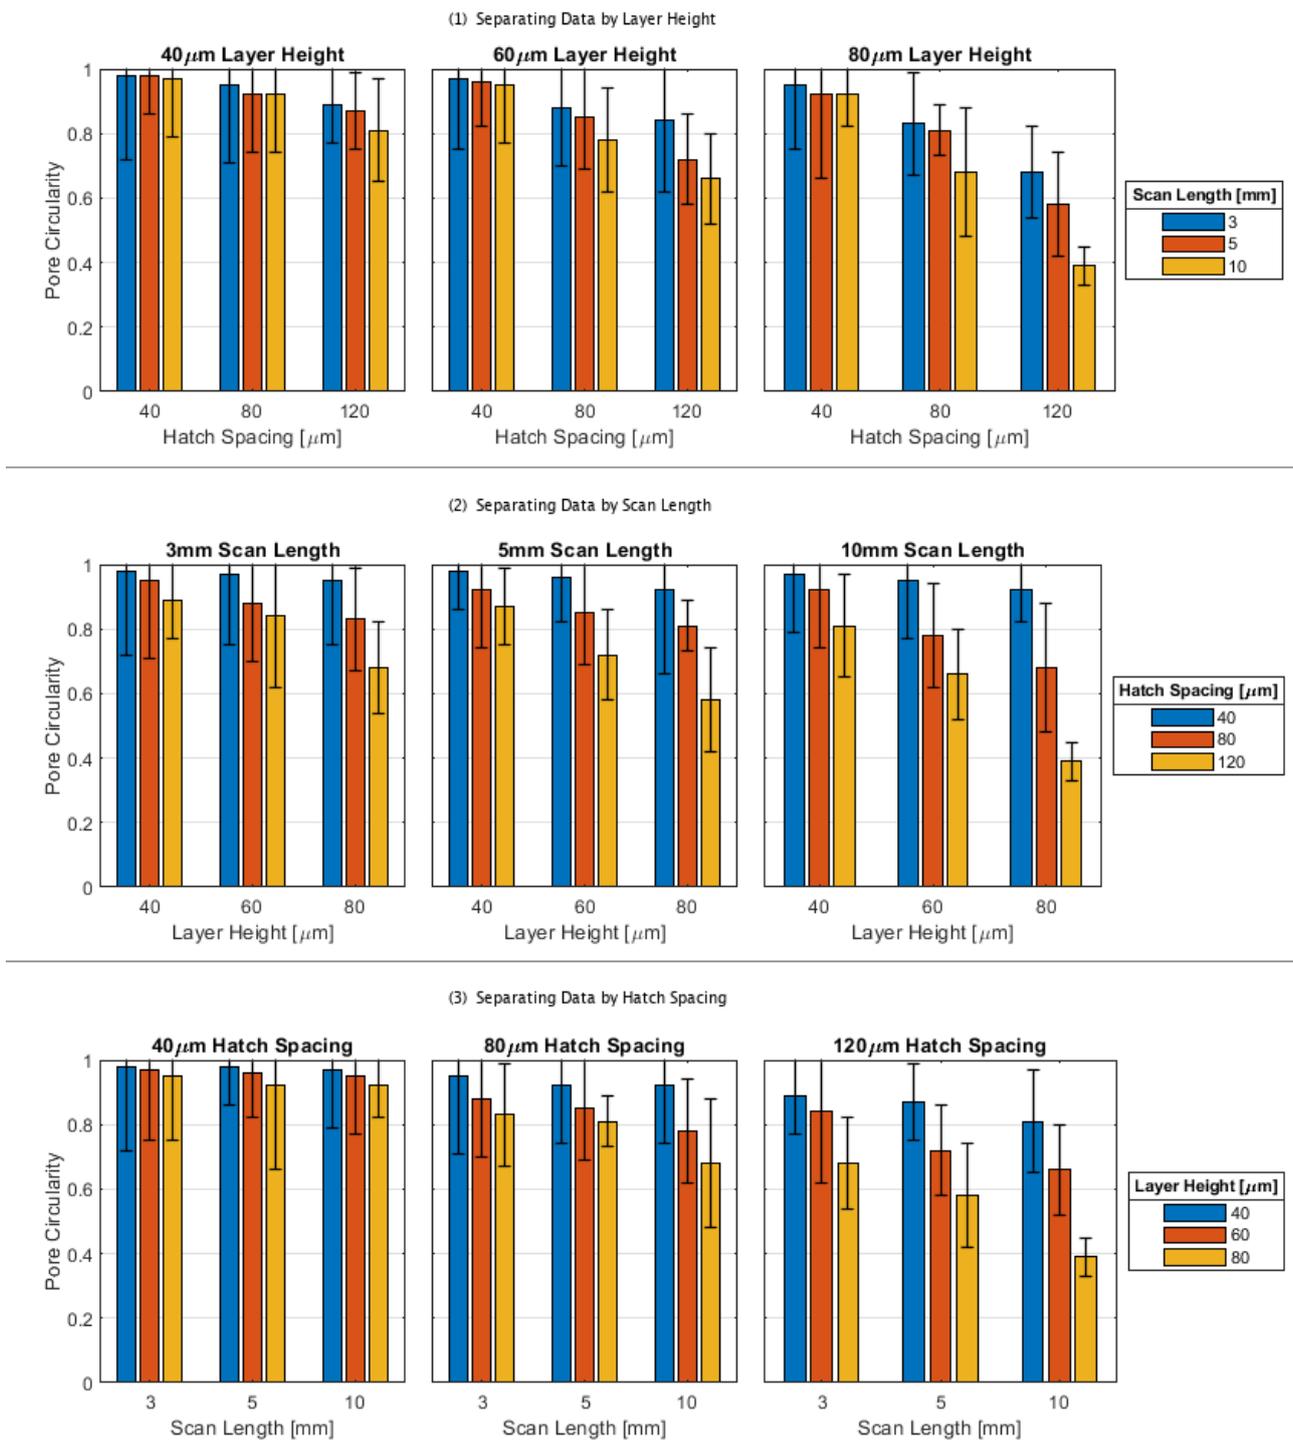


Figure 9.11

The effects of machine parameters on pore circularity

Effect of Machine Parameters on Std. Dev. of Cell Aspect Ratio

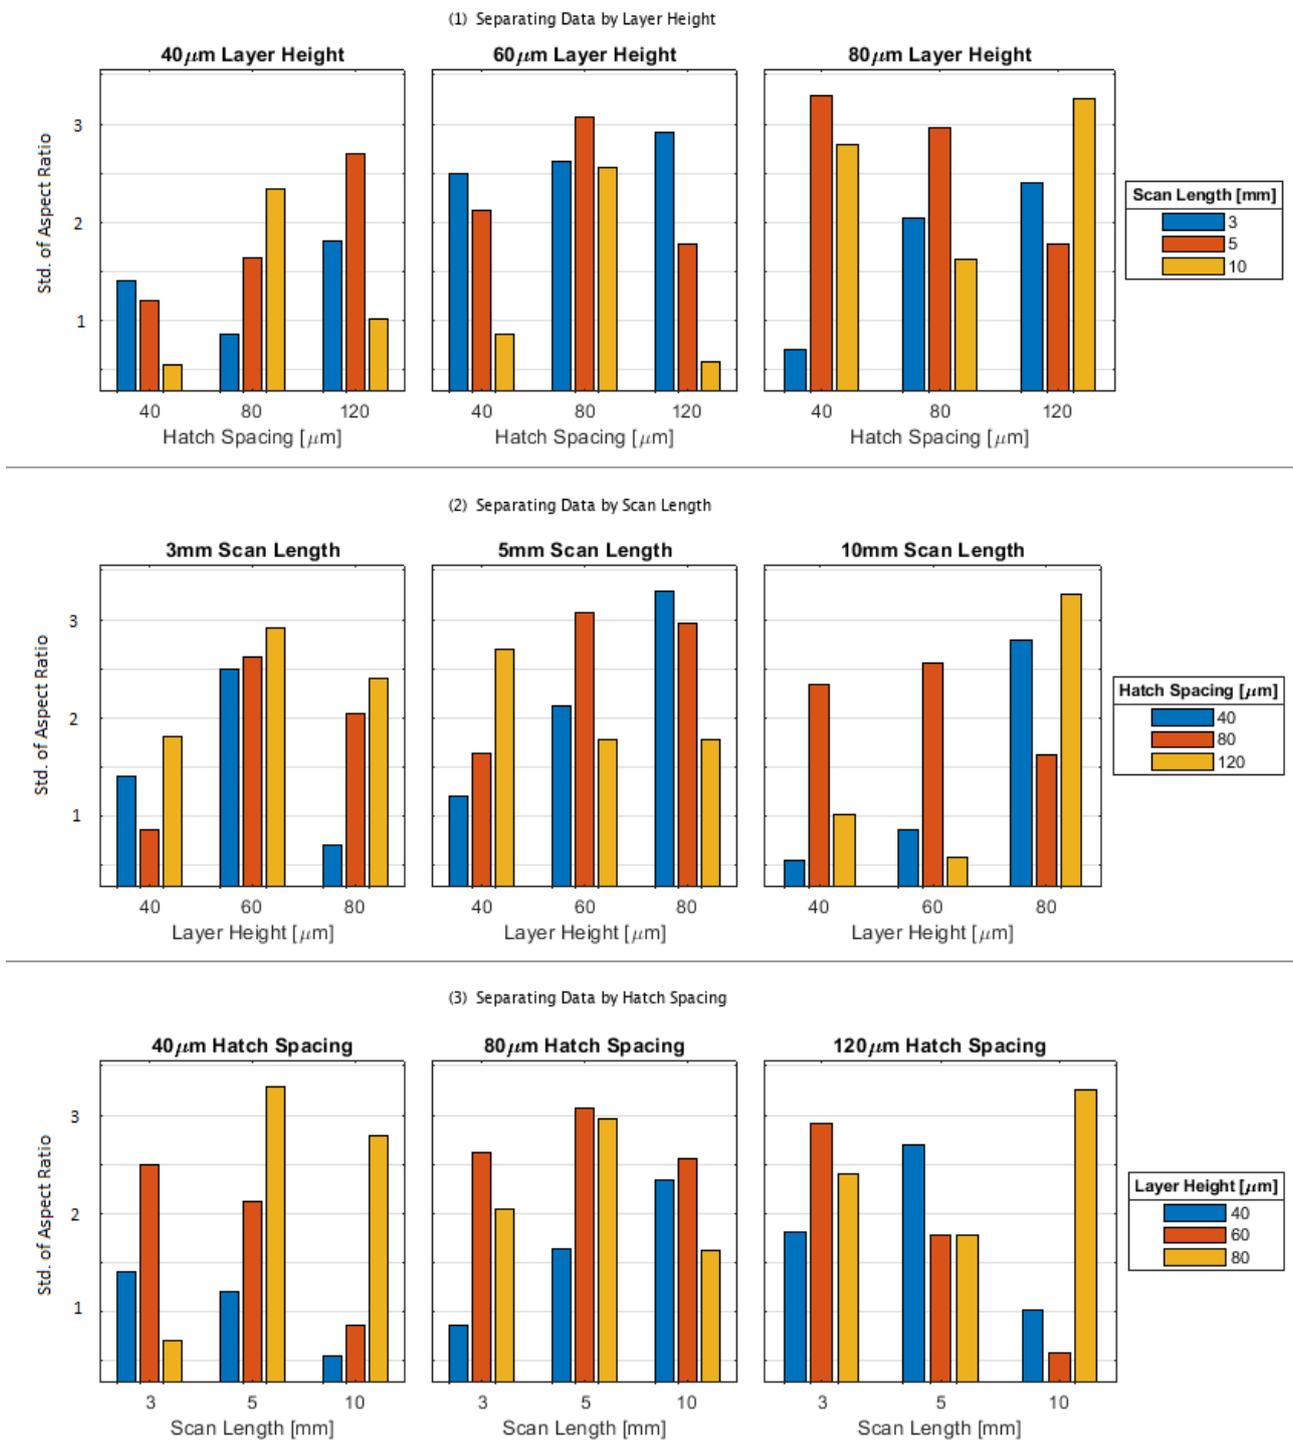


Figure 9.12

The effects of machine parameters on the standard deviation of cell aspect ratio

Effect of Thermal Gradient on Std. Dev. of Cell Aspect Ratio

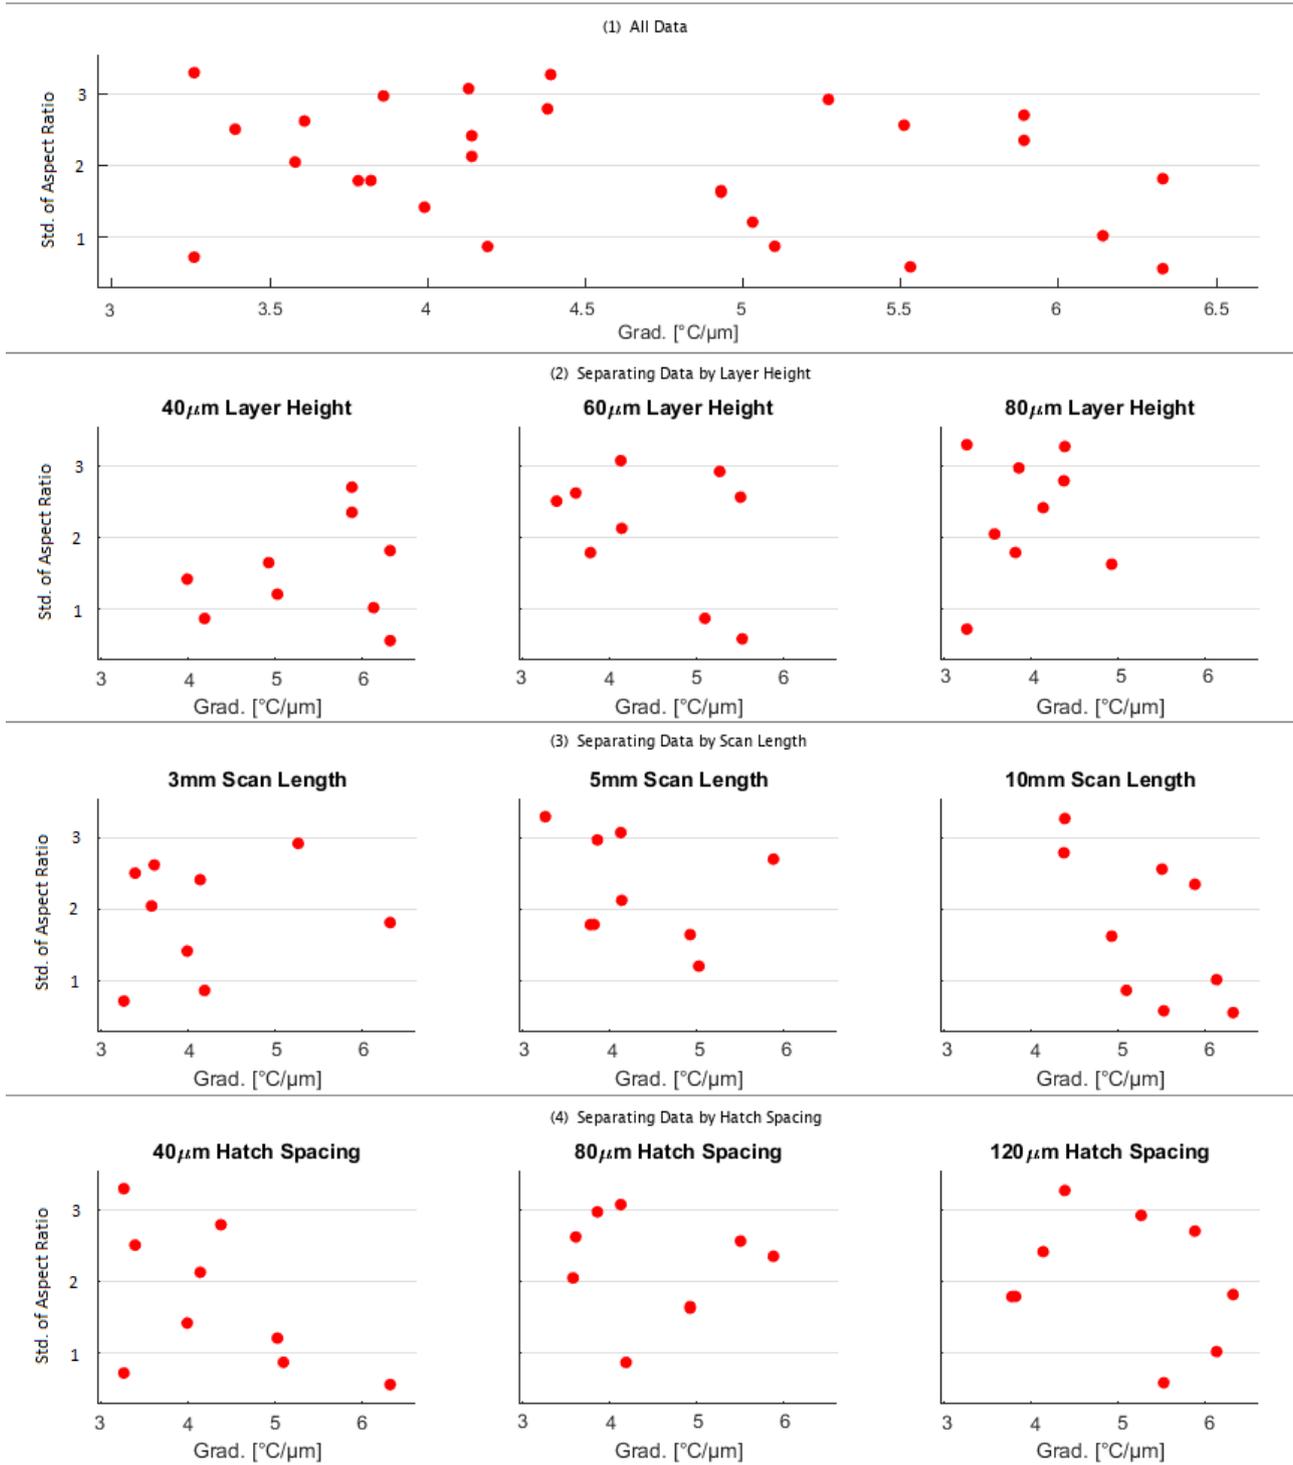


Figure 9.13

The effects of thermal gradient on the standard deviation of cell aspect ratio

9.2.4 - EBSD Data

9.2.4.1 - Grain Aspect Ratio

The area-weighted mean aspect ratios of the ellipse that Aztec fitted to each grain, are shown in Figure 9.14.

As can be seen, no clear trends exist. It was expected that as scan length, hatch spacing, and layer height increased, grain aspect ratio would increase too, as the region of solidification shifted upwards on the GR diagram, from equiaxed towards columnar or cellular. The fact that this trend, or any other trend, is not observed, suggests that attempting to predict microstructural morphology based on machine settings is not reliable.

The area-weighted means of the grain aspect ratios were plotted against the thermal gradients, as shown in Figure 9.15.

As can be seen, no clear trends exist. This finding provides further evidence that simply analysing thermal gradient (or at least thermal gradient as observable from thermal signatures) is not an accurate method of predicting grain morphology.

9.2.4.2 - Grain Average Misorientation (GAM)

The intra-grain misorientation was calculated using the GAM method, whereby the misorientation between each pixel and its adjacent neighbours was calculated, to give a mean misorientation for each grain. The area-weighted mean of these means was calculated for all grains in each sample, as shown in Figure 9.16.

As can be seen, no clear trends exist, when GAM is plotted against any machine parameter. It was expected that, as scan length, hatch spacing, and layer height increased, and therefore thermal gradient increased, the resulting more rapid solidification would cause grains to form with more internal defects, due to insufficient atomic diffusion time, which would in turn lead to greater GAM. The fact that no trends are seen to exist, suggests that attempting to predict the sufficiency of atomic diffusion, based on machine parameters, may not be a useful endeavour.

Area-weighted mean GAM was plotted against thermal gradient, as shown in Figure 9.17.

As can be seen, no clear trends exist. These findings further demonstrate the inability of thermal gradient measurements from the thermal camera to predict solidification mechanisms.

9.2.4.3 - Grain Orientation

The angle of orientation, of the longest axis of the ellipse fitted to each grain, was found. Since each sample had been oriented inside the EBSD device with the build direction in the Y direction, most grains had a fitted-ellipse angle close to 90°. Therefore, the fact that angles close to 0° or 180° over-

contribute to the standard deviation of the angle measurements, did not impact the statistics significantly. The area-weighted mean of these orientation angles is shown in Figure 9.18.

Effect of Machine Parameters on Grain Aspect Ratio

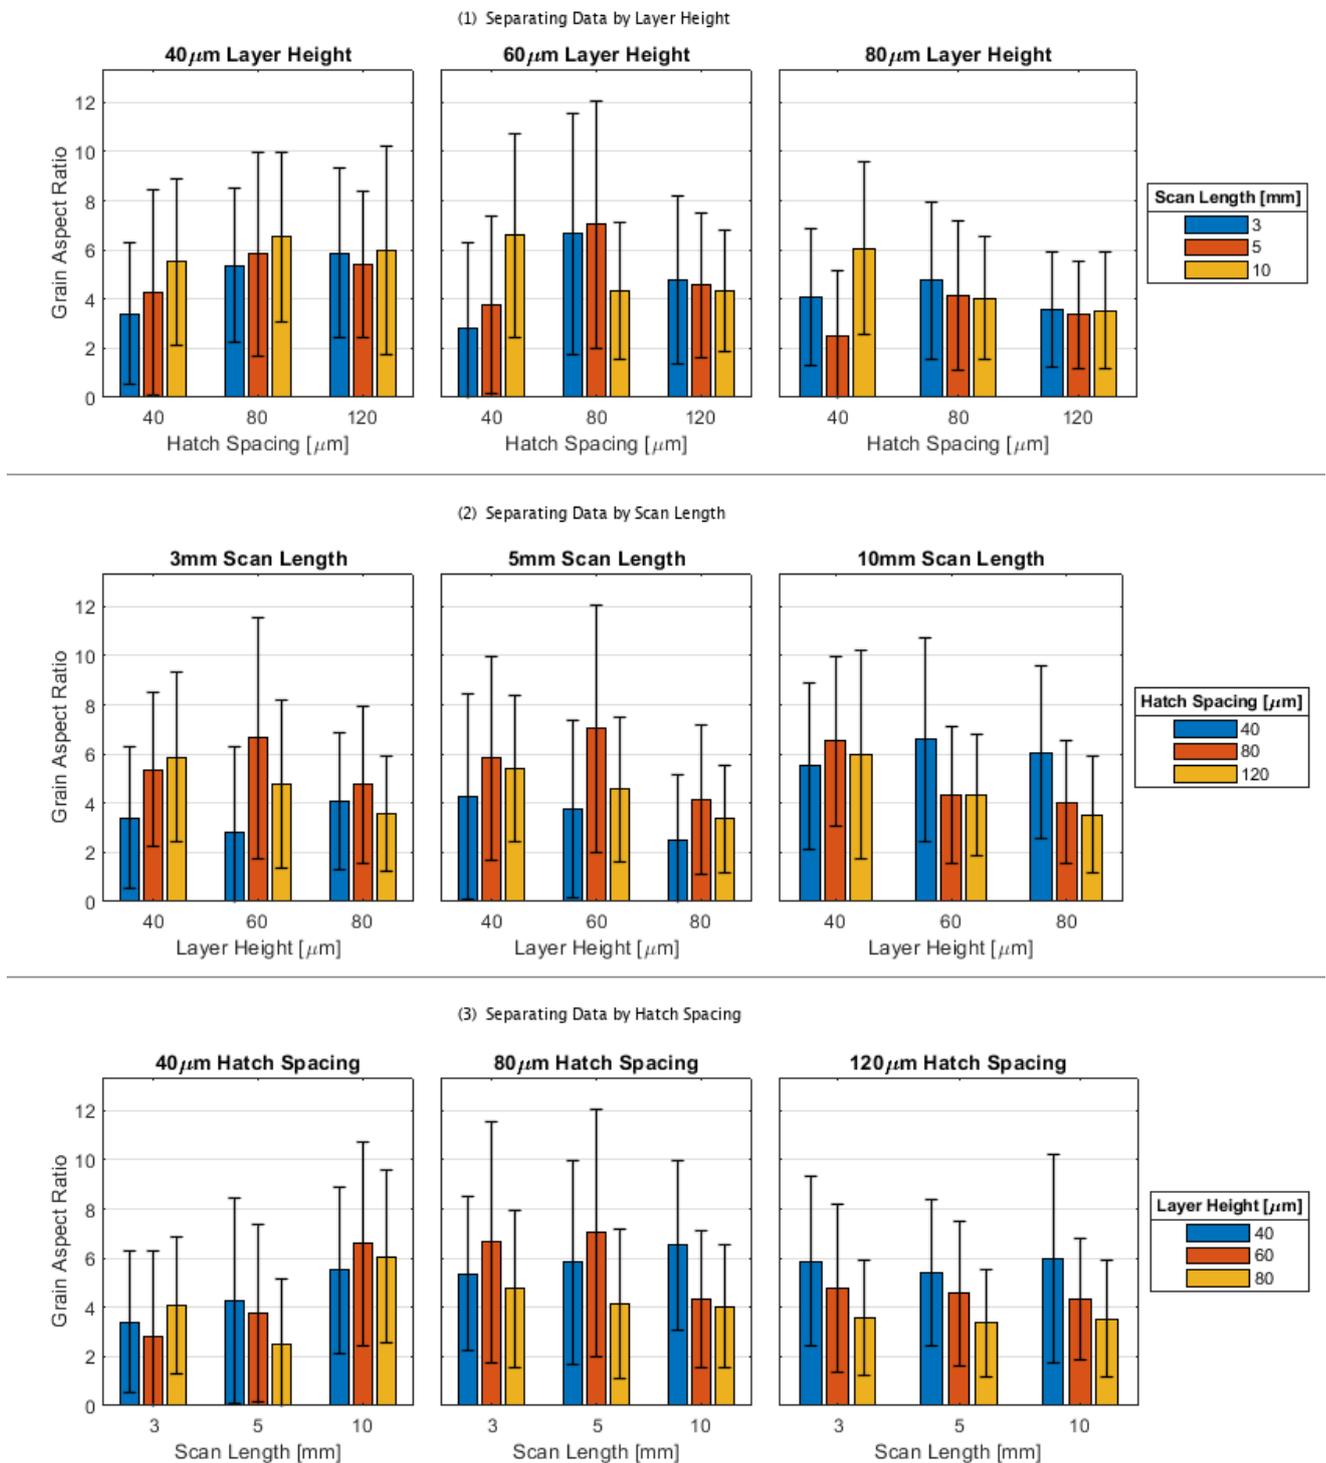


Figure 9.14

The effects of machine parameters on grain aspect ratio

Effect of Thermal Gradient on Grain Aspect Ratio

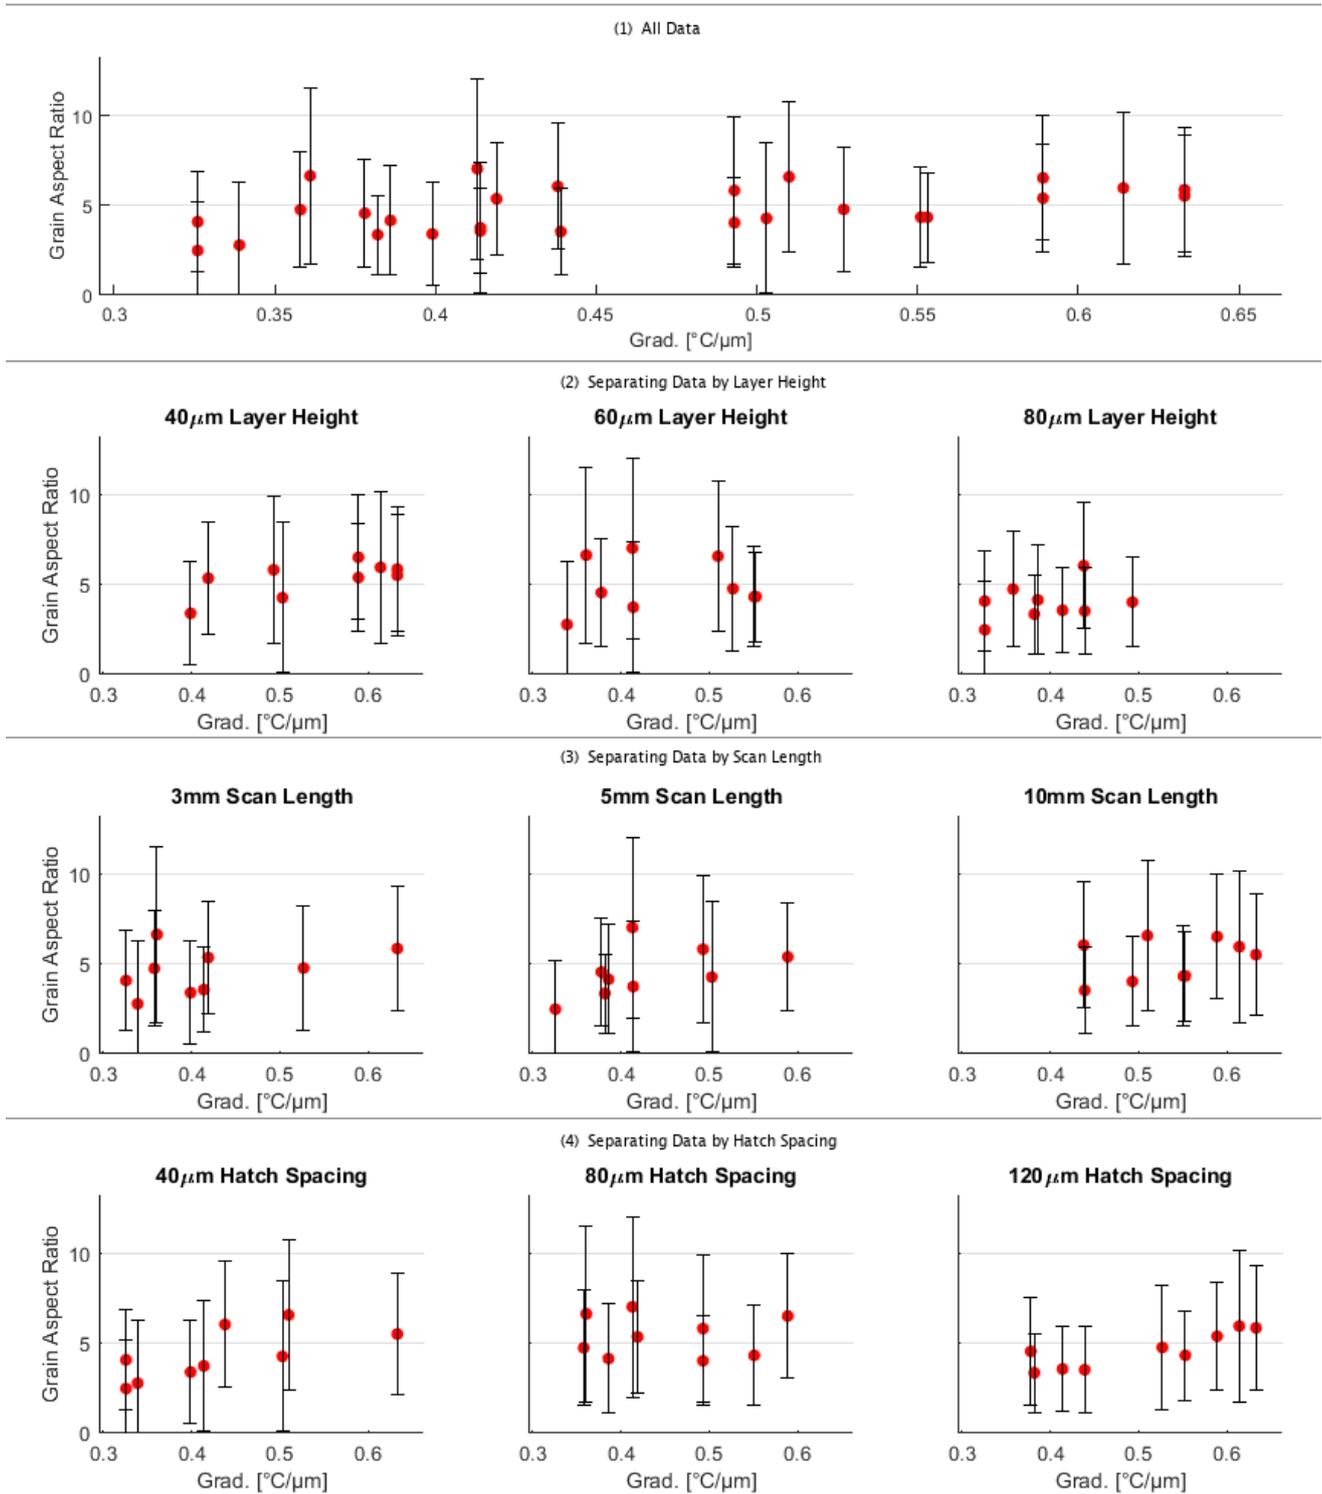


Figure 9.15

The effects of thermal gradient on grain aspect ratio

Effect of Machine Parameters on Mean GAM

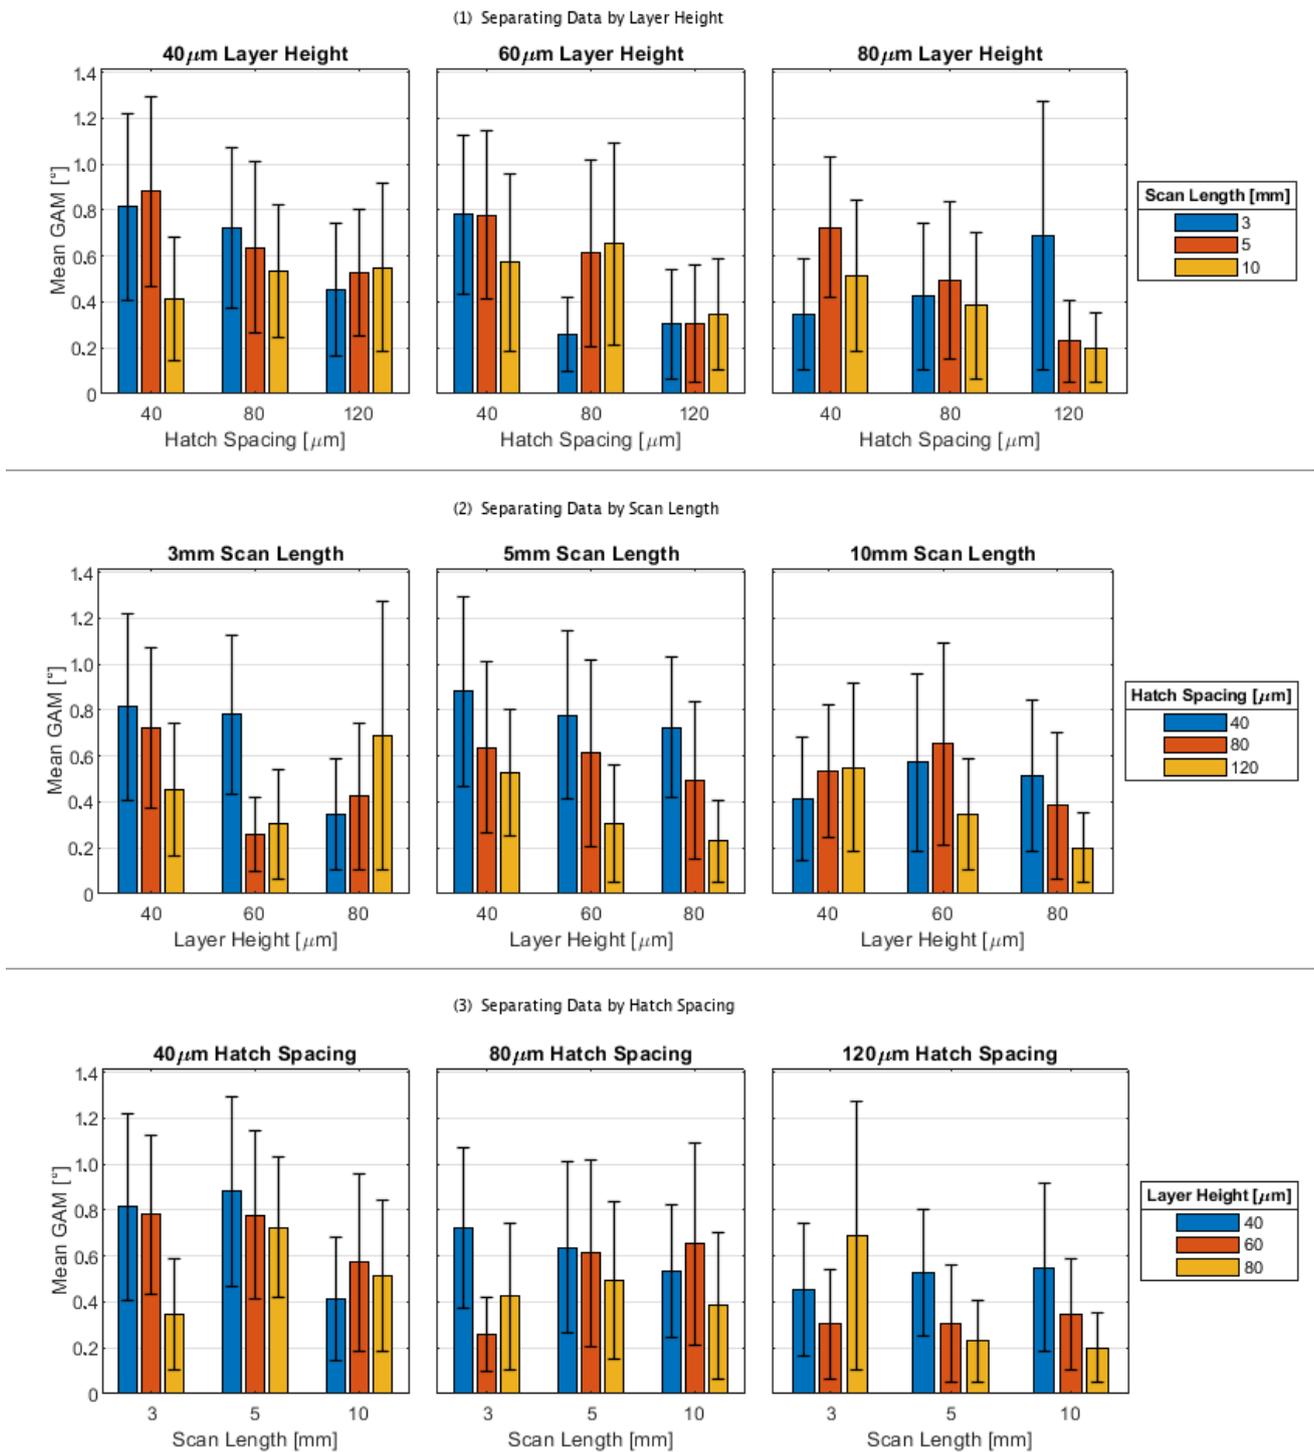


Figure 9.16

The effects of machine parameters on mean GAM

Effect of Thermal Gradient on Mean GAM

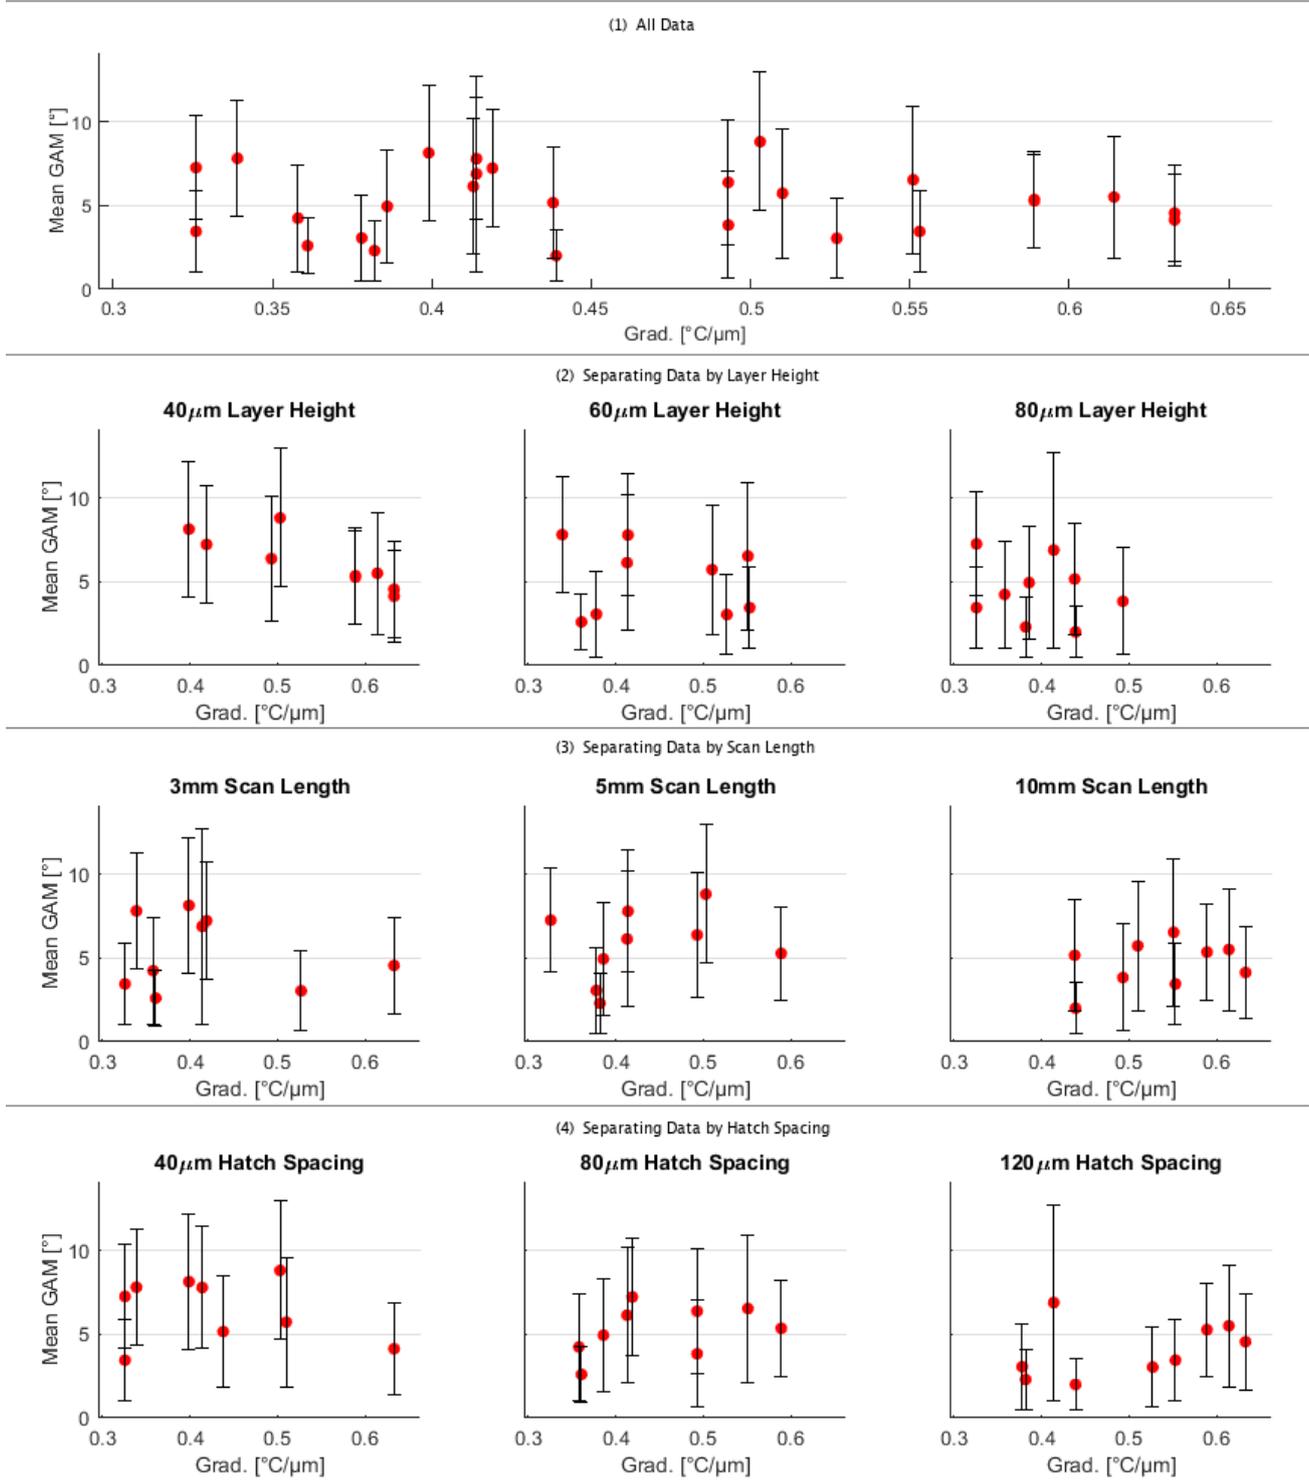


Figure 9.17

The effects of thermal gradient on mean GAM

Effect of Machine Parameters on Grain Orientation

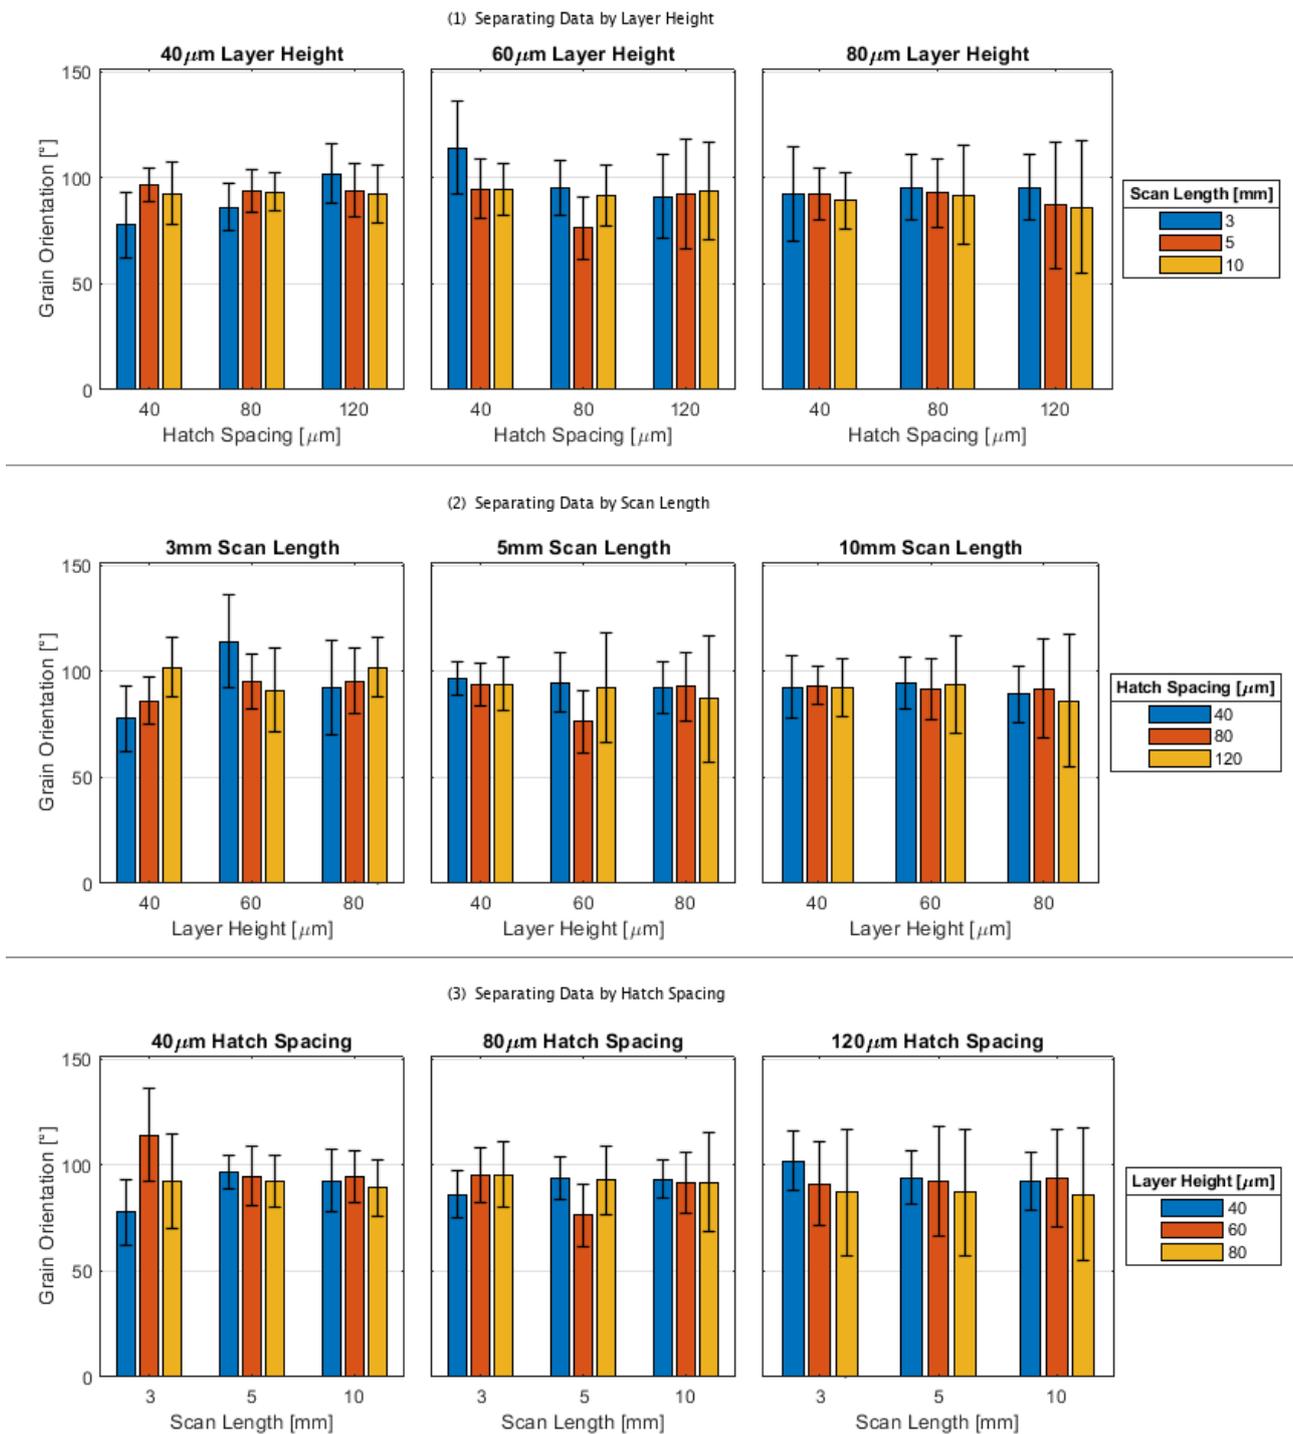


Figure 9.18

The effects of machine parameters on grain orientation

As can be seen, all grains in every sample were aligned strongly to the build direction, with little significant deviation. This result proves that the samples were highly anisotropic, with the major direction of heat flow in each melt pool being downward, and the major direction of solidification being upwards. It appears that no alternation of machine parameters, within the ranges tested in this experiment, could change these fundamentals.

9.3 - In-Situ Monitoring and the Mesostructure

The following charts and discussion were originally included in Section 6.2 - In-Situ Monitoring and Deliberate Porosity, but have been moved due to their lacking in significant insight.

9.3.1 - Melt Pool Mean Temperature Above the LoF Channels

The effect of LoF channels on melt pool mean temperature are shown in Figure 9.20.

There is an extremely strong similarity between the results for mean temperature, and the previously-discussed results for peak temperature. The same contrary-to-expected findings, and the same possible reasons for these findings, are again relevant here.

9.3.2 - Melt Pool Mean Temperature Above the Keyhole Line

The effect of parameters on melt pool mean temperature are shown in Figure 9.21.

As can be seen in row (1), the response of melt pool mean temperature to the triangle section width, is highly inconsistent. At some power and line width combinations, mean temperature increases as triangle width increases. At other combinations, it decreases.

Interestingly, the results from row (2) do not show the expected and trivial trends that appeared in the previous sets of charts. It would be expected that as laser power decreased, mean temperature would decrease too. The fact that this trend does not consistently exist across various line-width and triangle-width combinations, suggests that the large periphery of a melt pool, which exists at close to the liquidus temperature, dilutes its overall mean temperature to the extent that the metric may not be very meaningful.

As can be seen in row (3), no coherent trends exist between the width of the line, and the mean temperature of the melt pool. This result again suggests that the mean temperature may not be a useful characteristic in melt pool analysis.

Effect of LoF Channel on Melt Pool Mean Temps

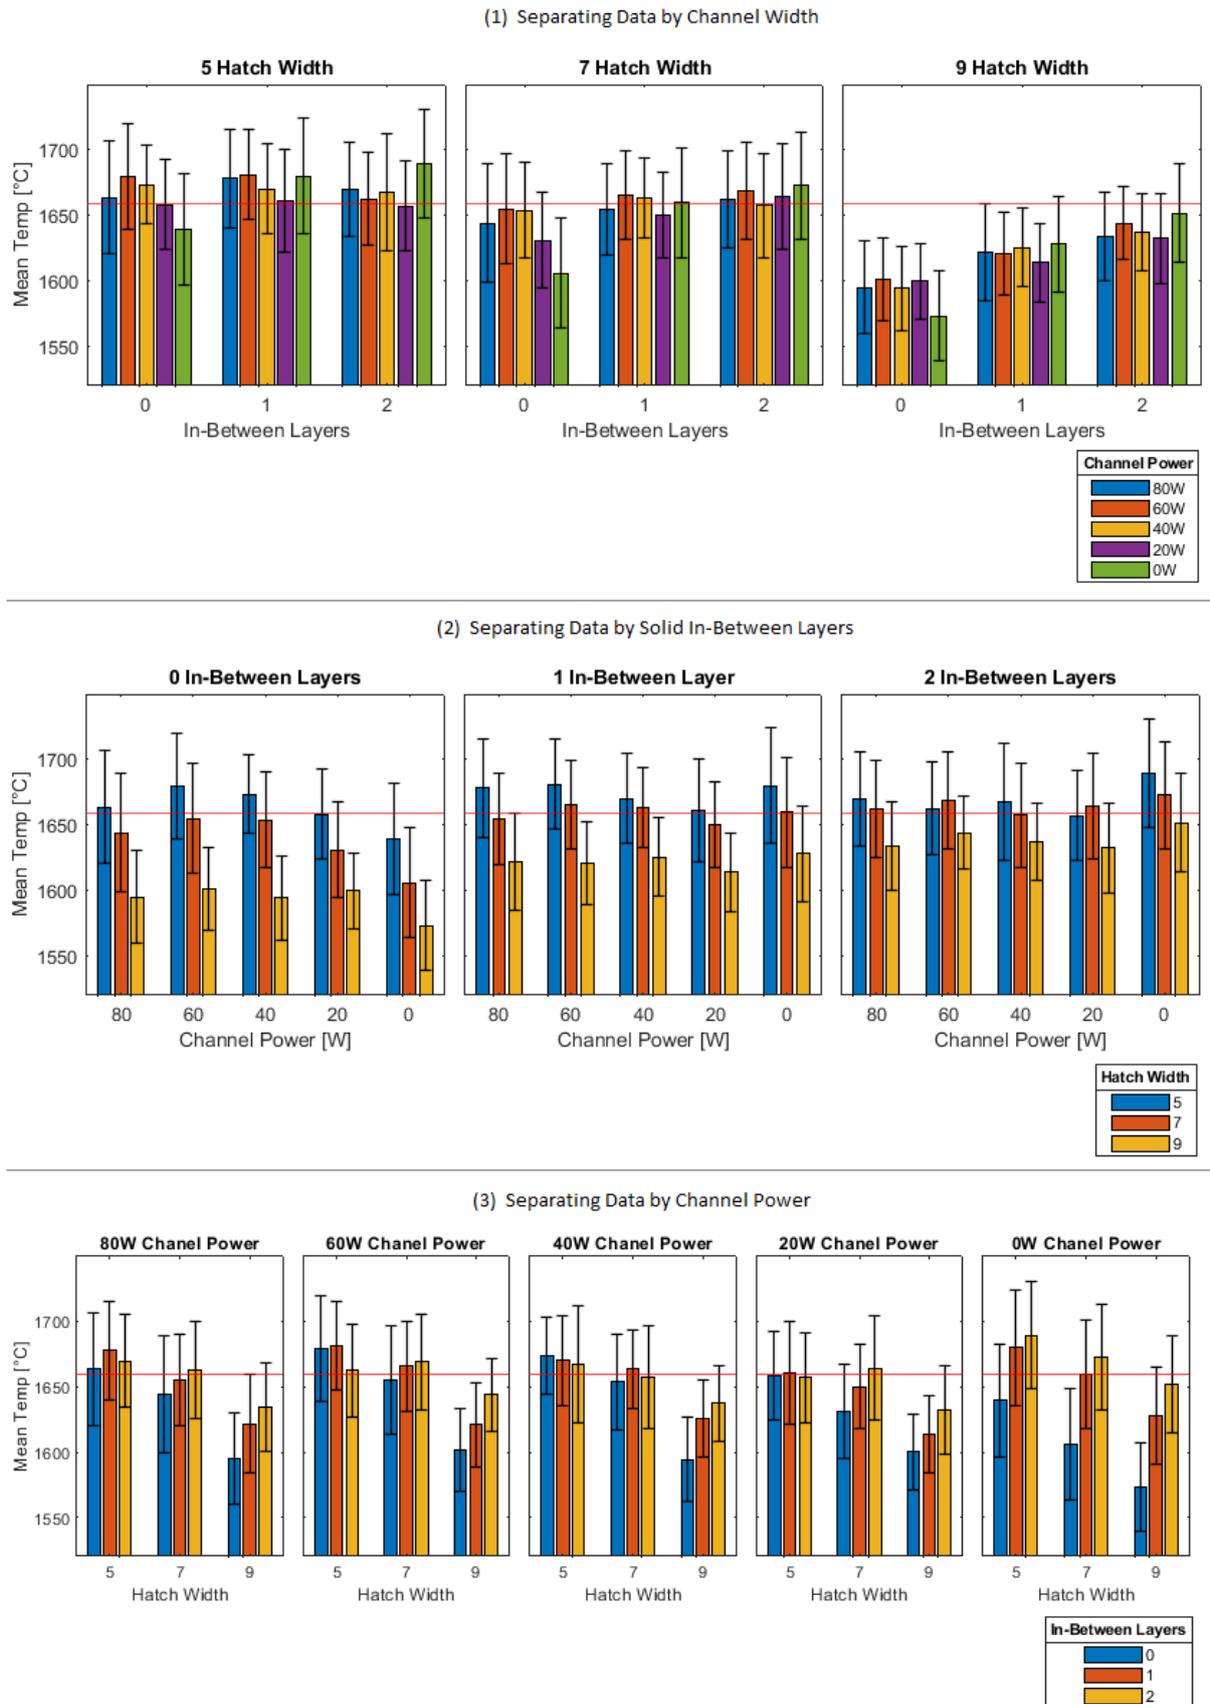


Figure 9.20

The effect of LoF channels on melt pool mean temperature

Effect of Parameters on Melt Pool Mean Temperature

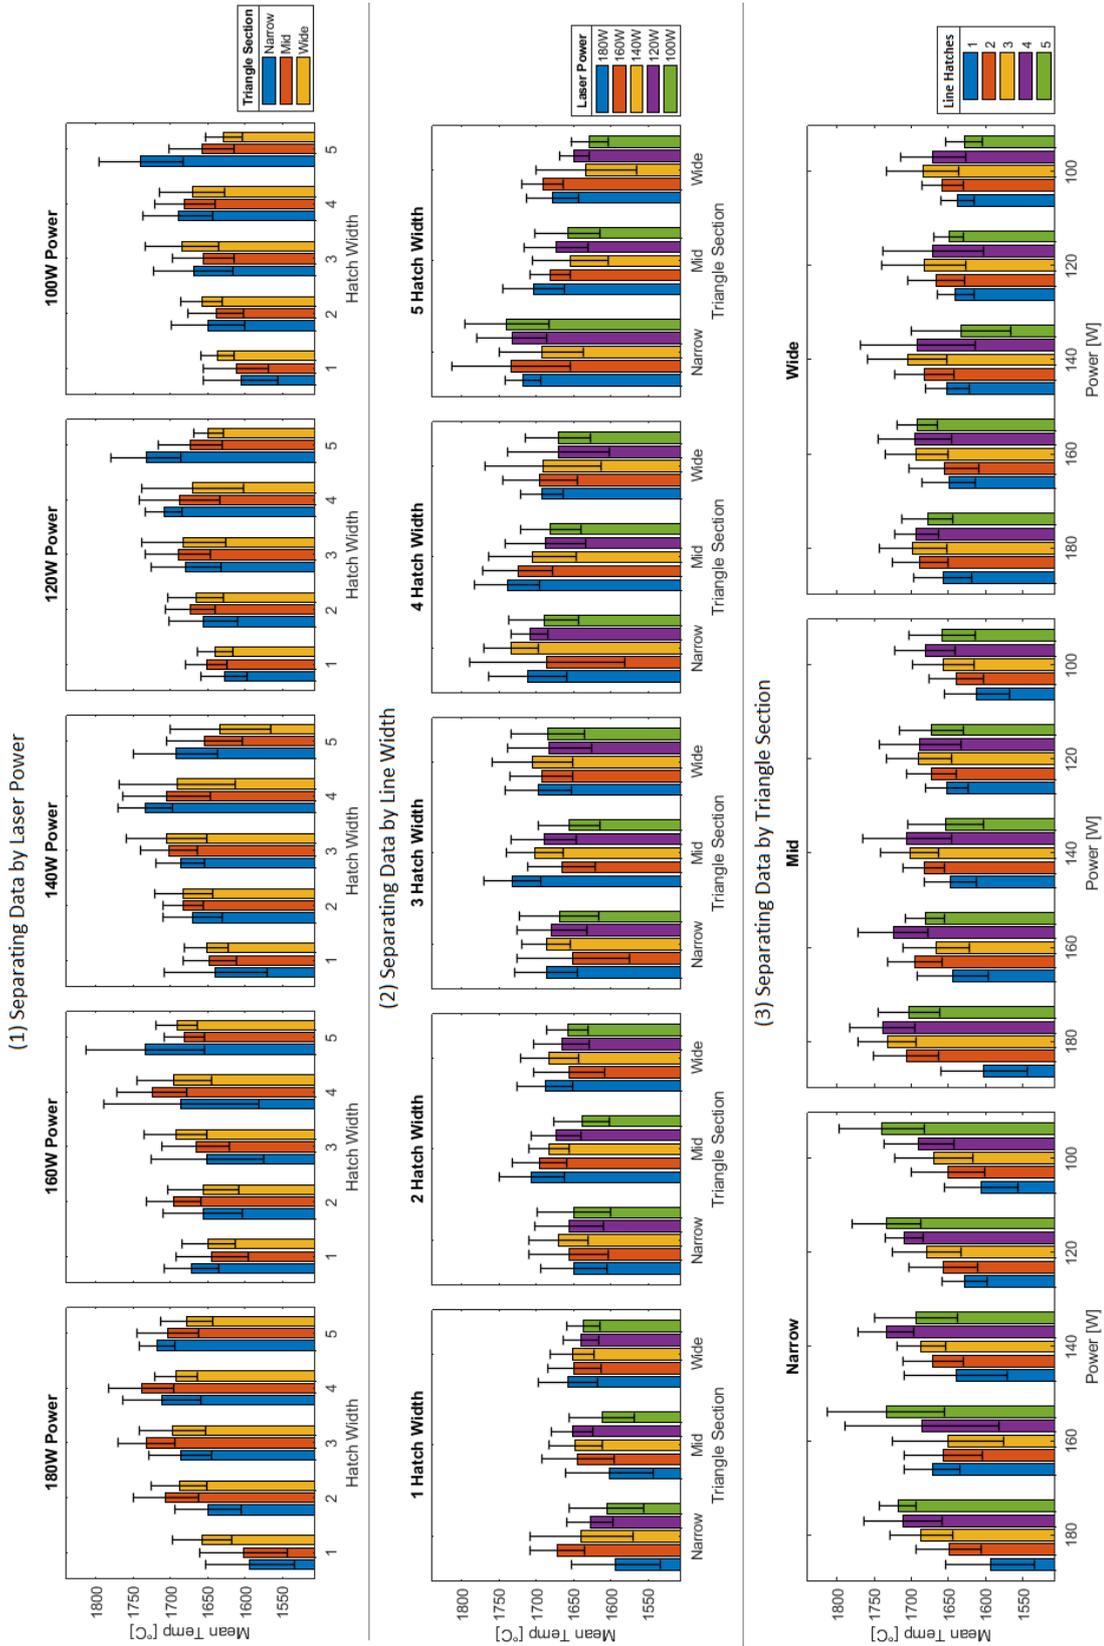


Figure 9.21

The effect of parameters on melt pool mean temperature

10 - Bibliography

1. Choi, J.Y., et al., *Accelerating phase-field simulation of three-dimensional microstructure evolution in laser powder bed fusion with composable machine learning predictions*. Additive Manufacturing, 2024. **79**: p. 103938.
2. Rodgers, T.M., et al., *Simulation of powder bed metal additive manufacturing microstructures with coupled finite difference-Monte Carlo method*. Additive Manufacturing, 2021. **41**: p. 101953.
3. Węgrzyn, N., *The use of additive manufacturing for production of commercial airplane power plants components: A review*. Safety & Defense, 2022. **2**.
4. Jiang, C.-P., et al., *Enhancing internal cooling channel design in Inconel 718 turbine blades via laser powder bed fusion: a comprehensive review of surface topography enhancements*. International Journal of Precision Engineering and Manufacturing, 2024: p. 1-25.
5. Yadroitsava, I., A. Du Plessis, and I. Yadroitsev, *Bone regeneration on implants of titanium alloys produced by laser powder bed fusion: A review*. Titanium for Consumer Applications, 2019: p. 197-233.
6. Equbal, A., A.K. Sood, and M. Shamim, *Rapid tooling: A major shift in tooling practice*. Manufacturing and Industrial Engineering, 2015. **14**(3-4).
7. Mazur, M. and P. Selvakannan, *Laser powder bed fusion—principles, challenges, and opportunities*. Additive manufacturing for chemical sciences and engineering, 2022: p. 77-108.
8. Kvaratskheliya, A., et al., *Making parts on Mars: Laser processing of iron contaminated by regolith simulant*. Materials Today Advances, 2025. **25**: p. 100566.
9. Chowdhury, S., et al., *Laser powder bed fusion: a state-of-the-art review of the technology, materials, properties & defects, and numerical modelling*. Journal of Materials Research and Technology, 2022. **20**: p. 2109-2172.
10. Gibbons, D.W., J.-P.L. Serfontein, and A.F. Van der Merwe, *Mapping the path to certification of metal laser powder bed fusion for aerospace applications*. Rapid prototyping journal, 2021. **27**(2): p. 355-361.
11. Papatheodorou, A., et al., *Recent Advances in Sensor Fusion Monitoring and Control Strategies in Laser Powder Bed Fusion: A Review*. Machines, 2025. **13**(9): p. 820.
12. ISO, *Additive manufacturing*, in *General principles*. 2021.
13. Jiménez, M., et al., *Additive Manufacturing Technologies: An Overview about 3D Printing Methods and Future Prospects*. Complexity, 2019. **2019**(1): p. 9656938.
14. Wong, K.V. and A. Hernandez, *A Review of Additive Manufacturing*. International Scholarly Research Notices, 2012. **2012**(1): p. 208760.
15. Bourell, D.L., *Perspectives on Additive Manufacturing*. Annual Review of Materials Research, 2016. **46**(Volume 46, 2016): p. 1-18.
16. Pérez, M., et al., *Current advances in additive manufacturing*. Procedia Cirp, 2020. **88**: p. 439-444.
17. Klahn, C., B. Leutenecker, and M. Meboldt, *Design strategies for the process of additive manufacturing*. Procedia Cirp, 2015. **36**: p. 230-235.
18. Ribeiro, I., et al., *Framework for Life Cycle Sustainability Assessment of Additive Manufacturing*. Sustainability, 2020. **12**: p. 929.
19. Wakjira, Y., A. Akessa, and H. Lemu, *Overview study on challenges of additive manufacturing for a healthcare application*. IOP Conference Series: Materials Science and Engineering, 2021. **1201**: p. 012041.
20. Singh, R., et al., *Powder bed fusion process in additive manufacturing: An overview*. Materials Today: Proceedings, 2020. **26**: p. 3058-3070.
21. Dev Singh, D., T. Mahender, and A. Raji Reddy, *Powder bed fusion process: A brief review*. Materials Today: Proceedings, 2021. **46**: p. 350-355.
22. Leung, C.L.A., et al., *Laser-matter interactions in additive manufacturing of stainless steel SS316L and 13-93 bioactive glass revealed by in situ X-ray imaging*. Additive Manufacturing, 2018. **24**: p. 647-657.
23. Lee, H., et al., *Lasers in additive manufacturing: A review*. International Journal of Precision Engineering and Manufacturing-Green Technology, 2017. **4**: p. 307-322.

24. Assuncao, E. and S. Williams, *Comparison of continuous wave and pulsed wave laser welding effects*. Optics and Lasers in Engineering, 2013. **51**(6): p. 674-680.
25. Horn, M., et al., *Laser powder bed fusion recoater selection guide—Comparison of resulting powder bed properties and part quality*. Powder Technology, 2024. **434**: p. 119356.
26. Khorasani, M., et al., *Laser subtractive and laser powder bed fusion of metals: review of process and production features*. Rapid Prototyping Journal, 2023. **29**(5): p. 935-958.
27. Khorasani, A., et al., *A review of technological improvements in laser-based powder bed fusion of metal printers*. The International Journal of Advanced Manufacturing Technology, 2020. **108**(1): p. 191-209.
28. Jackson, M. and R.R. Boyer, *Titanium and its alloys: processing, fabrication and mechanical performance*. Encyclopedia of Aerospace Engineering, 2010.
29. Marques, A., et al., *Inconel 718—copper parts fabricated by 3D multi-material laser powder bed fusion: a novel technological and designing approach for rocket engine*. The International Journal of Advanced Manufacturing Technology, 2022. **122**(3): p. 2113-2123.
30. Wildgoose, A.J., et al., *Impact of Additive Manufacturing on Internal Cooling Channels With Varying Diameters and Build Directions*. Journal of Turbomachinery, 2021. **143**(7).
31. Tan, C., et al., *Design and additive manufacturing of novel conformal cooling molds*. Materials & Design, 2020. **196**: p. 109147.
32. Wang, Y., et al., *The process planning for additive and subtractive hybrid manufacturing of powder bed fusion (PBF) process*. Materials & Design, 2023. **227**: p. 111732.
33. Sing, S.L. and W.Y. Yeong, *Laser powder bed fusion for metal additive manufacturing: perspectives on recent developments*. Virtual and Physical Prototyping, 2020. **15**(3): p. 359-370.
34. Wohler, *Wohler Report 2024*. 2024.
35. Awad, A., et al., *Advances in powder bed fusion 3D printing in drug delivery and healthcare*. Advanced Drug Delivery Reviews, 2021. **174**: p. 406-424.
36. Childerhouse, T. and M. Jackson, *Near Net Shape Manufacture of Titanium Alloy Components from Powder and Wire: A Review of State-of-the-Art Process Routes*. Metals, 2019. **9**: p. 689.
37. Priarone, P.C., et al., *Laser powder bed fusion (L-PBF) additive manufacturing: On the correlation between design choices and process sustainability*. Procedia CIRP, 2018. **78**: p. 85-90.
38. Gruber, K., et al. *Evaluation of Inconel 718 Metallic Powder to Optimize the Reuse of Powder and to Improve the Performance and Sustainability of the Laser Powder Bed Fusion (LPBF) Process*. Materials, 2021. **14**, DOI: 10.3390/ma14061538.
39. Heiden, M.J., et al., *Evolution of 316L stainless steel feedstock due to laser powder bed fusion process*. Additive Manufacturing, 2019. **25**: p. 84-103.
40. Tsai, C.-Y., et al., *Synchronized multi-spot scanning strategies for the laser powder bed fusion process*. Additive Manufacturing, 2019. **27**: p. 1-7.
41. Wong, H., et al., *Multi-Laser Powder Bed Fusion Benchmarking—Initial Trials with Inconel 625*. The International Journal of Advanced Manufacturing Technology, 2019. **105**(7): p. 2891-2906.
42. Aconity3D, *AconitySYSTEMS LASER POWDER BED FUSION*. 2024.
43. Sing, S.L., et al., *Perspectives of using machine learning in laser powder bed fusion for metal additive manufacturing*. Virtual and Physical Prototyping, 2021. **16**(3): p. 372-386.
44. Chen, M., et al., *A quantitative study of thermal cycling along the build direction of Ti-6Al-4V produced by laser powder bed fusion*. Materials & Design, 2023. **225**: p. 111458.
45. Sabzi, H.E., et al., *Grain refinement in laser powder bed fusion: The influence of dynamic recrystallization and recovery*. Materials & Design, 2020. **196**: p. 109181.
46. Dejene, N. and H. Lemu, *Current Status and Challenges of Powder Bed Fusion-Based Metal Additive Manufacturing: Literature Review*. Metals, 2023. **13**.
47. Dowling, L., et al., *A review of critical repeatability and reproducibility issues in powder bed fusion*. Materials & Design, 2020. **186**: p. 108346.
48. Slotwinski, J.A., et al., *Interlaboratory mechanical property study for Cobalt-Chromium alloy Made by Laser Powder-Bed-Fusion Additive Manufacturing*. 2018.
49. Mishchenko, M.I., L.D. Travis, and A.A. Lacia, *Scattering, absorption, and emission of light by small particles*. 2002: Cambridge university press.

50. Allmen, M.V. and A. Blatter, *Laser-beam interactions with materials: physical principles and applications*. Vol. 2. 2013: Springer Science & Business Media.
51. Santecchia, E., S. Spigarelli, and M. Cabibbo *Material Reuse in Laser Powder Bed Fusion: Side Effects of the Laser—Metal Powder Interaction*. *Metals*, 2020. **10**, DOI: 10.3390/met10030341.
52. Comins, N.R., *The optical properties of liquid metals*. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 1972. **25**(4): p. 817-831.
53. Ye, J., et al., *Energy Coupling Mechanisms and Scaling Behavior Associated with Laser Powder Bed Fusion Additive Manufacturing*. *Advanced Engineering Materials*, 2019. **21**(7): p. 1900185.
54. Trapp, J., et al., *In situ absorptivity measurements of metallic powders during laser powder-bed fusion additive manufacturing*. *Applied Materials Today*, 2017. **9**: p. 341-349.
55. Patel, S. and M. Vlasea, *Melting modes in laser powder bed fusion*. *Materialia*, 2020. **9**: p. 100591.
56. Ahmed, N., S. Darwish, and A.M. Alahmari, *Laser Ablation and Laser-Hybrid Ablation Processes: A Review*. *Materials and Manufacturing Processes*, 2016. **31**(9): p. 1121-1142.
57. Xi, C. and W. Hai-Xing, *A calculation model for the evaporation recoil pressure in laser material processing*. *Journal of Physics D: Applied Physics*, 2001. **34**(17): p. 2637.
58. Lee, D. and S. Jeong, *Analysis of recoil force during Nd: YAG laser ablation of silicon*. *Applied Physics A*, 2004. **79**: p. 1341-1344.
59. Han, L. and F.W. Liou, *Numerical investigation of the influence of laser beam mode on melt pool*. *International Journal of Heat and Mass Transfer*, 2004. **47**(19): p. 4385-4402.
60. Wang, T., et al., *Relationships between the characteristics of porosity, melt pool and process parameters in laser powder bed fusion AlZn alloy*. *Journal of Manufacturing Processes*, 2021. **68**: p. 1236-1244.
61. Mayi, Y.A., et al., *Physical mechanisms of conduction-to-keyhole transition in laser welding and additive manufacturing processes*. *Optics & Laser Technology*, 2023. **158**: p. 108811.
62. Balbaa, M., et al., *On selective laser melting of Inconel 718: Densification, surface roughness, and residual stresses*. *Materials & Design*, 2020. **193**: p. 108818.
63. Xianfeng, X., et al., *Progress on Experimental Study of Melt Pool Flow Dynamics in Laser Material Processing*, in *Liquid Metals*, C. Samson Jerold Samuel, S. Gnanasekaran, and M. Suresh, Editors. 2021, IntechOpen: Rijeka. p. Ch. 3.
64. Le, T.-N. and Y.-L. Lo, *Effects of sulfur concentration and Marangoni convection on melt-pool formation in transition mode of selective laser melting process*. *Materials & Design*, 2019. **179**: p. 107866.
65. Fuhrich, T., P. Berger, and H. Hügel, *Marangoni effect in laser deep penetration welding of steel*. *Journal of laser applications*, 2001. **13**(5): p. 178-186.
66. Ahmadi, E. and A.R. Ebrahimi, *Welding of 316L Austenitic Stainless Steel with Activated Tungsten Inert Gas Process*. *Journal of Materials Engineering and Performance*, 2015. **24**(2): p. 1065-1071.
67. Yang, J., et al., *Experimental quantification of inward Marangoni convection and its impact on keyhole threshold in laser powder bed fusion of stainless steel*. *Additive Manufacturing*, 2024. **84**: p. 104092.
68. Sun, S., M. Brandt, and M. Easton, *Powder bed fusion processes: An overview*. *Laser additive manufacturing*, 2017: p. 55-77.
69. Guo, Q., et al., *In-situ characterization and quantification of melt pool variation under constant input energy density in laser powder bed fusion additive manufacturing process*. *Additive Manufacturing*, 2019. **28**: p. 600-609.
70. Lane, B., et al., *Thermographic measurements of the commercial laser powder bed fusion process at NIST*. *Rapid Prototyping Journal*, 2016. **22**(5): p. 778-787.
71. Hooper, P.A., *Melt pool temperature and cooling rates in laser powder bed fusion*. *Additive Manufacturing*, 2018. **22**: p. 548-559.
72. Furumoto, T., et al., *Evaluating the thermal characteristics of laser powder bed fusion*. *Journal of Materials Processing Technology*, 2022. **299**: p. 117384.
73. Vecchiato, F.L., et al., *Melt pool microstructure and morphology from single exposures in laser powder bed fusion of 316L stainless steel*. *Additive Manufacturing*, 2020. **36**: p. 101401.
74. Criales, L.E., et al., *Laser powder bed fusion of nickel alloy 625: Experimental investigations of effects of process parameters on melt pool size and shape with spatter analysis*. *International Journal of Machine Tools and Manufacture*, 2017. **121**: p. 22-36.

75. Scime, L. and J. Beuth, *Melt pool geometry and morphology variability for the Inconel 718 alloy in a laser powder bed fusion additive manufacturing process*. Additive Manufacturing, 2019. **29**: p. 100830.
76. Yamagata, H., et al., *The effect of average cooling rates on the microstructure of the Al–20% Si high pressure die casting alloy used for monolithic cylinder blocks*. Journal of Materials Processing Technology, 2008. **203**(1): p. 333-341.
77. Rasouli, D., et al., *Effect of cooling rate on the microstructure and mechanical properties of microalloyed forging steel*. Journal of Materials Processing Technology, 2008. **206**(1): p. 92-98.
78. Blecher, J.J., T.A. Palmer, and T. DebRoy, *Solidification Map of a Nickel-Base Alloy*. Metallurgical and Materials Transactions A, 2014. **45**(4): p. 2142-2151.
79. Liu, H., *Numerical analysis of thermal stress and deformation in multi-layer laser metal deposition process*. 2014: Missouri University of Science and Technology.
80. Scipioni Bertoli, U., et al., *In-situ characterization of laser-powder interaction and cooling rates through high-speed imaging of powder bed fusion additive manufacturing*. Materials & Design, 2017. **135**: p. 385-396.
81. Averardi, A., et al., *Effect of particle size distribution on the packing of powder beds: A critical discussion relevant to additive manufacturing*. Materials Today Communications, 2020. **24**: p. 100964.
82. Agazhanov, A.S., D.A. Samoshkin, and Y.M. Kozlovskii, *Thermophysical properties of Inconel 718 alloy*. Journal of Physics: Conference Series, 2019. **1382**(1): p. 012175.
83. Desai, P. and C. Ho, *Thermal linear expansion of nine selected AISI stainless steels*. Cindas report, 1978. **51**.
84. Gumbleton, R., et al., *Evaluating the coefficient of thermal expansion of additive manufactured AlSi10Mg using microwave techniques*. Additive Manufacturing, 2019. **30**: p. 100841.
85. Freeman, F., *Structuring Difference: The Additive Manufacture of Spatially & Functionally Differentiated Microstructures*. 2018, University of Sheffield.
86. Sarkar, D., A. Kapil, and A. Sharma, *Advances in computational modeling for laser powder bed fusion additive manufacturing: A comprehensive review of finite element techniques and strategies*. Additive Manufacturing, 2024. **85**: p. 104157.
87. Kaess, M., M. Werz, and S. Weihe, *Residual stress formation mechanisms in laser powder bed fusion—a numerical evaluation*. Materials, 2023. **16**(6): p. 2321.
88. Qiu, C., et al., *On the role of melt flow into the surface structure and porosity development during selective laser melting*. Acta Materialia, 2015. **96**: p. 72-79.
89. Chen, Y., et al., *Effects of process parameters on the microstructure of Inconel 718 during powder bed fusion based on cellular automata approach*. Virtual and Physical Prototyping, 2023. **18**(1): p. e2251032.
90. Chen, L.-Q., *Chemical potential and Gibbs free energy*. Mrs Bulletin, 2019. **44**(7): p. 520-523.
91. Greskovich, C. and J. Rosolowski, *Sintering of covalent solids*. Journal of the American Ceramic Society, 1976. **59**(7-8): p. 336-343.
92. Jacobs, P. and F. Tompkins, *Ionic conductance in solid salts*. Quarterly Reviews, Chemical Society, 1952. **6**(3): p. 238-261.
93. White, J.C. and E.R. Davidson, *An analysis of the hydrogen bond in ice*. The Journal of chemical physics, 1990. **93**(11): p. 8029-8035.
94. Wang, J., et al., *Crystalline behavior of paraffin wax*. The Journal of Physical Chemistry B, 2022. **126**(4): p. 985-995.
95. Lang, P.F., *Is a metal “ions in a sea of delocalized electrons?”*. Journal of Chemical Education, 2018. **95**(10): p. 1787-1793.
96. Englezos, P. and P. Bishnoi, *Gibbs free energy analysis for the supersaturation limits of methane in liquid water and the hydrate-gas-liquid water phase behavior*. Fluid Phase Equilibria, 1988. **42**: p. 129-140.
97. Kalup, A., et al., *Liquidus and solidus temperatures and latent heats of melting of steels*. Journal of thermal analysis and calorimetry, 2017. **127**(1): p. 123-128.
98. Yoshioka, H., Y. Tada, and Y. Hayashi, *Crystal growth and its morphology in the mushy zone*. Acta Materialia, 2004. **52**(6): p. 1515-1523.

99. Edler, F., et al., *Guide on secondary thermometry: specialised fixed points above 0°C*. 2017.
100. Liu, X., *Heterogeneous nucleation or homogeneous nucleation?* The Journal of Chemical Physics, 2000. **112**(22): p. 9949-9955.
101. Atiyah, A., *LECTURE 2 HOMOGENEOUS NUCLEATION*. 2016.
102. Sear, R.P., *Heterogeneous and homogeneous nucleation compared: Rapid nucleation on microscopic impurities*. The Journal of Physical Chemistry B, 2006. **110**(10): p. 4985-4989.
103. Perepezko, J.H. and M.J. Uttormark, *Undercooling and nucleation during solidification*. ISIJ international, 1995. **35**(6): p. 580-588.
104. Herlach, D., *UNDERCOOLING AND SOLIDIFICATION OF PURE METALS*. Annu. Rev. Mater. Sci, 1991. **23**: p. 44.
105. Nakae, H., *Influence of inoculation on solidification in cast iron*. International Journal of Cast Metals Research, 2008. **21**(1-4): p. 7-10.
106. Chalmers, B., *Principles of solidification*, in *Applied solid state physics*. 1964, Springer. p. 161-170.
107. Avazkonandeh-Gharavol, M.H., M. Haddad-Sabzevar, and H. Fredriksson, *Effect of partition coefficient on microsegregation during solidification of aluminium alloys*. International Journal of Minerals, Metallurgy, and Materials, 2014. **21**(10): p. 980-989.
108. Kurz, W., B. Giovanola, and R. Trivedi, *Theory of microstructural development during rapid solidification*. Acta metallurgica, 1986. **34**(5): p. 823-830.
109. StJohn, D., et al., *The contribution of constitutional supercooling to nucleation and grain formation*. Metallurgical and Materials Transactions A, 2015. **46**(11): p. 4868-4885.
110. Mohanty, U.K. and H. Sarangi, *Solidification of metals and alloys*. Casting Processes and Modelling of Metallic Materials, 2020: p. 19-40.
111. Coriell, S.R., G.B. McFadden, and R.F. Sekerka, *Cellular growth during directional solidification*. Annual Review of Materials Science, 1985. **15**(1): p. 119-145.
112. University, I.S. *Solidification*. Available from: <https://www.nde-ed.org/Physics/Materials/Structure/solidification.xhtml>.
113. Trivedi, R. and W. Kurz, *Dendritic growth*. International Materials Reviews, 1994. **39**(2): p. 49-74.
114. Kurz, W., C. Bezençon, and M. Gäumann, *Columnar to equiaxed transition in solidification processing*. Science and technology of advanced materials, 2001. **2**(1): p. 185.
115. Akbari, M. and R. Kovacevic, *Closed loop control of melt pool width in robotized laser powder-directed energy deposition process*. The International Journal of Advanced Manufacturing Technology, 2019. **104**(5): p. 2887-2898.
116. Doherty, J., A. Giamei, and B. Kear, *The importance of grain boundary morphology and cohesion on intergranular strength*. Canadian Metallurgical Quarterly, 1974. **13**(1): p. 229-236.
117. Herlach, D.M., *Non-equilibrium solidification of undercooled metallic metals*. Materials Science and Engineering: R: Reports, 1994. **12**(4-5): p. 177-272.
118. Yu, H.-S., *A generalized scheil equation for the dendritic solidification of binary alloys*. Transactions of the Korean Society of Mechanical Engineers B, 1996. **20**(7): p. 2367-2374.
119. Elliott, R., *Eutectic solidification*. International Metals Reviews, 1977. **22**(1): p. 161-186.
120. Tiller, W., et al., *The redistribution of solute atoms during the solidification of metals*. Acta metallurgica, 1953. **1**(4): p. 428-437.
121. Coura, P., O. Mesquita, and B. Costa, *Molecular-dynamics simulation of directional growth of binary mixtures*. Physical Review B, 1999. **59**(5): p. 3408.
122. Broderick, T., et al., *The effect of cooling conditions on the microstructure of rapidly solidified Ti-6Al-4V*. Metallurgical Transactions A, 1985. **16**(11): p. 1951-1959.
123. Rajan, T., C. Sharma, and A. Sharma, *Heat treatment: principles and techniques*. 2023: PHI Learning Pvt. Ltd.
124. Mittemeijer, E.J., *Recovery, recrystallization and grain growth*, in *Fundamentals of materials science*. 2010, Springer. p. 463-496.
125. Adedayo, A., S. Ibitoye, and O. Oyetoyan, *Annealing heat treatment effects on steel welds*. Journal of mineral, materials characterization and engineering, 2010. **9**(6): p. 547-557.
126. Wang, L. and J.G. Speer, *Quenching and partitioning steel heat treatment*. 2013, Springer.

127. Akhtar, M., S.Z. Qamar, and R. Khan, *Vacuum heat treatment: Process, applications, and parameters*, in *Comprehensive Materials Processing: Reference Module in Materials Science and Materials Engineering*. 2024, Elsevier.
128. Doi, M., *Elasticity effects on the microstructure of alloys containing coherent precipitates*. Progress in Materials Science, 1996. **40**(2): p. 79-180.
129. Wang, X., et al., *Precipitation strengthening of the aluminum alloy AA6111*. Metallurgical and Materials Transactions A, 2003. **34**(12): p. 2913-2924.
130. Handoko, W., et al., *Effect of selective-precipitations process on the corrosion resistance and hardness of dual-phase high-carbon steel*. Scientific reports, 2019. **9**(1): p. 15631.
131. Dawes, J., R. Bowerman, and R. Trepleton, *Introduction to the additive manufacturing powder metallurgy supply chain*. Johnson Matthey Technology Review, 2015. **59**(3): p. 243-256.
132. Apelian, D., *Particulate Processing (Powder Metallurgy)*, in *Encyclopedia of Materials: Science and Technology*, K.H.J. Buschow, et al., Editors. 2001, Elsevier: Oxford. p. 6761-6768.
133. Fedina, T., et al., *A comparative study of water and gas atomized low alloy steel powders for additive manufacturing*. Additive Manufacturing, 2020. **36**: p. 101675.
134. Sanchez, S., et al., *Powder Bed Fusion of nickel-based superalloys: A review*. International Journal of Machine Tools and Manufacture, 2021. **165**: p. 103729.
135. Volpato, G.M., U. Tetzlaff, and M.C. Fredel, *A comprehensive literature review on laser powder bed fusion of Inconel superalloys*. Additive Manufacturing, 2022. **55**: p. 102871.
136. Shahwaz, M., P. Nath, and I. Sen, *A critical review on the microstructure and mechanical properties correlation of additively manufactured nickel-based superalloys*. Journal of Alloys and Compounds, 2022. **907**: p. 164530.
137. De Terris, T., et al. *Analysis of As-Built Microstructures and Recrystallization Phenomena on Inconel 625 Alloy Obtained via Laser Powder Bed Fusion (L-PBF)*. Metals, 2021. **11**, DOI: 10.3390/met11040619.
138. Ghiaasiaan, R., et al., *Superior tensile properties of Hastelloy X enabled by additive manufacturing*. Materials Research Letters, 2021. **9**(7): p. 308-314.
139. Atabay, S.E., et al., *Microstructure and mechanical properties of rene 41 alloy manufactured by laser powder bed fusion*. Materials Science and Engineering: A, 2020. **773**: p. 138849.
140. Seidel, A., et al., *Additive Manufacturing of Powdery Ni-Based Superalloys Mar-M-247 and CM 247 LC in Hybrid Laser Metal Deposition*. Metallurgical and Materials Transactions A, 2018. **49**(9): p. 3812-3830.
141. Shaikh, A.S., et al., *Microstructure and mechanical properties of Haynes 282 superalloy produced by laser powder bed fusion*. Materials Today Communications, 2021. **26**: p. 102038.
142. Jena, A., S.E. Atabay, and M. Brochu, *Microstructure and mechanical properties of crack-free Inconel 738 fabricated by laser powder bed fusion*. Materials Science and Engineering: A, 2022. **850**: p. 143524.
143. Aydinöz, M.E., et al., *On the microstructural and mechanical properties of post-treated additively manufactured Inconel 718 superalloy under quasi-static and cyclic loading*. Materials Science and Engineering: A, 2016. **669**: p. 246-258.
144. Sindhura, D., M.V. Sravya, and G. Murthy, *Comprehensive microstructural evaluation of precipitation in Inconel 718*. Metallography, Microstructure, and Analysis, 2019. **8**(2): p. 233-240.
145. Sweet, J.N., E.P. Roth, and M. Moss, *Thermal conductivity of Inconel 718 and 304 stainless steel*. International Journal of Thermophysics, 1987. **8**(5): p. 593-606.
146. Zhang, Y., et al., *Tribological behavior of IN718 superalloy coating fabricated by laser additive manufacturing*. Lasers in Manufacturing and Materials Processing, 2017. **4**(4): p. 153-167.
147. Ghiban, B., et al., *Requirements of Inconel 718 alloy for aeronautical applications*. AIP Conference Proceedings, 2018. **1932**(1): p. 030016.
148. Barros, R., et al. *Laser Powder Bed Fusion of Inconel 718: Residual Stress Analysis Before and After Heat Treatment*. Metals, 2019. **9**, DOI: 10.3390/met9121290.
149. Bridges, A., et al., *Creep analysis and microstructural evaluation of a novel additively manufactured nickel-base superalloy (ABD®-900AM)*. Journal of Engineering for Gas Turbines and Power, 2023. **145**(6): p. 061005.
150. Duval, A., *PDS_231024_ABD*

@900AM-Stellar_EN_V1.

151. Duval, A. *Stellar ABD-900AM*. 2023; Available from: https://www.aubertduval.com/wp-content/uploads/2024/12/231024_ABD-900AM_Stellar_AM_V1_GB.pdf.
152. Metelkova, J., et al., *Texture of inclined up-facing surfaces in laser powder bed fusion of metals*. Additive Manufacturing, 2021. **42**: p. 101970.
153. Maconachie, T., et al., *SLM lattice structures: Properties, performance, applications and challenges*. Materials & Design, 2019. **183**: p. 108137.
154. Bhatia, S., *Effect of Machine Positional Errors on Geometric Tolerances in Additive Manufacturing*. 2014, University of Cincinnati.
155. Moylan, S., et al. *Proposal for a standardized test artifact for additive manufacturing machines and processes*. 2012. Proceedings of the Solid Freeform Fabrication Symposium, Austin, TX.
156. Ning, Y., et al., *An approach to minimize build errors in direct metal laser sintering*. IEEE Transactions on Automation Science and Engineering, 2006. **3**(1): p. 73-80.
157. Wang, X., *Calibration of shrinkage and beam offset in SLS process*. Rapid Prototyping Journal, 1999. **5**(3): p. 129-133.
158. Chen, C., et al., *Deformation and control method of thin-walled part during laser powder bed fusion of Ti-6Al-4V alloy*. The International Journal of Advanced Manufacturing Technology, 2020. **110**(11): p. 3467-3478.
159. Paul, R., *Modeling and Optimization of Powder Based Additive Manufacturing (AM) Processes*. 2013, University of Cincinnati: United States -- Ohio. p. 204.
160. Venuvinod, P.K. and W. Ma, *Rapid Prototyping: Laser-based and Other Technologies*. 2013: Springer US.
161. Das, P., et al., *Optimum Part Build Orientation in Additive Manufacturing for Minimizing Part Errors and Support Structures*. Procedia Manufacturing, 2015. **1**: p. 343-354.
162. Paul, R., S. Anand, and F. Gerner, *Effect of Thermal Deformation on Part Errors in Metal Powder Based Additive Manufacturing Processes*. Journal of Manufacturing Science and Engineering, 2014. **136**(3).
163. Hauser, C., *Selective laser sintering of a stainless steel powder*. 2003, University of Leeds.
164. Gong, H., et al., *Analysis of defect generation in Ti-6Al-4V parts made using powder bed fusion additive manufacturing processes*. Additive Manufacturing, 2014. **1-4**: p. 87-98.
165. Zeng, K., et al., *Layer by layer validation of geometrical accuracy in additive manufacturing processes*. 2013.
166. Krishnan, M., et al., *On the effect of process parameters on properties of AlSi10Mg parts produced by DMLS*. Rapid Prototyping Journal, 2014. **20**(6): p. 449-458.
167. Kruth, J.-P., E. Yasa, and J. Deckers. *Roughness Improvement in selective laser melting*.
168. Yasa, E. and J.P. Kruth, *Microstructural investigation of Selective Laser Melting 316L stainless steel parts exposed to laser re-melting*. Procedia Engineering, 2011. **19**: p. 389-395.
169. Kruth, J.P., et al., *Consolidation phenomena in laser and powder-bed based layered manufacturing*. CIRP Annals, 2007. **56**(2): p. 730-759.
170. Lian, Z., *New Trends in Technologies: Devices, Computer, Communication and Industrial Systems*. 2010: IntechOpen.
171. Aiyiti, W., et al., *Investigation of the overlapping parameters of MPAW-based rapid prototyping*. Rapid Prototyping Journal, 2006. **12**(3): p. 165-172.
172. Matache, G., et al., *Edge and corner effects in selective laser melting of IN 625 alloy*. Manufacturing Review, 2020. **7**: p. 8.
173. Chechik, L., *Process monitoring and control during additive manufacturing*. 2022, University of Sheffield.
174. Yasa, E., et al., *Investigation on occurrence of elevated edges in selective laser melting*. 2009.
175. Kleszczynski, S., et al., *Error detection in laser beam melting systems by high resolution imaging*. 2012.
176. Wischeropp, T.M., et al., *Measurement of actual powder layer height and packing density in a single layer in selective laser melting*. Additive Manufacturing, 2019. **28**: p. 176-183.
177. Tang, P., et al., *The Formation of Humps and Ripples During Selective Laser Melting of 316L Stainless Steel*. JOM, 2020. **72**(3): p. 1128-1137.
178. Wang, W., J. Ning, and S.Y. Liang, *Analytical prediction of balling, lack-of-fusion and keyholing thresholds in powder bed fusion*. Applied Sciences, 2021. **11**(24): p. 12053.

179. Das, S., *Physical Aspects of Process Control in Selective Laser Sintering of Metals*. Advanced Engineering Materials, 2003. **5**(10): p. 701-711.
180. Li, R., et al., *Balling behavior of stainless steel and nickel powder during selective laser melting process*. The International Journal of Advanced Manufacturing Technology, 2012. **59**(9): p. 1025-1035.
181. Tang, C., K.Q. Le, and C.H. Wong, *Physics of humping formation in laser powder bed fusion*. International Journal of Heat and Mass Transfer, 2020. **149**: p. 119172.
182. Gu, D. and Y. Shen, *Balling phenomena in direct laser sintering of stainless steel powder: Metallurgical mechanisms and control methods*. Materials & Design, 2009. **30**(8): p. 2903-2910.
183. Balhara, H., et al., *Ripple formation and its effect on the multi-scale microstructure of Directed Energy Deposition (DED)-printed 316L components*. arXiv preprint arXiv:2111.11601, 2021.
184. Caprio, L., A.G. Demir, and B. Previtali, *Observing molten pool surface oscillations during keyhole processing in laser powder bed fusion as a novel method to estimate the penetration depth*. Additive Manufacturing, 2020. **36**: p. 101470.
185. Jia, Q. and D. Gu, *Selective laser melting additive manufactured Inconel 718 superalloy parts: High-temperature oxidation property and its mechanisms*. Optics & Laser Technology, 2014. **62**: p. 161-171.
186. Lu, C., et al., *An investigation on the oxidation behavior of spatters generated during the laser powder bed fusion of 316L stainless steel*. Applied Surface Science, 2022. **586**: p. 152796.
187. Asgharzadeh, H. and A. Simchi, *Effect of sintering atmosphere and carbon content on the densification and microstructure of laser-sintered M2 high-speed steel powder*. Materials Science and Engineering: A, 2005. **403**(1): p. 290-298.
188. Özsoy, K., Z.A. Çevik, and A. Erçetin, *THE EFFECT OF MACHINING PROCESSES ON THE PHYSICAL AND SURFACE MORPHOLOGY OF Ti6Al4V SPECIMENS PRODUCED THROUGH POWDER BED FUSION ADDITIVE MANUFACTURING*. International Journal of 3D Printing Technologies and Digital Industry, 2021. **5**(2): p. 187-194.
189. Lesyk, D.A., et al., *Surface Shot Peening Post-processing of Inconel 718 Alloy Parts Printed by Laser Powder Bed Fusion Additive Manufacturing*. Journal of Materials Engineering and Performance, 2021. **30**(9): p. 6982-6995.
190. Pitrmuc, Z., et al. *Mechanical and Microstructural Anisotropy of Laser Powder Bed Fusion 316L Stainless Steel*. Materials, 2022. **15**, DOI: 10.3390/ma15020551.
191. Zhu, J., et al., *Suppression of epitaxial grain growth in laser powder bed fusion fabricated Inconel 625 through intermittent variation of the build layer thickness*. Materials Science and Engineering: A, 2023. **886**: p. 145679.
192. Ren, G., et al., *Effect of subgrain microstructure on the mechanical properties of Invar 36 specimens prepared by laser powder bed fusion*. Journal of Alloys and Compounds, 2024. **1004**: p. 175839.
193. Mantri, S.A., et al., *Effect of micro-segregation of alloying elements on the precipitation behaviour in laser surface engineered Alloy 718*. Acta Materialia, 2021. **210**: p. 116844.
194. Paul, M.J., et al., *Fracture resistance of AlSi10Mg fabricated by laser powder bed fusion*. Acta Materialia, 2021. **211**: p. 116869.
195. Krakhmalev, P., et al. *Microstructure, Solidification Texture, and Thermal Stability of 316 L Stainless Steel Manufactured by Laser Powder Bed Fusion*. Metals, 2018. **8**, DOI: 10.3390/met8080643.
196. Wang, X., et al., *Investigation on microsegregation of IN718 alloy during additive manufacturing via integrated phase-field and finite-element modeling*. Journal of Materials Engineering and Performance, 2019. **28**(2): p. 657-665.
197. Ghosh, S., et al., *Uncertainty analysis of microsegregation during laser powder bed fusion*. Modelling and Simulation in Materials Science and Engineering, 2019. **27**(3): p. 034002.
198. Song, Z., et al., *Effects of non-equilibrium microstructures on microstructure evolution and mechanical properties of laser powder bed fusion IN625 Ni-based superalloy during long-term thermal exposure at 700 °C and 750 °C*. Materials Science and Engineering: A, 2022. **856**: p. 143883.
199. Zhang, W., et al., *Controlling of microstructures and mechanical properties based on the non-equilibrium microstructures of a nickel-based superalloy fabricated by laser powder bed fusion*. Materials Science and Engineering: A, 2024. **900**: p. 146487.

200. Gu, D. and Y. Shen, *Effects of processing parameters on consolidation and microstructure of W–Cu components by DMLS*. Journal of Alloys and Compounds, 2009. **473**(1): p. 107-115.
201. Gruber, K., et al., *Impact of high temperature stress relieving on final properties of Inconel 718 processed by laser powder bed fusion*. Materials Science and Engineering: A, 2021. **813**: p. 141111.
202. Zhou, L., et al., *Microstructure, precipitates and mechanical properties of powder bed fused inconel 718 before and after heat treatment*. Journal of Materials Science & Technology, 2019. **35**(6): p. 1153-1164.
203. Mercelis, P. and J.P. Kruth, *Residual stresses in selective laser sintering and selective laser melting*. Rapid prototyping journal, 2006. **12**(5): p. 254-265.
204. Li, C., et al., *Residual Stress in Metal Additive Manufacturing*. Procedia CIRP, 2018. **71**: p. 348-353.
205. Parry, L., I.A. Ashcroft, and R.D. Wildman, *Understanding the effect of laser scan strategy on residual stress in selective laser melting through thermo-mechanical simulation*. Additive Manufacturing, 2016. **12**: p. 1-15.
206. Bartlett, J.L. and X. Li, *An overview of residual stresses in metal powder bed fusion*. Additive Manufacturing, 2019. **27**: p. 131-149.
207. Mukherjee, T., et al., *Mitigation of thermal distortion during additive manufacturing*. Scripta Materialia, 2017. **127**: p. 79-83.
208. Wu, S., et al., *Hot cracking evolution and formation mechanism in 2195 Al-Li alloy printed by laser powder bed fusion*. Additive Manufacturing, 2022. **54**: p. 102762.
209. Zhang, X., H. Zheng, and W. Yu, *A review on solidification cracks in high-strength aluminum alloys via laser powder bed fusion*. Materials Today: Proceedings, 2022. **70**: p. 465-469.
210. Jeong, S.G., et al., *Liquation cracking in laser powder bed fusion-fabricated Inconel718 of as-built, stress-relieved, and hot isostatic pressed conditions*. Materials Science and Engineering: A, 2023. **888**: p. 145797.
211. Ploshikhin, V., et al., *Integrated mechanical-metallurgical approach to modeling of solidification cracking in welds*, in *Hot cracking phenomena in welds*. 2005, Springer. p. 223-244.
212. Agarwal, G., *Study of solidification cracking during laser welding in advanced high strength steels: A combined experimental and numerical approach*. 2019.
213. Robinson, J. and M. Scott, *Liquation cracking during the welding of austenitic stainless steels and nickel alloys*. Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences, 1980. **295**(1413): p. 105-117.
214. Deng, Q., et al., *Exceptional strength paired with increased cold cracking susceptibility in laser powder bed fusion of a Mg-RE alloy*. Journal of Materials Science & Technology, 2025. **213**: p. 300-314.
215. Wei, S., et al., *Suppressing cold cracking in laser powder bed fused Al-Fe-V-Si alloy using top-hat laser profile*. Materials Science and Engineering: A, 2025. **924**: p. 147858.
216. Li, Y., et al., *Multi-scale study of ductility-dip cracking in nickel-based alloy dissimilar metal weld*. Journal of Materials Science & Technology, 2019. **35**(4): p. 545-559.
217. Liu, T., Q. Wang, and K. Guan, *Analysis of the ductility dip cracking in the nickel-base alloy welding overlay on heat exchanger tube sheet*. Engineering Failure Analysis, 2023. **147**: p. 107158.
218. Zhang, C., et al., *On the Fabrication of Defect-Free Nickel-Rich Nickel–Titanium Parts Using Laser Powder Bed Fusion*. Journal of Manufacturing Science and Engineering, 2022. **144**(9).
219. Garibaldi, M., *Laser additive manufacturing of soft magnetic cores for rotating electrical machinery: materials development and part design*. 2018.
220. Okunkova, A.A., et al. *On Defect Minimization Caused by Oxide Phase Formation in Laser Powder Bed Fusion*. Metals, 2022. **12**, DOI: 10.3390/met12050760.
221. Kivirasi, E., et al., *Detecting spattering phenomena by using high speed imaging in L-PBD of 316 L*. Procedia CIRP, 2020. **94**: p. 398-403.
222. King, W.E., et al., *Observation of keyhole-mode laser melting in laser powder-bed fusion additive manufacturing*. Journal of Materials Processing Technology, 2014. **214**(12): p. 2915-2925.
223. Bidare, P., et al., *Fluid and particle dynamics in laser powder bed fusion*. Acta Materialia, 2018. **142**: p. 107-120.
224. Huang, Y., et al., *Keyhole fluctuation and pore formation mechanisms during laser powder bed fusion additive manufacturing*. Nature communications, 2022. **13**(1): p. 1170.

225. O'Brien, N., et al., *Computational analysis and experiments of spatter transport in a laser powder bed fusion machine*. Additive Manufacturing, 2024. **84**: p. 104133.
226. Esmaeilzadeh, R., et al., *On the effect of spatter particles distribution on the quality of Hastelloy X parts made by laser powder-bed fusion additive manufacturing*. Journal of Manufacturing Processes, 2019. **37**: p. 11-20.
227. Li, Z., et al. *A Review of Spatter in Laser Powder Bed Fusion Additive Manufacturing: In Situ Detection, Generation, Effects, and Countermeasures*. Micromachines, 2022. **13**, DOI: 10.3390/mi13081366.
228. Craeghs, T., et al., *Online quality control of selective laser melting*. 22nd Annual International Solid Freeform Fabrication Symposium - An Additive Manufacturing Conference, SFF 2011, 2011.
229. Nudelis, N. and P. Mayr, *A novel classification method for pores in laser powder bed fusion*. Metals, 2021. **11**(12): p. 1912.
230. Tang, M., P.C. Pistorius, and J.L. Beuth, *Prediction of lack-of-fusion porosity for powder bed fusion*. Additive Manufacturing, 2017. **14**: p. 39-48.
231. Wang, W., J. Ning, and S.Y. Liang, *Prediction of lack-of-fusion porosity in laser powder-bed fusion considering boundary conditions and sensitivity to laser power absorption*. The International Journal of Advanced Manufacturing Technology, 2021. **112**(1): p. 61-70.
232. Snow, Z., et al., *Observation of spatter-induced stochastic lack-of-fusion in laser powder bed fusion using in situ process monitoring*. Additive Manufacturing, 2023. **61**: p. 103298.
233. Zhu, W., et al., *A multi-scale experimental investigation for fatigue limit and fatigue crack initiation behavior of powder bed fusion-laser beam 316L stainless steel*. Materials Science and Engineering: A, 2023. **866**: p. 144692.
234. Aliyu, A.A.A., et al., *Lack of fusion-induced cracking effect on tensile and fatigue behaviours of laser powder-bed fusion-processed Ti-6Al-4V implant*. Engineering Failure Analysis, 2025. **168**: p. 109095.
235. Ronneberg, T., C.M. Davies, and P.A. Hooper, *Revealing relationships between porosity, microstructure and mechanical properties of laser powder bed fusion 316L stainless steel through heat treatment*. Materials & Design, 2020. **189**: p. 108481.
236. Yu, T. and J. Zhao, *Quantifying the mechanisms of keyhole pore evolutions and the role of metal-vapor condensation in laser powder bed fusion*. Additive Manufacturing, 2023. **72**: p. 103642.
237. Shrestha, S. and K. Chou, *Formation of keyhole and lack of fusion pores during the laser powder bed fusion process*. Manufacturing Letters, 2022. **32**: p. 19-23.
238. Guo, L., et al., *Understanding keyhole induced-porosities in laser powder bed fusion of aluminum and elimination strategy*. International Journal of Machine Tools and Manufacture, 2023. **184**: p. 103977.
239. Gan, Z., et al., *Universal scaling laws of keyhole stability and porosity in 3D printing of metals*. Nature Communications, 2021. **12**(1): p. 2379.
240. Montalbano, T., et al., *Uncovering the coupled impact of defect morphology and microstructure on the tensile behavior of Ti-6Al-4V fabricated via laser powder bed fusion*. Journal of Materials Processing Technology, 2021. **294**: p. 117113.
241. Kan, W.H., et al., *A critical review on the effects of process-induced porosity on the mechanical properties of alloys fabricated by laser powder bed fusion*. Journal of Materials Science, 2022. **57**(21): p. 9818-9865.
242. Sola, A. and A. Nouri, *Microstructural porosity in additive manufacturing: The formation and detection of pores in metal parts fabricated by powder bed fusion*. Journal of Advanced Manufacturing and Processing, 2019. **1**(3): p. e10021.
243. du Plessis, A., *Effects of process parameters on porosity in laser powder bed fusion revealed by X-ray tomography*. Additive Manufacturing, 2019. **30**: p. 100871.
244. Vukcum, V. and R. Gupta, *Review on corrosion performance of laser powder-bed fusion printed 316L stainless steel: Effect of processing parameters, manufacturing defects, post-processing, feedstock, and microstructure*. Materials & Design, 2022. **221**: p. 110874.
245. Atkinson, H.V. and S. Davies, *Fundamental aspects of hot isostatic pressing: An overview*. Metallurgical and Materials Transactions A, 2000. **31**(12): p. 2981-3000.
246. Kluczyński, J., et al., *Hot isostatic pressing influence on the mechanical properties of selectively laser-melted 316L steel*. Bulletin of the Polish Academy of Sciences Technical Sciences, 2020. **68**(No. 6): p. 1413-1424.

247. Tillmann, W., et al., *Hot isostatic pressing of IN718 components manufactured by selective laser melting*. Additive Manufacturing, 2017. **13**: p. 93-102.
248. Herzog, D., et al., *Design guidelines for laser powder bed fusion in Inconel 718*. Journal of Laser Applications, 2021. **34**(1): p. 012015.
249. Wu, Z., S.P. Narra, and A. Rollett, *Exploring the fabrication limits of thin-wall structures in a laser powder bed fusion process*. The International Journal of Advanced Manufacturing Technology, 2020. **110**(1): p. 191-207.
250. Zea Pérez, J.M., et al., *On the manufacturability of Inconel 718 thin-walled honeycomb structures by laser powder bed fusion*. Rapid Prototyping Journal, 2022. **28**(2): p. 307-316.
251. Zhou, H., et al., *Formation and Evolution of Surface Morphology in Overhang Structure of IN718 Superalloy Fabricated by Laser Powder Bed Fusion*. Acta Metallurgica Sinica (English Letters), 2023. **36**(9): p. 1433-1453.
252. Sit, C.K., et al., *A novel experimental approach to exploring process parameters for fabricating support-free horizontal overhangs in LPBF of Inconel 718*. Virtual and Physical Prototyping, 2024. **19**(1): p. e2422394.
253. Chen, Q., *Multiscale Process Modeling of Residual Deformation and Defect Formation for Laser Powder Bed Fusion Additive Manufacturing*. 2021, University of Pittsburgh.
254. Barale, G., *Characterization of Inconel 718 fabricated by Selective Laser Melting (SLM) to achieve more productive parameters*. 2019, Politecnico di Torino.
255. Abo Znemah, R.A., *Size Effects in Additively Manufactured Inconel-718 Thin Elements and Cellular Structures*. 2023.
256. Du, B., et al., *The heterogeneous microstructure in laser powder bed fabricated Inconel 718 pillar and its influence on mechanical properties*. Materials Science and Engineering: A, 2023. **872**: p. 144953.
257. Kasperovich, G., et al., *The effect of build direction and geometric optimization in laser powder bed fusion of Inconel 718 structures with internal channels*. Materials & Design, 2021. **207**: p. 109858.
258. Qin, Y., et al., *Automatic determination of part build orientation for laser powder bed fusion*. Virtual and Physical Prototyping, 2021. **16**(1): p. 29-49.
259. Brika, S.E., et al., *Multi-Objective Build Orientation Optimization for Powder Bed Fusion by Laser*. Journal of Manufacturing Science and Engineering, 2017. **139**(11).
260. Watring, D.S., et al., *Effects of laser-energy density and build orientation on the structure–property relationships in as-built Inconel 718 manufactured by laser powder bed fusion*. Additive Manufacturing, 2020. **36**: p. 101425.
261. Wang, Y.C., et al., *Building height-related creep properties of Inconel 718 superalloy fabricated by laser powder bed fusion*. Materials Science and Engineering: A, 2022. **854**: p. 143861.
262. Javidrad, H.R. and F. Javidrad, *Review of state-of-the-art research on the design and manufacturing of support structures for powder-bed fusion additive manufacturing*. Progress in Additive Manufacturing, 2023. **8**(6): p. 1517-1542.
263. Marques, A., et al., *Inconel 718 produced by laser powder bed fusion: an overview of the influence of processing parameters on microstructural and mechanical properties*. The International Journal of Advanced Manufacturing Technology, 2022. **121**(9): p. 5651-5675.
264. Mukherjee, T., et al., *Dimensionless numbers in additive manufacturing*. Journal of Applied Physics, 2017. **121**(6): p. 064904.
265. Rubenchik, A.M., W.E. King, and S.S. Wu, *Scaling laws for the additive manufacturing*. Journal of Materials Processing Technology, 2018. **257**: p. 234-243.
266. Ion, J.C., H.R. Shercliff, and M.F. Ashby, *Diagrams for laser materials processing*. Acta Metallurgica et Materialia, 1992. **40**(7): p. 1539-1551.
267. Thomas, M., G.J. Baxter, and I. Todd, *Normalised model-based processing diagrams for additive layer manufacture of engineering alloys*. Acta Materialia, 2016. **108**: p. 26-35.
268. Phan, M.A.L., O. Dew, and I. Todd, *Predictive process diagram for parameters selection in laser powder bed fusion to achieve high-density and low-cracking built parts*. Additive Manufacturing, 2024. **85**: p. 104145.
269. Malekipour, E. and H. El-Mounayri, *Scanning Strategies in the PBF Process: A Critical Review*. 2020.

270. Ravichander, B.B., et al., *Experimental investigation of laser scan strategy on the microstructure and properties of Inconel 718 parts fabricated by laser powder bed fusion*. *Materials Characterization*, 2022. **186**: p. 111765.
271. Liu, L., et al., *Effect of Scanning Strategies on the Microstructure and Mechanical Properties of Inconel 718 Alloy Fabricated by Laser Powder Bed Fusion*. *Advanced Engineering Materials*, 2023. **25**(5): p. 2200492.
272. Stegman, B.T., et al. *The Influence of Powder Particle Size Distributions on Mechanical Properties of Alloy 718 by Laser Powder Bed Fusion*. *Metals*, 2023. **13**, DOI: 10.3390/met13081384.
273. Sendino, S., et al., *Effect of powder particle size distribution on the surface finish of components manufactured by laser powder bed fusion*. *The International Journal of Advanced Manufacturing Technology*, 2023. **124**(3): p. 789-799.
274. Ruan, G., et al., *A comparative study on laser powder bed fusion of IN718 powders produced by gas atomization and plasma rotating electrode process*. *Materials Science and Engineering: A*, 2022. **850**: p. 143589.
275. Castillo, S.J., et al., *Characterization of recycled Inconel 718 metal powder for assessing its reusability in the laser powder bed fusion process*. *Clean Technologies and Recycling*, 2022. **2**(1): p. 32-46.
276. Rock, C., et al., *The Influence of Powder Reuse on the Properties of Nickel Super Alloy ATI 718™ in Laser Powder Bed Fusion Additive Manufacturing*. *Metallurgical and Materials Transactions B*, 2021. **52**(2): p. 676-688.
277. Grasso, M., et al., *In-situ measurement and monitoring methods for metal powder bed fusion: an updated review*. *Measurement Science and Technology*, 2021. **32**(11): p. 112001.
278. Grasso, M., F. Gallina, and B.M. Colosimo, *Data fusion methods for statistical process monitoring and quality characterization in metal additive manufacturing*. *Procedia CIRP*, 2018. **75**: p. 103-107.
279. Scime, L., et al., *Layer-wise anomaly detection and classification for powder bed additive manufacturing processes: A machine-agnostic algorithm for real-time pixel-wise semantic segmentation*. *Additive Manufacturing*, 2020. **36**.
280. Lu, Q.Y., et al., *Identification and evaluation of defects in selective laser melted 316L stainless steel parts via in-situ monitoring and micro computed tomography*. *Additive Manufacturing*, 2020. **35**: p. 101287.
281. Zhang, B., et al., *In situ surface topography of laser powder bed fusion using fringe projection*. *Additive Manufacturing*, 2016. **12**: p. 100-107.
282. Kalms, M., et al., *New approach to evaluate 3D laser printed parts in powder bed fusion-based additive manufacturing in-line within closed space*. *Additive Manufacturing*, 2019. **26**: p. 161-165.
283. Tan Phuc, L. and M. Seita, *A high-resolution and large field-of-view scanner for in-line characterization of powder bed defects during additive manufacturing*. *Materials & Design*, 2019. **164**: p. 107562.
284. Barrett, C., et al., *Micron-Level Layer-Wise Surface Profilometry to Detect Porosity Defects in Powder Bed Fusion of Inconel 718*. *JOM*, 2018. **70**(9): p. 1844-1852.
285. DePond, P.J., et al., *In situ measurements of layer roughness during laser powder bed fusion additive manufacturing using low coherence scanning interferometry*. *Materials & Design*, 2018. **154**: p. 347-359.
286. Grasso, M., et al., *In-Process Monitoring of Selective Laser Melting: Spatial Detection of Defects Via Image Data Analysis*. *Journal of Manufacturing Science and Engineering*, 2016. **139**(5).
287. Lough, C.S., et al., *Correlation of SWIR imaging with LPBF 304L stainless steel part properties*. *Additive Manufacturing*, 2020. **35**: p. 101359.
288. Mitchell, J.A., et al., *Linking pyrometry to porosity in additively manufactured metals*. *Additive Manufacturing*, 2020. **31**: p. 100946.
289. Yang, L., et al., *Monitoring and detection of meltpool and spatter regions in laser powder bed fusion of super alloy Inconel 625*. *Progress in Additive Manufacturing*, 2020. **5**(4): p. 367-378.
290. Eschner, E., T. Staudt, and M. Schmidt, *Correlation of spatter behavior and process zone formation in powder bed fusion of metals*. *CIRP Annals*, 2020. **69**(1): p. 209-212.
291. Grasso, M. and B.M. Colosimo, *A statistical learning method for image-based monitoring of the plume signature in laser powder bed fusion*. *Robotics and Computer-Integrated Manufacturing*, 2019. **57**: p. 103-115.

292. Shevchik, S.A., et al., *Deep Learning for In Situ and Real-Time Quality Monitoring in Additive Manufacturing Using Acoustic Emission*. IEEE Transactions on Industrial Informatics, 2019. **15**(9): p. 5194-5203.
293. Forien, J.-B., et al., *Detecting keyhole pore defects and monitoring process signatures during laser powder bed fusion: A correlation between in situ pyrometry and ex situ X-ray radiography*. Additive Manufacturing, 2020. **35**: p. 101336.
294. Samaei, A., et al., *Understanding coaxial photodiode-based multispectral pyrometer measurements at the overhang regions in laser powder bed fusion for part qualification*. Additive Manufacturing, 2024. **93**: p. 104398.
295. Bisht, M., et al., *Correlation of selective laser melting-melt pool events with the tensile properties of Ti-6Al-4V ELI processed by laser powder bed fusion*. Additive Manufacturing, 2018. **22**: p. 302-306.
296. Kwon, O., et al., *A deep neural network for classification of melt-pool images in metal additive manufacturing*. Journal of Intelligent Manufacturing, 2020. **31**(2): p. 375-386.
297. Vasileska, E., et al., *Layer-wise control of selective laser melting by means of inline melt pool area measurements*. Journal of Laser Applications, 2020. **32**(2): p. 022057.
298. Lough, C.S., et al., *In-situ optical emission spectroscopy of selective laser melting*. Journal of Manufacturing Processes, 2020. **53**: p. 336-341.
299. Ziefuss, A.R., et al., *In-situ composition analysis during laser powder bed fusion of Nd-Fe-based feedstock using machine-integrated optical emission spectroscopy*. Materials & Design, 2024. **244**: p. 113211.
300. Leung, C.L.A., et al., *The effect of powder oxidation on defect formation in laser additive manufacturing*. Acta Materialia, 2019. **166**: p. 294-305.
301. Calta, N.P., et al., *Pressure dependence of the laser-metal interaction under laser powder bed fusion conditions probed by in situ X-ray imaging*. Additive Manufacturing, 2020. **32**: p. 101084.
302. Paulson, N.H., et al., *Correlations between thermal history and keyhole porosity in laser powder bed fusion*. Additive Manufacturing, 2020. **34**: p. 101213.
303. Schmeiser, F., et al., *Experimental observation of stress formation during selective laser melting using in situ X-ray diffraction*. Additive Manufacturing, 2020. **32**: p. 101028.
304. Eschner, N., et al., *Classification of specimen density in Laser Powder Bed Fusion (L-PBF) using in-process structure-borne acoustic process emissions*. Additive Manufacturing, 2020. **34**: p. 101324.
305. Hehr, A., et al. *Smart Build-Plate for Metal Additive Manufacturing Processes*. Sensors, 2020. **20**, DOI: 10.3390/s20020360.
306. Mercado Rivera, F.J. and A.J. Rojas Arciniegas, *Additive manufacturing methods: techniques, materials, and closed-loop control applications*. The International Journal of Advanced Manufacturing Technology, 2020. **109**(1): p. 17-31.
307. Wang, R., et al., *Real-time process monitoring and closed-loop control on laser power via a customized laser powder bed fusion platform*. Additive Manufacturing, 2023. **66**: p. 103449.
308. Nahr, F., et al., *Advanced process control in laser-based powder bed fusion—Smart Fusion feedback-loop control as a path to uniform properties for complex structures?* Journal of Materials Research and Technology, 2025. **34**: p. 604-618.
309. Wenger, R., et al. *Design of a research system for process monitoring and closed-loop control of laser powder bed fusion*. in *Proceeding Lasers in Manufacturing Conference*. 2023.
310. Lhuissier, P., et al., *In situ 3D X-ray microtomography of laser-based powder-bed fusion (L-PBF)—A feasibility study*. Additive Manufacturing, 2020. **34**: p. 101271.
311. Higginson, R.L. and C.M. Sellars, *Worked Examples in Quantitative Metallography*. 2003: Taylor & Francis Group.
312. ASTM, *Standard Test Method for Apparent Density of Free-Flowing Metal Powders Using the Hall Flowmeter Funnel*. 2025.
313. Aconity3D, *AconityLAB Technical Specifications*.
314. THORLABS. *FL905-10 - Ø1" Soft-Coated Laser Line Filter*. Available from: <https://www.thorlabs.com/thorproduct.cfm?partnumber=FL905-10>.
315. THORLABS. *NENIR510B - Ø1/2" Unmounted NIR Absorptive ND Filter*. Available from: <https://www.thorlabs.com/thorproduct.cfm?partnumber=NENIR510B>.
316. THORLABS, *FESH1000 - Ø25.0 mm Shortpass Filter*.

317. Tec-Science. *Planck's law and Wien's displacement law*. 2020; Available from: <https://www.tec-science.com/thermodynamics/temperature/plancks-law-of-blackbody-radiation/>.
318. THORLABS. *R3L3S1PR - Positive Reflective 1951 USAF Test Target, 3" x 3"*. Available from: <https://www.thorlabs.com/thorproduct.cfm?partnumber=R3L3S1PR>.
319. Pichler, P., et al., *Surface tension and thermal conductivity of NIST SRM 1155a (AISI 316L stainless steel)*. International Journal of Thermophysics, 2022. **43**(5): p. 66.
320. Cheng, B., et al., *Melt pool sensing and size analysis in laser powder-bed metal additive manufacturing*. Journal of Manufacturing Processes, 2018. **32**: p. 744-753.
321. Heigel, J.C. and B.M. Lane, *Measurement of the Melt Pool Length During Single Scan Tracks in a Commercial Laser Powder Bed Fusion Process*. Journal of Manufacturing Science and Engineering, 2018. **140**(5).
322. Hagenlocher, C., et al., *The Effect of Heat Accumulation on the Local Grain Structure in Laser-Directed Energy Deposition of Aluminium*. Metals, 2022. **12**(10): p. 1601.
323. Bosio, F., et al., *Strengthening strategies for an Al alloy processed by in-situ alloying during laser powder bed fusion*. Materials & Design, 2021. **212**: p. 110247.
324. Serrano-Munoz, I., et al., *The residual stress in as-built Laser Powder Bed Fusion IN718 alloy as a consequence of the scanning strategy induced microstructure*. Scientific Reports, 2020. **10**(1): p. 14645.
325. Bayat, M., S. Mohanty, and J.H. Hattel, *Multiphysics modelling of lack-of-fusion voids formation and evolution in IN718 made by multi-track/multi-layer L-PBF*. International Journal of Heat and Mass Transfer, 2019. **139**: p. 95-114.
326. Khorasani, M., et al., *Benchmark models for conduction and keyhole modes in laser-based powder bed fusion of Inconel 718*. Optics & Laser Technology, 2023. **164**: p. 109509.
327. Sheridan, L., J.E. Gockel, and O.E. Scott-Emuakpor, *Primary Processing Parameters, Porosity Production, and Fatigue Prediction for Additively Manufactured Alloy 718*. Journal of Materials Engineering and Performance, 2019. **28**(9): p. 5387-5397.
328. Basak, D., R.A. Overfelt, and D. Wang, *Measurement of Specific Heat Capacity and Electrical Resistivity of Industrial Alloys Using Pulse Heating Techniques*. International Journal of Thermophysics, 2003. **24**(6): p. 1721-1733.
329. Terpiłowski, J., et al. *Thermal Diffusivity Characteristics of the IN718 Alloy Tested with the Modified Pulse Method*. Materials, 2022. **15**, DOI: 10.3390/ma15227881.
330. Ankudinov, V., et al., *Numerical Simulation of Thermal Conductivity of Stainless Steel and Al-12Si Powders for Additive Manufacturing*. Journal of Heat Transfer, 2022. **144**(5).
331. Zhang, S., et al., *On thermal properties of metallic powder in laser powder bed fusion additive manufacturing*. Journal of Manufacturing Processes, 2019. **47**: p. 382-392.
332. Charles, A., et al., *Elucidation of dross formation in laser powder bed fusion at down-facing surfaces: Phenomenon-oriented multiphysics simulation and experimental validation*. Additive Manufacturing, 2022. **50**: p. 102551.
333. Charles, A., et al., *Down-facing surfaces in laser powder bed fusion of Ti6Al4V: Effect of dross formation on dimensional accuracy and surface texture*. Additive Manufacturing, 2021. **46**: p. 102148.
334. Calignano, F., *Investigation of the accuracy and roughness in the laser powder bed fusion process*. Virtual and Physical Prototyping, 2018. **13**(2): p. 97-104.
335. Ravichander, B.B., et al., *Experimental investigation of laser scan strategy on the microstructure and properties of Inconel 718 parts fabricated by laser powder bed fusion*. Materials Characterization, 2022. **186**: p. 111765.
336. Nadammal, N., et al., *Critical role of scan strategies on the development of microstructure, texture, and residual stresses during laser powder bed fusion additive manufacturing*. Additive Manufacturing, 2021. **38**: p. 101792.
337. Hantke, N., et al. *Influence of Recoater Speed on Powder Bed and Part Quality in Powder Bed Fusion of Metals Using a Laser Beam*. Metals, 2025. **15**, DOI: 10.3390/met15030225.
338. Fox, H., A.B. Kamaraj, and D. Drake, *Investigating the effect of powder recoater blade material on the mechanical properties of parts manufactured using a powder-bed fusion process*. Manufacturing Letters, 2022. **33**: p. 561-568.

339. Wenzler, D.L., et al., *reAM250 — An open-source research platform for process monitoring and control in Powder Bed Fusion of Metals using a laser beam*. Procedia CIRP, 2024. **124**: p. 295-300.
340. Uriati, F. and G. Nicoletto, *A comparison of Inconel 718 obtained with three L-PBF production systems in terms of process parameters, as-built surface quality, and fatigue performance*. International Journal of Fatigue, 2022. **162**: p. 107004.
341. Reijonen, J., et al., *Effect of hard and soft re-coater blade on porosity and processability of thin walls and overhangs in laser powder bed fusion additive manufacturing*. The International Journal of Advanced Manufacturing Technology, 2024. **130**(5): p. 2283-2296.
342. Doris, A., *Effect of Multiple Machine Configurations and Wall Thickness on Microstructure and Microhardness of Laser Powder Bed Fusion (L-PBF) Additively Manufactured Heat-Treated Inconel 718 Products*. 2023, The University of Texas at El Paso.
343. Pauzon, C., et al. *Effect of the Process Atmosphere Composition on Alloy 718 Produced by Laser Powder Bed Fusion*. Metals, 2021. **11**, DOI: 10.3390/met11081254.
344. Deckers, T., et al., *Impact of processing gas composition on process stability and properties of PBF-LB/M processed alloy 718*. Journal of Manufacturing Processes, 2024. **120**: p. 712-718.
345. Abeyta, A., *An Experimental Approach for Characterizing Gas Flow in the Build Chamber of Laser Powder Bed Fusion Systems Utilizing Particle Image Velocimetry*. 2024, University of Washington: United States -- Washington. p. 78.
346. Chien, C.-Y., et al., *Numerical and experimental investigation into gas flow field and spattering phenomena in laser powder bed fusion processing of Inconel 718*. Materials & Design, 2021. **210**: p. 110107.
347. Ullah, A., et al., *Additive manufacturing of ceramics via the laser powder bed fusion process*. International journal of applied ceramic technology, 2025. **22**(3): p. e15087.
348. Cozzolino, E., et al., *Energy efficiency of Gaussian and ring profiles for LPBF of nickel alloy 718*. The International Journal of Advanced Manufacturing Technology, 2024. **132**(5): p. 3093-3104.
349. Mirzabeigi, N., et al., *Tailored microstructure in laser-based powder bed fusion of IN718 through novel beam shaping technology*. Journal of Laser Applications, 2024. **36**(4): p. 042070.
350. Zhong, Q., et al., *High power laser powder bed fusion of Inconel 718 alloy: Effect of laser focus shift on formability, microstructure and mechanical properties*. Journal of Materials Processing Technology, 2023. **311**: p. 117824.
351. Park, J.-H., et al., *Effect of Preheating Temperature on Microstructural and Mechanical Properties of Inconel 718 Fabricated by Selective Laser Melting*. Metals and Materials International, 2022. **28**(11): p. 2836-2848.
352. Shrivastava, A., et al., *Exploring How LPBF process parameters impact the interface characteristics of LPBF Inconel 718 deposited on Inconel 718 wrought substrates*. Optics & Laser Technology, 2024. **174**: p. 110571.
353. Chen, W.-Y., et al., *Laser powder bed fusion of Inconel 718 on 316 stainless steel*. Additive Manufacturing, 2020. **36**: p. 101500.
354. Yu, S., et al., *Effect of static magnetic field on the molten pool dynamics during laser powder bed fusion of Inconel 718 superalloy*. International Journal of Thermal Sciences, 2024. **197**: p. 108851.
355. Liu, F., et al., *Control of microstructure and mechanical properties of laser solid formed Inconel 718 superalloy by electromagnetic stirring*. Optics & Laser Technology, 2018. **99**: p. 342-350.
356. Wang, H., et al., *Ultrasonic vibration-assisted laser engineered net shaping of Inconel 718 parts: Effects of ultrasonic frequency on microstructural and mechanical properties*. Journal of Materials Processing Technology, 2020. **276**: p. 116395.
357. Wu, M., et al., *Plasma evolution in laser powder bed fusion using a double pulse format: Time-resolved measurements and physics-based modeling*. Additive Manufacturing, 2024. **84**: p. 104067.