Surrogate modelling strategies for the prediction of near-field blast impulse

Jordan J. Pannell

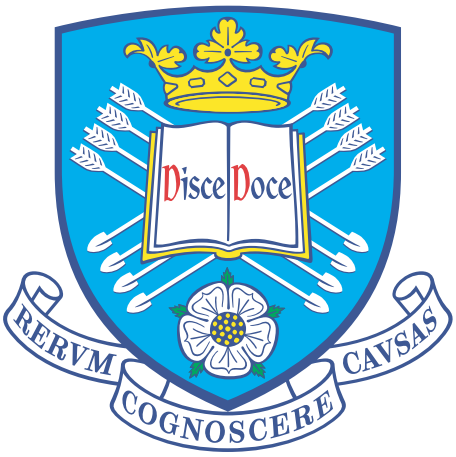

This thesis is submitted for consideration towards the degree of Doctor of Philosophy in

The Department of Civil and Structural Engineering, and Automatic Controls and Systems Engineering

at the University of Sheffield

Saturday 5th March, 2022

Abstract

The detonation of a high explosive results in the rapid release of energy as the explosive charge undergoes a rapid change in state and is converted into a high pressure, high temperature gas. As the gas expands, the surrounding air is displaced, resulting in a high pressure shock discontinuity – a shock wave. As this shock wave propagates away from the charge, it can cause severe damage to any structure that it impacts on. Structural blast engineers are tasked with designing infrastructure in a way that it is robust enough to withstand extreme loading, whilst dealing with several constraints such as time, cost and space. Due to the variability in initiation conditions (such as charge shape, charge location, chemical composition of charge and localised point of detonation), and the subsequent variability in loading produced, it becomes impractical to perform numerical simulations or experiments for all possible scenarios, though an understanding of the loading is required to accurately model structural response. Predictive models are therefore required that can predict the blast load parameters of interest (impulse) given certain input parameters that are fast to run and accurate – predictive models such as these are known as surrogate models.

The blast protection community, when tasked with assessing the viability and safety of structures (the structural response), need an accurate picture of what exactly the loading is. This loading information is composed of both a magnitude and location of a load across a structure, and therefore any predictive approach must predict both these constituent parts of the loading. However, obtaining this loading information, especially within a blast engineering context when the distance between a charge and target is small, is expensive and physical or numerical experiments are costly in both time and money.

Current predictive approaches are severely limited in this regard, in that they do not provide sufficient accurate information, nor are they flexible to handle more than the most simple scenarios. This thesis proposes strategies for surrogate model development in a blast protection engineering context, that allow the rapid evaluation of structural load, given input conditions for a range of scenarios. Furthermore, this thesis demonstrates three applications of strategies that increase the utility of data and knowledge already obtained, that address the fundamental issue of data being expensive to obtain. To achieve this end three approaches are presented: firstly, data transformation procedures, that reduce the dimensionality of the data enabling the use of simpler surrogate models; secondly, the use of directly including known physics into the objective function when model training as a regularisation procedure; and finally, implementing transfer learning by embedding learned knowledge into the architecture of a neural network. These three applications provide statistically significant improvements to model performance and training efficiency, and provide justification to their use in surrogate modelling strategies generally within blast protection engineering.

The results of this thesis should be used to guide surrogate model development for the prediction of peak specific impulse in the near-field for spherical and cylindrical charges. It presents frameworks for creating surrogate models and demonstrates how prior knowledge can be used to improve the performance of surrogate models, or the efficiency when training surrogate models in a new domain, and thereby drastically reducing the need for new data to be obtained. It is shown extensively that machine learning methods can reliably be used in surrogate model development. The findings presented within this thesis have the potential to be implemented into load prediction software which would be of

great utility to the blast protection community and insurance industry.

Acknowledgements

Firstly, I would like to thank my PhD supervisors, Dr. Sam Rigby and Prof. George Panoutsos for their support, guidance and friendship throughout my PhD. The quality of the work owes much to their dedication and attention to detail and I hope we will keep in touch.

Also I would like to thank Prof. Andy Tyas and Prof. Dan Pope, from DSTL, for their input on the direction of this thesis. Thanks also to the rest of the Blast and Impact Dynamics group, I will hold fond memories of the spontaneous trips to the Red Deer!

This research was made possible from the financial support of an EPSRC studentship, to whom I am grateful. Also a thank you to the various academics who have given me teaching and research opportunities that I have welcomed and enjoyed, including (but not limited to!): Sam, Dr Matteo Di Benedetti, Dr Hadi Arbabi, Dr Giuliano Punzo and Dr Peter Rockett.

I would like to thank my friends, Martin and James, and my family for their support. Particularly my Mum and eldest brother, Mat. Living up t'north in Sheffield gave me the opportunity to visit Mat much more regularly, and his support has been invaluable to me – especially the weekly burger and beer at Bungalows and Bears!

I would also like to thank my future parents-in-law, Steve and Jane, for their support throughout. Finally, I would like to thank my fiancée, Lizzie. My decision to move up to Sheffield to study a PhD meant we would be returning to a long-distance relationship. But when I told you I wanted to accept the PhD opportunity, you were wholeheartedly happy for me and have been constantly enthusiastic and supportive. I couldn't have got this far without you and I will miss those weekends spent wandering the peak district with you!

Declaration

I, Jordan James Pannell, certify that all the material contained within this thesis is my own work, except where it is clearly referenced to others.

Saturday 5th March, 2022

Contents

List of Figures

List of Tables

Nomenclature

$\bar{i}$	Scaled slant specific impulse
$\bar{S}$	Slant stand-off distance
$\bar{Z}$	Scaled slant distance
$\Delta\lambda$	Chapman-Jouguet progression rate
Δt	Timestep
Δx	Cell width
$\dot{\epsilon}$	Uniform stretch rate
$\dot{q}_h$	Time rate of heat transfer per unit mass
γ	Specific heat ratio
κ	Thermal diffusivity gradient
λ	Model hyper-parameter
$\lambda_{Phy,1,2}$	PGNN hyper-parameters
$\mathbb{R}$	Set of real numbers
$\mathcal{D}$	Domain
$\mathcal{T}$	Task

$\mathcal{X}$	Feature space
$\mathcal{Y}$	Label space
Ω	Finite control volume
ϕ	Slope limiter
ρ_0	Density of ambient air
ρ_s	Shock Density
ρ_{DP}	Density of detonation products
ρ_{HE}	Density of high explosive
ρ	Density
$\overline{\overline{\tau}}$	Viscous stress tensor
$\overline{\overline{F}}_C$	Convective flux tensor
$\overline{\overline{F}}_D$	Diffusive flux tensor
$\overline{\overline{I}}$	Unit tensor
$\overline{\overline{Q}}_S$	Surface sources tensor
$\mathbf{w}$	Model parameter weights
U	Conservative quantity
θ	Angle of incidence
$\theta_{1,2}$	Chapman-Jouguet trigger functions
ε	SVR allowable loss
$\vec{U}$	Conservative vector quantity

$\vec{F}_C$ Conservative flux vector

$\vec{F}_D$ Diffusive flux vector

$\vec{f}_e$ Body force per unit volume

$\vec{n}$ Outward unit normal vector

$\vec{Q}_S$ Surface sources

$\vec{Q}$ Source term

$\vec{Q}_V$ Volume sources vector

$\vec{v}$ Material velocity

$\vec{W}$ Conservative variables vector

ξ_ε SVR loss function

a_0 Sound speed in ambient air

A_G Gaussian width parameter

A_R Slant distance coefficient

A_S Slant distance coefficient

C_P Heat capacity at constant pressure

C_r Reflection Coefficient

C_V Heat capacity at constant volume

$C_{1,2}$ JWL constants

C_{cfl} Courant-Friedricks-Lewy number

$C_{v1,2}$ Specific heat constants

E	Specific Energy
E	Total energy per unit mass
H	Total enthalpy
h	Ultimate cell size
i	specific impulse
i_r	Reflected specific impulse
i_{so}	Incident specific impulse
k	Chapman-Jouguet sensitivity coefficient
L	Zone length
N	Resolution level
n	Amount of substance
P	Pressure
p	Pressure
p_0	Ambient air pressure
p_r	Reflected overpressure
p_{so}	Incident overpressure
q_s	Dyanmic pressure
Q_V	Volume sources
R	Ideal gas constant
R	Standoff

r	Ratio of successive gradients on CFD solution mesh
$R_{1,2}$	JWL constants
T	Temperature
t_a	Arrival time
t_d	Phase duration
u	x-component of velocity
u_1, u_2	Wave velocities
u_s	Particle velocity
V	Contravariant velocity
V	Volume
v	y-component of velocity
V_r	Contravariant velocity relative to motion of the grid
V_t	Contravariant velocity of face of control volume
V_{burn}	Chapman-Jouguet Local burning velocity
W	Charge Mass
w	z-component of velocity
Z	Scaled Distance
z	Scaled Distance

Chapter 1

Introduction

1.1 Background and motivation

On April 19$^{\text{th}}$ 1995 at 9:02am, a pickup truck containing over 2200kg of explosives detonated in front of the north side of the nine-story Alfred P. Murrah federal building in Oklahoma City, Oklahoma, USA (Irving 1995). As a result, the entire front face of the building collapsed, claiming 168 lives, of which 19 were children. The blast destroyed or caused damage to 324 buildings within a 4 block radius (Oklahoma City Police Department 2007). This event ranks as the second most damaging terrorist attack on U.S. soil with a total of 842 people injured or killed, and costs of damage exceeding 510 million dollars (Zinne & Williams 2013). The use of high explosives for malicious attacks has become more common in recent history, in a similar attack on a Norwegian government buildings in Oslo, on 22nd July 2011, approximately 8 people were killed in the explosion and a further 209 people received physical injuries from the blast and debris (The Guardian 2012).

The problem of blast protection engineering, however, is not just limited to malicious attacks. Several large scale *accidental* explosions

have occurred, such as Hemel Hempstead, U.K. (2005); Tianjin, China (2015) and Beirut, Lebanon (2020), resulting in significant losses of life and large costs incurred from property and business damage. Therefore, the need for civilian infrastructure to be able to resist the intense loading produced from a high explosive event is paramount to ensure the safety of the buildings occupants. Yet, it is not just the engineering community that have been forced to adapt to mitigating these threats. The insurance industry have had to acknowledge that standard insurance policies are not suitable in these circumstances and need to be modelled independently through catastrophe risk modelling.

The loading experienced from blast events typically exceeds the forces and actions a building will be designed to resist during normal operation by several orders of magnitude. Methods to design against these are beyond the scope of traditional civil engineering. Presently, it is common for buildings deemed "at risk" to have had some level of blast security considered during the design process. Typical solutions involve upgrades to the structural frame, façade glazing or cladding and increasing the distance between the traffic and building (through installing bollards for example).

An engineer might need to ascertain whether a column would fail under a certain loading scenario, or an insurer may be interested in estimating the cost and likelihood of any potential damage. But to do this the ability to predict the loading produced by explosive events is key for the design and modelling process. Considering the first application, Rigby et al. (2019b) show that in order to accurately predict the structural response knowledge of both the magnitude and distribution of the loading is critical.

There exist three possible approaches to model blast loading effects:

empirical (experimental), essentially correlations with experimental data; *semi-empirical*, based on simplified models of underlying physical phenomena or *numerical simulation*, methods that use first-principles to describe the basic laws of physics governing a problem. The blast protection community is equipped with some well-established engineering tools, such as the US (DoD 2008) design manual UFC-3-340-02, *Structures to Resist the Effects of Accidental Explosions*, and the ConWep program (Hyde 1991) that allows for rapid evaluation of blast wave parameters from a set of given explosive events. But these methods are unsuitable when considering explosives located extremely close to a structure, where complex interactions between the explosive and target result in, spatially, highly non-uniform loading.

Experimental analysis is the most reliable method of measuring blast loading, but it can be costly and impractical, especially if a large number of studies are required. Numerical simulation offers the potential to produce highly accurate results, but can be considerably expensive also in time and cost. Expensive, deterministic methods such as these also raise the question of *how useful is the analysis?* An engineer may be investigating the loading imparted on a structural column, but how do they decide on the input parameters? Is the charge likely to be a perfect sphere or some other shape? What would its mass be? How far away from the structure is it likely to be? What are the chances of this event even occurring – therefore is it reasonable to design against it?

Clearly, there are a large number of uncertainties and variables that influence the blast loading, and there are often not clear definitive answers to the questions raised previously. Therefore there is a need for a mode of assessment that can rapidly assess a large amount of different input scenarios, a "probabilistic" mode of assessment. Data-driven mod-

elling approaches offer the ability to build a surrogate model, trained on valid experimental or numerical data, that encompasses a broad range of input scenarios and that can provide an (almost) instant output which would enable such a probabilistic framework, provided that distributions of the input parameters are assumed. It would be highly useful, particularly at the earlier design stage where significant cost savings can occur. This thesis demonstrates the development of such a deterministic tool (the surrogate model) that maps the input to output, and also presents surrogate modelling strategies that increase the utility of data and knowledge already obtained which addresses the issue of operating data-driven models in a sparse-data environment. It is the hope of this author that an awareness of blast-resilient design can be built into the fabric of our cities.

1.2 Scope and objectives of thesis

This thesis aims to develop a suite of rapid predictive methods for near-field blast load predictions, leveraging data from high-fidelity numerical modelling, and incorporating cutting edge data-driven machine learning techniques that can increase the utility of data already obtained and reduce the requirement for new data to be generated. This thesis, therefore, has the following main objectives, related to the engineering challenges introduced above:

1. To review the current literature on predictive approaches for blast loading and discuss the limitations of existing approaches.

2. To establish a reliable numerical framework for generating loading distributions that has been validated against near-field and far-field experimental data.

3. To rigorously assess the sensitivities of numerical approaches for mesh effects and develop consistent rules for mesh sizing to be used throughout the thesis study.

4. To investigate data transformations that can be applied to datasets as a whole that reduce the dimensions of the dataset and allow the rapid development of surrogate models.

5. To investigate physics-guided machine learning approaches in a blast loading domain through incorporating known, or learned, physical knowledge directly into the objective function as a physics-based regularisation procedure.

6. To further investigate transfer learning approaches from previously trained surrogate models applied to new datasets, to determine whether knowledge learned from previous tasks and domains can be applied in a new task and domain and improve the efficiency in training, and performance of, the new surrogate models.

1.3 Thesis outline

A probabilistic framework of blast assessment would consist of a pipeline beginning with input parameters and ending with structural response. To achieve a probabilistic framework a fast-running method needs to be defined that maps a broad range of inputs to an output, this will allow distributions of input parameters to be assumed which could be rapidly evaluated. This thesis addresses the first part of the probabilistic pipeline: the relationship between sampled input parameters and the explosive loading (output) that they produce – and is therefore deterministic in nature. The second stage would be then to link the loading to the

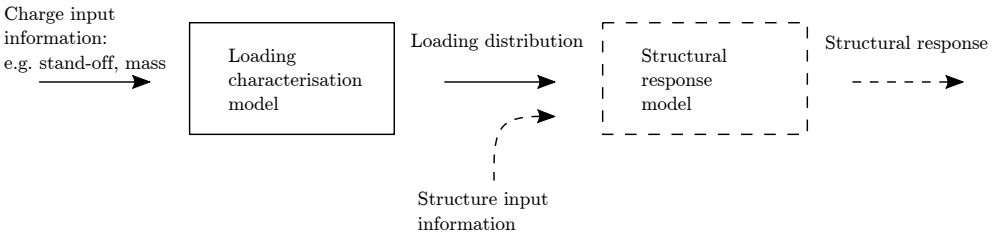

Figure 1.1: Proposed probabilistic framework achievable with fast-running data-driven surrogate models. Solid lines indicate the focus of this thesis. Charge input information would be sampled from an assumed distribution to account for the uncertainty in input parameters.

subsequent structural response, but that challenge is outside the scope of the thesis. The high level schematic of this framework is presented in Figure 1.1.

In Chapter 2, background information for explosive blast loading and current data-driven modelling approaches in this domain are provided, as well as a discussion on the limitations of such current approaches.

Chapter 3 demonstrates how the blast loading data is generated numerically. It provides a discussion of how computational fluid dynamics (CFD) software functions generally, as well as discussing how the modelling techniques in the chosen software can be used, such as adaptive mesh refinement techniques and domain resizing. Additionally, the chosen CFD software is used to generate numerical data for far-field explosive scenarios and is compared against experimental data in a validation exercise.

The investigation in Chapter 4 is concerned with developing a preliminary surrogate model for near-field explosive scenarios. Initially, the chosen CFD software is validated in the near-field against experimental data, and then a numerical dataset is generated to build a surrogate model from. A novel data transformation is found that reduces the dimensionality of the dataset, allowing a regression analysis and produces a novel predictive equation (the surrogate model) to be derived that can

predict the explosive loading for a range of input parameters and is compared to unseen numerical and experimental data.

The development of the surrogate modelling approach in Chapter 4 is suitable provided a data transformation that can remove a dimension of the data within the dataset exists or is known. When this is not the case, more complex data-driven methods are required that can handle this complexity. Instead of using such methods as unknowable "black-box" models, Chapter 5 investigates how to incorporate prior domain knowledge into the model development process and provide interpretable machine learning that the user trusts and understands. Through incorporating prior knowledge of physics it is shown that the amount of data required for effective data-driven models can be reduced. To achieve this end, a monotonic loss constraint is added into the objective function of a neural network due to specialist prior knowledge of the problem domain as a physics based regularisation procedure. The results are compared to a traditional neural network architecture, and some other black-box models as a benchmark and stress-tested through various approaches of restricting the available data to evaluate its generalisation ability.

To increase the capabilities of a predictive surrogate model, a large range of parameters need to be investigated. The final investigations in Chapter 6 investigates how prior data and domain knowledge can be leveraged to speed up learning or model performance in a new domain or task through implementing transfer learning, especially useful when data may be sparse or expensive to obtain as it is in blast engineering. In the first part, the discovered data transformations presented in Chapter 4 are utilised to rapidly develop similar surrogate models for 5 further explosive types. In the final part, a new numerical dataset (produced from cylindrical charges) is created and two different neural network modelling

approaches are compared. The first approach models the new numerical data (produced from cylindrical charges) with no transfer learning implemented whilst the second model uses a neural network architecture trained on numerical data obtained from spherical charges (from Chapter 5) and then trained again on the new numerical dataset produced from cylindrical charges. These two implementations are stress-tested at various levels of data restriction and their performances compared to assess the viability of this application of transfer learning.

1.4 Published work

The work detailed in this thesis has been published in peer-reviewed academic journals or presented in academic conferences and is listed below:

- Pannell, Rigby, Panoutsos, Tyas, Cooke & Pope (2019), Predicting near-field specific impulse distributions using machine learning, *in* '18^{th} International Symposium on Interaction of the Effects of Munitions with Structures (ISIEMS18), Panama City Beach, Florida, USA.

- Pannell, Rigby & Panoutsos (2020) A physics guided machine learning approach to understanding loading distributions from explosive events, *in* Young Researchers Conference 2020, Institution of Structural Engineers, London, U.K.

- Dennis, Pannell, Smyl & Rigby (2021), 'Prediction of blast loading in an internal environment using artificial neural networks', *International Journal of Protective Structures* 12(3), 287–314.

- Pannell, Panoutsos, Cooke, Pope & Rigby (2021), 'Predicting specific impulse distributions for spherical explosive in the extreme

near-field using a gaussian function', *International Journal of Protective Structures* 12(4), 437–459.

- Pannell, Rigby & Panoutsos (2022a), 'Physics-informed regularisation procedure in neural networks: an application in blast protection engineering', accepted for publication in *International Journal of Protective Structures*

- Pannell, Rigby & Panoutsos (2022b), 'Application of transfer learning for the prediction of blast impulse', accepted for publication in *International Journal of Protective Structures*

Chapter 2

Literature review and theoretical background

2.1 Introduction

This chapter provides background information for explosive blast loading and current techniques for predicting this loading, including specific machine learning applications. The main focus of this chapter is to provide theoretical knowledge of shock wave formation, and the subsequent structural loading that it produces. Additionally, discussion on current predictive methods for blast wave loading is provided. This chapter highlights gaps in the knowledge base that this thesis aims to address.

2.2 Explosive blast loading

2.2.1 Overview

An explosion is the result of rapid expansion in volume of gases occurring from a physical, or mechanical, change in a material. This rapid expansion is associated with a large release of energy into the surround-

ing medium. The explosive blast loading process can be considered in these phases:

- Initial energy release from detonated material
- Disturbance of molecules within the air between explosion source and structure
- Interaction of air molecules with the structure

2.2.2 Detonation and air shock formation

Generally, explosions can be classified into three categories: physical, chemical and nuclear. Though the physical impact on the surrounding medium is similar, the detonation process is different. Physical explosions are those that are not caused by any chemical reactions, an example would be the impact of a a high velocity meteor into the earth's atmosphere. Chemical explosions release large amounts of energy into the environment through the oxidation process. In typical chemical explosions, the explosive material experiences a rapid change in state and the oxygen required is present within the explosive material itself, resulting in the detonation process (see Figure 2.1). Nuclear explosions are outside the scope of this thesis and are not typically considered within the remit for civilian blast protective engineering.

Figure 2.1 shows the mechanism of detonation. A high velocity wave is instigated at the point of initiation by a detonator, which compresses the surrounding material. As the surrounding material is oxidised, the exothermic reaction begins that increases the density, temperature and pressure behind the wave front. This pressure gradient causes the wave front to accelerate away from the initiation point until it reaches its detonation velocity, typically between 6000-8000 m/s.

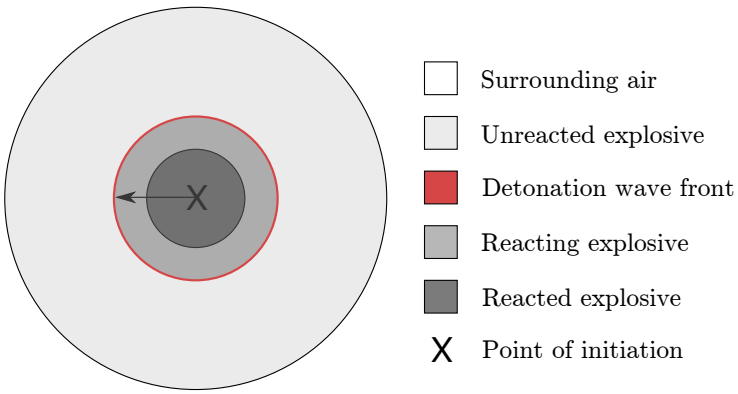

Figure 2.1: Detonation mechanism within a spherical explosive.

The detonation product gases can be at pressures between 10-30 GPa, and at a temperature of 3000-4000°C (Cormie et al. 2009). As the wave front propagates it is continually reinforced by the oxidation of the unreacted explosive ahead of it until the boundary between explosive and surrounding medium is reached. At this point, the explosive material then begins to expand and displace the surrounding air at supersonic speeds.

Air is a compressible material and the wave velocity is dependent on the pressure of the detonation products. Wave velocity is directly proportional to the gradient of the stress-strain relationship for the material it is reacted within[i] (in this case, air) shown graphically in Figure 2.2, where the gradient is greater at p_2 than at p_1 and where u_1 and u_2 are representative of wave velocities at each point.

Therefore, after an air disturbance has been formed, higher pressure components of the pressure disturbance will travel at a greater velocity relative to the lower pressure components, as shown in Figure 2.3. At a distance away from the point of detonation, there will be a point where the highest pressure (therefore the fastest) component will have

[i] When considering fluids, it is typical to refer to pressure-volume curves which are analogous to this.

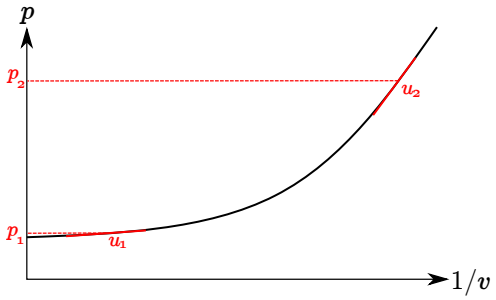

Figure 2.2: Pressure-volume relationship of air and corresponding wave velocity.

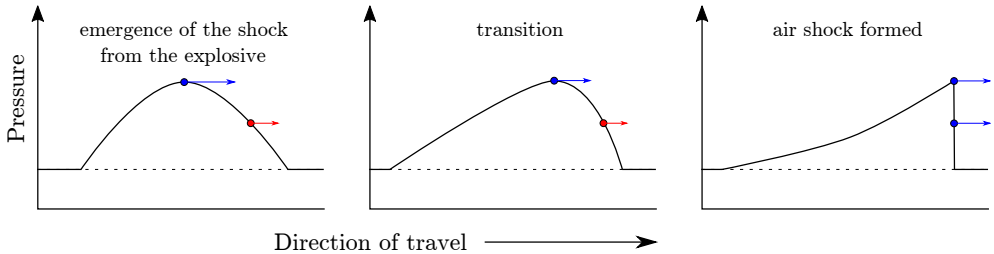

Figure 2.3: Development of explosive shock, adapted from Kinney & Graham (1985)

migrated to the front of the disturbance. This results in a discontinuous step change in pressure, temperature and density. This discontinuity represents the shock front (Baker 1973). Once the detonation process is completed, the total energy within the system remains constant. The volumetric expansion of the blast is associated with an increase in energy absorbed by the surrounding air. This causes the blast wave front to decrease in pressure and density as it expands away from the point of initiation.

2.2.3 Blast waves in free air

As air is a compressible material, the pressure disturbance caused by the detonation of an explosive material creates a shock front, the near discontinuous increase in pressure and density travelling away from the point of initiation. A blast wave is therefore a special form of stress wave:

a longitudinal wave in a fluid that propagates faster than the local speed of sound and therefore "shocks up".

If a point at a fixed distance away from a detonation is considered, the pressure at that point will experience a sudden increase in pressure from ambient air pressure p_0 to peak overpressure $p_{so,max}$ followed by a temporal decay back to ambient pressure. The "so" subscript represents incident pressure values, the pressure that is measured by a transducer that offers no resistance to flow behind the shock front. It is also common in blast engineering to refer to pressures as "overpressures", where overpressure refers to the pressure increase above normal, atmospheric conditions (101 kPa).

Subsequently after peak overpressure, $p_{so,max}$, is reached there is a pressure decay to ambient pressure, with duration t_d followed by a period of negative overpressure. Negative overpressure is a result of the expansion of air following the shock front (and the compression of air that results) with peak magnitude $p_{so,min}$ and duration t_d-. After the negative phase of overpressure, ambient pressure is restored. Figure 2.4 presents this idealised pressure time profile, known as the Friedlander waveform. If the pressure-time history is integrated with respect to time, the specific impulse-time history is obtained, where specific impulse, i, is the change in momentum (therefore the loading) a given area experiences.

2.2.4 Blast wave reflection

When the unimpeded blast wave (known as the incident blast wave) impinges upon a target, conservation of mass, momentum and energy at the collision interface result in changes in pressure, density and temperature of the blast wave. Overpressure and specific impulse at this surface are defined with subscript values r. The majority of blast protection engi-

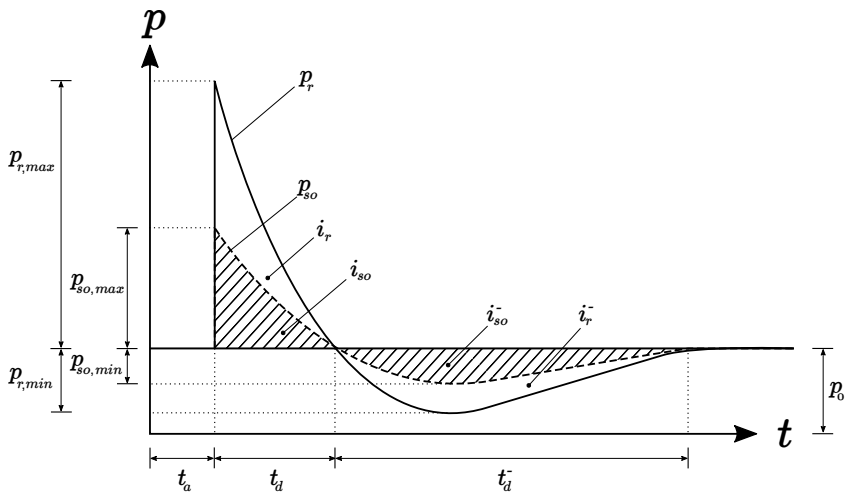

Figure 2.4: Friedlander waveform.

neering is interested in the reflected components of the blast wave as this constitutes the structural loading.

The Rankine-Hugoniot 'jump' conditions describe the relationship between states of compressed air at either side of the shock wave (Anderson 2010). Conserving mass, momentum and energy across the discontinuity the particle velocity, u_s (immediately behind the shock front) and density ρ_s can be calculated:

$$\rho_s = \rho_0 \frac{(\gamma+1)p_{so} + 2\gamma p_0}{(\gamma-1)p_{so} + 2\gamma p_0} \tag{2.1}$$

$$u_s = p_{so} a_0 \sqrt{\frac{2}{\gamma p_0[(\gamma+1)p_{so} + 2\gamma p_0]}} \tag{2.2}$$

Here, ρ_0, p_0, a_0 are density, pressure and sound speed of undisturbed air respectively. γ is the specific heat ratio (usually 1.4 for air at atmospheric conditions). For normal reflection, the reflected pressure p_r, can be expressed in terms of the incident pressure, p_{so} and dynamic pressure, q_s

$$p_r = 2p_{so} + (\gamma + 1)q_s \tag{2.3}$$

The first term in Equation 2.3 is the acoustic term, relating to the reflection of the incident pressure in the acoustic regime whilst the second term represents the pressure increase that occurs when bringing the compressed fluid to rest at the reflecting surface, given by

$$q_s = \frac{1}{2}\rho_s {u_s}^2. \tag{2.4}$$

By substituting Equations 2.1 and 2.2 into Equation 2.4 and then into Equation 2.5 leads to

$$p_r = 2p_{so}\frac{7p_0 + 4p_{so}}{7p_0 + p_{so}}. \tag{2.5}$$

The reflection coefficient, C_r, is defined as the ratio of reflected pressure to the incident pressure, $C_r = p_r/p_{so}$. For weak shocks, where the overpressure is small in comparison with atmospheric pressure ($p_0 >> p_{so}$), the reflection coefficient, $C_r \approx 2$. In strong shocks, where $p_{so} >> p_0$ the reflection coefficient reaches the upper limit of $C_r = 8$ as the reflection is dominated by the dynamic term. However, it is important to consider that Equation 2.5 assumes air behaves as an ideal gas at extremely high pressure and temperatures. An ideal gas is one that obeys the ideal gas law, often written as

$$PV = nRT \tag{2.6}$$

where P, V and T are pressure, volume and temperature respectively, n is the amount of substance (in moles) and R is the ideal gas constant. If real gas effects such as dissociation and ionization of the air molecules are taken into account, the reflection coefficient can be as high as 20 (Baker

1973).

2.2.5 Scaling laws

Hopkinson-Cranz (or 'cube-root') scaling, proposed independently by Hopkinson (1915) and Cranz (1926) is a highly useful tool used in the blast protection community that certain parameters for a given scenario can be scaled to another scenario, provided similitude is obeyed. This states similarity exists between blast waves produced at identical scaled distances from two explosive charges of the same geometry but different masses. Therefore, the blast pressure generated at a distance R (stand-off) from an explosive mass W will be similar to the blast pressure created at a distance KR from a mass of K^3W. Shown schematically in Figure 2.5.

The concept of *scaled distance, Z*, with units $m/kg^{1/3}$, is expressed mathematically as

$$Z = \frac{R}{W^{1/3}}, \tag{2.7}$$

where W is the TNT equivalent charge mass in kilograms. TNT (trinitrotoluene) is widely used as the "reference" explosive for historical reasons, due to being the standard explosive at the time of the Kingery & Bulmash (1984) compilation. To establish the scaled distance of blast waves from sources other than TNT, the TNT equivalence must be established. At present, there is no universally accepted way of calculating TNT equivalence (Cormie et al. 2009), however the most simple is to apply a TNT equivalency factor. There is substantial evidence however that TNT equivalence factors are dependent on scaled distance (Shin et al. 2015).

In Hopkinson-Cranz scaling, pressures are identical between scaled and actual values, and all times are scaled by the same value as the length

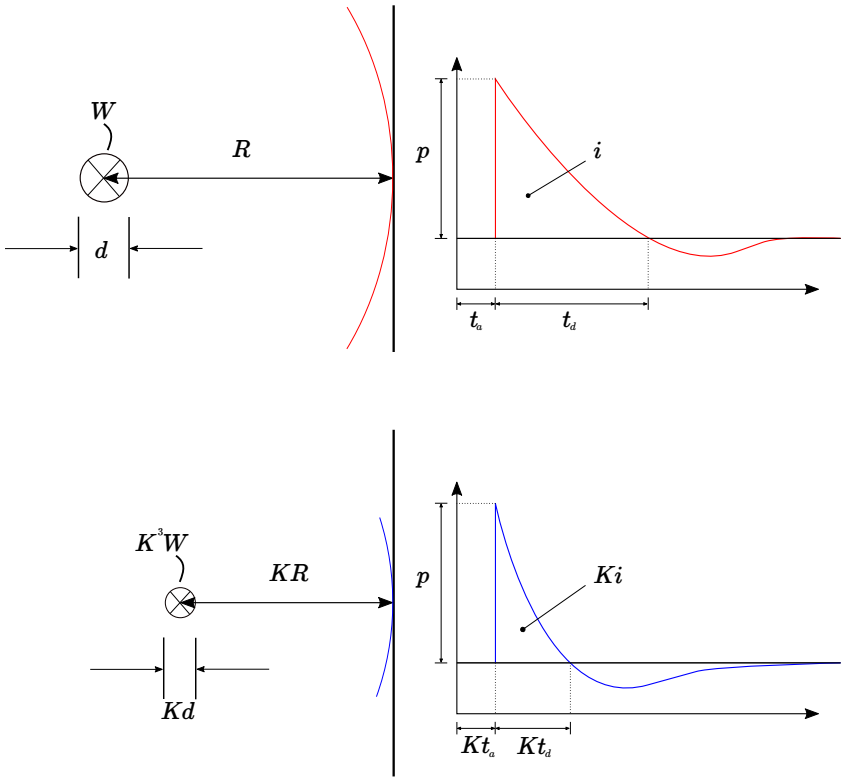

Figure 2.5: Hopkinson-Cranz blast wave scaling

scale factor, K, or the cube-root of the charge mass (Kinney & Graham 1985). To summarise:

$$p_{actual} = p_{scaled},$$
$$t_{actual} = t_{scaled} W^{1/3},$$
$$i_{actual} = i_{scaled} W^{1/3}.$$

2.2.6 Empirical and semi–empirical predictive methods in blast engineering

To provide predictions of structural loading from an explosion, positive blast wave parameters need to be quantified as a function of a given input. Notable early research for blast wave parameter prediction by Brode (1955), Taylor (1950), Granström (1956), are based on several assumptions (such as instantaneous energy release, negligible atmospheric pressure, point source explosive), and are not directly applicable for use in design and research of structures subject to explosive loads (Rigby 2014).

Quantification of blast loading became a particular topic of interest to the engineering community in the years following the second world war. One such study by Kingery (1966), later revisited by Kingery & Bulmash (1984) – hereafter abbreviated as KB – compiled existing blast pressure measurements and used geometric scaling laws, formulated independently by Hopkinson (1915) and Cranz (1926), to develop a semi-empirical tool for predicting normally reflected and incident blast wave parameters for hemispherical surface bursts and spherical free air bursts. Relationships are provided from effectively contact explosions, out to $Z \approx 40$ m/kg$^{1/3}$. However, as noted by Esparza (1986, p. 2), direct measurements of blast wave parameters in close proximity to the charge were either "*non-*

existent or very few", and near-field semi-empirical predictive data were inferred from non direct measurements such as smoke trails (Dewey 1964) or rudimentary numerical analyses (Brode 1955).

The possible explosive scenarios that can be modelled by the KB equations are spherical, free air bursts or hemispherical surface bursts, whereby a hemisphere is placed on a rigid, perfectly reflecting surface. The high-order polynomial curve fits are ungainly, and are therefore often summarised as a set of curves, shown in Figure 2.6. The blast parameters presented in Figure 2.6 form the basis of design guidance such as DoD (2008) and the computer code Hyde (1991).

Examples of further simplified relationships for calculating blast wave parameters are given in Baker (1973), Kinney & Graham (1985), Swisdak (1994). For a comprehensive review of current methods for predicting blast wave parameters, detailed discussion is provided by Remennikov (2003) and Shin et al. (2014b).

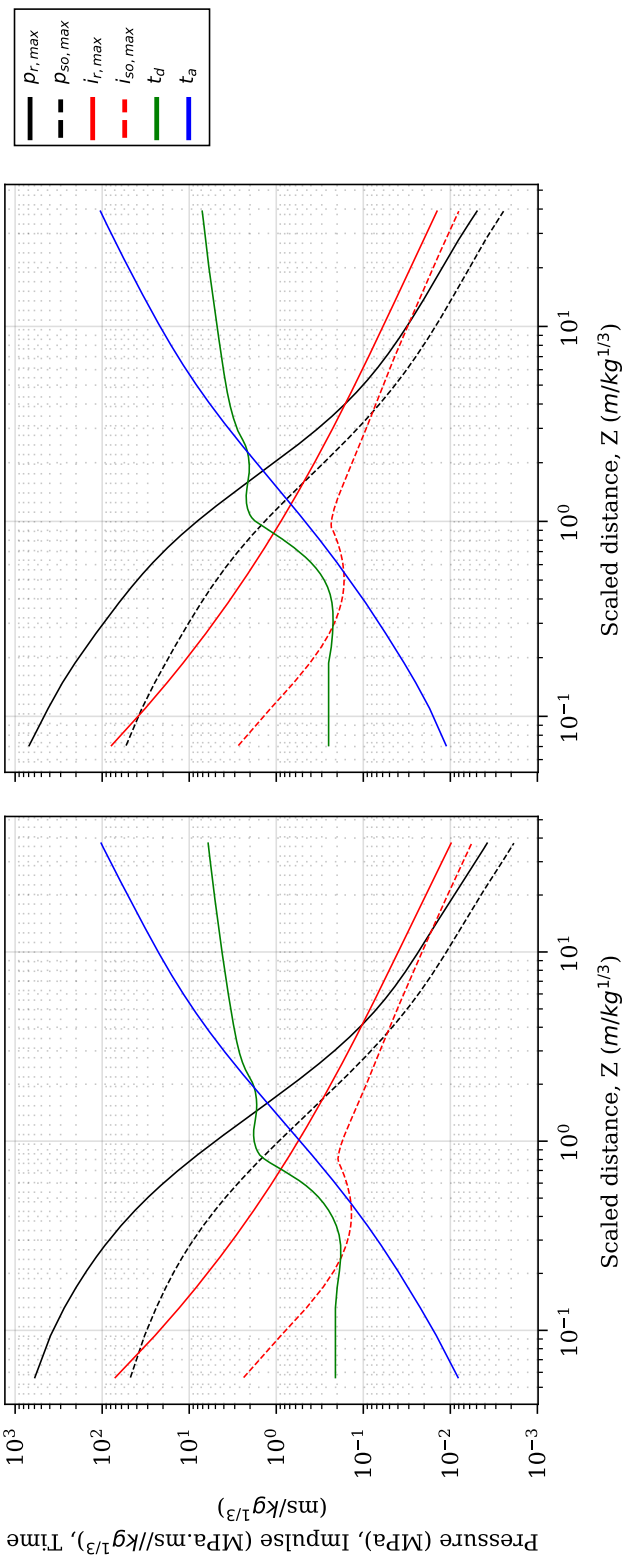

Figure 2.6: Incident and positive phase blast wave parameters from UFC-3-340-02 (DoD 2008). Spherical charges of TNT in free air (left) and hemispherical charges of TNT on the surface (right).

2.2.7 Near-field blast loading

When a blast wave strikes a target, the pressure imparted from the blast wave can be in the order of several Megapascals acting over sub-millisecond durations. If the source of the blast is located close to the target, the interaction between the blast wave and structure, and therefore the imparted load, is much more complex than in far-field events.

The KB semi-empirical predictions, that are built from the compilation of a number of large-scale blast trials, parametric studies and early numerical simulations have been shown to agree with closely controlled experimental tests for medium to far-field scaled distances (Rigby et al. 2014a, Rickman & Murrell 2007, Tyas et al. 2011, Cheval et al. 2010, 2012). When reviewing the research produced by Bogosian et al. (2002), Shin et al. (2014b, p. 179) state *"the lack of direct measurements of overpressure, impulse, arrival time and positive phase duration in the near-field, as described in prior sections of this chapter, and the significant discrepancy at the face of the charge noted above, calls into question the accuracy of the KB charts, especially in the near-field"*. Luccioni et al. (2006, p. 8) state *"empirical expressions are not applicable with confidence in the near-field because of the complexity of the flow processes involved in forming the blast wave"* or that *"accuracy of empirical relations in the near-field is not guaranteed"* (Luccioni et al. 2006, p. 4), citing work by Hetherington & Smith (2014). As a possible explanation for the large variations in loading obtained in the late near-field ($0.5\ \mathrm{m/kg^{1/3}} < Z < 2\ \mathrm{m/kg^{1/3}}$), Tyas (2019) suggests that it is due to the development of Rayleigh-Taylor (RT) and Richtmyer-Meshkov (RM) instabilities in the expanding fireball. Additionally, when using the KB semi-empirical predictions for explosives other than TNT, the choice of TNT equivalence factor is crucial. There is, however, no universally ac-

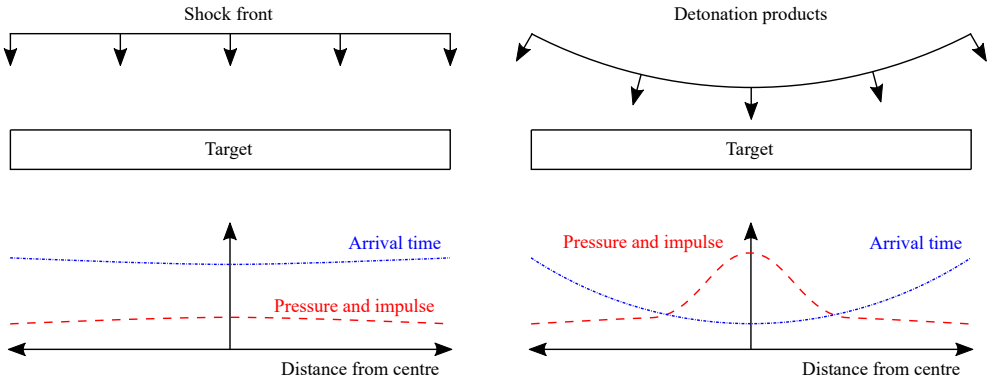

Figure 2.7: Blast interaction diagram and loading parameters associated with: far-field (left) and near-field (right). Figure adapted from Rigby et al. (2014c)

cepted way of calculating TNT equivalence factors (Cormie et al. 2009) which has been shown to be dependent on scaled distance (Shin et al. 2015).

Near-field blast events consist of extremely high magnitude, highly spatial and temporally non-uniform loads (Rigby et al. 2015c) and are generally associated with scaled distances of $Z < 1$ m/kg$^{1/3}$, or when the shock wave has not "detached" from the explosive fireball. A comparison between an example far-field blast wave and near-field blast wave striking a target is shown schematically in Figure 2.7, where the complexity of near-field loading is demonstrated.

In this near-field region, experimental methods can be used to directly measure near-field reflected specific impulse and specific impulse distributions. Close-in experimental blast parameter measurements are typically conducted by measuring residual momentum – therefore the imparted impulse – of a small, rigid metal plug embedded within a larger target surface (Huffington & Ewing 1985, Nansteel et al. 2013). The "impulse–plug" method gives an indication of the imparted load at discrete points on the target and is shown to reduce test-to-test scatter. However, this approach does not provide sufficient temporal resolution

of the applied blast overpressure, a necessary requirement to ascertain the mechanism of near-field loading scenarios and to provide validation data (Rigby et al. 2015c). In recent times, large scale experimental approaches for the direct measurement of spatial and temporal variation in loading resulting from an explosive event has been developed (Rigby et al. 2014c, Fay et al. 2014, Clarke et al. 2015). In this approach, a fixed target plate through which Hopkinson pressure bars (HPBs) are inserted such that the bottom of the bars are flush with the plate. Due to the location of these bars on the plate, overpressure-time histories can be obtained for the array of bars that provide temporal and spatial distribution of blast overpressure acting on a rigid target (a similar method with flush-mounted pressure gauges is presented by Aune et al. (2016)). However, near-field blast load measurements are onerous for a number of reasons. Firstly, the high magnitude of loading necessitates the use of robust support structures and protective housing for sensitive equipment (which itself is required to record in the MHz frequency range). Secondly, the measurements themselves are highly variable owing to the presence of surface instabilities in the early stages of expansion of the detonation products (Rigby et al. 2020), this variability in test repeatability means the semi-empirical predictions are less accurate when used deterministically (Bogosian et al. 2002).

It is therefore not practical to develop a predictive approach based solely on physical testing, however experimental data remains a fundamental requirement for validation of numerical modelling schemes. Finite element (FE) and computational fluid dynamics (CFD) approaches have been the primary tool of previous researchers when investigating air-shock propagation and shock-structure interaction in the near-field, with the results generally demonstrating good agreement with experi-

mental data where it is available (Shin et al. 2014a, Rigby et al. 2018, Whittaker et al. 2019). From these studies, authors have researched the complex interaction between shock wave and expanding detonation products (Edwards et al. 1992), the limits of representing the high explosive as an ideal gas (Wilkinson et al. 2013), scaled distance relationships for near-field explosions (Cormie et al. 2013), mesh sensitivity effects Hashemi & Bradford (2014) and the complex 3D waveform of an expanding shock wave (Shin et al. 2014a). Despite research into near-field blast loading currently being hampered by a lack of well-controlled experimental validation data (Tyas 2019), FE/CFD analyses can provide data at considerably higher spatial and temporal resolution than experimental studies and are therefore suitable tools with which to develop a refined predictive approach. However, physics-based models have a relatively high computational demand, and are unsuitable for probabilistic, risk-based analyses.

An outstanding challenge for the blast protection community is the provision of techniques for predicting near-field specific impulse distributions in a way which is both accurate and computationally inexpensive. An appropriate technique, which is adopted in this thesis, is to use validated CFD analyses to create datasets from which predictive methods can be developed. This process, known as data-driven surrogate modelling, is widely used in the field of computational dynamics, e.g. Loy et al. (2017), Both et al. (2019), but has not yet been successfully applied in the field of near-field blast load prediction.

2.2.8 Angle of incidence effects

Some common assumptions in semi-empirical predictive methods are that the blast wave impinges normally on a reflecting surface, in other words,

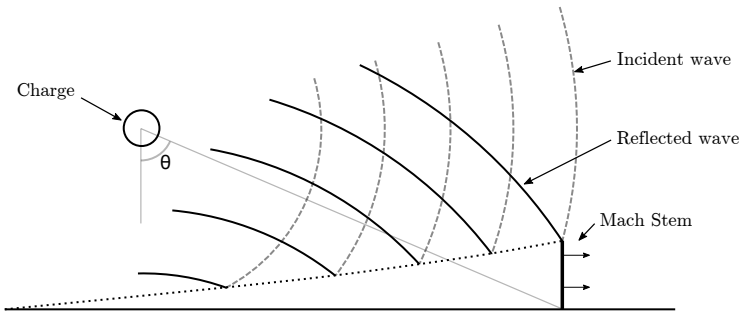

Figure 2.8: Mach Stem formation, recreated from Tyas & Rigby (2018).

the angle of incidence is 0°. The angle of incidence of a point on a surface is defined as the angle between the outward normal and the direct vector from the explosive charge to that point. However, an important phenomenon occurs at angle of incidences of approximately $\theta > 40°$ known as "Mach reflection" which illustrates the importance of considering angle of incidence effects in predictive approaches.

Consider a detonation of an explosive charge at an arbitrary distance from a rigid reflecting surface. For angles of incidence of approximately $\theta > 40°$, upon reflection the incidence wave no longer 'bounces' off the reflecting surface (as it typically does at smaller angles of incidence), but instead it 'skims' across the reflecting surface. This reflected wave will propagate with a large component of velocity parallel to the reflecting surface, as it coalesces with the incoming incident wave, a third wave is formed known as the "Mach Stem". This process is shown in Figure 2.8.

The point where the mach stem, incident wave and reflecting wave intersect is known as the "triple point". Between the surface and the triple point, the Mach Stem forms a vertical wave which can result in large amplifications in the oblique pressure which often exceed the normally reflected pressure at that point (Kinney & Graham 1985).

Angle of incidence is known to affect the pressure and impulse differently (Rigby et al. 2015a). A common method to account for angle

of incidence effects on positive phase parameters involves calculating the incident pressure at scaled slant distance, where slant distance is the distance from explosion to point of interest, and multiply this by the oblique reflection coefficient as shown in Figure 2.9. Remennikov (2003) notes that the recommended minimum scaled distance range for the reflection coefficients in Figure 2.9 is $Z_{\min} = 1.2\text{m/kg}^{1/3}$, so only coefficients corresponding to peak side-on pressures of below 1000 kPa are applicable. The scaled oblique impulse can then be calculated using Figure 2.9 for a particular value of incident overpressure.

Alternative semi-empirical approaches to calculate oblique impulse were investigated by Henrych (1979), later developed by Remennikov et al. (2017), based on the model of instantaneous detonation of a spherical explosive and subsequent expansion of the gaseous detonation products into a vacuum. In this scenario the loading is assumed to develop entirely from momentum transferred by the detonation products as they impact the loaded face. After cancelling constant terms, the Henrych (1979) equations can be expressed in normalised form:

$$\bar{i}(\theta) = \cos^4 \theta \tag{2.8}$$

where $\bar{i}$ is specific impulse divided by normally reflected specific impulse, and θ is angle of incidence as previously.

Randers-Pehrson & Bannister (1997) present a trigonometric expression to account for angle of incidence effects, which after some manipulation can be presented as a function of normalised specific impulse and angle of incidence:

$$\bar{i}(\theta) = A_R(\theta, Z)\cos^2 \theta + A_S(\theta, Z)[1 + \cos^2 \theta - 2\cos\theta] \tag{2.9}$$

Here, $A_R = i_r(\bar{Z})/i_r(Z)$ and $A_S = i_s(\bar{Z})/i_r(Z)$, where i_s is incident specific impulse, Z is normal scaled distance as previously, and $\bar{Z}$ is *slant* scaled distance to the point of interest: $\bar{Z} = Z/\cos(\theta)$. Similarly, slant stand-off distance can be defined as $\bar{S} = S/\cos(\theta)$. For a given normal scaled distance, A_R and A_S vary non-linearly with slant distance and hence these coefficients must be calculated separately for each and every angle of incidence under consideration. Furthermore, this relationship is also scaled-distance dependent.

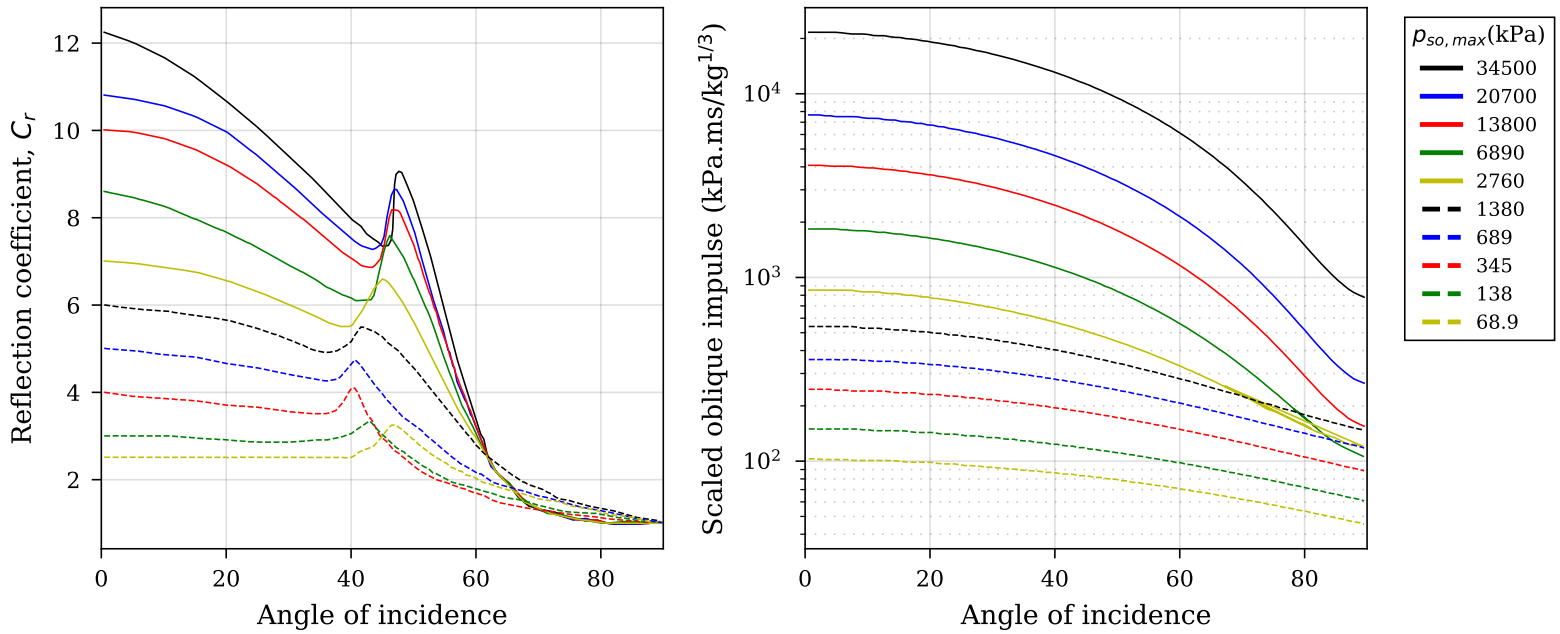

Figure 2.9: Reflection coefficient vs angle of incidence (left) and oblique scaled impulse vs angle of incidence (right). Figure adapted from Figures 2-193 and 2-194(a) in DoD (2008) and converted from psi in this Figure, digitized using software by Rohatgi (2020).

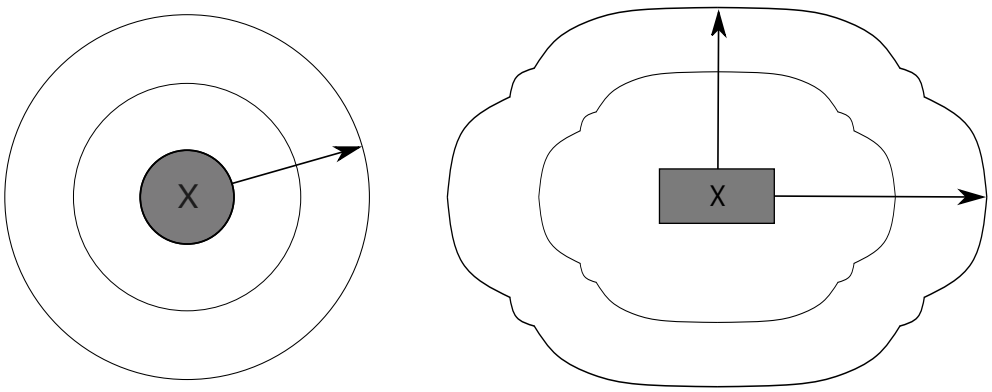

Figure 2.10: Early stage detonation comparison between sphere (left) and cylinder (right). The detonation products from the spherical charge expand spherically, whilst the detonation products from the cylindrical charge expand axially and radially, with a "bridging wave" (Knock & Davies 2013) formed in between. Adapted from Wisotski & Snyer (1965).

2.2.9 Charge shape effects

Existing methods for predicting blast loading on a structure, such as those mentioned previously based on the analytical work of Brode (1955) or the semi-empirical relations of Kingery (1966) and Kingery & Bulmash (1984) are based on the assumption that the explosive is a sphere detonated in free air or a hemi-sphere placed on the ground. However, it has been shown that the shape of the explosive charge has a considerable impact on the expansion of the detonation products and characteristics of the resulting blast wave (Rigby et al. 2021, Langran-Wheeler et al. 2021, Fan et al. 2022).

As a spherical explosive detonates, the detonation wave front reaches all edges of the explosive simultaneously and the blast wave will then expand outward into the air in a uniform and symmetric way. This is not true for a non-spherical charge such as a cylinder however, in this situation the detonation wave front will not reach all edges of the explosive simultaneously and therefore the subsequent expansion of the detonation products will be highly complex and non-uniform. The detonation of a

cylinder results in two types of primary waves: one set emitted from each end of the cylinder that expand axially and another set emanating radially from the curved surface. These two sets of wave travel in directions orthogonal to each other, and at the point where they meet, a secondary "bridging wave" (Knock & Davies 2013) is formed as shown in Figure 2.10. Due to this non-uniform blast wave formation, it is well known that peak overpressure and specific impulse are dependent on azimuth angle (Clutter & Stahl 2014).

The initial shape and subsequent expansion of a blast wave from a cylindrical explosive is largely dependent on its aspect ratio (L/D), i.e. the ratio between the length and diameter of the charge (Esparza 1992). The influence of aspect ratio on blast loading has been investigated in (Stoner & Bleakney 1948, Plooster 1977, 1982) and more recently in (Wu et al. 2010) where it was observed that for large values of L/D most of the energy was focussed in the radial direction, whereas for small values of L/D most of the energy was focussed in the axial direction. Similar observations were reported by (Sherkar et al. 2016, Artero-Guerrero et al. 2017, Langran-Wheeler et al. 2017, 2019).

2.3 Data modelling and machine intelligence

2.3.1 Overview

Broadly speaking data modelling and machine intelligence is concerned with creating representations of the world from observations; learning by induction in order to understand, predict or provide decisions. Of this, there are four elements of data modelling and machine intelligence (Harrison 2018):

- *the assumption*

  - "What we think the world is like".
  - Where $y_{\text{true}} = f_{\text{true}}(x)$ demonstrates some unknown process of input x. When this ouput variable, y_{true}, is observed, some observation error, ϵ_o, is introduced as $y_o = y_{\text{true}} + \epsilon_o$ where y_o is the observed output variable.

- *The model*
  - "A way of expressing the thought mathematically".
  - $f_{\text{true}}(x) \approx f(x)$, where f is the model, or approximator. So $y_{\text{true}} = f_{\text{true}}(x) = f(x) + \epsilon_m$ where ϵ_m is model error.
  - $f(x)$ has adjustable parameters, w.
  - Some typical models include: neural networks (multi-layer perceptron, radial basis function networks, functional link networks), kernel machines (support vector machines), basis functions (sine waves, polynomials, splines).

- *The inference paradigm*
  - "A framework of matching the model to the world".
  - Process of estimating model parameters (such as least squares, maximum likelihood, point estimation). Concerned with finding the 'best' values for w.

- *The inference engine*
  - "A means of doing the matching".
  - Specific method of optimisation, typically branched into the following methods: direct methods (calculus such as linear least squares), numerical methods (gradient descent) or simulation based (genetic algorithms, particle swarms).

When building predictive models, some assumptions about the data are often required. It is assumed that a pre-existing dataset of a set of input-target pairs exists $\mathcal{U} = \{(x_1, y_1), ..., (x_n, y_n)\}$. During the training process, the data is fixed and only the model parameters vary and this dataset is assumed to be representative of the problem, with sufficient data to build accurate models with. It is also typically useful to know how the data was collected, was it experimental (with controlled conditions where data is "designed") or observational (data is "found").

The central problem within data modelling and machine intelligence occurs due to resolution of the sampled data points from the sample space. As a finite number of data samples are collected, there will be gaps between sampled points, meaning a predictive model must infer a continuous function from a finite sample set. Therefore a choice must be made in how to create model predictions for these 'gaps' between points (interpolation or generalisation)? And what happens beyond the limits of the dataset that the model was trained on (extrapolation)?

This becomes an inverse problem, and introduces the problem of induction. The problem of induction is the philosophical question, originally proposed by David Hume in the mid-18$^{\text{th}}$ century, of the justifications (or lack thereof) for inductive reasoning. In particular the problem of generalising about the properties of a class of objects based on historical observations of that class, widely known as the "black swan" problem (Taleb 2007, Vickers 2011). The black swan problem is the inference that *all swans we have seen are white, and, therefore, all swans are white.* The second danger of inductive reasoning is in presupposing that a sequence of events in the future will occur as it always has in the past. A real example of this warning is given by the Fincancial Conduct Authority (FCA), who enforce legal obligations onto the financial services industry

in the UK, that there is clear communication to prospective clients of investment products that "*past performance is not a reliable indicator of future returns*[ii]" (FCA 2021). Model outputs, therefore, always require a level of scepticism, and knowledge of the problem domain is highly useful in determining how to handle extrapolation and interpolation scenarios when modelling.

Machine learning is the name given to describe the capability of an intelligent machine system (AI) to acquire its own knowledge, by extracting patterns from raw data. It is the scientific study concerned with statistical models and algorithms that the computers use to perform a task without explicit input, relying on pattern recognition and inference instead. Generally, machine learning can be separated into three categories:

- *supervised learning*
  - where a dataset is split into inputs and outputs and the goal is to predict the value of an output measure based on a number of input measures.
  - Typical predictive methods include regression (continuous data) and classification (discrete data).
- *Unsupervised learning*
  - where there is no outcome measure, and instead the purpose is in discovering the structure of the data and the associations and patterns amongst input measures.
  - Typical descriptive methods include clustering and association.

[ii] Section COBS 4.6.2

- *Reinforcement learning*

  - where a reward is associated with achieving a set goal.

The introduction of machine learning has allowed computers to make decisions based on real world knowledge that appear subjective. For example, a machine learning algorithm called logistic regression can determine whether to recommend caesarean delivery (Mor-Yosef et al. 1990) whilst an algorithm called naive Bayes can flag spam e-mail from incoming e-mail. The performance of these machine learning algorithms depends on the representation of the data they are given. In the caesarean example, the AI system will not examine the patient directly, but is given information deemed relevant, such as presence (or absence) of a uterine scar. The information included within this representation is known as a feature (Goodfellow et al. 2016, chap. 1), and the machine learning algorithm, in this case logistic regression, learns how these features correlate with various outcomes. Machine learning algorithms, therefore, not only are dependent on the data they are given, but also on the *representation* of this data. For example a non-English speaker would find it difficult to navigate from London Victoria station to Covent Garden station given a textual description of the route, but given an underground map drawn with a highlighted route and an arrow suggesting "you are here!", the task would be much easier.

Machine learning, and other data-driven approaches are therefore highly useful tools to build predictive models that can provide accurate predictions even when the relation between and input and output is not explicitly known. For an extensive overview of machine learning, see Goodfellow et al. (2016) and Hastie et al. (2009).

2.3.2 Data-driven approaches in blast and impact events

Surrogate modelling is an engineering method that establishes a structure (a surrogate model) that provides an output, given some input, to a variable that would otherwise be expensive to obtain experimentally or numerically. Constructing surrogate models is a data-driven, bottom-up process, and therefore machine learning approaches have been shown to produce accurate surrogate models in a variety of specialisms within civil engineering.

A specific type of machine learning model, artificial neural networks (ANNs) (Agatonovic-Kustrin & Beresford 2000), have been used to provide fast running surrogate models that predict: the failure of structural columns subject to dynamic loads (Stewart 2010, Stewart & Morrill 2015); impulse and mid plate displacement based on plate dimensions from blasts generated by PE4 charges (Bortorlan Neto et al. 2020), with a direct application in assessing the vulnerability of naval vessels (Bortorlan Neto et al. 2017); fatigue crack propagation in steel by considering the loading ratio (Iacoviello 2004); the corrosion-fatigue behaviour of steel (Haque 2001); determine ballistic performance of armour solutions (KılıÇ et al. 2015); predict the loading produced from boiling liquid expanding vapour explosions in an open environment (Li et al. 2021) and used alongside other models such as support vector machines (Noble 2006) to provide a classification of whipple shield performance subject to a hypervelocity impact event, providing a failure or non-failure (boolean) result (Ryan et al. 2016).

Considering predictive methods of blast loading on a larger scale, particularly in city geometries, is particularly challenging due to the complex wave interactions that occur. Remennikov & Rose (2005) highlight

this issue when presenting a framework to determine load enhancement factors from adjacent buildings in urban terrain, where the authors state *"the actual blast loads can either be reduced due to shadowing by other buildings or can be enhanced due to the presence of other buildings in the vicinity"*. For simple geometries and city streets, Smith & Rose (2006) suggest empirical rules can be formulated to predict blast resultants on building façades, though as the complexity of city street layouts increase, such rules become inaccurate and the authors recommend numerical simulation instead. Remennikov & Rose (2007) and Bewick et al. (2011) use an ANN to accurately predict the blast environment (peak overpressure and peak scaled impulse) behind a vertical wall barrier for various blast wall scenarios. To train and validate the ANN, the authors develop a database through a series of experiments and the results demonstrate the feasibility to use ANNs for prediction of non-ideal blast loading in situations with rigid obstacles such as blast walls. Bewick et al. (2011) note a core issue within surrogate modelling, in that it can be difficult to construct an accurate, data-driven surrogate model when the data generation itself is expensive. Extending the blast wall analysis Remennikov & Mendis (2012) use ANNs to predict blast loads in complex city street environments, where street configuration parameters were the principle inputs and peak pressures and impulses were the outputs and further demonstrate how ANNs can be used as surrogate models for blast loading prediction in complex environments.

Flood et al. (2009) also replicate the blast wall scenario modelled by Remennikov & Rose (2007), also including a target building behind the blast wall, and propose a concept of course-grain modelling, previously shown to work in modelling dynamic heat transfer in buildings (Flood et al. 2004). This concept is similar to traditional computational fluid

dynamics or finite element analyses in that the environment is discretised and the state of each element advances in time. The difference being that no physical equations are solved at each time step, a trained ANN instead implements the output of these physical equations. This proposed coarse-grain modelling approach is efficient and accurate in a one-dimensional framework but is not yet capable of modelling higher spatial dimensions or predicting impulse and is still in early stages of development (Flood et al. 2012). A coarse grain modelling approach was demonstrated in Löhner & Baum (2004) to model blast loads on generic building configurations which differed in that it uses theoretically derived functions to drive a simulation as opposed to the empirically derived functions proposed by Flood et al. (2012). It is suggested that the loss of information resulting from the coarse spatial resolution in this approach is compensated by sampling in the time-domain and gaining information about the temporal progress of the blast wave.

Alternative approaches identified for simplifying the challenge of blast load prediction in complex environments is to use ray tracing algorithms, specifically shortest path ray tracing. Ray tracing methods search for the shortest path a wave can follow from the point of detonation to any target points, considering reflection and diffraction and has been implemented by Frank et al. (2008) to produce a fast running tool to calculate air blast pressure waveforms in urban environments. However, although a promising methodology this tool does not provide accurate enough solutions to be used reliably by blast protection engineers. An alternative approach to simplifying the challenge lies in considering the trade-off between mesh resolution, computational time and accuracy. Frank et al. (2008) compare their proposed fast running model (with ray tracing) with a coarse mesh finite element code (Löhner et al. 1987) where

they state reasonable solutions are provided, however, more accuracy is certainly required to be used in civil engineering applications.

Löhner & Baum (2004) argue that coarse CFD simulations can provide a much higher accuracy than line-of-site (ray tracing style) calculations, particularly in more complex geometries that entail more complex wave interactions. The authors provide some guidance on necessary requirements of running these coarse CFD simulations and provide three important conclusions: peak impulses are adequately captured in coarse meshes, peak pressures show discrepancies between $25 - 50\%$ between coarse and fine meshes and finally coarse grids do not capture weaker shock reflections between buildings. Klomfass (2018) perform a similar set of analyses, with the intention of investigating the discretisation error associated with the finite spatial resolution of their own developed CFD code. An important conclusion from this work finds a similar conclusion to Löhner & Baum (2004) in that maximum impulse does not need an error correction, but it also concludes that practically useful predictions of maximum overpressures can be obtained on a standard computer within a few minutes provided the discretisation error is considered.

A common focus in the aforementioned applications is to produce a surrogate model that can predict blast wave output for some given complex scenario. Considering modelling at a larger scale, creating surrogate models becomes difficult due to the high number of variables that influence the evolution of a blast wave and the non-linear result, and therefore the volume of data required increases so fast the available data becomes sparse, termed the *curse of dimensionality* (Bellman 1966). Similarly when considering surrogate models focussed on structural response subject to localised blast, large numbers of variables are introduced and large amounts of data are required. Further, in these applications the knowl-

edge of structural loading is 'lost' within the model and cannot be used elsewhere, say to model a concrete RC beam instead of a steel I beam.

2.3.3 Surrogate modelling strategies to increase data utility

The data-driven nature of building a surrogate model naturally requires large amounts of data. However, a particular issue in a blast engineering domain is that experimental data can be difficult to obtain due to the loading magnitudes and sub-millisecond durations of blast events. Combining data-intensive deep learning approaches and scientific theory is one promising avenue of research to address this issue and is considered to be a crucial step to improve model predictive performance whilst respecting natural laws (Reichstein et al. 2019). One method to achieve this integration involves guiding the learning of a machine learning model though introducing a physical consistency penalty as a regularisation procedure (Karpatne et al. 2017, Daw et al. 2020, Jia et al. 2019, Stewart & Ermon 2017). Physics-based approaches such as these have been used effectively in a variety of domains such as power-flow research (Hu et al. 2020), seismic response modelling (Zhang et al. 2020) and the geo-sciences community (Muralidhar et al. 2018, Karpatne et al. 2017).

An alternative strategy to increase data utility involves knowledge transfer between domains, known in the machine learning as "transfer learning". Traditionally, data mining and machine learning algorithms provide predictions on future data using statistical models trained on previously collected labelled or unlabelled data (Pan & Yang 2010, Ramon et al. 2007, Taylor & Stone 2007). A core assumption underpinning many machine learning methods is that the training and test data belong to the

same distribution. Therefore when this distribution changes, most models need to be re-trained on newly collected training data. In such cases, it would be highly desirable to reduce the need to re-collect training data and therefore transfer learning between task domains would be highly useful. Many examples exist where transfer learning can be beneficial across a broad range of domains such as: smart buildings (Pinto et al. 2022); Alzheimer disease detection (Taher M. Ghazal 2022); Parkinson's disease classification (Rezaee et al. 2022); web-document classification (Mahmud & Ray 2007, Blitzer et al. 2008, Xing et al. 2007); sentiment classification (Li et al. 2009); image classification (Lee et al. 2007); WiFi localisation models (Yin et al. 2005, Raina et al. 2006, Pan et al. 2007, 2008, Zheng et al. 2008) and web-page translation (Ling et al. 2008).

The research challenge, therefore, is how to effectively implement data-driven, bottom-up surrogate modelling strategies in a domain in which it is considerably expensive to obtain data, experimentally or numerically, such as in explosive blast load prediction? Further, what methods can be used that increase the utility of data or models already obtained through incorporating prior knowledge and embedding it into the learning process?

2.4 Summary

This chapter has focussed on the theoretical knowledge behind the shock formation of blast waves and current predictive methods used in blast protection engineering to predict these blast wave parameters, with a particular focus on the near-field region. Though some of the current predictive approaches, particularly the semi-empirical KB method, are useful, their accuracy within the near-field region is shown repeatedly in

the literature to be severely limited.

The near-field region is an extremely challenging domain in which to obtain data, numerically or experimentally. Therefore an understanding of approaches that can leverage knowledge previously obtained from a domain or task and applied to a new domain or task such that it improves the efficiency or accuracy of the new surrogate model would be highly beneficial. Thereby increasing the utility of previously obtained data is of paramount importance. Furthermore, an understanding of the particular data points required to build an accurate surrogate model would have considerable benefits in both time and cost in experimental design.

It has been demonstrated that to accurately model structural response, knowledge of the magnitude and distribution of loading is critical (Rigby et al. 2019b). Furthermore, to enable a probabilistic approach to blast protection engineering, accurate and fast running surrogate models are urgently required. There is therefore a clear need for accurate, fast-running surrogate models in the near-field and extreme near-field region to enable such a probabilistic framework that would ultimately allow engineers better tools to quantify risk in a blast engineering context.

Chapter 3

Numerical modelling of air blast loading

3.1 Introduction

Computational fluid dynamics (CFD) is a branch of fluid mechanics that uses numerical analysis and data structures to simulate fluid processes. Fluid dynamics is a branch of physics concerned with the mechanics of fluids and the effect of forces on them.

Within this chapter a discussion of the theoretical groundwork underpinning CFD is provided. It begins by discussing the equations that govern the motion of an inviscid fluid (Euler equations) and of a viscous fluid (Navier-Stokes equations) and provides a high level view on the approaches that commerical CFD codes can take to simulate fluid flow. Numerical simulations within this thesis were performed using *Apollo blastsimulator* (shortened as '*Apollo*' hereafter). *Apollo* is an explicit CFD software which specialises in the simulation of explosions, blast waves and gas dynamics (Fraunhofer EMI 2018, Klomfass 2016, 2018). The conservation equations for transient flows of compressible, inviscid

and non-heat conducting, inert or chemically reacting fluid mixtures are solved using a second-order finite-volume scheme with explicit time integration.

A discussion of specific modules in *Apollo* that improve computational efficiency is provided, such as dynamic mesh adaptation (DMA), 1D-to-3D mapping, and 3D-to-3D staged mapping. This is followed by a mesh sensitivity study and then a validation exercise in the far-field against experimental data. This chapter begins with a discussion of the governing equations and their derivations. For an extensive overview into the theory of CFD, including derivations, see Blazek (2015).

3.2 Governing equations

3.2.1 Mathematical description of flow and the finite control volume

"Fluid dynamics" itself is the interactive investigation of a large number of individual particles, specifically molecules or atoms (Blazek 2015, chap. 2). By assuming the density of the fluid is sufficiently high, it can be approximated as a continuum, where matter is modelled from a macroscopic view rather than a microscopic view. The implication of this is that an infinitesimally small element of the fluid contains a sufficient number of particles such that a mean velocity and mean kinetic energy can be applied. Thereby at each point in the fluid: velocity, pressure, density, temperature and other important quantities can be defined (Blazek 2015, chap. 2).

The derivation of the governing equations of fluid dynamics is grounded on the fact that the dynamic behaviour of a fluid is determined by the following conservation laws:

1. conservation of mass,

2. conservation of momentum,

3. and conservation of energy.

The conservation of a specific flow quantity states that its total variation inside some arbitrary volume can be expressed as the net effect of the amount of the quantity being transported across the boundary, of any internal forces and sources, and of external forces acting on the volume. Where the amount of the quantity moving across the boundary is called *flux*. The flux has two constituent parts: the first due to convective transport and the second due to molecular motion in the fluid at rest (diffusive) (Blazek 2015, chap. 2).

By considering a general flow field, shown by the streamlines in Figure 3.1, a finite control volume can be defined to allow the modelling of flow properties in a bounded region. An arbitrary finite region of the flow, bounded by a closed surface $d\Omega$ and fixed in space, defines the control volume, Ω, with surface element, dS, and the associated outward unit normal vector, $\vec{n}$. Through application of the conservation law to a scalar quantity per unit volume U, its variation in time within Ω is

$$\frac{\partial}{\partial t} \int_{\Omega} U \; d\Omega$$

is equal to the sum of the contributions due to *convective flux*. Convective flux is the amount of quantity U entering the control volume through the boundary with velocity $\vec{v}$.

$$- \oint_{\partial\Omega} U(\vec{v} \cdot \vec{n}) \; dS$$

further due to *diffusive flux*, expressed by the generalised Fick's gra-

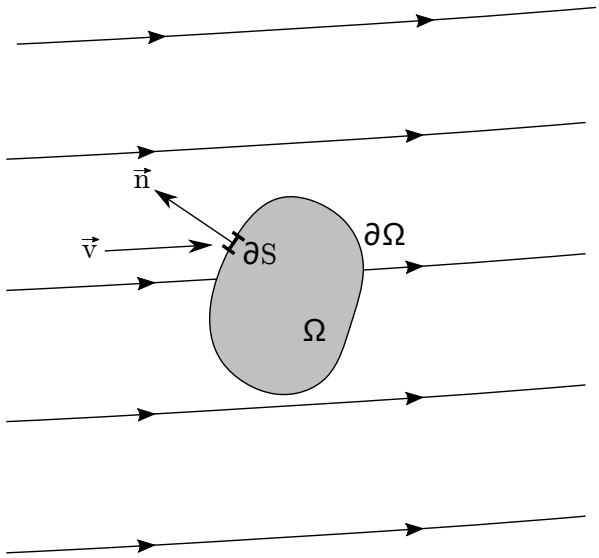

Figure 3.1: Finite control volume fixed in space, Figure recreated from Blazek (2015, chap. 2)

dient law

$$\oint_{\partial\Omega} \kappa\rho[\nabla(U/\rho)\cdot\vec{n}]\ dS$$

with κ as thermal diffusivity coefficient. With volume and surface sources, Q_V, $\vec{Q_S}$,

$$\int_{\Omega} Q_V d\Omega\ + \oint_{\partial\Omega} (\vec{Q_S}\cdot\vec{n})\ dS$$

therefore summing up the above contributive terms, the general form of the conservation law for a scalar quantity, U is obtained

$$\frac{\partial}{\partial t}\int_{\Omega} U\ d\Omega + \oint_{\partial\Omega} [U(\vec{v}\cdot\vec{n}) - \kappa\rho(\nabla(U^*\cdot\vec{n})]\ dS = \int_{\Omega} Q_V d\Omega\ + \oint_{\partial\Omega} (\vec{Q_S}\cdot\vec{n})\ dS \tag{3.1}$$

in the above formulation, U^* denotes the quantity U per unit mass (U/ρ). If the conserved quantity would be a vector instead of a scalar, the convective and diffusive flux terms become tensors instead of vectors, the convective flux tensor, $\overline{\overline{F}}_C$, and diffusive flux tensor, $\overline{\overline{F}}_D$. The volume sources would become vector $\vec{Q}_V$, and surface sources would become $\overline{\overline{Q}}_S$.

So the conservation law for a general vector quantity, $\vec{U}$ becomes

$$\frac{\partial}{\partial t}\int_{\Omega}\vec{U}\ d\Omega + \oint_{\partial\Omega}\left[(\overline{\overline{F}}_C - \overline{\overline{F}}_D)\cdot\vec{n})\right]\ dS = \int_{\Omega}\vec{Q}_V d\Omega\ + \oint_{\partial\Omega}(\overline{\overline{Q}}_S\cdot\vec{n})\ dS \tag{3.2}$$

the above formulations in Equation 3.1 and 3.2 are the integral formulations of the general conservation law for a control volume (Blazek 2015, chap. 2). This formulation has two important qualities: firstly, if there are no volume sources present, variation of U depends only on flux across the boundary $\partial\Omega$ and not on the flux inside the control volume Ω; secondly, this form remains valid in the presence of shocks or discontinuities in the flow field (Lax 1954). The differential form of the general conservation law (in Equation 3.2) requires the solution to be differentiable. At shock discontinuities the differential form breaks down, hence taking a finite volume approach which requires the integral form is more robust and generally favoured in CFD codes.

By using the integral form in Equation 3.2 the corresponding expressions for the three conservation laws in fluid dynamics can be derived: conservation of mass, conservation of momentum and conservation of energy.

3.2.2 Conservation laws

The continuity equation

Considering first the law of mass conservation for single-phase fluids, this states that mass cannot be created nor destroyed in such a fluid system. Furthermore there is no diffusive flux contribution here as for a fluid at rest, any variation of mass implies displacement of fluid particles (Blazek 2015, chap. 2).

Considering the finite control volume in Figure 3.1, the conserved quantity is density, ρ. For time rate of change of the total mass inside the finite volume Ω

$$\frac{\partial}{\partial_t}\int_\Omega \rho \ d\Omega$$

the mass flow of a fluid through some surface boundary is the product of density $\times$ surface area $\times$ perpendicular velocity (normal to the surface). So convective flux across each surface element, dS is

$$\rho(\vec{v}\cdot\vec{n})dS$$

with $\vec{n}$ pointing away from the surface by definition. As there are no volume or surface sources present, from the general form in Equation 3.2 the integral form of the continuity equation is obtained (Blazek 2015, chap. 2)

$$\frac{\partial}{\partial_t}\int_\Omega \rho \ d\Omega + \oint_{\partial\Omega} \rho(\vec{v}\cdot\vec{n})dS = 0 \tag{3.3}$$

The momentum equation

Newton's second law states the rate of change of momentum of a body over time is directly proportional to the force applied, and occurs in the same direction as the applied force. For the momentum of an infinitesimally small portion of the control volume (in Figure 3.1), Ω

$$\rho\vec{v}d\Omega.$$

The rate of change of momentum within the control volume equals

$$\frac{\partial}{\partial_t}\int_\Omega \rho\vec{v}d\Omega.$$

So the conserved variable is the product of density and velocity, where

$$\rho\vec{v} = [\rho u, \rho v, \rho w]^T.$$

The convective flux tensor, which describes the transfer of momentum across the boundary of the control volume, exists in the Cartesian coordinate system of the following three components:

$$\begin{aligned} &\text{x-component:} \quad \rho u\vec{v} \\ &\text{y-component:} \quad \rho v\vec{v} \\ &\text{z-component:} \quad \rho w\vec{v} \end{aligned}$$

the contribution from the convective flux tensor to the conservation of momentum is given as

$$-\oint_{\partial\Omega} \rho\vec{v}(\vec{v}\cdot\vec{n})dS.$$

Diffusive flux is zero for a fluid at rest. Body force per unit volume, $\rho\vec{f_e}$, corresponds to the volume sources in Equation 3.2. Thus the contribution of external body forces to the momentum conservation is

$$\int_{\Omega} \rho\vec{f_e}d\Omega$$

Surface sources consist of two parts: an isotropic pressure component and a viscous stress tensor $\overline{\overline{\tau}}$

$$\overline{\overline{Q}}_S = -p\overline{\overline{I}} + \overline{\overline{\tau}}$$

where $\overline{\overline{I}}$ is the unit tensor. By summing these contributions accord-

ing to Equation 3.2

$$\frac{\partial}{\partial_t}\int_\Omega \rho\vec{v}d\Omega + \oint_{\partial\Omega} \rho\vec{v}(\vec{v}\cdot\vec{n})dS = \int_\Omega \rho\vec{f_e}d\Omega - \oint_{\partial\Omega} p\vec{n}dS + \oint_{\partial\Omega} (\overline{\overline{\tau}}\cdot\vec{n})dS \quad (3.4)$$

for momentum conservation inside an arbitrary control volume, Ω, fixed in space (Blazek 2015, chap. 2).

The energy equation

Considering the first law of thermodynamics, applied to the control volume in Figure 3.1, it states that any changes in time of the total energy inside the volume are caused by the rate of work of forces acting on the volume and the net heat flux into it (Blazek 2015, chap. 2). Total energy per unit mass, E of a fluid is calculated by adding internal energy per unit mass, e (sometimes called specific energy), to its kinetic energy per unit mass. Therefore

$$E = e + \frac{|\vec{v}|^2}{2} = e + \frac{u^2 + v^2 + w^2}{2}.$$

The conserved variable in this instance is total energy per unit volume, ρE. Where its variation in time in the volume Ω is

$$\frac{\partial}{\partial_t}\int_\Omega \rho E d\Omega.$$

The contribution of convective flux is

$$-\oint_{\partial\Omega} \rho E(\vec{v}\cdot\vec{n})dS.$$

Unlike the previous continuity and momentum derivations, there is now a diffusive heat flux $\vec{F_D}$, which is proportional to the gradient of the

conserved quantity per unit mass (Fick's law).

$$\vec{F}_D = -\gamma\rho\kappa\nabla e,$$

where $\gamma = c_p/c_v$ is the ratio of specific heat constants. Diffusion flux represents one part of the heat flux, heat transfer due to temperature gradients and can be re-written in Fourier's law of heat conduction

$$\vec{F}_D = -\kappa\nabla T$$

where T is absolute static temperature. The second part of net heat flux is due to volumetric heating, from absorption or emission of radiation, or chemical reactions. The time rate of heat transfer per unit mass is denoted as $\dot{q}_h$, together with the rate of work done by body forces $\vec{f_e}$ means that the volume sources become

$$Q_V = \rho\vec{f_e}\cdot\vec{v} + \dot{q}_h.$$

The final contribution to conservation of energy are surface source Q_S, they correspond to the time rate of work done by pressure, along with shear and normal stresses

$$\vec{Q}_S = -p\vec{v} + \overline{\overline{\tau}}\cdot\vec{v}.$$

Therefore combining the above contributions, the energy conservation equation is obtained (Blazek 2015, chap. 2)

$$\frac{\partial}{\partial_t}\int_\Omega \rho E d\Omega + \oint_{\partial\Omega}\rho E(\vec{v}\cdot\vec{n})dS = \oint_{\partial\Omega}\kappa(\nabla T\cdot\vec{n})dS + \int_\Omega(\rho\vec{f_e}\cdot\vec{v} + \dot{q}_h)d\Omega - \oint_{\partial\Omega}p(\vec{v}\cdot\vec{n})dS + \oint_{\partial\Omega}(\overline{\overline{\tau}}\cdot\vec{v})\cdot\vec{n}dS. \quad (3.5)$$

using the relation between total enthalpy, total energy and pressure

$$H = h + \frac{|\vec{v}|^2}{2} = E + \frac{p}{\rho}.$$

the convective and pressure terms in the energy equation in 3.5 can be collected to be re-written in a more concise way (Blazek 2015, chap. 2)

$$\frac{\partial}{\partial_t}\int_\Omega \rho E d\Omega + \oint_{\partial\Omega} \rho H(\vec{v}\cdot\vec{n})dS = \oint_{\partial\Omega} \kappa(\nabla T\cdot\vec{n})dS + \int_\Omega (\rho\vec{f_e}\cdot\vec{v} + \dot{q}_h)d\Omega + \oint_{\partial\Omega} (\overline{\overline{\tau}}\cdot\vec{v})\cdot\vec{n}dS. \quad (3.6)$$

3.2.3 The complete Navier-Stokes equations

The previous sections, recreating the derivations demonstrated in (Blazek 2015, chap. 2), presents derivations of the conservation laws of mass (Equation 3.3), momentum (Equation 3.4) and energy (Equation 3.6). By collecting into one system of equations the complete system of the Navier-Stokes equations can be obtained. Introducing two flux vectors for brevity: the vector of convective fluxes, $\vec{F}_C$, and the vector of viscous fluxes, $\vec{F}_V$. Additionally a source term, $\vec{Q}$, comprises all volume sources due to body forces and volumetric heating. Combining Equation 3.2 with Equations 3.3, 3.4 and 3.6

$$\frac{\partial}{\partial_t}\int_\Omega \vec{W}\ d\Omega + \oint_{\partial\Omega} (\vec{F}_C - \vec{F}_V)\ dS = \int_\Omega \vec{Q}\ d\Omega \quad (3.7)$$

From this, the vector of conservative variables, $\vec{W}$, consists in three dimensions of five components

$$\vec{W} = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho E \end{bmatrix},$$

for the vector of conservative fluxes,

$$\vec{F}_C = \begin{bmatrix} \rho V \\ \rho u V + n_x p \\ \rho v V + n_y p \\ \rho w V + n_z p \\ \rho H V \end{bmatrix},$$

with V as contravariant velocity, the velocity normal to the surface element, dS. It is defined as the scaler product of velocity vector and unit normal vector

$$V \equiv \vec{v} \cdot \vec{n} = n_x u + n_y v + n_z w$$

and the source term is

$$\vec{Q} = \begin{bmatrix} 0 \\ \rho f_{e,x} \\ \rho f_{e,y} \\ \rho f_{e,z} \\ \rho f_e \cdot \vec{v} + \dot{q}_h \end{bmatrix}.$$

The vector of viscous fluxes has been omitted for brevity, but can be found in (Blazek 2015, chap. 2). The Navier-Stokes equations in Equation 3.7 describe the behaviour of a viscous fluid. In many applications, it is

a useful approximation to remove viscous effects entirely, so the vector of viscous fluxes, $\vec{F}_V$, is ommitted, leaving

$$\frac{\partial}{\partial_t}\int_\Omega \vec{W}\ d\Omega + \oint_{\partial\Omega} \vec{F}_C\ dS = \int_\Omega \vec{Q}\ d\Omega \tag{3.8}$$

the above formulation in Equation 3.8 is known as the *Euler Equations*, and they describe pure convection of flow quantities in an inviscid fluid. The integral formulation above, allow for accurate representation of shocks (Blazek 2015, chap. 2).

It is sometimes required to solve the Euler equations on a moving or deforming grid. The two most popular approaches for this, as suggested by Blazek (2015, chap. 2) are: Arbitrary Lagrangian Eulerian (ALE) formulations (Hirt et al. 1974, Pracht 1975, Donea et al. 1982) and Dynamic grids (Batina 1990). Where both approaches make an alteration to Equation 3.8 that allows relative motion between the grid and fluid

$$\frac{\partial}{\partial_t}\int_\Omega \vec{W}\ d\Omega + \oint_{\partial\Omega} (\vec{F}_C^M - V_t\,\vec{W})\ dS = \int_\Omega \vec{Q}\ d\Omega. \tag{3.9}$$

where V_t is contravariant velocity of the face of the control volume and is given by

$$V_t = \vec{g}\cdot\vec{n} = n_x\frac{\partial x}{\partial t} + n_y\frac{\partial y}{\partial t} + n_z\frac{\partial z}{\partial t}$$

and n_x, n_y and n_z are components of outward facing unit normal vector of surface $\partial\Omega$. Adapting the previous vector of convective fluxes, $\vec{F}_c$ to account for the relative fluid motion

$$\vec{F_C} = \begin{bmatrix} \rho V_r \\ \rho u V_r + n_x p \\ \rho v V_r + n_y p \\ \rho w V_r + n_z p \\ \rho H V_r + V_t p \end{bmatrix},$$

with V_r being the contravariant velocity relative to the motion of the grid (Blazek 2015, chap. 2)

$$V_r = V - V_t.$$

If $V_t = V$ then the *Lagrange formulation* results, if $V_t = 0$ the *Euler formulation* results, and if $V_t =$ other the *Arbitrary Lagrangian Eulerian (ALE) formulation* results.

3.2.4 Equations of state

The Navier-Stokes equations (Equation 3.7) and Euler equations (Equation 3.8) are systems of five conservative variables ρ, ρu, ρv, ρw and ρE, but contain seven unknown flow field variables ρ, u, v, w, E, p and T. To close out this system of equations, two additional equations of state need to be provided. Two common examples are expressing pressure as a function of density and temperature, the thermal equation of state, and expressing internal energy as a function of density and temperature, the caloric equation of state (Blazek 2015, chap. 2).

In modelling high explosives in *Apollo*, an equations of state data file provides the simulation with both caloric equation of state and thermal equation of state information. The caloric equation of state model used in *Apollo* for all fluids takes the form of a quadratic polynomial function

$$e = e_0 + c_{V1}T + c_{V2}T^2$$

where c_{V1} and c_{V2} are specific heat constants, which become temperature dependent at higher temperatures. All of the simulations within this thesis consist of an explosive reacting in air. Air is modelled as a thermally perfect gas where it can be assumed

$$p = \rho RT,$$

where R is the specific gas constant, such that under ambient conditions (288 °K), ambient pressure is p = 101.3kPa. The ideal gas caloric equation of state provides the relationship between thermodynamic variable:

$$e = \frac{p}{(\gamma - 1)\rho}$$

where γ is the adibiatic index (ratio of specific heats). Explosives are modelled in *Apollo* with a linear model

$$p = a_0^2(\rho - \rho_0) + p_0,$$

where a_0 is sound speed, and p_0 is ambient pressure (101.3kPa). Detonation products are modelled with the Jones Wilkins Lee equation of state (Lee et al. 1968).

$$p(\rho, T) = C_1 e^{-R_1\rho_0/\rho} + C_2 e^{-R_2\rho_0/\rho} + \rho RT$$

where ρ and ρ_0 are density and initial density, R is a gas constant, and T is absolute temperature. The constants C_1, C_2, R_1, and R_2 are defined within the equation of state file for a given fluid.

3.2.5 Detonations

The detonation model used in *Apollo* relies on the Chapman-Jouguet model. This uses a local initiation condition and a local state dependent burning velocity. This results in the reaction front progressing in a physically meaningful way (Fraunhofer EMI 2018). The detonation product mass generated in a time step is calculated by:

$$\Delta\rho_{\mathrm{DP}} = \Delta\lambda\rho_{\mathrm{uHE}}$$

where the progression rate is calculated from:

$$\Delta\lambda = \Delta t \frac{V_{\mathrm{burn}}}{\Delta x}\mathrm{Min}(1, \mathrm{Max}(0, \theta_1, \theta_2)),$$

and trigger functions, θ_1 and θ_2 are calculated from

$$\begin{aligned}\theta_1 &= k\rho_{\mathrm{DP}}/\rho_{DP}^{\mathrm{Max}},\\ \theta_2 &= k(\rho_{\mathrm{uHE}}/\rho_{\mathrm{uHE}}^{0} - 1)\\ \theta_{1,2} &\in [0, 1].\end{aligned}$$

where ρ_{DP} is the actual detonation product density in the cell and ρ_{uHE} is density of unburned explosive with reference density, ρ_{uHE}^{0}, of the unburned explosive. k is a sensitivity coefficient fixed to 3 (Fraunhofer EMI 2018). In the CJ model a cell is initiated only through inflow of detonation products. Local burning velocity is determined according to the CJ condition:

$$V_{\mathrm{burn}} = V_{\mathrm{det,\ CJ}} = a + |v|,$$

which represents the sum of local sound velocity, a and material velocity, v.

3.3 The numerical scheme

3.3.1 Spatial discretization – the finite volume method

The previous section described the derivation of the Euler Equations (Equation 3.8) and the introduction of the equations of state to provide thermodynamic relations between variables. There remains two key steps to solve these system of equations: *spatial discretization* of the domain (the space where the flow is being computed), and *temporal discretization*, to calculate the time-wise progress of flow variables in the domain.

For spatial discretization, *Apollo* implements a finite volume scheme. For a comprehensive overveiw of other spatial discretization schemes, see (Blazek 2015, chap. 3). The finite volume method directly implements the integral form of the Euler Equations by dividing the physical domain into a number of arbitrary polyhedral control volumes. The surface integral on the right hand side of Equation 3.8 is approximated by the sum of the fluxes crossing each face of the control volume (Blazek 2015, chap. 3). Two basic approaches for defining the shape and position of the control volumes are given by (Blazek 2015, chap. 3) as: cell-centered schemes, where flow quantities are stored at the centroid of the grid cells so that control volumes are identical to grid cells; cell-vertex schemes, where flow variables are stored at grid points or some volume around a grid point.

A common approach utilised in numerical methods is a separate discretization in both space and time, called the method of lines (Richtmyer & Morton 1967). In this method, the grid is either used to construct control volumes and to evaluate flux integrals, or to approximate spatial derivatives of flow quantities. Then in a further time step the time-dependent equations are advanced in time (Blazek 2015, chap. 3).

The finite volume approach demonstrates the general choice of spatial discretization, various numerical schemes exist that can implement the approach, and can be classified into *central* or *upwind* schemes. Central schemes rely on central difference or central averaging formulas where the cell states are averaged on the left and right side of a boundary to establish the flux at the vertex in between, whilst upwind schemes account for both upstream and downstream influences. For more detail on both approaches see Blazek (2015, chap. 3).

3.3.2 Temporal discretization – explicit time integration

The code used in *Apollo* uses explicit time integration. This is performed in two steps: firstly a Lagrange step, with acceleration and deformation of a material volume evaluated and a subsequent remap step in which updated material volume is remapped back onto the grid cells (Klomfass 2018).

For the Lagrange step, a characteristics based approximation of pressure and velocity at cell interfaces is used, which is based on the solution of the locally one-dimensional Euler equations for discontinuous states at an interface (Blazek 2015, chap. 3) and falls into the class of "flux-difference splitting schemes". This scheme solves the acoustic approximation of the one-dimensional Reimann problem for material velocity and pressure on the contact surface between the left and right cell states, and can be considered a Godunov-type method (Godunov 1959). Considering one-dimensional Euler equations, the Reimann problem in this instance is the hyperbolic equation with a set of piecewise constant initial data with a single discontinuity at $x = 0$ and $t = 0$.

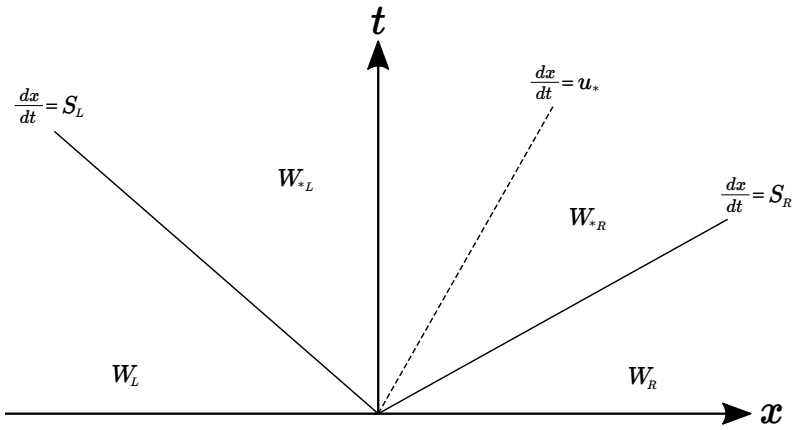

Figure 3.2: Solution of the Reimann problem in x, t space.

Initial data:

$$\boldsymbol{u}(x,0) = \begin{cases} \boldsymbol{u}_L, & \text{if } x < 0, \\ \boldsymbol{u}_R, & \text{if } x > 0. \end{cases}$$

where $\boldsymbol{u}_L$ and $\boldsymbol{u}_R$ are states either side of the discontinuity. The solution of the Reimann problem is shown in Figure 3.2. The complete solution is constant along lines emanating from the origin, depending only on the variable x/t.

There are three waves: the middle wave is the contact discontinuity, u. The left (S_L) and right (S_R) non-linear waves are shocks or rarefactions. Contacts and shocks are discontinuous solutions whereas rarefaction waves are continuous solutions. Four different states in this instance are denoted as: W_L, W_{*L}, W_{*R}, W_R. The region between non-linear waves has constant pressure and constant velocity, known as the star region (denoted by the "*"). Density and internal energy change discontinuously across the contact surfaces. Solving the Reimann problem begins with obtaining the star states using the appropriate equations.

Considering first the rarefaction wave, the isentropic law

$$p = C\rho^{\gamma},$$

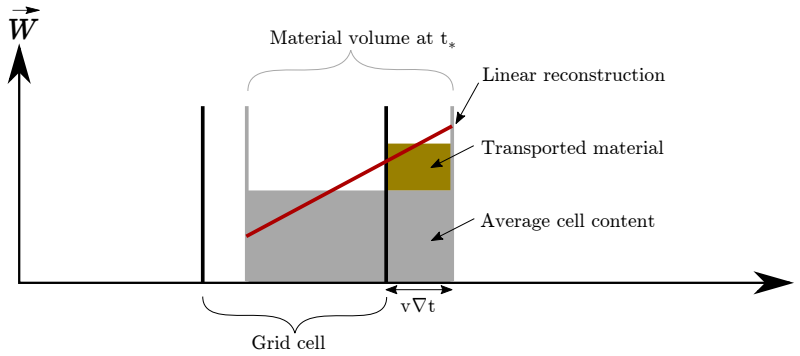

Figure 3.3: Remap stage: the one-dimensional donor cell method. Figure adapted from Fraunhofer EMI (2018).

and the generalised Reimann invariant

$$C = u + \frac{2a}{\gamma - 1}$$

with u as speed, C is a constant and a is the sound speed. These are used to relate the states across the rarefaction wave to obtain the star state.

Specifically considering shock, the Rankine-Hugoniot conditions are solved across the wave in the shock frame of reference:

$$\rho\hat{u} = \rho_*\hat{u}_*,$$
$$\rho\hat{u}^2 + p = \rho_*\hat{u}_*^2 + p_*,$$
$$\hat{u}(\hat{E} + p) = \hat{u}_*(\hat{E}_* + p_*),$$

Non starred quantities portray the left or right initial data state. $\hat{u} = u - s$, $\hat{u}_* = u_* - s$ with s as shock speed.

The remap phase transports the updated material volume back onto the mesh cells using the one-dimensional donor cell method. See Figure 3.3 for a schematic demonstrating this process.

Both stages (the Lagrangian and remap phase) are computed with interpolated states based on tri-linear reconstructions of the distribution

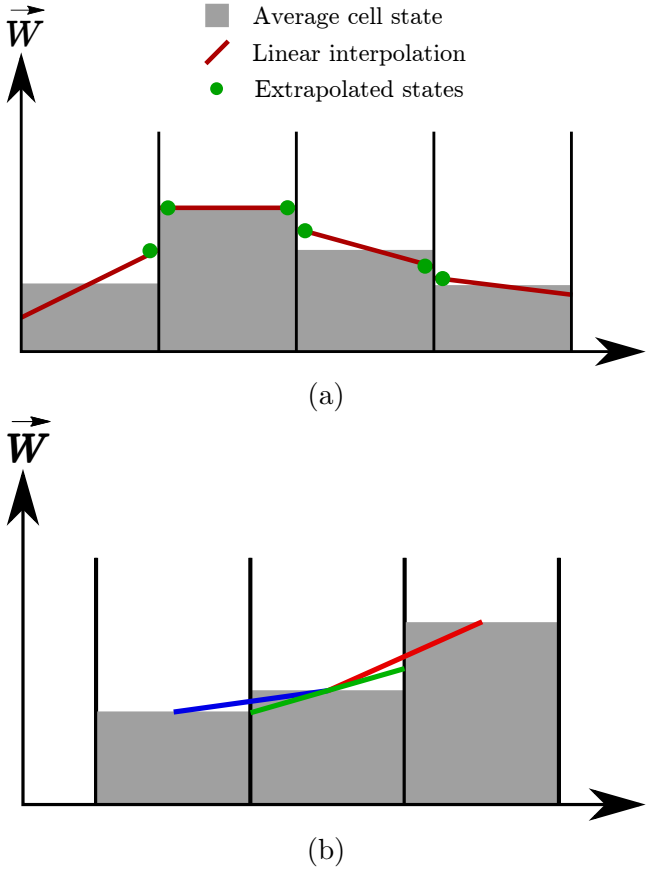

Figure 3.4: Extension to second order accuracy. (a) shows linear interpolation of averaged cell states, and (b) shows how the linear interpolation values are obtained, as a function of left and right sided difference. Figure adapted from Fraunhofer EMI (2018).

of conservative variables within the grid cells in order to achieve second order accuracy (Klomfass 2018), and is shown schematically in Figure 3.4.

These linear reconstructions, of which there is one for each spatial direction, are controlled individually through a slope limiter resembling the UMIST formulation (Lien & Leschziner 1994) given as

$$\phi(r) = \max[0, \min(2r, (0.25 + 0.75r), (0.75 + 0.25r), 2)]; \quad \lim_{r\to\infty} \phi(r) = 2$$

where r represents the ratio of successive gradients on the solution mesh. Implementation of a slope limiter ensures that no new extrema are created and the extrapolated states are used in flux calculations at the interface.

In explicit time-stepping schemes, initial values are known and used to calculate the solution at time $t + \Delta t$. Therefore the new solution relies solely on values already known. To ensure that the solution remains stable, a restriction on the time step is required. This restriction prevents information travelling further than a cell width, Δx, in a time step, Δt, which means the solution calculated in a given cell is only affected by the information that is able to affect it physically.

The time step equation takes the form:

$$\Delta t < C_{\text{cfl}} \frac{\Delta x}{|v - v_t| + a}, \tag{3.10}$$

where Δx is the smallest cell size, $|v - v_t|$ is relative velocity between material and (moving) cell, a is sound speed and C_{cfl} is a dimensionless quantity called the Courant-Friedricks-Lewy number, in *Apollo* this is defined as 0.4 but can be decreased if instabilities in a model occur.

3.4 Additional features in *Apollo*

3.4.1 One-dimensional simulation

A significant feature of *Apollo* is the one-dimensional (B1D) module, it works with uniform one-dimensional meshes and permits the first stage of the simulation to be completed in one-dimension, provided a spherically symmetric solution exists (it cannot therefore be used when modelling cylinders). Whenever a pressure discontinuity (such as the shock front)

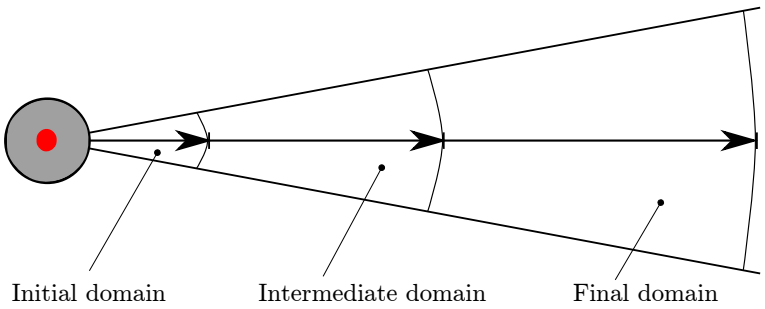

Figure 3.5: Global mesh adaptation process in B1D. Recreated from Fraunhofer EMI (2018).

approaches the outer boundary of the present domain, it is resized, as in Figure 3.5.

In each adaptation step, cell size and the domain are doubled so that the total number of cells remains approximately constant throughout the B1D simulation. Sensitivity of the re-meshing can be controlled via the input file, and is dependent on the gradients of variables in the flow field (with maximum sensitivity meaning that all zones are set to maximum level of resolution). This stage is terminated when user-inputted termination criteria are met, such as number of timesteps or distance.

3.4.2 Dynamic mesh adaptation

The Dynamic Mesh Adaptation (DMA) stage is based on Cartesian meshes which have an inconsistent resolution globally. Specific regions of the computational domain (such as those that require a greater resolution, e.g. the shock front) will require different resolutions as the simulation progresses temporally. This has the benefit of saving the computational power for the areas that require it through reducing the computational demand from areas where it is not needed and is shown schematically in Figure 3.6.

The highest mesh resolution to be used is directed in the input file.

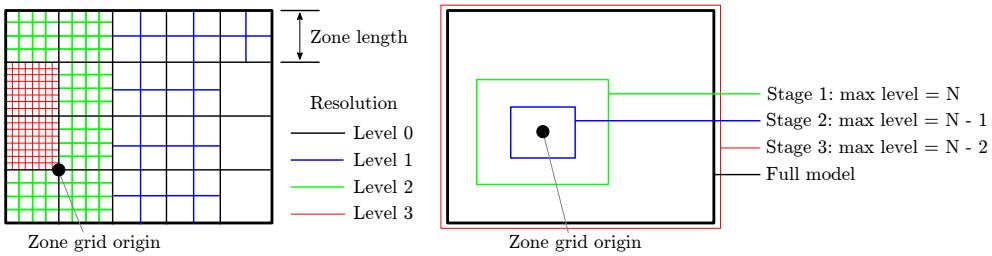

Figure 3.6: Dynamic mesh adaptation (left) and multi-stage concept for global adaptation shown (right). Recreated from Fraunhofer EMI (2018).

The resolution level represents the size of the smallest mesh cell used in the simulation, determined through resolution level "N" and zone length, "L". Ultimate cell size, "h", then equals:

$$h = \frac{L}{2^N}$$

The DMA algorithm works in such a way that always at least one zone (the zone with the largest gradients[i] at the time) is resolved at the highest level. When the entire flow field is uniformly at ambient state, and no objects are embedded then all zones will remain at level 0 (Fraunhofer EMI 2018). Sensitivity of DMA with respect to the distribution of cell resolution can be defined in the simulation (similar to the B1D module). A maximum sensitivity value would mean all cells are resolved to the highest possible resolution, and leads to a significant increase in computational time.

A stability level for the computations can also be defined in the input file. This becomes useful if any non-physical states or instabilities in the computations occur (noticeable by significantly reduced time steps). The stability level acts on the time step size by reducing the CFL number and the combination of parameters that act on the tri-linear reconstruction of state variables.

[i] A focus can be applied to particular variables in the keyword file.

An option in the DMA module exists to perform an ALE simulation via a mesh expansion. This means the mesh is continuously stretched according to a pre-defined velocity-time function. If this option is selected, the computational scheme changes from Eulerian (mesh fixed in space) to ALE. A uniform mesh expansion about a central point is used by applying a globally uniform stretch rate given by:

$$\dot{\epsilon}(t) = \frac{1}{L} \times \left(\frac{V_1}{(1 - V_1(t - t_0) \div L)^{\beta}} + V_{\text{end}} \right); V_1 = V_{\text{beg}} - V_{\text{end}}$$

where V_{beg} and V_{end} are initial and final velocities respectively, β is 0.75, L is the distance to centre point (X, Y, Z), t_0 and t_{end} are initial and end times, ϵ_{end} is the ultimate stretch ratio and all are constants to be defined. Motion starts for $t = t_0$ and ends at t_{end} or the ultimate stretch ratio is reached, calculated by:

$$\epsilon_{\text{end}} = \frac{L(t)}{L(t_0)}.$$

3.4.3 Stages

In *Apollo* there is the option to split up the overall simulation into a multi-stage simulation for computational efficiency, as shown in Figure 3.6. The initial stage radius, termination criteria and growth rate can be defined by the user, or the default "auto-staging" procedure can be used. The "auto-staging" procedure consists of 6 stages defined with scaled termination radii of 0.2, 0.5, 0.8, 1.25, 5 and infinity (in $\text{m/kg}^{1/3}$). Additional resolution levels are defined for each stage, beginning with an additional 6 levels, therefore as the simulation progresses, when the disturbance reaches some termination criteria, the total (additional) res-

olution level will drop by 1, and the new domain size will grow by the defined growth rate. This process repeats until the additional resolution levels drop to 0 and the entire domain is then included in the simulation. Successive stages are linked together through remapping of the flow field reached at the end of the prior stage onto the mesh of the successive stage.

3.5 Validation of *Apollo blastsimulator* in in the far-field

3.5.1 Mesh sensitivity study

The first priority when validating *Apollo* is to ascertain the confidence that can be placed from numerical data produced for a less challenging scenario (i.e. far-field scenarios) before completing near-field procedures (in Section 4.2.2). An initial mesh sensitivity study was therefore completed in order to determine required element sizes to achieve convergence. A series of numerical simulations were completed for 0.35 kg hemispheres of PE4 (modelled as a 0.7kg sphere in *Apollo*), located at a stand-off distance, S, of 6 m (scaled distance $Z = 8.51$ m/kg$^{1/3}$) from a rigid reflecting surface. PE4 is a plastic explosive comprising 87% RDX and 13% mineral oil binder (Tyas 2019) and was used throughout the mesh sensitivity analysis and validation using *Apollo*'s in-built model for C4 since the two explosives are nominally identical (Bogosian et al. 2016). The model parameters are given in Table **??** in the appendix. In all numerical simulations the explosives were centrally detonated using the Chapman-Jouguet model and the mass of the detonator was not included.

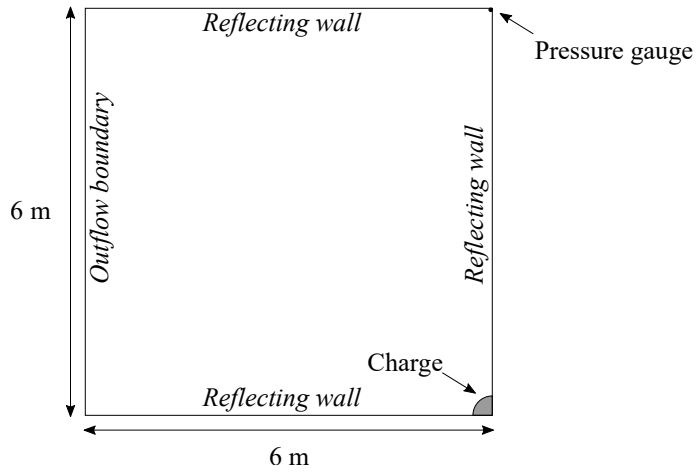

Figure 3.7: Model set-up for mesh sensitivity study. The floor is modelled as a reflecting wall and the ceiling is modelled as an outflow boundary. Eighth-symmetry is used.

The domain size was $6 \times 6 \times 6$ m, and eighth-symmetry was used, with symmetry planes located in the directions orthogonal and opposite to the reflecting wall, originating at the centre of the charge and with a ground plane set as a reflecting surface. The use of symmetry planes technically models the explosive as if it was situated between two walls, spaced 12 m apart, however the assumption is that the analysis will be terminated before any waves from the artificial wall would reach the true wall in the model and begin to affect the results. In all models an integer number of zones was specified for each of the domain side lengths. Outflow boundaries were defined at the roof of the domain and the remaining boundary. A numerical pressure gauge was placed at the base of the wall directly opposite the charge centre. This model set-up is shown in Figure 3.7. *Apollo*'s 1D module was used with the limit defined at one zone length from the gauge (for a zone length of 300 mm, the 1D stage extended 5700 mm from the charge centre), and the auto-staging module was used throughout to improve the simulation time. In total 24 simulations were completed with the mesh information summarised

Table 3.1: Ultimate cell length (element size at highest resolution level) and number of elements (between charge centre and normal gauge location) for initial mesh sensitivity study, Z = 8.5 m/kg$^{1/3}$. Note, this does not consider the additional resolution provided by the staging modules.

	Zone length (mm)				
Res. level	**1000**	**600**	**400**	**300**	**200**
0	1000	600	400	300	200
1	500	300	200	150	100
2	250	150	100	75	50
3	125	75	50	37.5	25
4	62.5	37.5	25	18.75	

(a) Ultimate cell length (mm)

	Zone length (mm)				
Res. level	**1000**	**600**	**400**	**300**	**200**
0	6	10	15	20	30
1	12	20	30	40	60
2	24	40	60	80	120
3	48	80	120	160	240
4	96	160	240	320	

(b) Number of cells between charge and target

in Table 3.1.

Results from the mesh sensitivity analysis are shown in Figure 3.8, each sub-figure shows: peak overpressure; peak specific impulse and total analysis time (termed 'wall time'), all plotted against the ratio of stand-off distance, S, to ultimate cell length. The solid black line shows the average experimental value from Rigby et al. (2015a) and the dashed line shows a 10% variation from the experimental value included as a benchmark for comparison.

The study suggests that an ultimate cell length of $h = S/240$ and $h = S/40$ are required to be within 10% error (from experimental benchmark) in overpressure and specific impulse respectively. As this research focusses on predicting specific impulse values, the less-stringent conver-

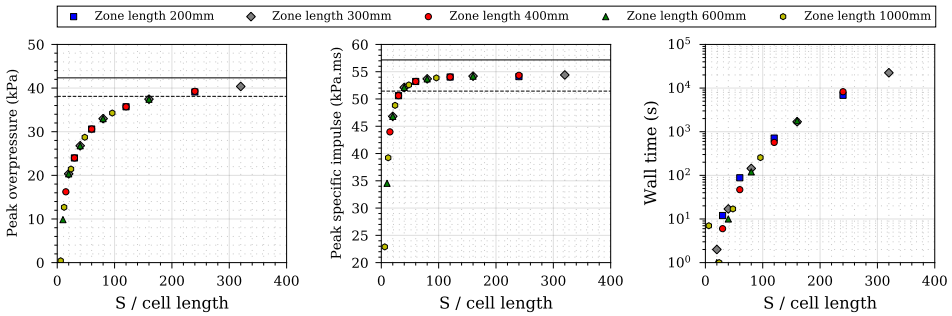

Figure 3.8: Mesh sensitivity study for 0.35 kg PE4 at 6 m stand-off from a rigid reflecting wall. Solid line indicates average experimental value (Rigby et al. 2015a) and dashed line indicates 10% variation from the experimental value. Note, *Apollo's* staging procedure was used so S/cell length is greater than shown, but would apply uniformly to all models.

gence requirement of $h = S/40$ is implemented in later analyses. It is worth noting here that at $h = S/40$, peak specific impulse is within 4% of its converged value, considered to be acceptable as an additional 4% in accuracy would result in a 2000000% increase in simulation time taken. Further, it is important to note that this does not include the additional cells from the first stage being resolved in one-dimension.

The benefits of the DMA module can be seen when wall time is considered: at $h = S/60$ a noticeable reduction in wall time occurs when going from 200 mm to 400 mm zone length. Higher resolution levels allow *Apollo* to efficiently allocate computational resource, therefore a minimum level of 3 and a minimum ultimate cell length of $h = S/40$ are recommended, provided an integer number of zones are specified for each of the domain side lengths.

3.5.2 Experimental validation

Rigby et al. (2015a) present a series of experimental trials where pressure gauges, embedded flush with the surface of a masonry wall, were used

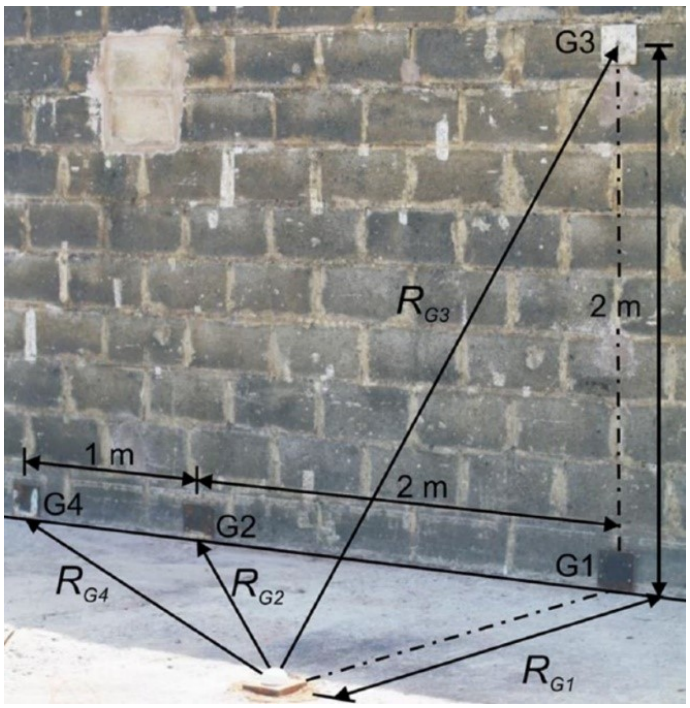

Figure 3.9: Pressure gauge location and general test arrangement of experimental data (Rigby et al. 2015a)

to record reflected pressure histories from 0.18–0.35 kg PE4 hemispheres located 2–10 m from the bunker wall. The experimental set-up is shown in Figure 3.9. For this validation exercise, only data from the normally reflected gauge, 'G1', was used.

The experimental dataset consists of 19 tests at 9 unique scaled distances, therefore a further 8 numerical analyses were performed (in addition to the 0.35 kg at 6 m test presented in Section 3.5.1). In each, the domain size was a regular cube with side-length equal to the stand-off distance in each case. As in the mesh sensitivity study, eighth-symmetry was used, with symmetry planes located in the directions orthogonal and opposite to the reflecting wall, originating at the centre of the charge. Spheres of PE4 were modelled with the ground modelled as a reflecting surface. Outflow boundaries were defined at the roof of the domain and the remaining side. A numerical pressure gauge was placed at the base of

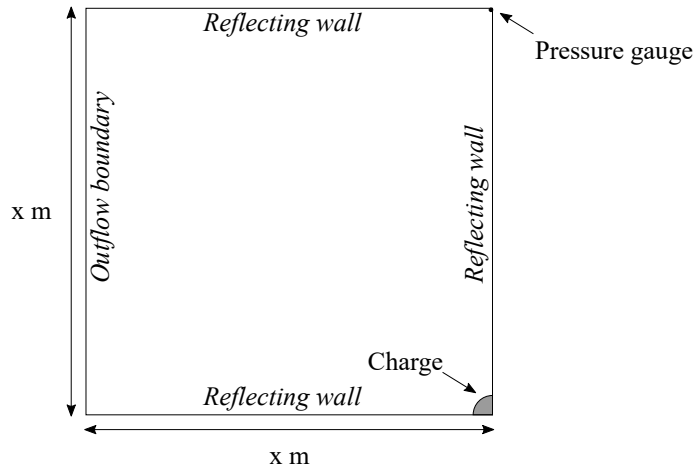

Figure 3.10: General model set-up for far-field validation study. The floor is modelled as a reflecting wall and the ceiling is modelled as an outflow boundary. Eighth-symmetry is used. 'x' is one of the stand-off distances given in Table 3.2

the wall directly opposite the charge centre. Further information for each of the validation models is presented in Table 3.2, including resolution level, ultimate cell length, and resulting stand-off/cell length (S/h) ratio, and a model schematic is given in Figure 3.10.

Results from example numerical analyses are compared to experimental data in Figure 3.11 for all the experimental scenarios outline in Table 3.2. It can be seen in all experiments that the magnitudes and general form of the numerical pressure and impulse histories closely match the experiments giving confidence that *Apollo* is correctly modelling the mechanisms of normal reflection in far-field blast scenarios. In all cases, the well-known secondary shock (Rigby et al. 2016) can be seen to arrive later in the numerical models, indicating that afterburning of the explosive detonation products is being slightly underestimated (Schwer & Rigby 2017, 2018). However, the effect of on peak specific impulse is negligible since the secondary shock consistently arrives during the negative phase.

Additionally, a comparison of numerical and experimental scaled

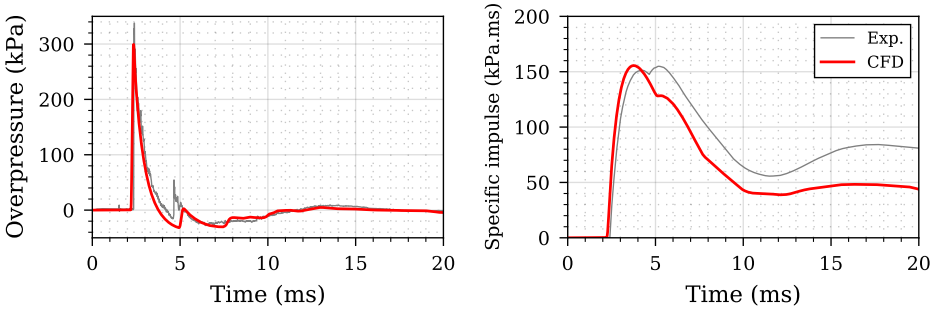

(a) 250g at 2m stand-off

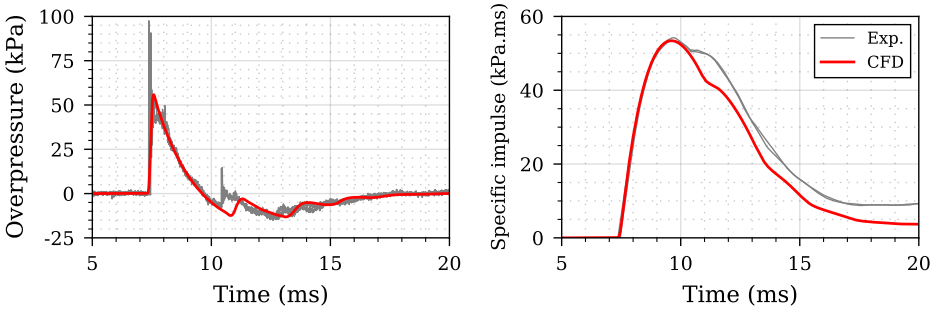

(b) 180g at 4m stand-off

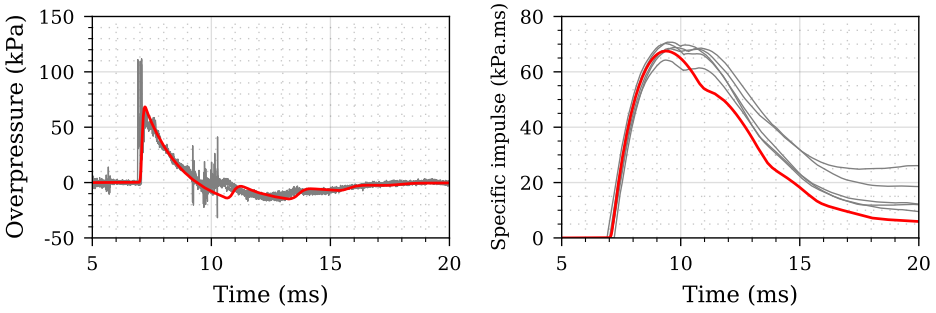

(c) 250g at 4m stand-off

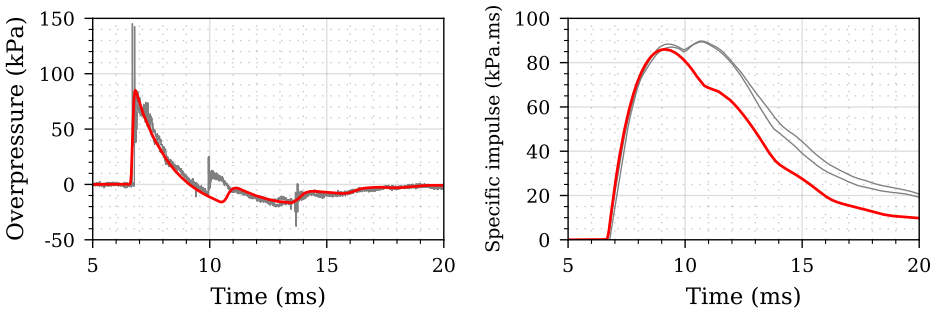

(d) 350g at 4m stand-off

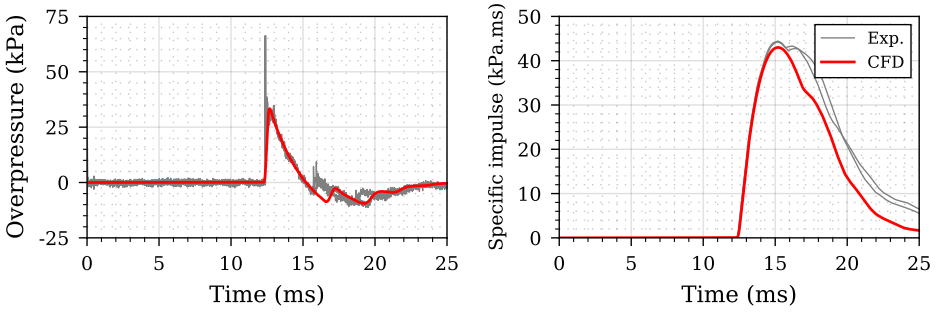

(e) 250g at 6m stand-off

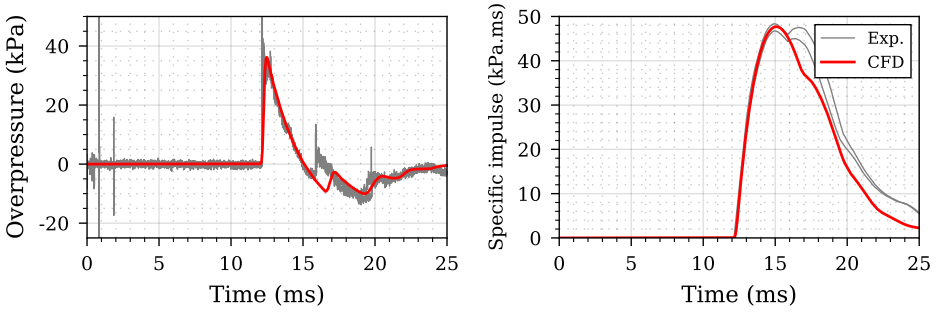

(f) 290g at 6m stand-off

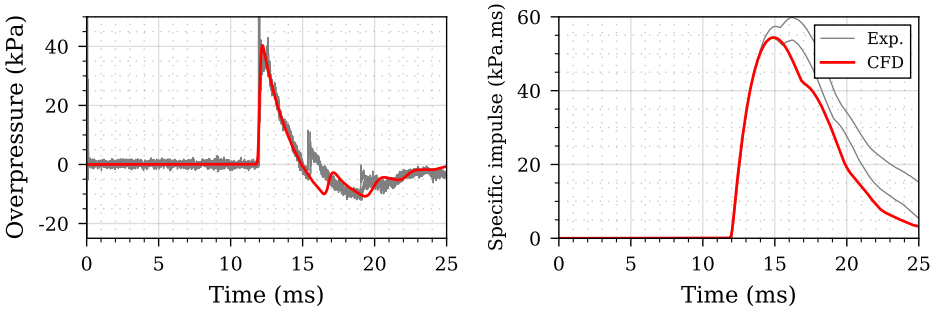

(g) 350g at 6m stand-off

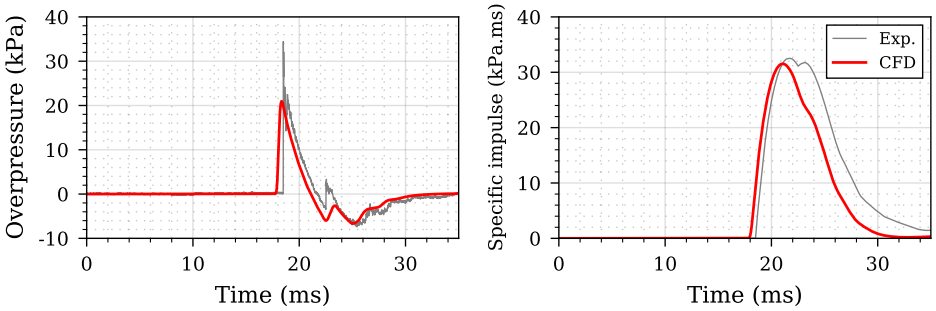

(h) 250g at 8m stand-off

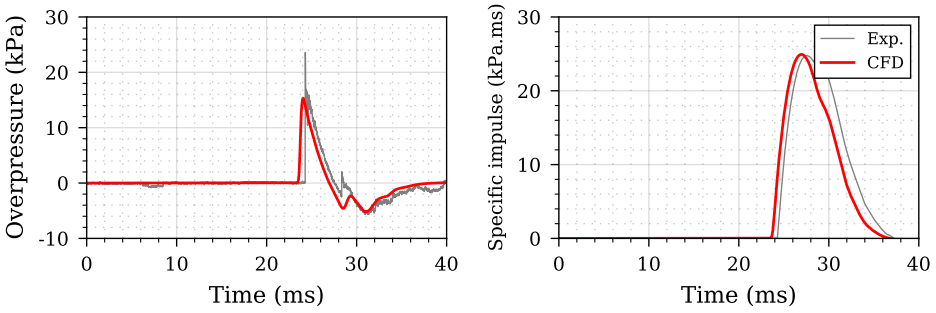

(i) 250g at 10m stand-off

Figure 3.11: Experimental validation of numerical overpressure and specific impulse histories for the 9 different experimental scenarios as summarised in Table 3.2. All experiments are PE4 hemispheres.

Table 3.2: Input parameters and meshing strategy used for validation models. Note, the B1D module was used as previously.

Stand-off (m)	**Charge mass (kg)**	**Zone length (mm)**	**Res. level**	**Ultimate cell length (mm)**	**S/cell length**
2	0.25	200	4	12.50	160
4	0.18	200	4	12.50	320
	0.25	200	4	12.50	320
	0.35	200	4	12.50	320
6	0.25	300	4	18.75	320
	0.29	300	4	18.75	320
	0.35	300	4	18.75	320
8	0.25	400	4	25.00	320
10	0.25	400	4	25.00	400

peak specific impulse values (divided by the cube-root of the charge mass) is presented in Figure 3.12. The numerically generated peak specific impulses can be seen to closely match the experimental data consistently across the entire range of scaled distance and therefore *Apollo* can be considered to provide accurate specific impulse values in this region of scaled distances, provided the mesh requirements outlined previously are satisfied.

3.6 Summary

This chapter has focussed on the theoretical background of fluid dynamics, and demonstrates typical approaches taken in computational fluid dynamics (CFD) codes. Within this thesis, the chosen CFD code was *Apollo*, this chapter demonstrates the numerical scheme that *Apollo* implements and describes additional features that improve computational

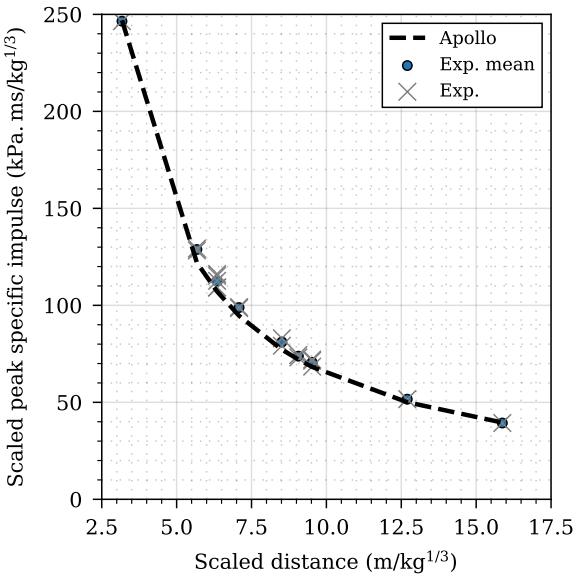

Figure 3.12: Validation of *Apollo* scaled peak specific impulse against experiments (Rigby et al. 2015a)

efficiency that are utilised in subsequent analyses to demonstrate why *Apollo* is a suitable choice.

Additionally, *Apollo* is validated in the far-field extensively against a range of experimental data following the presentation of a mesh sensitivity analysis where it can be seen *Apollo* provides highly accurate solutions that show good agreement with experimental data.

Chapter 4

Developing a preliminary surrogate model with data transformations

4.1 Introduction

As previously mentioned in Chapter 2, accurate appraisal of the viability of structures and protective systems following the close-in detonation of a high explosive is only possible if the analyst has an accurate model detailing both the distribution and magnitude of the imparted load (Rigby et al. 2019b). In these so-called extreme near-field scenarios, the loading is highly localised, whilst deformation and damage occur on timescales which are orders of magnitude longer than the duration of loading. It is clear, therefore, that the *temporal* form of the loading is of lesser importance, yet a description of the *spatial* distribution of the peak loading, i.e. peak specific impulse, is critical. Due to the limitations of using experimental approaches to inform a data-driven surrogate model (as mentioned in Section 2.2.7) an appropriate technique is proposed in this

thesis of using validated CFD analyses to create datasets from which surrogate models can be developed for near-field blast load prediction.

This chapter addresses the near-field validation of *Apollo* that is used to generate the datasets used throughout this thesis. It then demonstrates how this data can be used to create a preliminary surrogate model through the use of data transformations. To create this surrogate model it is found that: the peak perpendicular specific impulse can be modelled as a power law with respect to scaled distance and the normalised specific impulse with respect to angle of incidence can be modelled as a modified gaussian equation. The product of these two components allows the prediction of peak specific impulse for a given range of scaled distance and angle of incidence. The accuracy of this new approach is rigorously tested against the existing CFD dataset, unseen CFD analyses, and available experimental data from the literature. It is demonstrated that localised peak specific impulse and area-integrated total impulse (for a range of target sizes) are typically predicted to within 10%.

4.2 Numerical modelling

4.2.1 *Apollo Blastsimulator*

The in-built "auto-staging" procedure in *Apollo* was used which consists of 6 stages defined with scaled termination radii of 0.2, 0.5, 0.8, 1.25, 5 and infinity (in $\mathrm{m/kg^{1/3}}$). Additional resolution levels are defined for each stage, beginning with an additional 6 levels, therefore as the simulation progresses, when the disturbance reaches some termination criteria, the total (additional) resolution level will drop by 1, and the new domain size will grow by the defined growth rate. This process repeats until the additional resolution levels drop to 0 and the entire

domain is then included in the simulation. Successive stages are linked together through remapping of the flow field reached at the end of the prior stage onto the mesh of the successive stage.

Apollo has the option to progressively expand a computational domain according to a pre-defined velocity-time function (e.g. the rate of expansion of a freely expanding spherical blast wave), however this is not usually required for close-in detonations. PE4 is a plastic explosive comprising 87% RDX and 13% mineral oil binder (Tyas 2019). The explosives were modelled as 100 g PE4 spheres throughout the mesh sensitivity study, validation, and data harvesting aspect of this study, using *Apollo*'s in-built model for C4 explosive, as the two explosives are of nominally identical composition (Bogosian et al. 2016). The pressure-density-temperature relationship of the post-detonation explosive products is given by the Jones Wilkins Lee equation of state (Lee et al. 1968):

$$p(\rho, T) = C_1 e^{-R_1 \rho_0/\rho} + C_2 e^{-R_2 \rho_0/\rho} + \rho R T$$

where ρ and ρ_0 are density and initial density, R is a gas constant ($R = 365$ J/(kg.K) for PE4), and T is absolute temperature. The constants C_1, C_2, R_1, and R_2 are assigned the values 734.60 (GPa), 8.86 (GPa), 4.79 (-), and 1.06 (-) respectively for PE4. The air is modelled as a thermally perfect gas, $p = \rho RT$, such that under ambient conditions (288 °K), ambient pressure is $p = 101.3$ kPa. Ambient pressure has been subtracted from all pressure histories in this article to present *over*-pressures and impulses only. Afterburn has been shown to influence the development of loading in near-field blast scenarios (Tyas et al. 2016), and so was modelled using the Klomfass Afterburning (KAB) Model. The KAB model defines an empirical relationship for global combustion rate based on volume and rate of change of volume of the explosive fireball,

whereby the total mass of combustion products generated in a time-step is obtained (Fraunhofer EMI 2018). In all numerical simulations the explosives were centrally detonated and the mass of the detonator was not included. The Chapman-Jouguet detonation model was used, with the progression rate a function of: local burning velocity, defined by the Chapman-Jouguet condition; local detonation product density; and local un-reacted explosive density. This model allows the detonation front to propagate in a physically meaningful manner (Fraunhofer EMI 2018).

4.2.2 Mesh sensitivity and validation

Mesh sensitivity analysis

Prior to validating *Apollo* results against available experimental data, a mesh sensitivity study was conducted with the aims of determining the required element size to achieve convergence and identifying suitable combinations of zone length and resolution level.

A series of simulations were run for a 100 g PE4 sphere (24.3 mm radius) located at a stand-off distance, S, of 0.05 m from the centre of the charge to a rigid reflecting wall. This corresponds to a scaled distance of $Z = 0.108$ m/kg$^{1/3}$, which is the shortest scaled distance to be used in the subsequent modelling study. The domain was $0.35 \times 0.35 \times 0.35$ m, and quarter symmetry was utilised, with symmetry planes located in the directions orthogonal to the reflecting wall, originating at the centre of the charge. Numerical pressure gauges were placed along the rigid surface out to a distance of 0.28 m from the point normal to the explosive. 200 gauges in total were placed at evenly-spaced angles of incidence, θ, where the angle of incidence is defined as the angle between the outward normal of the reflecting surface and the direct vector from the explosive charge to that point (Rigby et al. 2015a). The gauge which is normal to the

Table 4.1: Ultimate cell length (element size at highest resolution level) and number of elements (between charge centre and normal gauge location) for initial mesh sensitivity study, Z = 0.108 m/kg$^{1/3}$. Note, the × 4 multiple accounts for the additional cells from the staging procedure.

	Zone length (mm) × 4		
Resolution Level	**50**	**20**	**10**
0	50	20	10
1	25	10	5
2	12.5	5	2.5
3	6.25	2.5	1.25
4	3.125		

(a) Ultimate cell length (mm). Ultimate cell length values should be divided by 4 to account for *Apollo*'s auto-staging procedure.

	Zone length (mm)		
Resolution Level	**50**	**20**	**10**
0	4	10	20
1	8	20	40
2	16	40	80
3	32	80	160
4	64		

(b) Number of cells between charge and target

explosive therefore has an angle of incidence of $\theta = 0°$, and the most remote gauge has an angle of incidence of $\theta = 80°$. A total of 13 models were run in the initial sensitivity analysis, with details of the meshes provided in Table 4.1. *Apollo*'s auto-staging was used throughout.

Results from the mesh sensitivity analysis are shown in Figure 4.1, where each sub-figure shows: peak specific impulse for the normal gauge ("peak" refers to the maximum specific impulse from the specific impulse-time history); area-integrated impulse out to a radius of 0.28 m from the normal gauge; and total analysis time (termed 'wall time' in *Apollo*), all plotted against the ratio of stand-off distance to ultimate cell length. The results suggest that an ultimate cell length of $S/50$ is sufficient to achieve

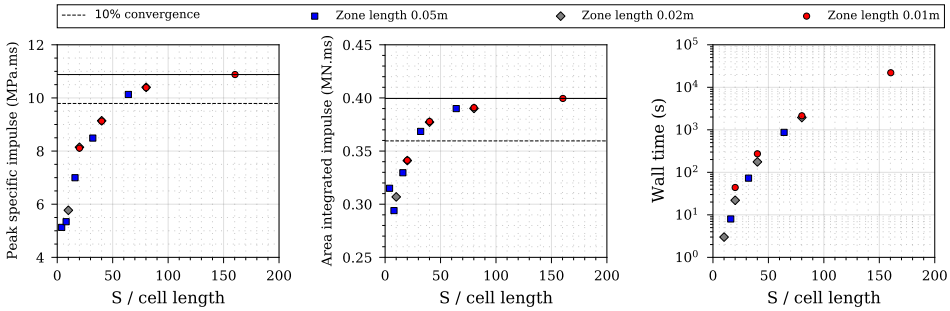

Figure 4.1: Mesh convergence study for $Z = 0.108$ m/kg$^{1/3}$

convergence for both specific and total impulse, whilst maintaining a practical analysis time: a decrease in cell length from $S/64$ to $S/160$ increases the total impulse by $< 3\%$, whilst increasing the wall time by approximately a factor of 25. The benefits of DMA, however, can be seen when considering wall time: at $S/40$, a clear reduction in analysis time is observed when going from a zone length of 0.01 m to 0.02 m. Therefore a meshing strategy of 50 mm zone length and resolution level of 4 will be implemented, extending to a 100 mm zone length and resolution level of 5 (maintaining an ultimate cell length of $> S/50$ throughout) to accommodate a number of larger domain sizes required to populate the CFD dataset in subsequent sections of this paper.

Experimental validation

The *Characterisation of Blast Loading* apparatus (Clarke et al. 2015) was designed to measure spatial and temporal distribution of loading from buried (Rigby et al. 2016) and free air explosions (Rigby et al. 2015b). The apparatus utilises two perpendicular arrays of Hopkinson (1914) pressure bars (HPBs), mounted such that their faces lie flush with the underside of a large (100 mm thick, 1400 mm diameter) steel plate, which acts as a nominally rigid reflecting surface under which the

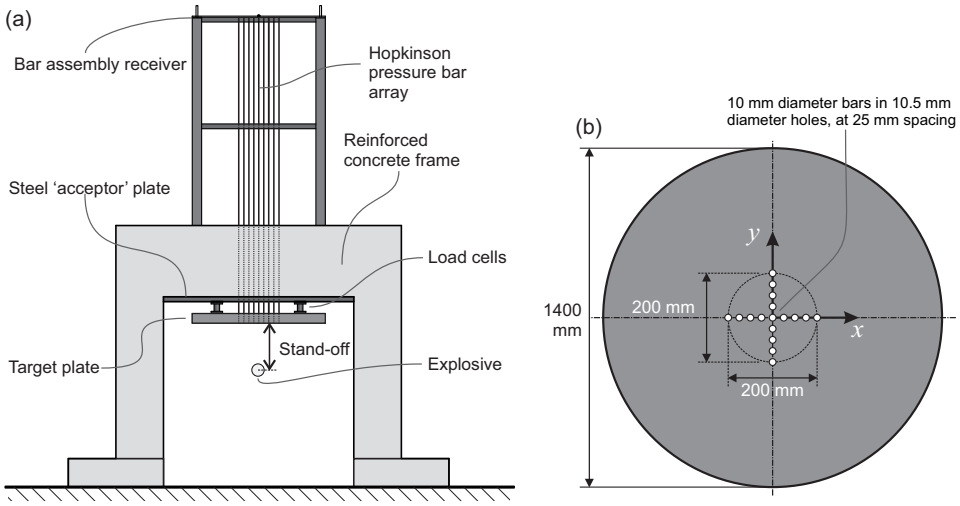

Figure 4.2: Schematic of the University of Sheffield's *Characterisation of Blast Loading* apparatus (not to scale): (**a**) elevation; (**b**) detailed plan view of target plate showing bar arrangement and coordinate axes. Figure taken from Rigby et al. (2019a)

explosive charges are detonated. 17 HPBs are used in total: one bar is located directly above the charge (which itself is aligned with the centre of the target plate); and four bars are located at each distance of 25 mm, 50 mm, 75 mm and 100 mm from the centre, an experimental schematic for this arrangement is shown in Figure 4.2. Numerical integration of the recorded pressure signals enables peak specific impulse distributions over the central 200 mm diameter region of the target to be generated. Two different configurations were used to validate *Apollo*: 100 g PE4 spheres were used throughout, with three tests at 80 mm stand-off (Rigby et al. 2019a) and three tests at 380 mm stand-off (Rigby et al. 2020), corresponding to scaled distances of 0.172 m/kg$^{1/3}$ and 0.819 m/kg$^{1/3}$ respectively.

Apollo models were run with 50 mm zone length and resolution level of 4 as informed by findings from the mesh sensitivity analysis, and utilising the auto-staging procedure with the first stage resolved in 1D. Thereby satisfying the S/50 constraint as mentioned in the mesh

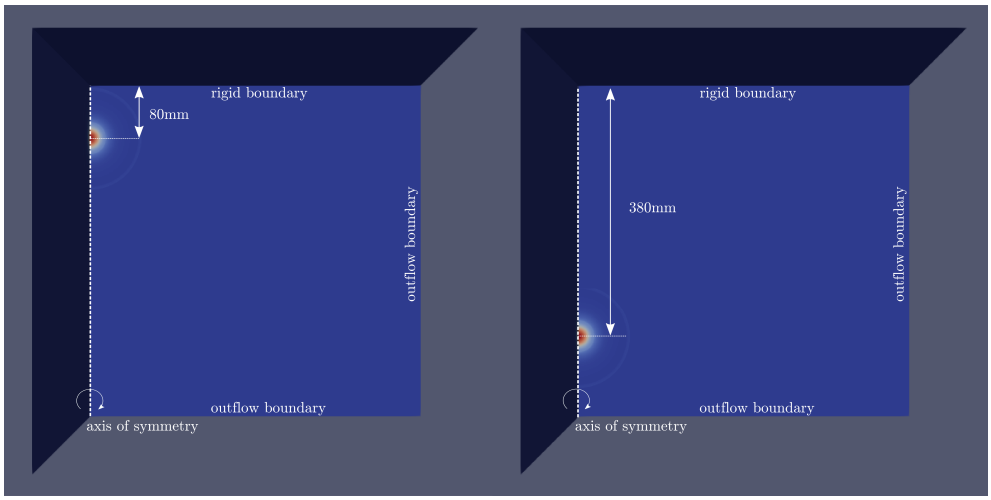

Figure 4.3: Experimental configuration of 100g spherical charge scenarios modelled in *Apollo*. 80mm standoff (left) and 380mm standoff (right).

sensitivity study. The reflecting surface was modelled as a rigid wall and quarter symmetry was used as before. The domain was extended to $1.5 \times 1.5 \times 1.5$ m to prevent expansion waves from the edge of the domain reaching the gauge locations, and to establish a consistent domain size for use in subsequent models. Experimental configurations are shown in Figure 4.3. Numerical overpressure and specific impulse histories at 0, 25, and 50 mm from the target centre are compared against those from one of the 80 mm stand-off experiments in Figure 4.4. It can be seen that the magnitudes and general form of the numerical pressure and impulse histories match the experiments closely, which gives confidence that *Apollo* is correctly modelling the mechanisms of normal and oblique reflection in near-field blast scenarios.

Additionally, Figure 4.5 shows compiled peak specific impulse distributions for the 80 mm stand-off and 380 mm stand-off tests with the CFD data evaluated at 0.4° increments of angle of incidence. Also shown are the mean experimental recordings at each measurement location for ease of comparison (three data points at the target centre and 12 data

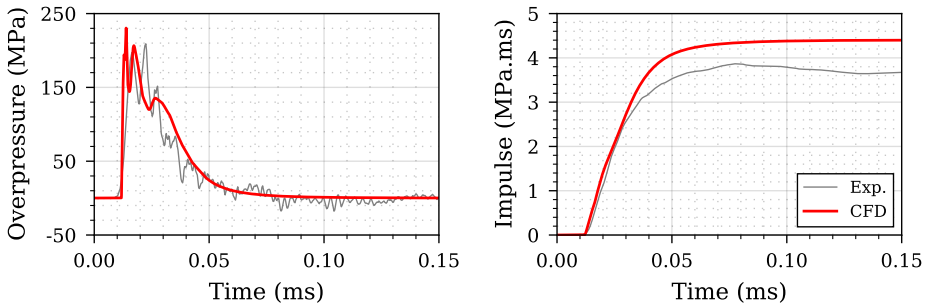

(a) 0 mm from centre, $\theta = 0°$

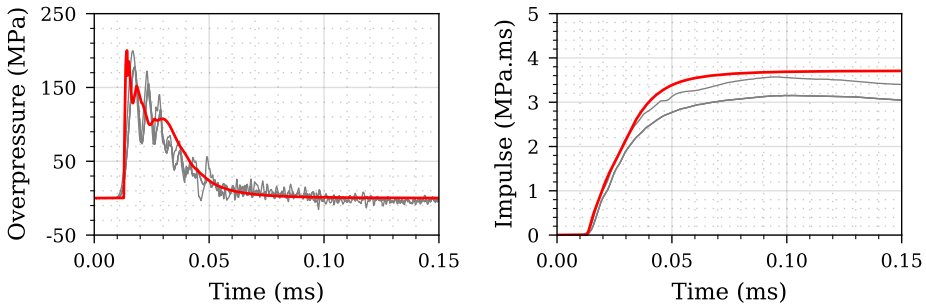

(b) 25 mm from centre, $\theta = 17°$

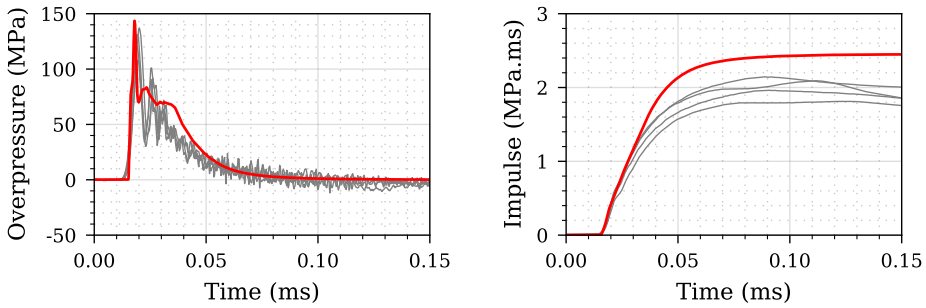

(c) 50 mm from centre, $\theta = 32°$

Figure 4.4: Experimental validation of numerical overpressure and specific impulse histories for $Z = 0.172$ m/kg$^{1/3}$ as 0, 25, and 50 mm distance from the target centre

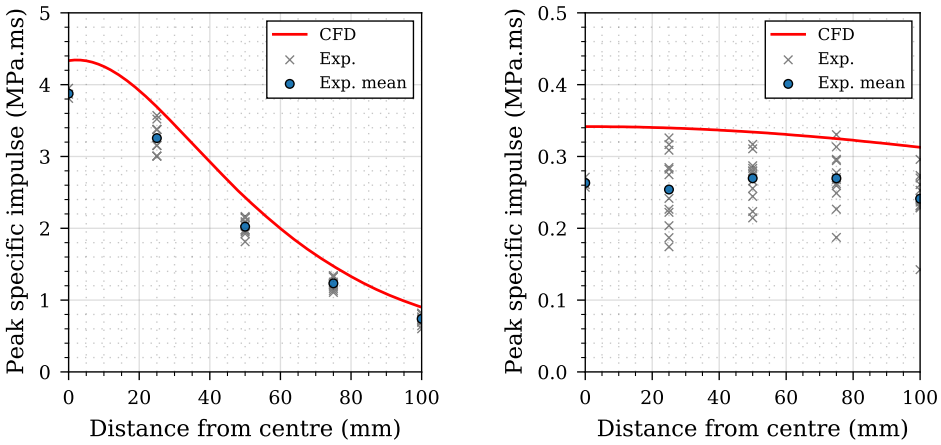

(a) 80 mm stand-off, $Z = 0.172$ m/kg$^{1/3}$ (b) 380 mm stand-off, $Z = 0.819$ m/kg$^{1/3}$

Figure 4.5: Experimental validation of numerical specific impulse distributions at 0–100 mm distance from the target centre

points at each other distance from the centre). Again, the numerical results closely match the experimental data, in both form and magnitude, although the results appear to be in better agreement with the upper bound of the experimental results. Deriving a model based on *Apollo* CFD results will therefore show conservatism, however this is still representative of the typical variations seen in blast loading from close-in detonations, as exhibited here.

4.2.3 Generation of dataset

Apollo was used to generate specific impulse distributions along a rigid reflecting surface, at angles of incidence between 0–80°, with 200 numerical gauges placed at regular increments of angle of incidence as per the validation exercise. The analyses were formed into two groupings. Firstly, an initial dataset was generated for extreme near-field conditions, 0.11–0.21 m/kg$^{1/3}$, at increments of ∼0.025 m/kg$^{1/3}$. Here, the domain size was fixed at 1.5 × 1.5 × 1.5 m, using quarter symmetry, with 100 mm

zone length and resolution level of 5. This results in the same ultimate cell length as the validation model and recommendations from the mesh sensitivity study whilst maximising computational efficiency through the use of the DMA feature in *Apollo* and the auto-staging procedure.

The dataset was then expanded to incorporate intermediate near-field conditions, with analyses run between 0.24–0.55 m/kg$^{1/3}$, again at increments of ∼0.025 m/kg$^{1/3}$. The full dataset comprised 18 analyses in total, with 3600 data points between 0.11–0.55 m/kg$^{1/3}$ and 0–80°. For the $Z > 0.21$ m/kg$^{1/3}$ analyses, a separate domain was specified for each model, with equal side lengths of an integer number of zone lengths, given by: $\lceil 1.3S\tan(80°)/L \rceil$, where S and L are stand-off and zone length as previously, such that $S\tan(80°)$ is the distance from the centre of the target to the most remote gauge. The domain length was set a factor of 1.3 greater than this distance (rounded up to the nearest multiple of zone length), again to prevent edge expansion waves from reaching the most remote gauge location during the analysis. As with previous models, quarter symmetry was used, with 100 mm zone length and resolution level 5 and *Apollo's* auto-staging procedure was used.

CFD results are shown in detail in Figure 4.6 for the reduced and full datasets. Here, a Savitzky-Golay filter (Savitzky & Golay 1964) has been applied to the specific impulse distributions to remove small spurious oscillations induced by DMA, which were more apparent in the analyses at larger scaled distances. The Savitzky-Golay filter is a low-pass data-smoothing method based on local least-squares polynomial approximation, and is particularly suited for removing high frequency noise with the advantage of preserving the original shape and features of the signal (Acharya et al. 2016).

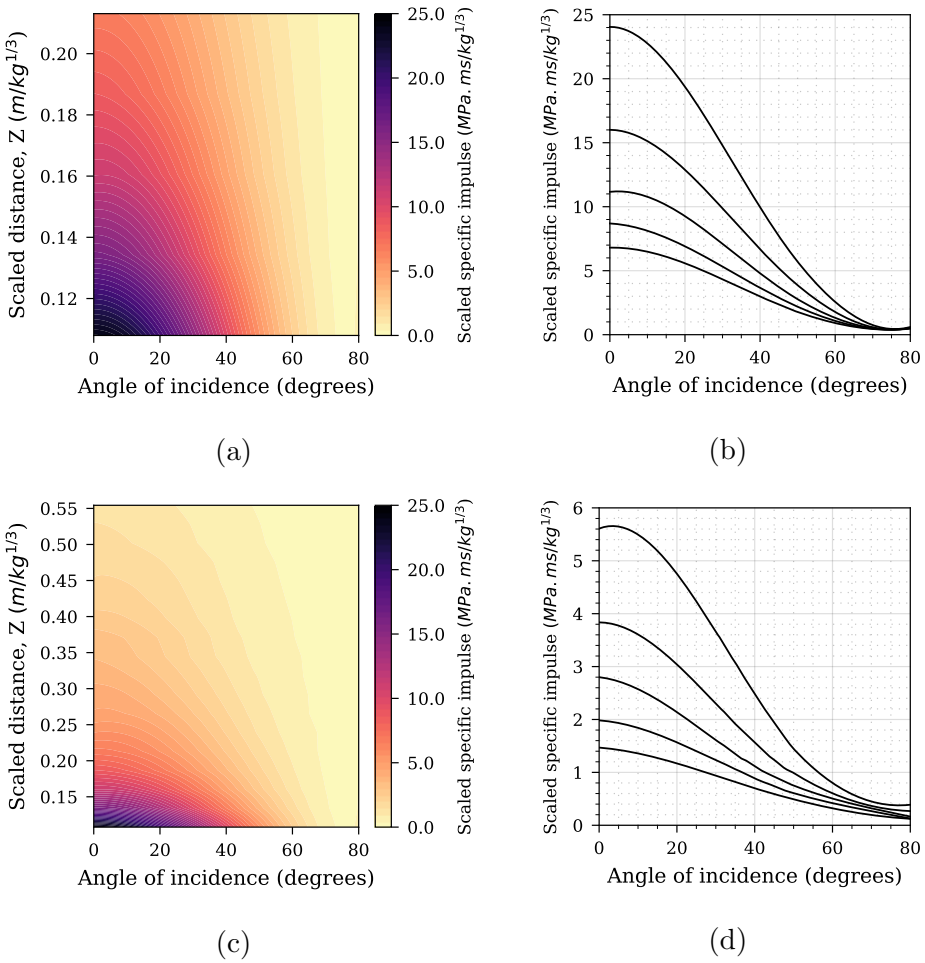

Figure 4.6: CFD results: (a) filled contours of scaled peak specific impulse for reduced dataset, 0.11–0.21 m/kg$^{1/3}$; (b) distributions of peak specific impulse for 0.11, 0.13, 0.16, 0.19 and 0.21 m/kg$^{1/3}$; (c) filled contours of scaled peak specific impulse for full dataset, 0.11–0.55 m/kg$^{1/3}$; (d) distributions of peak specific impulse for 0.24, 0.32, 0.40, 0.48 and 0.55 m/kg$^{1/3}$

4.3 Development of a data-driven predictive approach

4.3.1 Overview

This section details the development of a data-driven predictive approach for computing near-field specific impulse distributions from spherical PE4 charges. Firstly, the relationship between peak specific impulse and angle of incidence is examined, and three different curve types are trialled to establish the most suitable candidate for representing the spatial variation of loading. Secondly, the relationship between scaled distance and peak specific impulse is considered and a power law fit is determined. The proposed predictive method is given as the product of these two relationships: spatial variation, and variation with scaled distance.

It was noted by Wilson et al. (2018) that "*further research into analysis scenarios at varying angles of incidence are required... at a scaled distance less than 0.16 m/kg*$^{1/3}$", hence initial efforts are focussed on the reduced dataset, and then extended to the full dataset.

4.3.2 Spatial variation of normalised impulse

From examination of Figure 4.6b and 4.6d, it appears as though that the relationship between peak specific impulse and angle of incidence is similar for all scaled distances considered in this article. Dividing each through by the respective peak specific impulse at 0°, i.e. the normally reflected peak specific impulse, i_r, allows the normalised peak specific impulse distribution to be plotted, as shown in Figure 4.7. Here, it can be seen that the curves form a consistent grouping, particularly for the reduced dataset, which suggests that it is possible to define a single curve

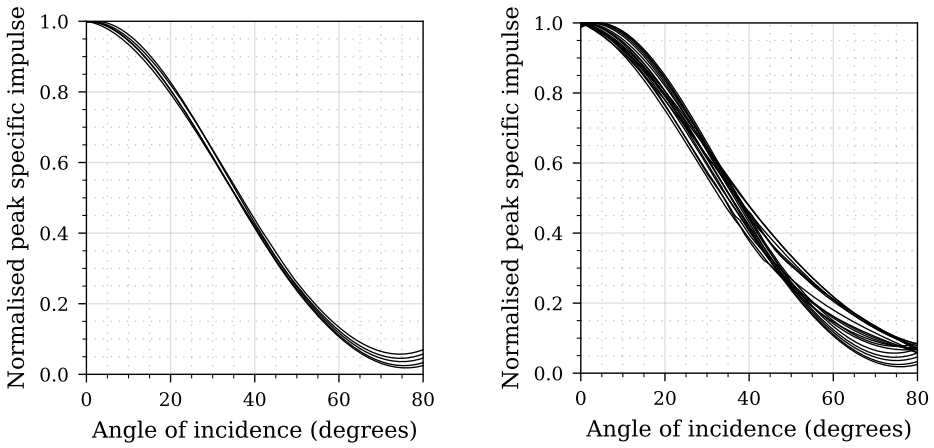

(a) (b)

Figure 4.7: Normalised peak specific impulse distributions: (a) reduced dataset, 0.11–0.21 m/kg$^{1/3}$; (b) full dataset, 0.11–0.55 m/kg$^{1/3}$

to describe the shape of the distributed specific impulse with respect to angle of incidence. Three candidate curve-types have been identified: the semi-analytical trigonometric functions of Henrych (1979), and Randers-Pehrson & Bannister (1997); and the Gaussian function proposed by Pannell et al. (2019).

The semi-analytical work of Henrych (1979), later developed by Remennikov et al. (2017), is based on the model of instantaneous detonation of a spherical explosive and subsequent expansion of the gaseous detonation products into a vacuum. Here the loading is assumed to develop entirely from momentum transferred by the detonation products as they impact the loaded face. With some cancelling of constants, the Henrych (1979) equations can be expressed in normalised form:

$$\bar{i}(\theta) = \cos^4 \theta \tag{4.1}$$

where $\bar{i}$ is specific impulse divided by normally reflected specific impulse, and θ is angle of incidence as previously.

Randers-Pehrson & Bannister (1997) present a trigonometric expression to account for angle of incidence effects, which after some manipulation can be presented as a function of normalised specific impulse and angle of incidence:

$$\bar{i}(\theta) = A_R(\theta, Z)\cos^2\theta + A_S(\theta, Z)[1 + \cos^2\theta - 2\cos\theta] \tag{4.2}$$

Here, $A_R = i_r(\bar{Z})/i_r(Z)$ and $A_S = i_s(\bar{Z})/i_r(Z)$, where i_s is incident specific impulse, Z is normal scaled distance as previously, and $\bar{Z}$ is *slant* scaled distance to the point of interest: $\bar{Z} = Z/\cos(\theta)$. Similarly, slant stand-off distance can be defined as $\bar{S} = S/\cos(\theta)$. For a given normal scaled distance, A_R and A_S vary non-linearly with slant distance and hence these coefficients must be calculated separately for each and every angle of incidence under consideration. Furthermore, this relationship is also scaled-distance dependent.

Finally, Pannell et al. (2019) presented work on using machine learning to fit Gaussian probability density functions to near-field specific impulse distributions. Here, a modified version of the equations as a function of angle of incidence is presented:

$$\bar{i}(\theta) = \exp\left(\frac{-\theta^2}{A_G}\right) \tag{4.3}$$

where A_G is the Gaussian width parameter, which controls the rate of decay of the Gaussian curve with respect to angle of incidence. Two other parameters form a traditional Gaussian function, i.e. those which determine the height and position of the centre of the peak respectively. Due to the normalised and rotationally-symmetric nature of the loading considered here, the height coefficient is unit throughout, and the position coefficient is zero as each Gaussian curve is centred around $\theta = 0°$. Whilst

the loading is considered rotationally-symmetric, the method can still be used to model off-centre charge placements relative to the target surface, as detailed in Section 4.4.2.

Figure 4.8 shows an assessment of the quality of fit for each of the normalised predictions compared against the reduced CFD dataset. For the Randers-Pehrson & Bannister (1997) model, A_R and A_S were determined at every angle of incidence (200 values between 0 and 80°) for each scaled distance using the 5th-order and 11th-order polynomial equations for $i_r(Z)$ and $i_s(Z)$ provided in Table 4-2 of Shin et al. (2014b). Negligible difference was found between the A_R and A_S vs. θ relationships averaged across all scaled distances, and the A_R and A_S vs. θ relationship for the *median* scaled distance ($Z = 0.16$ m/kg$^{1/3}$) of the reduced dataset. This low dependency of A_R and A_S with Z is to be expected for this particular grouping of scaled distances owing to the similarity of the normalised impulse curves as shown in Figure 4.7a. Even given this simplification, however, the process for determining A_R and A_S is relatively involved and required a minimum of 200 computations.

For the Pannell et al. (2019) Gaussian model, A_G was determined through a non-linear regression analysis in which the root mean squared error (RMSE) was minimised using a trust region reflective method (Branch et al. 1999). For the reduced dataset, the regression-estimated best fit yielded an A_G value of 1829.

Also shown in Figure 4.8 are the model residuals (CFD value minus model prediction), with mean absolute error (MAE) and RMSE provided for each fit against the mean CFD distribution. The Henrych (1979) and Randers-Pehrson & Bannister (1997) trigonometric identities generally under-predict specific impulse for intermediate angles of incidence (which would result in a degree of unconservatism when considering to-

tal impulse) whereas the Gaussian function (Figure 4.8c) can be seen to closely match the CFD data for all angles of incidence. The Gaussian model also exhibits the best quantitative agreement when considering both MAE and RMSE and therefore has been selected to represent the spatial distribution of specific impulse in the predictive methodology proposed in this article.

The Gaussian function was subsequently assessed against the full dataset, Figure 4.9, again with A_G determined through non-linear regression ($A_G = 2007$ for the full dataset). Although the spread of normalised CFD curves is larger for the full dataset, the Gaussian function can again be seen to represent the mean CFD data to a high degree of accuracy and is therefore suitable for use in the entire range of near-field scaled distances, $0.11 \leq Z \leq 0.55$ m/kg$^{1/3}$.

4.3.3 Variation with scaled distance and final models

With a method for predicting distribution of normalised peak specific established, it still remains to quantify the variation of peak normally reflected specific impulse with scaled distance. Rather than making use of the existing higher order polynomial expressions of Kingery & Bulmash (1984), Swisdak (1994), or Shin et al. (2014b), the aim here is to develop a simple equation with a small number of coefficients over a more focussed range of scaled distances. Accordingly, the relationship between scaled distance and peak normally reflected specific impulse (i.e. the peak specific impulse at $\theta = 0^\circ$ for each scaled distance) is plotted in Figure 4.10. Here, it can be seen that a power law represents the data well, with statistically robust fits of $R^2 = 0.999$ and $R^2 = 0.995$ for the reduced and full datasets respectively.

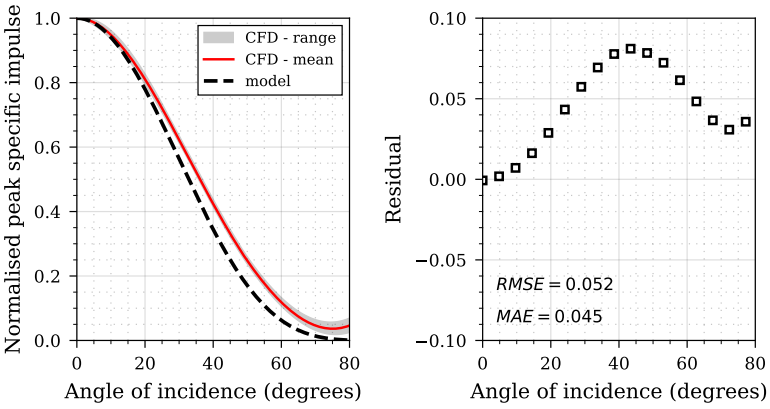

(a)

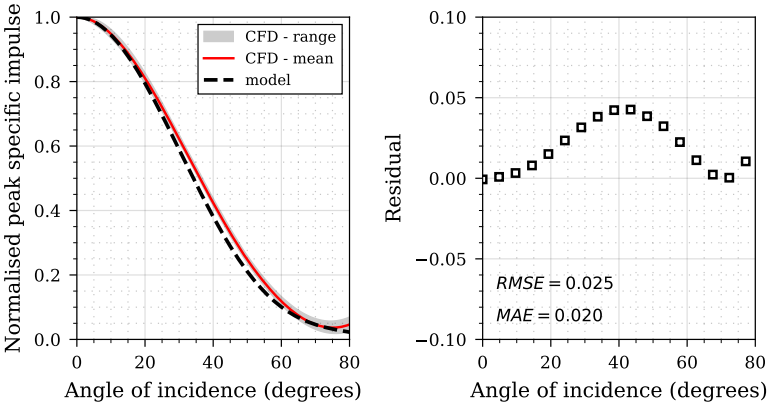

(b)

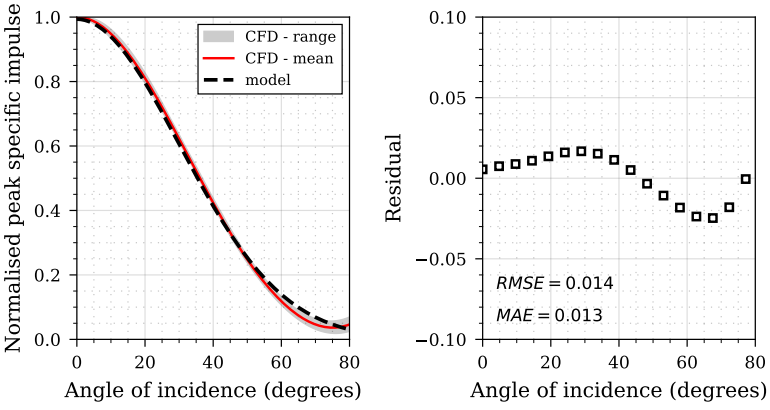

(c)

Figure 4.8: Normalised specific impulse distributions for reduced dataset (0.11–0.21 m/kg$^{1/3}$) with associated residuals and error assessment: (a) Henrych (1979) trigonometric model, equation 4.1; (b) Randers-Pehrson & Bannister (1997) trigonometric model, equation 4.2; (c) Pannell et al. (2019) Gaussian model, equation 4.3

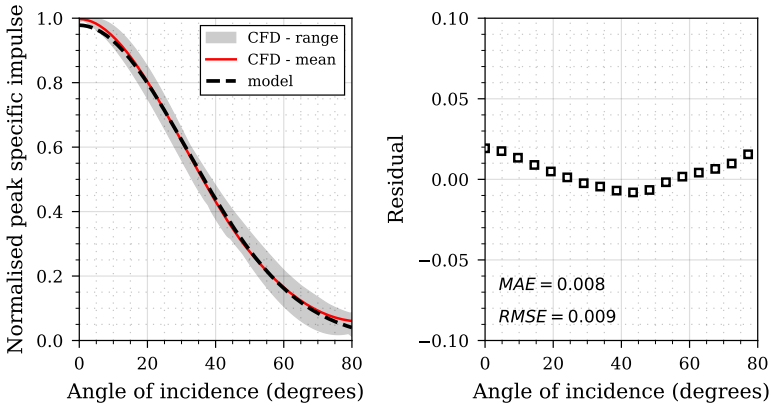

Figure 4.9: Normalised specific impulse distributions for full dataset (0.11–0.55 m/kg$^{1/3}$) with associated residuals and error assessment for the Pannell et al. (2019) Gaussian model, equation 4.3

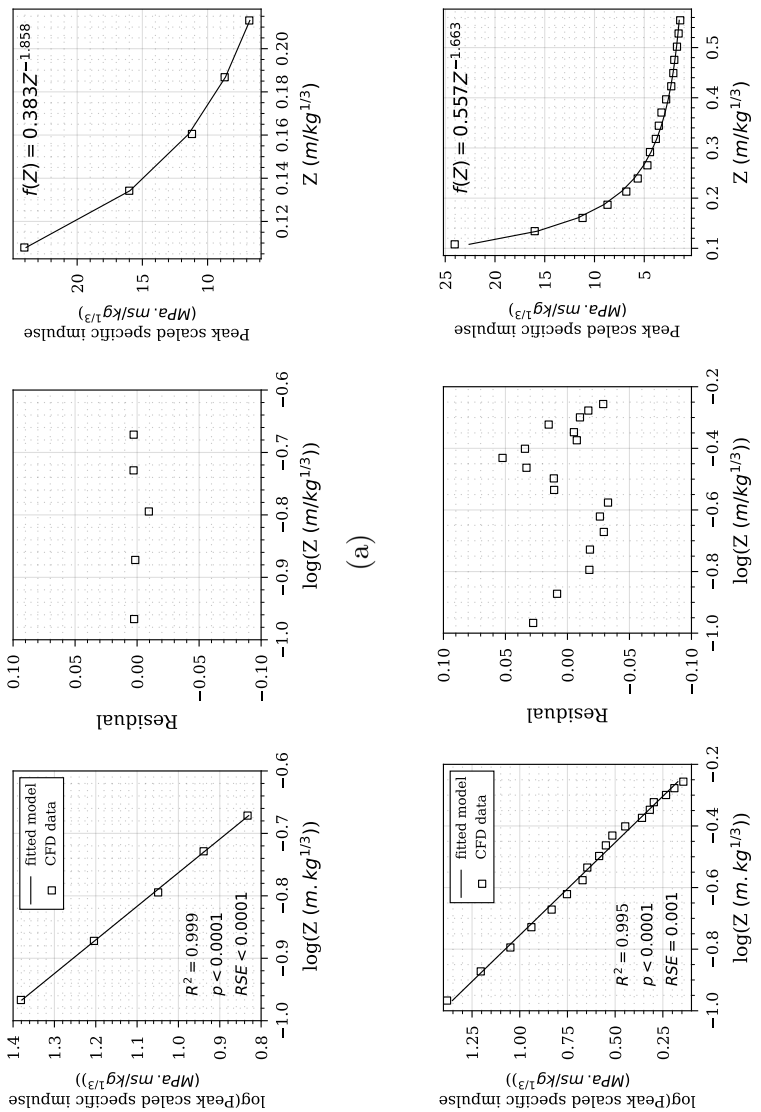

Figure 4.10: Relationship between peak normally reflected specific impulse and scaled distance, residuals, and proposed fit: (a) reduced dataset, 0.11–0.21 m/kg$^{1/3}$; (b) full dataset, 0.11–0.55 m/kg$^{1/3}$

Accordingly, the derived relationship for specific impulse, scaled using Hopkinson (1915) and Cranz (1926) scaling to express the relations as a function of charge mass, W (kg), are:

$$i(Z,\theta,W) = 0.383Z^{-1.858}\exp\left(\frac{-\theta^2}{1829}\right)W^{1/3} \tag{4.4}$$

for $0.11 \leq Z \leq 0.21$ m/kg$^{1/3}$, and:

$$i(Z,\theta,W) = 0.557Z^{-1.663}\exp\left(\frac{-\theta^2}{2007}\right)W^{1/3} \tag{4.5}$$

for $0.11 \leq Z \leq 0.55$ m/kg$^{1/3}$, where both expressions return specific impulse in units of MPa.ms.

4.4 Assessment of proposed method

4.4.1 Evaluation against numerical dataset

This section presents a rigorous assessment of the quality of the proposed predictive method by comparing the results against the full CFD dataset. Equations 4.4 and 4.5 were solved for $0° \leq \theta \leq 80°$ between the limits of $0.11 \leq Z \leq 0.21$ m/kg$^{1/3}$ and $0.11 \leq Z \leq 0.55$ m/kg$^{1/3}$ respectively, in order to evaluate the accuracy of the predicted specific impulses against the entire CFD dataset. Filled contours of scaled peak specific impulse are shown for the reduced dataset in Figure 4.11a, and the full dataset in Figure 4.11c, with associated residuals shown in Figure 4.11b and 4.11d.

It can be seen that the residuals are generally positive, particularly for the reduced dataset. This is acceptable given the slight conservatism shown in the CFD data itself (see Figure 4.5). The predictions for the full dataset show slight under-predictions for small scaled distances and $\theta < 30°$, which is deemed acceptable given the availability of the re-

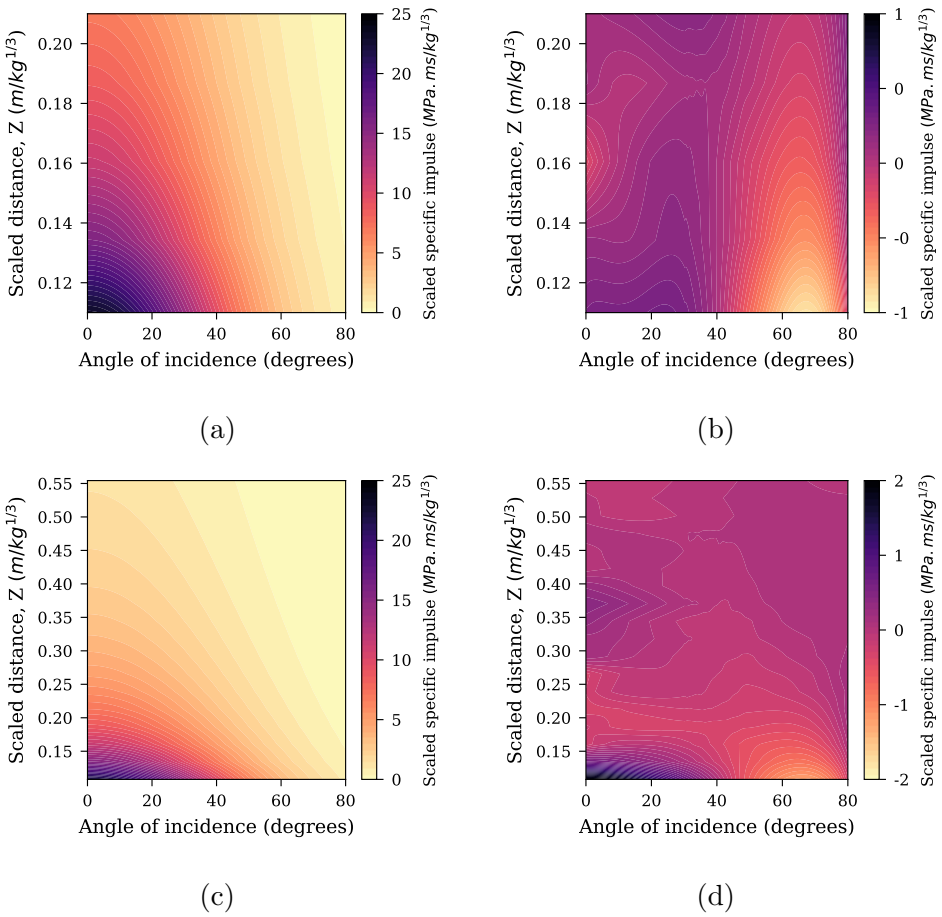

Figure 4.11: Predicted specific impulse distributions: (a) filled contours of scaled peak specific impulse for reduced dataset, 0.11–0.21 m/kg$^{1/3}$; (b) residual errors for reduced dataset; (c) filled contours of scaled peak specific impulse for full dataset, 0.11–0.55 m/kg$^{1/3}$; (d) residual errors for full dataset

duced dataset predictions which encompass a smaller range of scaled distances but at higher accuracy. The maximum absolute residuals are 0.922 MPa.ms/kg$^{1/3}$ for the reduced dataset, and 1.944 MPa.ms/kg$^{1/3}$ for the full dataset, and the mean absolute residuals are 0.198 MPa.ms/kg$^{1/3}$ for the reduced dataset, and 0.184 MPa.ms/kg$^{1/3}$ for the full dataset. The residual errors are considerably smaller than the maximum specific impulse (25 MPa.ms/kg$^{1/3}$), and indicate that the predictive equations are representative of the entire CFD dataset.

Whilst knowledge of the distributed specific impulse is a key requirement for determining structural response (Rigby et al. 2019b), total impulse (integral of specific impulse with respect to area, often termed 'area-integrated' impulse) is a useful metric particularly when considering global momentum transfer. Equations 4.4 and 4.5 were numerically integrated with respect to area for scaled target radii of 0.10, 0.25, 0.50, and 1.00 m/kg$^{1/3}$, and for scaled distances of 0.11, 0.20, 0.30, 0.40, and 0.50 m/kg$^{1/3}$. These values are provided in Table 4.2, alongside area-integrated CFD data and percentage differences between the two methods. For the two cases where $Z = 0.11$ m/kg$^{1/3}$ and $Z = 0.20$ m/kg$^{1/3}$, predictions were generated for both equations (4.4 and 4.5) in order to comment on the relative accuracy of the two approaches.

Generally the predictions are in excellent agreement with the CFD results, particularly for $Z \leq 0.20$ m/kg$^{1/3}$ when using equation 4.4, where the maximum deviation is 4% from the CFD data and the typical deviation is $\pm 1\%$. When the relationship in equation 4.5 is used to calculate total impulse for $Z \leq 0.20$ m/kg$^{1/3}$ the results are less accurate, with a maximum deviation of 12% from the CFD data. However, when equation 4.5 is used to make predictions for $Z > 0.20$ m/kg$^{1/3}$ a similar level of accuracy to equation 4.4 is obtained, with a maximum deviation of 6% from the CFD data, and a typical deviation of $\pm 4\%$. Thus, the predictive methods are considered to be highly representative of the CFD data from which they are derived.

4.4.2 Comparison against unseen CFD data

The predictive model was compared against two sets of unseen CFD data (i.e. data not used in the development of the predictive model). The first set comprises two models: a 5 kg spherical PE4 charge detonated at

Table 4.2: Comparison of total area-integrated impulse from numerical integration of CFD data and predicted peak specific impulse distributions for various scaled target radii

Scaled distance ($m/kg^{1/3}$)	**Scaled target radius ($m/kg^{1/3}$)**	**Total scaled impulse**			**Model eqn.**
		CFD (MN.ms/kg)	**Prediction (MN.ms/kg)**	**% diff.**	
0.11	0.10	0.719	0.710	-1	(4.4)
	0.25	3.964	3.905	-1	
	0.50	17.88	17.95	0	
	1.00	22.57	23.48	4	
0.20	0.10	0.236	0.233	-1	
	0.25	1.299	1.283	-1	
	0.50	6.044	5.886	-3	
	1.00	7.812	7.705	-1	
0.11	0.10	0.719	0.664	-8	(4.5)
	0.25	3.964	3.690	-7	
	0.50	17.883	17.820	0	
	1.00	22.57	23.78	5	
0.20	0.10	0.234	0.244	4	
	0.25	1.295	1.356	5	
	0.50	6.039	6.556	9	
	1.00	7.790	8.747	12	
0.30	0.10	0.130	0.124	-4	
	0.25	0.705	0.692	-2	
	0.50	3.212	3.340	4	
	1.00	4.245	4.458	5	
0.40	0.10	0.082	0.077	-6	
	0.25	0.446	0.429	-4	
	0.50	2.095	2.070	-1	
	1.00	2.886	2.763	-4	
0.50	0.10	0.052	0.053	3	
	0.25	0.283	0.296	5	
	0.50	1.455	1.428	-2	
	1.00	2.033	1.906	-6	

0.291 m normal stand-off distance (to charge centre) from a rigid surface ($Z = 0.17$ m/kg$^{1/3}$); and a 250 kg spherical PE4 charge detonated at 2.52 m normal stand-off distance (to charge centre) from a rigid surface ($Z = 0.40$ m/kg$^{1/3}$). Note that these models were not used to form the CFD datasets outlined previously, and hence serve as unseen data for comparative purposes. In the first model, a $2.1 \times 2.1 \times 2.1$ m domain was discretised into 0.1 m zone lengths with resolution level 4. In the second model, a $19.0 \times 19.0 \times 19.0$ m domain was discretised into 1.0 m zone lengths with resolution level 5. Both meshes resolve the first stage in 1D and use the auto-staging module, therefore satisfy the limit of $S/50$, as determined in the mesh sensitivity study.

Figure 4.12 shows a comparison of the CFD and predicted specific impulse distributions out to 0.25 m and 2.50 m from the target centre, with select values presented in Table 4.3 alongside percentage differences between CFD and predicted values. Here, predictions for the $Z = 0.17$ m/kg$^{1/3}$ case were determined using equation 4.4, and predictions for the $Z = 0.40$ m/kg$^{1/3}$ case were determined using equation 4.5. It can be seen that the predicted distributions closely match those from the CFD analyses, albeit with the predictive model being slightly conservative by design. The quantitative agreement is excellent throughout: the maximum deviation between CFD and prediction is 11% of the CFD specific impulse, at 1875 mm from the centre of the target for the $Z = 0.40$ m/kg$^{1/3}$ case ($\theta = 36.7°$). This exercise also validates the use of Hopkinson (1915) and Cranz (1926) scaling.

Detonation of explosive charges aligned with the centre of a target plate presents the worst-case scenario in terms of imparted loading and resulting plate deformation/failure (Chung Kim Yuen & Nurick 2005, Chung Kim Yuen et al. 2016). However, there may be situations where

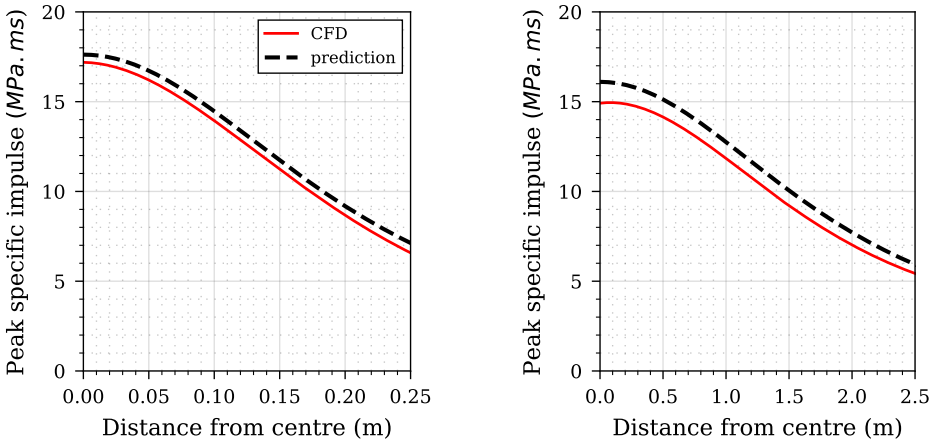

(a) (b)

Figure 4.12: CFD and predicted peak specific impulse distributions: (a) 5 kg at 0.291 m stand-off, $Z = 0.17$ m/kg$^{1/3}$; (b) 250 kg at 2.52 m stand-off, $Z = 0.40$ m/kg$^{1/3}$

Table 4.3: Comparison of unseen CFD data and predicted peak specific impulses at various distances from the centre of the target

Charge mass (kg)	**Stand-off (m)**	**Distance from centre (mm)**	**Angle of incidence (°)**	**Peak specific impulse**		
				CFD (MPa.ms)	**Prediction (MPa.ms)**	**% diff.**
5	0.291	0	0	17.19	17.62	3
		62.5	12.5	15.65	16.25	4
		125.0	23.2	12.58	13.11	4
		187.5	32.8	9.23	9.78	6
		250.0	40.7	6.48	7.12	10
250	2.520	0	0	14.93	16.10	8
		625	13.9	13.63	14.60	7
		1250	26.4	10.46	11.39	9
		1875	36.7	7.42	8.25	11
		2500	44.8	5.37	5.94	10

an off-centre detonation is of interest to the analyst, and therefore the purpose of the following exercise is to demonstrate generality of the predictive model in terms of charge placement.

According to the terminology in Figure 4.13a, angle of incidence at any point on the target plate can be calculated knowing the lateral distance from the point of interest to the normal impingement point, r, and the normal stand-off distance, S:

$$\theta = \arctan\left(r/S\right) \tag{4.6}$$

Where, in the previous examples described in this article, the normal impingement point was coincident with the plate centre, thus $r = \theta = 0$ at that point.

The second set of unseen CFD comprises one model: a 5 kg spherical PE4 charge detonated at 291 mm normal stand-off distance (normal scaled distance, $Z = 0.17$ m/kg$^{1/3}$) from a 200 mm square plate, with the charge offset from the plate centre by 50 mm in both in-plane directions, as in Figure 4.13b. Here, the mesh specifications were the same as those in the previous 5 kg model. Noting that the plate centre is labelled as $(0, 0)$, and the normal impingement point is labelled as $(50, 50)$, $r = 70.7$ mm and $\theta = 13.7°$ at the plate centre.

Figure 4.14 shows a comparison between the CFD results and Gaussian model predictions using the Gaussian model for the off-centre charge example, evaluated on a grid of 1×1 mm elements. It can be seen that a high level of qualitative agreement is attained, which demonstrates the suitability of the method detailed herein for predicting the loading from non-centrally aligned explosive charges.

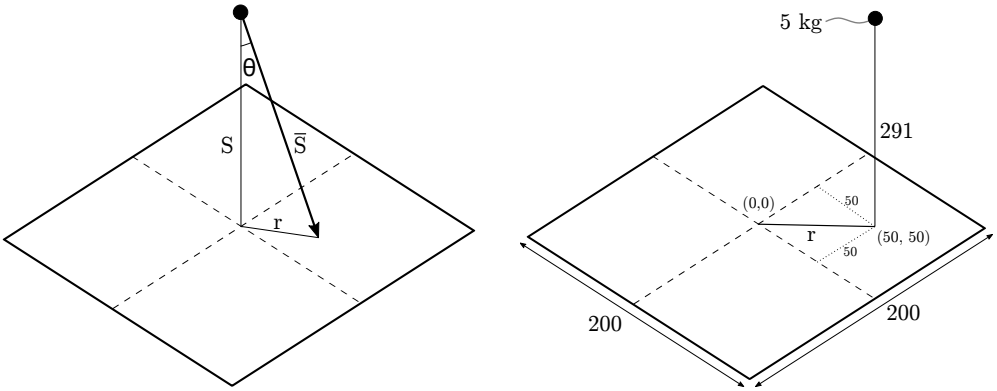

(a) (b)

Figure 4.13: Experimental schematics for a spherical charge detonating above a square target: (a) general case of a centrally located charge; (b) example off-centre case of 5 kg PE4 at 291 mm normal stand-off distance, offset by 50 mm in both in-plane directions from the centre of a 200 mm square plate. Note: dimensions in mm

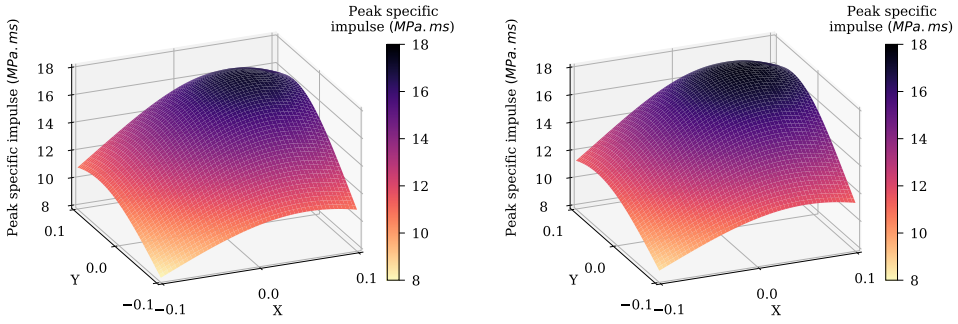

(a) (b)

Figure 4.14: Peak specific impulse distributions for 5 kg spherical PE4 charge at 291 mm normal stand-off distance, offset by 50 mm in both in-plane directions ('X' and 'Y') from the centre of a 200 mm square plate: (a) CFD model (b) predictive model using equation 4.4

4.4.3 Comparison against unseen experimental data

Finally, the experimental data presented in Table 2 of Geretto et al. (2015) is used to assess the accuracy of the method in calculating total impulse. Here spheres of PE4 explosive with masses ranging from 10–70 g were detonated at a constant stand-off distance of 100 mm (to charge centre, $0.243 \leq Z \leq 0.464$ m/kg$^{1/3}$) from a target plate with exposed area of 200 × 200 mm (70.5° angle of incidence to the corner of the plate). 34 experiments were performed altogether in eight different configurations. Total impulse was measured using a horizontal ballistic pendulum. Figure 4.15 shows the experimental results (individual test data and the mean at each scaled distance) and predictions from numerical integration of equation 4.5 over a 200×200 mm target area on a grid of 0.1×0.1 mm elements. The predictions can be seen to closely match the overall trend of the experimental results. In particular, the predictions seem to better match the upper bound of the experimental data, which again is to be expected on account of the inherent conservatism of the CFD analyses and predictive method.

A direct comparison is provided in Table 4.4. It can be seen that, with the exception of the furthest scaled distance (16% difference), the predictions are all within 9% of the maximum recorded experimental area-integrated impulse. In all instances, the predictions are greater than the experimental values, confirming the slight conservatism of the method presented in this article. It is tentatively suggested that, in cases where a conservative model is not required, a factor of 0.9 may be applied to predicted specific impulses. Applying this factor here would shift the predictions from closely matching the maximum experimental values to closely matching the experimental mean.

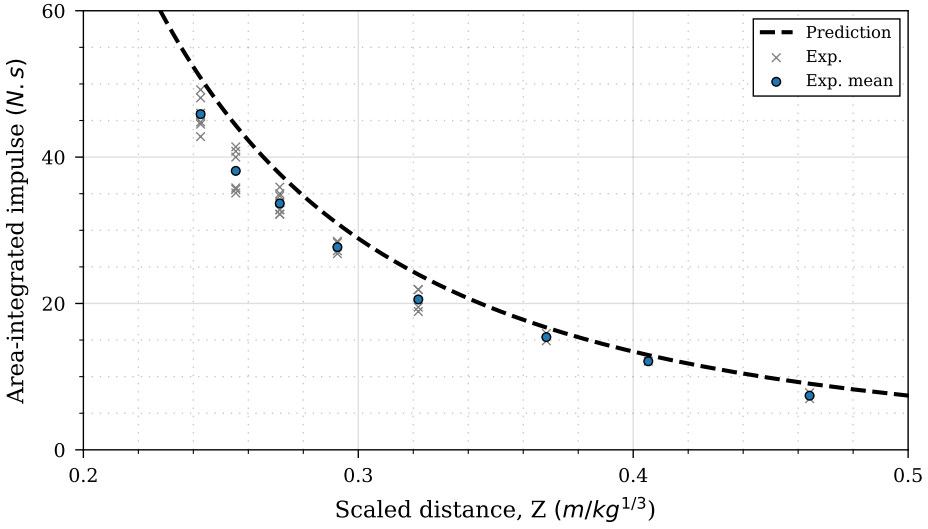

Figure 4.15: Comparison of predicted area-integrated impulse and experimental data from Geretto et al. (2015)

Table 4.4: Comparison of experimental area-integrated impulses (Geretto et al. 2015) and predictions

Charge mass (g)	Stand-off (mm)	Scaled distance ($m/kg^{1/3}$)	Area-integrated impulse			
			Experiment (N.s)		Prediction (N.s)	% diff. (from max.)
			Mean	Max.		
70	100	0.243	45.87	49.20	50.80	3
60	100	0.255	38.12	41.40	44.30	7
50	100	0.271	33.68	35.90	37.68	5
40	100	0.292	27.70	28.50	30.91	8
30	100	0.322	20.55	21.90	23.95	9
20	100	0.368	15.40	15.90	16.71	5
15	100	0.405	12.10	12.10	12.94	7
10	100	0.464	7.40	7.80	9.03	16

4.5 Summary

Peak specific impulse is the primary parameter that governs structural deformation under short-duration loading (Rigby et al. 2019b). Predicting the distribution of specific impulse acting on a reflecting surface following detonation of a high explosive charge remains a key challenge to engineers involved in transport security, infrastructure assessment, and defence. This is a particular requirement in regions extremely close to the charge, where a historic lack of experimental data has precluded both the development of accurate semi-empirical predictive approaches and validation of high fidelity, physics-based numerical schemes. In this chapter, a new data-driven predictive model is developed to address a key issue facing the blast protection engineering community: the lack of an accurate and fast-running engineering tool for blast load prediction in the extreme near-field.

Newly-available experimental near-field specific impulse data is used to validate *Apollo*, an explicit CFD software specialising in problems involving blast wave propagation. A CFD dataset is populated with *Apollo* analyses, and the results are used to derive the novel data-driven surrogate model presented in this chapter. The approach is formed of two aspects: firstly, a Gaussian function is used to describe the relationship between normalised specific impulse and angle of incidence (shown in this chapter to form a consistent grouping for the range of scaled distances considered). Secondly, a power law is fit to values of peak normally reflected specific impulse, and Hopkinson-Cranz scaling (Hopkinson 1915, Cranz 1926) is used to express the predictions in scaled form. Separate equations are derived for the full dataset ($0.11 \leq Z \leq 0.55$ m/kg$^{1/3}$) and a reduced dataset ($0.11 \leq Z \leq 0.21$ m/kg$^{1/3}$) offering the analyst improved accuracy for extreme close-in detonations.

Finally, accuracy of the approach is rigorously assessed against: the full CFD dataset; two new CFD analyses using 5 kg and 250 kg explosives (which also demonstrates the applicability of the implemented scaling laws and the generality of the approach for off-centre charge placement); and experimentally-recorded impulse acting on a 200 × 200 mm plate from a range of explosives (10–70 g) placed at 100 mm stand-off (Geretto et al. 2015). The spatial distribution of impulsive loading and the total impulsive load are shown to be in good agreement throughout, with the predictive method demonstrating slight conservatism and variations ranging between 1–16% of the experimental/test examples, with a typical variation of <5%.

The simple equations presented herein (three coefficients and three input variables: scaled distance, angle of incidence, and charge mass) enable the fast and accurate prediction of near-field blast loads in situations where previous methods are known to be unsuitable or inaccurate. This chapter has demonstrated the applicability of data-driven predictive approaches for blast load prediction. The methodology outlined in this chapter has the potential to be supplemented with more advanced, machine learning-assisted approaches to model a more complex suite of scenarios and input parameters, such as explosive type (TNT equivalence), charge shape effects, and blast wave clearing.

Chapter 5

Physics-based regularisation for near-field spherical charges

5.1 Introduction

As established in Chapter 2 obtaining data in the field of blast protection engineering is considerably expensive in both time and cost. Yet accurate appraisal of the viability of structures following close-in detonation of a high explosive is only possible with the availability of accurate models detailing both the distribution and magnitude of the imparted load. To provide accurate appraisals of the viability of protective structures following close in detonations, it is not only necessary to have accurate fast-running models, but also to adopt a probabilistic rather than deterministic approach when modelling blast scenarios. Such approaches (Netherton & Stewart 2016, Alterman et al. 2019, Campidelli et al. 2015) allow for the uncertainties related to the blast event to be modelled which accurately support a quantitative calculation of risk relating to blast

events.

In order to adopt probabilistic approaches, it is necessary to provide loading information for a large range of potential loading scenarios in a reasonable time-frame. In Chapter 2 it is shown that existing empirical and semi-empirical models are fast and can be accurate, but lack flexibility to consider anything other than scenarios with few variables and simple geometries, and typically only focus on far-field scenarios. Conversely, artificial neural networks (ANNs or NNs) can accommodate more variables when acting as a vector mapping model and provide the additional flexibility required to handle highly non-linear behaviour, as expected in extreme near-field blast events. They have been shown to accurately predict explosive loading in confined internal environments (Dennis et al. 2021), on a building behind a blast wall (Remennikov & Rose 2007, Flood et al. 2009), or along simple city streets (Remennikov & Mendis 2006).

However, there exist limitations in these 'black-box' approaches (such as NNs) in that they are independent to the underlying scientific principles that drive real-world physical phenomena and therefore do not always provide interpretable characteristics. Due to this nature, they often show poor generalisability when extrapolating beyond the realms of the training data, or indeed interpolating between gaps within it as they possess no knowledge of the underlying physics – an issue raised in Section 2.3.1. This issue can be exacerbated in supervised learning problems when the sample training sets may be small and, or, not evenly distributed as the model is more likely to learn spurious, non-physical relationships.

A fundamental challenge in machine learning (ML) is for the trained algorithm to perform well on *new, unseen data*, referred to as 'test' data.

The ability for a model to accurately predict previously unobserved inputs is called generalisation. The main factors that determine how well a ML algorithm will perform are its ability to: a) make the training error small; b) make the gap between training and test error small (Goodfellow et al. 2016, chap. 5). These two factors also refer to two fundamental challenges in ML, preventing overfitting and underfitting. The former occurs when the gap between the training error and test error is too large, whilst the latter occurs when the model does not achieve a low enough error on the training dataset. The likelihood of a model overfitting or underfitting can be controlled to some extent by changing its capacity (or also known as complexity). A model's capacity refers to its ability to fit a wide variety of functions. A useful analogy to demonstrate model capacity is to imagine a linear, quadratic and fifth order predictive model attempting to fit to data that is sampled from a quadratic function. In this instance, the linear model's capacity will be too low and will struggle to capture the curve in the underlying data (underfitting). The fifth order model will be able to represent the correct function and many more in between (overfitting), whilst the quadratic predictor will have the optimum capacity. The relationship between model capacity and error is summarised in Figure 5.1, where the underfitting and overfitting zones can be seen either side of the 'optimal capacity' region.

Enhancing the model's capacity can be achieved in several different ways. One way is by restricting the set of functions an algorithm can choose as its solution through selection of its hypothesis space. In the context of the previous example this would be allowing higher order polynomials rather than just linear functions for the quadratic and fifth order predictors. Once a family of representative functions have been chosen, there is the issue of finding the optimum function within this, termed

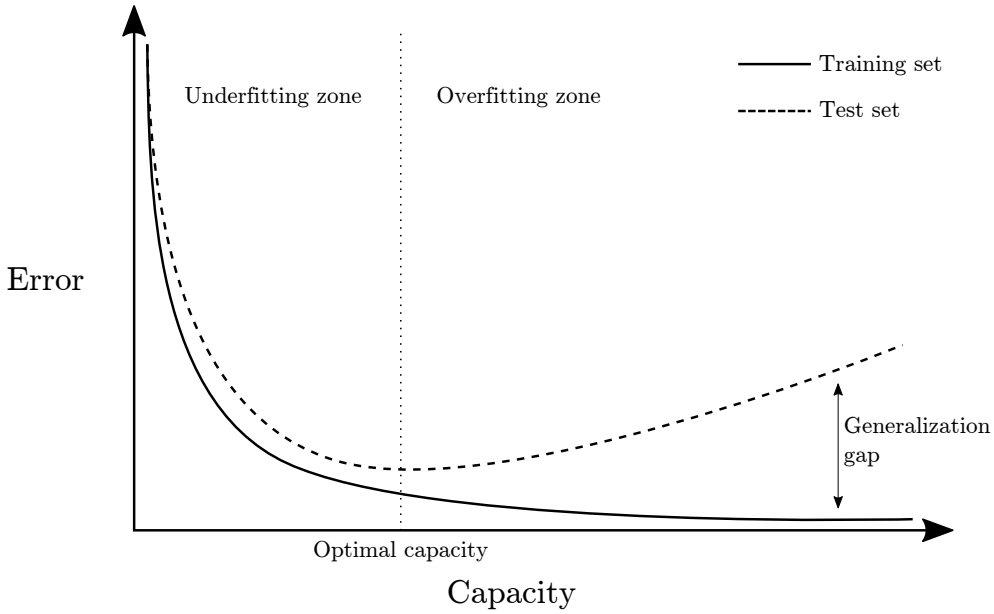

Figure 5.1: Typical relationship between model capacity and error. Training error and test error show different behaviour. Once a model reaches the 'optimal' capacity, its performance on unseen test data decreases, whereas its performance on the training dataset increases, this is the *overfitting zone* and this gap in performance is known as the *generalization gap*. Conversely the *underfitting zone* is where the model does not have sufficient capacity to accurately model the training set.

the representational capacity of the model. In many cases, finding the optimum function proves to be a difficult optimisation problem.

Recent advancements in improving generalisability suggest that, amongst competing hypotheses that explain the data equally well, the 'simplest' should be chosen (Goodfellow et al. 2016, Vapnik & Chervonenkis 1971, Vapnik 2006, Blumer et al. 1989, Vapnik 2013). From the 'no free lunch' theorem (Wolpert 1996), it is impossible to discover the overall optimum ML algorithm but that instead, ML algorithms should be designed to perform well on a specific task by building preferences into the learning algorithm (Goodfellow et al. 2016, chap. 5). Previously-mentioned methods of modifying the model's representational capacity have focused on controlling the hypothesis space of possible solutions, however there are other ways of controlling a preference for certain solutions, known generally as regularisation. Regularisation is "*any modification we make to a learning algorithm that is intended to reduce its generalisation error but not its training error*" (Goodfellow et al. 2016, chap. 5), in other words, the chosen learning algorithm can be forced to express a preference for certain solutions in the hypothesis space. Alternative approaches to regularisation exist that make use of data transformation procedures, particularly useful if the data is well understood. These approaches transform the data to a space where it is well-defined and then train on this new transformed distribution, and have been shown to be a promising avenue of research (Notley et al. 2021).

Integrating deep learning approaches, which are data intensive, and scientific theory is considered to be a crucial step to improve model predictive performance whilst respecting natural laws (Reichstein et al. 2019). A promising avenue of research to achieve this involves guiding the learning of a ML model though introducing a physical consistency

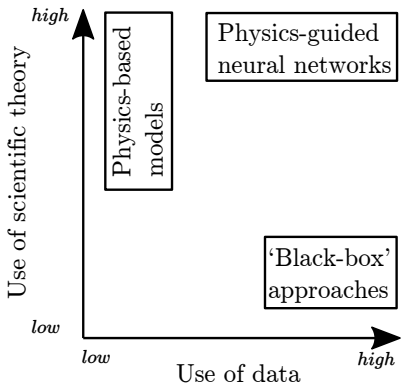

Figure 5.2: Schematic representation of physics-guided neural networks in the context of other knowledge discovery approaches. Adapted from Karpatne et al. (2017).

penalty as a regularisation procedure (Karpatne et al. 2017, Daw et al. 2020, Jia et al. 2019, Stewart & Ermon 2017). This framework aims to combine the use of scientific theory and leveraging the power of data-intensive approaches to develop a hybrid approach (as shown in Figure 5.2) that improves generalisability by guiding the learning process to bias physically valid solutions. Physics-based approaches such as these have been used effectively in a variety of domains such as the geo-sciences community (Muralidhar et al. 2018, Karpatne et al. 2017), power-flow research (Hu et al. 2020) and seismic response modelling (Zhang et al. 2020).

This chapter presents a novel use case of this framework in modelling peak specific impulse distributions of near-field blast loading distributions with a comparison against a standard neural network. Firstly, the dataset generation is discussed, then an overview into the development of a neural network architecture with a sensitivity study of the network parameters followed by the implementation of a physics-based regularisation. Then the generalisation ability of two neural network architectures (PGNN vs NN) is evaluated by stress-testing the models through withholding specific fractions of the variable distribution, and is

compared to some other "black-box" machine learning methods as an additional means of comparison. A discussion and concluding remarks are then provided where results show that of the seven stress tests performed, four of the five statistically significant performance premiums were shown by the PGNN. These results indicate that physics-based regularisation procedures can be implemented to provide statistically significant performance premiums when the modellers have specialist knowledge of the problem domain.

5.2 Surrogate modelling and model order reduction

5.2.1 Data-driven modelling

Surrogate modelling and model order reduction are strategies that aim to reduce the computational complexity of mathematical models in numerical simulations. Generally model order reduction in fluid dynamics applications (and subsequently in blast engineering applications) is particularly challenging due to the presence of highly non-linear behaviour (Lassila et al. 2014), though have been used when considering a material or structural response to a dynamic blast load (Jiang et al. 2020, Iliopoulos et al. 2017). An example of model order reduction is demonstrated in Chapter 4 which made use of normalising peak specific impulse with respect to the given angle of incidence.

A surrogate model is an engineering method when an output of interest cannot be directly measured with relative ease. Particularly in engineering, most practical problems require extensive experimental or numerical data to evaluate design decisions. For instance, in aerospace

engineering the design of an optimal aerofoil shape for an aircraft wing, requires knowledge of airflow parameters for a variety of variables (length, material, angle, etc.), typically determined through simulation (Mack et al. 2007). However, due to the computational expense of running an accurate numerical simulation, obtaining complete information for the design space would carry a considerable time cost and be unsuitable in practice, so a surrogate model is constructed that can evaluate points in the design space that may not necessarily have been explicitly sampled. Surrogate models are constructed using a data-driven, bottom-up approach whereby data sampled from particular simulations or experiments is used to inform a surrogate model (also known as a metamodel or response surface model). The surrogate model then essentially becomes a tool to interpolate in the input–output space of the different design variables.

The data-driven, bottom-up nature of building a surrogate model therefore requires large amounts of data. However, a particular issue in blast engineering is that experimental data can be prohibitively expensive or difficult to obtain due to the loading magnitudes and sub-millisecond durations of blast events, as mentioned previously in Chaptere 2. Computational fluid dynamics software provides engineers with the ability to generate the data required to inform an input–output space by solving the physical equations governing fluid flow through discretisation in space and time, and can be validated against known experimental data. The CFD models provide a response for a given set of input parameters, and these are used to fill a parameter space, requiring next the choice of surrogate model to interpolate within this parameter space. A detailed discussion on different types of surrogate model is given in Jin (2005). One such surrogate model is an artificial neural network (ANN)

and has been used extensively for this purpose in a variety of disciplines (Papadopoulos et al. 2018, Kim et al. 2015, Ahmadi 2015, White et al. 2019) and in specific blast engineering applications (Flood et al. 2009, Remennikov & Mendis 2006, Remennikov & Rose 2007, Dennis et al. 2021).

5.2.2 ANNs as a surrogate model

Artificial neural networks (ANNs) are collections of connected nodes (also called neurons), inspired by a simplification of neurons in a brain (example shown in Figure 5.3). They have been used extensively in a wide variety of disciplines due to their ability to act as universal approximators for complex functions. A linear model, which would map inputs to outputs through matrix multiplication, can only represent linear functions, or summations thereof. The universal approximation theorem (Hornik et al. 1989, Cybenko 1989) states that any feed-forward network with a linear output layer and at least one hidden layer with any 'squashing' activation function can approximate any Borel measurable function from one finite-dimensional space to another with any desired non-zero amount of error, provided the network is given enough hidden units (Goodfellow et al. 2016). Similarly, an ANN can approximate any function mapping from any finite dimensional discrete space to another. This ability to approximate non-linear functions makes ANNs a suitable choice for a wide variety of procedures, particularly suitable for recreating the non-linear physical phenomena seen in near-field explosive events.

The hidden layers in an ANN consist of interconnected neurons. Information is conveyed through these neurons with a series of weights and biases, transmitted with non-linear activation functions. To train an ANN, the optimal values for weights and biases must be found, which

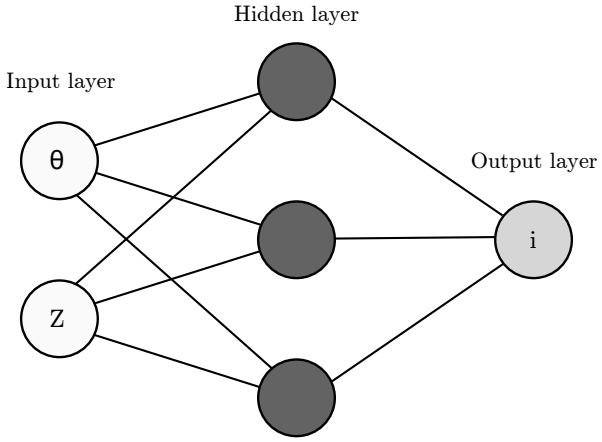

Figure 5.3: Example neural network architecture with two input nodes, 3 nodes in a hidden layer and one output node.

typically involves two stages: 'feedforward propagation' and 'error back-propagation'. Considering Equation 5.1:

$$y_j = f\left(\sum_{i=1}^{m} W_{ji}x_i + b_j\right) \tag{5.1}$$

where y_j is the output for the jth neuron in the current hidden layer, the output of each neuron is obtained by inputting into activation function, f, the product of the weight connecting the two neurons, W_{ji} and the input, x_i, then adding a bias term, b_j. This summation is continued for all hidden units until reaching the output layer, finalising the feedforward propagation. The prediction from the network can be compared to a target output and an error can be calculated (such as mean squared error). Once the error has been calculated, the learning algorithm can adjust the weights and biases after a set number of trials (the batch size) to minimise the error between model predictions and target values. This is termed 'error backpropagation' and is repeated for a given number of complete passes through a dataset (the epoch number).

As shown previously in Figure 5.1, as the capacity of an ML model

increases, as does the chance of overfitting. Specific to an ANN, there are several variables that can be changed in order to find the 'optimal capacity'. First, considering the structure of the network itself, the capacity of the model can be altered by increasing the number of layers, the nature of these layers, number of neurons in each layer, and the connectivity of these neurons. If every neuron in each layer is connected to every other neuron in adjacent layers, this is known as 'fully-connected' or 'dense' and is the structure of traditional multi-layer perceptrons. An alternative class of neural networks are convolutional neural networks, which also consist of input layers, hidden layers and output layers, but the hidden layers perform convolutions on input data, and have been shown to be useful in many image analysis tasks (Albawi et al. 2017). Finding the optimal structural configuration of an ANN is a challenging task and has been shown to be case-dependent and largely data-driven in nature (D'souza et al. 2020).

A central problem in ML is improving the ability of a model to generalise to unseen data. One such approach that focuses on reducing the test error (typically at the cost of an increased training error) is regularisation, as mentioned previously. There are several approaches to regularisation, some target the parameters of the model, by restricting or constraining their values in some way or including additional penalties, thereby encoding some prior knowledge into the model. One common example of a regularisation procedure is 'weight decay', where a sum is minimised that involves an additional criterion that adds a preference for model parameter weights to have a smaller L^2 norm, as shown in Equation 5.2 (Goodfellow et al. 2016):

$$J(\mathbf{w}) = MSE_{train} + \lambda \mathbf{w}^T \mathbf{w}, \tag{5.2}$$

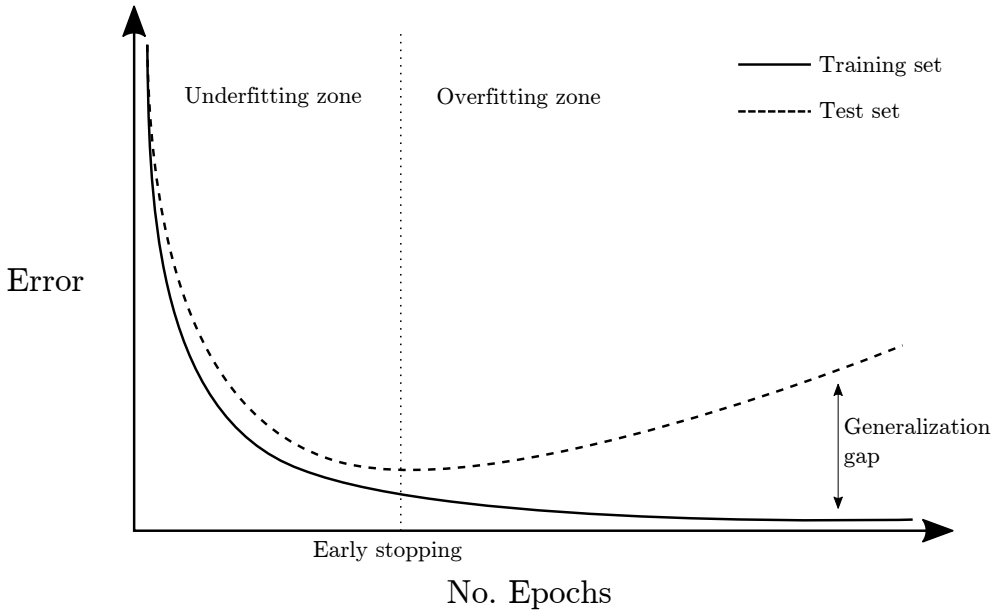

Figure 5.4: Example training history of a neural network. As the model continually trains, the model performs better on the training set but worse on the test set, increasing the generalisation gap. By early stopping the generalisability of the model is improved.

where $J(\mathbf{w})$ is the cost function to minimise, MSE_{train} is the mean squared error of training data, $\mathbf{w}$ is the parameter weights and λ is a hyper-parameter. This is an obtrusive form of regularisation as the objective function of the minimisation procedure is directly modified. Because of this, the choice of the hyper-parameter λ is crucial: too small and the model will overfit to the data, whilst too large and it will underfit.

Another common regularisation procedure is known as 'early-stopping', and is simple to implement and highly effective. Consider the training history shown in Figure 5.4, after a certain point the model would begin to overfit. This approach can be considered an algorithm for selecting the optimum model parameters, as throughout training the model weights and biases are being changed. By implementing early stopping the model parameters that show the best generalisability are returned.

Further regularisation techniques rely on combining several different

trained models, in the hope that their average, aggregate performance shows better generalisability, known as bagging (Breiman 1996). This general strategy is known as model averaging and similar methods are known as ensemble methods. Neural networks often benefit from model averaging of some kind due to the differences in parameter initialisation or other hyper-parameters (Goodfellow et al. 2016). On a similar theme, another particularly useful form of regularising ANNs is known as 'dropout' (Srivastava et al. 2014). When a fully-connected layer has a large number of neurons, there is likely to be some co-adaption. Co-adaption refers to the scenario where multiple neurons extract the same information or features from a set of input data, the detriment of this is an inefficient use of computational resources and an increased risk of overfitting by adding more significance to certain model features. In dropout, the ensemble of all the model subnetworks that can be formed by removing units from a base network are trained. For an extensive overview of further regularisation procedures, see Goodfellow et al. (2016).

5.2.3 Alternative data-driven surrogate models

As an additional means of comparison for the development of a surrogate model, two alternative data-driven models were studied. Support vector regressors, a specific implementation of a support vector machine (Boser et al. 1992, Guyon et al. 1993, Cortes & Vapnik 1995, Schölkopf et al. 1995, 1996, Vapnik et al. 1997), map the input data's original space into a higher dimensional spaces through a kernel function. In this higher dimensional space, a hyperplane is fitted that can be used for classification and regression tasks. In ϵ-SV regression the goal is to find a function $f(x)$ that has at most ε deviation from the targets y_i for all the training data, and is simultaneously as flat as possible.

First considering the case for linear functions, f of the form

$$f(x) = \langle w, x \rangle + b \text{ with } w \in X, b \in \mathbb{R} \tag{5.3}$$

where $\langle ., . \rangle$ is the dot product in X and flatness is a small w.

Given training vectors $x_i \in \mathbb{R}^p$, $i = 1, ..., n$, and a vector $y_i \in \mathbb{R}^n$, support vector regression (ϵ-SVR) solves the following problem, from Vapnik (2013):

$$\begin{aligned} \text{minimise} \quad & \frac{1}{2}||w||^2 + C\sum_{i=1}^{l}(\xi_i + \xi_i^*) \\ \text{subject to} \quad & \begin{cases} y_i - \langle w, x_i \rangle - b & \leq \varepsilon + \xi_i \\ \langle w, x_i \rangle + b - y_i & \leq \varepsilon + \xi_i^* \\ \xi_i, \xi_i^* & \geq 0 \end{cases} \end{aligned} \tag{5.4}$$

where the constant, $C > 0$ determines the trade-off between the flatness of f and the amount up to which deviations larger than ϵ are tolerated Smola & Schölkopf (2004). This leads to a ε-insensitive loss function $|\xi|_\varepsilon$ as:

$$|\xi|_\varepsilon := \begin{cases} 0 & \text{if } |\xi| \leq \varepsilon \\ |\xi| - \varepsilon & \text{otherwise.} \end{cases} \tag{5.5}$$

Figure 5.5 represents this graphically. Only points outside of an allowable error contribute to the 'cost'. The dual formulation of this optimisation problem can often be solved more easily, and is the key for extending support vector machines to nonlinear functions.

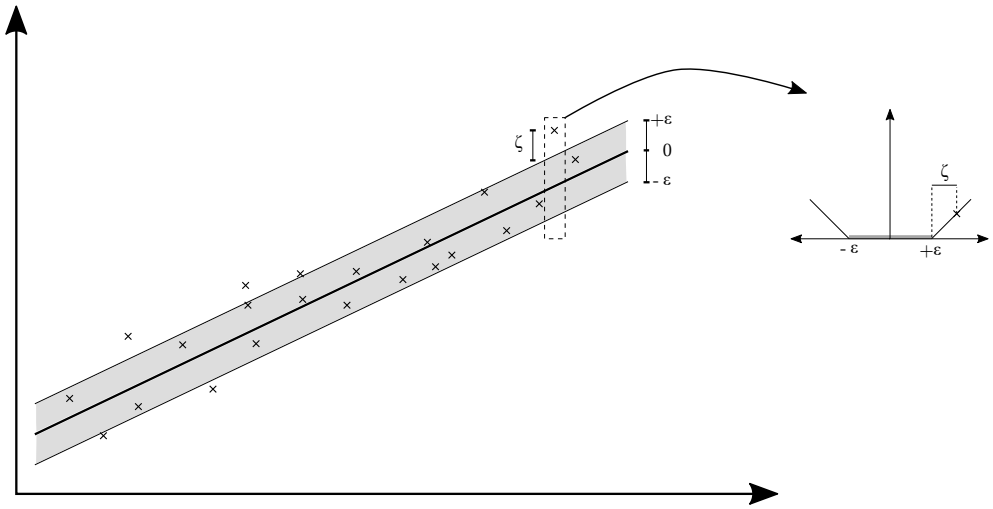

Figure 5.5: An example of the soft margin loss setting for a linear support vector machine. Recreated from Schölkopf et al. (2002).

The dual formulation is given as

$$\begin{aligned} &\text{minimise} \quad \frac{1}{2}(\alpha - \alpha^*)^T Q(\alpha - \alpha^*) + \varepsilon e^T(\alpha - \alpha^*) - y^T(\alpha - \alpha^*) \\ &\text{subject to} \quad \begin{cases} e^T(\alpha - \alpha^*) & = 0 \\ 0 \leq \alpha_i, \alpha_i^* \leq C, i & = 1, ..., n \end{cases} \end{aligned} \tag{5.6}$$

where e is the vector of all ones, $C > 0$ the upper bound, Q is an n by n positive semidefinite matrix, $Q_{ij} \equiv K(x_i, x_j) = \phi(x_i)^T \phi(x_j)$ is the kernel. Training vectors are mapped into a higher dimensional space through the function ϕ. The decision function is given by:

$$\sum_{i=1}^{n} (\alpha_i - \alpha_i^*) K(x_i, x) + \rho \tag{5.7}$$

with ρ as the intercept term. For a more detailed explanation of support vector machines, see Smola & Schölkopf (2004).

The second alternative model is a specific type of ensemble method, known as a gradient boosting regressor (GBR). It is built from an en-

semble of decision trees, where each subsequent tree is optimised to fit to the residuals of the previous tree (Friedman 1999). It considers additive models of the following form:

$$F(x) = \sum_{m=1}^{M} \gamma_m h_m(x) \tag{5.8}$$

with $h_m(x)$ as the basis functions, deemed *weak learners* in the context of boosting. Similar to other boosting algorithms, GBR builds the additive model in a forward stagewise approach:

$$F_m(x) = F_{m-1}(x) + \gamma_m h_m(x) \tag{5.9}$$

where at each stage the decision tree $h_m(x)$ is chosen to minimize the loss function L given the current model F_{m-1} and its fit $F_{m-1}(x_i)$

$$F_m(x) = F_{m-1}(x) + \underset{h}{\operatorname{argmin}} \sum_{i=1}^{n} L(y_i, F_{m-1}(x_i) + h(x)) \tag{5.10}$$

In gradient boosting, the minimisation problem is solved numerically by steepest descent. The steepest descent direction is the negative gradient of the loss function evaluated at the current model F_{m-1} which can be calculated for any differentiable loss function:

$$F_m(x) = F_{m-1}(x) - \gamma_m \sum_{i=1}^{n} \nabla_F L(y_i, F_{m-1}(x_i)) \tag{5.11}$$

where the step length γ_m is chosen using line search:

$$\gamma_m = \underset{\gamma}{\operatorname{argmin}} \sum_{i=1}^{n} \left(L(y_i, F_{m-1}(x_i) - \gamma \frac{\partial L(y_i, F_{m-1}(x_i))}{\partial F_{m-1}(x_i)} \right) \tag{5.12}$$

For a more detailed overview of gradient boosting regression, see Pe-

dregosa et al. (2011), Friedman (1999, 2001), Hastie et al. (2009).

5.2.4 Dataset overview

The dataset was generated from CFD simulations using *Apollo* (Fraunhofer EMI 2018) consisting of 100g spherical charges of PE4 located between 0.05m to 0.26m normal distance from the surface of the target, termed 'stand-off' distance. This is equivalent to a *scaled* distance range of 0.11-0.55$m/kg^{1/3}$. For each sample, 150 gauges are linearly spaced along the target surface at angles of incidence between 0 and 60 degrees, where angle of incidence is defined as the angle between the outward normal of the surface and the direct vector from the explosive charge to that point (Rigby et al. 2015a). Each gauge outputs pressure-time histories at that location, which are numerically integrated (with respect to time) in postprocessing to yield specific impulse-time histories.The maximum of each of these is taken to provide the distribution of peak specific impulse.

In summary, there are 18 CFD experiments with 150 values of peak specific impulse for each, resulting in a dataset of 2700 data points. An overview of this dataset is given in Figure 5.6. For more information on the validation of *Apollo* including mesh sensitivity studies, experimental validation and the construction of the dataset, see Chapters 3 and 4.

Each input for the NN is shown in Table 5.1, with an example entry: X represents the input vectors and Y represents the labelled output so that in its entirety the dataset is a (2700, 3) array.

The variables X_1 and X_2 are minmax scaled across the entire dataset with a feature range of $[0, 1]$. The output, Y, has a log-normal distribution and was scaled via a power transform; a family of parametric, monotonic transformations applied to provide normality and homoscedasticity to the training data for ease of model training, using the method

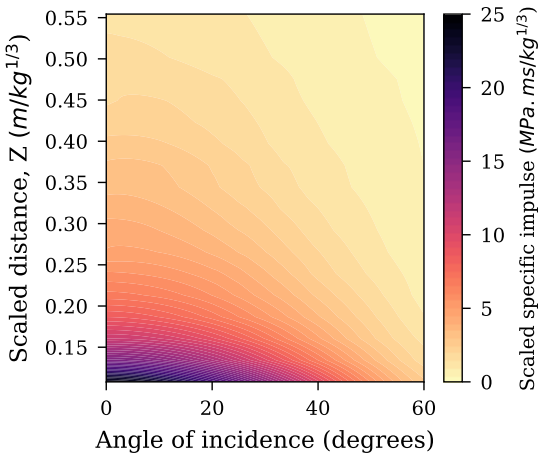

Figure 5.6: CFD dataset: filled contours of scaled peak specific impulse for full dataset, 0.11–0.55 m/kg$^{1/3}$.

Table 5.1: Example dataset information

X_1	X_2	Y
Scaled distance $(m/kg^{1/3})$	Angle of incidence	Peak specific impulse (MPa.ms)
0.1	15	45.87

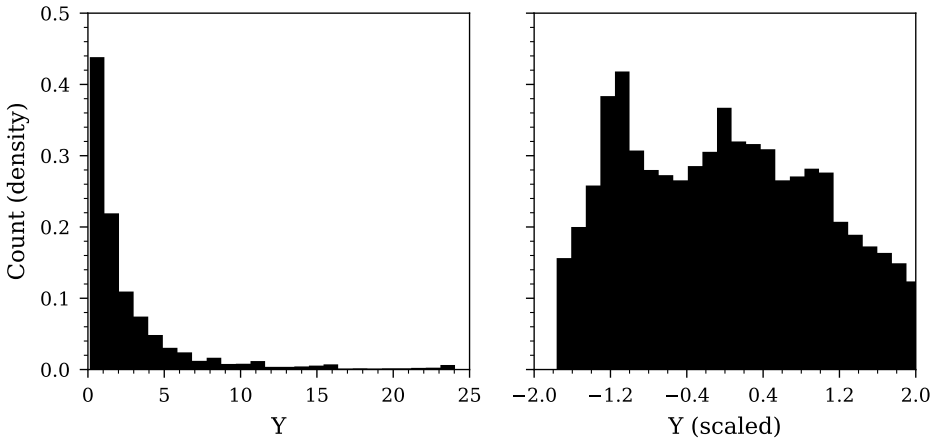

Figure 5.7: Unscaled Y dataset showing a log-normal distribution (left) and the resulting power transformation (right)

described in Yeo & Johnson (2000). The result of the power transform is given in Figure 5.7. The data transformations are completed after any data splitting procedures are undertaken to ensure there is no data leakage from the transformation.

5.2.5 Training and network architecture

To create and train the networks, the open-source Keras package with TensorFlow backend (Chollet et al. 2015) was used. The Adadelta (Zeiler 2012) gradient descent algorithm was chosen which has the added benefit of not requiring a default learning rate. 10000 epochs with early stopping set at 500 epochs was implemented to prevent over-fitting, with a batch size of 32. Activation functions for the hidden units were set as hyperbolic tangent, and layer weights initialised with the Glorot normal initialiser (Glorot & Bengio 2010). K-fold cross validation was implemented with 5 splits, following an initial data split of 25% data randomly removed, this data splitting schematic is summarised in Figure 5.8. The networks to be examined all had 1 hidden layer, and are fully connected.

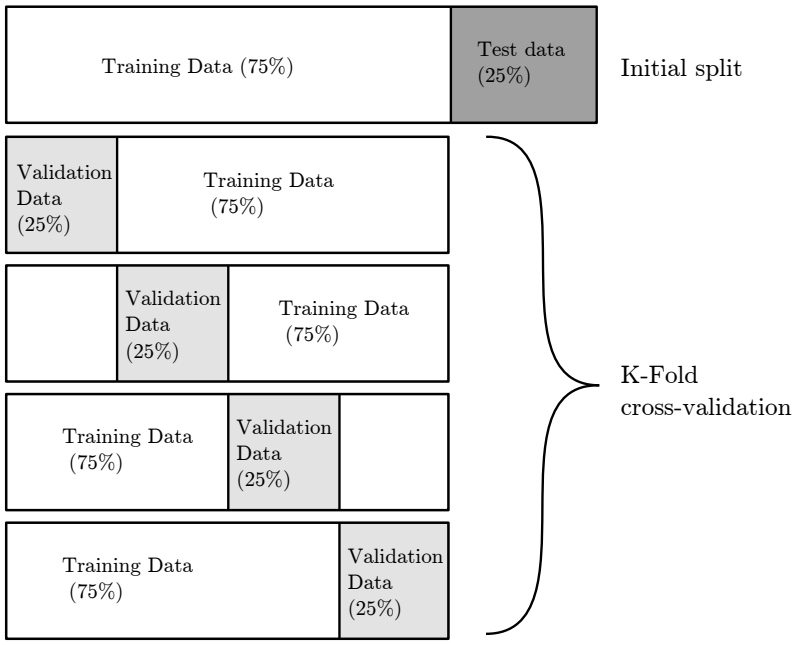

Figure 5.8: Example schematic of splitting the dataset prior to and during training and the subsequent k-fold cross-validation. The initial 25% of test data removed is never seen by the model during training.

Initially, a varying number of hidden units were examined, from 1 to 8. The results of these analyses are presented in Figure 5.9. Ideally, the structural optimisation of the network would be performed in parallel to the parametric optimisation via an overarching, many-objective process (Jin & Sendhoff 2009, Loghmanian et al. 2012), though this is considered to be beyond the scope of this thesis. Three separate sub-figures provide different metrics: mean squared error (Figure 5.9a), mean absolute error (Figure 5.9b) and coefficient of determination (Figure 5.9c). For all analyses, the metrics are assessed for the three separate data portions: test, train and validation as previously summarised in Figure 5.8. If large discrepancies occur between the different data types, this can be indicative of over-fitting issues, as the models may perform far better on 'seen' data rather than 'unseen' data. As shown, there is a general trend of error reducing as the number of neurons increases, though this plateaus after 4

neurons, suggesting that 4 neurons provide sufficient predictive capability. Furthermore, from 2 neurons and upwards, there are no significant discrepancies between the different data types, suggesting that the model is not over-fitting. For later analyses, a network with four hidden units was chosen to take forward.

With the network architecture chosen, it is useful to gain an insight into the model learning performance of this network to check further for over-fitting and that the model is learning as expected. Training histories from the cross-validation are shown in Figure 5.10 presenting the loss from validation and training data during training. As shown, the performance consistently improves as the number of epochs increases, and also improves for each fold of the cross-validation, suggesting that the network is learning effectively. The generalisation gap between the training and validation set is small and remains fairly constant throughout which is demonstrative of a good fit. On initial inspection there appears to be some stochastic behaviour in the later folds, however considering the magnitude of the losses here these are not considered to be a cause for concern and it can be seen the model shows a good fit.

5.2.6 A physics-based regularisation procedure

All regularisation methods are attempting to improve the generalisability of a ML model, that is how it performs on unseen data. Regularisation would not be required if the optimum structure, learning algorithm, and data transformation for a given NN were known ahead of time. Typical approaches for training models without physical constraints involve minimising empirical loss (mean square error, MSE) of model predictions ($\hat{Y}$) on the training data (Y) whilst maintaining low model complexity

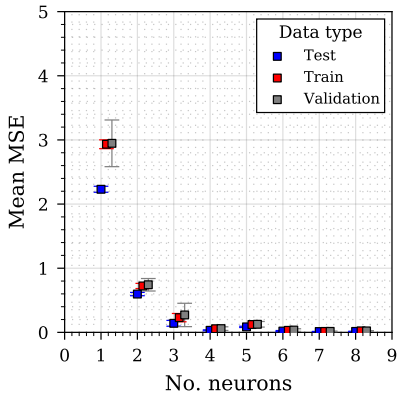

(a) Mean absolute error of model predictions for varying number of neurons in hidden layer.

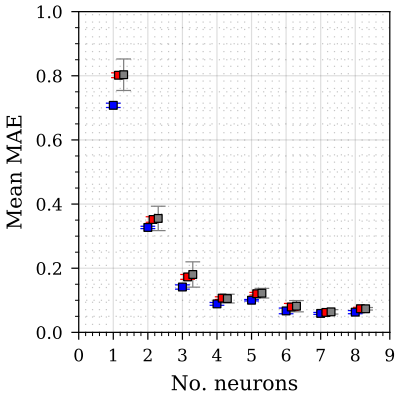

(b) Mean absolute error of model predictions for varying number of neurons in hidden layer.

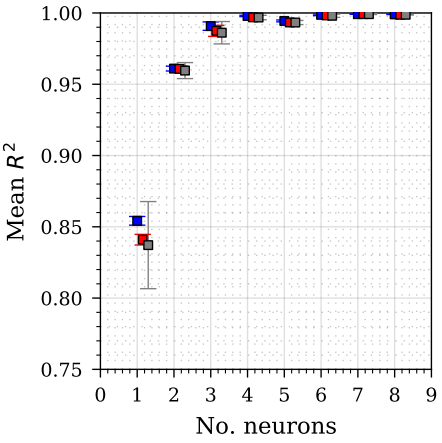

(c) Mean absolute error of model predictions for varying number of neurons in hidden layer.

Figure 5.9: Hyper-parameter configuration with various performance metrics shown on the unseen test data. Error bars are standard deviation.

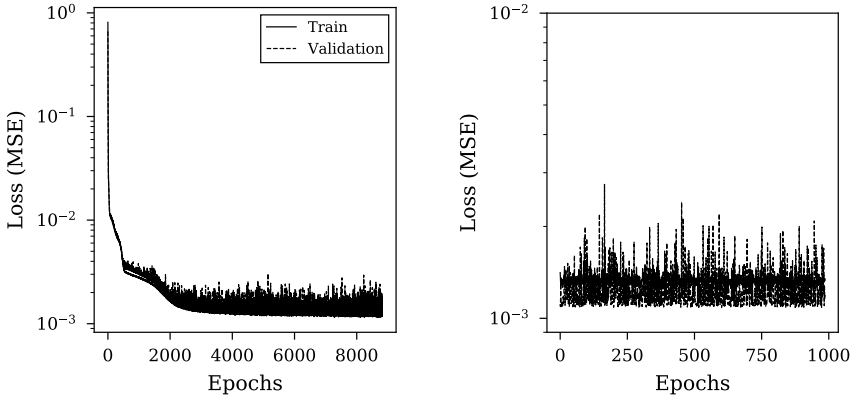

(a) First fold.

(b) Second fold.

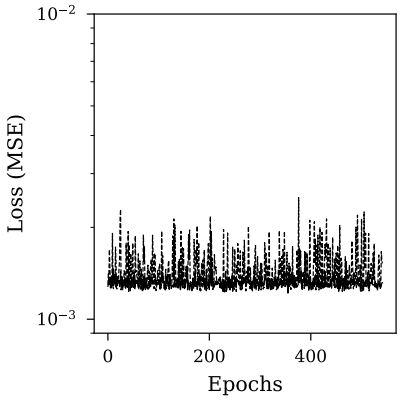

(c) Third fold.

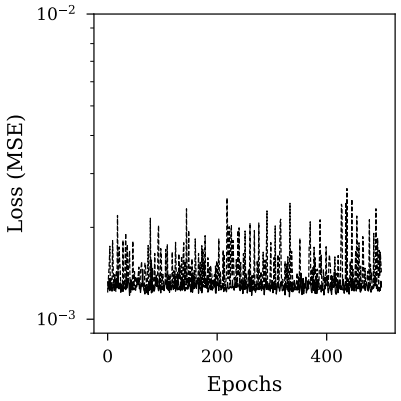

(d) Fourth fold.

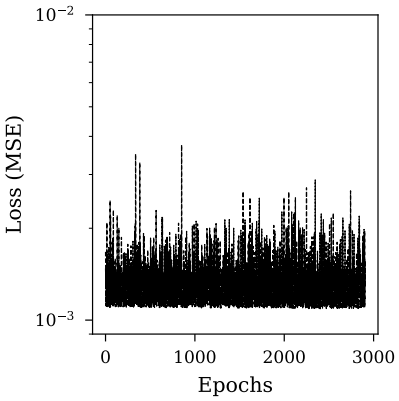

(e) Fifth fold.

Figure 5.10: Training loss histories for each of the 5-fold cross validation. Analyses represent the chosen network with four hidden units.

as follows in Equation 5.13:

$$\underset{f}{\operatorname{argmin}}\, Loss(\hat{Y}, Y), \tag{5.13}$$

where Y refers to target values. Through an additional physical inconsistency term the model will now punish predictions that are not valid with respect to the physical law that is incorporated. This is demonstrated in Equation 5.14 (Karpatne et al. 2017) with the addition of a physical loss function, $Loss.Phy(\hat{Y})$, and the corresponding hyper-parameter (λ_{Phy}):

$$\underset{f}{\operatorname{argmin}}\, Loss(\hat{Y}, Y) + \lambda_{Phy} Loss.Phy(\hat{Y}). \tag{5.14}$$

In this thesis a physics-based regularisation procedure based on each of the input features is implemented. Firstly, the prior knowledge that peak specific impulse decays monotonically with scaled distance for a given angle of incidence (from Chapter 4) is utilised, which allows the implementation of a monotonic loss constraint. This monotonic loss constraint is in place to only penalise values that do not satisfy the requirement, implemented through the use of a rectified linear unit activation function (ReLU). The ReLU function is represented mathematically as $f(x) = \max(0, x)$. Then the physical loss function returns the following, adapted from Karpatne et al. (2017):

$$\text{Loss.Phy}_1(\hat{Y}) = \frac{1}{n} \sum_{i=1}^{n-1} \text{ReLU}(f(X_{1,i+1}, X_2) - f(X_{1,i}, X_2)), \tag{5.15}$$

where X_1 refers to scaled distance and hence $X_{1,0}$ being the smallest value in the dataset with $f(X_{1,i}, X_2)$ representing the model prediction for a given angle of incidence, X_2. The subscript i indicates the index of the vector X_1. It is important to remember here that the minimum

X_1 value ($X_{1,0}$) will result in the maximum output (Y), as increasing the scaled distance between charge and target will increase X_1 and decrease Y. Therefore if the model predicts a greater output at a larger scaled distance ($X_{1,i+1}$) than it does for a smaller scaled distance ($X_{1,i}$), this result is non-physical and is punished by the monotonic loss constraint.

Similarly, prior knowledge that peak specific impulse decays monotonically with angle of incidence for a given scaled distance (where scaled distance is calculated from the hypotenuse distance between the charge centre and the measurement location) allows the implementation of a second monotonic loss constraint:

$$\text{Loss.Phy}_2(\hat{Y}) = \frac{1}{n}\sum_{i=1}^{n-1} \text{ReLU}(f(X_1, X_{2,j+1}) - f(X_1, X_{2,j})), \qquad (5.16)$$

where for a given X_1, the maximum output will be obtained at $X_{2,0}$, decaying monotonically as X_2 increases. Therefore if the model predicts a greater output at a larger angle of incidence ($X_{2,j+1}$) than it does for a smaller angle of incidence ($X_{2,j}$) this result is also deemed non-physical and is again punished by the model. Combining the two additional physical loss constraints (Equations 5.15 and 5.16) into Equation 5.13 the final model loss function is obtained:

$$\underset{f}{\text{argmin}}\, \text{Loss}(\hat{Y}, Y) + \lambda_{\text{Phy},1}\text{Loss.Phy}_1(\hat{Y}) + \lambda_{\text{Phy},2}\text{Loss.Phy}_2(\hat{Y}), \quad (5.17)$$

where $\lambda_{\text{Phy},1}\text{Loss.Phy}_1(\hat{Y})$ refers to the monotonic loss constraint with respect to scaled distance and $\lambda_{\text{Phy},2}\text{Loss.Phy}_2(\hat{Y})$ refers the monotonic loss constraint with respect to angle of incidence.

Implementation of the physical loss function requires values for the

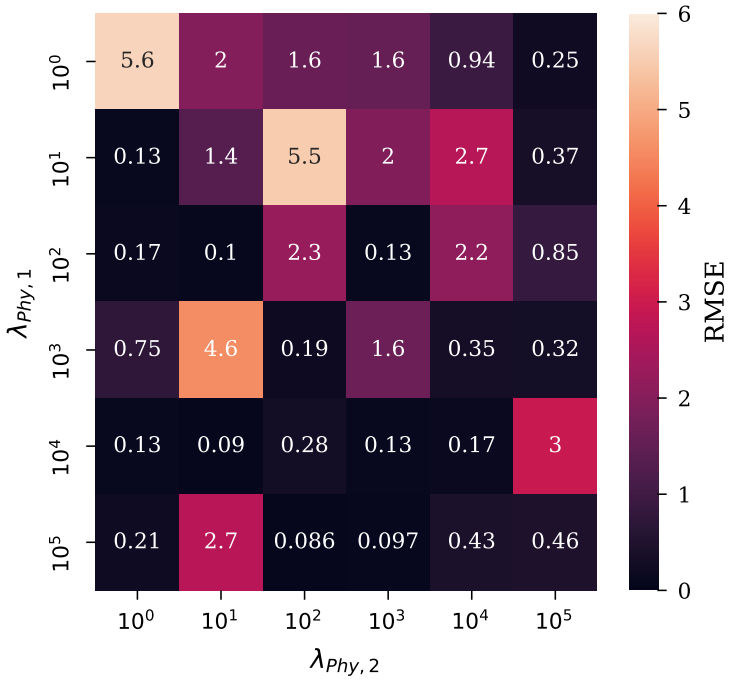

Figure 5.11: Grid search of $\lambda_{Phy,1}$ and $\lambda_{Phy,2}$

hyper-parameters to be found. A sensitivity study was conducted and with the performance metric set as root-mean-square-error (RMSE) on unseen, test data. The sensitivity study for these hyper-parameters ($\lambda_{\text{Phy},1}$ and $\lambda_{\text{Phy},2}$) assessed six logarithmically spaced values varying between 10^0 and 10^5. The results from this sensitivity study are given in Figure 5.11, with $\lambda_{\text{Phy},2}$ shown on the x-axis and $\lambda_{\text{Phy},1}$ on the y-axis, with the colour presenting the mean value for test RMSE (from the five folds of the cross-validation). The optimum values from this exercise were seen to be in a region of $\lambda_{\text{Phy},1} = 10^4$ and $\lambda_{\text{Phy},2} = 10^3$. However on closer inspection of the training history data, the values taken forward were $\lambda_{\text{Phy},1} = 10^4$ and $\lambda_{\text{Phy},2} = 10^4$.

5.3 Stress testing evaluation

5.3.1 Overview of stress-testing procedure

It is useful for predictive models in blast loading to be able to generalise both beyond the limits of the dataset and between points within the dataset. A model which is able to generalise in this way would be particularly beneficial to bypass physical constraints relating to availability of equipment, placement of gauges, and limitations of the recording equipment. For example, there may be practical limitations which prevent tests from being performed at certain scaled distances (limitations on robustness and survivability of recording equipment at smaller scaled distances, excessive signal to noise ratios at larger scaled distances), as well as potential issues relating to limited availability of diagnostics which may result in a sparse dataset. Understanding how accurately a model can extrapolate and interpolate has significant implications for the design of future experiments: such knowledge will provide experimentalists with a clear steer on typical areas of high sensitivity and therefore those with significant contributions to the overall accuracy of ML approaches, as well as areas where detailed measurements may not be required.

The stress-testing evaluation procedure therefore aims to investigate two themes of estimation ability: interpolation and extrapolation, as well as their dependencies on the availability of data to make accurate predictions. Herein, the dataset comprises three variables: X_1 (Z, scaled distance), X_2 (θ, angle of incidence) and Y (peak specific impulse). By restricting the availability of certain regions of each variable distribution, the ability of the model to generalise can be assessed for each given variable. Seven types of stress-test were performed, assessing either extrapolation or interpolation ability with 25% of the dataset removed. For

each stress-test, each network was trained to 100 epochs, and this procedure repeated 25 times. The information for each type of stress-test is summarised below:

1. Interpolation
    (a) Mean X_1 values removed
    (b) Mean X_2 values removed
    (c) Random data removed
2. Extrapolation
    (a) Maximum X_1 values removed
    (b) Minimum X_1 values removed
    (c) Maximum X_2 values removed
    (d) Minimum X_2 values removed

The architecture of the neural networks (NN and PGNN) follows that summarised previously, with the two additional hyper-parameters $\lambda_{Phy,1} = 10^4$ and $\lambda_{Phy,2} = 10^4$ included for the PGNN. Before the training process, various proportions of the data were removed, dependent on the mode of stress-testing, and the model performance was evaluated in each case by recording the root-mean-square-error (RMSE) on the removed, unseen test data. By also analysing the training history in each simulation, a physical inconsistency term can be quantified as the proportion of total epochs where physically invalid predictions are made, this will provide an insight into how 'physically valid' each network is. For clarity if physically inconsistent predictions are made in 4 of the 100 epochs, then the physical inconsistency measure will be 0.04, where a physically inconsistent prediction is defined as a positive result from Equation 5.15 or 5.16.

5.3.2 Results

Interpolation

The first theme of stress-testing aims to assess the estimation ability of the models to interpolate between the upper and lower limits of the dataset. Firstly, removing the region of the variable distribution around the mean of scaled distance (X_1), and angle of incidence (X_2), and finally removing data randomly across the distribution space. For each analysis, the distribution of the dataset with 25% of the variable removed is presented in Figure 5.12, with the root mean squared error (RMSE) on unseen test data and physical inconsistency values presented in Figure 5.13. Due to the cross-validation procedure and the simulation repeats, more than one value for RMSE, and physical inconsistency, is provided for each NN architecture, therefore the standard deviation of these values are plotted as error bars whilst the mean indicated as the scatter point for each architecture.

Extrapolation

The second theme of stress-testing aims to assess the estimation ability of the models to extrapolate beyond the upper and lower limits of the dataset. Firstly, removing the maximum values of scaled distance (X_1), and maximum values of angle of incidence (X_2). It is important to note here that the minimum values of peak specific impulse correspond to the maximum values in scaled distance and angle of incidence, as physically the distance between the charge and target is increased, hence the smaller impulse value. For each analysis, the distribution of the dataset with 25% of the variable removal is presented in Figure 5.14, with the root mean squared error (RMSE) on unseen test data and physical inconsistency values presented in Figure 5.15. As mentioned previously, due to the

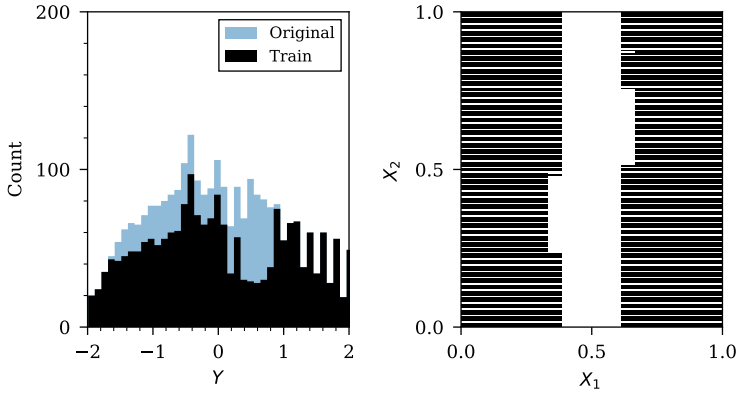

(a) Removing 25% of data around the mean X_1 value.

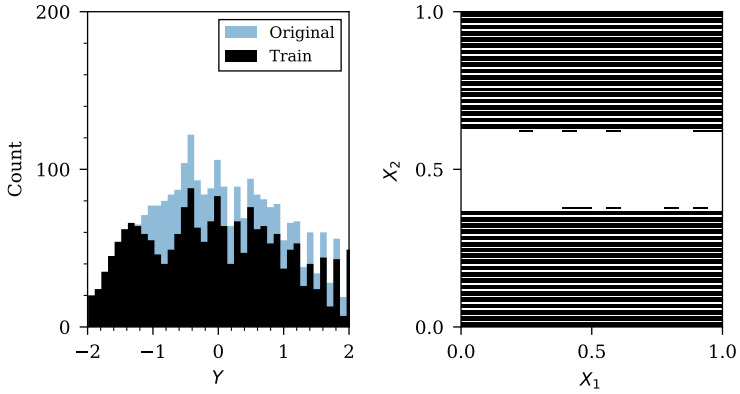

(b) Removing 25% of data around the mean X_2 value.

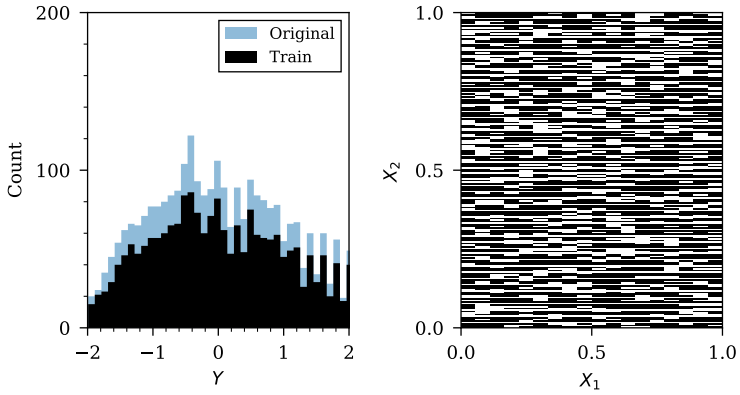

(c) Removing the 25% of X values randomly.

Figure 5.12: Interpolation stress tests, effects of data removal on the data distributions. Distribution of Y (peak specific impulse) before and after data removal (left); distribution of X_1 (scaled distance) and X_2 (angle of incidence) after data removal (right).

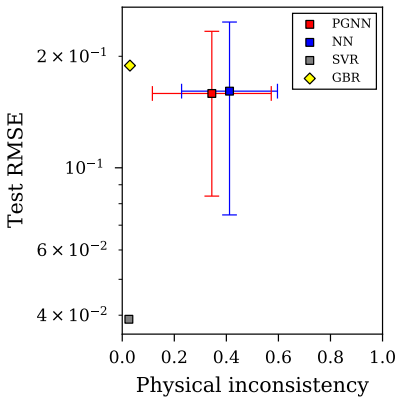

(a) Removing 25% of data around the mean of X_1 (scaled distance).

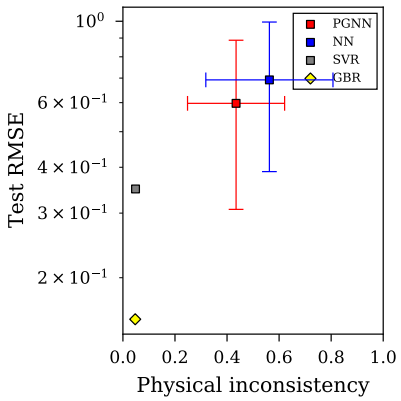

(b) Removing 25% of data around the mean of X_2 (angle of incidence) data.

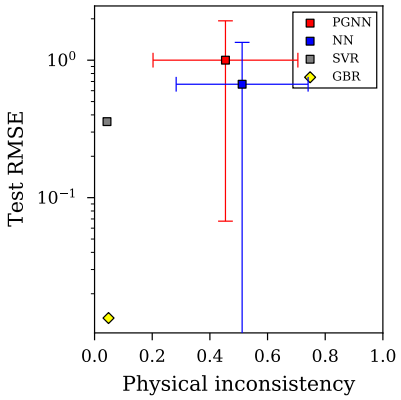

(c) Removing the 25% of X data randomly.

Figure 5.13: Results showing root mean square error (RMSE) of removed test data and physical inconsistency for various interpolation stress-test procedures.

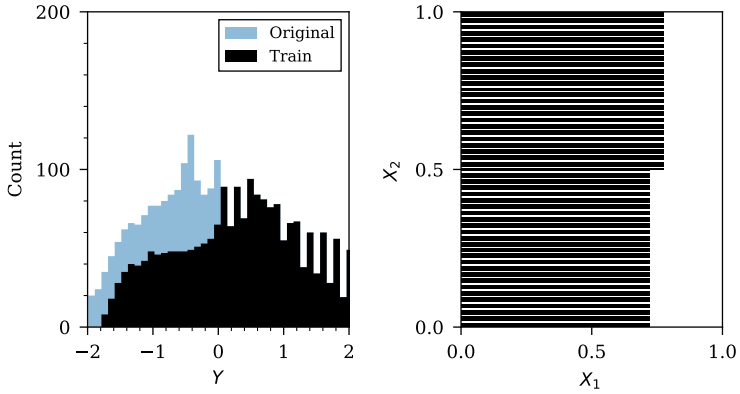

(a) Removing the maximum 25% of X_1 values.

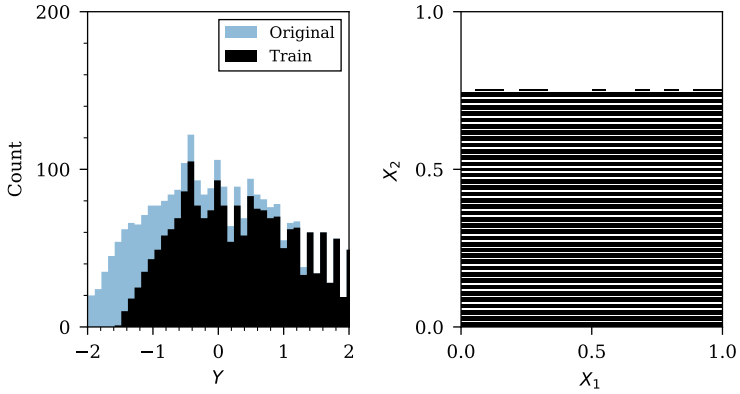

(b) Removing the maximum 25% of X_2 values.

cross-validation procedure and the simulation repeats, more than one value for RMSE, and physical inconsistency, is provided for each neural network architecture, therefore the standard deviation of these values are plotted as error bars whilst the mean indicated as the scatter point for each architecture.

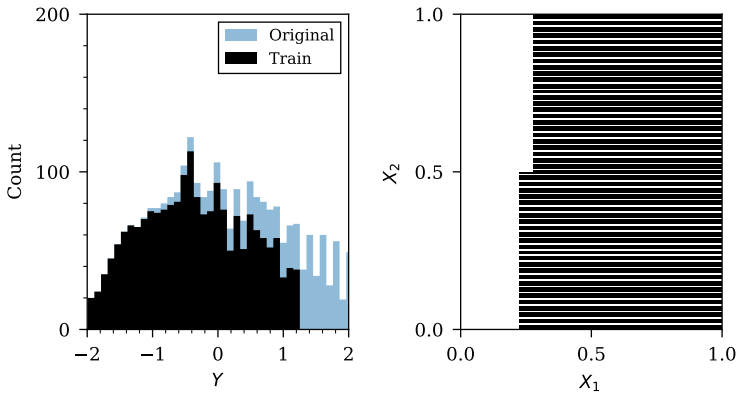

(c) Removing the minimum 25% of X_1 values.

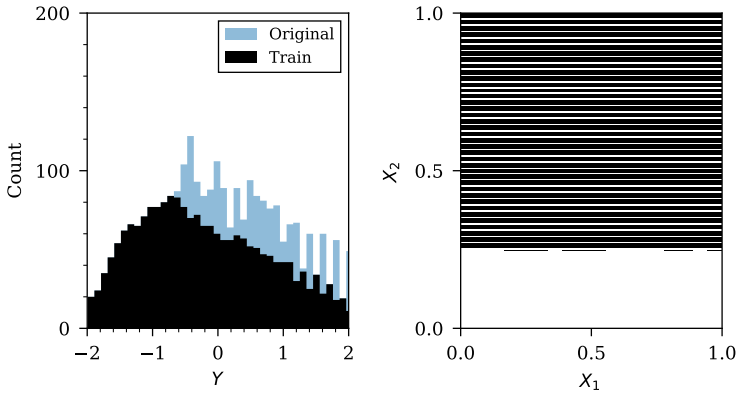

(d) Removing the minimum 25% of X_2 values.

Figure 5.14: Extrapolation stress tests, effects of data removal on the data distributions. Distribution of Y (peak specific impulse) before and after data removal (left); distribution of X_1 (scaled distance) and X_2 (angle of incidence) after data removal (right).

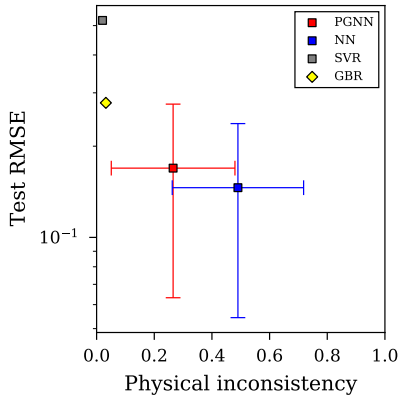

(a) Removing the maximum 25% of X_1 (scaled distance) data.

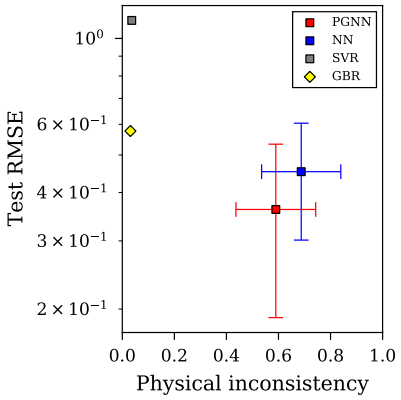

(b) Removing the maximum 25% of X_2 (scaled distance) data.

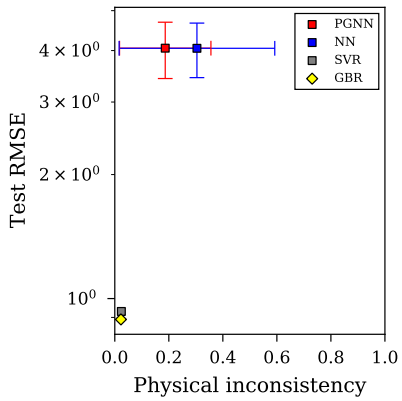

(c) Removing the minimum 25% of X_1 (scaled distance) data

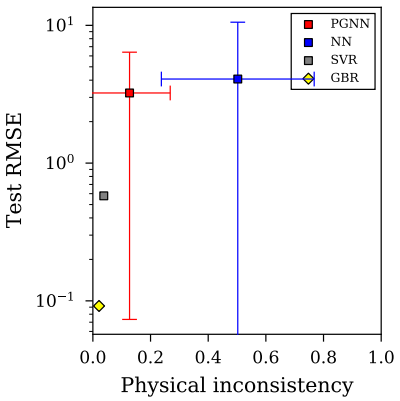

(d) Removing the minimum 25% of X_2 (scaled distance) data

Figure 5.15: Results showing root mean square error (RMSE) of removed test data and physical inconsistency for various extrapolation stress-test procedures.

5.3.3 Discussion

The results from the stress-testing evaluation for the NN and PGNN are compared in Table 5.2. In order to check the statistical significance of the results from each stress-test, two-tailed Kolmogorov-Smirnov tests (Hodges 1958) have been performed in each case. This is a two-tailed test for the null hypothesis that two independent samples are drawn from the same continuous distribution. In the perspective of these analyses, the test checks whether the RMSE or physical inconsistency values recorded by each model are drawn from the same distribution, therefore suggesting no statistically significant difference.

Interestingly, the GBR shows optimal performance in 4 of the 7 recorded stress-tests, and would suggest a promising model to include in further research. Both the GBR and SVR appear to show superior performance with regards to physical inconsistency from the previous Figures 5.13, 5.15, however this may be a slightly misleading result. The physical inconsistency term for the SVR and GBR is based on the trained models. In comparison, the neural networks physical inconsistency term include the amount of training iterations that produce physically inconsistent results, and provides a useful indicator into the usefulness of the imposed monotonic loss constraint, therefore it would be invalid to compare the neural networks to the SVR and GBR in respect of the physical inconsistency term. It is useful however that the trained SVR and GBR models show a low physical inconsistency, and the performance with regard to RMSE is still a valid comparison. However, the focus in this section is on the merits of the monotonic loss constraint and therefore comparison between NN and PGNN is more suitable.

Whilst extrapolation and interpolation with respect to X_1 appears insensitive to network type, it is clear that the PGNN outperforms the

Table 5.2: Mean results from each stress-testing evaluation of NN and PGNN models. Entries in bold indicate a statistically significant difference ($p < 0.10$) from the Kolmogorov-Smirnov two-tailed test statistic (Hodges 1958) for both performance premium and physical inconsistency (two-tailed p-value) for comparison between NN and PGNN only. SVR and GBR are included for additional benchmark comparison. "P.I.", "Max." and "Min." are abbreviations for physical inconsistency, maximum and minimum respectively.

Stress-test		**Test RMSE**				**P.I.**	
		NN	**PGNN**	**SVR**	**GBR**	**NN**	**PGNN**
Extrapolate	Max. X_1	0.15	0.17	0.52	0.28	0.49	**0.27**
	Max. X_2	0.45	**0.36**	1.11	0.58	0.69	0.59
	Min. X_1	4.05	4.05	0.93	0.89	0.30	0.19
	Min. X_2	4.08	3.23	0.58	0.09	0.50	**0.19**
Interpolate	Mean X_1	0.16	0.16	0.04	0.19	0.41	0.34
	Mean X_2	0.69	**0.60**	0.35	0.35	0.56	0.43
	Random	**0.67**	1.00	0.36	0.36	0.51	0.45

NN when extrapolating and interpolating with respect to X_2. There were three statistically significant findings with respect to test RMSE values, the results show that the PGNN significantly outperformed the NN when extrapolating beyond the maximum of X_2, and removing data around the mean of X_2, whereas the NN outperformed the PGNN when data was removed randomly.

The PGNN showed a statistically significant performance premium when extrapolating beyond the maximum of X_1, and beyond the minimum of X_2 with respect to physical inconsistency. When considering performance with respect to physical inconsistency on aggregate, however, it is shown that the PGNN outperforms the NN across all seven stress-tests. If the null hypothesis of equal mean performance between the PGNN and NN with respect to physical inconsistency is considered, the observed outperformance in all seven modes of testing has a probability of 0.007 ($p < 0.05$), a clear performance premium.

These results are highly promising from an engineering perspective.

The PGNN model will never under-predict when extrapolating. Essentially, the networks contain inbuilt conservatism in model predictions. This ensures that when analysing structures subjected to blast loading following close-in detonation of a high explosive, solutions informed by PGNNs will always be conservative; an approach that whilst not maximising efficiency will ensure safety. Extrapolating beyond data limits should always be handled with caution in predictive models, however the PGNN approach explored in this article provides additional conservatism and maintains physical consistency if extrapolation is ever required.

5.4 Assessment of surrogate model

5.4.1 Evaluation against numerical dataset

This section assesses the predictive ability of the studied surrogate models by comparing the results against the full CFD dataset. Filled contours of scaled peak specific impulse are shown in Figure 5.16 with the residuals, representing the difference between model predictions and actual data, featured alongside. The previously studied model presented in Chapter 4 was included as an additional surrogate model for comparison in Figures 5.16a and 5.16b. All the surrogate models have the same configurations as mentioned previously, with the neural network architectures being trained to 1000 epochs.

The strongest overall performance was shown by the GBR (Figures 5.16e and 5.16f) and the NN (Figures 5.16g and 5.16h), as shown by the residual plots. The PGNN (Figures 5.16i, 5.16j) showed a minor region of weak performance at extremely low values of Z and angle of incidence but otherwise performed similarly to the GBR and NN. The SVR showed a large area of weak performance for low values of Z generally. Importantly

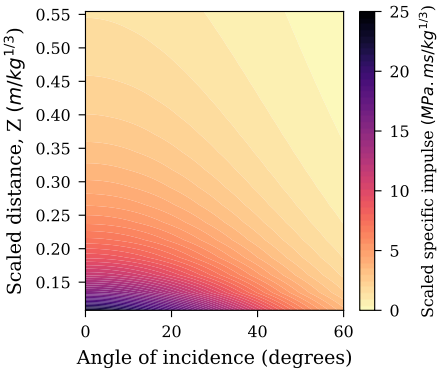

(a) Surrogate model predictions from Chapter 4.

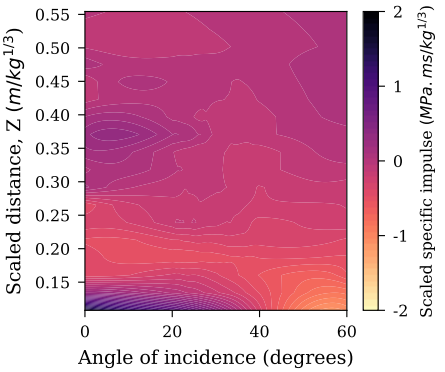

(b) Surrogate model residuals from Chapter 4.

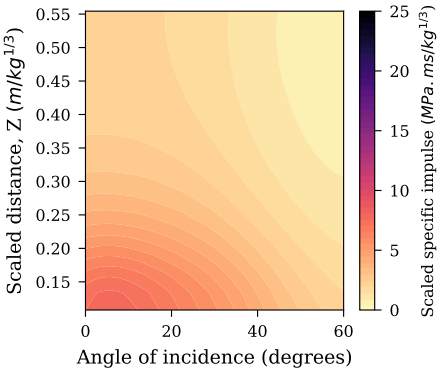

(c) SVR model predictions.

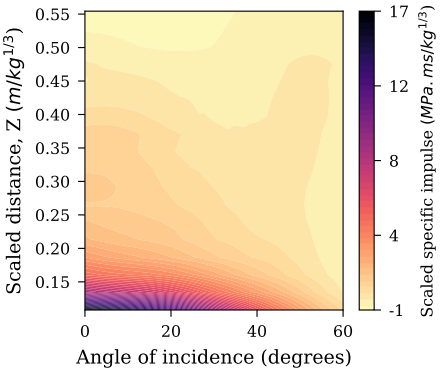

(d) SVR model residuals.

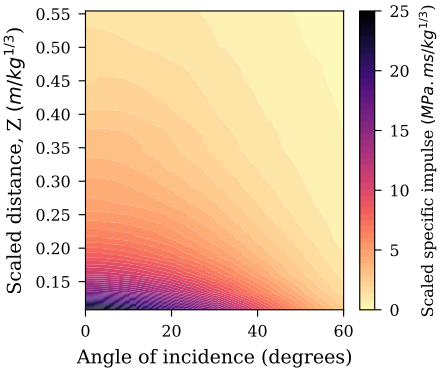

(e) GBR model predictions.

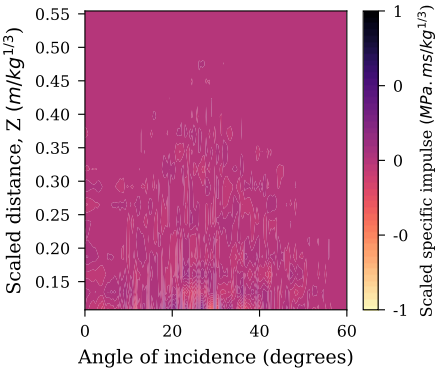

(f) GBR model residuals.

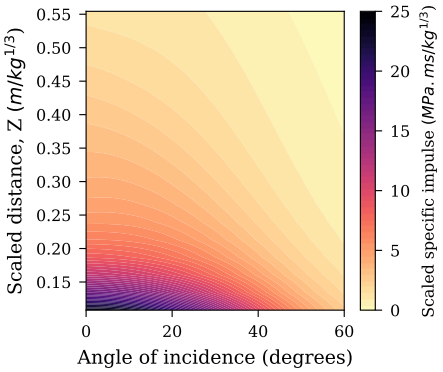

(g) NN model predictions.

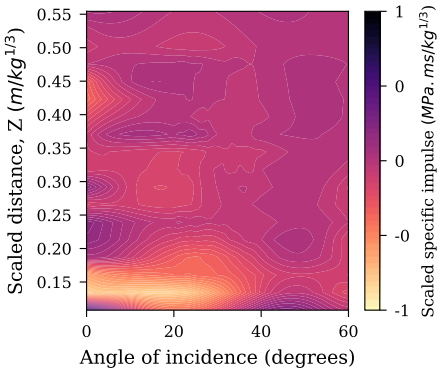

(h) NN model residuals.

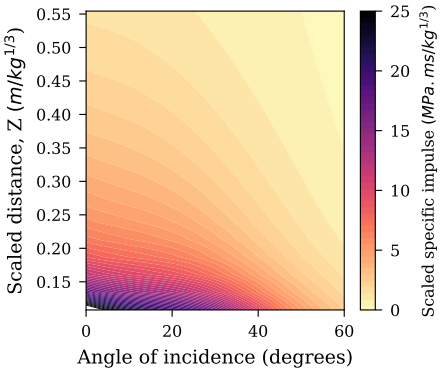

(i) PGNN model predictions.

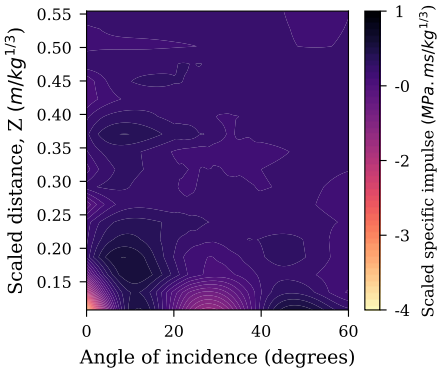

(j) PGNN model residuals.

Figure 5.16: Predicted specific impulse distributions for Z =0.11–0.55 m/kg$^{1/3}$: filled contours of scaled peak specific impulse for full dataset (a, c, e, g, i); residual errors for full dataset (b, d, f, h, j).

however, the GBR and the NN showed a better performance than the model presented in Chapter 4 and set a precedent for use of machine learning surrogate models that can handle greater complexity.

5.4.2 Comparison against unseen CFD data

The predictive model was compared against two sets of unseen CFD data (i.e. data not used in the development of the predictive model) as in Chapter 4. The first set comprises two models: a 5 kg spherical PE4 charge detonated at 0.291 m normal stand-off distance (to charge centre) from a rigid surface ($Z = 0.17$ m/kg$^{1/3}$); and a 250 kg spherical PE4 charge detonated at 2.52 m normal stand-off distance (to charge centre) from a rigid surface ($Z = 0.40$ m/kg$^{1/3}$). Note that these models were not used to form the CFD datasets outlined previously, and hence serve as unseen data for comparative purposes. The CFD information for these simulations is given in Chapter 4 and not repeated here for brevity.

Figure 5.17 shows a comparison of the CFD and predicted specific impulse distributions out to 0.5 m and 4 m from the target centre. Here, predictions for the $Z = 0.17$ m/kg$^{1/3}$ case were determined using equation 4.4, whilst predictions for the $Z = 0.40$ m/kg$^{1/3}$ case were determined using equation 4.5. In both cases the previously studied surrogate models were included as an additional means of comparison, excluding the SVR, omitted for its poor performance in the previous analysis. It can be seen that the predicted distributions closely match those from the CFD analyses. The minor region of slightly poorer performance previously mentioned for the PGNN at low Z and theta is demonstrated here where the PGNN shows some fluctuations at < 0.1 m in Figure 5.17a and < 1 m in Figure 5.17b. This is likely due to an interaction between the physical inconsistency penalty and the MSE where the physical incon-

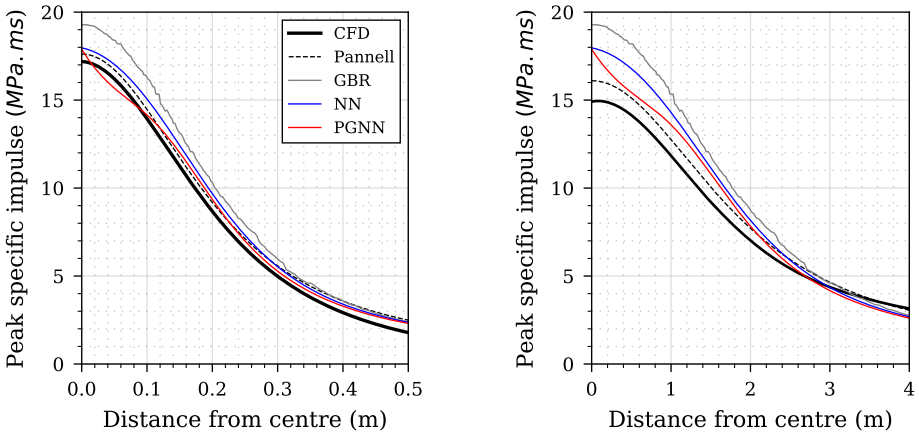

Figure 5.17: CFD and predicted peak specific impulse distributions: (a) 5 kg at 0.291 m stand-off, $Z = 0.17$ m/kg$^{1/3}$; (b) 250 kg at 2.52 m stand-off, $Z = 0.40$ m/kg$^{1/3}$. "Pannell" refers to the models proposed in 4.

sistency penalty may have influenced more than ideally required at this point. The regularisation hyper-parameter that controls this influence was chosen from the prior grid search, and a more resolute search could be completed within the chosen sub-region to find a more suitable value. Alternatively an additional second order differential term could be added to the objective function to "smooth" the profile, or more complex loss terms than a monotonic loss could be employed generally (such as the predictive equation in Chapter 4). However, this would further add complexity to the procedure, and considering this is a small region of weak performance should be taken in context with the entire surface where strong performance is generally shown. Furthermore, when calculating the area-integrated total impulse from this, the discrepancy will have an exceedingly minor influence. Of the remaining two machine learning models the NN shows good performance, with the weakest performance shown by the GBR which noticeable over-predicts at < 0.1 m in Figure 5.17a and < 1 m in Figure 5.17b.

5.5 Summary

This chapter presents a novel application of a physics-based regularisation framework for neural networks in the application of predicting blast loading following detonation of an explosive charge. The physics-based regularisation is shown to provide statistically significant performance improvements over traditional neural network procedures when generalising beyond the dataset through extrapolation and interpolation with regards to accuracy (shown via test RMSE) and physical validity (via physical inconsistency metric). This finding provides positive evidence to the theme of research in using physics-based regularisation when training ML models, and the use of ML models generally, especially in applications involving the prediction of blast loads resulting from the detonation of high explosives when operating in a sparse-data domain.

An initial parameter search study was completed for a traditional neural network (NN), with the chosen parameters also used in the physics-guided neural network (PGNN). Both models were stress-tested through various data holdout approaches, and their ability to predict high explosive blast loading were comprehensively evaluated and compared. It is found that the PGNN out-performed the NN in (3/7) of the stress-tests (lower RMSE), and exhibited lower physical inconsistency in the training process in 100% (7/7) of the stress-tests ($p < 0.05$). With respect to the specific mode of stress-testing, of the PGNN's 10 performance premiums, four were found to be statistically significant (RMSE when extrapolating with respect to maximum X_2 and interpolating with respect to mean X_2, and physical inconsistency when extrapolating with respect to maximum X_1 and minimum X_2). Two additional data-driven surrogate models are tested (a GBR and SVR) where it can be seen that in the optimum performance with respect to RMSE is shown by the GBR in 4 of the 7

cases. Therefore there is a benefit to implementing a physics based regularisation procedure in neural networks, but it needs to be considered in the context of other machine learning models and the performance of those. It is also worth considering that neural networks and other data-driven methods would have a number of hyper-parameters: hard (such as number of hidden units) or soft (number of data points and distribution). These hyper-parameters are not considered in an overarching multi-objective optimisation process within this thesis (as it is considered to be beyond the scope of this thesis), meaning that the "optimum" trained models presented are, in essence, still sub-optimal. The fundamental concept, however, of embedding a desired control (enforcing physically valid solutions) through an obtrusive form of regularisation, a direct modification to the objective function, is shown to produce the desired effect – as shown by the smaller physical inconsistency value in 7/7 stress tests.

Additionally, this research has considered the importance of careful consideration of the parameters which govern a network's ability to generalise, and subsequent implications for experimental design. The results demonstrate the feasibility and enhanced accuracy achievable when developing ML models to predict near-field blast loading. As a continuation of this research, future work will involve incorporating different charge types, shape effects, and blast wave clearing as features in the model. An additional research direction is to perform an overarching multi-objective optimisation process when training data-driven models. Further, incorporating more complex physics (through implementing a non-linear loss constraint instead of the monotonic loss constraint) will enable the possibility of improved PGNN performance through transfer learning. Developments in this area will facilitate more accurate probabilistic-based

approaches to engineering design and risk mitigation that encompass a more complex suite of scenarios than is capable presently.

Chapter 6

Transfer learning – what information can be learned from a spherical charge?

6.1 Introduction

As established by Rigby et al. (2019b) and discussed in Chapter 4, to perform the accurate appraisal of structures and protective systems, knowledge of the distribution and magnitude of loading is required. A surrogate model allows the analyst to instantly obtain the loading information, within the parameters of the surrogate model, for a multitude of scenarios (that would otherwise be costly to ascertain) and is the first step towards a probabilistic mode of risk assessment. The preliminary surrogate model presented in Chapter 4 is an equation made of three separate terms and is suitable for a specific charge shape, type and range of scaled distances.

However, to increase the capabilities of the surrogate model, and therefore the situations an analyst can simulate, a model that can handle

the additional complexity is required. Integrating data-driven methods with scientific theory is considered crucial in order to improve surrogate model performance whilst respecting natural laws (Reichstein et al. 2019). Chapter 5 investigates this by implementing a physics-based regularisation procedure when training a machine learning model through adding a monotonic loss constraint to the loss function. This could be considered a form of transfer learning, by implementing the discoveries from Chapter 4 it has informed the learning procedure described in Chapter 5 (the physics based regularisation) to enhance performance in a specific task (producing physically valid predictions).

Traditional data mining and machine learning algorithms provide predictions on future data using statistical models trained on previously collected labelled or unlabelled data (Pan & Yang 2010, Ramon et al. 2007, Taylor & Stone 2007). Many machine learning methods work under the assumption that the training and test data belong to the same distribution. When this distribution changes, most statistical models need to be re-trained on newly collected training data. Though some methods do exist that model non-stationary data where the "data-drift" is parameterised and modelled, alternatively there are heuristic methods for continuous learning (Panoutsos & Mahfouf 2008). In the context of blast protection engineering, obtaining data is considerably expensive in time and cost (as discussed in Chapter 2) and therefore any method that increases the utility of this data is of paramount importance, as it would be in many other applications. In these cases, it would be highly useful to reduce the need to re-collect training data and therefore knowledge transfer, or transfer learning, between task domains is highly desirable. Many examples exist where transfer learning can be beneficial such as web-document classification (Mahmud & Ray 2007, Blitzer et al. 2008,

Xing et al. 2007), sentiment classification (Li et al. 2009), image classification (Lee et al. 2007), WiFi localisation models (Yin et al. 2005, Raina et al. 2006, Pan et al. 2007, 2008, Zheng et al. 2008) and web-page translation (Ling et al. 2008). For an insight into the benefit transfer learning can explicitly bring over traditional machine learning approaches, see Table 5 in Pan & Yang (2010).

This chapter presents two novel applications of transfer learning for the prediction of peak specific impulse. Firstly to model alternative spherical charge types and secondly to model cylinders, both using knowledge already gained when modelling spheres of a single charge type. Discussion on dataset generation is provided in each case and assessments of the proposed models are presented where it is shown clearly that by implementing transfer learning, the need for new training data is drastically reduced.

6.2 Transfer learning

Transfer learning, and domain adaptation, refer to the situation where what has been learned in one scenario is exploited to improve generalisation in a second scenario. The inherent assumption is that the factors that influence variations in the first scenario also apply, to some level, to the second. In the real world, there are many clear examples of transfer learning. For example, one may find that learning to play the organ will facilitate learning the piano. The field of transfer learning is motivated by this awareness that people can apply previously learned knowledge intelligently when faced with a new problem and can solve it more quickly or with better solutions (Pan & Yang 2010).

To aid understanding of transfer learning it is useful to have some

formal notation and definitions. Firstly, the definitions of "*domain*" and "*task*". A *domain*, $\mathcal{D}$, consists of a feature space $\mathcal{X}$ and a marginal probability distribution $P(X)$, where $X = \{x_1, ..., x_i\} \in \mathcal{X}$. Consider as an example the learning task of document classification where each term is taken as a binary feature, $\mathcal{X}$ is the vector space for all terms, x_i is the ith term vector (corresponding to some documents) and X is a particular sample. It can generally be considered that if two domains are different, then they may have different feature spaces or marginal probability distributions (Pan & Yang 2010).

Given a specific domain, $\mathcal{D} = \{\mathcal{X}, P(X)\}$ a task consists of two components: label space $\mathcal{Y}$ and an objective predictive function $f : \mathcal{X} \rightarrow \mathcal{Y}$ and is denoted $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$. This predictive function is learned from the training data, pairs $\{x_i, y_i\}$ where $x \in X$ and $y \in \mathcal{Y}$ and can be used to predict new labels $f(x)$ from an instance x (Pan & Yang 2010).

A definition of transfer learning is given as follows: "*Given a source domain $\mathcal{D}_S$ and learning task $\mathcal{T}_S$, a target domain $\mathcal{D}_T$ and learning task $\mathcal{T}_T$, transfer learning aims to help improve the learning of the target predictive function $f_T(\cdot)$ in $\mathcal{D}_T$ using the knowledge in $\mathcal{D}_S$ and $\mathcal{T}_S$, where $\mathcal{D}_S \neq \mathcal{D}_T$, or $\mathcal{T}_S \neq \mathcal{T}_T$.*".

In the above definition, from Pan & Yang (2010), a domain is a pair $\mathcal{D} = \{X, P(X)\}$. So the condition that $D_S \neq D_T$ has the implication that $X_S \neq X_T$ or $P_S(X) \neq P_T(X)$. Likewise, a task is defined as a pair $\mathcal{T} = \{\mathcal{Y}, P(Y|X)\}$, therefore the condition $\mathcal{T}_S \neq \mathcal{T}_T$ implies $\mathcal{Y}_S \neq \mathcal{Y}_T$ or $P(\mathcal{Y}_S|X_S) \neq P(\mathcal{Y}_T|X_T)$. If the source and target domains are the same $\mathcal{D}_S = \mathcal{D}_T$, and their learning tasks are the same $\mathcal{T}_S = \mathcal{T}_T$, the problem then becomes a classical machine learning problem.

There are considered to be three main research questions in the field of transfer learning: 1) what to transfer, 2) how to transfer, and 3) when

to transfer. "What to transfer" is concerned with ascertaining which part of knowledge from the source can be transferred, and what may be useful knowledge to transfer for improving performance in the target domain or task. "How to transfer" is concerned with choosing a learning algorithm that can transfer the knowledge from the source to the task, and "when to transfer" considers when the transfer of knowledge should be implemented. An important point for consideration here is that it is equally useful in knowing when not to transfer as in when to transfer. When transfer learning takes place and is harmful to performance in the target, it is referred to as negative transfer (Pan & Yang 2010).

The overall objective, therefore, of transfer learning is to take advantage of knowledge from the source domain ($\mathcal{D}_S$), and use this to improve performance when learning, or making predictions in the target domain ($\mathcal{D}_T$) (Goodfellow et al. 2016). There are clear advantages to allowing more accurate predictions in the target domain, but a unique benefit to transfer learning is that accuracy of predictions can be high, even when data is sparse or severely limited (in $\mathcal{D}_S$). This has the practical benefit that gathering new data is not as important and there can be substantial savings in cost and time. For an extensive overview of advancements in transfer learning across a wide range of settings and implementations, see Pan & Yang (2010).

6.3 Modelling charge composition effects with transfer learning

6.3.1 Dataset overview

The first application of transfer learning implements the surrogate model presented in Chapter 4 for spherical charges, in a similar range of scaled distance but for different charge compositions. In total, five new datasets were generated from CFD simulations using *Apollo* for 1kg spherical charges of five varying charge compositions: TNT, HMX, RDX, PETN and COMPB. For each dataset, *Apollo* was used to generate the specific impulse distribution along a rigid reflecting surface at angles of incidence between 0-60° with 150 gauges linearly spaced (with respect to angle of incidence) along the rigid reflecting surface. Six different stand-off distances were analysed for each charge type from 0.2 m–0.6 m in increments of 0.1 m. The ultimate cell size adheres to the minimum S/50 in all cases whilst the DMA module was used throughout.

For all analyses, a separate domain was specified for each model, with equal side lengths of an integer number of zone lengths, given by: $\lceil 1.2S\tan(60°)/L \rceil$, where S and L are stand-off and zone length as previously, such that $S\tan(60°)$ is the distance from the centre of the target to the most remote gauge. The domain length was set a factor of 1.2 greater than this distance (rounded up to the nearest multiple of zone length), again to prevent edge expansion waves from reaching the most remote gauge location during the analysis. If this distance was less than 1 m, then 1 m was used as the side lengths. *Apollo*'s staging procedure was used where a spherical cutout region of the entire domain is specified to resolve the detonation in greater detail but conserve computational efficiency. The first stage was solved in one-dimension until a boundary

is reached (equivalent to the standoff) before implementing the DMA module. The custom DMA conditions were defined as a 0.2 m spherical region around the charge, with an additional 2 resolution levels and a growth rate of 2. When the shock front reaches the defined distance (0.2 m), the resolution level is dropped by 1 and the new domain is the previous radius multiplied by the growth rate, so becomes 0.4 m. This continues until the additional resolution levels reach 0 and the entire domain is then included in the simulation. Quarter symmetry was used, with 100 mm initial zone length and intial resolution level of 3.

For all explosive types, *Apollo*'s in-built model parameters were used. The pressure-density-temperature relationship of the post-detonation explosive products is given by the Jones Wilkins Lee equation of state (Lee et al. 1968):

$$p(\rho, T) = C_1 e^{-R_1 \rho_0/\rho} + C_2 e^{-R_2 \rho_0/\rho} + \rho R T \tag{6.1}$$

where ρ and ρ_0 are density and initial density, R is a gas constant, and T absolute temperature. The constants C_1, C_2, R_1, and R_2 are assigned the values given in Table 6.1. Air is modelled $p = \rho RT$, such that under ambient conditions ($288°K$), ambient pressure is $p = 101.3$kPa. Afterburn was modelled using the Klomfass Afterburning (KAB) model, and the Chapman-Jouguet detonation model was used in all cases. In all numerical simulations, the explosives were centrally detonated and the mass of the detonator was not included. A Savitzky-Golay filter (Savitzky & Golay 1964) has been used to remove spurious oscillations induced by the DMA procedure.

In summary, there are 30 different CFD models (6 different stand-off distances for each of the 5 different charge types) with 150 values of peak specific impulse within each, therefore a total of 4500 data points.

Table 6.1: Equation of state information for the five newly studies charge compositions, including the previously studied PE4.

Charge type	**Gas constant** $J/(kg.K)$	**Initial density** (kg/m^3)	C_1 **(GPa)**	C_2 **(GPa)**	R_1	R_2
PE4	365	1660	734.60	8.86	4.79	1.06
TNT	315	1630	527.28	6.30	4.71	1.07
HMX	365	1905	1215.93	12.30	4.77	1.12
RDX	345	1805	1053.18	11.29	4.77	1.12
PETN	300	1778	946.18	9.83	4.76	1.08
COMPB	345	1725	732.74	8.83	4.68	1.09

The six stand-off distances represent a scaled distance, Z, range of 0.2 m/kg$^{1/3}$–0.6 m/kg$^{1/3}$ (ignoring any TNT equivalence), this range was chosen as it partly lies within the range the original model was built (0.11 m/kg$^{1/3}$–0.55 m/kg$^{1/3}$) but also extends slightly beyond.

6.3.2 Model development

As shown in Chapter 4, the surrogate model for a dataset of spherical charges of PE4 in the near-field consists of three different parts: a term to capture the angle of incidence effects, a term to determine the peak perpendicular specific impulse value and a term to scale the charge mass. The new datasets cover a similar range of scaled distance values and consist of different charge compositions, but crucially are the same shape, spheres. It would make physical sense therefore, that the term that handles the angle of incidence effects would be similar for different charge compositions (due to equivalent charge shape), and the term that would vary the most would be the term that determines the peak normal specific impulse value. Hence the knowledge of the angular distribution of peak specific impulse in spherical charges is transferred to other charge compositions, and the only new part that needs to be "learned" is the

part of the model that defines peak normal impulse (although this is still known to be a power law).

To evaluate whether the angle of incidence component can be transferred in this way, two models were fitted for each new dataset to the same format as the model proposed in Chapter 4. The first of these models fits new parameters for both the peak impulse component and angle of incidence component and are presented below in Equations 6.2. The second model only uses the new peak scaled impulse component for the different charge type, and transfers the original angle of incidence component for PE4 spheres from Chapter 4, these models are summarised in Equations 6.3. The model fit process for the peak impulse component and assessment of this fit is shown in Figure 6.1 where it is seen that a power law still provides a good model even when the charge composition is varied. This is demonstrated from the low values obtained for residual standard error and the high values for coefficient of determination in each case.

$$\text{TNT: } i(Z,\theta,W) = 0.383Z^{-1.683}\exp\left(\frac{-\theta^2}{2282}\right)W^{1/3} \tag{6.2a}$$

$$\text{HMX: } i(Z,\theta,W) = 0.435Z^{-1.664}\exp\left(\frac{-\theta^2}{2301}\right)W^{1/3} \tag{6.2b}$$

$$\text{RDX: } i(Z,\theta,W) = 0.426Z^{-1.671}\exp\left(\frac{-\theta^2}{2288}\right)W^{1/3} \tag{6.2c}$$

$$\text{PETN: } i(Z,\theta,W) = 0.419Z^{-1.660}\exp\left(\frac{-\theta^2}{2308}\right)W^{1/3} \tag{6.2d}$$

$$\text{COMPB: } i(Z,\theta,W) = 0.421Z^{-1.668}\exp\left(\frac{-\theta^2}{2298}\right)W^{1/3} \tag{6.2e}$$

$$\text{TNT}_{\theta\text{-transfer}}\text{: } i(Z,\theta,W) = 0.383Z^{-1.683}\exp\left(\frac{-\theta^2}{2007}\right)W^{1/3} \quad (6.3a)$$

$$\text{HMX}_{\theta\text{-transfer}}\text{: } i(Z,\theta,W) = 0.435Z^{-1.664}\exp\left(\frac{-\theta^2}{2007}\right)W^{1/3} \quad (6.3b)$$

$$\text{RDX}_{\theta\text{-transfer}}\text{: } i(Z,\theta,W) = 0.426Z^{-1.671}\exp\left(\frac{-\theta^2}{2007}\right)W^{1/3} \quad (6.3c)$$

$$\text{PETN}_{\theta\text{-transfer}}\text{: } i(Z,\theta,W) = 0.419Z^{-1.660}\exp\left(\frac{-\theta^2}{2007}\right)W^{1/3} \quad (6.3d)$$

$$\text{COMPB}_{\theta\text{-transfer}}\text{: } i(Z,\theta,W) = 0.421Z^{-1.668}\exp\left(\frac{-\theta^2}{2007}\right)W^{1/3} \quad (6.3e)$$

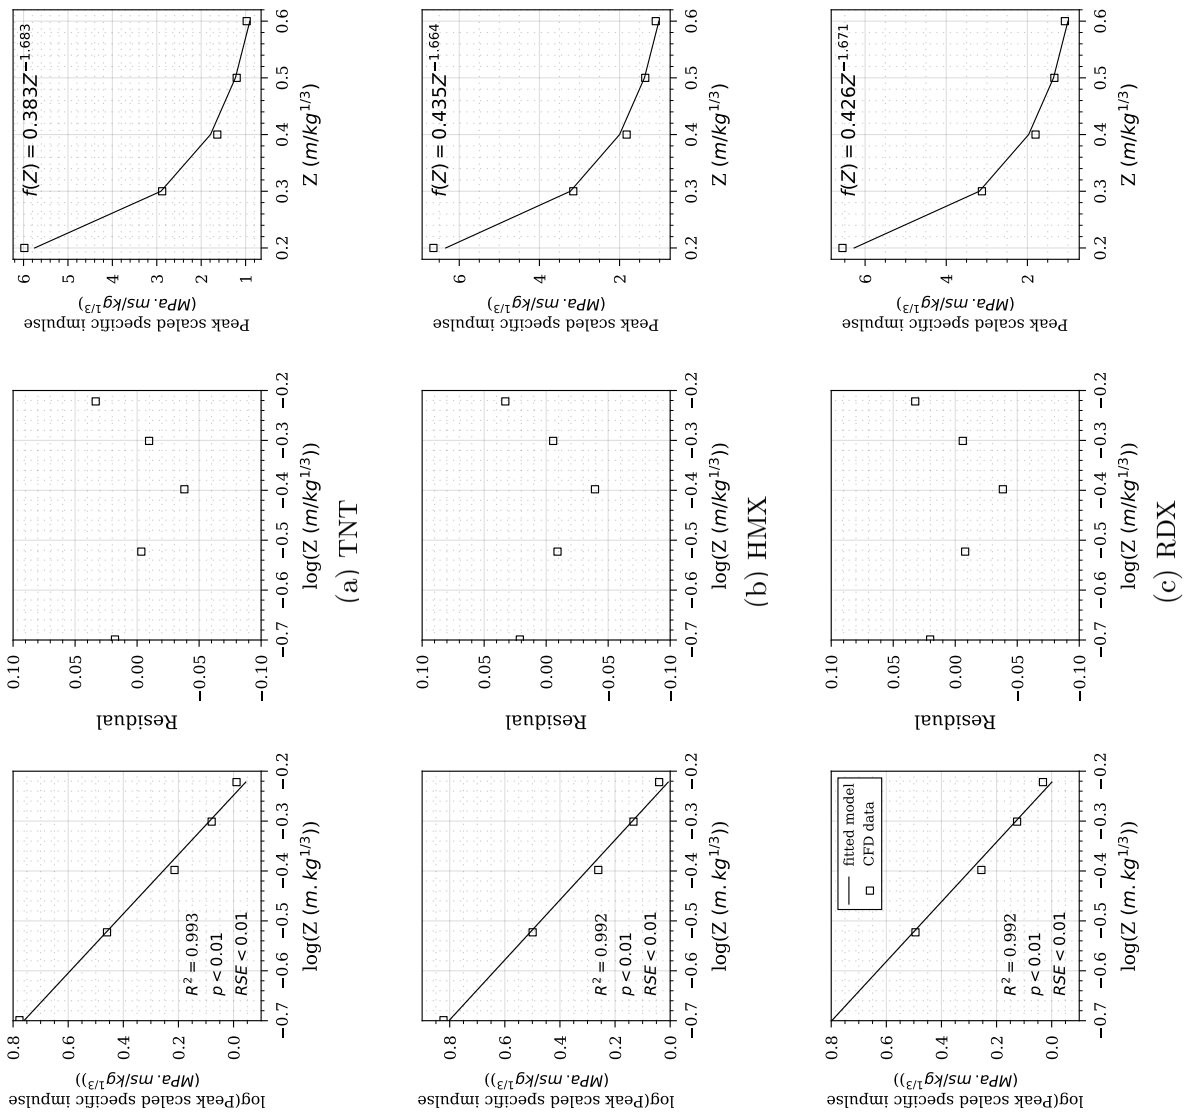

(a) TNT (b) HMX (c) RDX

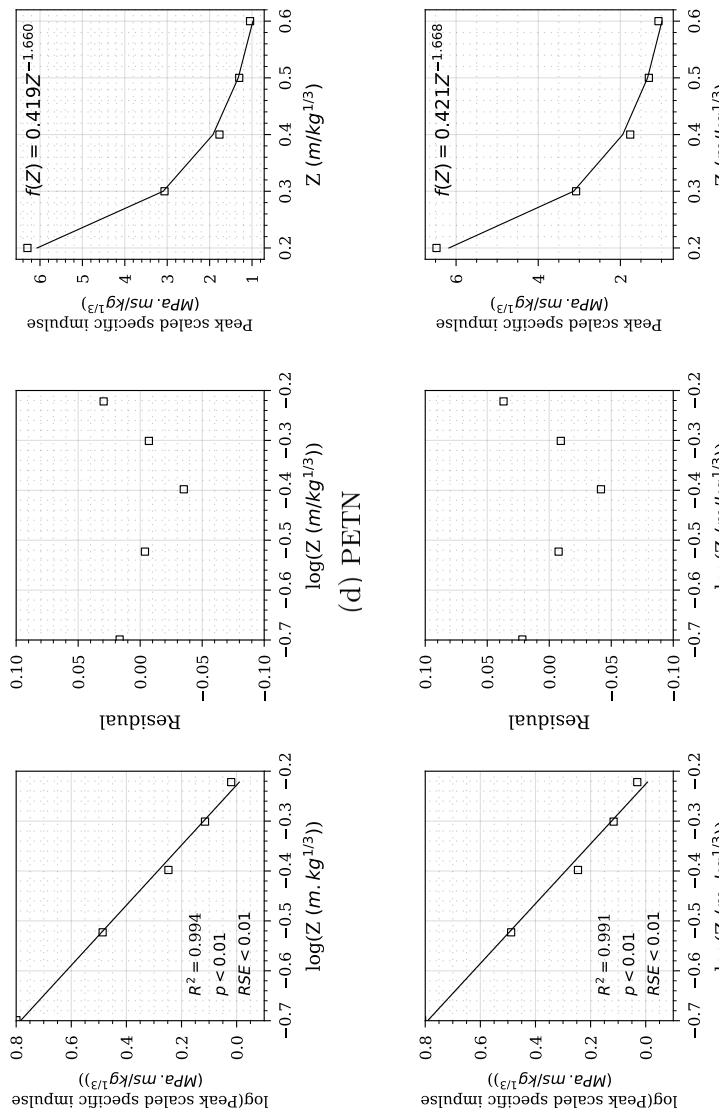

(d) PETN

(e) COMPB

Figure 6.1: Relationship between peak normally reflected specific impulse and scaled distance, residuals, and proposed fit for five different charge types.

6.3.3 Model verification

The evaluation of the new models and the transfer learning process for each charge composition is shown in Figures 6.2 to 6.6. In each case: the peak impulse data and its model fit is shown; followed by the angular component of the model which includes the transferred model, the new fit and the normalised CFD data. Then the entire dataset is predicted using both the transferred angle component and the new fitted angle component; with the residuals presented by subtracting the predictions from the true CFD values, this demonstrates how accurately each model fits to the true CFD data.

For each new dataset, the mean absolute error of the residuals was calculated for the transferred model and the new model and presented in Table 6.2. There is negligible difference between the two models for each charge composition which confirms that the angle of incidence component of PE4 spheres can be transferred successfully to other charge compositions. This is a highly useful finding as it means for new charge compositions, only a few normally reflected CFD models are required (to obtain the peak specific impulse values for different values of scaled distance) and then the entire dataset can be constructed from this.

This exercise has demonstrated that previously learned knowledge obtained during modelling of near-field spherical charges of a certain charge composition and scaled distance range can be used to rapidly speed up the learning in a new situation. Knowledge from the source domain $\mathcal{D}_S$ and the source task $\mathcal{T}_S$ was used to improve performance in the new domain $\mathcal{D}_T$ and task $\mathcal{T}_T$ where $\mathcal{D}_S \neq \mathcal{D}_T$.

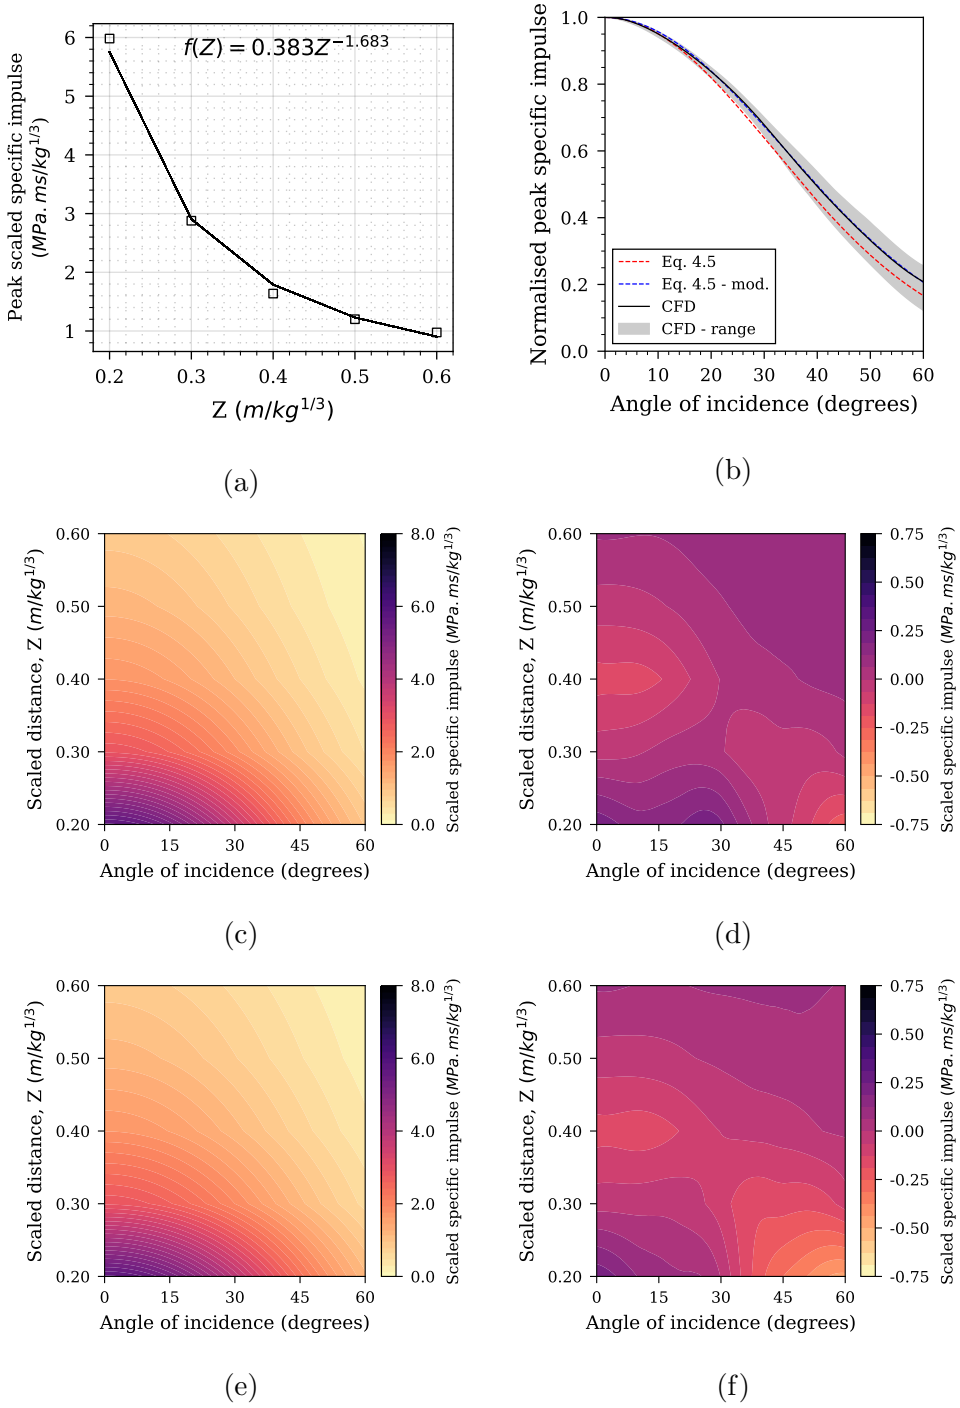

Figure 6.2: Transfer learning of spherical TNT charges. (a) newly acquired peak scaled impulse data, (b) the transferred angular distribution from PE4 spheres and the updated angular distribution built from new TNT data, (c) predicted peak specific impulse surface from the transferred model, (d) transfer model residuals with respect to true data, (e) predicted peak specific impulse surface from the new model, (f) new model residuals with respect to true data

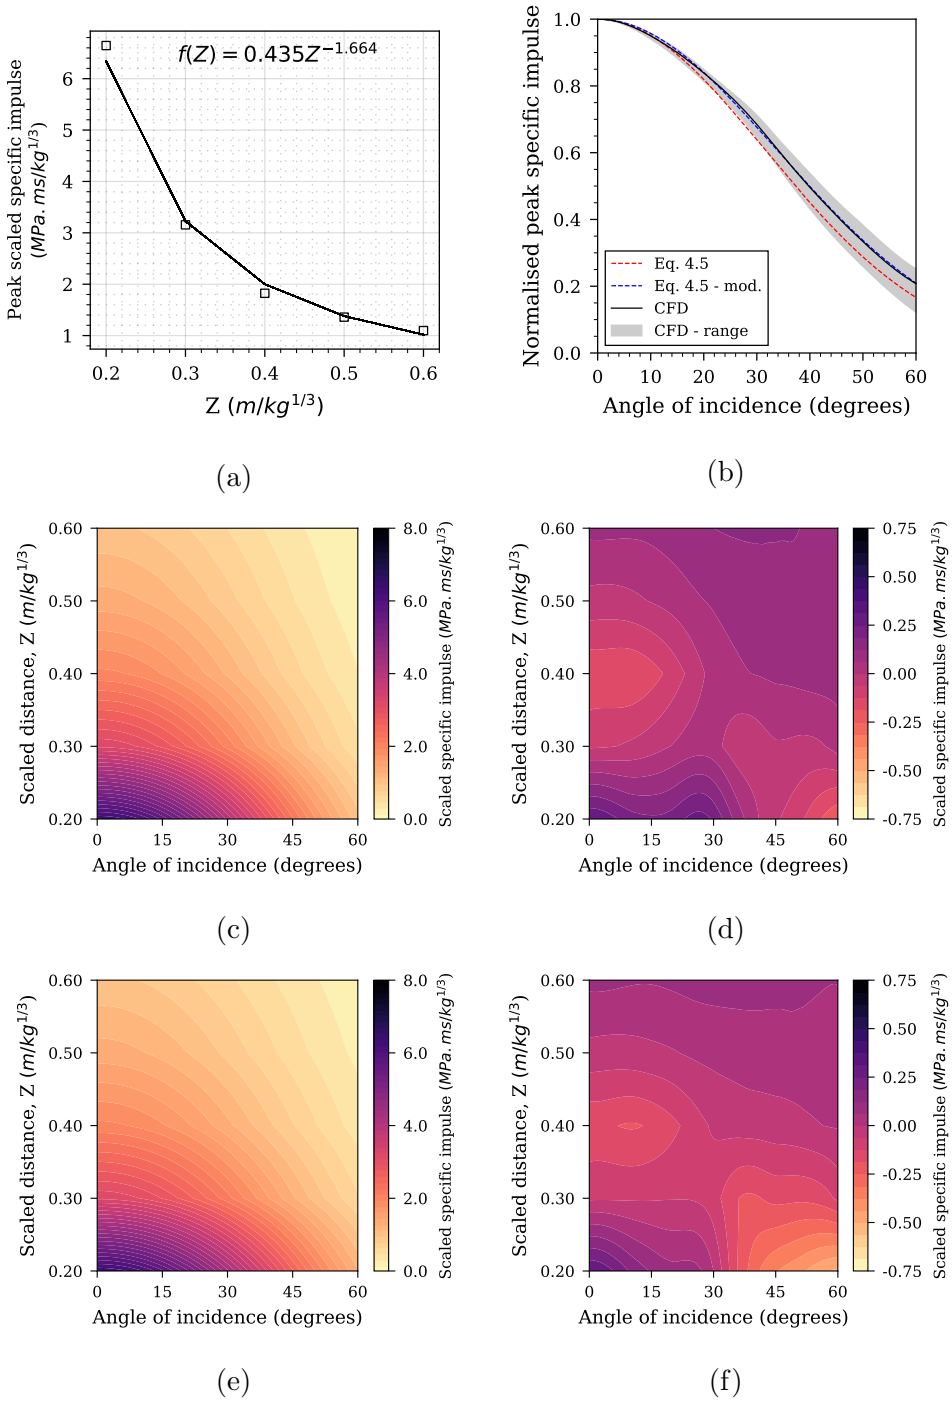

Figure 6.3: Transfer learning of spherical HMX charges. (a) newly acquired peak scaled impulse data, (b) the transferred angular distribution from PE4 spheres and the updated angular distribution built from new HMX data, (c) predicted peak specific impulse surface from the transferred model, (d) transfer model residuals with respect to true data, (e) predicted peak specific impulse surface from the new model, (f) new model residuals with respect to true data

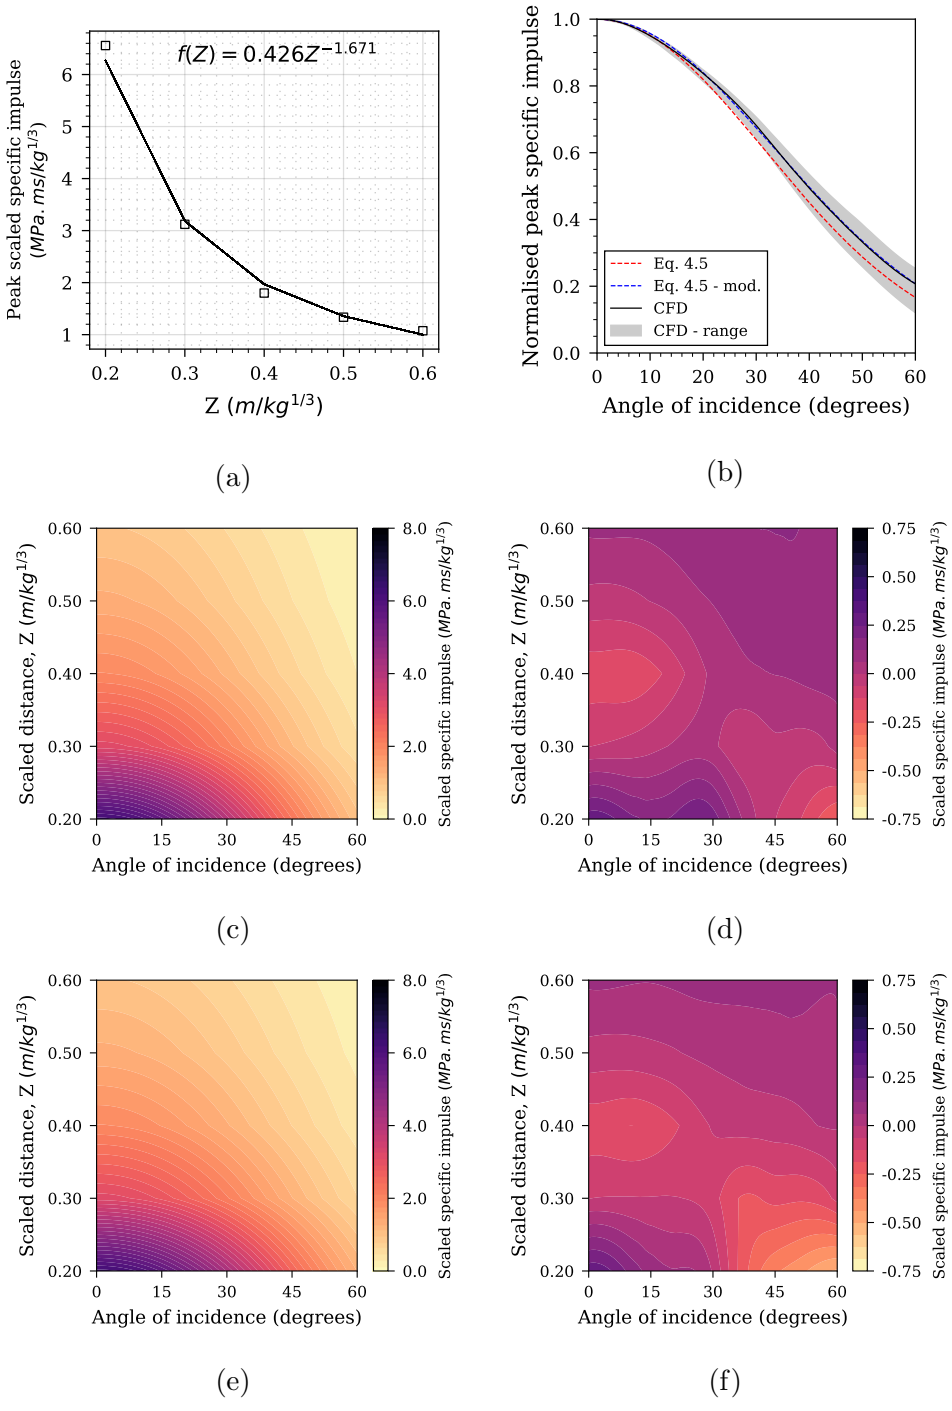

Figure 6.4: Transfer learning of spherical RDX charges. (a) newly acquired peak scaled impulse data, (b) the transferred angular distribution from PE4 spheres and the updated angular distribution built from new RDX data, (c) predicted peak specific impulse surface from the transferred model, (d) transfer model residuals with respect to true data, (e) predicted peak specific impulse surface from the new model, (f) new model residuals with respect to true data

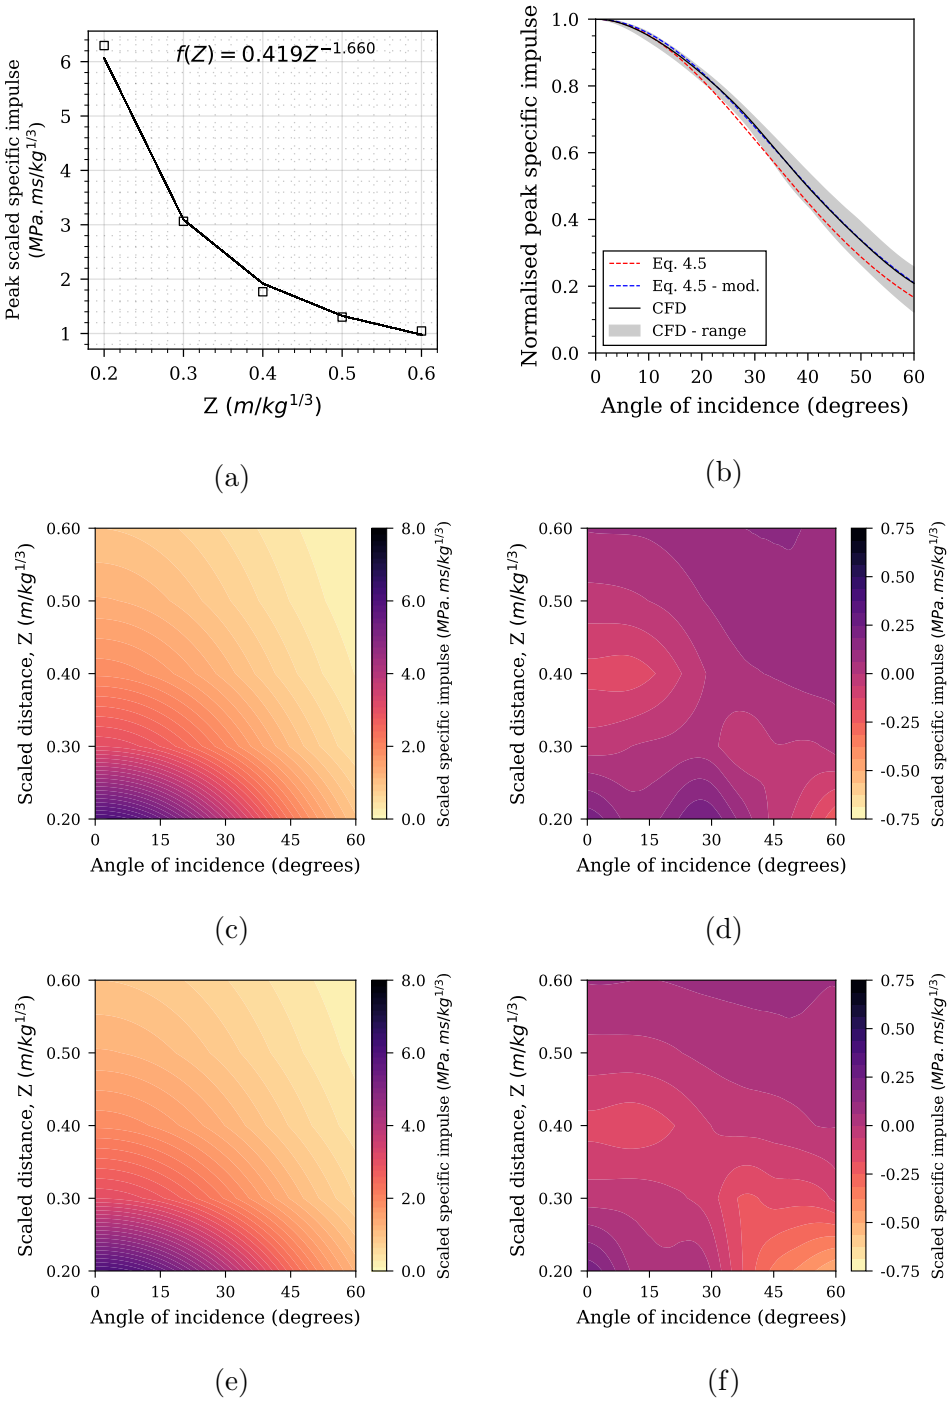

Figure 6.5: Transfer learning of spherical PETN charges. (a) newly acquired peak scaled impulse data, (b) the transferred angular distribution from PE4 spheres and the updated angular distribution built from new PETN data, (c) predicted peak specific impulse surface from the transferred model, (d) transfer model residuals with respect to true data, (e) predicted peak specific impulse surface from the new model, (f) new model residuals with respect to true data

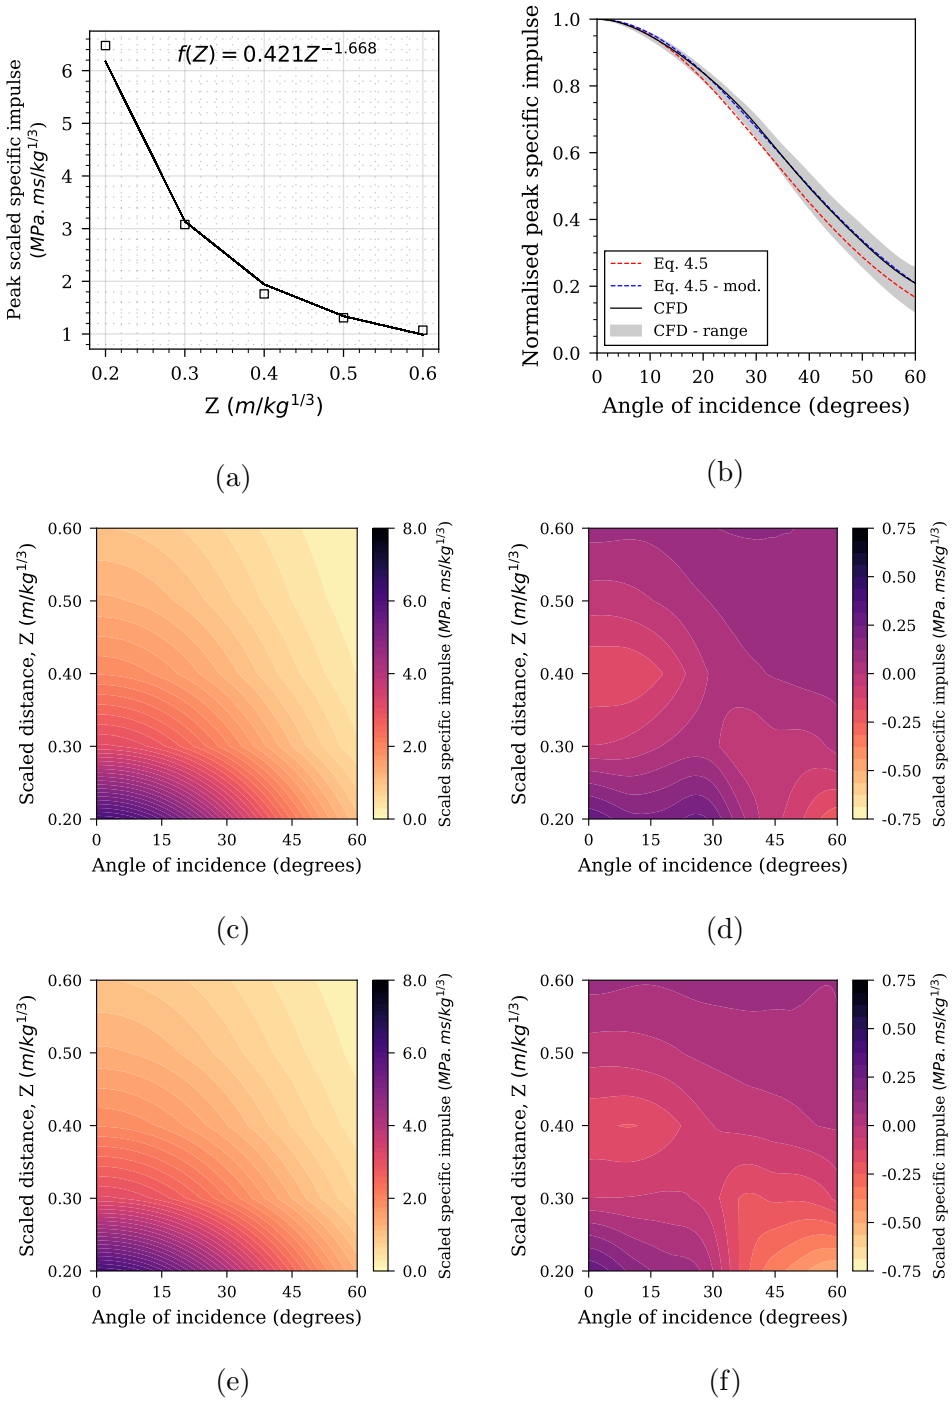

Figure 6.6: Transfer learning of spherical COMPB charges. (a) newly acquired peak scaled impulse data, (b) the transferred angular distribution from PE4 spheres and the updated angular distribution built from new COMPB data, (c) predicted peak specific impulse surface from the transferred model, (d) transfer model residuals with respect to true data, (e) predicted peak specific impulse surface from the new model, (f) new model residuals with respect to true data

Table 6.2: Mean absolute error of residual values from Figures 6.2 to 6.6, given to 2 dp.

Charge composition	**MAE (MPa.ms/kg$^{1/3}$)**	
	Transfer	**New**
TNT	0.08	0.09
HMX	0.10	0.11
RDX	0.09	0.11
PETN	0.08	0.09
COMPB	0.09	0.11

6.4 Modelling charge shape effects with transfer learning

6.4.1 Numerical modelling of cylinders in *Apollo Blastsimulator*

Prior to validating Apollo results against experimental data, a mesh sensitivity study was conducted with the aims of determining the required element size to achieve convergence and identifying suitable combinations of zone length and resolution level for cylindrical explosives. The chosen model set-up was a 0.078kg PE4 squat cylinder (L/D = 1/3) at 0.1774 m stand-off, which was chosen due to having experimental data for this case from Rigby et al. (2019b) and is presented schematically in Figure 6.7. The results of the mesh sensitivity study are shown in Figure 6.8. Figure 6.9 presents the studied meshes compared to the experimental peak specific impulse distribution where it can be shown that a mesh with a S/cell length from 168 can be considered suitable.

The CFD model with S/cell length of 336 was chosen for further analysis, with the overpressure-time histories and impulse-time histories compared with experimental data and presented in Figure 6.10 where

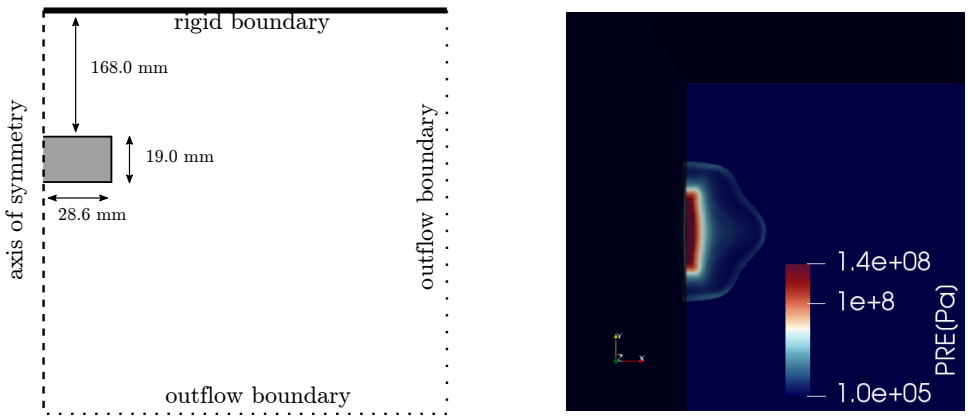

(a) (b)

Figure 6.7: 0.078kg PE4 squat cylinder (L/D = 1/3), (a) CFD model set-up and (b) blast wave development moments after detonation.

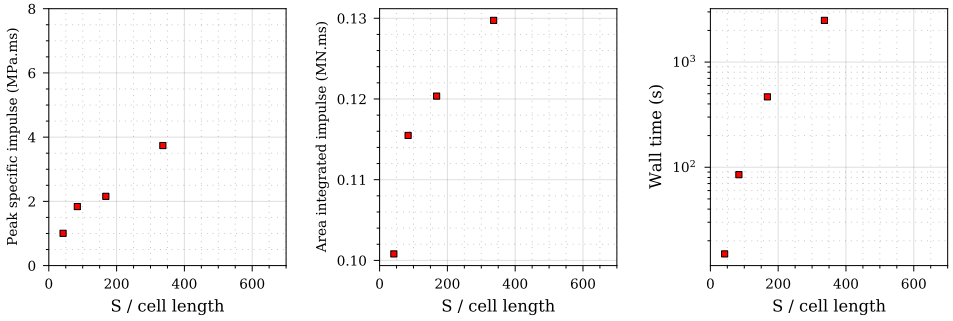

Figure 6.8: Mesh convergence study for 0.078kg PE4 cylinder , Z = 0.415 m/kg$^{1/3}$, standoff from charge centre = 0.1774m, L/D = 1/3.

good agreement between CFD and experimental data is shown.

6.4.2 Dataset overview

The dataset was generated from CFD simulations using *Apollo* consisting of centrally-detonated 100g cylinders of PE4 located between 0.09m-0.23m from the centre of the charge to a target. This is equivalent to a scaled distance range of $0.2-0.5$m/kg$^{1/3}$ where 5 linearly spaced Z values were chosen to sample from. Four different L/D ratios were chosen of 1/5, 1/3, 1/2, and 1, where L/D represents length/diameter ratio. Axi-

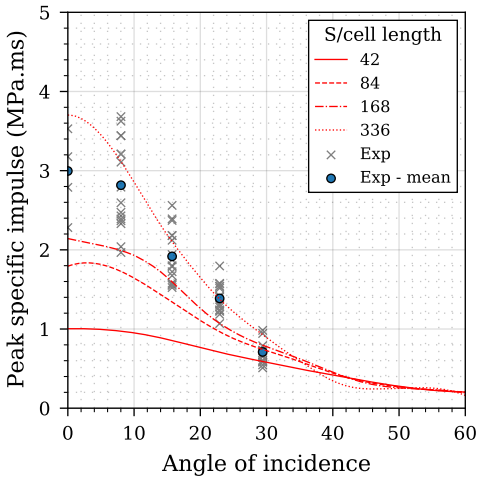

Figure 6.9: Mesh sensitivity analysis - comparison of different CFD models with experimental data.

Table 6.3: Example dataset information for cylindrical dataset

X_1	X_2	X_3	Y
Scaled distance $(m/kg^{1/3})$	Angle of incidence	L/D ratio	Peak specific impulse (MPa.ms)
0.1	15	1	45.87

symmetry was used, with a domain of 2 m × 2 m × 2 m, with 100 mm zone length and resolution level 3. *Apollo*'s auto-staging procedure was used and otherwise, the model set up was identical to those previously.

In summary, there are 20 CFD models (representing the 5 different stand-off distances analysed for each of the 4 different L/D ratios) with 150 values of peak specific impulse recorded for each, resulting in a dataset of 3000 samples. Each input in the dataset is shown in Table 6.3, with an example entry. There are 3 input features scaled distance (X_1), angle of incidence (X_2) and L/D ratio (X_3), the labelled values Y are peak specific impulse.

The variables X_1 and X_2 are minmax scaled across the entire dataset

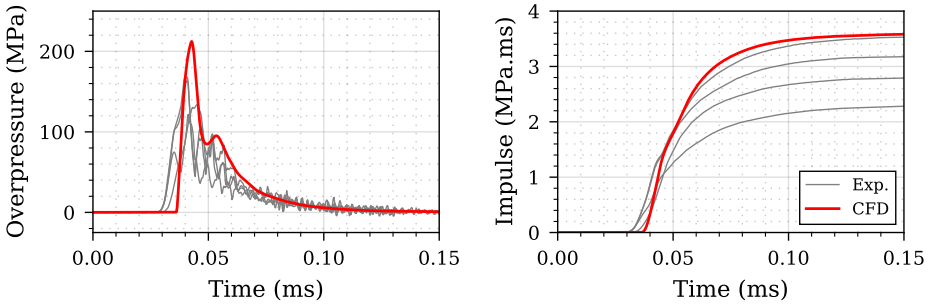

(a) 0 mm from centre, $\theta = 0°$

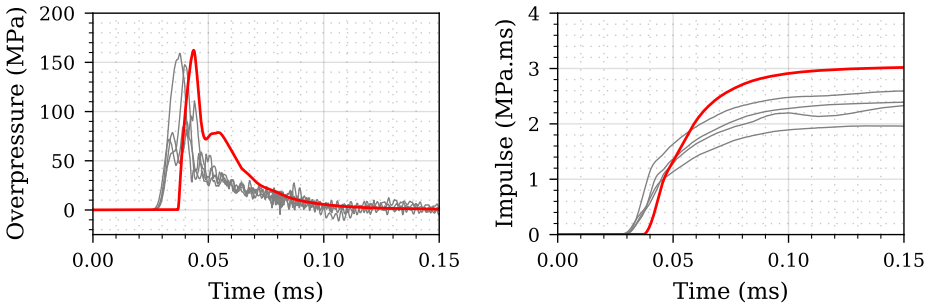

(b) 25 mm from centre, $\theta = 8°$

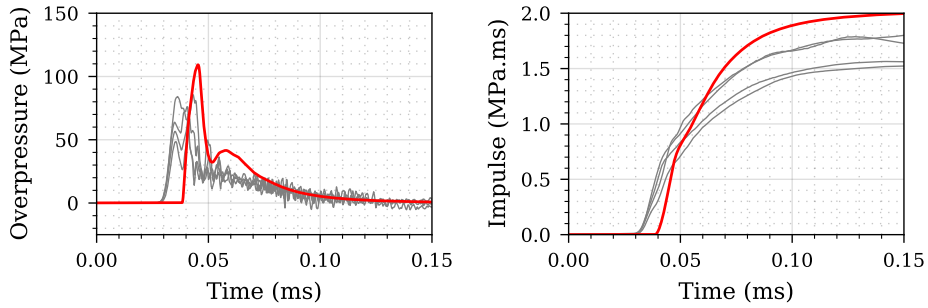

(c) 50 mm from centre, $\theta = 16°$

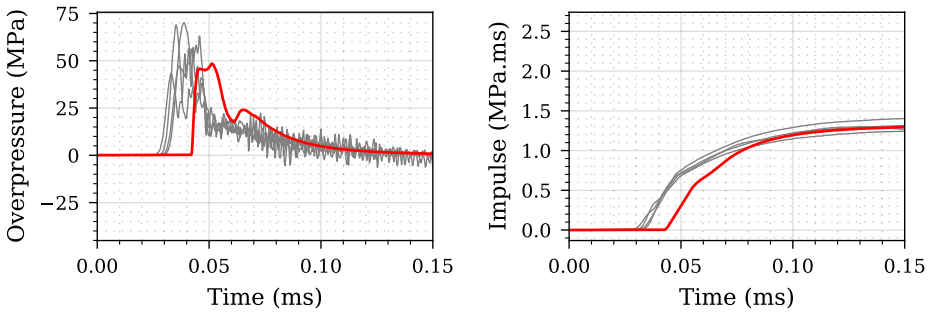

(d) 75 mm from centre, $\theta = 23°$

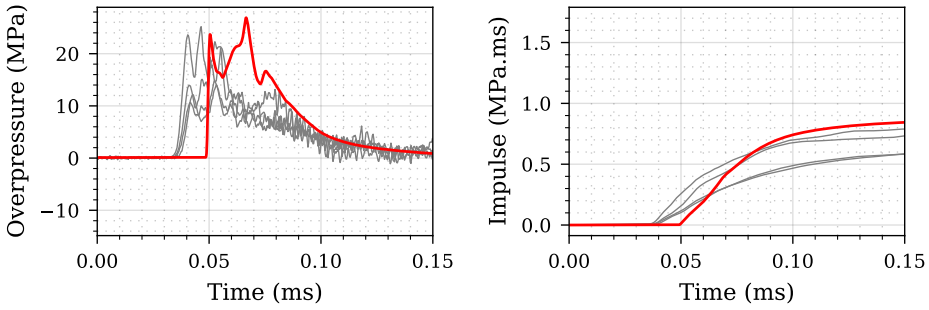

(e) 100 mm from centre, $\theta = 30°$

Figure 6.10: Experimental validation of numerical overpressure and specific impulse histories for $Z = 0.415$ m/k$g^{1/3}$ as 0, 25, 50, 75 and 100 mm distance from the target centre

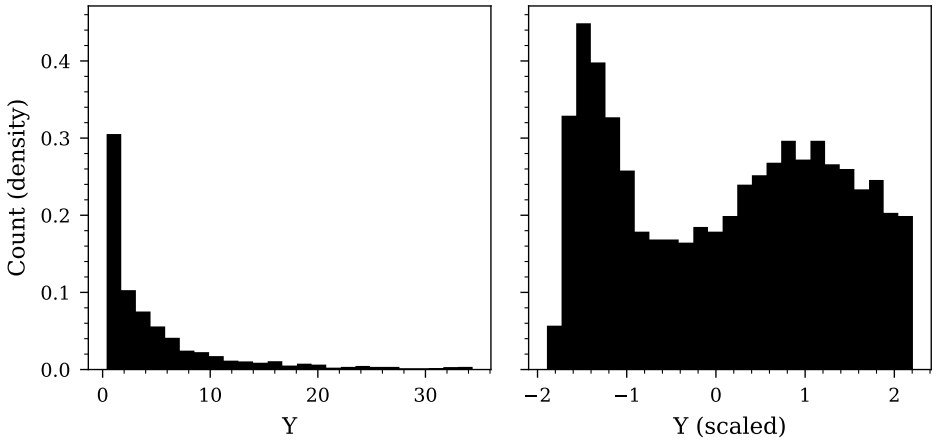

Figure 6.11: Unscaled Y dataset (left) and the resulting power transform (right)

with a feature range of [0,1] (using the scalars from the spherical model in Chapter 5), whilst X_3 is left unchanged, and varies between 0.2-1. The vector of labels, Y, has a log-normal distribution and is scaled via a power transform using the method described in Yeo & Johnson (2000) and again uses the same scaler that scaled the spherical dataset used in Chapter 5. The result of this data transformation in presented in Figure 6.11, and the transformation is applied prior to model training to allow for the knowledge transfer.

Previously in Section 6.3 it was demonstrated how prior knowledge from PE4 spheres can be leveraged to improve model development in a new domain. The specific method of applying a charge shape effect component was suitable as this applied across different charge compositions, and could be captured in a single exponential equations (as shown in Chapter 4). Figure 6.12 demonstrates why this would not be a suitable approach, as the normalised impulse curves for cylinders are drastically different to a sphere. Furthermore, the normalised profile for each L/D ratio is not always monotonically decreasing.

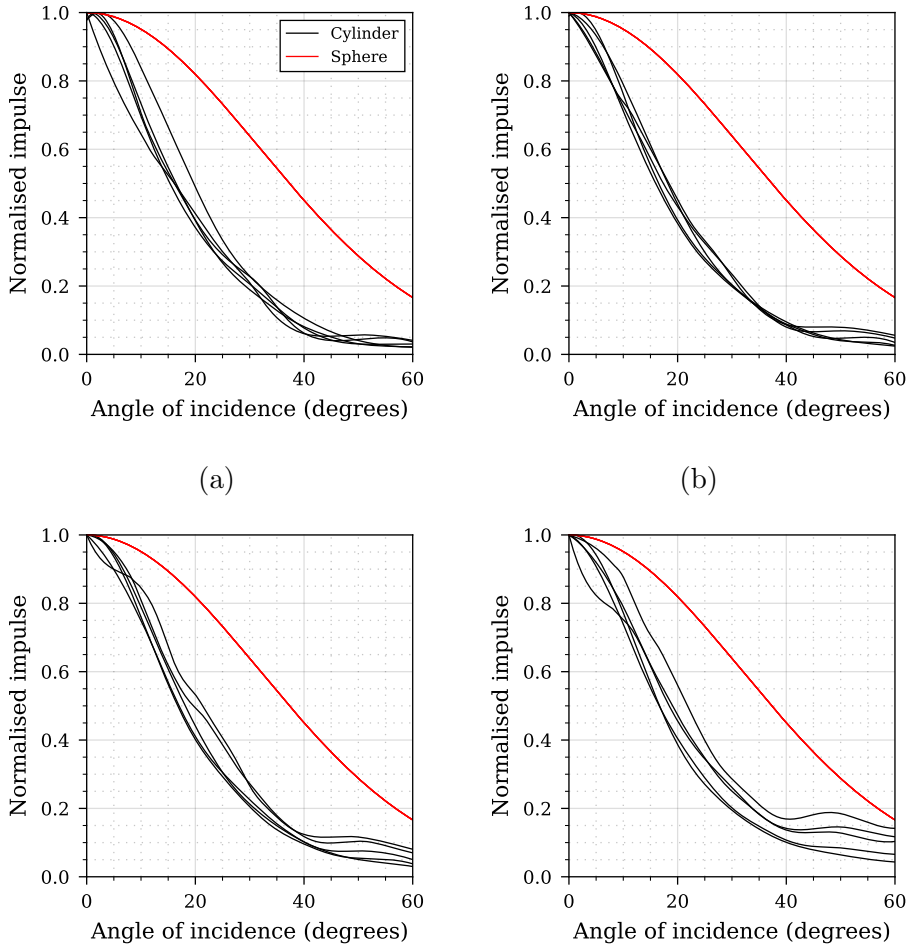

Figure 6.12: Normalised peak specific impulse comparison for four different L/D ratios: (a) 1/5, (b) 1/3, (c) 1/2 and (d) 1. In each case there are five normalised impulse curves corresponding to each scaled distance sample modelled.

6.4.3 Network architecture study

To model charge shape effects with transfer learning, two alternative network architectures were compared using the Keras package with Tensorflow backend (Chollet et al. 2015). The first model (NN) is shown in Figure 6.13a and does not utilise any transfer learning. It is only trained on the cylindrical dataset and provides a benchmark to compare a transfer network (TNN) to. It consists of 3 input nodes, 1 hidden layer and 1 output layer and is fully connected. The number of epochs was set at 1000, with early stopping, where the patience value was set as 50 epochs, to prevent over-fitting.

The TNN structure is shown schematically in Figure 6.13b, it consists of a pre-trained spherical model that was trained on the spherical dataset and is the network produced in Chapter 5, and a "bolt-on" network that handles the additional feature, X_3 (L/D). The output from both the spherical model and the bolt-on model are summed to provide an overall model output for the TNN. During model training, the spherical model is "frozen" so that the parameters are not updated during back-propagation, and the only parameters that are updated are in the bolt-on network. After the initial 1000 epochs training, the TNN is fine-tuned by un-freezing the spherical model, reducing the learning rate by an order of magnitude and training for 100 further epochs (with early-stopping implemented again).

In both cases the activation functions for the hidden units were set as hyperbolic tangent, with layer weights initialised with the Glorot normal initialiser (Glorot & Bengio 2010). The "Adam" algorithm was chosen as the optimiser, a stochastic gradient descent method that is based on adaptive estimation of first-order and second-order moments. K-fold cross-validation was implemented with 5 splits, following an initial data

split of 25% data randomly removed. The batch size was set at 32.

A varying number of hidden units are examined ranging from 1 to 10 in increments of 1. The hidden units apply to the hidden layer in Figure 6.13a and the hidden layer in the bolt-on network in Figure 6.13b. The results of these analyses are presented in Figure 6.14 where three separate sub-figures provide different metrics: mean absolute error (Figure 6.14a), mean squared error (Figure 6.14b) and coefficient of determination (Figure 6.14c). For all analyses, the metrics are evaluated for the three separate data portions: train, validation and test data. If any large discrepancies occur between data types, this can be indicative of overfitting issues. For the NN the global minimum mean MAE and MSE occur with 6 units in the hidden layer, suggesting that this capacity provides adequate predictive capability whilst in the TNN, a sufficient capacity is provided by 6 hidden units in the bolt-on sub-network. These two model architectures were taken forward for modelling.

6.4.4 Stress-testing

As previously established, obtaining data is expensive within a blast engineering context (and commonly other domains). A useful assessment for the utility of transfer learning would be how the models cope when data is increasingly sparse, any model or modelling framework that would improve the performance in a sparse data environment would be highly beneficial. To test this, the cylindrical dataset has been restricted in three separate levels of random data removal: a low threshold representing 20% data removal, a medium threshold representing 55% data removal, and a high threshold representing 90% data removal. The effects of this data restriction on the dataset is shown in Figure 6.16.

The modelling procedure follows that set out in Section 6.4.3, with

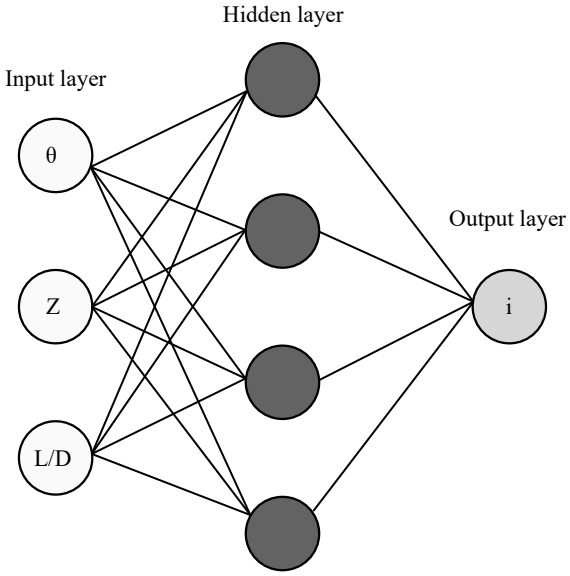

(a)

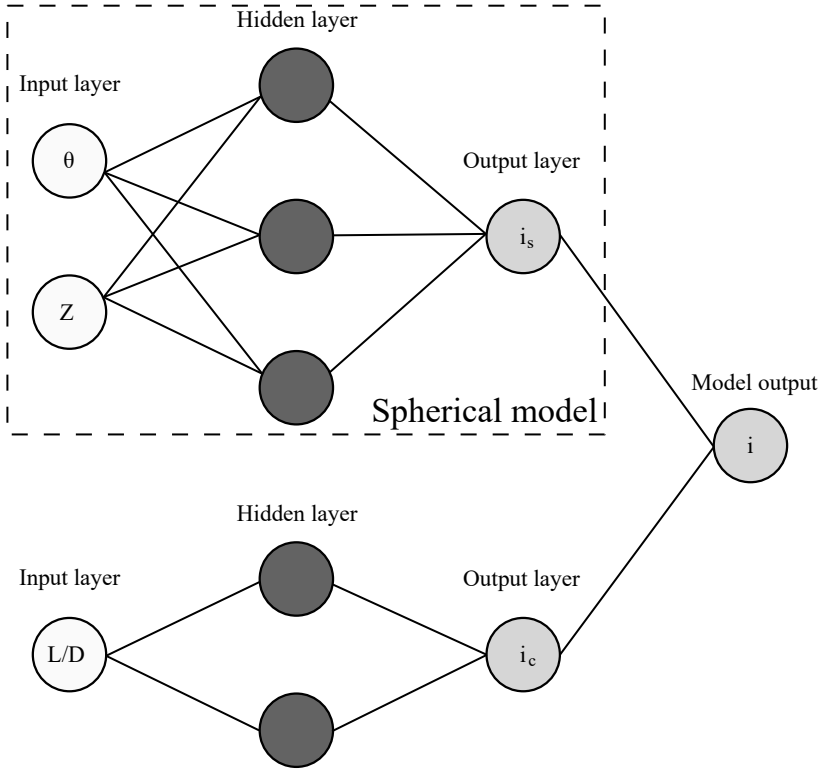

(b)

Figure 6.13: Model architectures for (a) NN architecture, model trained with no transfer learning and (b) Transfer neural network (TNN). In the TNN the previously trained "Spherical model" is used and an additional "bolt-on" network is added to handle the additional L/D input. The output of the spherical model (i_s) and bolt-on network (i_c) are summed to produce the overall model output.

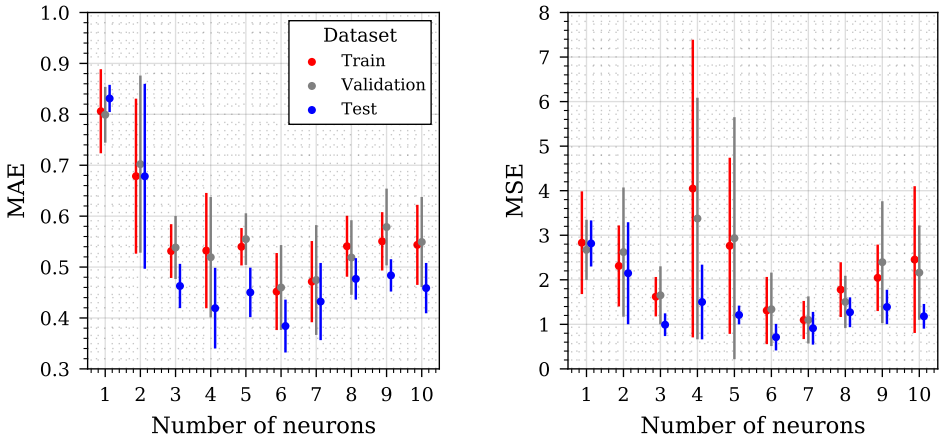

(a) Mean absolute error, (b) mean squared error,

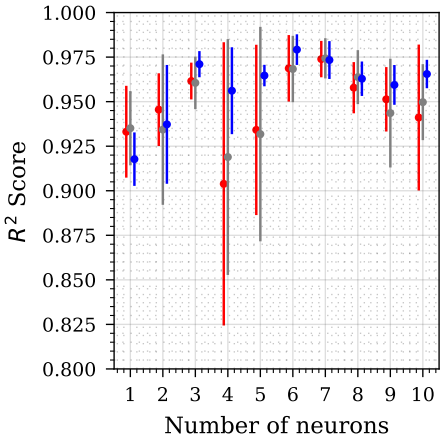

(c) coefficient of determination.

Figure 6.14: Hyper-parameter configuration of NN with various performance metrics shown. Error bars are standard deviation.

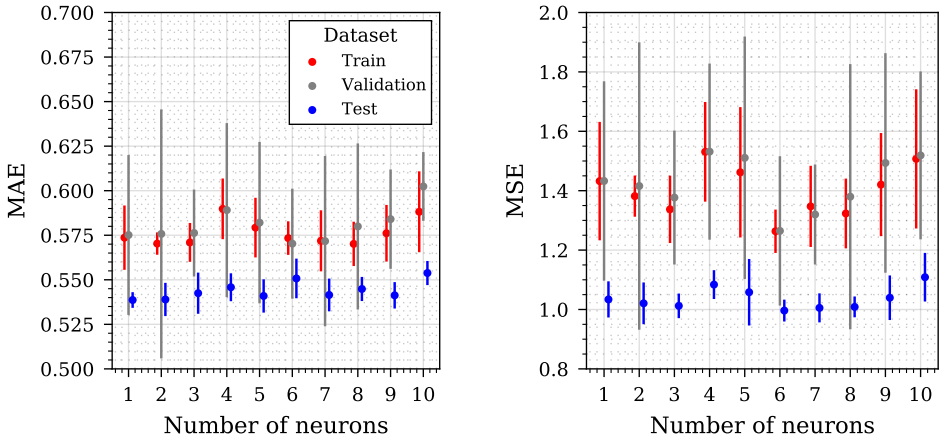

(a) Mean absolute error, (b) mean squared error,

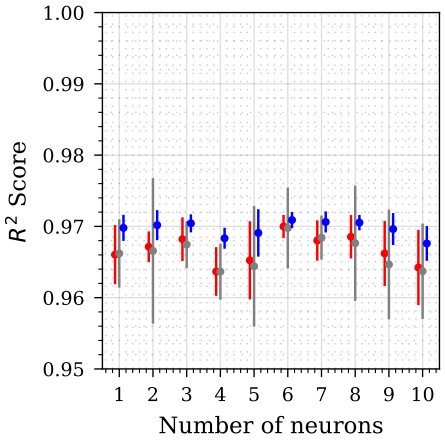

(c) coefficient of determination.

Figure 6.15: Hyper-parameter configuration of TNN with various performance metrics shown. Error bars are standard deviation.

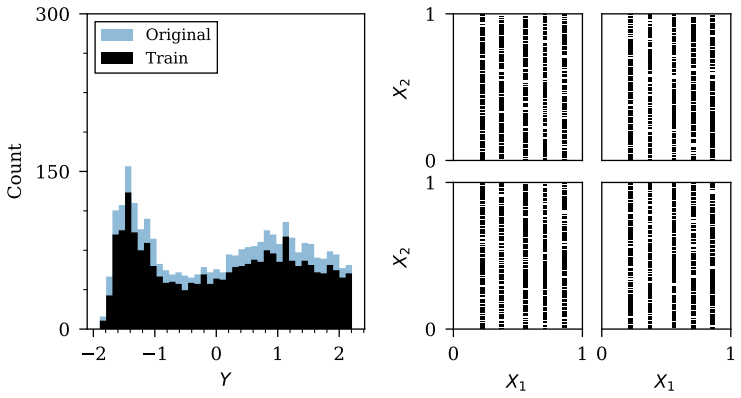

(a)

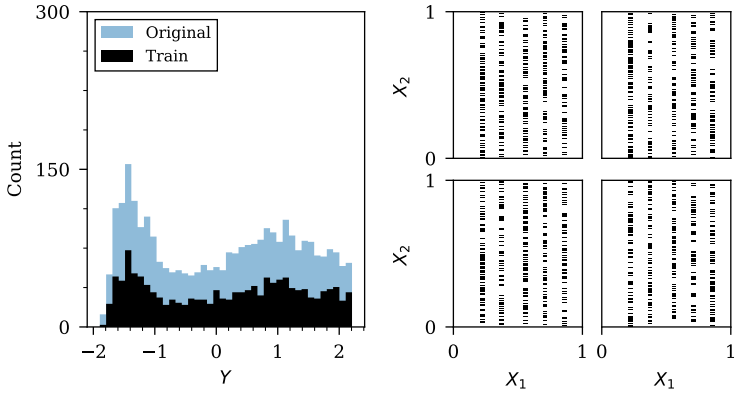

(b)

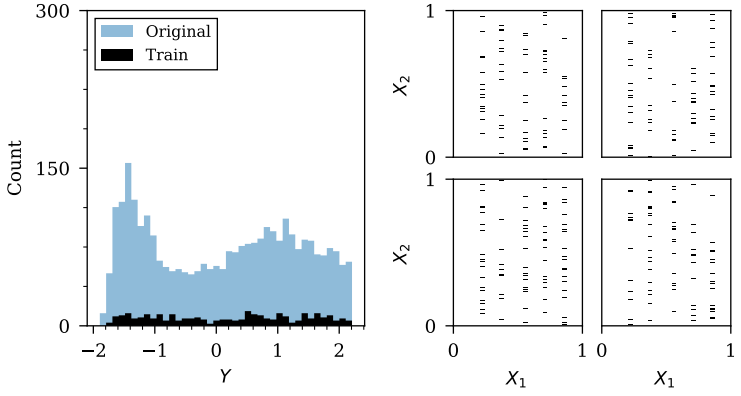

(c)

Figure 6.16: Distributions for (a) 20%, (b) 55% and (c) 90% random data removal. Each of the four plots in the right hand side represent a one of the four L/D ratios. The features are scaled using the fitted scalers from the dataset of spherical data in Chapter 5.

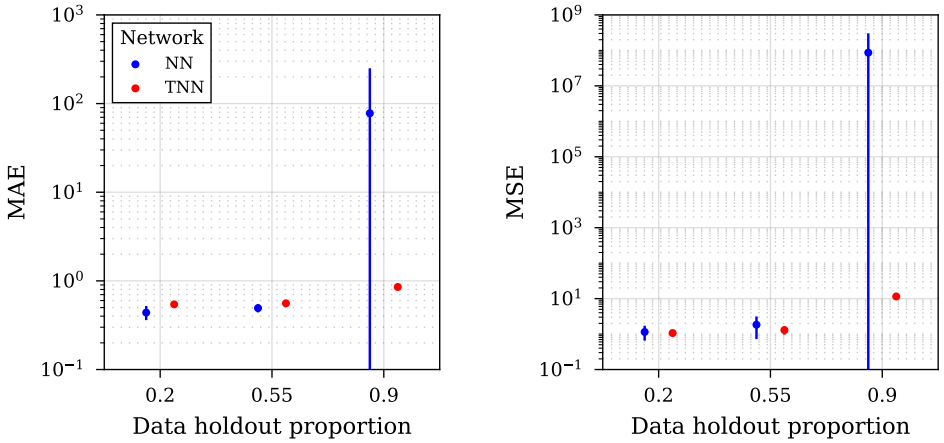

(a) Mean absolute error, (b) mean squared error,

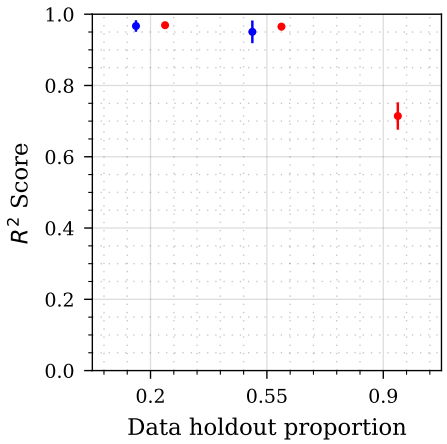

(c) Coefficient of determination.

Figure 6.17: Stress-test results from three data holdout proportions.

the exception that the K-fold cross-validation procedure is repeated 3 times for 5 splits. The results of these analyses are shown in Figure 6.17, with sub-figures for each of the three assessment metrics: mean absolute error (Figure 6.17a), mean squared error (Figure 6.17b) and coefficient of determination (Figure 6.17c). The points plotted are mean values, whilst the error bars represent standard deviation.

To better understand how each model is learning, a training history from the high threshold removal (90% data removal) case is presented in

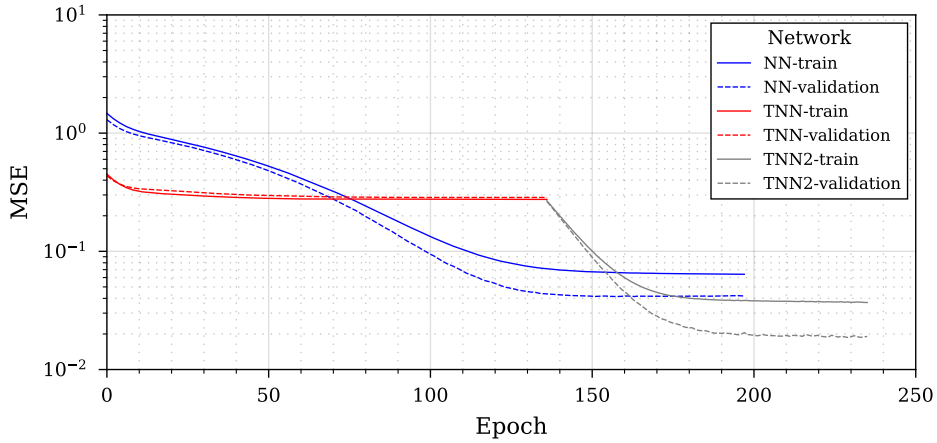

Figure 6.18: Training history for 90% data holdout of the two different networks (NN and TNN). "TNN-2" represents the fine-tuning of the TNN.

Figure 6.18 comparing the two different models NN and TNN. The TNN has been included as two separate parts in the legend, the initial training when the spherical model is frozen, followed by the "fine-tuning", where the entire model can be updated.

For the critical case of 90% data removed, a closer inspection of the models was completed and an overview of the transfer learning process shown in Figures 6.19, 6.20, 6.21 and 6.22, representing each of the four L/D ratios. In each figure, the training and unseen data is presented, alongside the original PE4 spherical dataset that is used for the transfer learning in the TNN. The NN and TNN are trained on this information and then are subsequently used to predict the entire dataset, with the predictions also presented. A closer inspection of these predictions is given by presenting the relationship between predicted and true values. Finally, the true dataset is presented in its entirety as a means of comparison. It can be seen in these Figures what information is available to each model, how it influences the accuracy of model predictions and provides a succinct overview of the stress-testing process.

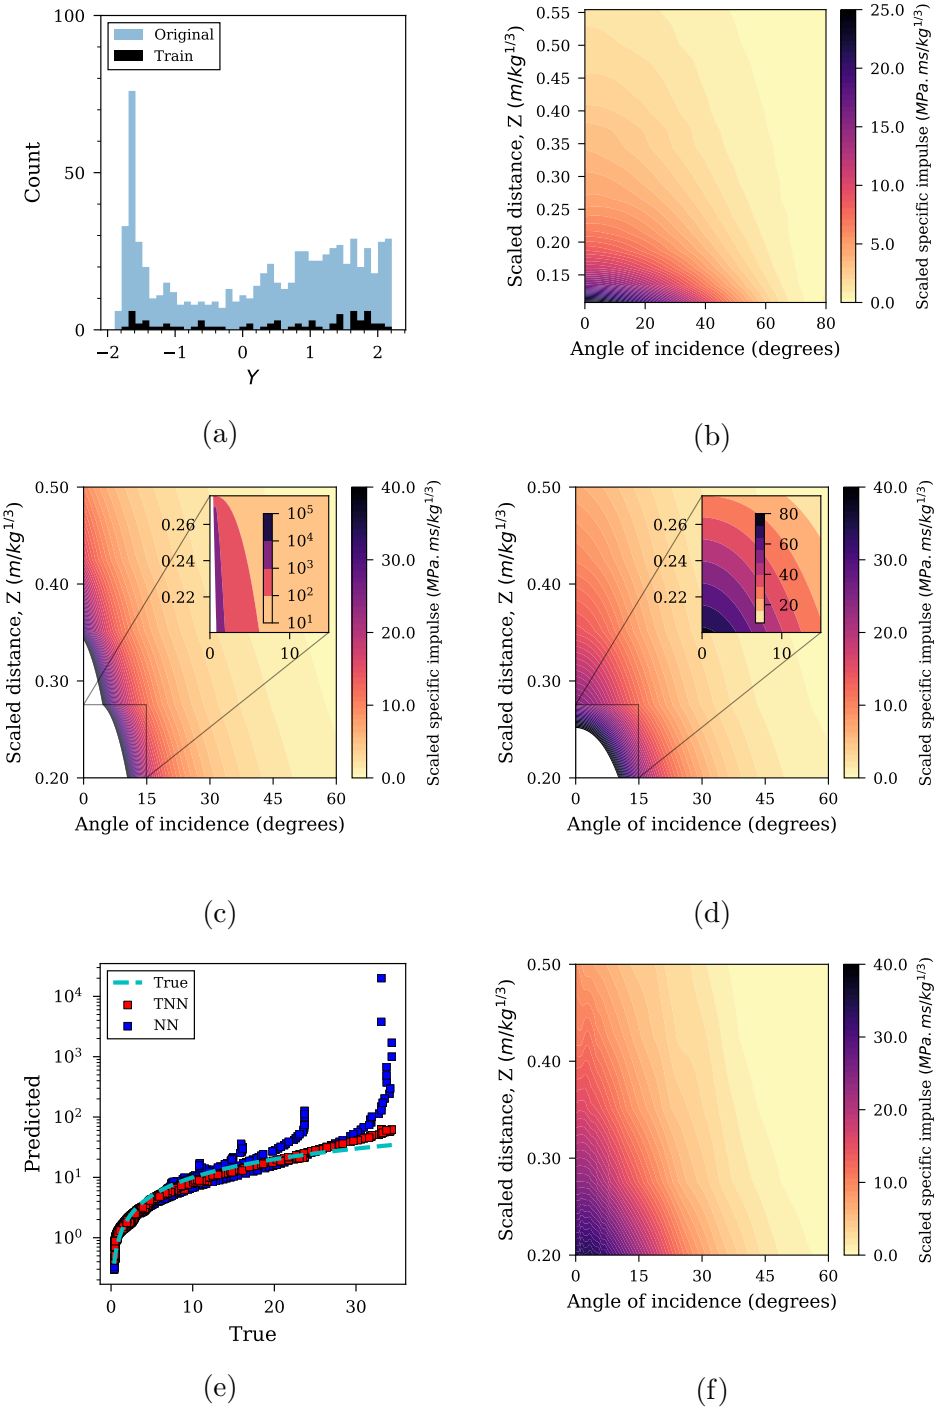

Figure 6.19: Stress-testing of 100g PE4 cylinder, L/D =1/5 with 90% of data removed. (a) histogram of original and training data, (b) transferred spherical PE4 dataset from Chapter 4, (c) NN predicted surface, (d) TNN predicted surface, (e) predicted vs true unseen data and (f) true CFD dataset.

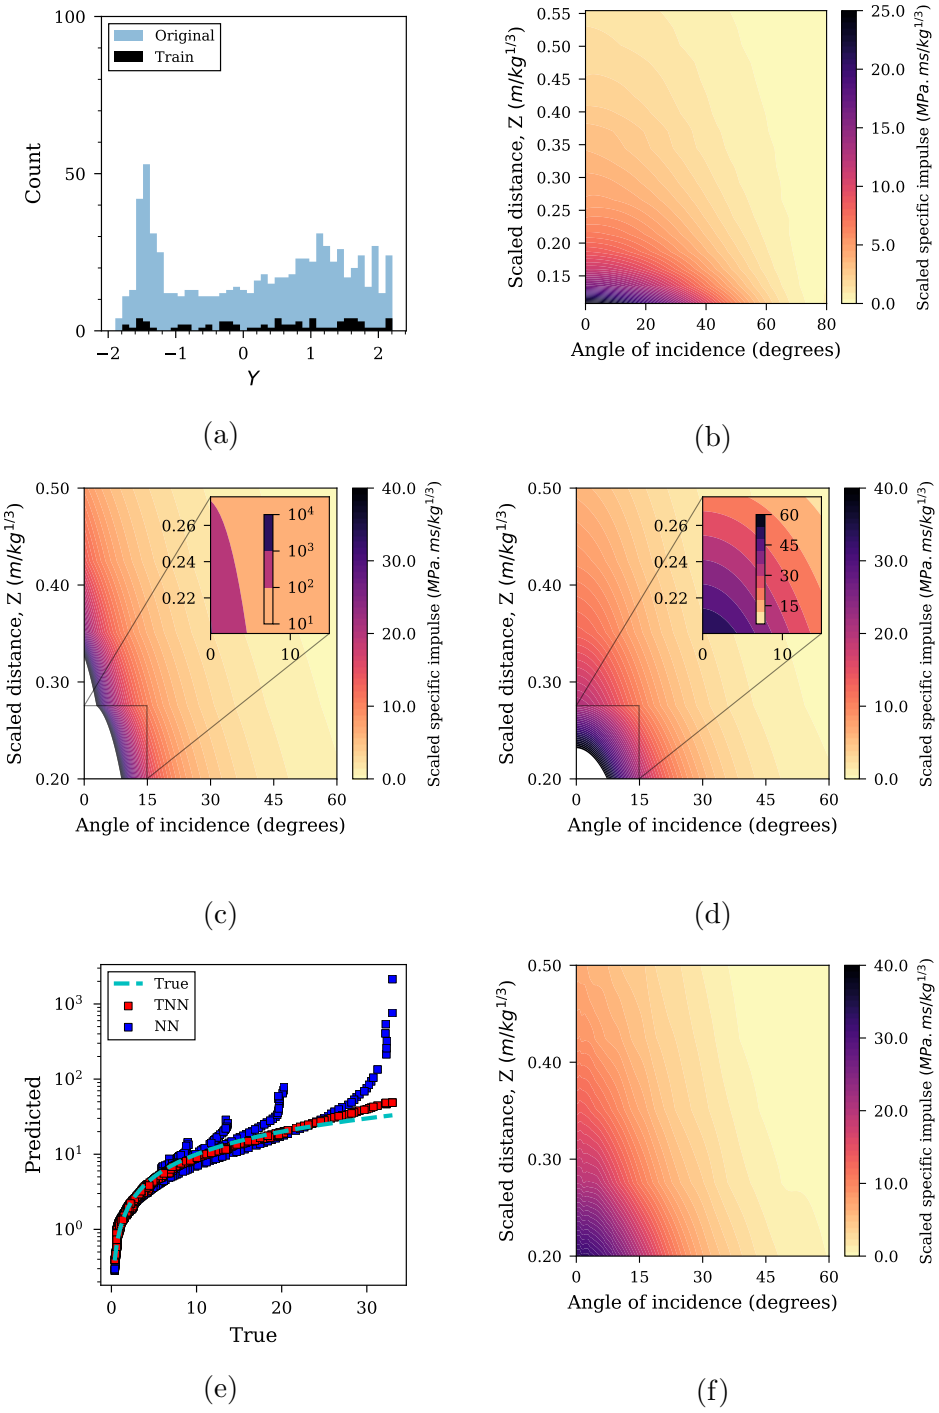

Figure 6.20: Stress-testing of 100g PE4 cylinder, L/D =1/3 with 90% of data removed. (a) histogram of original and training data, (b) transferred spherical PE4 dataset from Chapter 4, (c) NN predicted surface, (d) TNN predicted surface, (e) predicted vs true unseen data and (f) true CFD dataset.

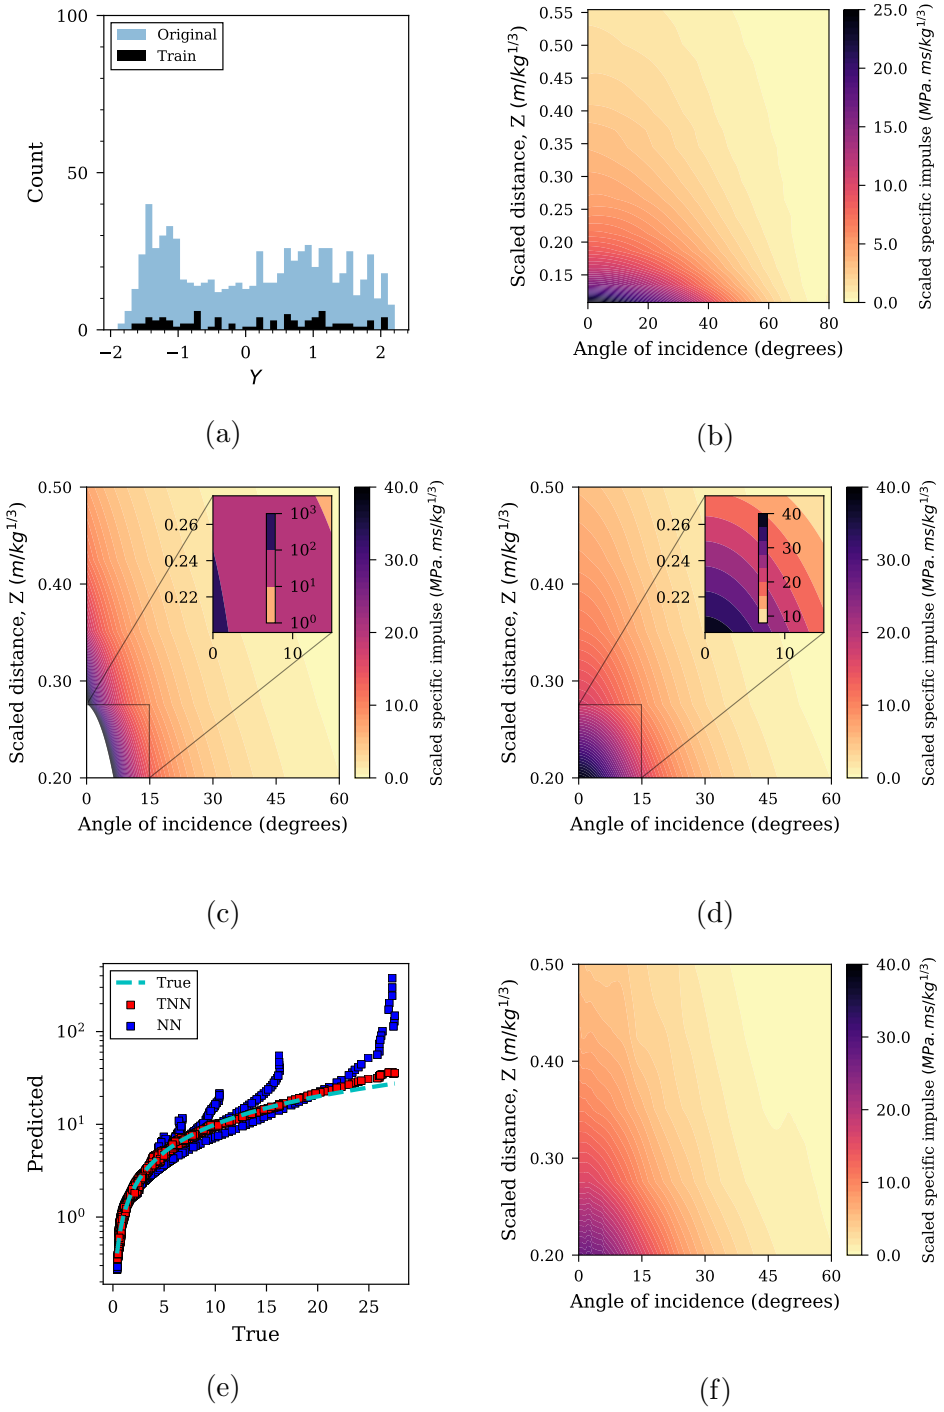

Figure 6.21: Stress-testing of 100g PE4 cylinder, L/D =1/2 with 90% of data removed. (a) histogram of original and training data, (b) transferred spherical PE4 dataset from Chapter 4, (c) NN predicted surface, (d) TNN predicted surface, (e) predicted vs true unseen data and (f) true CFD dataset.

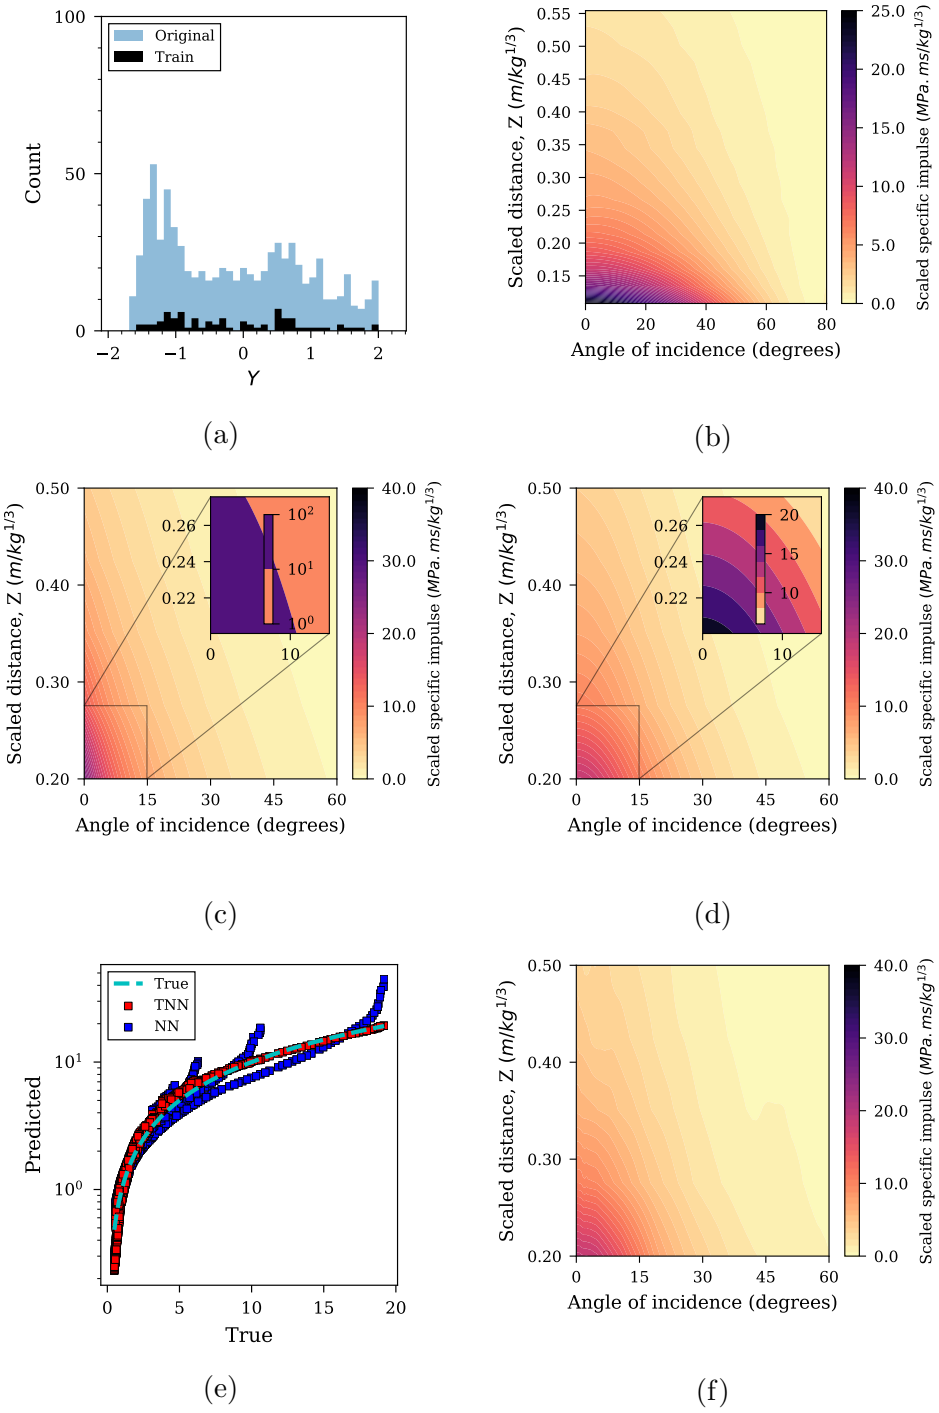

Figure 6.22: Stress-testing of 100g PE4 cylinder, L/D =1 with 90% of data removed. (a) histogram of original and training data, (b) transferred spherical PE4 dataset from Chapter 4, (c) NN predicted surface, (d) TNN predicted surface, (e) predicted vs true unseen data and (f) true CFD dataset.

6.5 Discussion

The results from the stress testing evaluation for the NN and TNN are compared in Table 6.4. To check the statistical significance of the results from each test, two-tailed Kolmogorov-Smirnov (KS) tests (Hodges 1958) have been performance in each case. This is a two-tailed test to test the hypothesis that both independent samples are drawn from the same continuous distribution. In this analysis, it establishes if any performance premium is statistically significant.

For every metric and for each data holdout proportion it can be seen that the TNN shows a performance premium over the NN, and this performance premium widens as the data holdout proportion increases. This performance premium is shown to be statistically significant for data holdout values of 20% and 90%. As the performance premium widens as the proportion of data removed increases, it suggests that the transfer learning has more utility as data becomes scarce. It is also shown that the TNN shows drastically less variability than the NN in all cases as shown by the considerably smaller standard deviation values.

A further insight into this performance premium can be seen in the training history from Figure 6.18. Initially it would appear that the NN

Table 6.4: Mean RMSE results from each stress-testing evaluation of NN and TNN models, with standard deviation given in brackets. RMSE values are from the unseen, test data. Entries in bold indicate a statistically significant difference ($p < 0.10$) from the Kolmogorov-Smirnov two-tailed test statistic (Hodges 1958) for performance premium (two-tailed p-value).

Data holdout	**Mean RMSE ± (s.d)**	
	NN	**TNN**
20%	1.047 (± 0.230)	**1.033 (± 0.037)**
55%	1.314 (± 0.335)	1.133 (± 0.106)
90%	3841.733 (± 8458.127)	**3.384 (± 0.212)**

is learning well, and there is a clear gap between the NN and TNN, until the TNN enters the fine-tuning stage. After the fine-tuning stage the TNN shows a clear performance premium over the NN and would appear to be a crucial element in the transfer learning implementation.

The stress-testing overview in Figures 6.19, 6.20, 6.21 and 6.22 further demonstrate the effectiveness of transfer learning. As shown, when predicting values at the minimum values of angle of incidence and scaled distance, the NN often over-predicts, quite drastically in some instances by up to 10000% from the true value, this can be seen particularly in the Subfigure (e) in all cases, where the TNN remains closer to the "true" line plotted and the NN lifts away from this when predicting the maximum values. It suggests that the knowledge gained from the PE4 dataset is considerably useful in preventing such drastic over-predictions, even though there is considerable difference in charge shape.

These results are highly promising, particularly from an engineering perspective. It has been established that knowledge of the source domain ($\mathcal{D}_S$) and task ($\mathcal{T}_S$) can be used to provide more accurate predictions and improve learning in the target domain ($\mathcal{D}_T$) and task ($\mathcal{T}_T$). In this instance the source domain is the spherical dataset and the target domain is the cylindrical dataset. The important practical implication of this finding can improve the efficiency of experimental design and improve the accuracy of predictive models in a blast engineering setting.

6.6 Summary

This chapter presents two novel applications of transfer learning for the prediction of peak specific impulse in a blast engineering setting. The first implementation used the predictive surrogate model proposed in

Chapter 4 where it is shown to speed up the learning of new predictive models for similar charge shapes, and scaled distance ranges but different charge compositions. Specifically, the part of the model in Chapter 4 responsible for the normalised impulse with respect to angle of incidence was transferred to newly studied charge compositions.

The second implementation aimed to investigate if knowledge obtained when modelling spherical explosives could be used to improve the learning when modelling cylinders. An initial architecture study was completed for two separate network architectures to determine a model that had a sufficient capacity to model the cylindrical charges. The first model (NN) did not implement any transfer learning and was included as a benchmark for comparison, this network did not have knowledge of the spherical dataset. The second network (TNN) did implement transfer learning through incorporating the trained spherical model proposed in Chapter 5, with an additional "bolt-on" network to handle the new L/D parameter. The models were stress-tested for three levels of random data removal, where it is shown the TNN outperforms the NN for every level, with this out-performance increasing as the percentage of data removed increases and showing statistical significant results for the low and high threshold. The TNN also shows drastically less variability in each case shown by the far smaller standard deviation values.

In a domain where data is expensive to obtain, a method is proposed here that improves the utility of data already obtained and demonstrates how this can be used when modelling a new, but related, domain. The implications of this research can directly affect how experiments are designed and will facilitate more accurate probabilistic-based approaches to experimental design and risk mitigation that encompass a more complex suite of scenarios than is capable presently.

Chapter 7

Summary and conclusions

7.1 Summary

This thesis has aimed to investigate data-driven modelling approaches for surrogate model development, specifically to create a "loading characterisation model", that can produce peak specific impulse distributions given parameters about the explosive and its configuration.

The detonation of a high explosive results in the rapid release of energy as the explosive charge undergoes a rapid change in state and is converted into a high pressure, high temperature gas. As the gas expands, the surrounding air is displaced, resulting in a high pressure shock discontinuity (shock wave). As this shock wave propagates away from the charge, it can cause severe damage to any structure that it impacts on. Blast protection engineers are tasked with designing infrastructure in a way that it is robust enough to withstand extreme loading, whilst dealing with several constraints such as time, cost and space. Due to the variability in initiation conditions (charge shape, position etc.), and the subsequent variability in loading produced, it becomes impractical to perform numerical simulations or experiments for all possible scenarios,

though an understanding of the loading is required to accurately model structural response. To this end, this thesis presents surrogate modelling strategies that provide the crucial loading information an engineer will require to design safe structures.

In Chapter 2, the current literature in the field of surrogate modelling in a blast context is discussed, as well as the theoretical background of shock wave formation and the detonation process. It also provides a general overview into machine learning as a field of study, and how its application is used in a blast engineering context. It therefore fully meets objective 1 as defined on page 4:
"*1. To review the current literature on predictive approaches for blast loading and discuss the limitations of existing approaches.*"

In Chapter 3 this thesis provides a theoretical background of computational fluid dynamics (CFD), discussing the numerical scheme used in the chosen software (*Apollo*) and additional modules that make *Apollo* a suitable choice. Further, *Apollo* is then validated in the far-field using experimental data obtained at the University of Sheffield.

Chapter 4 demonstrates preliminary approaches into surrogate modelling of peak specific impulse. First it provides a validation of *Apollo* in the near-field against experimental data, and then outlines a surrogate modelling framework – from data generation to surrogate model production. An important insight gained within Chapter 4 is the discovery of a suitable data transformation that can reduce a dimension of the dataset, by modelling the "normalised peak specific impulse", therefore improving surrogate modelling efficiency. A surrogate model is proposed and is extensively evaluated and compared to other models in the literature, where it shown to be more accurate and considerably easier to evaluate. Chapters 3 and 4 both fully address objectives 2, 3 and 4:

"*2. To establish a reliable numerical framework for generating loading distributions that has been validated against near-field and far-field experimental data.*

3. To rigorously assess the sensitivities of numerical approaches for mesh effects and develop consistent rules for mesh sizing to be used throughout the thesis study.

4. To investigate data transformations that can be applied to datasets as a whole that reduce the dimensions of the dataset and allow the rapid development of surrogate models."

In Chapter 5, this thesis addresses the question of *how we can incorporate prior domain knowledge into the model training process to speed up performance or efficiency of model training?*. To do this, insights gained from the dataset in Chapter 4 are used and directly implemented as an additional regularisation term into the objective function that is minimised during model training, to penalise "physically invalid" predictions. This approach is compared to models without additional regularisation and both models are stress-tested by withholding various portions of the dataset which allows the specific evaluation of the generalisation ability of both models. The results demonstrate statistically significant benefits to the novel regularisation procedure in both model performance by an accuracy metric (RMSE), and model performance with respect to physical inconsistency and can be considered a suitable approach in surrogate modelling within this domain. This chapter fully meets objective 5:

"*5. To investigate physics-guided machine learning approaches in a blast loading domain through incorporating known, or learned, physics directly into the objective function as a physics-based regularisation procedure.*"

Finally, in Chapter 6 this thesis further investigates transfer learning strategies to improve the utility of data already obtained, and demon-

strate how prior knowledge can be used to improve both model performance and learning efficiency in a new domain. Firstly, five additional datasets of spherical charges are developed of varying charge compositions. Surrogate models, of the format presented in Chapter 4 are developed for each new dataset, where crucially a component of the model (the angular distribution component) is transferred from the original model to the five new models. This application of transfer learning is shown to be suitable and demonstrates the utility of the process as considerably less new data is required to be obtained to create additional subsequent surrogate models for alternative charge compositions. Then in a second exercise of transfer learning, this thesis demonstrates how knowledge of a spherical dataset can be used to improve the efficiency of surrogate model development for a different charge shape entirely. Initially, four new datasets of cylindrical explosives are developed, of varying length/diameter ratios and modelled in two different network structures. The first is a classical neural network structure, and does not implement transfer learning. The second novel model is a network containing the trained spherical model in Chapter 5, with an additional "bolt-on" network to handle the new L/D parameter. Both models are stress-tested by withholding three different levels of cylindrical training data and the results of their predictive performance on this unseen data is evaluated. It is shown that the transfer learning implementation provides a considerable statistically significant performance benefit and drastically reduces the need to obtain new data. Chapter 6 fully addresses objective 6:

"*6. To further investigate transfer learning approaches from previously trained surrogate models applied to new datasets, to determine whether knowledge learned from previous tasks and domains can be applied in a new task and domain and improve the efficiency in training, and perfor-*

mance of, the new surrogate models. ”

The results of this thesis should be used to guide surrogate model development for the prediction of peak specific impulse in the near-field. Surrogate modelling frameworks for spheres and cylinders are presented and it is demonstrated how prior domain knowledge, and data, can be used to improve the performance of models in a new domain, and improve the efficiency when training models in a new domain. It is shown that machine learning methods can reliably be used in surrogate model development. The findings presented within this thesis have the potential to be implemented into load prediction software which would be of great utility to structural engineers working in the blast protection and resilience community.

7.2 Conclusions

The main conclusions of this thesis can be summarised in the following, listed in their order of appearance within this thesis:

- *Apollo* can be used to simulate blast events in the far-field region to a good level of agreement with data collected from a series of experimental trials by Rigby et al. (2015a).

- *Apollo* can be used to simulate near-field blast events to a good level of agreement with experimental data. Peak specific impulse requires an ultimate cell size of S/50, where S is stand-off to the nearest gauge.

- In a scaled distance range of 0.11–0.55 m/kg$^{1/3}$ for spherical charges, a data transformation can be applied to the specific impulse distributions by dividing through by the impulse value at the perpen-

dicular gauge. The result is that all distributions collapse into a "normalised specific impulse" profile.

- In a scaled distance range of 0.11–0.55 m/kg$^{1/3}$ for spherical charges, the "normalised specific impulse" profile can be well approximated by a modified form of a Gaussian function.

- In a scaled distance range of 0.11–0.55 m/kg$^{1/3}$ for spherical charges, the peak perpendicular impulse can be well approximated as a power law, with scaled distance as the independent variable.

- Current literature guidance for modelling angle of incidence effects in the near-field is assessed (Randers-Pehrson & Bannister 1997, Henrych 1979) where they are shown to be less accurate than the model proposed herein.

- Two extensively validated surrogate models are proposed to predict the peak specific impulse produced from the detonation of a spherical PE4 explosive between a scaled distance range 0.11–0.55 m/kg$^{1/3}$ in Equation 4.5 on page 101 and between a scaled distance range 0.11–0.21 m/kg$^{1/3}$ in Equation 4.4 on page 101.

- A fully connected neural network with four units in the hidden layer can be used to accurately predict peak specific impulse given scaled distance (0.11–0.55 m/kg$^{1/3}$) and angle of incidence (0-60^o) as the inputs and a sufficient level of training. Activation functions for the hidden units were set as hyperbolic tangent functions and the layer weights initialised with the Glorot normal initialiser.

- A monotonic loss constraint can be directly included in the objective function as a physics-based regularisation procedure when

training a neural network to produce statistically significant performance benefits when generalising for the prediction of peak specific impulse produced by a spherical charge of PE4 given scaled distance (0.11–0.55 m/kg$^{1/3}$) and angle of incidence (0-60^o).

- The component of the surrogate model that predicts normalised peak specific impulse given angle of incidence presented in Equations 4.5 and 4.4 can be transferred to other charge compositions but similar shapes and range of scaled distances. These are presented in Equations 6.2 on 165 for TNT, HMX, RDX, PETN and COMPB for a scaled distance range of 0.2–0.6 m/kg$^{1/3}$.

- Knowledge of spherical charges can be used to reduce the need to obtain new data when modelling cylindrical charges of L/D ratios 1/5, 1/3, 1/2, 1/1 for a scaled distance range of 0.2–0.6 m/kg$^{1/3}$. This knowledge transfer is established by adding an additional "bolt-on" sub-network of 6 hidden units in a hidden layer onto the neural network architecture presented for PE4 spheres in Chapter 5 (the "spherical model"). During model training, the "spherical model" is frozen and only the "bolt-on" model parameters are updated during back-propagation, however for the final 100 epochs, the learning rate is dropped and the entire model is un-frozen to allow fine tuning. This implementation of transfer learning is shown to drastically improve model performance (compared to a model that did not implement transfer learning), with this outperformance increasing as the available training data becomes increasingly limited.

7.3 Evaluation and future work

Further work within this thesis falls within two broad categories: the first, improving the existing loading characterisation models; whilst the second is concerned with the future structural response model. These two categories correspond to the framework suggested in Figure 1.1 on Page 6.

To improve the loading characterisation model, suggested avenues of work include an investigation into sampling methods for surrogate modelling. For example, if the modeller wanted to create a peak specific impulse surrogate model for a certain charge composition, and shape, in the near-field domain, how many CFD simulations would they need to run? To achieve this, a study could be completed of various modelling strategies and the results compared. The work presented herein would suggest a modelling strategy based on a complexity metric could be promising, as the region close in to the explosive is much more complex than the region further out, which would suggest more samples are required from this complex region.

The introduction of additional regularisation terms into the objective function is a fairly underdeveloped area of the literature. Within this thesis an application of a monotonic loss constraint is demonstrated, but there is no reason why this could not be a more complex term. For example, in Chapter 4 it is demonstrated that the peak specific impulse varies with scaled distance as a power law, so it would be reasonable to assume if the complexity of the regularisation term was extended from a monotonic loss to a power law the results would be more accurate, particularly when generalising.

Finally, the second part of the probabilistic framework remains untouched within this thesis, the connection between loading and structural

response. However, significant developments have been made herein that encourage the use of this framework. The conclusions presented within this thesis should be used as positive evidence towards the theme of surrogate modelling in a blast engineering context, and justifies the use of machine learning models and transfer learning strategies to improve model predictive performance and training efficiency. The results of this thesis have direct impact guiding experimental design (due to an awareness of the amount of data required) which will correspond to significant savings in time and cost.

Bibliography

Acharya, D., Rani, A., Agarwal, S. & Singh, V. (2016), 'Application of adaptive Savitzky-Golay filter for EEG signal processing', *Perspectives in Science* **8**, 677–679.

Agatonovic-Kustrin, S. & Beresford, R. (2000), 'Basic concepts of artificial neural network (ann) modeling and its application in pharmaceutical research', *Journal of pharmaceutical and biomedical analysis* **22**(5), 717–727.

Ahmadi, M. A. (2015), 'Developing a robust surrogate model of chemical flooding based on the artificial neural network for enhanced oil recovery implications', *Mathematical Problems in Engineering* **2015**.

Albawi, S., Mohammed, T. A. & Al-Zawi, S. (2017), Understanding of a convolutional neural network, *in* '2017 International Conference on Engineering and Technology (ICET)', Ieee, pp. 1–6.

Alterman, D., Stewart, M. G. & Netherton, M. D. (2019), 'Probabilistic assessment of airblast variability and fatality risk estimation for explosive blasts in confined building spaces', *International Journal of Protective Structures* **10**(3), 306–329.

Anderson, J. D. (2010), *Fundamentals of aerodynamics*, Tata McGraw-Hill Education.

Artero-Guerrero, J., Pernas-Sánchez, J. & Teixeira-Dias, F. (2017), 'Blast wave dynamics: The influence of the shape of the explosive', *Journal of hazardous materials* **331**, 189–199.

Aune, V., Fagerholt, E., Hauge, K. O., Langseth, M. & Børvik, T. (2016), 'Experimental study on the response of thin aluminium and steel plates subjected to airblast loading', *International Journal of Impact Engineering* **90**, 106–121.

Baker, W. E. W. E. (1973), *Explosions in Air*, University of Texas Press, Austin.

Batina, J. T. (1990), 'Unsteady euler airfoil solutions using unstructured dynamic meshes', *AIAA journal* **28**(8), 1381–1388.

Bellman, R. (1966), 'Dynamic programming', *Science* **153**(3731), 34–37.

Bewick, B., Flood, I. & Chen, Z. (2011), 'A neural-network model-based engineering tool for blast wall protection of structures', *International Journal of Protective Structures* **2**(2), 159–176.

Blazek, J. (2015), *Computational fluid dynamics : principles and applications*, third edition. edn, Elsevier, Boston, Massachusetts.

Blitzer, J., Crammer, K., Kulesza, A., Pereira, F. & Wortman, J. (2008), 'Learning bounds for domain adaptation'.

Blumer, A., Ehrenfeucht, A., Haussler, D. & Warmuth, M. K. (1989), 'Learnability and the vapnik-chervonenkis dimension', *Journal of the ACM (JACM)* **36**(4), 929–965.

Bogosian, D., Ferritto, J. & Shi, Y. (2002), Measuring uncertainty and conservatism in simplified blast models, *in* '30th Explosives Safety Seminar, Atlanta, GA, USA'.

Bogosian, D., Yokota, M. & Rigby, S. (2016), TNT equivalence of C-4 and PE4: a review of traditional sources and recent data, *in* '24th International Conference on Military Aspects of Blast and Shock (MABS24), Halifax, Nova Scotia, Canada'.

Bortorlan Neto, L., Saleh, M., Pickerd, V., Yiannakopoulos, G., Mathys, Z. & Reid, W. (2017), Rapid vulnerability assessment of naval structures subjected to localised blast, *in* 'Proceedings of Pacific IMC, The Royal Institution of Naval Architects, The Institute of Marine Engineering, Science and Technology and Engineers Australia, Sydney, Australia', p. 12.

Bortorlan Neto, L., Saleh, M., Pickerd, V., Yiannakopoulos, G., Mathys, Z. & Reid, W. (2020), 'Rapid mechanical evaluation of quadrangular steel plates subjected to localised blast loadings', *International Journal of Impact Engineering* **137**, 103461.

Boser, B. E., Guyon, I. M. & Vapnik, V. N. (1992), A training algorithm for optimal margin classifiers, *in* 'Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory', pp. 144–152.

Both, A.-L., Hisken, H., Rückmann, J.-J. & Steihaug, T. (2019), 'Surrogate-based model parameter optimization based on gas explosion experimental data', *Engineering Optimization* **51**(2), 301–316.

Branch, M. A., Coleman, T. F. & Li, Y. (1999), 'A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems', *SIAM Journal on Scientific Computing* **21**(1), 1–23.

Breiman, L. (1996), 'Bagging predictors machine learning 24 (2), 123-140 (1996) 10.1023', *A: 1018054314350* .

Brode, H. L. (1955), 'Numerical solutions of spherical blast waves', *Journal of Applied Physics* **26**(6), 766–775.

Campidelli, M., Tait, M. J., El-Dakhakhni, W. W. & Mekky, W. (2015), 'Inference of Blast Wavefront Parameter Uncertainty for Probabilistic Risk Assessment', *Journal of Structural Engineering* **141**(12), 04015062.

Cheval, K., Loiseau, O. & Vala, V. (2010), 'Laboratory scale tests for the assessment of solid explosive blast effects. part i: Free-field test campaign', *Journal of Loss Prevention in the Process Industries* **23**(5), 613–621.

Cheval, K., Loiseau, O. & Vala, V. (2012), 'Laboratory scale tests for the assessment of solid explosive blast effects. part ii: Reflected blast series of tests', *Journal of Loss Prevention in the Process Industries* **25**(3), 436–442.

Chollet, F. et al. (2015), 'Keras', `https://keras.io`.

Chung Kim Yuen, S., Nurick, G., Langdon, G. & Iyer, Y. (2016), 'Deformation of thin plates subjected to impulsive load: Part III – an update 25 years on', *International Journal of Impact Engineering* **107**, 108–117.

Chung Kim Yuen, S. & Nurick, G. N. (2005), 'Experimental and numerical studies on the response of quadrangular stiffened plates. Part I: subjected to uniform blast load', *International Journal of Impact Engineering* **31**(1), 55–83.

Clarke, S. D., Fay, S. D., Warren, J. A., Tyas, A., Rigby, S. E. & Elgy, I. (2015), 'A large scale experimental approach to the measure-

ment of spatially and temporally localised loading from the detonation of shallow-buried explosives', *Measurement Science and Technology* **26**, 015001.

Clutter, J. & Stahl, M. (2014), Effect of charge shape on damage from potential explosive attacks, *in* '6th international conference on protection of structures against hazards, Tianjin, China', pp. 16–17.

Cormie, D., Mays, G. & Smith, P. (2009), *Blast Effects on Buildings*, ICE Publishing, London.

Cormie, D., Wilkinson, W., Shin, J. & Whittaker, A. (2013), Scaled-distance relationships for close-in detonations, *in* '15th International Symposium on the Interaction of the Effects of Munitions with Structures (ISIEMS)'.

Cortes, C. & Vapnik, V. (1995), 'Support-vector networks', *Machine learning* **20**(3), 273–297.

Cranz, C. (1926), *Lehrbuch der Basllistik*, Springer, Berlin, Germany.

Cybenko, G. (1989), 'Approximation by superpositions of a sigmoidal function', *Mathematics of control, signals, and systems* **2**, 303–314.

Daw, A., Thomas, R. Q., Carey, C. C., Read, J. S., Appling, A. P. & Karpatne, A. (2020), Physics-guided architecture (pga) of neural networks for quantifying uncertainty in lake temperature modeling, *in* 'Proceedings of the 2020 siam international conference on data mining', SIAM, pp. 532–540.

Dennis, A. A., Pannell, J. J., Smyl, D. J. & Rigby, S. E. (2021), 'Prediction of blast loading in an internal environment using artificial neural

networks', *International Journal of Protective Structures* **12**(3), 287–314.

Dewey, J. M. (1964), 'The air velocity in blast waves from T.N.T. explosions', *Proceedings of The Royal Society of London Series A, Mathematical and Physical Sciences* **279**, 366–385.

DoD, U. (2008), Structures to resist the effects of accidental explosions, Technical Report UFC-3-340-02, US Department of Defence, Washington DC, U.S.A.

Donea, J., Giuliani, S. & Halleux, J.-P. (1982), 'An arbitrary lagrangian-eulerian finite element method for transient dynamic fluid-structure interactions', *Computer methods in applied mechanics and engineering* **33**(1-3), 689–723.

D'souza, R. N., Huang, P.-Y. & Yeh, F.-C. (2020), 'Structural analysis and optimization of convolutional neural networks with a small sample size', *Scientific reports* **10**(1), 1–13.

Edwards, D. H., Thomas, G. O., Milne, A., Hooper, G. & Tasker, D. (1992), 'Blast wave measurements close to explosive charges', *Shock Waves* **2**, 237–243.

Esparza, E. D. (1986), 'Blast measurements and equivalency for spherical charges at small scaled distances', *International Journal of Impact Engineering* **4**(1), 23–40.

Esparza, E. D. (1992), Spherical equivalency of cylindrical charges in free-air, Technical report, Southwest Research Institute.

Fan, Y., Chen, L., Li, Z., Xiang, H.-b. & Fang, Q. (2022), 'Modeling

the blast load induced by a close-in explosion considering cylindrical charge parameters', *Defence Technology* .

Fay, S., Clarke, S., Tyas, A., Warren, J., Rigby, S., Bennett, T., Elgy, I. & Gant, M. (2014), Measuring the spatial and temporal pressure variation from buried charges, *in* 'Proceedings of the 23rd International Symposium on Military Aspects of Blast and Shock', Sheffield.

FCA (2021), 'Fca handbook'.

Flood, I., Bewick, B. T., Dinan, R. J. & Salim, H. A. (2009), 'Modeling blast wave propagation using artificial neural network methods', *Advanced Engineering Informatics* **23**(4), 418–423.

Flood, I., Bewick, B. T. & Rauch, E. (2012), 'Rapid Simulation of Blast Wave Propagation in Built Environments Using Coarse-Grain Simulation', *International Journal of Protective Structures* **3**(4), 431–448.

Flood, I., Issa, R. R. A. & Abi-Shdid, C. (2004), 'Simulating the Thermal Behavior of Buildings Using Artificial Neural Networks-Based Coarse-Grain Modeling', *Journal of Computing in Civil Engineering* **18**(3), 207–214.

Frank, S., Löhner, R., Camelli, F., Frank, R. & Amini, A. (2008), Comparison of a fast running engineering model and a coarse 3D mesh 3-D Euler predictions for airblast in an urban environment, *in* '20th Proceedings Of', Oslo, p. 12.

Fraunhofer EMI (2018), *APOLLO Blastsimulator manual, version: 2018.2*, Fraunhofer Institute for High-Speed Dynamics, Ernst-Mach-Institut, Freiburg, Germany.

Friedman, J. (1999), Stochastic gradient boosting, Technical report, Stanford University, San Francisco, CA.

Friedman, J. H. (2001), 'Greedy function approximation: a gradient boosting machine', *Annals of statistics* pp. 1189–1232.

Geretto, C., Chung Kim Yuen, S. & Nurick, G. N. (2015), 'An experimental study of the effects of degrees of confinement on the response of square mild steel plates subjected to blast loading', *International Journal of Impact Engineering* **79**, 32–44.

Glorot, X. & Bengio, Y. (2010), Understanding the difficulty of training deep feedforward neural networks, *in* Y. W. Teh & M. Titterington, eds, 'Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics', Vol. 9 of *Proceedings of Machine Learning Research*, PMLR, Chia Laguna Resort, Sardinia, Italy, pp. 249–256.

Godunov, S. K. (1959), 'A difference scheme for numerical solution of discontinuous solution of hydrodynamic equations', *Math. Sbornik* **47**, 271–306.

Goodfellow, C., Bengio, Y. & Courville, A. (2016), *Deep learning*, Adaptive computation and machine learning, The MIT Press, Cambridge, Massachusetts.

Granström, S. A. (1956), Loading characteristics of air blasts from detonating charges, Technical Report 100, Royal Institute of Technology, Stockholm, Sweden.

Guyon, I., Boser, B. & Vapnik, V. (1993), Automatic capacity tuning of very large vc-dimension classifiers, *in* 'Advances in neural information processing systems', pp. 147–155.

Haque, M. (2001), 'Prediction of corrosion–fatigue behavior of DP steel through artificial neural network', *International Journal of Fatigue* **23**(1), 1–4.

Harrison, R. (2018), 'Lecture notes in data modelling and machine intelligence'.

Hashemi, S. & Bradford, M. A. (2014), Numerical simulation of free-air explosion using ls-dyna, *in* 'Applied Mechanics and Materials', Vol. 553, Trans Tech Publ, pp. 780–785.

Hastie, T., Tibshirani, R. & Friedman, J. (2009), *The elements of statistical learning: data mining, inference, and prediction*, Springer Science & Business Media.

Henrych, J. (1979), *The Dynamics of Explosion and Its Use*, Elsevier Scientific Publishing Company, North-Holland, New York.

Hetherington, J. & Smith, P. (2014), *Blast and ballistic loading of structures*, CRC Press.

Hirt, C. W., Amsden, A. A. & Cook, J. (1974), 'An arbitrary lagrangian-eulerian computing method for all flow speeds', *Journal of computational physics* **14**(3), 227–253.

Hodges, J. L. (1958), 'The significance probability of the smirnov two-sample test', *Arkiv för Matematik* **3**(5), 469–486.

Hopkinson, B. (1914), 'A method of measuring the pressure produced in the detonation of high explosives or by the impact of bullets', *Philiosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **213**(1914), 437–456.

Hopkinson, B. (1915), 'Board minutes'.

Hornik, K., Stinchcombe, M. & White, H. (1989), 'Multilayer feedforward networks are universal approximators', *Neural Networks* **2**(5), 359–366.

Hu, X., Hu, H., Verma, S. & Zhang, Z. (2020), 'Physics-Guided Deep Neural Networks for PowerFlow Analysis', *arXiv:2002.00097 [cs, eess, stat]* .

Huffington, N. J. & Ewing, W. O. (1985), *Reflected Impulse Near Spherical Charges*, BRL-TR-2678, Ballistic Research Laboratories, MD, USA.

Hyde, D. (1991), 'Conventional Weapons Effect Program (CONWEP)', *US Army* .

Iacoviello, F. (2004), 'Analysis of stress ratio effects on fatigue propagation in a sintered duplex steel by experimentation and artificial neural network approaches', *International Journal of Fatigue* **26**(8), 819–828.

Iliopoulos, A. P., Michopoulos, J. G., Avery, P., Farhat, C., Teferra, K. & Qidwai, S. (2017), Towards model order reduction for uncertainty propagation in blast-induced traumatic brain injury, *in* 'International Design Engineering Technical Conferences and Computers and Information in Engineering Conference', Vol. 58110, American Society of Mechanical Engineers, p. V001T02A086.

Irving, C. (1995), *In Their Name: Dedicated to the Brave and the Innocent, Oklahoma City, April 1995*, Random House Incorporated.

Jia, X., Willard, J., Karpatne, A., Read, J., Zwart, J., Steinbach, M. & Kumar, V. (2019), Physics guided rnns for modeling dynamical systems: A case study in simulating lake temperature profiles, *in* 'Pro-

ceedings of the 2019 SIAM International Conference on Data Mining', SIAM, pp. 558–566.

Jiang, W., Bennett, A., Vlahopoulos, N. & Zhang, G. (2020), A reduced-order modeling based design and optimization for a lightweight multilayer armor plate against blast and impact, *in* 'Advanced Materials for Defense', Springer, pp. 79–93.

Jin, Y. (2005), 'A comprehensive survey of fitness approximation in evolutionary computation', *Soft computing* **9**(1), 3–12.

Jin, Y. & Sendhoff, B. (2009), 'A systems approach to evolutionary multiobjective structural optimization and beyond', *IEEE Computational Intelligence Magazine* **4**(3), 62–76.

Karpatne, A., Watkins, W., Read, J. & Kumar, V. (2017), 'Physics-guided neural networks (PGNN): An application in lake temperature modeling'.

KılıÇ, N., Ekici, B. & Hartomacıoğlu, S. (2015), 'Determination of penetration depth at high velocity impact using finite element method and artificial neural network tools', *Defence Technology* **11**(2), 110–122.

Kim, S.-W., Melby, J. A., Nadal-Caraballo, N. C. & Ratcliff, J. (2015), 'A time-dependent surrogate model for storm surge prediction based on an artificial neural network using high-fidelity synthetic hurricane modeling', *Natural Hazards* **76**(1), 565–585.

Kingery, C. N. (1966), *Airblast parameters versus distance for hemispherical TNT surface bursts*, ARBRL-TR-1344, US Army Ballistic Research Laboratory, Aberdeen Proving Ground, MD, USA.

Kingery, C. N. & Bulmash, G. (1984), *Airblast Parameters from TNT Spherical Air Burst and Hemispherical Surface Burst*, ARBRL-TR-02555, US Army Ballistic Research Laboratory, Aberdeen Proving Ground, MD, USA.

Kinney, G. & Graham, K. (1985), 'Explosive shocks in air'.

Klomfass, A. (2016), 'Improved explosion consequence analysis with combined cfd and damage models', *Chemical Engineering Transactions* **48**, 109–114.

Klomfass, A. (2018), Accuracy of CFD predictions for explosive far fields, *in* '25th International Conference on Military Aspects of Blast and Shock (MABS252), The Hague, Netherlands'.

Knock, C. & Davies, N. (2013), 'Blast waves from cylindrical charges', *Shock Waves* **23**(4), 337–343.

Langran-Wheeler, C., Rigby, S. E., Clarke, S. D., Tyas, A., Stephens, C. & Walker, R. (2021), 'Near-field spatial and temporal blast pressure distributions from non-spherical charges: Horizontally-aligned cylinders', *International Journal of Protective Structures* **12**(4), 492–516.

Langran-Wheeler, C., Tyas, A., Rigby, S., Stephens, C., Clarke, S. & Walker, R. (2019), Reflected blast loads from long cylinders in the near-field, *in* 'Proceedings of the 18th international symposium for the interaction of munitions with structures (ISIEMS18), Panama City, FL'.

Langran-Wheeler, C., Tyas, A., Rigby, S., Stephens, C., Clarke, S. & Warren, J. (2017), Characterisation of reflected blast loads in the very-near field from non-spherical explosive charges, *in* 'Proceedings of the

17th international symposium for the interaction of munitions with structures (ISIEMS17)', Vol. 16.

Lassila, T., Manzoni, A., Quarteroni, A. & Rozza, G. (2014), 'Model order reduction in fluid dynamics: challenges and perspectives', *Reduced Order Methods for modeling and computational reduction* pp. 235–273.

Lax, P. D. (1954), 'Weak solutions of nonlinear hyperbolic equations and their numerical computation', *Communications on pure and applied mathematics* **7**(1), 159–193.

Lee, E. L., Hornig, H. C. & Kury, J. W. (1968), *Adiabatic Expansion of High Explosive Detonation Products*, TID 4500-UCRL 50422, Lawrence Radiation Laboratory, University of California, CA, USA.

Lee, S.-I., Chatalbashev, V., Vickrey, D. & Koller, D. (2007), Learning a meta-level prior for feature relevance from multiple related tasks, *in* 'Proceedings of the 24th international conference on Machine learning', pp. 489–496.

Li, B., Yang, Q. & Xue, X. (2009), Can movies and books collaborate? cross-domain collaborative filtering for sparsity reduction, *in* 'Twenty-First international joint conference on artificial intelligence', Citeseer.

Li, J., Li, Q., Hao, H. & Li, L. (2021), 'Prediction of bleve blast loading using cfd and artificial neural network', *Process Safety and Environmental Protection* **149**, 711–723.

Lien, F. S. & Leschziner, M. A. (1994), 'Upstream monotonic interpolation for scalar transport with application to complex turbulent flows', *International Journal for Numerical Methods in Fluids* **19**(6), 527–548.

Ling, X., Xue, G.-R., Dai, W., Jiang, Y., Yang, Q. & Yu, Y. (2008), Can chinese web pages be classified with english data source?, *in* 'Proceedings of the 17th international conference on World Wide Web', pp. 969–978.

Loghmanian, S. M. R., Jamaluddin, H., Ahmad, R., Yusof, R. & Khalid, M. (2012), 'Structure optimization of neural network for dynamic system modeling using multi-objective genetic algorithm', *Neural Computing and Applications* **21**(6), 1281–1295.

Löhner, R. & Baum, J. D. (2004), Comparison of coarse and fine mesh 3D Euler predictions for blast loads on generic building configurations, *in* '18th International Symposium for The', Bad Reichenhall.

Löhner, R., Morgan, K., Peraire, J. & Vahdati, M. (1987), 'Finite element flux-corrected transport (FEM-FCT) for the euler and Navier-Stokes equations', *International Journal for Numerical Methods in Fluids* **7**(10), 1093–1109.

Loy, Y., Rangaiah, G. & Lakshminarayanan, S. (2017), 'Surrogate modelling for enhancing consequence analysis based on computational fluid dynamics', *Journal of Loss Prevention in the Process Industries* **48**, 173–185.

Luccioni, B., Ambrosini, D. & Danesi, R. (2006), 'Blast load assessment using hydrocodes', *Engineering Structures* **28**(12), 1736–1744.

Mack, Y., Goel, T., Shyy, W. & Haftka, R. (2007), Surrogate model-based optimization framework: a case study in aerospace design, *in* 'Evolutionary computation in dynamic and uncertain environments', Springer, pp. 323–342.

Mahmud, M. & Ray, S. R. (2007), Transfer learning using kolmogorov complexity: Basic theory and empirical evaluations, Technical report.

Mor-Yosef, S., Samueloff, A., Modan, B., Navot, D. & Schenker, J. G. (1990), 'Ranking the risk factors for cesarean: logistic regression analysis of a nationwide study.', *Obstetrics and gynecology* **75**(6), 944–947.

Muralidhar, N., Islam, M. R., Marwah, M., Karpatne, A. & Ramakrishnan, N. (2018), Incorporating Prior Domain Knowledge into Deep Neural Networks, *in* '2018 IEEE International Conference on Big Data (Big Data)', IEEE, Seattle, WA, USA, pp. 36–45.

Nansteel, M. W., Veldman, R. L., Chen, C. C.-T. & Lawrence, W. (2013), 'Impulse plug measurements of blast reflected impulse at close range', *Propellants, Explosives, Pyrotechnics* **38**(1), 120–128.

Netherton, M. D. & Stewart, M. G. (2016), 'Risk-based blast-load modelling: Techniques, models and benefits', *International Journal of Protective Structures* **7**(3), 430–451.

Noble, W. S. (2006), 'What is a support vector machine?', *Nature biotechnology* **24**(12), 1565–1567.

Notley, S. V., Chen, Y., Lee, P. D. & Panoutsos, G. (2021), Variance stabilised optimisation of neural networks: A case study in additive manufacturing, *in* '2021 International Joint Conference on Neural Networks (IJCNN)', IEEE, pp. 1–7.

Oklahoma City Police Department (2007), 'Alfred P. Murrah Building Bombing After Action Report', `https://web.archive.org/web/20070703233435/http://www.terrorisminfo.mipt.org/pdf/okcfr_App_C.pdf`. Archived from original: 03/07/2007.

Pan, S. J., Kwok, J. T., Yang, Q. & Pan, J. J. (2007), Adaptive localization in a dynamic wifi environment through multi-view learning, *in* 'AAAI', Vol. 7, pp. 1108–1113.

Pan, S. J., Shen, D., Yang, Q. & Kwok, J. T. (2008), Transferring localization models across space., *in* 'AAAI', pp. 1383–1388.

Pan, S. J. & Yang, Q. (2010), 'A survey on transfer learning', *IEEE Transactions on Knowledge and Data Engineering* **22**(10), 1345–1359.

Pannell, J. J., Panoutsos, G., Cooke, S. B., Pope, D. J. & Rigby, S. E. (2021), 'Predicting specific impulse distributions for spherical explosives in the extreme near-field using a gaussian function', *International Journal of Protective Structures* **12**(4), 437–459.

Pannell, J. J., Rigby, S. E. & Panoutsos, G. (2020), A physics guided machine learning approach to understanding loadingn distributions from explosive events, *in* '2020 Young Researchers Conference, Institution of Structural Engineers, London, U.K.'.

Pannell, J. J., Rigby, S. E. & Panoutsos, G. (2022a), 'Physics-informed regularisation procedure in neural networks: an application in blast protection engineering', *International Journal of Protective Structures* .

Pannell, J. J., Rigby, S. E. & Panoutsos, G. (2022b), 'Application of transfer learning for the prediction of blast impulse', *International Journal of Protective Structures* .

Pannell, J. J., Rigby, S. E., Panoutsos, G., Tyas, A., Cooke, S. B. & Pope, D. J. (2019), Predicting near-field specific impulse distributions using machine learning, *in* '18th International Symposium on Interaction of

the Effects of Munitions with Structures (ISIEMS18), Panama City Beach, FL, USA'.

Panoutsos, G. & Mahfouf, M. (2008), An incremental learning structure using granular computing and model fusion with application to materials processing, *in* 'Intelligent techniques and tools for novel system architectures', Springer, pp. 139–153.

Papadopoulos, V., Soimiris, G., Giovanis, D. & Papadrakakis, M. (2018), 'A neural network-based surrogate model for carbon nanotubes with geometric nonlinearities', *Computer Methods in Applied Mechanics and Engineering* **328**, 411–430.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. & Duchesnay, E. (2011), 'Scikit-learn: Machine learning in Python', *Journal of Machine Learning Research* **12**, 2825–2830.

Pinto, G., Wang, Z., Roy, A., Hong, T. & Capozzoli, A. (2022), 'Transfer learning for smart buildings: A critical review of algorithms, applications, and future perspectives', *Advances in Applied Energy* p. 100084.

Plooster, M. (1977), *Blast Front Pressure from Cylindrical Charges of High Explosive*, Denver Research Institute.

Plooster, M. N. (1982), Blast effects from cylindrical explosive charges: Experimental measurements, Technical report, Denver Research Institute.

Pracht, W. E. (1975), 'Calculating three-dimensional fluid flows at all speeds with an eulerian-lagrangian computing mesh', *Journal of Computational Physics* **17**(2), 132–159.

Raina, R., Ng, A. Y. & Koller, D. (2006), Constructing informative priors using transfer learning, *in* 'Proceedings of the 23rd international conference on Machine learning', pp. 713–720.

Ramon, J., Driessens, K. & Croonenborghs, T. (2007), Transfer learning in reinforcement learning problems through partial policy recycling, *in* 'European Conference on Machine Learning', Springer, pp. 699–707.

Randers-Pehrson, G. & Bannister, K. A. (1997), *Airblast Loading Model for DYNA2D and DYNA3D*, ARL-TR-1310, Army Research Laboratory, MD, USA.

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N. & Prabhat (2019), 'Deep learning and process understanding for data-driven Earth system science', *Nature* **566**(7743), 195–204.

Remennikov, A. M. (2003), 'A review of methods for predicting bomb blast effects on buildings', *Journal of battlefield technology* **6**(3), 8.

Remennikov, A. M. & Mendis, P. A. (2006), 'Prediction of airblast loads in complex environments using artificial neural networks', *WIT Transactions on the Built Environment* **87**, 269–278.

Remennikov, A. M. & Mendis, P. A. (2012), Prediction of airblast loads in complex environments using artificial neural networks, *in* S. Syngellakis, ed., 'WIT Transactions on State of the Art in Science and Engineering', first edn, Vol. 1, WIT Press, pp. 53–62.

Remennikov, A. M., Ngo, T., Mohotti, D., Uy, B. & Netherton, M. (2017), 'Experimental investigation and simplified modeling of response of steel plates subjected to close-in blast loading from spherical liquid explosive charges', *International Journal of Impact Engineering* **101**, 78–89.

Remennikov, A. M. & Rose, T. A. (2005), 'Modelling blast loads on buildings in complex city geometries', *Computers & Structures* **83**(27), 2197–2205.

Remennikov, A. M. & Rose, T. A. (2007), 'Predicting the effectiveness of blast wall barriers using neural networks', *International Journal of Impact Engineering* **34**(12), 1907–1923.

Rezaee, K., Savarkar, S., Yu, X. & Zhang, J. (2022), 'A hybrid deep transfer learning-based approach for parkinson's disease classification in surface electromyography signals', *Biomedical Signal Processing and Control* **71**, 103161.

Richtmyer, R. & Morton, K. (1967), *Difference Methods for Initial Value Problems*, London: Wiley-Interscience.

Rickman, D. D. & Murrell, D. W. (2007), 'Development of an improved methodology for predicting airblast pressure relief on a directly loaded wall'.

Rigby, S. E. (2014), Blast Wave Clearing Effects on Finite-Sized Targets Subjected to Explosive Loads, PhD thesis, University of Sheffield.

Rigby, S. E., Akintaro, O. I., Fuller, B. J., Tyas, A., Curry, R. J., Langdon, G. S. & Pope, D. J. (2019b), 'Predicting the response of plates subjected to near-field explosions using an energy equivalent impulse', *International Journal of Impact Engineering* **128**, 24–36.

Rigby, S. E., Fay, S. D., Tyas, A., Warren, J. A. & Clarke, S. D. (2015a), 'Angle of Incidence Effects on Far-Field Positive and Negative Phase Blast Parameters', *International Journal of Protective Structures* **6**(1), 23–42.

Rigby, S. E., Fuller, B. J. & Tyas, A. (2018), Validation of near-field blast loading in LS-DYNA, *in* '5th International Conference on Protective Structures (ICPS5), Poznan, Poland'.

Rigby, S. E., Knighton, R., Clarke, S. D. & Tyas, A. (2020), 'Reflected near-field blast pressure measurements using high speed video', *Accepted for publication in Experimental Mechanics* .

Rigby, S. E., Tyas, A., Bennett, T., Fay, S. D., Clarke, S. D. & Warren, J. A. (2014a), 'A Numerical Investigation of Blast Loading and Clearing on Small Targets', *International Journal of Protective Structures* **5**(3), 253–274.

Rigby, S. E., Tyas, A., Clarke, S. D., Fay, S. D., Reay, J. J., Warren, J. A., Gant, M. & Elgy, I. (2015c), 'Observations from preliminary experiments on spatial and temporal pressure measurements from near-field free air explosions', *International Journal of Protective Structures* **6**(2), 175–190.

Rigby, S. E., Tyas, A., Curry, R. J. & Langdon, G. S. (2019a), 'Experimental measurement of specific impulse distribution and transient deformation of plates subjected to near-field explosive blasts', *Experimental Mechanics* **59**(2), 163–178.

Rigby, S., Fay, S., Clarke, S., Tyas, A., Reay, J., Warren, J., Gant, M. & Elgy, I. (2016), 'Measuring spatial pressure distribution from explosives buried in dry leighton buzzard sand', *International Journal of Impact Engineering* **96**, 89–104.

Rigby, S., Osborne, C., Langdon, G., Cooke, S. & Pope, D. (2021), 'Spherical equivalence of cylindrical explosives: Effect of charge shape

on deflection of blast-loaded plates', *International Journal of Impact Engineering* **155**, 103892.

Rigby, S., Tyas, A., Clarke, S. D., Fay, S. D., Warren, J. A., Elgy, I. & Gant, M. (2014c), Testing apparatus for the spatial and temporal pressure measurements from near-field free air explosions, *in* '6th International Conference on Protection of Structures Against Hazards (PSH14), Tianjin, China'.

Rigby, S., Tyas, A., Clarke, S., Fay, S., Reay, J., Warren, J., Gant, M. & Elgy, I. (2015b), 'Observations from Preliminary Experiments on Spatial and Temporal Pressure Measurements from Near-Field Free Air Explosions', *International Journal of Protective Structures* **6**(2), 175–190.

Rohatgi, A. (2020), 'Webplotdigitizer: Version 4.4'.
URL: *https://automeris.io/WebPlotDigitizer*

Ryan, S., Thaler, S. & Kandanaarachchi, S. (2016), 'Machine learning methods for predicting the outcome of hypervelocity impact events', *Expert Systems with Applications* **45**, 23–39.

Savitzky, A. & Golay, M. J. E. (1964), 'Smoothing and differentiation of data by simplified least squares procedures', *Analytical Chemistry* **36**(8), 1627–1639.

Schölkopf, B., Burges, C. & Vapnik, V. (1996), Incorporating invariances in support vector learning machines, *in* 'International Conference on Artificial Neural Networks', Springer, pp. 47–52.

Schölkopf, B., Smola, A. J., Bach, F. et al. (2002), *Learning with kernels: support vector machines, regularization, optimization, and beyond*, MIT press.

Schölkopf, P., Burgest, C. & Vapnik, V. (1995), Extracting support data for a given task, *in* 'Proceedings, First International Conference on Knowledge Discovery & Data Mining. AAAI Press, Menlo Park, CA', pp. 252–257.

Schwer, L. & Rigby, S. (2017), Reflected secondary shocks: Some observations using afterburning, *in* 'Proceedings of the 11th European LS-DYNA Conference, Salzburg, Austria'.

Schwer, L. & Rigby, S. (2018), Secondary and height of burst shock reflections: Application of afterburning, *in* 'Proceedings of the 25th Military Aspects of Blast and Shock (MABS25), The Hauge, Netherlands'.

Sherkar, P., Shin, J., Whittaker, A. & Aref, A. (2016), 'Influence of charge shape and point of detonation on blast-resistant design', *Journal of Structural Engineering* **142**(2), 04015109.

Shin, J., ttaker, A. S., Cormie, D. & Wilkinson, W. (2014a), 'Numerical modeling of close-in detonations of high explosives', *Engineering Structures* **81**, 88–97.

Shin, J., Whittaker, A. S., Aref, A. J. & Cormie, D. (2014b), *Air-Blast Effects on Civil Structures*, MCEER-14-0006, The University at Buffalo, NY, USA.

Shin, J., Whittaker, A. S. & Cormie, D. (2015), 'Tnt equivalency for overpressure and impulse for detonations of spherical charges of high explosives', *International Journal of Protective Structures* **6**(3), 567–579.

Smith, P. D. & Rose, T. A. (2006), 'Blast wave propagation in city streets—an overview', *Progress in Structural Engineering and Materials* **8**(1), 16–28.

Smola, A. J. & Schölkopf, B. (2004), 'A tutorial on support vector regression'.

Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. (2014), 'Dropout: a simple way to prevent neural networks from overfitting', *J. Mach. Learn. Res.* **15**, 1929–1958.

Stewart, L. (2010), Testing and Analysis of Structural Steel Columns Subjected to Blast Loads, PhD thesis, University of California, San Diego.

Stewart, L. & Morrill, K. (2015), 'Residual capacity prediction of blast-loaded steel columns using physics-based fast running models', *International Journal of Safety and Security Engineering* **5**(4), 289–303.

Stewart, R. & Ermon, S. (2017), Label-free supervision of neural networks with physics and domain knowledge, *in* 'Thirty-First AAAI Conference on Artificial Intelligence'.

Stoner, R. G. & Bleakney, W. (1948), 'The attenuation of spherical shock waves in air', *Journal of Applied Physics* **19**(7), 670–678.

Swisdak, M. M. (1994), Simplified Kingery airblast calculations, *in* '26th Department of Defense Explosives Safety seminar, Miami, FL, USA'.

Taher M. Ghazal, Sagheer Abbas, S. M. M. A. K. M. A. G. F. I. S. B. Z. M. A. K. M. K. H. (2022), 'Alzheimer disease detection empowered with transfer learning', *Computers, Materials & Continua* **70**(3), 5005–5019.

Taleb, N. N. (2007), *The black swan: The impact of the highly improbable*, Vol. 2, Random house.

Taylor, G. I. (1950), 'The formation of a blast wave by a very intense explosion i. theoretical discussion', *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **201**(1065), 159–174.

Taylor, M. E. & Stone, P. (2007), Cross-domain transfer for reinforcement learning, *in* 'Proceedings of the 24th international conference on Machine learning', pp. 879–886.

The Guardian (2012), 'Anders Behring Breivik: the indictment', `https://www.theguardian.com/world/2012/apr/16/anders-behring-breivik-indictment`. Accessed: 13/07/2021.

Tyas, A. (2019), Blast loading from high explosive detonation: what we know and what we don't know, *in* '13th International Conference on Shock and Impact Loads on Structures', Guangzhou, China.

Tyas, A., Reay, J. J., Fay, S. D., Clarke, S. D., Rigby, S. E., Warren, J. A. & Pope, D. J. (2016), 'Experimental studies of the effect of rapid afterburn on shock development of near-field explosions', *International Journal of Protective Structures* **7**(3), 456–465.

Tyas, A. & Rigby, S. E. (2018), 'Civ4415: Blast and impact effects on structures – course notes', University of Sheffield.

Tyas, A., Warren, J. A., Bennett, T. & Fay, S. (2011), 'Prediction of clearing effects in far-field blast loading of finite targets', *Shock Waves* **21**(2), 111–119.

Vapnik, V. (2006), *Estimation of dependences based on empirical data*, Springer Science & Business Media.

Vapnik, V. (2013), *The nature of statistical learning theory*, Springer science & business media.

Vapnik, V. & Chervonenkis, A. (1971), 'On the uniform convergence of relative frequencies of events to their probabilities', *Theory of Probability and its Applications* **16**, 264–280.

Vapnik, V., Golowich, S. E., Smola, A. et al. (1997), 'Support vector method for function approximation, regression estimation, and signal processing', *Advances in neural information processing systems* pp. 281–287.

Vickers, J. (2011), 'Can induction be justified', *The Stanford Encyclopedia of Philosophy. Stanford: The Metaphysics Research Lab* .

White, D. A., Arrighi, W. J., Kudo, J. & Watts, S. E. (2019), 'Multiscale topology optimization using neural network surrogate models', *Computer Methods in Applied Mechanics and Engineering* **346**, 1118–1135.

Whittaker, M. J., Klomfass, A., Softley, I. D., Pope, D. J. & Tyas, A. (2019), Comparison of numerical analysis with output from precision diagnostics during near-field blast evaluation, *in* '18th International Symposium on Interaction of the Effects of Munitions with Structures (ISIEMS18), Panama City Beach, FL, USA'.

Wilkinson, W., Cormie, D., Shin, J. & Whittaker, A. (2013), Modelling close-in detonations in the near-field, *in* 'Proceedings of the 15th International Symposium on Interaction of the Effects of Munitions with Structures (ISIEMS)'.

Wilson, D., Blass, D. & Noli, S. (2018), Implementation of MCEER TR 14-0006 blast load curves in LS-DYNA and benchmark to commonly

practiced blast loading application methods, *in* '15th International LS-DYNA Users Conference, Detroit, MI, USA'.

Wisotski, J. & Snyer, W. H. (1965), *Characteristics of Blast Waves Obtained from Cylindrical High Explosive Charges (U).*, Denver Research Institute, University of Denver.

Wolpert, D. H. (1996), 'The lack of a priori distinctions between learning algorithms', *Neural computation* **8**(7), 1341–1390.

Wu, C., Fattori, G., Whittaker, A. & Oehlers, D. J. (2010), 'Investigation of air-blast effects from spherical-and cylindrical-shaped charges', *International Journal of Protective Structures* **1**(3), 345–362.

Xing, D., Dai, W., Xue, G.-R. & Yu, Y. (2007), Bridged refinement for transfer learning, *in* 'European Conference on Principles of Data Mining and Knowledge Discovery', Springer, pp. 324–335.

Yeo, I.-K. & Johnson, R. A. (2000), 'A new family of power transformations to improve normality or symmetry', *Biometrika* **87**(4), 954–959.

Yin, J., Yang, Q. & Ni, L. (2005), Adaptive temporal radio maps for indoor location estimation, *in* 'Third IEEE international conference on pervasive computing and communications', IEEE, pp. 85–94.

Zeiler, M. (2012), 'ADADELTA: An Adaptive Learning Rate Method', *arXiv:1212.5701 [cs]* .

Zhang, R., Liu, Y. & Sun, H. (2020), 'Physics-guided convolutional neural network (PhyCNN) for data-driven seismic response modeling', *Engineering Structures* **215**, 110704.

Zheng, V. W., Xiang, E. W., Yang, Q. & Shen, D. (2008), Transferring localization models over time., *in* 'AAAI', Vol. 2008, pp. 1421–1426.

Zinne, E. & Williams, M. (2013), *When A Community Weeps: Case Studies In Group Survivorship*, Routledge.

Appendices

Appendix A

A.1 *Apollo blastsimulator* equation of state information

Table A.1: Equation of state information for the six studied charge compositions

Charge type	**Gas constant** $J/(kg.K)$	**Initial density** (kg/m^3)	C_1 **(GPa)**	C_2 **(GPa)**	R_1	R_2
PE4	365	1660	734.60	8.86	4.79	1.06
TNT	315	1630	527.28	6.30	4.71	1.07
HMX	365	1905	1215.93	12.30	4.77	1.12
RDX	345	1805	1053.18	11.29	4.77	1.12
PETN	300	1778	946.18	9.83	4.76	1.08
COMPB	345	1725	732.74	8.83	4.68	1.09