Time Frequency Analysis Techniques in Terahertz Pulsed Imaging

by

James William Handley

Submitted in accordance with the requirements
for the degree of Doctor of Philosophy.

The University of Leeds
School of Computing

December 2003

The candidate confirms that the work submitted is his own and the appropriate credit has been given where reference has been made to the work of others.

Abstract

Terahertz (THz) radiation is abundant in the natural world yet very hard to harness in the laboratory. Forming the boundary between ‘radio’ and ‘light’, the so called “terahertz gap” results from the failure of optical techniques to operate below a few hundred terahertz, and likewise the failure of electronic methods to operate above a few hundred gigahertz. However, recent advances in opto-electronic and semiconductor technology have enabled bright THz radiation to be coherently generated and detected, and THz imaging systems are now commercially available, if still very expensive. Terahertz pulsed imaging data are unusual in that an entire time series is ‘behind’ every pixel of the image. While resulting in rich data sets, this high dimensionality necessitates some form of distillation or extraction of pertinent features before images can be formed.

Within this thesis the technology of THz pulsed imaging is examined, together with the imaging modalities that are employed and the type of data that are acquired. The sources of noise are categorised, and it is demonstrated that this noise can be modelled by the family of stable distributions, but that it is neither normally distributed nor distributed according to a simple mixture of Gaussians. Joint time-frequency techniques such as those used in RADAR or ultrasound – windowed Fourier transforms and wavelet transforms – are applied to THz data, and are shown to be appropriate tools to use when analysing and processing THz pulses, particularly in signal compression. Finally, clustering algorithms in time, frequency, and time-frequency based feature spaces demonstrate that such tools have potential application in the segmentation of THz images into their constituent regions.

The analyses herein improve our understanding of the nature of THz data, and the techniques developed are steps along the road to move THz imaging into real world applications, such as dental and medical imaging and diagnosis.

Acknowledgements

I would like to thank my supervisors, Dr E Berry and Prof R Boyle, and additionally Dr A Fitzgerald, for all the assistance and guidance they offered during my research and writing.

I would like to thank the members of Teravision Workpackage 4; specifically CoMIR and IMP at the University of Leeds, Physikalisches Institüt, J.W. Goethe–Universität, Frankfurt, and TeraView (formerly Toshiba Research Europe Laboratory) in Cambridge for providing THz data.

I would like to thank all my fellow researchers at Leeds (both past and present), especially Dr N Cohen for her input on stable distributions. Particular thanks also to those within the Computer Vision and CoMIR research groups, for helping to create such a productive and pleasant working environment.

My thanks also to the EPSRC for funding this research.

Finally I would like to thank my family and friends for all their prayers and support, particularly my wife Anna, who went far beyond the call of duty in proof reading.

quia quod stultum est Dei sapientius est hominibus

Declarations

Some parts of the work presented in this thesis have been published in the following articles:

J.W. Handley, A.J. Fitzgerald, T. Löffler, K. Siebert, E. Berry, and R.D. Boyle, "Potential Medical Applications of THz Imaging", *Proceedings Medical Image Understanding and Analysis 2001*, pp 17–20.

J.W. Handley, A.J. Fitzgerald, E. Berry, and R.D. Boyle, "Approaches to Segmentation in Medical Terahertz Pulsed Imaging", *Proceedings Medical Image Understanding and Analysis 2002*, pp 157–160.

J.W. Handley, A.J. Fitzgerald, E. Berry, and R.D. Boyle, "Wavelet Compression in Medical Terahertz Pulsed Imaging", *Physics in Medicine and Biology*, **47** (21) pp 3885–3892, 2002.

J.W. Handley, A.J. Fitzgerald, E. Berry, and R.D. Boyle, "The Short Time Fourier Transform applied to Terahertz Pulsed Imaging", *Medical Physics*, **30** (6) p 1541, 2003 (abstract).

E. Berry, R.D. Boyle, A.J. Fitzgerald, and J.W. Handley, "Time and Frequency Analysis in Terahertz Pulsed Imaging", in B. Bhanu and I. Pavlidis, editors, *Computer Vision Beyond the Visible Spectrum*, Chapter 9, pages 290–329. Springer–Verlag, in press.

J.W. Handley, A.J. Fitzgerald, E. Berry, and R.D. Boyle, "Distinguishing between Materials in Terahertz Pulsed Imaging using Wide-Band Cross Ambiguity Functions", Digital Signal Processing, **14** (2) pp 99-111, 2004.

E. Berry, J.W. Handley, A.J. Fitzgerald, W.J. Merchant, R.D. Boyle, N.N. Zinov'ev, R.E. Miles, J.M. Chamberlain, and M.A Smith, "Multispectral Classification Techniques for Terahertz Pulsed Imaging: an Example in Histopathology", Medical Engineering and Physics, in press 2004.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Recent developments in semi-conductor technology have made the region of the electromagnetic spectrum previously known as the "terahertz gap", shown in Figure 1.1, accessible for imaging. Radiation around this frequency — 1 terahertz (THz) is 10^{12} cycles per second or 3 mm in wavelength — although abundant in nature, falls between the limits of optical and electronic technology, and had been impossible to coherently generate and detect until as recently as the late 1980s.

This band forms a very interesting region of the electromagnetic spectrum for several reasons, including the sensitivity of terahertz radiation to polar substances, such as water, and its insensitivity to non-polar substances, rendering dust, plastic, and even clothes almost transparent. Since the breakthrough in the 1980s, terahertz imaging technology has spread rapidly, and a terahertz scanner was built at the University of Leeds as part of the European Union "Teravision" project that ran from 2000-2003, alongside five other academic and commercial entities across Europe[1].

---

[1] For full details see the Teravision website — *http://www.teravision.org/* — Last visited 27th November 2003.

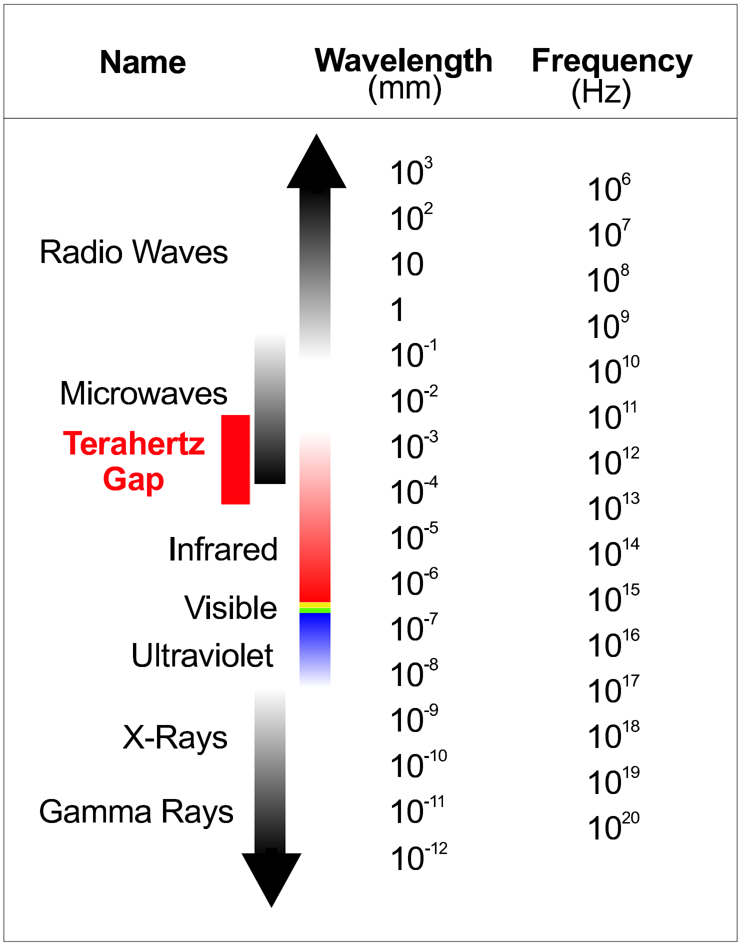

Figure 1.1: The electromagnetic spectrum, with the “terahertz gap” highlighted.

This thesis presents a number of novel analyses and analysis techniques which are suitable for terahertz images.

1.1 Imaging with Terahertz Pulses

The technology considered exclusively in this thesis is terahertz pulsed imaging, whereby broadband terahertz pulses are coherently generated, steered to interact with a sample, and then coherently detected in the time domain. An example of a terahertz pulse is shown in Figure 1.2. When a pulse such as this interacts with a sample, it undergoes changes which are dependent on the optical properties of the sample at terahertz frequencies. Typically the transmitted pulse will experience a delay, an attenuation, and a broadening, as the different component frequencies are phase shifted, absorbed, reflected, and scattered.

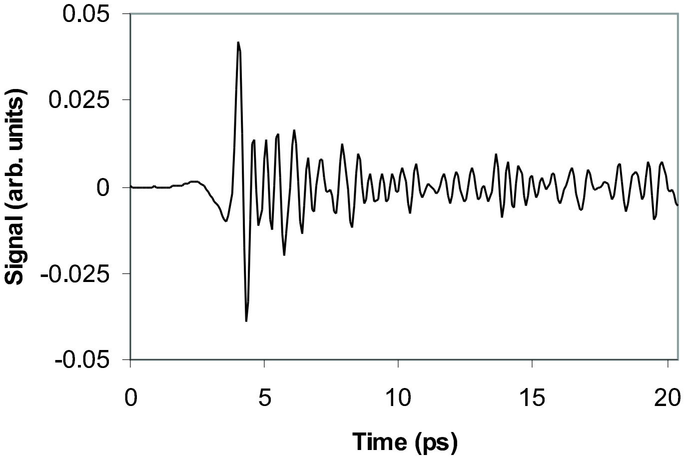

Figure 1.2: An example terahertz pulse, acquired at Leeds.

A key feature of terahertz imaging is that an entire time series, such as that shown in Figure 1.2, is acquired ‘behind’ every pixel of the terahertz ‘image’. On the one hand this causes visualisation problems, as a small number of features must somehow be extracted before a 2-D image can be formed. This is quite different from X-ray or MRI (for example), when generally a single value obtained at each pixel or voxel can be simply mapped to a greyscale value. On the other hand, the acquisition of a coherent time series enables frequency specific phase change and attenuation to be calculated, opening up a rich seam of data.

1.2 Aims and Motivations

Terahertz imaging is an immature technology — the main research emphasis is still firmly on instrumentation — and there are several issues where improvements in technology or understanding would be beneficial to its acceptance as a mainstream imaging technique:

- **Long acquisition times**

At the time of writing, the fastest terahertz system still requires up to 30 minutes to acquire a one hundred by one hundred pixel image with a 512 sample time series.

- **The high dimensionality of data acquired**

  A typical terahertz 'pixel' consists of 512 or 1,024 time samples, and obviously some form of parametric extraction (or other dimensionality reduction) must be carried out before an image can be formed. Even classifying or distinguishing between pixels is non-trivial in such a high dimensional space.

- **Large data sets**

  The high dimensionality of the terahertz data immediately raises the issues of the space needed for storage and bandwidth needed for transmission. A single terahertz image, one hundred by one hundred pixels, would usually be of the order of 40 Mb.

- **System instability and unfriendliness**

  Most terahertz systems are still built on large optical benches using delicate optical components that are prone to failure and drift, as well as being extremely sensitive to misalignment and even variation in atmospheric conditions. Just considering the medical domain, it is unthinkable to have a system where the patient has to be suspended in a harness over an optical bench for several hours without any motion. Systems that are more user friendly systems are starting to appear on the market, however.

- **Resolution**

  Both the spatial and temporal resolution of the terahertz scanners have room for improvement.

- **Noise**

  As with any real world system, there is noise present in terahertz images. The coherent generation and detection can give rise to excellent signal to noise ratios at an individual pulse level, but the modality is still noisy at an 'image' level. These high signal to noise ratios also tend to be reduced if acquisition times are shortened.

Clearly some of these items require electrical engineers and photonics experts, however this thesis endeavours to address some of the issues using signal processing and image processing techniques. In particular, novel strategies for managing the *noise*, the *high dimensionality of data*, the long *acquisition times*, and the large *volume of data* are explored.

Terahertz data naturally lends itself to analyses based in the time or frequency domains, and as the title of this thesis 'Time Frequency Analysis Techniques in Terahertz Pulsed Imaging' suggests, this work has been undertaken in the time domain, the frequency (or Fourier) domain, and in joint time/frequency domains such as those formed by using short time Fourier transforms and wavelet transforms. In this way, this thesis endeavours to take full advantage of the richness of terahertz pulsed imaging data while providing useful analysis techniques to the terahertz practitioner.

1.3 Overview of Thesis

Chapter 2 provides a review of the history and development of terahertz imaging, together with previous work carried out in this field. Chapter 3 contains a detailed description and analysis of the terahertz image acquisition process, followed by a theoretical review of the analysis techniques used throughout this thesis. The remaining chapters describe the original work of the thesis, organised as follows :-

- **Chapter 4**

  A detailed analysis of the noise in transmission mode terahertz pulsed imaging is presented, and original statistically reliable models of this noise are built. A novel exploration of how existing denoising techniques may be extended for data compression is additionally presented.

- **Chapter 5**

  The novel application of short time Fourier transforms and the wide-band cross ambiguity function to transmission mode terahertz data is presented. Both are used to determine the optical parameters of various materials, and are evaluated against each other and against the traditional analysis using the Fourier transform. The short time Fourier transform is also used indirectly to explore the effect that acquiring a shorter time series would have.

- **Chapter 6**

  Building on the previous chapter, the novel application of k-means clustering to terahertz data in order to segment an image into its constituent materials is presented. The clustering is evaluated using a variety of feature vectors drawn from the time domain, frequency domain, a discrete wavelet domain, and the domain of derived physical features (such as absorbance at a given frequency). The evaluation is against hand segmented images obtained in a different modality.

Finally, conclusions and future work are discussed in chapter 7.

Chapter 2

Terahertz imaging: an historical perspective

Electromagnetic radiation in the terahertz band, broadly 300 GHz to 10 THz, was first isolated in 1897 by Heinrich Rubens [56], but remained largely unexplored in the years following. Falling on the boundary between microwave and infrared, this so-called "terahertz gap" resulted from the failure of optical techniques to operate below a few hundred terahertz, and likewise the failure of electronic/radio methods to operate above a few hundred gigahertz.

Aside from the "because it's there" line of reasoning, the terahertz band is interesting because:

- The radiation is non-ionizing,
- The wavelength is shorter than for microwave wavelengths, with the associated improvement in spatial resolution, while still being long enough to experience less

of the Rayleigh scattering experienced by infrared,

- Terahertz radiation is highly sensitive to the presence of polar substances, such as water and thus hydration state,
- 'Dry' non-polar substances, such as plastics, fibres, and so on, are almost transparent to terahertz radiation,
- Light weight molecules have strong emission or absorption lines in this region for rotational and vibrational excitations, and finally,
- The universe is naturally bathed in terahertz radiation.

Recent advances in laser and electro-optical technologies have enabled bright terahertz radiation to be coherently generated and detected, making this band accessible. It should be noted that although there are exciting developments in the field of incoherent detection and generation, such as terahertz cameras and telescopes, this thesis is solely concerned with the coherent generation and detection of terahertz pulses and the processing of the data thus acquired.

2.1 The Development of Terahertz Systems

Before the advent of these bright sources, terahertz radiation was generated either using sources similar to those used in infrared Fourier Transform Spectroscopy, which generate weak and incoherent radiation, or by bulky complex equipment like free electron lasers or optically pumped gas lasers [1, 38]. The detection methods were also incoherent, recording only the intensity of incident radiation using a helium cooled bolometer for example. Unfortunately terahertz radiation is naturally present in abundance — black body radiation at 300 K is at 6.25 THz — making this technology very prone to noise. Advances in the fields of ultrashort pulsed lasers, non-linear optics and crystal growth techniques have enabled these limitations to be overcome through the introduction of terahertz time-domain spectroscopy [2, 60, 36, 61], which appeared in the late 1980s and early 1990s. In these systems, femtosecond pulses are used to both generate a coherent terahertz wave and to subsequently gate that wave's detection. In this way extremely bright radiation is coherently generated and detected, enabling systems to be several orders of magnitude more sensitive than those using bolometric methods. A further advantage to coherent detection is that it is possible to record the amplitude of the electric field in the time domain. This opens up the possibility of using Fourier transforms, wavelet transforms, and other time/frequency techniques in analysis.

The next step was the building of terahertz pulsed imaging systems using terahertz time-domain spectroscopy technology, such as that reported by van Exter et al. in the early 1990s [61], before the first real time imaging system was reported in 1995 by Hu and Nuss [36]. The technology has now reached the stage where terahertz pulsed imaging systems are commercially available[1], although these are still expensive (in excess of UK 200,000 pounds), mainly due to the cost of the lasers. The terahertz pulsed imaging process used to acquire the data that is used throughout this thesis is described in detail in chapter 3.

The hardware and instrumentation side of terahertz imaging is possibly still the main area of active research (see section 2.2, below), and there is every reason to expect acquisition times to drop, resolution and signal to noise ratios to improve, and for the systems to become cheaper and more compact.

2.1.1 Other Terahertz Technologies

This thesis is concerned solely with terahertz pulsed imaging, however other technologies have been and are being developed in parallel with pulsed systems. Continuous wave systems [42, 57] use monochromatic radiation that can be precisely tuned to a specific

---

[1] Teraview in the U.K. and Picometrix in the USA both sell "off the shelf" systems.

frequency, leading to correspondingly simpler data. Other advances include compact free electron laser systems [28], and terahertz microscopy using near-field techniques [52].

The passive detection of incoherent terahertz radiation is a more established field — for example the First International Symposium on Space TeraHertz Technology was held in 1990, and the IEEE proceedings devoted a special issue to terahertz technology in 1992. The interested reader is referred to Phillips and Keene [53] in that issue for a review of this field. Closer to home, passive detection of (incoherent) terahertz radiation is also being explored, with the 'first ever' picture of a human hand taken in September 2002 using a terahertz camera built by ESA's StarTiger project[2]. QinetiQ have also subsequently demonstrated a passive millimetre wave scanner being used to detect weapons or contraband hidden on a person's body.

2.2 Terahertz Imaging

Terahertz imaging is still a very immature field, with the majority of research focused on instrumentation and hardware. For example, at a recent Royal Society discussion meeting [55], 14 out of the 20 posters presented 'pure' instrumentation research, and a further 2 were on the boundary of instrumentation and application. The oral presentations are harder to classify, but of the 14 essentially research based presentations 4 solely dealt in the application of terahertz technology, compared with 7 instrumentation and hardware presentations. The remaining 3 tended to deal with both the technology and its application. With only a handful of terahertz imaging systems around the world, and most of them in physics or electrical engineering research groups, perhaps this is no surprise. Indeed the long acquisition times (a 30 by 30 pixel image currently takes around 32 hours to acquire on the Leeds system) and instability of the systems has meant that terahertz data has been quite scarce. Of course the balance is shifting — the technology is constantly

---

[2] *http://www.startiger.org/* — Last visited 25 November, 2003.

improving (the Teraview “TPI Scan™” system can acquire 100 by 100 pixel images in around 15 to 30 minutes each, for example) and at the time of writing there is a terahertz imaging clinical trial underway in a Cambridge hospital.

The imaging techniques initially tended toward being simple demonstrations of the capabilities of the technology, so, for example, images were found by acquiring a single time point at each pixel, leading to simple greyscale images based on amplitude. The acquisition of an entire time series enables potentially more useful parametric images to be generated ([34] for example). These have often been based on mature techniques from other fields. Other acquisition methods based on terahertz pulsed imaging are also being developed, such as dark field imaging [45]. More sophisticated analyses are emerging too, such as tomographic imaging [26, 62] and reflection geometry imaging [19, 65]. Some applications of these techniques are mentioned in the summary of this chapter.

2.2.1 Computer Vision

The application of computer vision techniques to terahertz pulsed imaging data is still in its infancy. Herrmann et al. have suggested the use of “display modes”, for example using parameters calculated from appropriate parts of the spectrum, such as those corresponding with absorption or emission lines of particular molecules [34]. Löffler et al. have demonstrated the range of parameters available for such displays [46]. These techniques result in relatively simple images, for example false colour images where three different parameters are mapped to the red, green, and blue components of a pixel’s colour. Mittleman et al. suggested the use of wavelet based techniques [49] — an idea taken up by Mickan et al. [48] and Ferguson et al. for denoising [25, 24]. This aspect of analysis is covered in detail in section 2.3.

In addition to denoising, Ferguson et al. have been using multi-spectral classification techniques to identify biological tissue [26]. In that work, chirped probe terahertz imaging[3] [8] is used, rather than the scanning pulse method used throughout this thesis, and the results need to be interpreted in the light of this. They show that a two parameter linear predictor, such as a finite impulse response model, can distinguish between samples of beef or bone, chicken, and air if the two parameters are used in a simple two dimensional Mahalanobis classifier. They found that their parameter based classifier correctly identified 297 out of the 300 pulses, whereas a simple feature vector based on two obvious spectral features only managed 283. It should be noted that 150 of the 300 pulses were used for training the parameter based classifier, and it is not clear how many were used for training the spectral feature classifier. Finally they demonstrate their classifier distinguishing between chicken and chicken bone, this time using a 5 parameter model (and hence 5 dimensional classifier). 10,000 pulses were obtained in a 100x100 image, and 150 of those pulses were chosen to train the classifier, 50 from each of chicken, bone, and air. The evaluation was entirely qualitative — a false colour image showed the classification of each pixel with a photograph of the original sample displayed alongside for comparison. This work is clearly in its early stages, although Ferguson et al. have the advantage of large datasets on which to train the classifiers. In Chapter 6 a different approach is taken, and unsupervised clustering algorithms are applied to images of a cross section of human tooth and to images of specially created phantoms using a variety of feature vectors. Clustering is appropriate because the richer time series data allow more flexibility in choice of vector, and the relative scarcity of data makes removing pulses for training purposes unrealistic, which is why the unsupervised route is taken. The phantoms define a ground truth against which these techniques may be quantitatively assessed.

---

[3] This is a very fast imaging technique that compromises temporal resolution in favour of speed.

2.3 Noise in Terahertz Data

The work to date on noise in terahertz imaging has taken one of two approaches; analysis and evaluation of the system components (such as the lasers and detectors) [61, 14, 54] and signal to noise ratio (SNR) [49, 70], or applying denoising techniques to try and improve the data [25, 24]. The former work concerns itself with noise from a systems engineering point of view, and considers signal to noise ratio on the basis of individual pulses. In chapter 4 this thesis builds upon this work by analysing the noise present across an entire image, and by building models of the ways in which pulses that have been passed through nominally the same material vary. This is noise in a much broader sense, but these models capture the actual variation that needs to be accounted for if algorithms are to be reliable.

Ferguson and Abbott applied Donoho's wavelet shrinkage algorithm for denoising [17], as suggested by Mittleman in 1998 [50]. Wavelets are a natural choice of tool for denoising, argued Mittleman, because of their "striking similarities" with terahertz pulses, and a review of wavelet transforms may be found in chapter 3 of this thesis. Ferguson and Abbott added white Gaussian noise to terahertz pulses in order to reduce the SNR. These 'noised' pulses were then denoised with various mother functions, and the improvement in SNR was quantitatively measured, by comparison with the original pulse. Additionally they employed a qualitative visual comparison method. Ferguson and Abbot take the approach that 'quick and dirty' imaging can be cleaned up using denoising and filtering techniques, and do indeed demonstrate that the noise they add can be significantly improved using their techniques. However, in chapter 4 it is demonstrated that the noise in terahertz data can not be accurately modelled simply using a white Gaussian process. Furthermore Ferguson and Abbott make no attempt to discover how far the denoising can be 'pushed' (or in other words what the minimum threshold value for the wavelet shrinkage can be) before errors are introduced. This second point has key application in the field of data compression, and an exploration into the impact that wavelet based compression has on the calculation of optical constants can be found in chapter 5.

2.4 Summary

It is sometimes said that terahertz imaging is a solution looking for a problem. Indeed when compared to mature technologies just in the medical domain like ultrasound, X-ray, and even newer technologies like CT scanning and MRI, it can be hard to see the need for terahertz imaging. It is certainly true that terahertz imaging is not a panacea, and neither will it replace these well established technologies — but it should be remembered that it is only 8 years since Hu and Nuss reported the first real time imaging system in 1995 [36].

It is an exciting time for the field, and in terms of application areas terahertz radiation has been used to characterise semi-conductors [30], assess the moisture content of a leaf [31], to identify gases [50, 37], to discover items hidden in powder [35], and possibly examine space shuttle tiles for defects [69]. In the biomedical domain, terahertz imaging has been applied to dental tissue [11], skin and skin cancers [65, 66], and DNA [47], amongst other things. Catalogues of the optical properties of human tissue have also been published [4, 27]. The interested reader is referred to recent special issues of journals on the biological application of terahertz radiation[4].

Building on previous research in the signal processing and computer vision domains, this thesis provides a selection of tools and techniques evaluated on terahertz data. This thesis is thus a timely contribution to the field of terahertz pulsed imaging, presenting a tool-kit of analysis techniques that are effective in this field, specifically the short time Fourier transform in chapter 5, and clustering techniques for segmentation in chapter 6. It further demonstrates the limitations of other techniques in the terahertz domain, such as wavelet compression and cross ambiguity functions (chapter 5). Finally, by way of chapter 4, a detailed analysis of the noise present in terahertz pulsed imaging is given,

---

[4] *Physics in Medicine and Biology* **47** (21), 2002, and *Journal of Biological Physics* **29** (2–3), 2003.

and it is shown that this noise may be modelled by distributions from the stable family.

Chapter 3

Theory

3.1 Introduction

This chapter forms a detailed review of the processes and techniques used in the remainder of the thesis. We will start with a detailed examination of the process of acquiring terahertz data — a necessary step in understanding what terahertz data actually consist of, and where noise and errors appear. This is followed by a review of the mathematical techniques used in the time and frequency analysis of terahertz data, namely Fourier and wavelet techniques. Finally there is a brief overview of clustering techniques.

3.2 Terahertz pulsed imaging systems

3.2.1 Hardware

In terahertz pulsed imaging, pulses of terahertz radiation are generated using either non-linear optics or a dipole antenna, steered to interact with a sample, and are subsequently detected in the time domain using similar techniques to the generation. These pulses are recorded via the selective amplification of a lock-in amplifier. The terahertz imager may be set up in either transmission or reflection modality, where the detector is either the diametrically opposite side of the sample to the transmitter (transmission mode) or is the same side of the sample as the transmitter, being positioned in such as way as to capture reflections from the sample (reflection mode). Figure 3.1 shows a schematic of a terahertz pulsed imaging system in transmission mode, which is described in detail below. Throughout this thesis systems employing non-linear optical rectification, as described below, were used although some data were acquired in reflection mode.

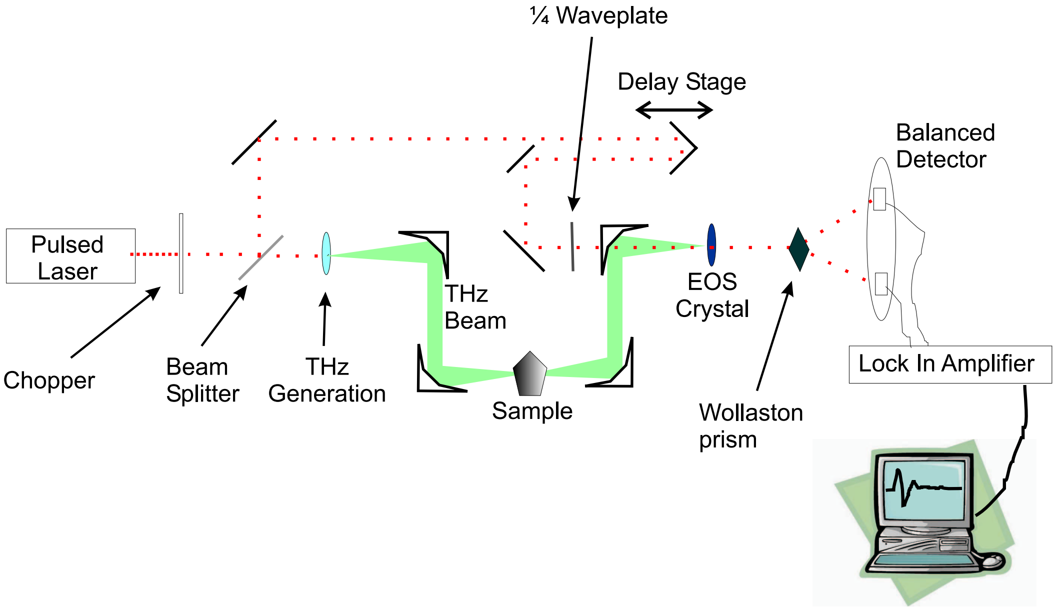

Figure 3.1: Schematic of a terahertz pulsed imaging system in transmission mode.

An ultra-fast Ti:Sapphire laser emits femtosecond pulses of wavelength around 775 nm (380 THz, in the near infrared), at a repeat frequency of around 80 MHz. These pulses pass through a ‘chopper’, which modulates the pulse train for the lock-in amplifier. The choppers operate at up to 5 kHz, but are typically used at around 200 Hz. The pulse train therefore is modulated into 2.5 ms sections of pulses followed by 2.5 ms of nothing.

Each individual laser pulse is converted into a pulse of broadband radiation with spectral content ranging across the low terahertz frequencies (typically 0.5 to 5 THz, where 1 THz corresponds to a wavelength of 300 μm and a period of 1 ps) either through optical rectification using a nonlinear crystal (for example Ga:As), or through a dipole antenna [20, 70]. In this way, a terahertz pulse train is generated, with 2.5 ms of terahertz pulses ('signal') followed by 2.5 ms of no signal, with each 'signal' part containing around 200,000 terahertz pulses.

The terahertz pulse interacts with the sample in some way, and then is focused on an electro-optical sampling (EOS) crystal, for example $\{110\}$ Zn:Te, for detection. The incident terahertz field creates an instantaneous birefringence in the Zn:Te, which is measured with a near-infrared beam. This probe beam is circularly polarised with a quarter waveplate, before the birefringence modulates how elliptical this polarsation is. The difference between the vertical and horizontal components are measured using a Wollaston polarization splitting prism, and two balanced photodiodes. This balanced detector employs common mode noise rejection by subtracting the 'vertical' signal from the 'horizontal' signal. If no modulation has occurred (i.e., no terahertz radiation has fallen on the crystal) then the result of this subtraction is zero. If some modulation of the elipticity has occurred, then the difference will be proportional to this modulation, which is in turn proportional to the incident terahertz radiation. Thus the difference on the balanced detector will be directly proportional to the radiation incident on the detector crystal. In this way the actual electric field is recorded. In practice, the original laser beam is split to provide both the pump beam and the probe beam.

Each laser pulse lasts only femtoseconds, whereas the terahertz pulse has a duration of picoseconds, and certainly a pulse may have been delayed many tens of picoseconds by its interaction with the sample. Thus the detected signal is a snapshot of the terahertz electric field at that (femtosecond) instant. In order to obtain the electric field of the entire pulse, an optical delay stage is used, which lengthens or shortens the path of the probe beam compared to the path of the terahertz pulses. This varies the time point at which the

Figure 3.2: Photograph of the terahertz imaging system at Leeds.

snapshot of the terahertz pulse is taken, and enables an entire time series to be collected. The delay stage must obviously be positioned in such a way that a meaningful window into the data is achieved, and this is position is known as the initial displacement – a value which is essentially arbitrary to a given system.

The final stage is the lock-in amplifier (LIA). This is a device which selectively amplifies an incoming signal at a given frequency, in order to boost it above the noise. The chopper makes the terahertz signal a square wave modulated at the chopping frequency, whereas the general noise is not so modulated. The LIA requires a time-constant to be set. This time constant dictates how long the LIA will spend acquiring each time-point, and a typical value in the set-up described is 200 ms. Over 200 ms, the LIA will acquire 40 values, since it will acquire a single value from each 5 ms ‘period’. It will average all of these to create a single value, after which the time-delay stage would typically move to the next part of the pulse.

The terahertz pulsed imaging machine is a physical device built on an optical bench, using mirrors, lenses, and other optical components. The equipment used is very delicate and must be positioned exactly — any misalignment causes a corruption of the data. Poorly aligned optics, for example, cause part of the signal to be undetected. If the setup remains unchanged throughout an experiment then this does not really cause a problem — it is mainly relative values and ratios that are used. However it does make direct comparison between scans acquired at different times unreliable. Figure 3.2 is a photograph of the optical part of the first terahertz pulsed imaging system at Leeds which shows the complexity of the system. This system was used May 2001 to November 2003.

Finally, Figure 3.3 shows example terahertz pulses transmitted through varying thicknesses of nylon, and the corresponding power spectra. The pulses are very well localised in time, and have a broadband frequency content. Notice that in the time domain the pulses are translated, attenuated, and dilated by vary amounts depending on the thickness of material. The power spectrum shows the pulse frequency content is centered around 1 THz, and that the higher frequency content is attenuated more by nylon than the lower frequencies. This is also suggested in the time domain — the pulse through 1 mm is a lot smoother, i.e., devoid of high frequency content.

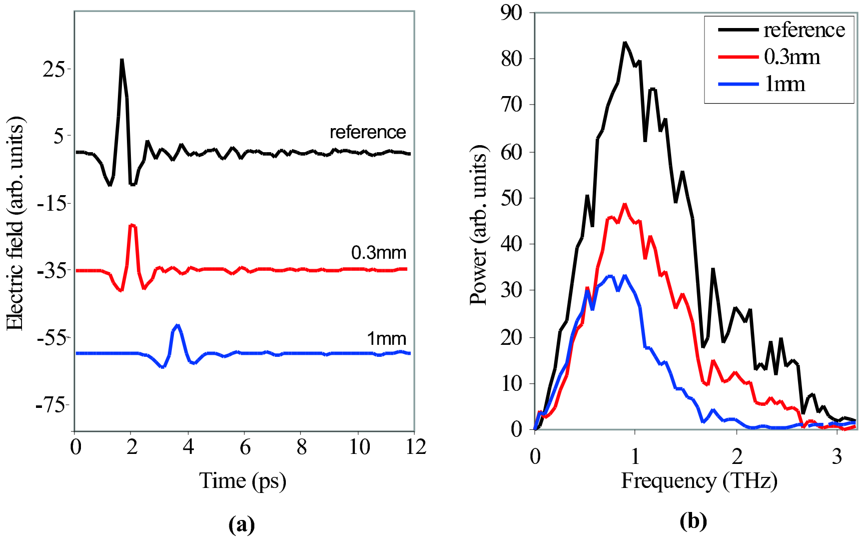

(a) **(b)**

Figure 3.3: Examples of terahertz pulses in (a) time domain, and (b) magnitude of Fourier coefficients.

3.3 A review of time and frequency analyses

3.3.1 The Fourier Transform

The Fourier transform (FT) needs no introduction as a tool for analysis, and full discussion of it is beyond the scope of this thesis. However it is still the most important signal processing tool, and as such warrants a brief overview. It is also a necessary background for the understanding of the STFT.

A real valued periodic function $f(t)$, with period T has Fourier representation

$$f(t) = \sum_{-\infty}^{+\infty} a_k e^{ik\omega t}$$

where $\omega = 2\pi/T$ is the *fundamental frequency* and the Fourier coefficients are given by

$$a_k = \frac{1}{T}\int_{t_0}^{t_o+T} f(t)e^{-ik\omega t}dt$$

This representation provides a decomposition of the function into frequency harmonics, whose contribution is given by the coefficients a_k.

For a non-periodic function, the Fourier Transform of $f(t)$, and its inverse are given by

$$\hat{f}(\omega) = \int_{-\infty}^{+\infty} f(t)e^{-i\omega t}dt \tag{3.1}$$

$$f(t) = \frac{1}{2\pi}\int_{-\infty}^{+\infty} \hat{f}(\omega)\, e^{i\omega t}d\omega \tag{3.2}$$

The Fourier transform may be discretised and applied to signals that have been discretely sampled (such as terahertz pulses). Each coefficient is a complex number whose amplitude represents the power of that frequency component, and whose angle represents the phase of that frequency component modulo 2π.

The theory of Fourier series and transforms is described in more depth elsewhere (for example [29]), and their application to image and signal processing is also described elsewhere (for example [59]).

3.3.2 Short Time Fourier Transform

Although the Fourier transform is the standard spectral analysis technique, it performs poorly at analysing non-stationary signals, since the frequency content is considered over all time. The short-time Fourier transform (STFT) was therefore developed to overcome this limitation. In STFT analysis, the signal is windowed using some window function $\phi(t)$ before the Fourier transform is applied. The complete transform is then acquired by translating this window along the signal, applying the Fourier transform to this windowed signal at each location. In this way a 2-D transform is created, with values defined at translation points β and spectral points ξ. These broadly correspond to t and ω respectively, although the correspondence is not exact due to the uncertainty principle. The STFT of a function $f(t)$ with respect to the window function $\phi(t)$ evaluated at the location (β, ξ) is therefore

$$STFT_{\phi} f(\beta, \xi) = \int_{-\infty}^{\infty} f(t) \phi_{\beta,\xi}^{*}(t) dt \tag{3.3}$$

where

$$\phi_{\beta,\xi}(t) = \phi(t - \beta) e^{j\xi t} \tag{3.4}$$

The window function has fixed time and frequency resolution, however, which is still a shortcoming [9]. The trade-off between the time resolution and frequency resolution is achieved through the width of the window function - the variance in the case of the Gaussian window function. Narrow window functions trade good temporal resolution against poor frequency analysis — an infinitesimal window width is time-domain analysis. Wide window functions trade good spectral resolution against poor temporal resolution — an infinite window width is frequency-domain analysis (an infinite windowed STFT is in fact a Fourier transform).

The choice of window function will also have an impact on the analysis. We might expect a window with sharp edges — for example a rectangular window — to have discontinuities in the frequency domain, and these will reflect in the analysis. On the other hand, a smooth window function – such as a Gaussian — will have a smoothing effect on the analysis (see Figure 3.9).

Windowing the pulse using a rectangular function at a fixed β also simulates the acquisition of fewer samples at the same temporal resolution, for instance acquiring only 10 samples either side of the response's peak. If fewer recorded samples give accurate results for the absorption coefficient, then the acquisition time can be shortened accordingly.

3.3.3 Application

Figures 3.4 – 3.6 show the application of the Fourier transform and the STFT to a variety of test signals. These figures show the strengths and weaknesses of these two techniques.

Figure 3.4 shows the analysis of two simple sine waves; one at 2.5 Hz, and one at 25 Hz. The FT precisely locates these frequencies and the STFT identifies the centre frequency, although the range of frequencies included is larger.

Figure 3.5 shows the analysis of two different combinations of the sine waves from Figure 3.4. These figures show up the limitation of the FT, since the power spectrum for both combinations looks essentially the same. The reason for this can be seen in (3.1) — the terms of the integral are over all time. On the other hand, the STFT correctly identifies the switch over between the 2.5 Hz and the 25 Hz section of the signal.

Finally, Figure 3.6 shows the analysis of a 'chirp' signal, where the frequency is linearly increasing between 0 Hz and 40 Hz over the 2.5 seconds of signal. This type of signal can typically be caused by a rotating device (such as an engine) starting up and spin-

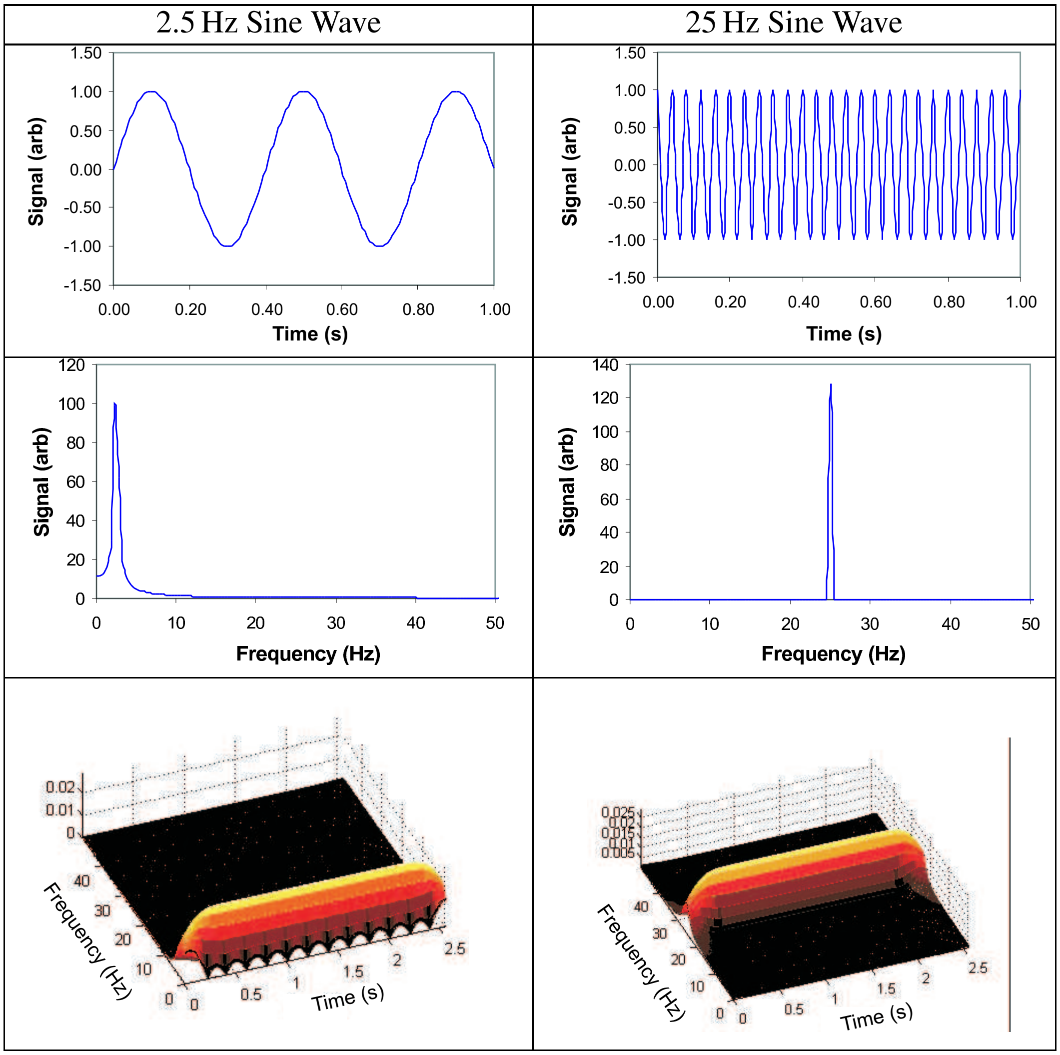

Figure 3.4: Signals presented in different domains; the time domain (top), the power spectrum from the FT (middle), and the spectrogram from the STFT (bottom).

ning up to speed. The FT correctly identifies that there is a range of frequencies present, but nothing beyond that. The STFT on the other hand shows the frequency changing linearly.

The effect of changing the width of the STFT windowing function can be demonstrated using the signal with the 2.5 Hz wave followed by the 25 Hz wave. Figure 3.7 shows the signal being analysed with both a relatively wide window (the top figure) and a relatively narrow window (the bottom figure). Note that the viewing angle of the surface has changed between the two figures in order to highlight the differences. With the wide

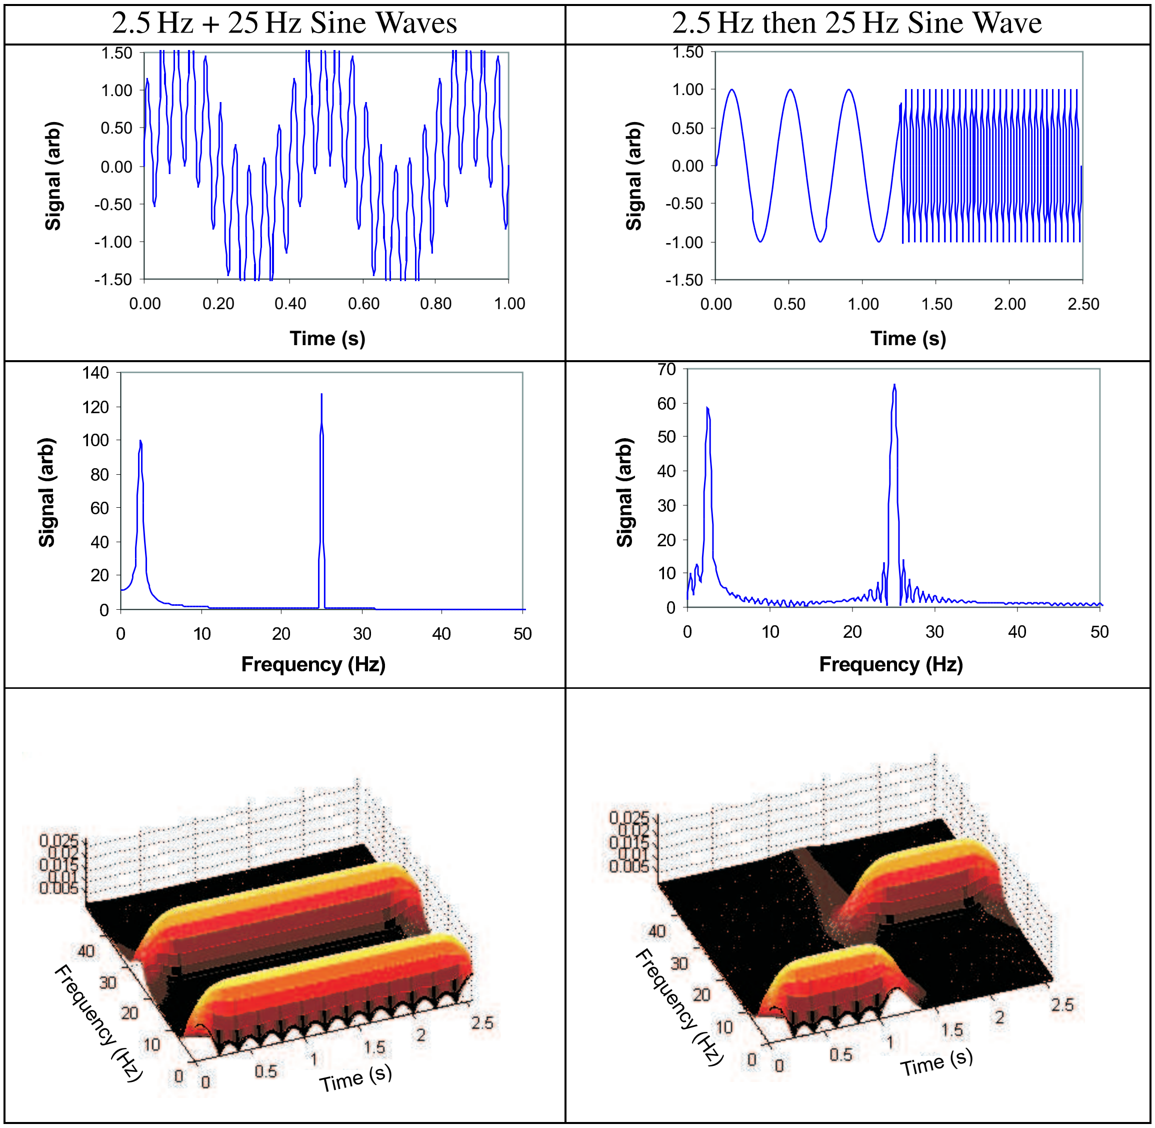

Figure 3.5: Signals presented in different domains: the time domain (top), the power spectrum from the FT (middle), and the spectrogram from the STFT (bottom).

window, the frequency resolution is very good, but the time resolution is poor, making it very hard to tell when the signal switches frequency. On the other hand, with the narrow window, the time resolution is excellent, and the distinction between the frequencies is obvious. However the frequency resolution is very poor, providing very little information about the frequency content.

The other parameter to consider is the window function. Figure 3.9 shows the effect of analysing the chirp from Figure 3.6 with a Gaussian window, a triangular window, and a rectangular window. For examples of these window functions, see Figure 3.8. The rectangular window function very clearly shows discontinuities in the frequency domain, the triangular window has them to a lesser extent, and the Gaussian window does not have any such artefacts. Note in these figures the intensity denotes the magnitude of the STFT coefficients, with black showing the largest magnitude, and white the smallest.

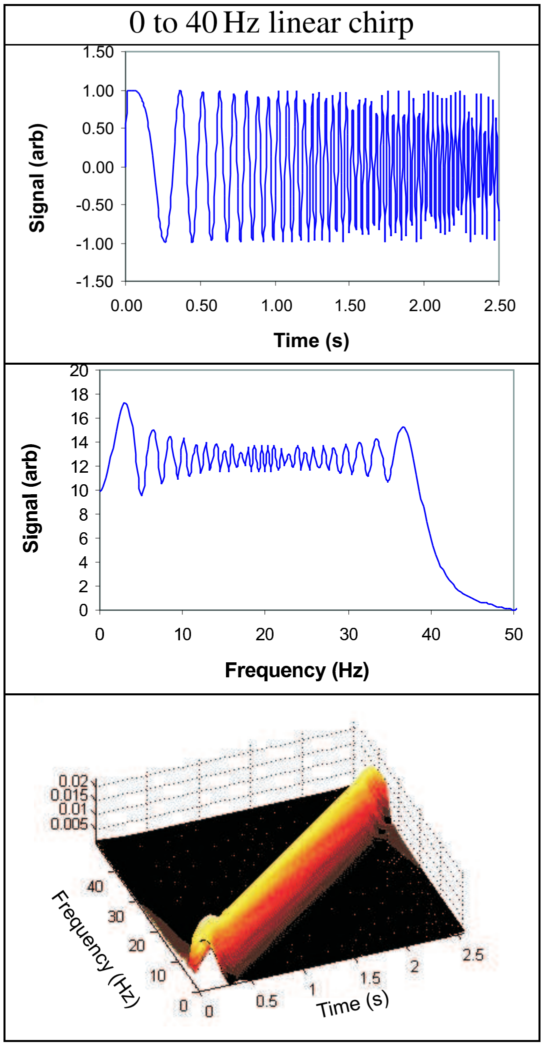

Figure 3.6: A signal presented in different domains: the time domain (top), the power spectrum from the FT (middle), and the spectrogram from the STFT (bottom).

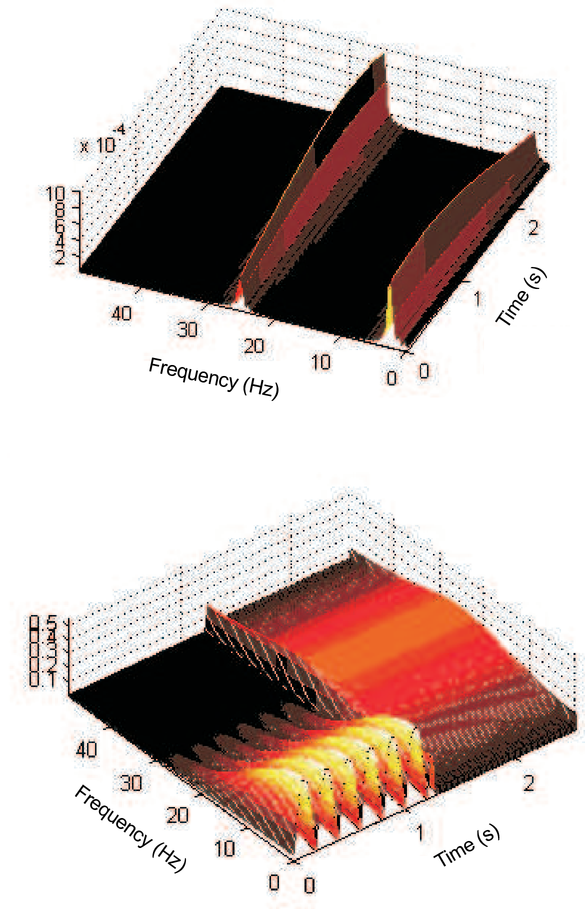

Figure 3.7: The STFT of a 2.5 Hz wave followed by a 25 Hz wave, with Gaussian window function of σ 100 s (top) and 100 ms (bottom). Note that the viewing angle changes for clarity.

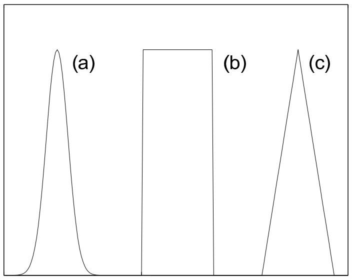

Figure 3.8: Examples of window functions — (a) Gaussian, (b) square, (c) triangular. All are shown with equal 'width', and have been normalised to the same height for ease of comparison.

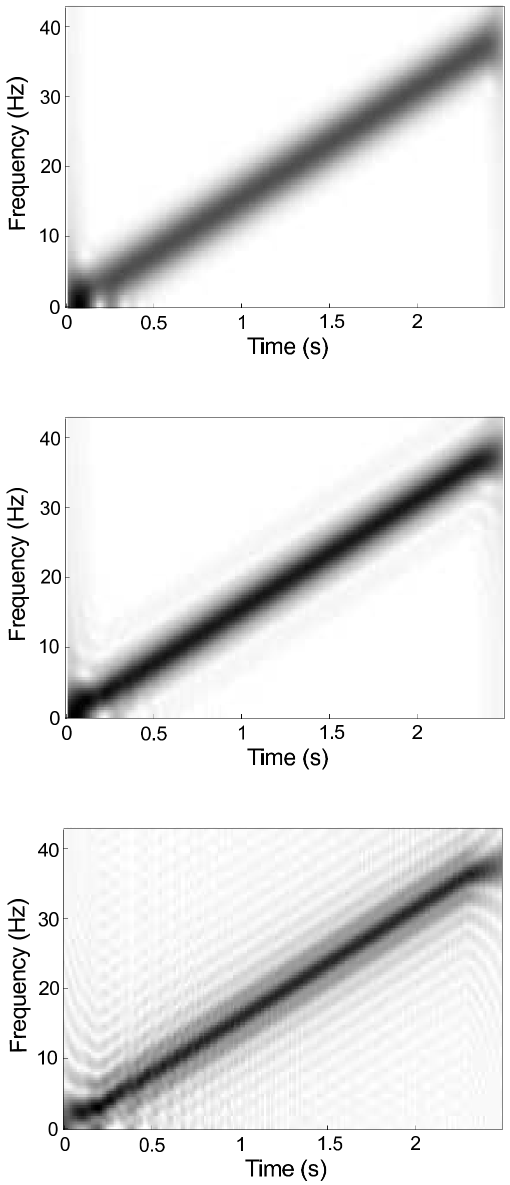

Figure 3.9: The introduction of frequency artefacts by using discontinuous window functions in the STFT analysis of a chirp signal. The coefficients are shown from analyses using (a) a Gaussian window function, (b) a triangular window function, and (c) a rectangular window function.

3.3.4 Wavelet Transforms

The continuous wavelet transform (CWT) of a 1-D function $x(t)$ with the mother function ψ is defined as

$$CWT_{\psi}x(\tau,\sigma) = \frac{1}{\sqrt{\sigma}} \int x(t)\,\psi^{*}\left(\frac{t-\tau}{\sigma}\right) dt \tag{3.5}$$

where τ and σ are the translation and scale parameters respectively. This corresponds to a correlation between the input signal and scaled/translated versions of the mother function. In this way a 2-D plot in translation/scale space is obtained, with translation being directly related to time, and scale being inversely related to frequency. It is not possible to define an exact relationship between translation and time, because each translation value actually corresponds to a time window, the width of which is dependent on scale. Similarly each scale value corresponds to a frequency window, the width of which also depends on the scale. Thus translation corresponds to a range of times, and scale corresponds to a range of frequency. This is an inevitable consequence of uncertainty, and the strength of wavelet analysis lies in the optimisation of these window widths. The centre of these windows is known however, as the centre of the time window is directly proportional to translation, and the centre of the frequency window is inversely proportional to scale. The width of the time window is directly proportional to scale (i.e., high frequencies, which correspond to small scales, are well localised in time), whereas the width of the frequency window is inversely proportional to scale (i.e., low frequencies, which correspond to large scales, are well localised in frequency).

Discrete Wavelet Transform

The CWT is straightforward to discretise, for implementation on a computer, however an efficient transform called the Discrete Wavelet Transform (DWT) may also be used.

The DWT has a similar expansion to the Fourier transform, defined as

$$f(t) = \sigma_j \sigma_k a_{j,k} \phi_{j,k}(t) \tag{3.6}$$

where j and k are integers and the functions $\phi_{j,k}(t)$ are the wavelet basis functions — these usually form an orthoganal basis. $a_{j,k}$ are then the DWT coefficients of $f(t)$, and are calculated using

$$a_{j,k} = \int f(t)\phi_{j,k}(t)dt \tag{3.7}$$

As with the CWT, the wavelet basis functions are a two-parameter family of functions related to a mother function thus

$$\phi_{j,k}(t) = 2^{j/2}\phi(2^j t - k) \tag{3.8}$$

k and j are called the translation and dilation parameters respectively, and so the wavelet basis is obtained from a single mother function through translating and scaling.

3.3.5 Wide Band Cross Ambiguity Functions

In the case of terahertz pulsed imaging, however, the interest is in the relative time-delay and spectral changes of a pulse compared with its reference pulse. We can extract these relative differences by using a cross ambiguity function between the sample pulse and the reference pulse [68, 64]. The wide-band cross ambiguity function (WBCAF) is defined as

$$WBCAF_{x_1}x_2(\tau, \sigma) = \frac{1}{\sqrt{|\sigma|}} \int_{-\infty}^{\infty} x_2(t) x_1^* \left(\frac{t - \tau}{\sigma} \right) dt \tag{3.9}$$

where $x_1(t)$ is the reference waveform, and $x_2(t)$ is the delayed and attenuated sample waveform. The similarities between this and the CWT of (3.5) are clear — we are effectively using our reference pulse as the mother function. It has been noted [49] that terahertz pulses exhibit similar properties as wavelets, namely the compact support and broadband content. The basis of the WBCAF will be non-orthogonal, and hence the transform will have redundancy. However we are interested in small changes to the scale parameter, so an orthogonal basis, which typically use dyadic scales, would not provide the resolution of scale that was hoped for.

In practical terms, the different scales are achieved by re-sampling the reference pulse, through digital interpolation, filtering, and decimation [13]. The filtering step is necessary to prevent aliasing and other artefacts as the sample rate is modified, and is simply a low-pass digital filter.

3.4 Clustering

Clustering is the technique of grouping 'similar' *n*-dimensional data points together in order to partition a dataset. Clustering techniques fall broadly into two categories — supervised and unsupervised.

Supervised clustering uses extensive training data to create the clusters, and subsequent data points are allocated to the most appropriate groups. With large data sets and good exemplars, supervised clustering is the best method to use, and examples range from simple linear discriminators, such as applied by Ferguson et al.to terahertz data [26], to trained artificial neural networks such as back propagation [22] and support vector machines [12].

Unsupervised clustering, on the other hand, is appropriate where either there is insufficient data for training, little *a priori* knowledge, or where the outcome (number of groups, exemplars, and so on) is not known in advance. In this approach 'similar' pixels are grouped together to create the clusters by minimising a cost function (the distance of each data point from its cluster's centre, for example). Examples of unsupervised clustering include *k*-means [33] and self organising networks [22, 43]. Variations include region merging and region splitting [33]. Of these, *k*-means clustering is perhaps the most well understood and established algorithm, and has been applied to terahertz pulsed imaging data in Chapter 6.

3.4.1 *K*-means Clustering

The *n*–dimensional data are considered to be vectors that form points in *n*–dimensional space. A standard distance metric, such as Euclidean distance, is used to assess similarity, where a small distance between two data points indicates a large similarity. Each cluster is typically described by the location of its centroid in the *n*–dimensional space, one such measure being the mean of the position of all its members. In *k*-means clustering clusters are formed by initialising the desired number of cluster centroids and then adding each data point to its 'nearest' centroid (which in turn updates the position of the centroid), until all the pixels have been clustered. The *k*-means algorithm is shown in Figure 3.10. This algorithm is deterministic for a given initialisation.

1. Initialisation:	Use some initialisation metric to establish the initial positions of each cluster.
2. Allocate points to cluster:	Allocate each data point to the cluster that is 'nearest'.
3. Update the cluster positions:	Calculate the new position of each cluster based on the mean position of all that cluster's members.
4. Repeat until finished:	Repeat **2** and **3** until the stopping criteria are met. Examples of stopping criteria are 'no more points change cluster' or 'no cluster centroid moves by more than a threshold'.

Figure 3.10: Standard *k*-means clustering.

When the *n*–dimensional space is formed from domains which are not directly comparable, care must be taken to ensure that differences in scale or units do not bias the outcome. In this case, feature vectors should be normalised to be uni-variate within a unit hypercube [33].

3.5 Summary

In this chapter mechanics of terahertz pulsed imaging have been reviewed in some detail, which will be of use throughout this thesis, but especially in chapter 4 when considering the noise present in terahertz data. We have looked at time/frequency analysis techniques, starting with traditional Fourier transforms, moving through short time Fourier transforms, and then reviewing wavelet transforms and cross ambiguity functions. These techniques are used in chapter 5 in the interpretation of terahertz pulsed imaging data. Finally, the areas of clustering were reviewed, specifically k-means clustering which is applied to terahertz pulsed data in Chapter 6.

Chapter 4

Noise in pulsed terahertz systems

4.1 Introduction

Undesired data and the corruption of signals, or ‘noise’, are issues that affect all ‘real world’ signals, and all signal processing applications must take it into account. However, before the noise can be reliably removed or allowed for, it must be characterised or modelled. Terahertz pulsed imaging is no exception to this rule, and in this study I estimated the noise present in a number of terahertz data-sets and found distributions that model the underlying distribution of the noise. The parameters of the distribution not only characterise the noise but could also aid the design of algorithms that either remove the noise or are robust against it — a process that requires good noise models [3, 6, 67]. A noise model may also enable synthetic images to be created with realistic simulated noise.

In this chapter distributions to model the noise in pulsed terahertz imaging are empirically built, with the aim of providing a model that may be used in both signal processing and in the creation of realistically noisy synthetic terahertz images. After a brief analysis of the theoretical noise sources, the noise is extracted from a number of terahertz data sets that demonstrate different noise characteristics from each other, before finding the distribution that best describes the noise in each case. I am exclusively considering noise in the terahertz time series, and not image noise.

Finally I consider the wavelet shrinkage denoising algorithm [17] applied to terahertz imaging [25, 24], assessing what level of denoising may be applied before statistically significant errors are introduced in the optical parameter calculations. More details on calculating the optical parameters may be found in Chapter 5.

4.1.1 Noise sources in terahertz imaging

In order to break down the noise components of terahertz pulsed imaging and aid the analysis, the noise can be split into two categories, signal noise and imaging noise.

Signal noise has three main components; Johnson noise, shot noise, and set–up noise, where set–up noise is that which is due to the atmospheric conditions, sample type, background radiation, set up of the machine, and so on. The signal noise is concerned with noise that has deformed the terahertz pulse in some undesired way, as opposed to noise which is an artefact of the imaging process. These are all discussed in more detail below.

Imaging noise, on the other hand, is noise that is introduced through the imaging process. For example the incoming analogue signal is sampled onto the computer, a process that invariably results in a digital approximation of the original analogue signal. Similarly, signal processing or rounding errors in the computers will introduce further errors. Finally when an entire image is being formed, if the raster scanner is misaligned this will result in pixels in the final image being in the ‘wrong place’. Of course the signal noise will impact the generation of images — a large amount of shot noise may be apparent in attenuation plots, for example — but I am not concerned with these effects of noise in this study.

Johnson noise

Johnson noise is due to the electron equivalent of Brownian motion. To quote Demir [15];

> "Random fluctuations in the short-circuit terminal of an arbitrary linear resistor, having resistance R and held at temperature T, are independent of all other parameters, and can be modelled with a wide sense stationary stochastic process with a spectral density given by
>
> $$S_i(f) = \frac{2kT}{R} \tag{4.1}$$
>
> where k is Boltzmann's constant."

The book further states that

> "It can be shown that a thermal noise of a linear resistor as a white wide sense stationary stochastic process with spectral density (4.1) is accurately modelled by a Gaussian process, ... , as a direct consequence of the central limit theorem."

This noise is present even when the laser is switched off. There is also background noise present, due to the abundance of incoherent terahertz radiation — black body radiation at 300 K is at 6.25 THz. However, the gated and coherent nature of the terahertz receiver makes it reject background noise, meaning this noise is a factor of 7 times smaller than the Johnson noise [61].

Shot noise

Existing studies have categorised the lasers and their noise sources [14, 54], and the main noise source that results from the laser generation process is shot or quantum noise [61].

When the photons of the laser hit the detector, they give rise to a constant photocurrent, as if a bias voltage had been applied. The mechanism for this is the induction of a charge transfer and hence random effective voltage in the photo-diodes of the detector. However, since the interactions of the photons (including their emission and detection) are stochastic processes, there will also be a 'quantum mottle' effect, leading to random fluctuations in the charge transfers. This noise is about the same order of magnitude as the Johnson noise, and can be modelled using a wide sense stationary Gaussian process [15, 61].

Other sources of noise specifically concerned with the laser include laser timing jitter and pulse energy fluctuation [14, 54], but these have an insignificant impact on the signal compared with the shot noise.

Previous studies have shown that the dominant noise source in these categories is the shot noise, giving a signal to noise ratio in the order of 100:1 [49].

Set–up noise

There are many sources of physical noise in terahertz imaging. Background radiation has been mentioned above, but in addition to this atmospheric conditions have a large impact on the signal acquired. The main factor of this is humidity, because terahertz radiation is so strongly absorbed by water molecules. Water molecules also absorb and re-emit radiation at this frequency. This type of activity is not noise in the traditional sense as it is deterministic if the conditions do not change. However varying conditions would make direct comparisons unreliable between pulses acquired on different days or over a long period of time.

If a sample is present, there will also tend to be a scattering effect. There are many different forms of scattering (Rayleigh, Raman, Brillouin, and so on) depending on the type of molecule, and the wavelength of the radiation. The effect this has on terahertz pulses is two-fold; firstly it introduces a loss as some radiation will be scattered away, and secondly some spectral content may end up phase shifted, as it takes a longer route through the material. The modelling of scatter of terahertz pulses through various materials is an ongoing research topic, but for the purpose of this analysis it is ignored. Techniques such as Monte Carlo modelling are both complex and computationally expensive, and the empirical nature of this investigation will automatically include the effects of scattering.

The other form of imaging noise arises from losses due to absorption and reflection. Absorption and reflection are what make imaging possible (if the radiation didn't interact with the material, there would be nothing to see), however inadequately allowing for reflection losses can cause errors in other calculations, such as when deriving the optical properties of the sample.

The main impact these losses have on this study is that a decreased level of radiation will be transmitted through the material. This will tend to decrease the signal to noise ratio, making it harder to be certain about what constitutes signal and what constitutes noise. Similarly the balance of the noise component will be shifted towards noise that is dependent on the detection technology and away from noise caused by the generation process — this latter noise will tend to be modulated along with the signal.

Imaging noise

As discussed in Chapter 3, the terahertz pulsed imaging machine is a very delicate system requiring careful optical alignment. Inevitable variations in the optics make direct comparisons between machines unreliable, and indeed scans acquired on different days may not be directly comparable. Similarly, the lock-in amplifier can drift over time, causing a d.c. bias to be introduced in the data acquired.

At an imaging level, as opposed to the pulse level, raster scanning errors can cause straight edges to appear staggered. The motors moving the sample may not displace it correctly, leading to a varying physical distance between pixels on the image.

It is expected that all these noise sources are insignificant and rectified as far as possible during the set-up and calibration of the scanner.

Finally, as the analogue waveform is captured to and processed on a computer, there is the potential for the introduction of further errors. Similarly computers have only finite accuracy on mathematical operations, and each time an operation is performed some accuracy may be lost. These are both unavoidable. It is assumed that appropriate sample rate, dynamic range, and accuracy values are chosen, but these still need to be considered when giving an error margin.

4.1.2 Summary

In summary, the main theoretical sources of pulse noise in terahertz imaging are Johnson noise and shot noise, both of which may be modelled as wide sense stationary Gaussian processes.

However the additional noise sources make direct comparisons between scans unreliable, and the effects of scattering and atmospheric conditions may contribute significantly to the noise levels, altering the underlying noise distribution from the simple Gaussian processes.

4.2 Methods

The noise models were built in a three phase process. Firstly noise was estimated from measured signals. Secondly various models were fitted to these noise estimates, using standard distributions. Finally each model was evaluated to see how good a fit was obtained, in order to assess if it indeed modelled the distribution. It was hoped that a generic model for noise would emerge through this process.

4.2.1 Estimating the noise

The noise was estimated using three different methods. The first was simply to block all the terahertz radiation from falling on the detector crystal. In this way, the entire signal recorded was taken as noise, as there was no terahertz radiation present. These pulses were acquired at the University of Leeds using a displacement of 22.5 mm, acquiring 256 time points, with a separation between time points of 80 fs. The time constant was set at 50 ms, 200 ms, and 500 ms in order to acquire three pulses. The LIA delay was twice the time constant in these instances. The variation in LIA time constants was to explore the effect that this setting has on noise.

The second method was to record free air scans several times under identical conditions. With no sample in place, any variation between the scans is entirely due to receiver, transmitter, and 'set–up' noise such as atmospheric conditions. Sixteen pulses were recorded at Leeds through free air using identical acquisition parameters within a short space of time. The displacement was 22.5 mm, and again 256 steps were acquired at 80 fs separation. The time constant was 100 ms, with a LIA delay time of 300 ms. The mean pulse in the time domain was calculated and assumed to represent a pulse free of random noise. The differences in the time domain between each of the pulses and the mean pulse was then used to create the noise distribution.

It was noticed that there appeared to be a correlation between the amplitude of the electric field and the level of noise in the free air scans. This was explored by plotting the root mean square difference between the mean pulse and the individual pulses at time t against various derivatives of the amplitude of the mean pulse at time t. Firstly the amplitude was used, followed by the absolute amplitude. Finally an 'historical absolute sum' over n (HAS_n) was used, where

$$\mathrm{HAS}_n(t) = \sum_{i=(t-(n-1))}^{t} |meanpulse(i)|. \tag{4.2}$$

The HAS was chosen as a metric because the detector cannot predict the future (what the next value of the electric field will be), but it might 'remember' the recent past (what the previous value, or values, were.) The value over which to sum, n, was chosen empirically based on which gave the best correlation.

The final method of estimating the noise was similar to the previous method, except that the pulses were recorded through homogeneous regions of an object. Two step wedge phantoms were manufactured — one ‘nylon’ and one ‘resin’ — with the step depth profiles shown in Table 4.1. The nylon wedges were made from Duraform polyamide (nylon 12) by the selective laser sintering process. The resin wedges were made by stereolithography in SL 5190 photopolymer resin from 3D Systems[1]. The wedges were made at REACT, Rotherham in Jan/Feb 2001. Figure 4.1 shows an idealised cross section of one of these wedges.

Nylon	Resin
0.3 mm	0.1 mm
1 mm	0.5 mm
2 mm	1.5 mm
4 mm	3 mm
6 mm	5 mm
7 mm	7 mm

Table 4.1: Depth profiles of the step wedge phantoms

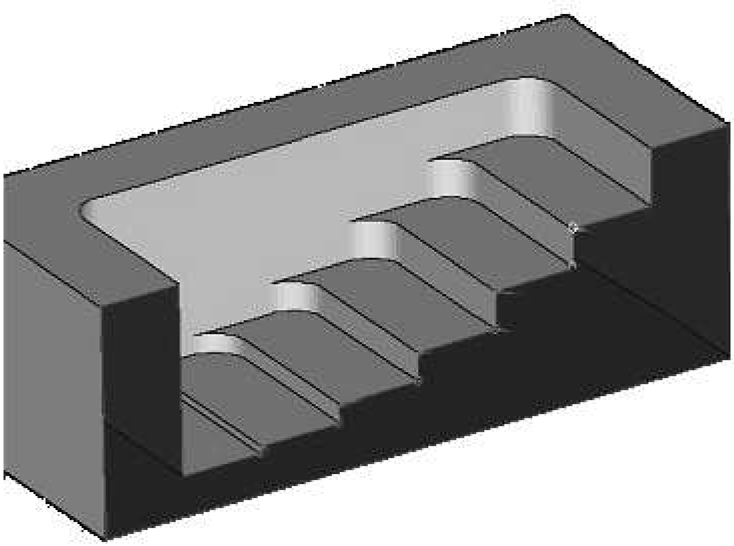

Figure 4.1: Idealised cross section of a step wedge phantom.

These phantoms were imaged in transmission modality on the pulsed terahertz imag-

---

[1] 3D Systems were supplied by Ciba Geigy at the time. Ciba Geigy have subsequently been taken over by Vantico (www.vantico.com).

ing systems at Physikalisches Institüt, J.W. Goethe-Universität, Frankfurt. Table 4.2 shows the acquisition parameters. The reference pulse in both cases was a terahertz pulse through free air.

Phantom	Nylon	Resin
Spacing of time points (fs)	150	200
Number of time points / pixel	128	128
Displacement	-32.5	-34
Number of pixels (x × y)	100 × 4	100 × 8
x range (mm)	-1 to 24	-2 to 23
y range (mm)	-3 to -1	0 to 4

Table 4.2: Acquisition parameters of the step wedge phantoms

A region of interest was then defined within each step, taking care to avoid the edges to minimise boundary effects, and the pulses within that region were assumed to have passed through the same thickness of homogeneous material. Around 40 terahertz pulses were then recorded through each region of interest of the nylon step wedge, with around 80 recorded similarly through the resin step wedge.

The average pulse from within each region was calculated, as before, and used as an estimate of a pulse free of random noise.

4.2.2 Building and fitting the model

The theoretical noise analysis suggests that the noise sources are wide sense stationary processes, and therefore should be distributed in a zero-mean approximately normal distribution. These types of distribution belong to the stable family, an infinite set that includes Gaussian, Cauchy, Lévy, and Gaussian mixtures [23, 39]. Broadly speaking the family of stable distributions are heavy tailed, skewed, or both heavy tailed and skewed.

In each case the first step was to plot the noise histogram whereby the approximate normality of the distribution could be visually confirmed.

Stable Distributions

Stable distributions are probability distributions that allow skewness and heavy tails, and were characterised by Paul Lévy in the 1920s. The family of distributions is exclusively such that if random variables x and y have stable distributions, then the random variable $x+y$ will also have a stable distribution. The Gaussian distribution is a special case of this family, as it is the only such distribution with a finite variance [23]. There are no closed formulae for densities and distribution functions for all the stable functions except Gaussian, Cauchy, and Lévy, which make them hard to use in practice. However there are now computer programs to compute stable densities, distribution functions and quantiles. I used Nolan's "stable" program in this work [51].

As stable distributions generally have no closed form, they are typically expressed through the characteristic function of a stable random variable X. In this representation, four parameters are required to describe a general stable distribution: α, β, γ, and δ. $\alpha \in (0,2]$ is the index of stability, $\beta \in [-1,1]$ is the skewness parameter (where $\beta = 0$ means the distribution is symmetric), $\gamma > 0$ is the scale parameter, and $\delta \in \Re$ is the location parameter. γ and δ are the stable equivalent of the standard deviation and variance of a Gaussian distribution. The shape of the probability density function is determined by α and β, and certain values give rise to specific distributions within the family, namely

$\alpha = 2, \beta = 0$ give Gaussian distributions,
$\alpha = 1, \beta = 0$ give Cauchy distributions, and
$\alpha = \frac{1}{2}, \beta = 1$ give Lévy distributions.

A stable random variable X has the characteristic function

$$E\{e^{iuX}\} = \begin{cases} e^{-\gamma^{\alpha}|u|^{\alpha}\ (1+i\beta(\tan\frac{\pi\alpha}{2}(\text{sign}(u))+i\delta u} & (\alpha \neq 1) \\ e^{-\gamma^{\alpha}|u|^{\alpha}\ (1+i\beta\frac{2}{\pi}(\text{sign}(u))ln(|u|))\ +i\delta u} & (\alpha = 1) \end{cases} \tag{4.3}$$

where

$$\mathrm{sign}(u) = \begin{cases} -1 & (u < 0) \\ 0 & (u = 0) \\ 1 & (u > 0) \end{cases}$$

Figure 4.2 shows some examples of stable distributions, together with the effects of varying the parameters.

Fitting the model

The stable family of distributions is difficult to use analytically, so before fitting a stable distribution the noise was first tested against a Gaussian model and a Gaussian mixture model, or GMM — a model built from the simple addition of two or more Gaussian distributions.

The noise distributions were all well sampled, and so the mean and standard deviation of this noise were used for the parameters of the hypothetical Gaussian distribution.

The Expectation Maximisation (EM) algorithm [16], as shown in Figure 4.3, was used to build the GMMs. This is a two stage iterative process that calculates the local maximum likelihood fit of a given arbitrary number of Gaussians to a data set. Although EM is sensitive to initialisation in many applications, in this context the distribution being fitted is a heavy tailed normal distribution, which means the Gaussians that fit it will have approximately the same mean. This in turn means the fit will not particularly depend on the initialisation metric. The Gaussians were initialised to have equal weighting and to be centered on the mean of the data. The standard deviation of the first Gaussian was initialised to the standard deviation of the data, and subsequent Gaussians initialised with increasingly larger standard deviations. EM can also have difficulties with sparse data sets causing singularities – however in this application the data are not sparse so there is no problem here. The algorithm stops when the parameters of the Gaussians remain constant for an iteration.

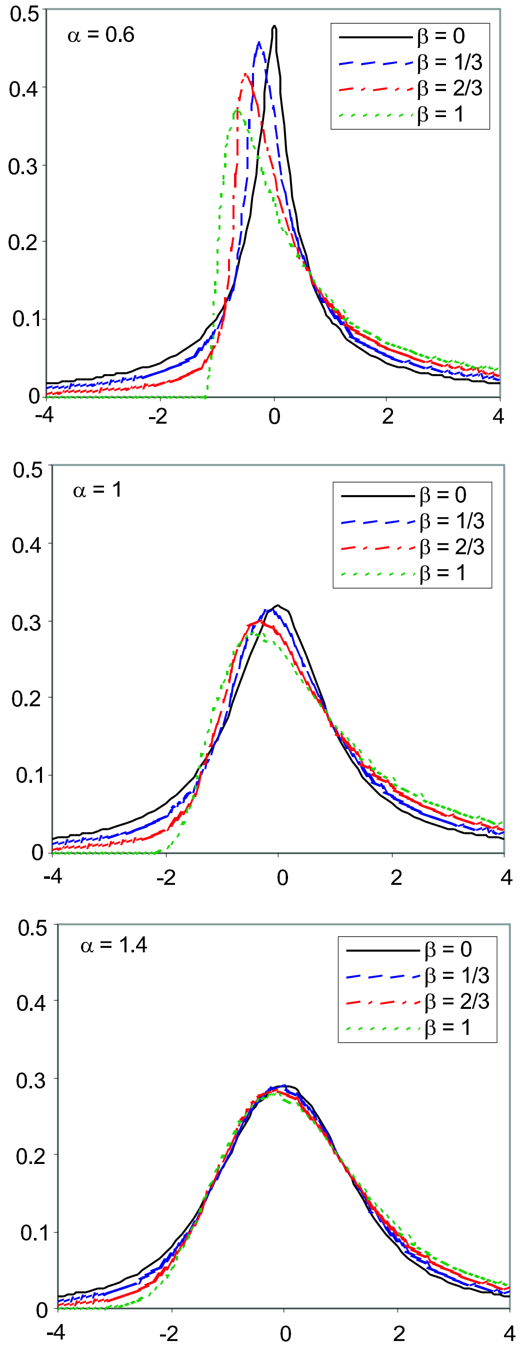

Figure 4.2: Some example probability density functions of stable distributions.

An alternative approach to building GMMs is to use a 'kernel method' [58], whereby each data point is modelled by a single Gaussian. However this approach is computationally expensive, and while it would provide a good model for the noise, it would not provide insight into the underlying noise distributions.

E-step: Compute the contribution of the i_{th} sample to the j_{th} Gaussian.

$$p_{ij} = \frac{w_j N(x_i : \mu_j, \sigma_j^2)}{\sum_{k=1}^{M} w_k N(x_i : \mu_k, \sigma_k^2)}$$

M-step: Compute the parameters of the Gaussians.

$$w_j = \frac{1}{n} \sum_{i=1}^{n} p_{ij}$$

$$\mu_j = \frac{\sum_{i=1}^{n} p_{ij} x_i}{\sum_{i=1}^{n} p_{ij}}$$

$$\sigma_j^2 = \frac{\sum_{i=1}^{n} p_{ij} (x_i - \mu_j)(x_i - \mu_j)^T}{\sum_{i=1}^{n} p_{ij}}$$

Figure 4.3: Standard EM algorithm to fit a mixture of m Gaussians to n samples x_i.

4.2.3 Evaluating the Model

Whichever model is derived, the result is a distribution with known parameters. This hypothetical distribution needs to be evaluated to see if it indeed models the noise from which it was derived. To test the model, a null hypothesis, H_0, was created that stated the data are sampled from this distribution of known parameters.

A probability plot was plotted for each of the models — the Gaussian, GMM, and stable distributions. The expected values from the model were plotted against the actual data values, and where the model is accurate a straight line of gradient 1 is produced. The correlation coefficient between the expected values and the actual values gives the probability plot correlation coefficient, PPCC [7].

A significance value of 0.05 (5%) was chosen, and the PPCCs were then compared to the critical value of the PPCC for that significance. Appendix A shows the critical values for different sample sizes for the 0.05 and 0.01 significance values. If the PPCC is greater than the critical value, then H_0 (that the data came from the distribution specified) can be accepted. The probability plot itself also gives a graphical indication of how the data distribution varies from the hypothetical model.

4.2.4 Denoising

The nylon step wedge data from Frankfurt was used in the denoising analysis. The data were denoised by transforming each waveform into its DWT coefficients, and setting to zero the smallest $p\%$ of coefficients. The error this process introduced was measured by reconstructing the waveforms from the truncated DWT coefficients, then recalculating the broadband refractive index and absorption coefficient analogues using the process described in Chapter 5. In the case of the refractive index, the value can be directly compared between the raw and denoised data. For the absorption coefficient, it is the profile of absorption coefficient against frequency which distinguishes materials, and so the profiles need to be compared. The raw and denoised data sets were compared pairwise in three ways, using Pearson correlation, root mean square difference, and Student's paired t-test.

4.3 Results

4.3.1 Blocked scans

Figure 4.4 shows a histogram and normal probability plot for the blocked terahertz signal with a time constant of 500ms. The straight line indicates the hypothetical line on which all the points would lie if the data were normally distributed.

The data shown have a mean of 4.65×10^{-4} and a standard deviation of 3.72×10^{-5}. Plotting the data as a normal probability plot gives a PPCC of 0.999, which is larger than the critical value of 0.9961 for 525 samples from a Gaussian distribution, which means the null hypothesis can be accepted that the data are from a Gaussian distribution. However, at the tail ends of the distribution there is the suggestion of a deviation from normality, particularly at the bottom left.

Figure 4.5 is a plot of the same properties as Figure 4.4, but with a time constant of 50ms instead of 500ms. In this instance the data shown have a mean of 4.64×10^{-4}

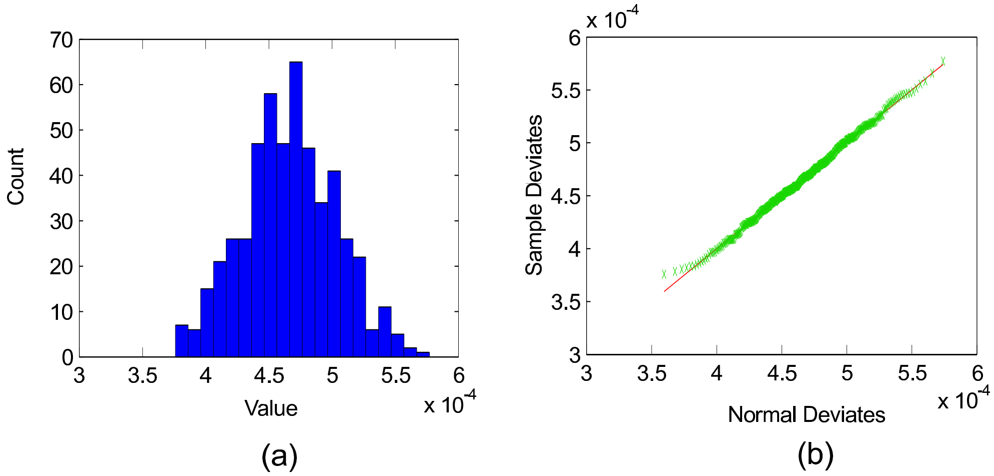

Figure 4.4: A blocked terahertz scan (noise only) with a time constant of 500ms shown as (a) a histogram, and (b) a normal probability plot.

and a standard deviation of 6.25×10^{-5}. Plotting the data as a normal probability plot gives a PPCC of 0.9994 which is even larger than the critical value of 0.9961 for 525 samples from a Gaussian distribution, which means the null hypothesis can be accepted in this case as well. Again a slight deviation from normality is suggested in the tails of the distribution. Table 4.3 summarises these results.

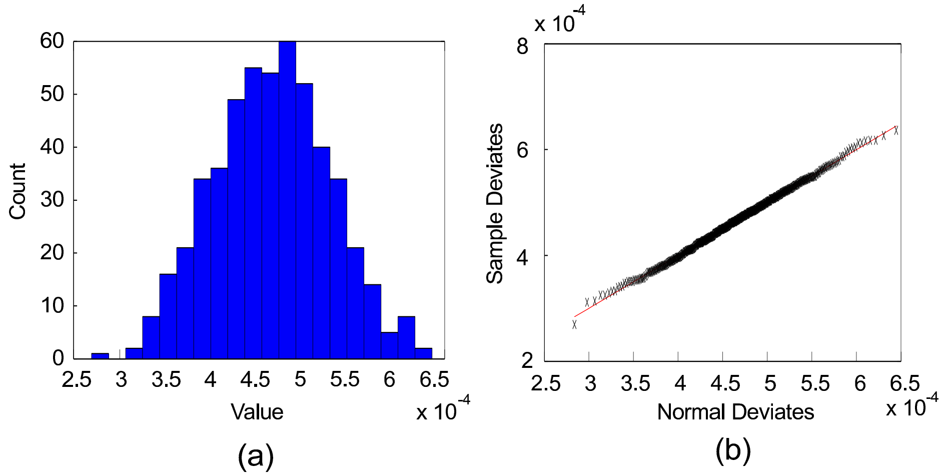

Figure 4.5: A blocked terahertz scan (noise only) with a time constant of 50ms shown as (a) a histogram, and (b) a normal probability plot.

Time Constant	Mean	Standard Deviation	Normal PPCC	Accept?
500ms	4.65×10^{-4}	3.72×10^{-5}	0.9990	$\surd$
50ms	4.64×10^{-4}	6.25×10^{-5}	0.9994	$\surd$

Table 4.3: Distribution of noise at different time constants

Time Constant	α	β	γ	δ	Stable PPCC	Accept?
500ms	2.0	0.0093	0.00002	0.00047	0.9995	$\surd$
50ms	2.0	-0.0561	0.00004	0.00047	0.9998	$\surd$

Table 4.4: Fit of stable distribution to noise at different time constants

The simple model of a Gaussian distribution correctly models these data, so there is no need to apply the GMM or stable distribution. However, for completeness a stable distribution was fitted to this noise, and gave the parameters shown in Table 4.4. Notice that the α value of 2 confirms this is a Gaussian distributio

4.3.2 Free air scans

The process from Section 4.3.1 was repeated with a set of 16 terahertz pulses transmitted through free air. All 16 were recorded within a short space of time and with identical acquisition parameters. In each case 256 time points were recorded, meaning the noise distribution will consist of 4096 samples. There was a very small deviation between each of the pulses, as demonstrated in Figure 4.6, which shows (a) the 16 pulses superimposed on top of one another, and (b) the root mean square difference in the time domain between the 16 pulses and the mean of the 16 pulses.

Figure 4.6(a) particularly demonstrates that the apparently noisy oscillations after the main pulse are not in fact random noise at all — the 16 pulses are virtually indistinguishable. There is also the strong suggestion of a correlation between the amplitude and the noise (as measured by the RMS difference). Figure 4.7 does not show that there is any direct correlation between either electric field (a) or the absolute value of the electric field

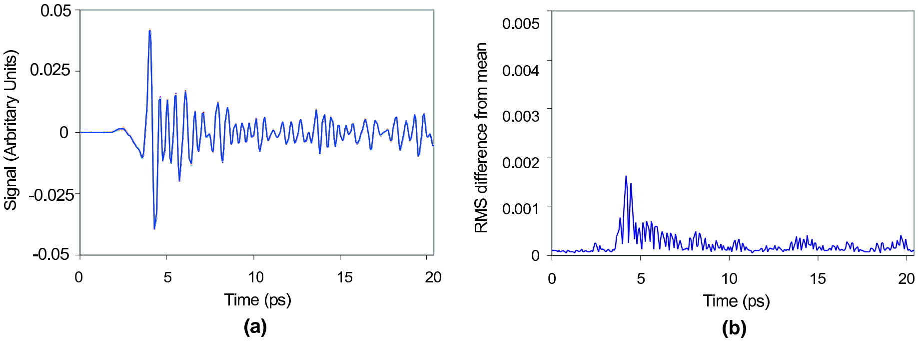

Figure 4.6: (a) Superimposed time domain of 16 pulses through free air acquired with identical parameters, and (b) the root mean square difference between the pulse and the mean pulse at each time point.

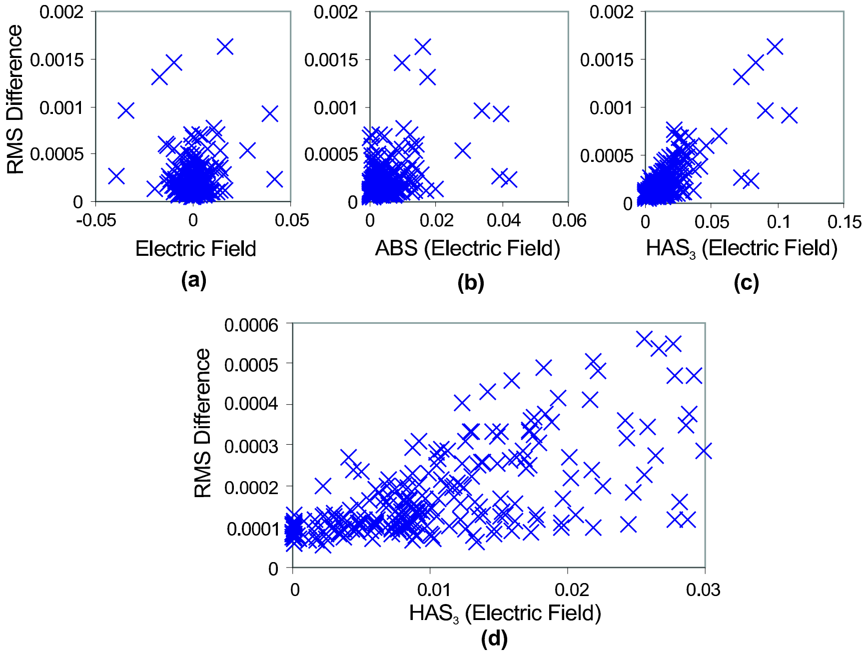

Figure 4.7: The RMS difference (noise) versus (a) electric field, (b) absolute value of the electric field, and (c) the historical sum of the absolute values of the electric field including the previous 2 values, making a total of 3 values included in the sum ($HAS3$). Plot (d) is rescaled plot of (c), showing the bottom left hand corner in more detail.

(b). However, if the absolute value of the electric field is summed over 3 values 'historically' there is a suggestion of a correlation. The sum was over 3 values because this was determined empirically to give the greatest correlation, with rapid dropping off of correlation with larger and smaller values.

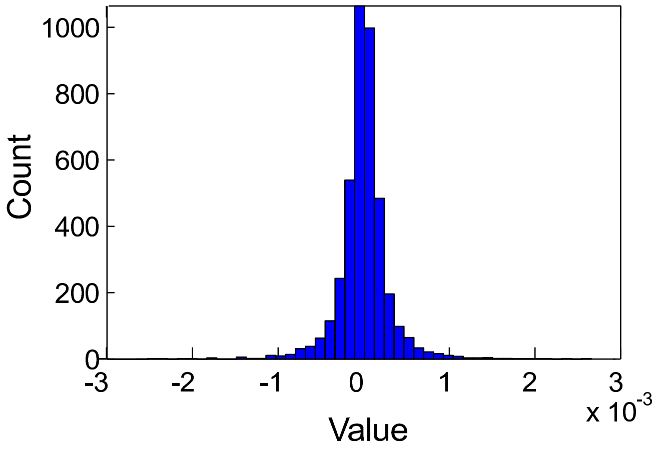

Figure 4.8: Histogram of noise from a free air scan

Gaussian	μ_j	σ_j^2	w_j
G_1	-1.23×10^{-6}	2.12×10^{-8}	0.79
G_2	4.6×10^{-6}	3.52×10^{-7}	0.21

Table 4.5: Gaussian mixture model created using EM for a free air scan.

Figure 4.8 shows the histogram of the estimated noise distribution.

The noise from these pulses had a mean of -7.73×10^{-21} and a variance of 9.11×10^{-8}. The normal probability plot had a normal PPCC of 0.9179, far below the critical value of over 0.9979 for 4096 samples, meaning this noise is not normally distributed. Applying a GMM to the distribution gave rise to two Gaussians with the parameters shown in Table 4.5. This GMM gave a correlation coefficient of 0.9882, a much better match, but still below the critical value.

Figure 4.9(a) shows the normal probability plot, and (b) shows the GMM probability plot. Notice that the noise clearly diverges from the hypothetical models at the tail ends of the distribution, suggesting that the noise is heavy tailed.

Fitting a stable distribution resulted in a model with the parameters and PPCC shown in Table 4.6. Figure 4.10 shows the probability plot from the stable distribution. The PPCC value of 0.9999 is above the critical value, and the probability plot confirms the goodness of fit.

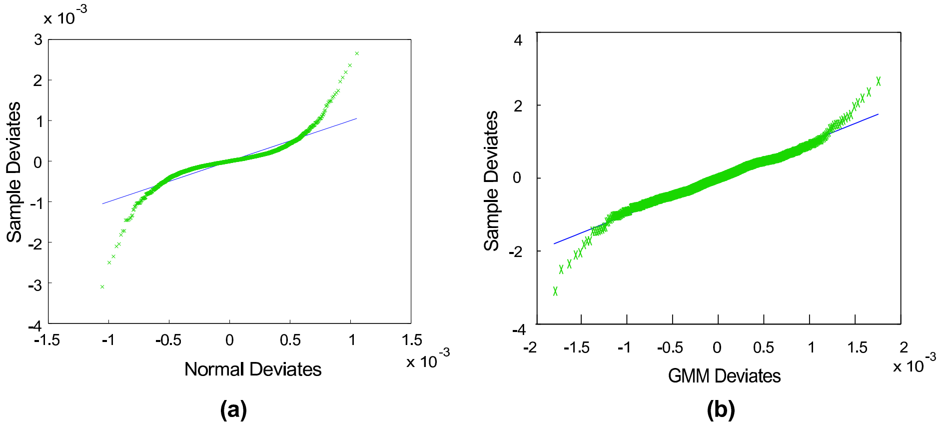

Figure 4.9: (a) Normal probability plot, and (b) GMM probability plot from a free air scan.

α	β	γ	δ	PPCC	accept?
1.404	0.008	0.00012	-0.9×10^{-6}	0.9999	√

Table 4.6: Results for a stable distribution fitted to noise from a free air scan.

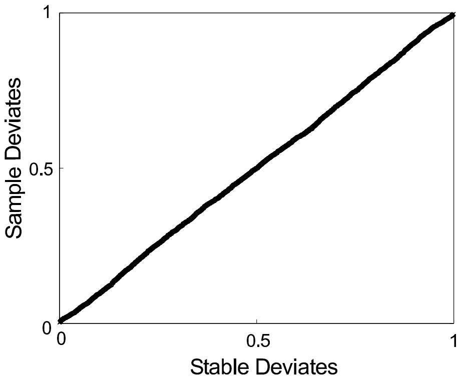

Figure 4.10: Stable distribution probability plot from a free air scan.

4.3.3 Mean pulse of step wedges

Table 4.7 shows the depth of each step of the wedges, and the number of pulses obtained from that step using the region of interest technique, with each pulse having 128 time points.

The average pulse for each particular step provided the estimate of a low-noise pulse through that depth of material. The differences between the pulses through a given step

Nylon	Count	Size of distribution	Resin	Count	Size of distribution
0.3 mm	32	4096	0.1 mm	43	5504
1 mm	53	6784	0.5 mm	95	12160
2 mm	40	5120	1.5 mm	59	7552
4 mm	44	5632	3 mm	49	6272
6 mm	31	3968	5 mm	79	10112
7 mm	13	1664	7 mm	60	7680

Table 4.7: Depth profiles and pulse counts of the step wedges.

and that step's average pulse provided the estimate of the noise distribution for that depth of material. These noise distributions were then assessed as before.

Figure 4.11 shows example normal probability plots for the thickest and thinnest nylon steps. Table 4.8 shows the mean, variance and correlation coefficient obtained for each of the 12 steps.

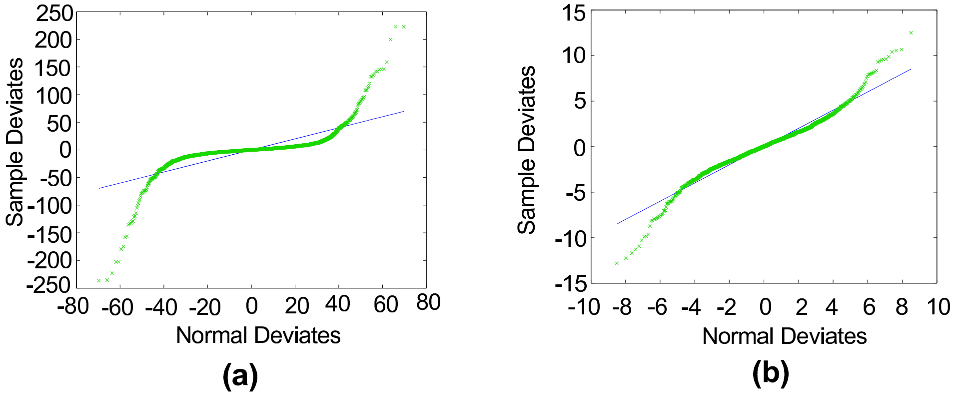

Figure 4.11: Normal probability plot of noise from a step of nylon of (a) 0.3 mm, and (b) 7 mm.

These results clearly show that one cannot have confidence in using a single Gaussian distribution to model this noise, as the critical value for these sizes of distribution is in excess of 0.9979. The only exception to this rule is for the 7 mm thickness of resin, where our critical value table does not extend to 7680 samples. However since it is only just above the critical value for 1000 samples, it seems very unlikely to be above it at 7680.

Using the Gaussian mixture model, a far better fitting model was built, as shown in

Step	Mean	Variance	PPCC	Accept?
Nylon 0.3 mm	0	399.83	0.7415	×
Nylon 1 mm	0	209.04	0.758	×
Nylon 2 mm	0	64.35	0.8491	×
Nylon 4 mm	0	17.08	0.9371	×
Nylon 6 mm	0	6.68	0.9912	×
Nylon 7 mm	0	6.90	0.9799	×
Resin 0.1 mm	0	2873.53	0.7238	×
Resin 0.5 mm	0	619.03	0.7382	×
Resin 1.5 mm	0	113.86	0.7786	×
Resin 3 mm	0	23.24	0.9922	×
Resin 5 mm	0	21.93	0.9978	×
Resin 7 mm	0	20.50	0.9985	?

Table 4.8: Normal probability plot correlation coefficient for nylon and resin step wedges.

Table 4.9. Figure 4.12 shows an example GMM probability plot for the thinnest nylon step, for comparison with Figure 4.11(a). Most of these correlation coefficient values are smaller then the critical value for this number of samples, which means it cannot be said that the noise distribution has been accurately modelled. In the cases where the critical value is above the largest value I have, the probability plot still showed signs of heavy tailedness. The EM algorithm was run with two and with three Gaussians, and it was found that the weighting on the third Gaussian became infinitesimal very rapidly, meaning that it was not contributing to the model. Only results for a mixture of two Gaussians are therefore shown.

Finally a stable distribution was fitted to this noise data, resulting in PPCC values that were well above the largest critical value I have. Since our critical values do not extend to this size of distribution, one cannot be certain that the stable distribution is modelling it. However the probably plots show an almost perfect line, similar to Figure 4.10. Table 4.10 shows the parameters fitted and the corresponding PPCC values.

These values were then plotted against step depth, as shown in Figure 4.13. For ease of comparison, the α value of the blocked and free air scans are shown as dotted lines in Figure 4.13. Note that the thickness axis in these two cases is meaningless. The γ values

Step	Gaussian	μ_j	σ_j^2	w_j	PPCC	Accept
Nylon 0.3 mm	G_1	-0.0873	30.7	0.87	–	
	G_2	0.593	2905	0.13	–	
	$\sum w_j G_j$	–	–	–	0.960251	×
Nylon 1 mm	G_1	-0.0986	13.8	0.86	–	
	G_2	0.599	1395	0.14	–	
	$\sum w_j G_j$	–	–	–	0.975396	×
Nylon 2 mm	G_1	-0.100	9.94	0.86	–	
	G_2	0.565	411	0.14	–	
	$\sum w_j G_j$	–	–	–	0.994245	×
Nylon 4 mm	G_1	-0.132	5.24	0.81	–	
	G_2	0.565	67.3	0.19	–	
	$\sum w_j G_j$	–	–	–	0.994299	×
Nylon 6 mm	G_1	-0.0101	4.36	0.83	–	
	G_2	0.049	18	0.17	–	
	$\sum w_j G_j$	–	–	–	0.998484	?
Nylon 7 mm	G_1	0.00357	3.33	0.80	–	
	G_2	-0.0141	21	0.20	–	
	$\sum w_j G_j$	–	–	–	0.960251	×
Resin 0.1 mm	G_1	-0.152	259	0.90	–	
	G_2	1.434	27512	0.10	–	
	$\sum w_j G_j$	–	–	–	0.966960	×
Resin 0.5 mm	G_1	-0.149	64.1	0.90	–	
	G_2	1.29	5425	0.10	–	
	$\sum w_j G_j$	–	–	–	0.960667	×
Resin 1.5 mm	G_1	0.0593	24.5	0.94	–	
	G_2	-0.992	1607	0.06	–	
	$\sum w_j G_j$	–	–	–	0.975454	×
Resin 3 mm	G_1	0.0591	16.1	0.86	–	
	G_2	-0.37	67.8	0.14	–	
	$\sum w_j G_j$	–	–	–	0.999650	?
Resin 5 mm	G_1	0.0367	18.3	0.90	–	
	G_2	-0.332	54.6	0.10	–	
	$\sum w_j G_j$	–	–	–	0.998484	?
Resin 7 mm	G_1	-0.0557	13	0.59	–	
	G_2	0.0787	31.1	0.41	–	
	$\sum w_j G_j$	–	–	–	0.999857	?

Table 4.9: Gaussian mixture models fitted to the noise distribution from nylon and resin step wedges. Numbers are shown to three significant figures, except the correlation coefficients.

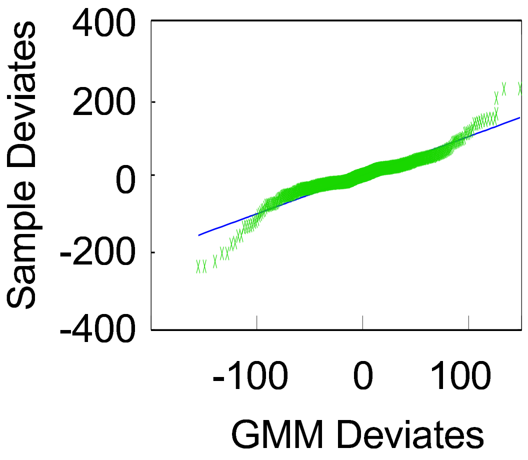

Figure 4.12: GMM probability plot of noise from a step of nylon of 0.3 mm

Step	α	β	γ	δ	PPCC
Nylon 0.3 mm	1.2167	-0.0055	3.8476	-0.062	0.99996
Nylon 1 mm	1.1784	-0.0795	2.6364	0.057	0.99996
Nylon 2 mm	1.3461	-0.01	2.3656	-0.05	0.99993
Nylon 4 mm	1.4874	0.101	1.8335	-0.135	0.99982
Nylon 6 mm	1.8571	0.0188	1.6535	-0.0078	0.99994
Nylon 7 mm	1.6862	0.075	1.4902	0.00215	0.99988
Resin 0.1 mm	1.2954	0.013	10.8478	-0.137	0.99998
Resin 0.5 mm	1.3366	-0.0065	5.5739	-0.111	0.99998
Resin 1.5 mm	1.5453	-0.0055	3.3658	0.044	0.99997
Resin 3 mm	1.8611	-0.0567	3.0958	0.0504	0.99992
Resin 5 mm	1.9267	0.0613	3.1259	0.0273	0.99997
Resin 7 mm	1.9407	0.0582	3.1088	-0.014	0.99989

Table 4.10: Stable distribution models fitted to the noise from nylon and resin step wedges.

for the blocked and free air scans are too small to be seen on the scale of Figure 4.13(c) — they are effectively both zero.

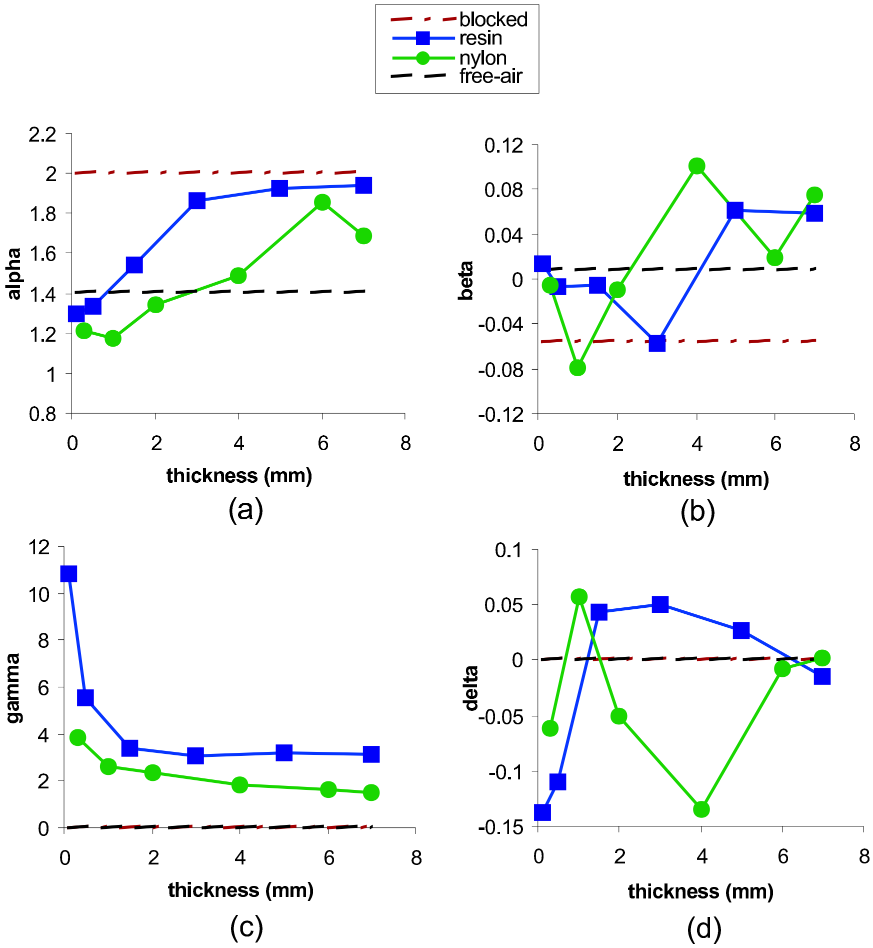

Figure 4.13: Stable distribution parameters against step thickness for nylon and resin step wedges: (a) α, (b) β, (c) γ, (d) δ. Values for the blocked and free air scans are shown for comparison.

4.3.4 Denoising

It was found that wavelet denoising has a negligible impact on the refractive index calculations, even down to keeping only 10% of coefficients. Figures 4.14(a)–(e) show the effect that wavelet shrinkage has on the pulse in the time domain. Notice that the main peak is unaffected, and by shrinkages to 10% nearly all the low-level oscillations have been removed.

Table 4.11 shows the time delays for each step at the various shrinkage levels, together

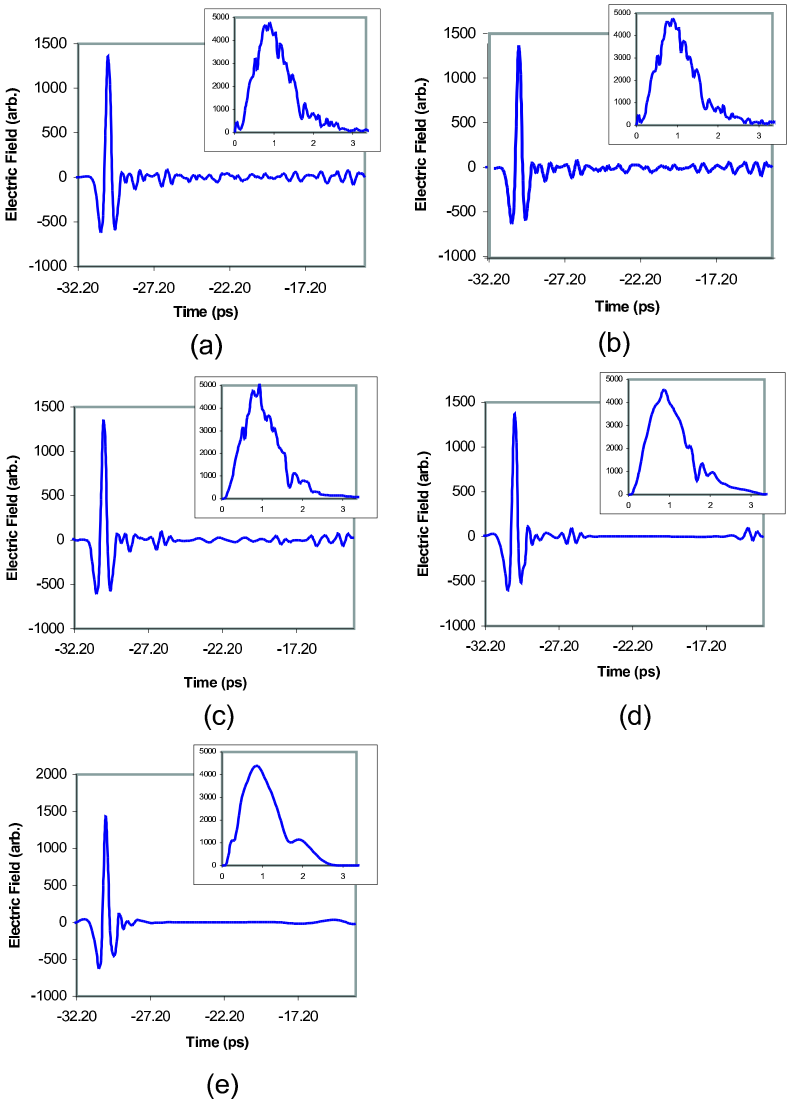

(a) (b) (c) (d) (e)

Figure 4.14: The time series samples of terahertz pulses with spectral content inset, at shrinkage to (a) 100 % (raw), (b) 50%, (c) 20%, (d) 10%, and (e) 5% of DWT coefficients.

with the final calculated refractive index. As can be seen, the extreme shrinkage at 5% does cause differences in the time-delay calculations, but even then only a small error is introduced to the refractive index value.

Similarly it was found that the absorption coefficient profile remained essentially unaltered with only 20% of the wavelet coefficients. Table 4.12 and Figure 4.15 show the absorption coefficient versus frequency profile at different shrinkage levels. On the plot, the absorption coefficient values have been vertically offset for ease of view. By 10% the

Step	Time delay (ps) at shrinkage level				
	100%	50%	20%	10%	5%
Nylon 0.3 mm	0.49 ±0.02	0.49 ±0.02	0.49 ±0.02	0.49 ±0.02	0.49 ±0.03
Nylon 1 mm	1.91 ±0.02	1.91 ±0.02	1.91 ±0.02	1.91 ±0.02	1.93 ±0.02
Nylon 2 mm	3.94 ±0.01	3.94 ±0.01	3.94 ±0.02	3.96 ±0.02	3.991 ±0.01
Nylon 4 mm	7.96 ±0.02	7.96 ±0.02	7.97 ±0.02	7.97 ±0.03	7.97 ±0.02
Nylon 6 mm	11.95 ±0.02	11.95 ±0.02	11.95 ±0.02	11.92 ±0.03	11.89 ±0.01
Nylon 7 mm	13.98 ±0.03	13.98 ±0.03	13.99 ±0.02	13.99 ±0.04	14.15 ±0.01
Refractive index	1.603 ±0.002	1.603 ±0.002	1.604 ±0.003	1.603 ±0.005	1.61 ±0.01

Table 4.11: Time delay and refractive index for different amounts of wavelet shrinkage.

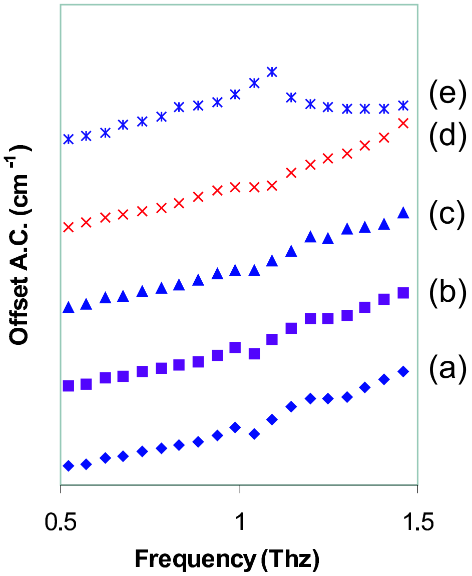

Figure 4.15: Absorption coefficient profile vs. frequency for a nylon step wedge at shrinkage to (a) 100% (raw), (b) 50%, (c) 20%, (d) 10%, and (e) 5%, offset vertically.

feature at 1.1 THz has been lost, and with only 5% of coefficients large errors have been introduced above 1 THz. The difference measures at the bottom are the Pearson correlation, the root mean square difference, and the Student's paired t-test probability. A t-test suggests that the two series have no statistically significant difference at the 95% level, even with the shrinkage to 5% of coefficients (Table 4.12). To ensure that important features in the shape of the absorption coefficient curve, including the 1.1 THz feature, are not removed, it would be prudent not to use the highest levels of shrinkage. A 'safe' level of shrinkage for this experiment would therefore be 20%.

Frequency (THz)	Absorption coefficient (cm^{-}1) at shrinkage level				
	100%	50%	20%	10%	5%
0.52	4.63	4.61	4.43	4.29	4.54
0.57	5.14	5.15	5.30	5.50	5.03
0.63	6.68	6.68	6.66	6.87	6.13
0.68	7.29	7.29	7.26	7.75	7.91
0.73	8.27	8.27	8.30	8.41	8.64
0.78	9.22	9.21	9.23	9.16	10.13
0.83	9.87	9.89	10.02	10.24	12.41
0.89	11.00	10.98	11.01	12.14	12.87
0.94	12.35	12.40	12.80	13.51	13.80
0.99	14.33	14.34	13.71	14.23	15.43
1.04	12.94	12.91	13.46	14.33	18.53
1.09	16.27	16.28	16.20	14.94	21.17
1.15	19.56	19.37	18.45	18.11	14.82
1.20	21.76	21.71	21.88	19.91	13.35
1.25	21.69	21.55	21.41	21.50	12.56
1.30	22.14	22.35	23.94	22.96	12.09
1.35	24.28	24.39	24.22	24.72	11.90
1.41	26.54	26.35	25.33	26.98	12.17
1.46	28.29	28.19	28.13	30.33	12.91
Pearson Correlation	1.0	0.9999	0.9966	0.9918	0.495
R.M.S.D	0	0.09	0.61	0.97	7.05
P-value from paired t-test	–	0.43	0.85	0.40	0.06

Table 4.12: Absorption coefficient (cm^{-1}) of nylon at various frequencies and shrinkage levels, with measures of difference between each shrinkage level and the raw data.

4.4 Discussion

The first observation to make is that the electric field values recorded in the terahertz pulses are arbitrary values, and so a direct comparison between the distributions obtained for different acquisitions is not possible.

The second observation is that, unlike what might have been naïvely expected, the noise distributions are not zero-mean. However, it is known that the terahertz scanner can introduce a d.c. offset to the data, and this would account for the small shift away from zero.

From the results of the blocked scan in Section 4.3.1 one can confidently say that the background and detector noise in terahertz imaging may be modelled as a Gaussian distribution, and altering the LIA time constant does not really change the mean of the noise but does affect the variance — a shorter time-constant leads to a larger variance. This is as expected, since the time constant is in effect a measure of how much averaging the hardware is applying. Thus shorter time constants, or less averaging, lead to larger variance in the noise distribution, or more noise. However with the blocked scan there is a suggestion of ‘heavy tailedness’, as shown in the probability plot (Figure 4.5). The points creep below the hypothetical straight line at the bottom left, and above it at the top right, and so the apparent good fit may be due to under-sampling from the distribution.

The free air scan results demonstrate that the oscillations after the main pulse were not random noise but rather were created by a deterministic process, probably absorption and re-emission by water vapour [49], although another possibility is reflection inside the crystals. However terahertz pulses through homogeneous regions of material showed considerably larger random variation in these oscillations, perhaps due to scattering, suggesting the working signal to noise ratio may actually be smaller. In any case, the random noise in all these signals is not distributed in either a Gaussian distribution or a simple combination of Gaussian distributions. There is also a relationship clearly visible between the arrival of the pulse and the noise level in Figure 4.6, even without a particular correlation. It is possible that the arrival of the pulse is causing a ‘ringing’ effect as there does appear to be some oscillation in the noise amplitude after the pulse. This is probably due to a physical effect such as reflection, or absorption and re-emission, but would require further investigation.

The examination of a homogeneous region did not lead to a Gaussian distribution, but did lead to a distribution which looked like a normal distribution with ‘heavy tails’ (Section 4.3.3). The application of a Gaussian Mixture Model to the heavy tailed distribution led to a far better match with the averaged region of interest data, although one still could not be confident this was the correct model based on the probability plot. If the Johnson noise forms a normal distribution, and the shot-noise is from a different normal distribution, then the overall noise will be modelled by a mixture of normal distributions, which might explain the improved fit.

A stable distribution modelled the noise almost perfectly in all of these cases. One would expect the noise distribution to be symmetrical, and hence the skewness parameter β to be zero. Instead of this, β fluctuated around 0, depending on which material, if any, had been imaged. β can take values from -1 to 1, and the largest β value fitted was 0.1 for 4 mm of nylon. It is possible that these fluctuations may be related to the scattering properties of nylon and resin, but this is impossible to tell without the inclusion of a scattering model. As previously discussed, one would also expect the noise process to zero mean, and hence for the location parameter δ to be zero. As before, the small deviations from zero are probably due to d.c. offset errors in the lock-in amplifier.

Figure 4.13 suggests that α and γ might vary with thickness — as the material thickness increases, α is tending to increase, and γ initially decreases, before flattening out. Neither β nor δ particularly demonstrate a relationship with thickness, but rather both fluctuate slightly around zero. Once again there may be a scattering process influencing these variations.

The results from applying wavelet shrinkage show that only 20% of the wavelet coefficients are actually required when calculating optical properties. This suggests that a wavelet shrinkage may actually be useful in a compression scheme, where only around 20% of the data need be transmitted and stored, which would represent a significant reduction in storage and transmission costs of data.

4.5 Conclusions

I have shown that the noise present in terahertz images can be modelled using distributions from the stable family, and have calculated the parameters needed to model the noise for a variety of terahertz images. I have also shown that, with the exception of the case when all terahertz radiation is blocked, the noise present in terahertz images can not be modelled by either a single or multiple Gaussian model.

Terahertz imaging is resilient to ‘classical’ noise such as shot noise, Johnson noise, and background radiation due to the coherent nature of the detection. However inevitable differences between and within samples, combined with other factors such as atmospheric conditions, lead to a large variation in the pulses. The distribution of noise in terahertz pulsed imaging is therefore more complex than was originally hypothesised. It had been hoped that a computationally simple noise model, such as a single or dual Gaussian distribution, would result from this analysis. The probability plots show that this is not the case, but the noise can accurately be modelled by other distributions from the stable family.

Furthermore, the noise model created demonstrates dependence on both material type and thickness, making it very difficult to build realistic models of noise, particularly for material which has not been well characterised under terahertz. This in turn makes the addition of entirely realistic noise to synthetic images an impossible goal until these noise processes are better understood.

The apparent dependence on thickness also makes the design of algorithms that are robust to terahertz noise more complex, as the thickness is often not known. Similarly the lack of closed forms for the probability density functions of most stable distributions make analytical methods impossible. However the distributions are essentially heavy tailed Gaussians. In practice this means that there are more data points at the edges of the distribution than would be expected from a Gaussian distribution, and it may be possible to heuristically allow for these 'extra' points.

It is worth considering that these parameters may provide some form of insight into the nature of the material being imaged. This may prove beneficial if a new material produces similar results to a material previously characterised in this way.

It has also been shown that the wavelet shrinkage algorithm can use as few as 20% of the wavelet coefficients from a step wedge image without introducing any significant errors in the calculation of the optical properties. This result is of interest for medical terahertz pulse imaging since it could enable a real reduction in the storage and transmission costs of scans, while not causing any change to results based on the physical properties of the subject.

Chapter 5

Terahertz Imaging: Using Optical Parameters as a Contrast Mechanism

---

5.1 Introduction

One of the most interesting aspects of terahertz pulsed imaging is the richness of data acquired by virtue of the coherent detection in the time-domain. Typically one would expect 512 or more samples to be acquired in the time-domain for each pixel in the image. In this way a high dimensional vector is stored 'behind' every pixel of the 'image', rather than just a single scalar value. Furthermore since the detection is coherent each pulse may be transformed to the frequency domain. Combining the data from these domains enables even higher dimensional vectors to be formed behind each pixel — just using the time domain data, and the amplitude and angle of Fourier coefficients gives a 1536 dimensional vector for each pixel.

Not every dimension of these vectors contain useful information – for example the signal to noise ratio at higher frequencies is too low for meaningful data to be extracted.

However there is scope for selecting smaller feature vectors or extracting other parameters in order to create images [34, 46]. The data acquisition process can also be quite slow, and although advances in the technology are expected to continue speeding up the imaging process, it is worth exploring whether an entire time series need be recorded, or only a subset thereof. Figures 5.1–5.4 show examples of simple parametric images. In each of these figures a parameter has been extracted and mapped to grey level in order to form an image.

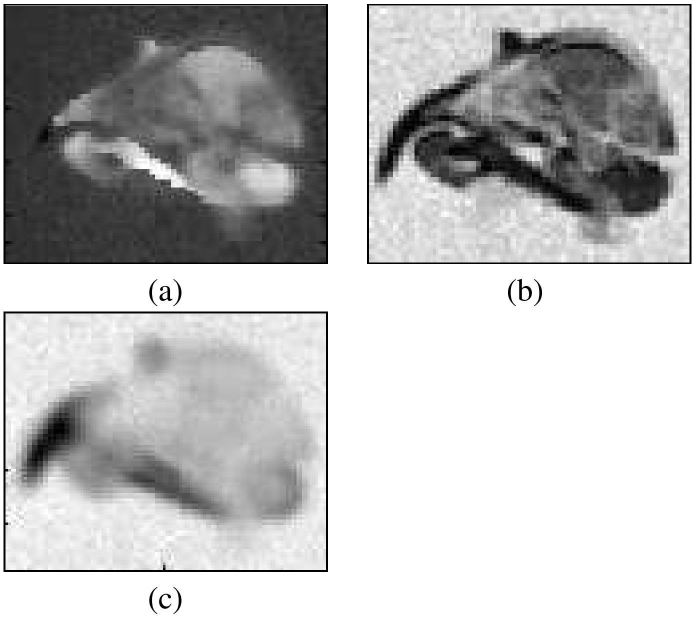

Figure 5.1: A cross section of a bird's skull imaged in Frankfurt. 32 time points were acquired at intervals of 150 fs, displacement -85. 60 × 48 pixels were acquired, with x range -4 – 26 mm and y range -12 – 12 mm, and no averaging. (a) Reference peak to signal peak time difference, white indicates largest delay. (b) & (c) Relative transmission at (b) 1 THz and (c) 0.41 THz, white indicates largest transmission.

The differences in grey level at each pixel in these images is caused by two factors — the thickness of material, and its optical parameters. The bird's head is a slice of uniform thickness, meaning the variation is entirely due to varying optical parameters. The nylon and resin step wedges are of homogeneous material, meaning the variation is entirely due to varying thickness, and the marzipan pig is somewhere between the two, with a mostly homogeneous material having varying thickness.

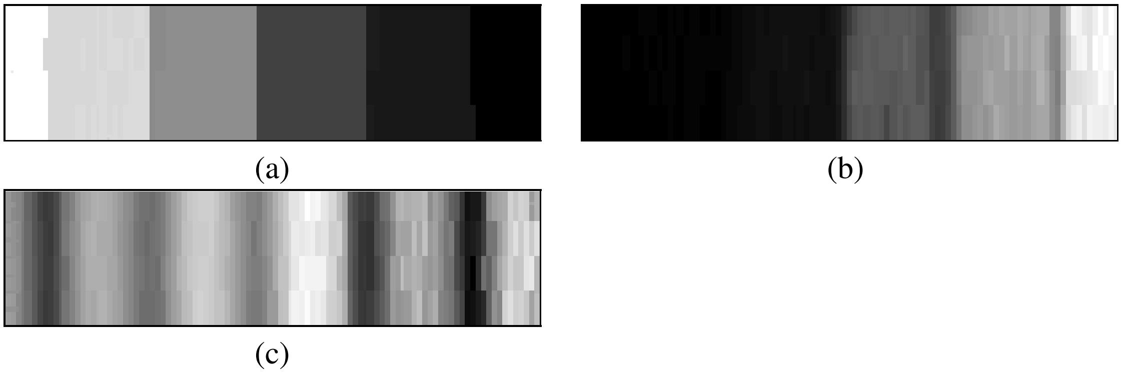
(a) (b) (c)

Figure 5.2: A nylon step wedge imaged in Frankfurt. The step wedges are fully described on page 40, and the thinnest step is to the right of the image. (a) Reference peak to signal peak time difference, white indicates largest delay. (b) & (c) Relative transmission at (b) 1 THz and (c) 0.31 THz, white indicates largest transmission.

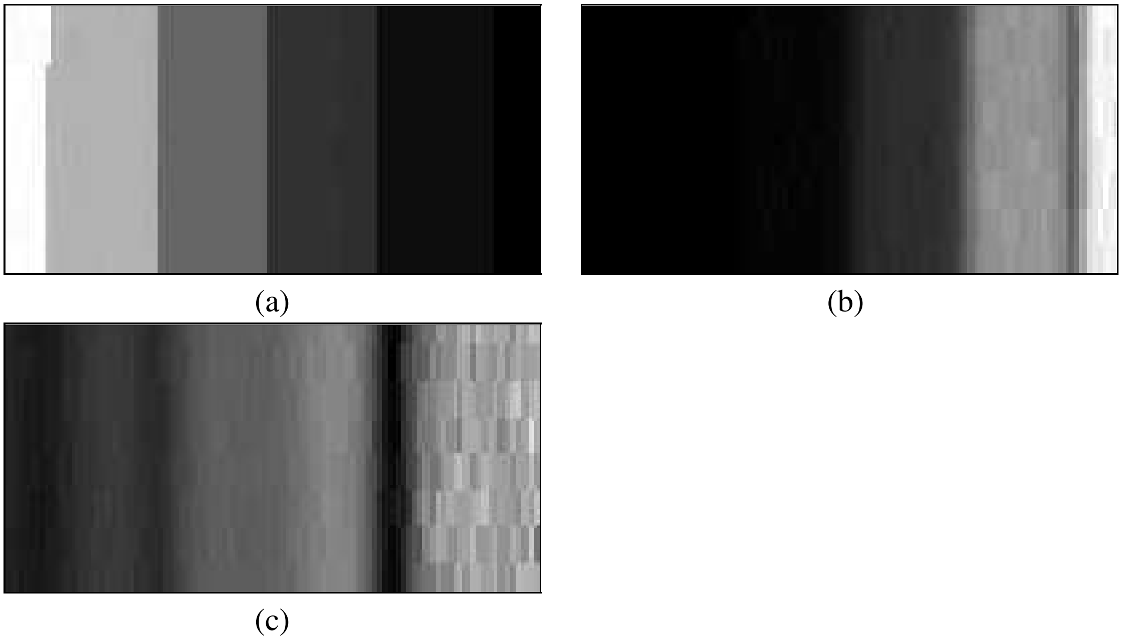
(a) (b) (c)

Figure 5.3: A resin step wedge imaged in Frankfurt. The step wedges are fully described on page 40, and the thinnest step is to the right of the image. (a) Reference peak to signal peak time difference, white indicates largest delay. (b) & (c) Relative transmission at (b) 1 THz and (c) 0.31 THz, white indicates largest transmission.

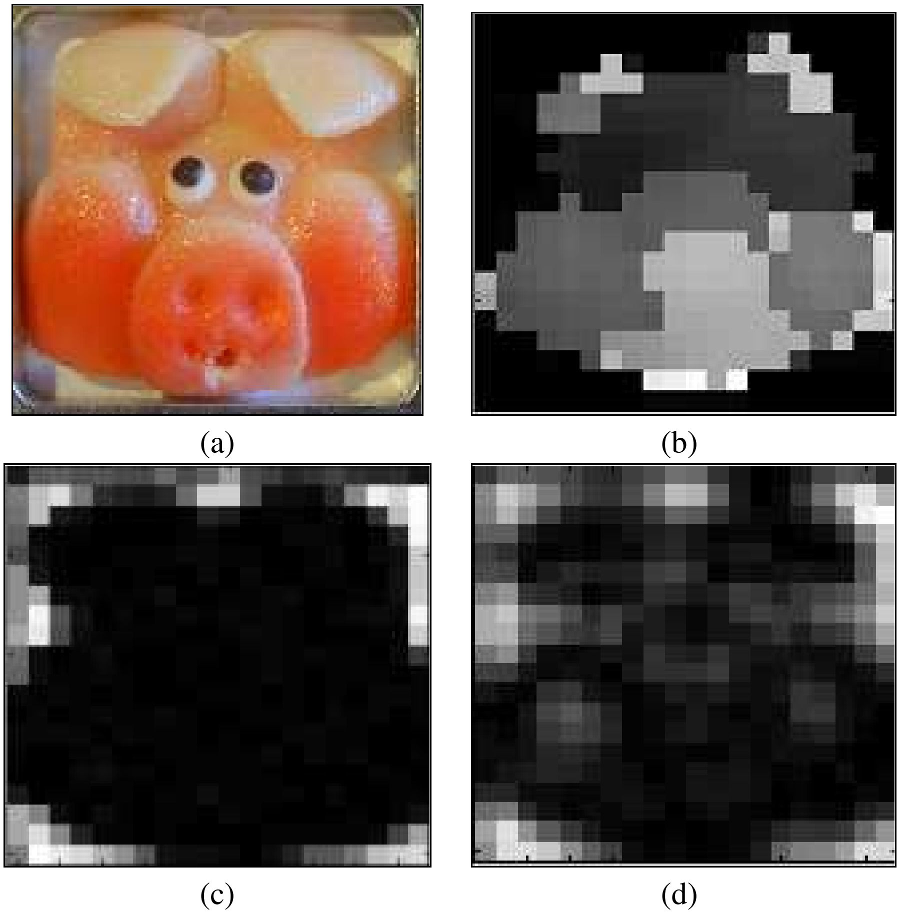
(a) (b)
(c) (d)

Figure 5.4: The head of a marzipan pig imaged in Frankfurt. 128 time points were acquired at intervals of 300 fs, displacement -77. 20 × 20 pixels were acquired, with x range -3 – 25 mm and y range -15 – 15 mm, and no averaging. (a) Photograph of the pig. (b) Reference peak to signal peak time difference, white indicates largest delay. (c) & (d) Relative transmission at (c) 1 THz and (d) 0.31 THz, white indicates largest transmission.

5.1.1 Complex Refractive Index

The complex refractive index, $\tilde{n}$, of a material gives the values of the two optical constants n (refractive index), and k (extinction coefficient) according to (5.1), where both parameters vary with wavelength λ.

$$\tilde{n}(\lambda) = n(\lambda) + ik(\lambda) \tag{5.1}$$

Together these constants describe the propagation of a given frequency of electromagnetic radiation through a material, and conversely the complex refractive index may be derived by observing the effect that a known thickness of material has on electromagnetic radiation of a known frequency.

These parameters n and k are defined below, and are considered separately for the rest of this study. Note that this thesis will also be using the absorption coefficient, α, instead of k, where (5.2) describes the conversion between α and k.

$$\alpha = \frac{4\pi k}{\lambda} \tag{5.2}$$

The optical constants are traditionally expressed in physics in terms of wavelength, as in (5.1) and (5.2). In the context of this thesis, however, frequency (ω) is a more natural expression and will be used for the rest of this chapter. The conversion between frequency and wavelength is trivial, as shown in (5.3), where c is the speed of light in a vacuum.

$$\omega = \frac{c}{\lambda} \tag{5.3}$$

Refractive Index

The refractive index, n, for any substance is the ratio of the velocity of electromagnetic radiation of a given frequency in a vacuum to its velocity in the substance. The relationship between the phase angle, ϕ, of electromagnetic radiation of frequency ω transmitted through a thickness x of material and the incident electromagnetic radiation, ϕ_0, is given in (5.4) [63].

$$\phi(\omega) = \phi_0(\omega) + \frac{2\pi\omega(n(\omega)-1)}{c}x \tag{5.4}$$

The refractive index provides a mechanism for contrast as radiation transmitted through the same thickness of different materials will experience differing phase shifts.

Absorption Coefficient

The absorption coefficient is a measure of how much electromagnetic radiation of a given frequency is absorbed by a given thickness of material. The Beer-Lambert law, shown in (5.5), defines the relationship between the absorption coefficient $\alpha(\omega)$, the depth of material x, and the relative transmission $\frac{I}{I_0}$ (where I_0 is the incident radiation and I is the transmitted radiation).

$$ln\left(\frac{I(\omega)}{I_0(\omega)}\right) = -\alpha(\omega)x \tag{5.5}$$

The absorption coefficient provides a mechanism for contrast as radiation transmitted through the same thickness of different material will experience different attenuation. This contrast is illustrated for two different frequencies in Figures 5.1(b) and (c).

5.1.2 Broadband Optical Properties

If a broadband signal, such as a pulse, is used, both the refractive index and absorption coefficient have analogous pulse specific broadband versions, calculated from the time series data. The contrast in Figures 5.1(a), 5.2(a), 5.3(a), and 5.4(b) arises from differences in the broadband refractive index. (5.6) shows the calculation of this broadband n using peak to peak time difference (TD).

$$TD = \frac{(n-1)x}{c} \tag{5.6}$$

(5.7) shows the calculation of this broadband α, I_0 and I being incident and transmitted peak amplitudes in the time domain.

$$ln\left(\frac{I}{I_0}\right) = -\alpha x \tag{5.7}$$

5.1.3 Summary

As highlighted in Chapter 2, terahertz radiation is a novel imaging field, and very few materials have been characterised under these frequencies of radiation. Consequently there are no predications about the expected outcome of these analyses. It is hoped that these data will prove useful for modelling the expected appearances of images and the feasibility of future investigations.

In this chapter I present the results of calculating the refractive indices and absorption coefficients of known thicknesses of material. Analogues to these parameters are calculated using broadband techniques, the short time Fourier transform (STFT) and the wide band cross ambiguity function (WBCAF).

I also use the STFT with a rectangular window to simulate shorter acquisition times and measure the effect that this has on the calculation of the complex refractive index.

5.2 Method

5.2.1 Data

In addition to the nylon and resin step wedges detailed on page 40, a number of biological samples were imaged at Leeds. These consist of two sets of slices taken from two amputated legs, with local Research Ethical Committee approval. Six different types of tissue were excised; skin, adipose tissue (‘fat’), striated muscle (‘muscle’), vein, artery, and nerve. Each tissue type was sliced into four different thicknesses; 50 μm, 100 μm, 150 μm, and 200 μm. Three pulses were acquired through each of the four thicknesses of tissue, leading to twelve pulses in total for each tissue type per set. The first set of slices was imaged twice over two days, having been stored in saline solution in a fridge overnight.

Each pulse was a point measurement of 256 time points, acquired 80 fs apart, with a displacement of 23.6. The time constant was 200 ms, with an LIA delay of 300 ms.

5.2.2 Traditional Analysis

The Fourier transform, introduced in Chapter 3, was used to calculate values of the dependent variables for the calculation of the refractive indices and absorption coefficients as a function of frequency. The techniques described here have been successfully applied to terahertz data of water and other polar liquids by Kindt and Schmuttenmaer [41]. It is also possible to calculate these parameters by fitting an expression that combines all the reflection and propagation to the observed signal [18, 21], but these methods are not robust.

Absorption Coefficient

The magnitude of each Fourier coefficient, $|\hat{f}(\omega)|$, from the transform of reference pulse, gives values for $I_0(\omega)$. Similarly the coefficients from the transform of a pulse transmitted through a given thickness of material gives $I(\omega)$ for that thickness of material. The mean of the logarithm of $\frac{I(\omega)}{I_0(\omega)}$ for each thickness of material can then plotted against thickness, which should result in a straight line, the gradient of which is directly proportional to $\alpha(\omega)$, as shown in (5.5). The error metric was the maximum variation in gradient permitted within one standard deviation of the data points. In each case the gradient was calculated using least squares linear regression.

Finally, it is necessary to apply inclusion criteria as at higher frequencies and thicker steps the absorption becomes such that no meaningful terahertz signal passes through the material. At any given thickness and frequency only values of $\frac{I(\omega)}{I_0(\omega)}$ greater than 0.05 were used. This corresponds to a relative transmission greater than 5%. Additionally, $\alpha(\omega)$ was only calculated at frequencies where there were $\frac{I(\omega)}{I_0(\omega)}$ values for at least 3 thicknesses.

Refractive Index

The angle of each Fourier coefficient gives values for $\phi(\omega)$ and $\phi_0(\omega)$ for (5.4), enabling essentially the same method to be applied for $n(\omega)$ as for $\alpha(\omega)$. The 5% relative transmission results from the $\alpha(\omega)$ were used to decide which pulses to include, and again $n(\omega)$ was only calculated where there were at least 3 thicknesses.

There is extra processing required however, as the angles of the Fourier coefficients are modulus 2π, i.e., range from $-\pi$ to π, a feature of the sine and cosine basis for this transform. These angles need to be ‘unwrapped’ to correctly reflect the phase data, and this is carried out by assuming that any jump in angle between adjacent coefficients greater than a threshold (usually π) is in fact ‘round the clock behaviour’. For example, if $\phi(m)$ is -3 and $\phi(m+1)$ is $+3$, $\phi(m+1)$ would be unwrapped to be $+3-2\pi$. Similarly if $\phi(m+2)$ was $+2.5$, this would be unwrapped to $2.5-2\pi$. Figure 5.5 shows an example of the phase data before and after this process.

There are two sources of error in this process. Firstly there is an inherent ambiguity in the unwrapping process — is a jump of exactly π due to modulus effects or not? Similarly if there are genuinely jumps of more than π these will be removed. This effect can be offset to some extent by padding out the time series with 0, which decreases the gaps between the Fourier coefficients, hence making large ambiguous jumps less likely, and then decimating the unwrapped coefficients back to the original length. This technique is not a panacea however; not only does padding the time series increase the computational load, but it is only of limited benefit in terms of reducing the errors in the $n(\omega)$ values.

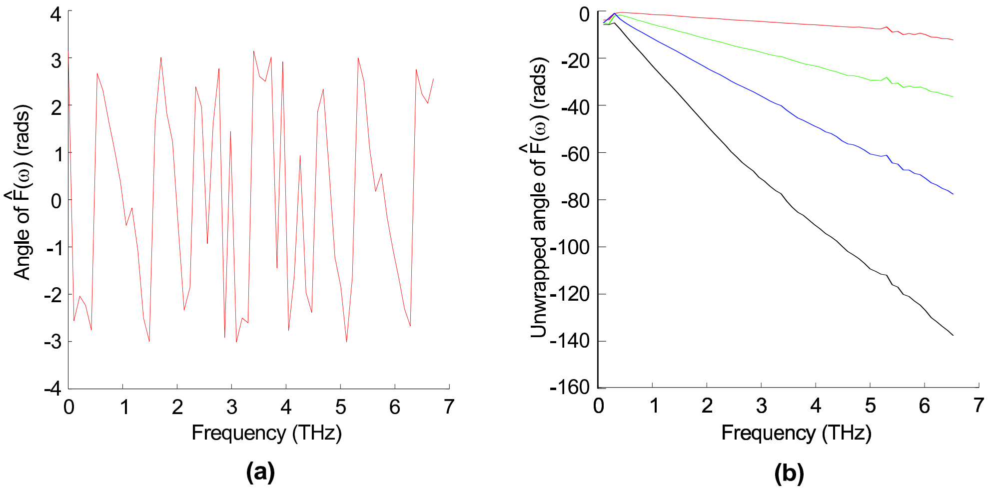

Figure 5.5: (a) The raw angle of Fourier coefficient. (b) The result of unwrapping the angle of the Fourier coefficients for the reference pulse from (a) (red), and the three thinnest steps of nylon.

Indeed, after a certain level of padding it was found that the errors started increasing. In this study the time series was padded until the errors in $n(\omega)$ were at a minimum. The nylon and resin data were padded from 64 samples to 256 samples, and the biological data were padded from 256 samples to 1024 samples.

The second source of error is cumulative error, whereby phase incorrectly unwrapped in the first few coefficients carries down through the entire profile. This leads to phase profiles for the same thickness of material proceeding down parallel paths that are multiples of 2π apart. The first few coefficients cannot be simply ignored as they establish the relative positions of the profiles for different thicknesses of materials. To ignore them would result in a change in separation between the profiles, which in turn would affect the gradiant of the phase thickness plot, leading to an incorrect $n(\omega)$ value. This is particularly an issue with noisy data — the spectral content of the terahertz pulses tends to be less for the first few coefficients, impacting the signal to noise ratio at these frequencies, leading to noisy phase values. This error was minimised by skipping the noisy coefficients and then taking the mean of all the phase profiles for that thickness of material. The profile closest to this mean was taken to be the correct interpretation, and the jumps in the noisy section were adjusted to remove any 2π offsets. It was found that skipping the first 10% of coefficients minimised the error in $n(\omega)$.

Broadband Analogues

The broadband pulse specific analogues of $n(\omega)$ and $\alpha(\omega)$, n and α respectively, were calculated for comparative problems. In this case, (5.6) and (5.7) were used for the linear regression plots. As these analogues are not frequency specific, there is just a single value obtained for any given material. Once again the 5% relative transmission inclusion criterion was applied, and every material gave values for at least 3 thicknesses.

These analogues are useful as a baseline comparison against which to compare the alternative techniques. n is particularly useful given the unwrapping issues involved when calculating $n(\omega)$. They are also the properties that influence image contrast for images based on parameters calculated in the time domain.

5.2.3 STFT

The short time Fourier transform, introduced in Chapter 3, was used to calculate $n(\omega)$ and $\alpha(\omega)$. The STFT results in a 2D array of complex coefficients, β and ξ. The magnitude of each coefficient represents the magnitude of signal present at that time (β) and frequency (ξ). For every given ξ, the maximum magnitude of coefficient was used for calculating $\alpha(\omega)$, and the β at which that coefficient occurred was used for calculating $n(\omega)$. By performing this operation on both the reference pulse and the sample pulses, comparable values are obtained which may be used in the complex refractive index calculations.

The STFT requires two parameters to be set before analysis — the window function and the width of that function. Three window functions were compared in this analysis — a Gaussian function, a rectangular function, and a triangular function. These three window functions are 'simple' and well understood analytically, and, particularly in the case of the last two, computationally 'light'. However they do capture a range of properties of windowing functions, from the smooth (Gaussian) to the sharply discontinuous (rectangular), and intermediate (triangular). The width of the Gaussian function was taken to be the standard deviation of the distribution, whereas the width of the other two distributions were taken to be the distance from the mid-point of the distribution to its edges, that is where its value is 0. No normalisation was used with the rectangular or triangular windows since relative values and ratios were used.

Finally, the short acquisition study was carried out by centering rectangular windows of various widths on the peak of the pulse on the time-domain for each pulse in the dataset. Following the windowing of the pulses, traditional frequency domain analysis was carried out in order to plot the absorption coefficient profiles. In this way a collection of absorption coefficient profiles using various window widths was created. The window widths were chosen from 1 sample up to 126 samples, in steps of 5, where the original pulses consisted of 128 samples in total.

5.2.4 WBCAF

A refractive index analogue was calculated using the wide band cross ambiguity function introduced in Chapter 3. The WBCAF was used to estimate the time-delay at scale 1, and this estimate was then used in (5.6) to give the WBCAF absorption parameter. The time delay was estimated by finding the translation τ which gave the maximum WBCAF value; this corresponds to the amount the reference pulse must be shifted in order to achieve the largest correlation, and hence best match, with the sample pulse. The WBCAF cannot be used to estimate time-delay at different scales, since the scaling operation causes a shift in the reference pulse.

The value of the WBCAF at any given scale and translation will be related to the intensity of waveform with respect to the mother function, and this value was normalised at a given scale σ using (5.8), giving an estimate of relative transmission at that scale.

$$\frac{MAX_{\tau}(WBCAF_{f_1}f_2[\tau,\sigma])}{MAX_{\tau'}(WBCAF_{f_1}f_1[\tau',\sigma])} \tag{5.8}$$

This estimate of relative transmission was plotted against step thickness in order to obtain an absorption parameter in a manner similar to the absorption coefficient calculations. In this way a profile of the WBCAF absorption parameter against scale can be determined. This absorption parameter differs significantly from the absorption coefficient in that there is no exact correspondence between scale and frequency. Each σ will correspond to a range of frequencies, and this range will vary with scale. This is unlike the STFT, where there is an exact correspondence between ξ and frequency. A further difference is that the band of frequencies represented varies inversely with σ - small σ corresponds to high frequency content, and large σ corresponds to low frequency content.

5.3 Results

5.3.1 Traditional Analysis

Refractive Index

Calculating the refractive index analogue using the peak to peak difference in the time domain as an estimate of time delay for (5.6) gives rise to a broadband refractive index, *n*. Figure 5.6 shows the plot of peak to peak time delay against thickness for the resin step wedge, where the gradient of the best fit line corresponds to the refractive index. The *n* values calculated in this way are shown in Table 5.1.

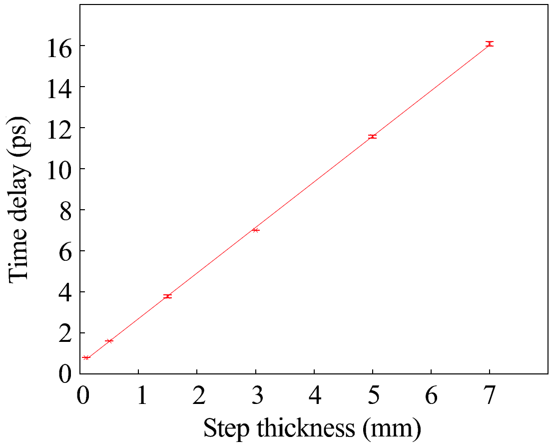

Figure 5.6: Time difference against step depth for a resin step wedge.

Table 5.1: The broadband refractive index value, calculated in the time domain.

Material	Refractive Index
Nylon	1.604 ± 0.004
Resin	1.666 ± 0.009
Skin	1.8 ± 0.2
Fat	1.49 ± 0.45
Muscle	2.1 ± 0.4
Vein	1.60 ± 0.77
Artery	2.75 ± 0.64
Nerve	2.1 ± 0.5

Finally, using the change of phase and (5.4) the profiles of $n(\omega)$ were created, and these profiles are shown in Figure 5.7.

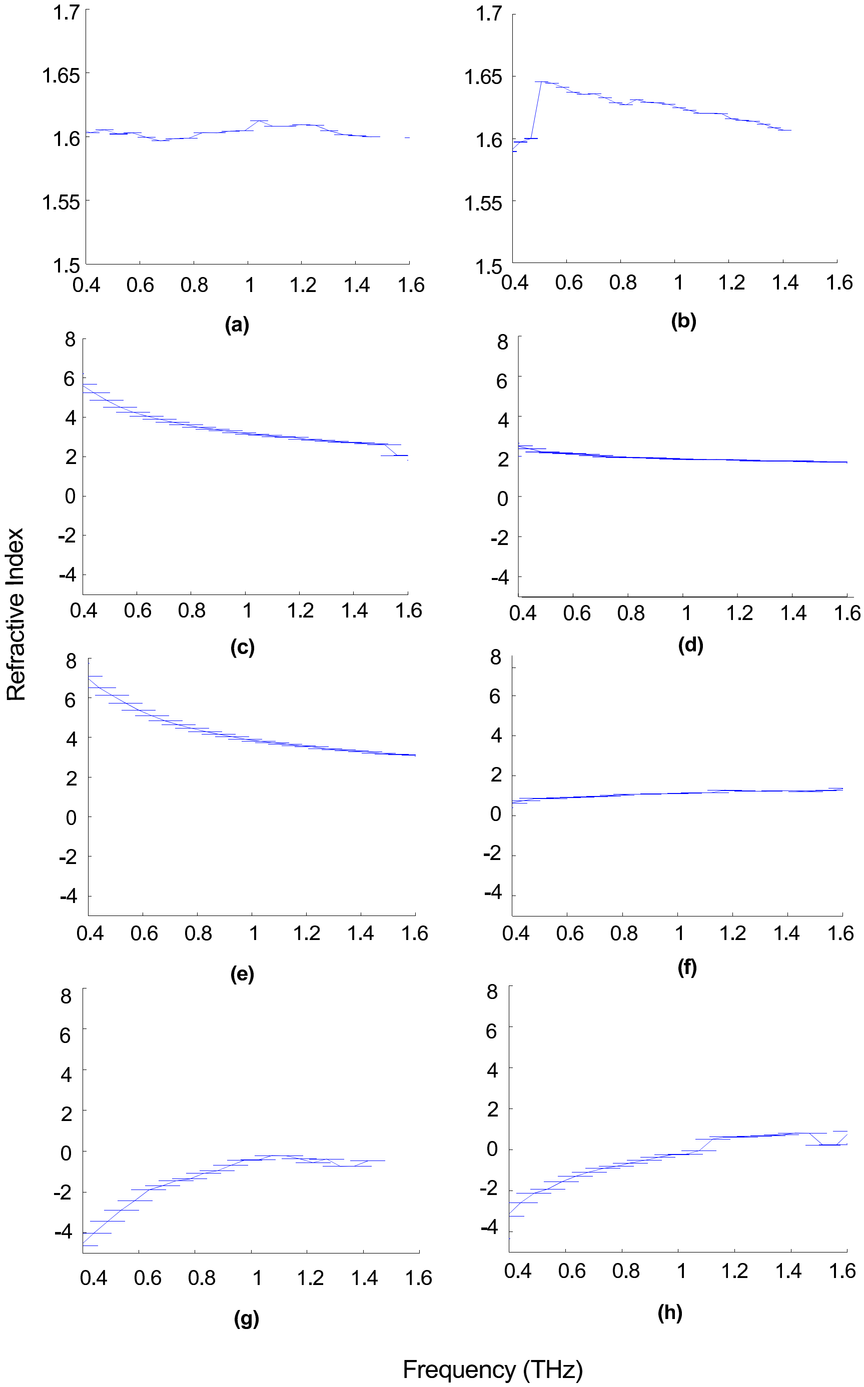

Figure 5.7: Refractive Index vs. frequency profiles for (a) nylon, (b) resin, (c) skin, (d) fat, (e) muscle, (f) vein, (g) artery, and (h) nerve.

In Figure 5.7 there are clearly problems with the biological samples. Although the standard error on the linear regression is fairly small, the refractive index values do not particularly correspond to the broadband values, and indeed in some of the values of n are negative! These inconsistencies are generally due to the extremely noisy nature of these data, and in particularly due to the difficulty of resolving the ambiguities when unwrapping the phase.

Absorption Coefficient

In a similar way to that of the broadband refractive index calculation, the peak to peak intensity ratio can be used as a measure of relative transmission in (5.7). However, this gives rise to a broadband value α, which may not be linear with thickness and thus not be described correctly by Beer's law, as shown in Figure 5.8.

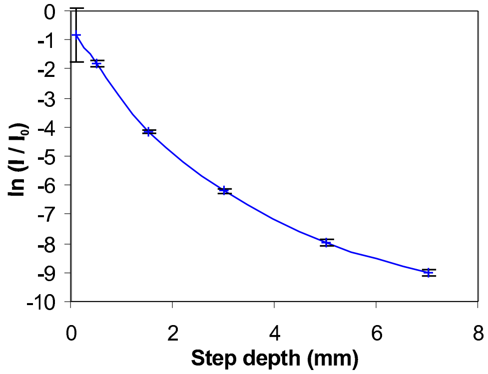

Figure 5.8: Broadband relative transmission against step depth for a resin step wedge.

The nonlinear relationship between the logarithm of broadband relative transmission and step depth shown in Figure 5.8 renders the linear regression technique inappropriate, and so this analogue was not calculated.

Figure 5.9 shows the logarithm of relative transmission against step thickness plotted for a high frequency (1.5 THz) and a low frequency (0.5 THz), together with the pulse height ratio plot from Figure 5.8 for ease of comparison.

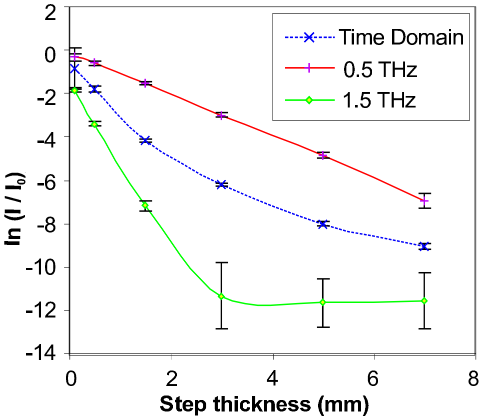

Figure 5.9: The natural logarithms of relative transmissions at 0.5 THz and 1.5 THz against step thickness for a resin step wedge, with the broadband time domain based values included for comparison.

Figure 5.9 clearly demonstrates that absorption changes with frequency, as different gradients result for 0.5 THz and 1.5 THz. This suggests $\alpha(\omega)$ varies across this frequency range in this material, as onewould expect, and indeed $\alpha(\omega)$ varies for all the samples in this frequency range. Figure 5.9 also highlights the need for inclusion criteria in these calculations - the thickest two resin steps no longer demonstrate Beer's law at 1.5 THz, since all the radiation of and above this frequency has effectively been absorbed.

The $\alpha(\omega)$ of each material was calculated, and are plotted in Figure 5.10. Below around 0.5 THz and above 1.5 THz the signal is swamped by noise, especially with the thicker samples of material, and so the profile is only shown within these boundaries. Results are reported for nylon for the slightly lower frequency band 0.13 to 0.48 THz by Birch et al. [5], where the values for $\alpha(\omega)$ are comparable with ours.

The large errors in the biological samples are believed to be the result of a combination of factors, including the normal biological variation between the two donors, the hydration state of the tissue varying as the imaging progressed, inevitable small differences in the thickness of tissue presented, and scattering losses.

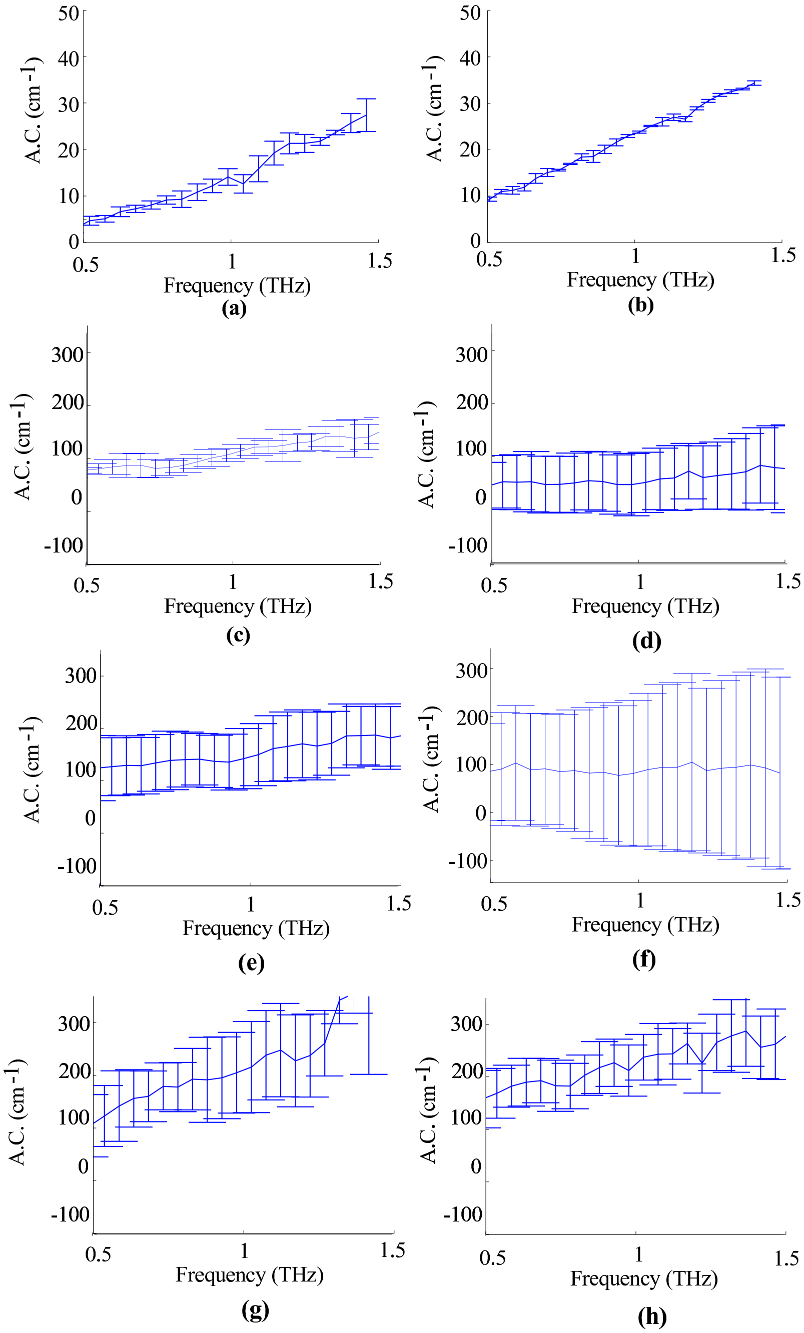

Figure 5.10: Absorption coefficient vs. frequency profiles for (a) nylon, (b) resin, (c) skin, (d) fat, (e) muscle, (f) vein, (g) artery, and (h) nerve.

5.3.2 STFT

Figure 5.11 shows the output from the STFT when it is applied to terahertz pulses. The Gaussian window function was used with a width of 9 ps (60 samples) to demonstrate the time/frequency nature of the STFT. Notice that the retardation of the pulse by the nylon is visible in the STFT spectrogram, as the larger coefficients (denoted by darker grey level) are further to the right than the reference pulse. The attenuation is also visible, as the coefficients are all smaller (lighter grey levels). Finally pulse broadening can be seen, as the widest part of the nylon pulse's STFT coefficients is wider than those of the reference pulse's STFT.

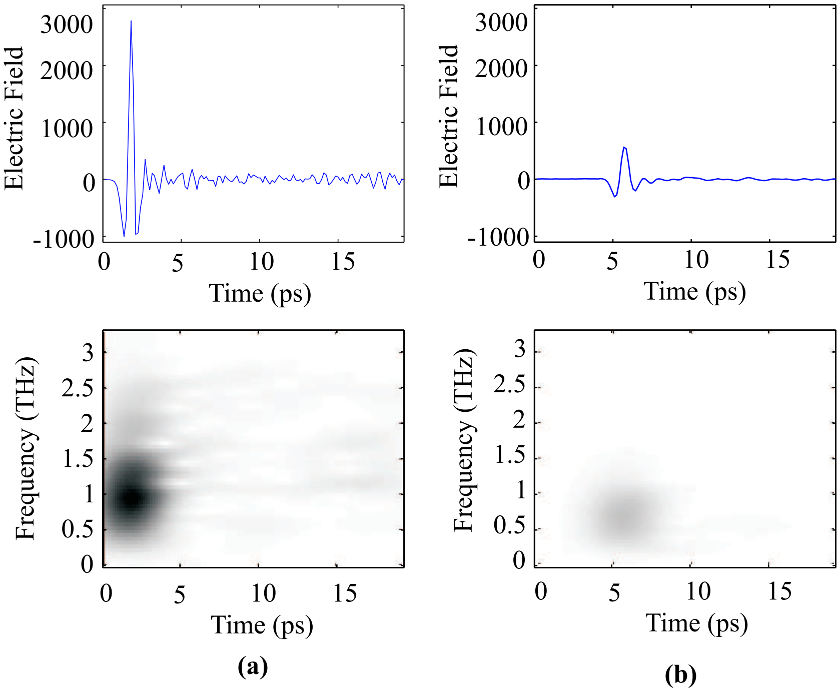

Figure 5.11: Time domain of terahertz pulse and corresponding STFT coefficients for (a) the reference pulse, and (b) a pulse through 2 mm of nylon. The grey level represents the mutually normalised magnitude of the coefficient, with black indicating 1 and white indicating 0.

Refractive Index

Figure 5.12 shows three examples of the graph of step thickness against pulse delay for different window width of the STFT. The STFT was performed with the Gaussian window function, and the delays taken were for 1 THz. Note how the error bars increase and the straight line fit is less good as window width increases. This is to be expected, as the resolution balance changes from being temporally biased to being spectrally biased.

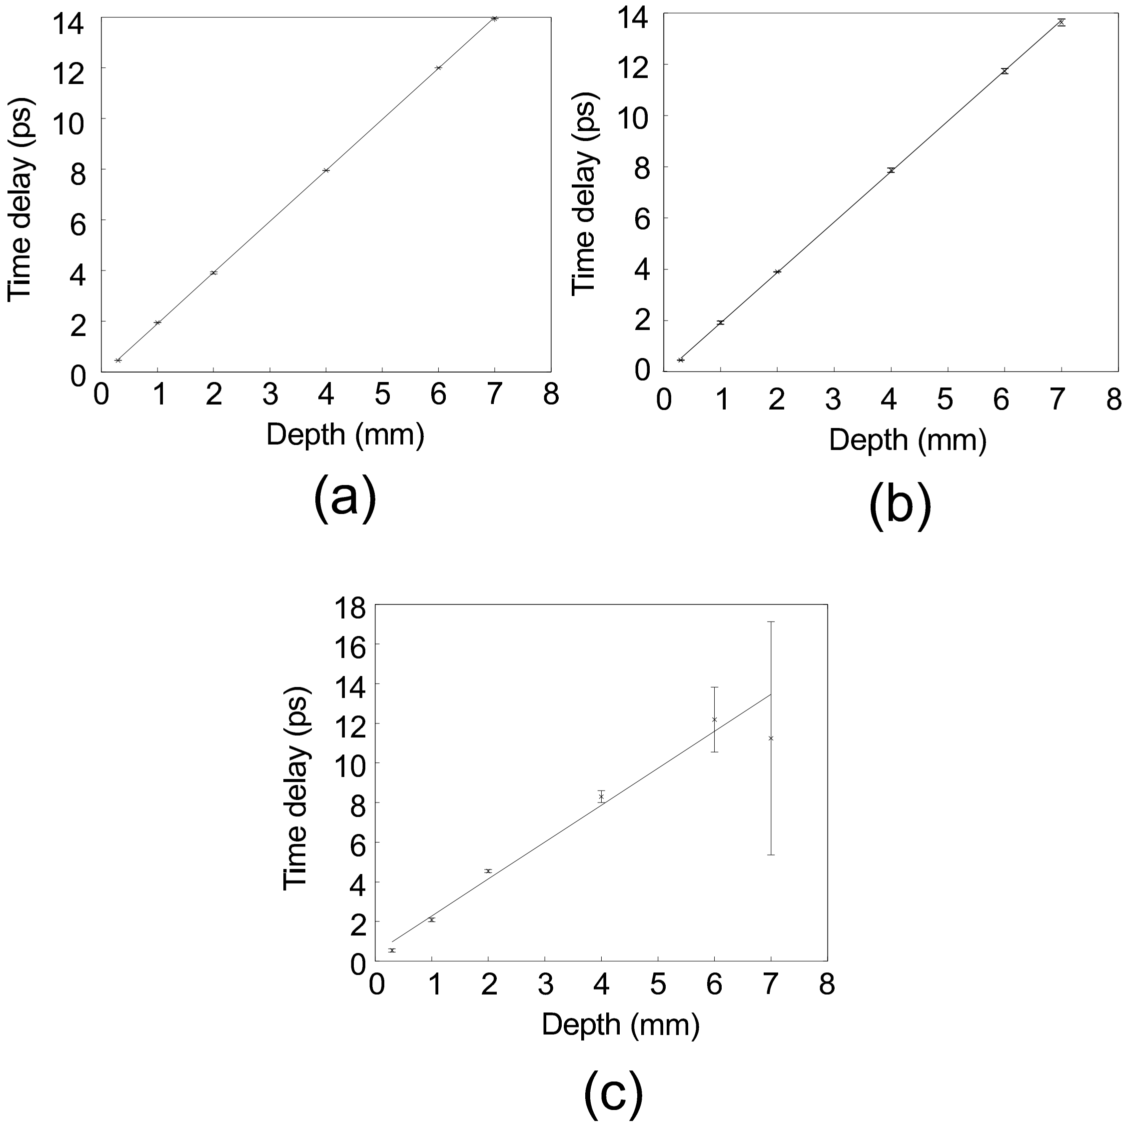

Figure 5.12: Time delay vs step thickness for a nylon step wedge at 1 THz. Calculated using the STFT with a Gaussian window of widths (a) 10^{-1} (0.015 ps) (b) 10^{1} (1.5 ps) (c) $10^{3.5}$ samples (474 ps).

Figure 5.13 shows the refractive index profiles of the nylon step wedge, calculated using the STFT with a Gaussian window, with widths from 10^{-1} to $10^{3.5}$. These widths correspond to around 0.015 ps and 474 ps respectively. The exponential changes are required to fully demonstrate the effect of changing the window width.

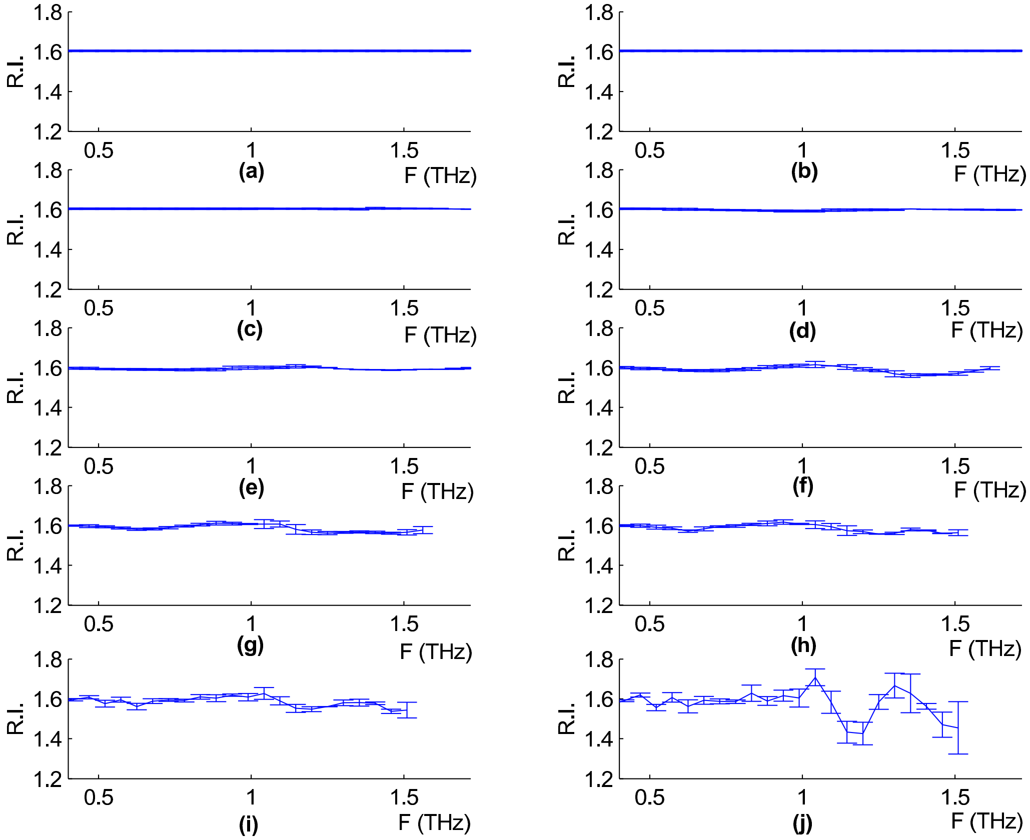

Figure 5.13: Refractive index profiles for a nylon step wedge using the Gaussian windowed STFT with widths (a) 10^{-1} (0.015 ps), (b) $10^{-0.5}$ (0.047 ps), (c) 10^{0} (0.15 ps), (d) $10^{0.5}$ (0.47 ps), (e) 10^{1} (1.5 ps), (f) $10^{1.5}$ (4.7 ps), (g) 10^{2} (15 ps), (h) $10^{2.5}$ (47 ps), (i) 10^{3} (150 ps), and (j) $10^{3.5}$ (474 ps).

Figures 5.14 and 5.15 show the equivalent figures for the triangular and rectangular windows respectively. Rectangular windows wider than 10 samples gave refractive index results that were completely swamped with noise, and so the profile is only shown for windows narrower than this.

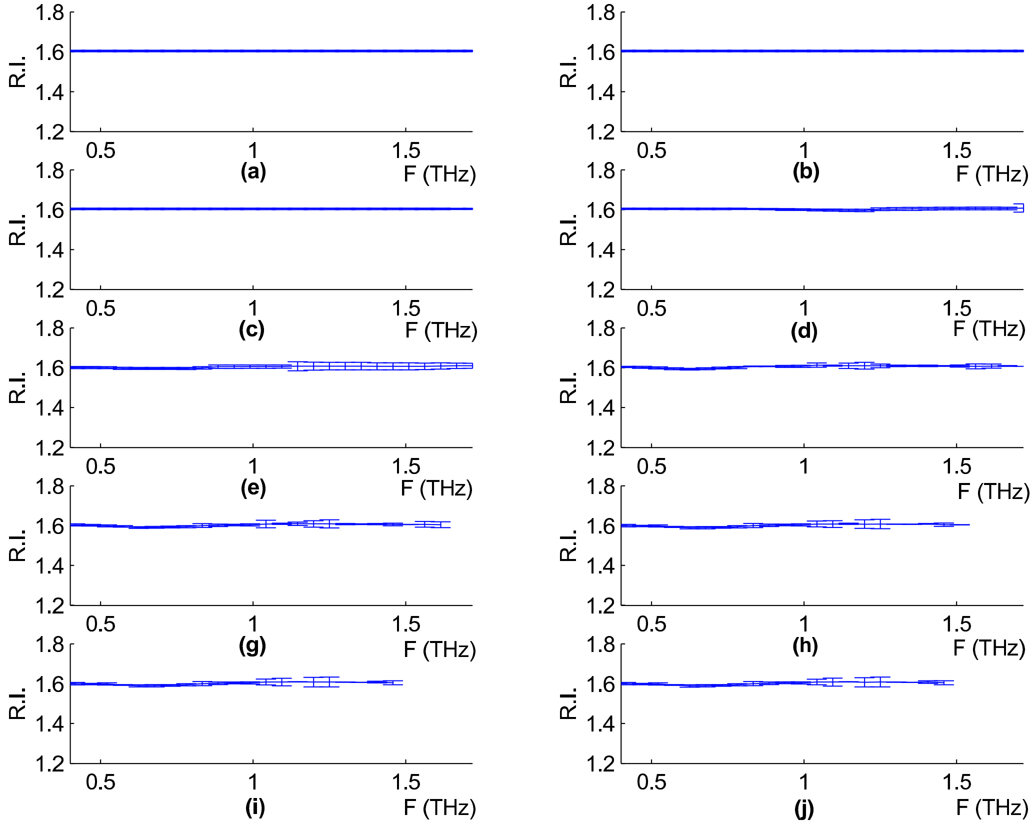

Figure 5.14: Refractive index profiles for a nylon step wedge using the triangular windowed STFT with widths (a) 10^{-1} (0.015 ps), (b) $10^{-0.5}$ (0.047 ps), (c) 10^{0} (0.15 ps), (d) $10^{0.5}$ (0.47 ps), (e) 10^{1} (1.5 ps), (f) $10^{1.5}$ (4.7 ps), (g) 10^{2} (15 ps), (h) $10^{2.5}$ (47 ps), (i) 10^{3} (150 ps), and (j) $10^{3.5}$ (474 ps).

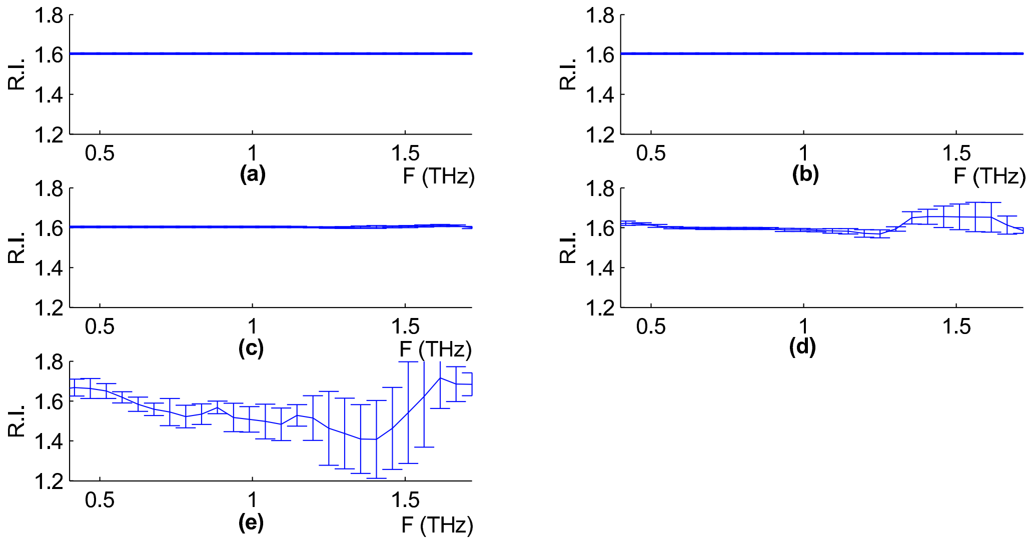

Figure 5.15: Refractive index profiles for a nylon step wedge using the rectangular windowed STFT with widths (a) 10^{-1} (0.015 ps), (b) $10^{-0.5}$ (0.047 ps), (c) 10^{0} (0.15 ps), (d) $10^{0.5}$ (0.47 ps), and (e) 10^{1} (1.5 ps).

As expected, the narrowest window results in a completely flat refractive index profile at the level of the broadband value, no matter which window shape is used. This is because the narrow limit of the STFT is effectively calculation in the time domain, with no frequency dependence. Of course, the Fourier based refractive index profile also demonstrates little variation against frequency, and this is reflected in the plots for most of the window widths. The triangular window function appears to be the most stable for this calculation, whereas the Gaussian and rectangular functions both demonstrate large fluctuations away from truth as the windows get wider. However, the Gaussian window function does demonstrate the transition from temporal to spectral emphasis. How the triangular window function introduces spectral artefacts has also been seen, and so only the Gaussian window has been used for the rest of the samples. These results are shown in Figure 5.16. The complete set of graphs showing how the refractive index profile varies with window width for each sample type may be found in Appendix B.

All of these graphs show refractive indices comparable with the broadband values calculated in the time domain. The skin, artery, and nerve analyses all give values that

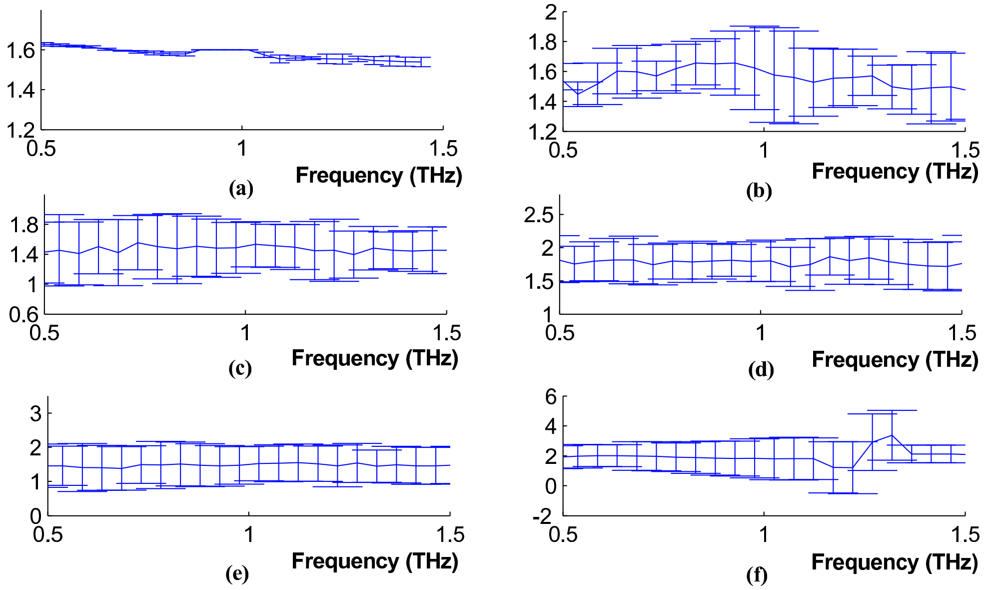

Figure 5.16: Refractive index profiles calculated using the Gaussian windowed STFT with width $10^{1.5}$ (6.3 ps) of (a) a resin step wedge, (b) excised skin, (c) excised fat, (d) excised muscle, (e) excised vein, and (f) excised artery

are somewhat lower than the broadband value, although still within the bounds of error from these calculations. Only the fat and vein refractive indices bear much resemblance to the Fourier calculated refractive indices. The dissimilarity on the other samples is almost certainly due to the problems in unwrapping the phase of these noisy biological data.

In every case, it seems to be windows of 10 samples or wider that give rise to refractive indices which are dependent on frequency, and it is generally true that within this frequency band the refractive index seems to be almost constant.

Absorption Coefficient

The absorption coefficient at a given frequency is calculated from the gradient of the logarithm of the ratio of incident to transmitted intensities versus step thickness. Figure 5.17 shows three examples of such gradients, all using the STFT with the Gaussian window function and calculated at 1 THz but with different widths of window function. By 1 THz, the deepest nylon steps have less than 5% transmittance, which results in a very low signal to noise ratio. To avoid this noise, only steps with transmittance greater than 5% are used. The straight line shows the gradient extracted by this method — the solid part indicates which steps thicknesses were used in calculating this gradient.

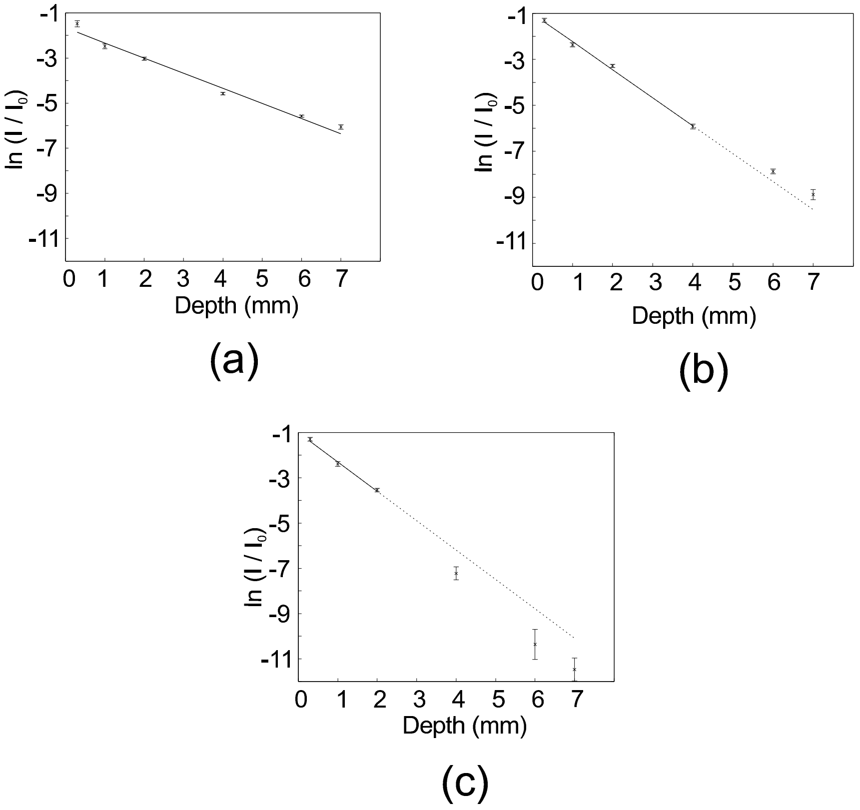

Figure 5.17: Logarithm of transmittance vs step thickness for a nylon step wedge at 1 THz. Calculated using the STFT with a Gaussian window of widths (a) 10^{-1} (0.015 ps) (b) 10^{1} (1.5 ps) (c) $10^{3.5}$ samples (474 ps). The points near the solid line are those that were used in calculating the gradient, the others were excluded by the noise criteria.

As before, applying the STFT to the terahertz data enables us to calculate the absorption coefficient profiles for different window widths and window functions. Figures 5.18 to 5.20 show the profiles of the nylon step wedge for a Gaussian, triangular, and rectangular window function respectively.

In the case of the absorption coefficients, the results become worse as the window gets narrower. This is entirely expected, since the absorption is calculated from the intensities of frequencies present. The cut-off point at which the 1 THz absorption feature is lost when using windows narrower than $10^{1.5}$ samples (4.7 ps) for the triangular or Gaussian window functions, and 10^2 samples (15 ps) for the rectangular window function. There is little to differentiate the three different window functions, except that the rectangular window has a smaller range of widths within which its results are meaningful.

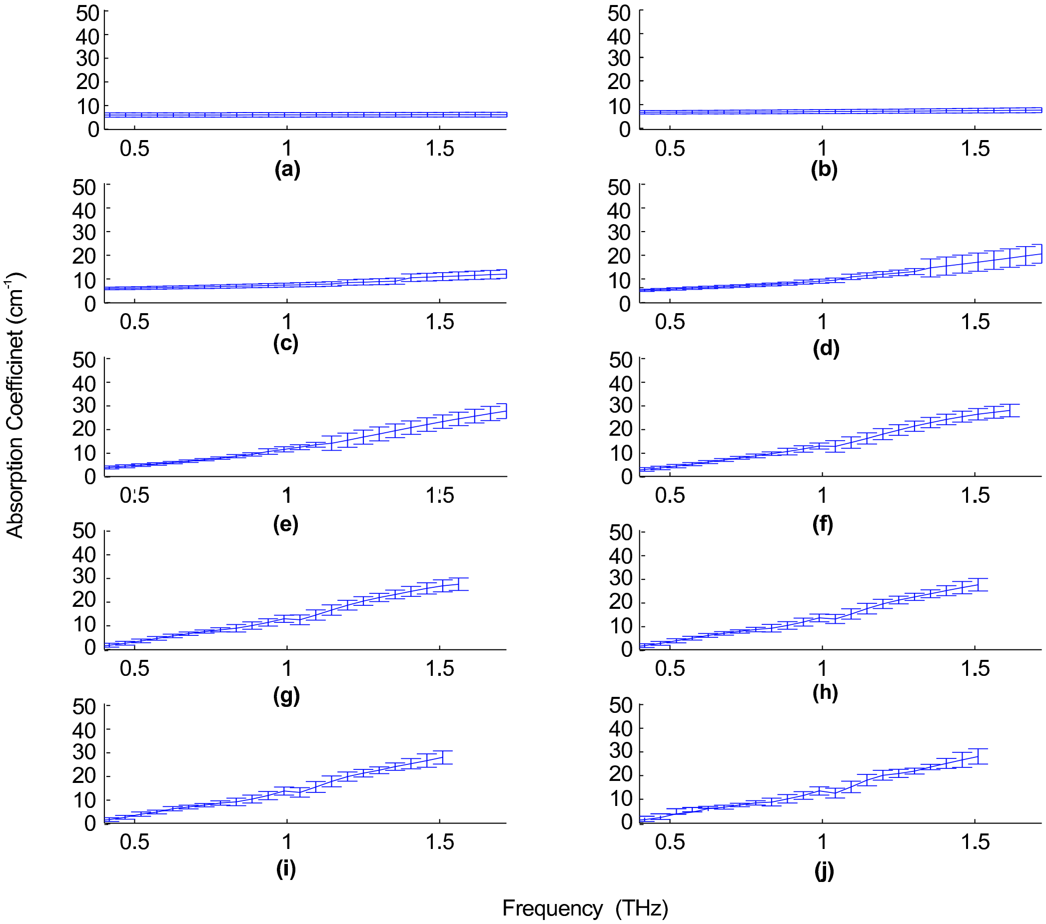

Figure 5.18: Absorption coefficient profiles for a nylon step wedge using the Gaussian windowed STFT with widths (a) 10^{-1} (0.015 ps), (b) $10^{-0.5}$ (0.047 ps), (c) 10^0 (0.15 ps), (d) $10^{0.5}$ (0.47 ps), (e) 10^1 (1.5 ps), (f) $10^{1.5}$ (4.7 ps), (g) 10^2 (15 ps), (h) $10^{2.5}$ (47 ps), (i) 10^3 (150 ps), and (j) $10^{3.5}$ (474 ps).

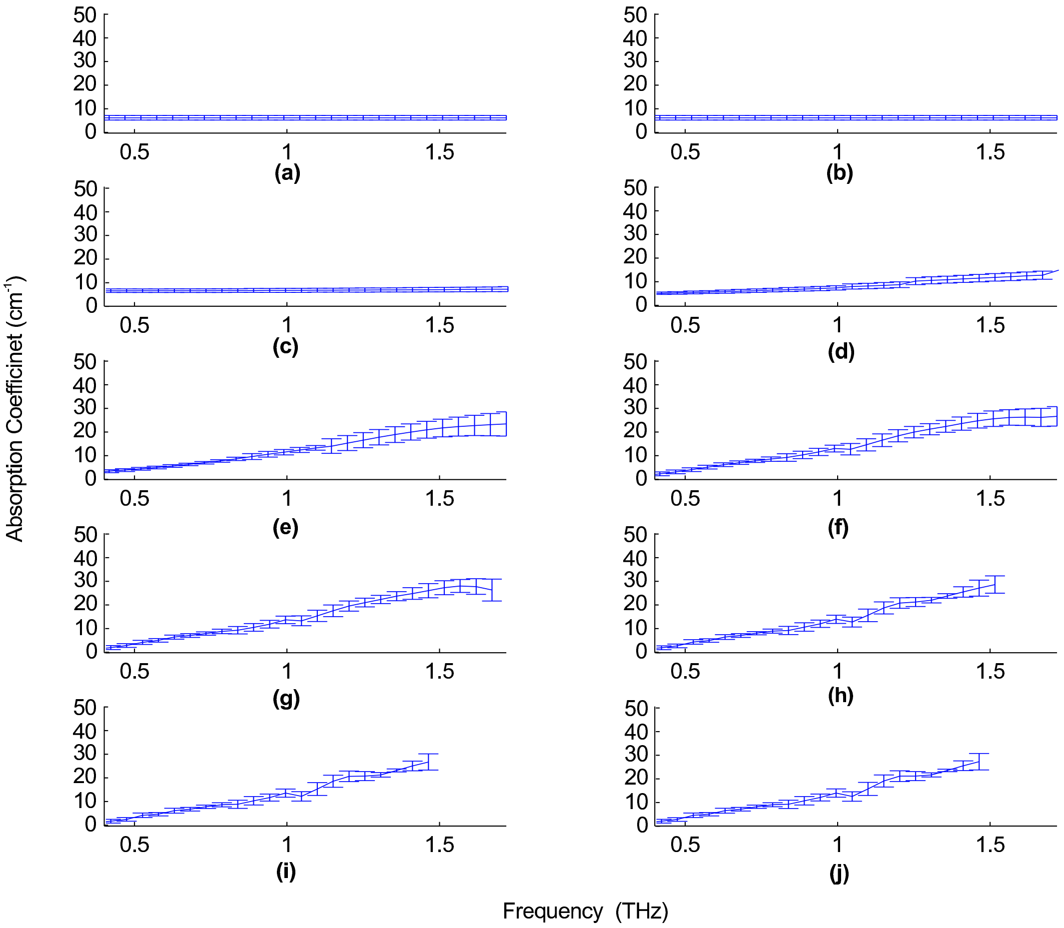

Figure 5.19: Absorption coefficient profiles for a nylon step wedge using the triangular windowed STFT with widths (a) 10^{-1} (0.015 ps), (b) $10^{-0.5}$ (0.047 ps), (c) 10^{0} (0.15 ps), (d) $10^{0.5}$ (0.47 ps), (e) 10^{1} (1.5 ps), (f) $10^{1.5}$ (4.7 ps), (g) 10^{2} (15 ps), (h) $10^{2.5}$ (47 ps), (i) 10^{3} (150 ps), and (j) $10^{3.5}$ (474 ps).

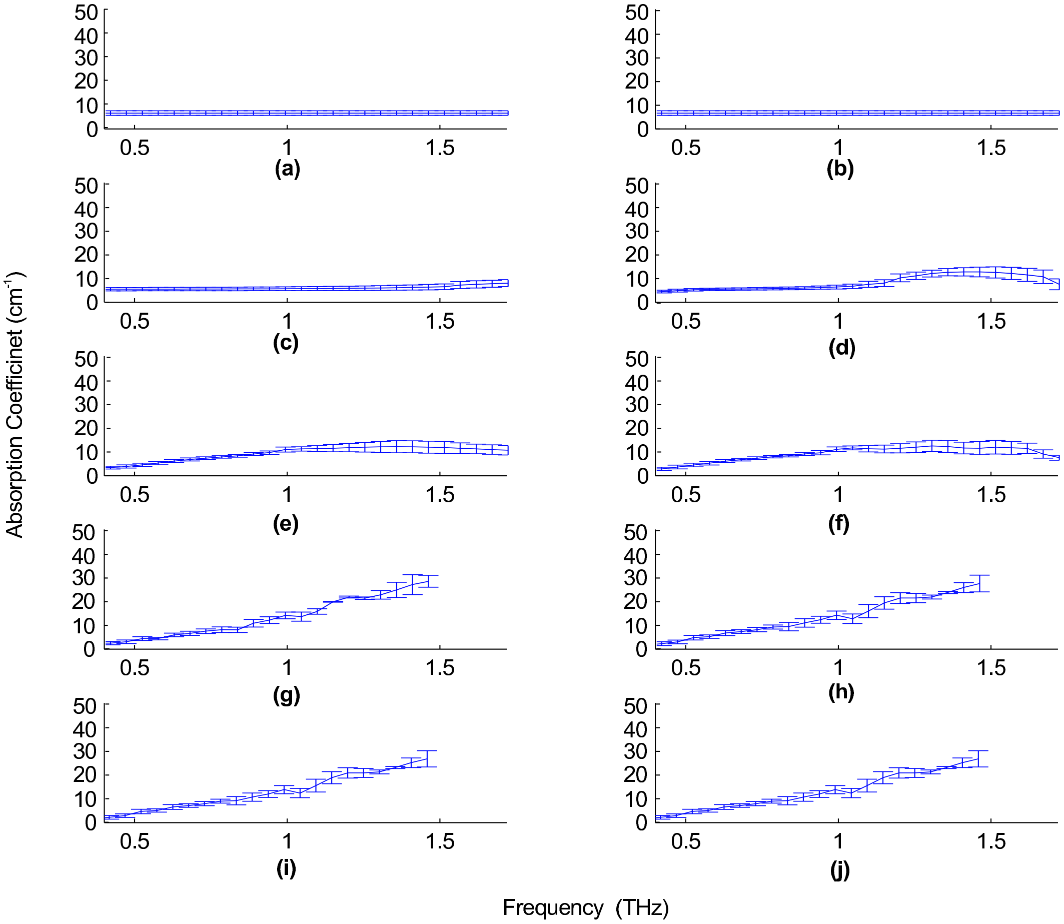

Figure 5.20: Absorption coefficient profiles for a nylon step wedge using the rectangular windowed STFT with widths (a) 10^{-1} (0.015 ps), (b) $10^{-0.5}$ (0.047 ps), (c) 10^{0} (0.15 ps), (d) $10^{0.5}$ (0.47 ps), (e) 10^{1} (1.5 ps), (f) $10^{1.5}$ (4.7 ps), (g) 10^{2} (15 ps), (h) $10^{2.5}$ (47 ps), (i) 10^{3} (150 ps), and (j) $10^{3.5}$ (474 ps).

Figure 5.21 shows how the absorption coefficient varies against frequency with a Gaussian window function of $10^{1.5}$ samples. The complete set of graphs showing how the absorption coefficient profile varies with window width for each sample type may again be found in Appendix B.

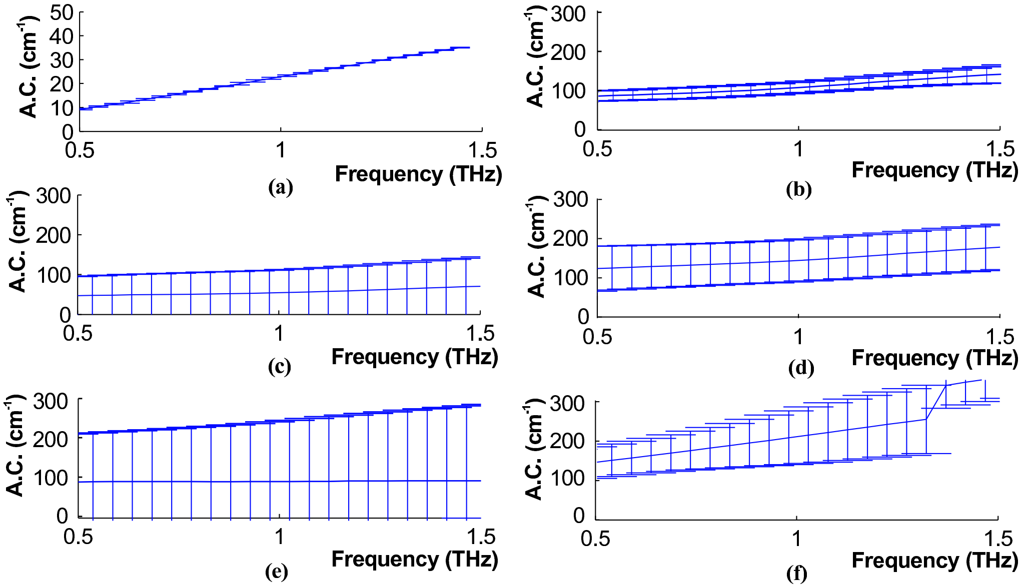

Figure 5.21: Absorption coefficient profiles calculated using the Gaussian windowed STFT with width $10^{1.5}$ (6.3 ps) of (a) a resin step wedge, (b) excised skin, (c) excised fat, (d) excised muscle, (e) excised vein, and (f) excised artery.

In these graphs the absorption coefficients mimic the refractive indices in that the narrowest window functions give an unchanging absorption coefficient against frequency, meaning that little spectral information is being used in the calculation. The absorption coefficient profile obtained using windows of $10^{1.5}$ samples and wider follow the basic trend of the Fourier based profiles, although in each case there is clearly a smoothing effect, as would be expected through the application of a Gaussian window.

These results together suggest that, for these acquisitions, an STFT analysis with a Gaussian window function of width of around 31 or 32 samples strikes an appropriate balance between spectral and temporal resolution, and thus enables a good approximation of the refractive indices and absorption coefficients to be calculated.

Short Acquisition Simulation

Figure 5.22 shows the absorption coefficient profiles of the nylon and resin step wedges when the pulses have been windowed with a rectangular window function of various widths. The widths are given in time rather than number of samples for ease of comparison between the two different scans, since the data were recorded with different sampling rates. The widths chosen demonstrate the narrowest window that gives comparable results to the unwindowed results, a medium width window, and a narrow window.

I found that a window width of at least 111 samples (16.5 ps) was required before all the features of the nylon step absorption profile were captured. Similarly 106 samples (21.20 ps) were required for the resin step wedge. The nylon step was originally imaged using 128 samples, or 19.2 ps. The resin step also used 128 samples, but since there was a larger gap between the sample points, this corresponds to 25.6 ps. Note these values are not the time taken to acquire the pulses, since each individual sample point takes many seconds to acquire. Instead, these figures refer to the amount of time it would take to capture the pulse if it could be acquired in real time.

However, the overall trend of the absorption coefficient against frequency is retained even with very small windows. Figure 5.22 demonstrates that the trend is broadly followed for even a 2.4 ps window (16 samples). With the resin step wedge, a 1.2 ps window (6 samples) results in an absorption coefficient profile that follows the true trend closely. These correspond to a total acquisition of 32 and 12 data points respectively, compared with 128 on the original pulses. This in turn corresponds to a reduction in acquisition time to about 25% and 10% of the original acquisition time respectively. Once smaller window functions than these were used, large errors were introduced in the α plots.

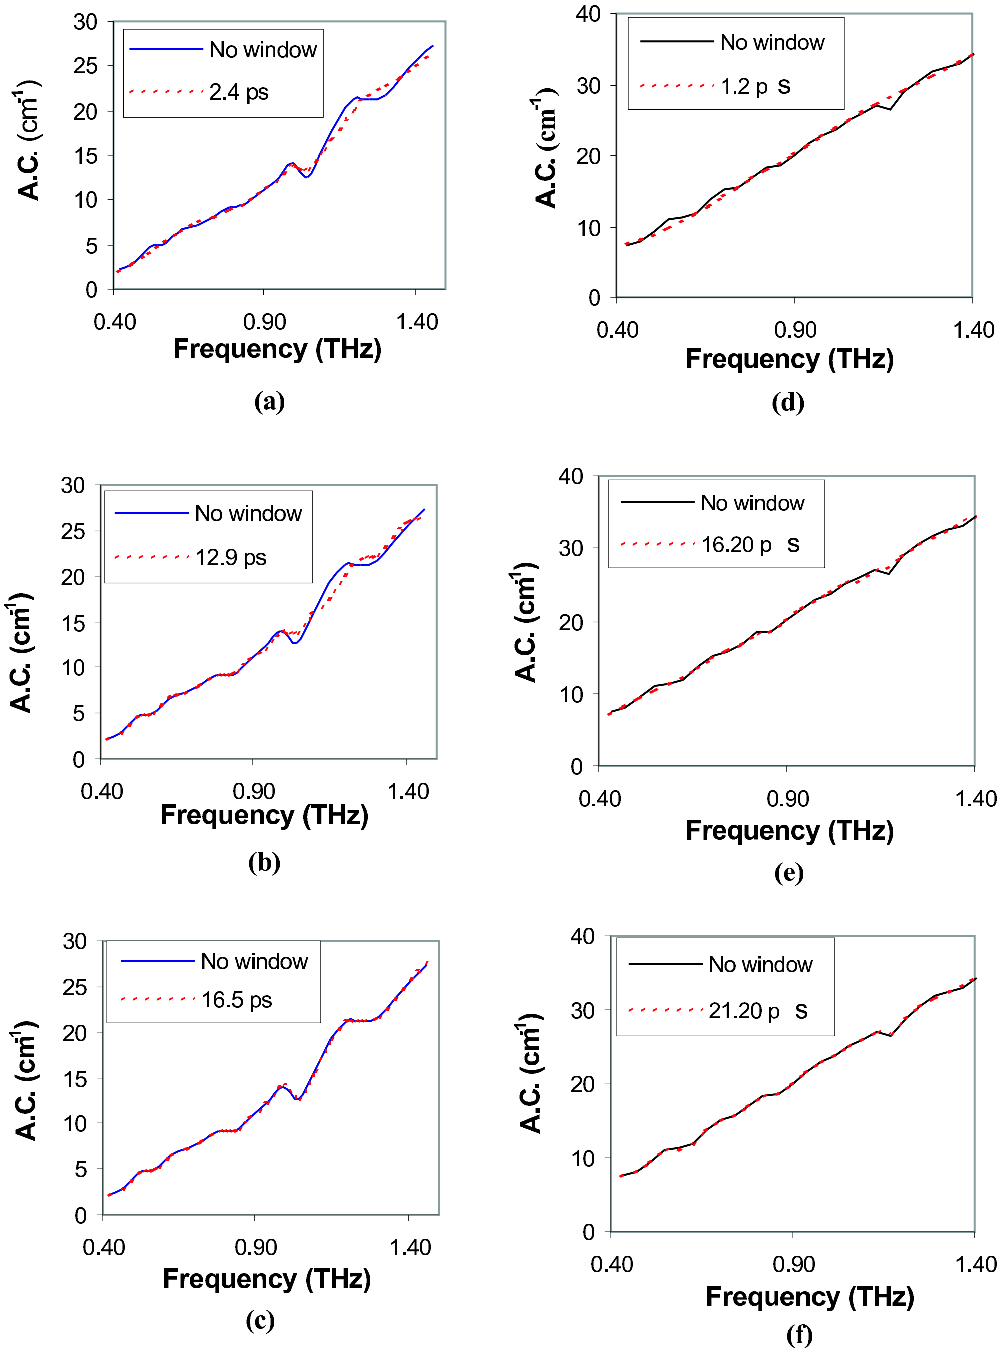

(a) (d)

(b) (e)

(c) (f)

Figure 5.22: The absorption coefficient profiles of fixed rectangular windowed data from (a)-(c) a nylon step wedge, and (d)-(f) a resin step wedge.

5.3.3 WBCAF

Figure 5.23 shows the coefficients when the WBCAF applies the reference pulse to a pulse through 2 mm of nylon. Notice that the coefficients show a drift to the left as the scale increases down the figure — this is due to the increasing scale having the side effect of shifting the pulse to the right.

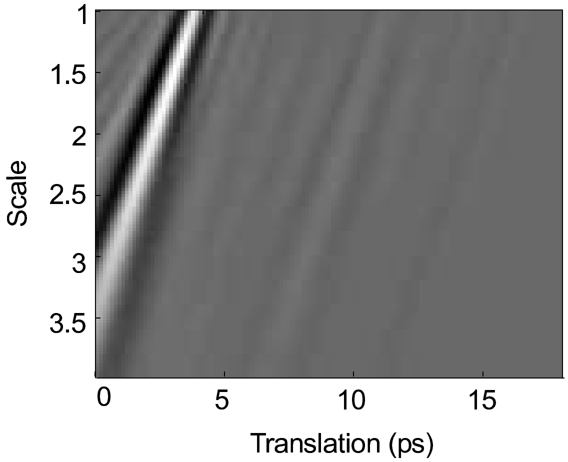

Figure 5.23: The WBCAF coefficients of a pulse though 2 mm of nylon. White pixels indicate the largest positive coefficients, black indicates the largest negative coefficients, and grey pixels in between, with most being almost zero.

5.3.4 Refractive Index

The refractive indices of nylon and resin were calculated using the translation of the maximum WBCAF parameter as an estimate of time delay, as described in Section 5.2.4. Figures 5.24(a) and (b) show plots of step depth against estimated time delay, together with the line of best fit, calculated using linear regression. Table 5.2 shows the values of refractive index calculated with the WBCAF, together with the broadband values for comparison. As with the broadband calculations, there is a very good fit achieved by linear regression, and this is reflected in the low error boundaries.

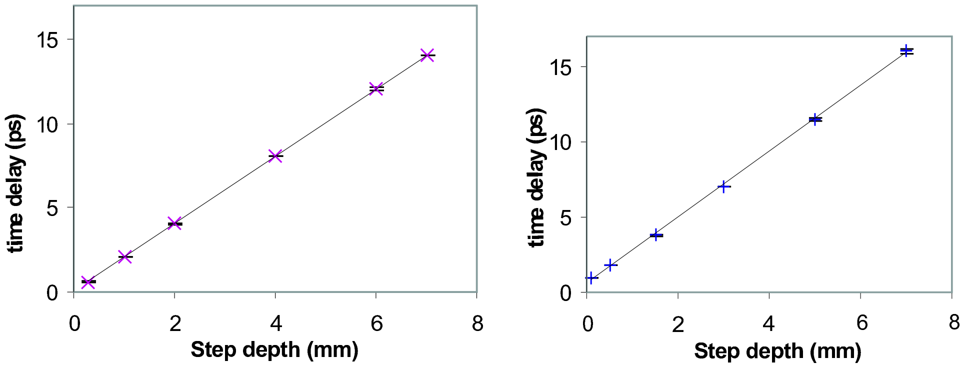

(a) (b)

Figure 5.24: WBCAF time delay estimate against step-depth with best fit line for (a) nylon, and (b) resin, both at scale 1.

	Nylon	Resin
Time Domain	1.603 ± 0.004	1.66 ± 0.01
WBCAF	1.597 ± 0.005	1.64 ± 0.02

Table 5.2: Refractive indices calculated by traditional methods and using a WBCAF.

5.3.5 WBCAF Absorption Parameter

The absorption parameter was calculated using the normalised maximum WBCAF coefficients from (5.8) to estimate relative transmission. Figure 5.25 shows the logarithm of this relative transmission for the nylon and resin step wedges.

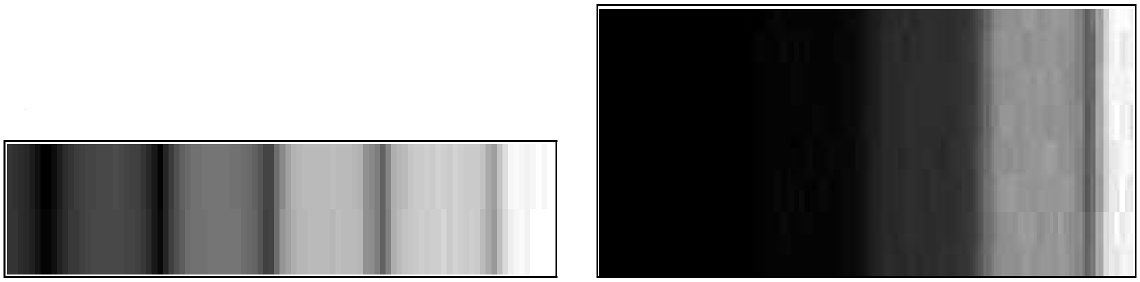

(a) (b)

Figure 5.25: Logarithm of the normalised WBCAF maxima at scale 1 for (a) nylon and (b) resin step wedges. Black indicates the smallest value, the most absorption, and white indicates the largest value, or least absorption.

Figure 5.26 shows the plots of the WBCAF estimated relative transmission against step depth for (a) resin, and (b) nylon, at scale 1. The points in the nylon plot approximate a straight line, however the resin plot displays a similar curve to the time domain based plot in Figure 5.8. This is to be expected given the larger absorption coefficient of resin compared to nylon, and the range of frequencies present within each individual WBCAF scale.

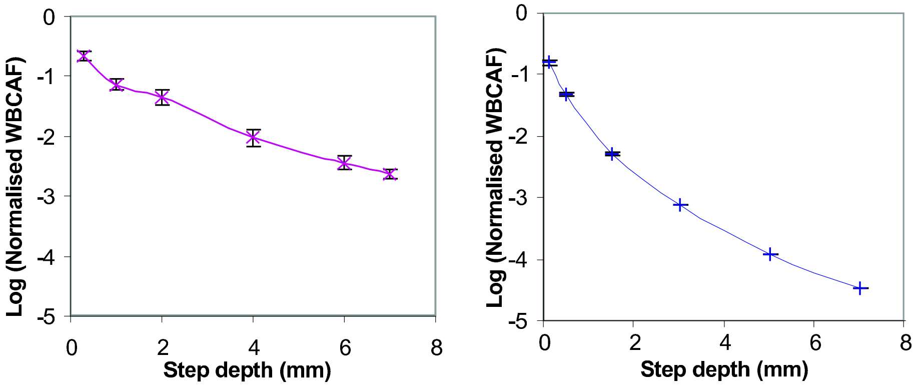

(a) (b)

Figure 5.26: Relative transmission, as estimated by the WBCAF, against step thickness for (a) nylon, and (b) resin.

Finally Figure 5.27 shows the profile of the WBCAF estimated absorption parameter against scale for both nylon and resin. The absorption coefficient values decrease with increasing scale as scale is inversely related to frequency content. This parameter clearly distinguishes between these two materials, although the degree of difference in absorption parameter is much larger then the difference in $\alpha(\omega)$.

5.4 Discussion

As the figures of refractive indices and absorption coefficients show, these parameters enable the different materials used in this work to be distinguished, at least at some frequencies.

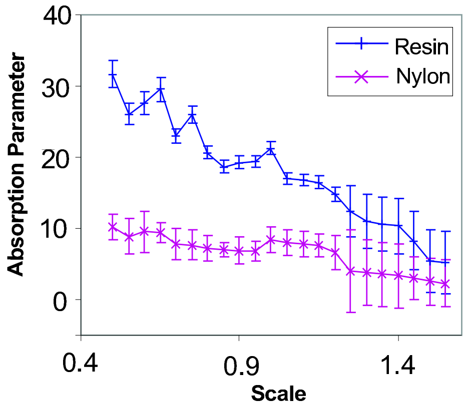

Figure 5.27: WBCAF absorption parameter against scale for nylon and resin.

5.4.1 STFT

The absorption coefficient profiles generated using the STFT are entirely consistent with the traditional analysis techniques. Naturally as the window function widens, the results come closer and closer to the Fourier based analysis. Similarly the refractive index results are consistent with the time-based values. As can be seen, the narrowest windows give refractive index values which are almost independent of frequency, but which correspond closely to the time domain calculated value. The STFT also enables the generation of $n(\omega)$ vs. frequency that doesn't depend on phase, and therefore avoids the difficulties associated with that method.

The rectangular window function performs most poorly, giving meaningful refractive index and absorption coefficient results with only a small range of window widths. There is no overlap between the two results — a window narrow enough to give meaningful refractive index results is too narrow to give absorption coefficient results, and a window wide enough to give absorption coefficient results is too wide to give refractive index results.

The right balance must be struck between time and frequency resolution. As has been shown, narrow windows achieve good temporal resolution at the expense of spectral resolution, while the opposite is true for wide windows. The obvious question is how wide (or narrow) a window must be in order to achieve meaningful results in both temporal and spectral spaces.

Fourier analysis gives 'perfect' results for $\alpha_{(}\omega)$, and so how $\alpha(\omega)$ varies with frequency is known. The α calculated by the STFT can therefore be evaluated against this ground truth. In this way, the *minimum* window width of a given window function needed to retain spectral resolution can be found. One can then be certain that using a window of this width (or wider) will give meaningful results in terms of spectral resolution.

The problem is then to find the *maximum* window width that can be used that still retains sufficient temporal resolution for the analysis of refractive index. We have no 'perfect' results for the refractive index against frequency, but if the refractive index values are constant against frequency, we can be confident that not enough spectral information is being used in that analysis, i.e., the window is too narrow. Similarly, large error bars and large fluctuations in the refractive indices (especially if it ever goes negative) indicate that not enough temporal information is being used, i.e., the window is too large. In this way an estimate can be obtained of the maximum possible window width at which temporal resolution is retained.

If the maximum window width usable to retain temporal resolution is smaller than the minimum width usable to retain spectral resolution (that is, there is no overlap of window widths), then the window function itself is unsuitable for joint time/frequency analysis. The rectangular window function falls into this category, although neither the triangular nor Gaussian functions do.

5.4.2 Simulation of short acquisition

The simulations of short acquisition results were disappointing; almost the entire pulse was needed to capture all the features of the absorption coefficient profile. The width of the rectangular window function refers to the distance from the centre of the window to its edge, that is when its value becomes zero. Thus the window is in reality twice the width of its ‘width’ parameter. This means that if the window is centered about the middle of the pulse then a window ‘width’ of 64 samples will capture the entire pulse for the nylon and resin data. If the window is centered on the peak of the pulse then the amount of the pulse that remains after the windowing will depend on the position of that peak. The thinnest nylon step has its peak at 15 samples, meaning that the 111 width window will capture the first 126 samples — effectively the entire pulse. The thinnest resin step has its peak at 17 samples, meaning that the 106 width window will capture the first 123 samples, just a little short of the entire pulse. In both cases, the thickest steps had pulses which were at 103 and 94 samples for the nylon and resin respectively. Both of these values mean that the entire pulse remains after windowing.

However, the overall trends of α_ω are captured, even when using very small windows (short acquisition times). Thus it may save valuable time if only a false colour dual- or tri-frequency image is being formed. An example of such an image might be colouring the pixels such that the red component corresponds to α at 0.5 THz, and the blue to α at 1.3 THz (Löffler et al. demonstrate this sort of visualisation in [46]). In such a situation the pulse would be rapidly scanned to locate the peak, and then the appropriate number of samples either side of that peak acquired.

The reduction in acquisition time to 10 or even 25 percent of the original acquisition time are also highly significant when it is considered that the original scans took over eight hours each.

5.4.3 WBCAF

The WBCAF results for the refractive index are entirely consistent with the broadband value, but the absorption parameters are harder to assess. Clearly the profiles for nylon and resin are quite different, and so this parameter shows promise as a contrast mechanism. The actual values are also the same order of magnitude as the Fourier based values.

Scale is not directly comparable to frequency, although frequency does correspond to inverse scale in some manner. With decreasing scale, and hence increasing frequency, the separation between the nylon and resin absorption increases, mainly due to the increase of the resin's absorption. This does not mirror exactly the true absorption coefficient values, although the absorption coefficient of resin is generally larger than that of nylon. This inconsistency with the true values for absorption coefficients does raise the obvious concern about whether the WBCAF values are actually meaningful and a useful metric. This is not a question that can be answered with the limited data currently available.

The WBCAF does also have a number of disadvantages over the traditional techniques. Firstly the method is very computationally expensive, as the reference pulse has to be interpolated and filtered for every scale in the analysis. Secondly, the effective shifting of the reference pulse causes difficulty with calculating the phase changes or time delays, and hence refractive indices. Finally the results are not wholly consistent with the traditional analyses.

5.5 Conclusions

The high dimensionality of terahertz data is a double edged sword — while providing a richness of data it also requires more sophisticated processing techniques than simple scalar data. In this chapter I have explored some techniques for reducing this high dimensionality through the calculation of the complex refractive index and its analogues.

The short time Fourier transform is the most promising of these, handling the noisy biological data while still giving results consistent with traditional analyses. This is a tool which might be especially useful in reflection mode imaging, when the ability to determine the spectral content at a given point in time becomes crucial. Further analysis needs to be carried out in order to determine whether the STFT analysis sufficiently distinguishes between materials to actually be useful. Similarly in this study a Gaussian window was shown to be an appropriate choice of window function, and a rectangular window was shown not to be.

The cross ambiguity function showed itself to be theoretically capable of distinguishing between nylon and resin based on refractive index and an absorption parameter, however the absorption parameter did not mirror the absorption coefficient, suggesting that it is not truly the absorption which has been calculated.

A joint time-frequency analysis has three distinct advantages over a traditional approach; it enables the refractive index to be unambiguously calculated as a function of frequency, it enables ‘one-shot’ signal processing, where both refractive indices and absorption coefficients are calculated in one pass over the data, and, perhaps most importantly, it compromises spectral resolution in order to gain temporal resolution, so that not only is the spectral content obtained, but also *when* that content occurred — a point that is of special interest in reflection mode imaging.

Finally, the effects that shortening the acquisition time have on traditional calculations of the optical properties were analysed. It was interesting to see that an extremely short acquisition time still led to an absorption coefficient that captured the general trend of the full-sample absorption coefficient. However, such short acquisition times do cause spectral features to be lost, and future work may demonstrate that these features are key in distinguishing between materials.

Chapter 6

Approaches to Segmentation of Terahertz Pulsed Imaging Data

6.1 Introduction

As discussed in Chapter 5, terahertz pulsed imaging delivers data of a high dimensionality. This high dimensionality raises the possibility of clustering techniques to segment terahertz images into their constituent regions. Groups of pixels which are 'similar' in some way are clustered together, and an image may then be formed by the assignment of different colours or greyscale levels to these groups. These images may then either be used as an end in themselves, or alternatively could be used as a first step that highlights potential areas of interest for further analysis.

The key advantage that this approach has over analyses based on the optical properties of the sample is that very little information about the sample is required. For instance, it is hard to calculate the complex refractive index if the depth of material is not known, or if several types of material are present in the sample. Clustering techniques, on the other hand, bypass these problems by just working on the theory that 'similar' pixels are likely to have similar constituency. Of course, a good similarity metric is vital if the clustering is to be successful, and this chapter explores some of the options in the choice of metric.

In this chapter, I apply *k*-means clustering to a variety of terahertz data, based on both high dimensional vectors formed from temporal, spectral, and combined domains, and on low dimensional vectors formed by extracting specific features such as absorption at a given frequency.

6.2 Method

6.2.1 Data

Synthetic Slice of Tooth

In order to assess the accuracy of the segmentations, a synthetic terahertz data set of a slice of tooth was created. A model of a slice of tooth was taken, consisting of three regions; air, enamel, and dentine, as shown in Figure 6.1(a). The number of pixels was then decreased to reflect the spatial resolution of terahertz, resulting in an image of 50x50 pixels, consisting of 1585 pixels of air, 201 pixels of enamel, and 714 pixels of dentine (Figure 6.1(b)). The boundary effects between the different tissue types were intentionally ignored in this image, since this is a first stage exploration of the segmentation techniques. A synthetic image built in this way leads to a perfect ground truth for the segmentation as each pixel in the image belongs to exactly one class, which is already known. In this way, each segmentation can be evaluated by counting the number of pixels which have been incorrectly clustered.

To create the teraherta data set, a terahertz time-series must be 'acquired' for each pixel — as if the model was an actual object that was being imaged. This was done by

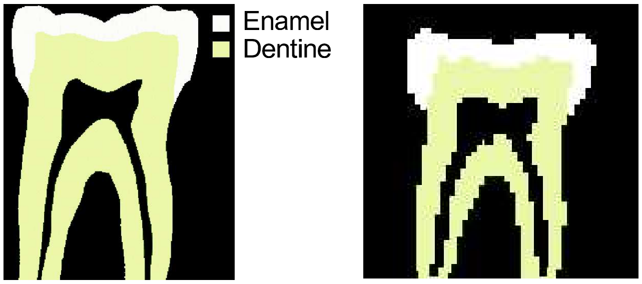

(a) (b)

Figure 6.1: (a) An example cross-section of a tooth, showing the enamel and dentine areas, and (b) the allocation of regions in the synthetic data.

using 3 exemplar terahertz pulses from a real scan of a slice of tooth — one per region — and using them to populate the terahertz data set as appropriate. Thus all the air regions within the terahertz data set consist of an indentical terahertz pulse that was actually acquired through air, and similarly for the dentine and enamel regions.

The disadvantage of a pure synthetic image is that segmentation becomes trivial — the intra-region differences are zero while the inter-region differences are large. To attempt to make a more realistic model random normally distributed noise was addedd to the time series. This study was undertaken before the noise investigation of Chapter 4 and without that knowledge, therefore a simple noise model was used.

Real Slice of Tooth

A 200 μm cross section of tooth sliced from a real tooth was imaged in transmission mode at Leeds. 64 time points were acquired at 146.7 fs intervals, with no averaging and an offset of 22.5. The time constant was set to 500 ms, with an LIA delay of 1 s. 56 scans were obtained in each of the x and y directions, with x ranging from 0 to 22.2 mm and y ranging from 0 to 9 mm.

An x-ray of the slice of tooth was taken, and was scaled and registered to a high frequency terahertz phase image, using the external edge of the tooth and the hole as the surfaces to be matched, in Analyze. Segmentation was done interactively in Analyze by a single observer to obtain the ground truth for comparison. Figure 6.2 shows the segmented radiograph of the tooth slice and a terahertz absorption image.

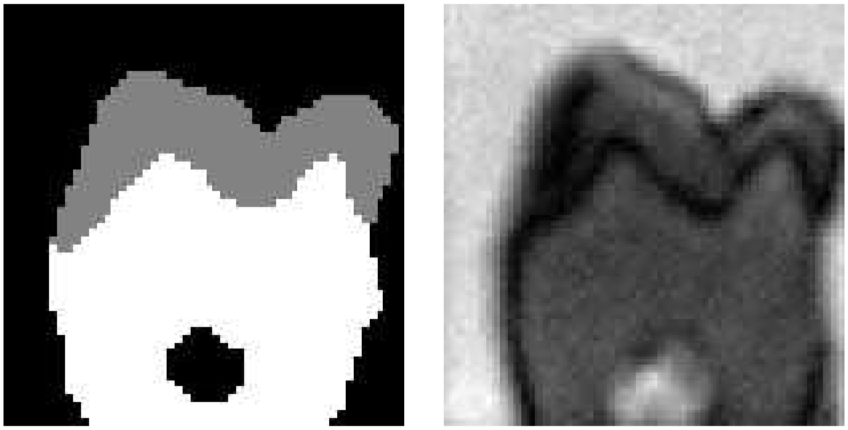
(a) (b)

Figure 6.2: (a) A hand segmented radiograph of a real slice of tooth, and (b) the relative transmission image of the real slice of tooth at 1.38 THz.

Phantoms

Finally, phantoms were created specifically for the purpose of developing and evaluating the segmentation techniques. The desireable properties are for the phantoms to have distinct, easily quantifiable regions of different, well-characterised materials. Unfortunately at the time of writing there were very few well characterised materials in the terahertz band, and those that were did not lend themselves well to the fashioning of clustering phantoms within the scope of this thesis.

It was decided to use a hollow (or outline) sticker whose regions could be filled with different materials because this immediately provides a phantom that has distinct well definied regions, that can also be easily segmented by hand from a visual image and subsequently registered to its terahertz image. Each region defined by the sticker was carefully painted with a domestic paint until it was completely filled. Various types and colours of paint were chosen in the hope that these would provide enough contrast in the terahertz band to make the segmentation non-trivial but still possible. Two further phantoms were manufactured - one based solely on a sticker, and one based solely on paint.

The phantoms were manufactured on TPX sample plates with four different domestic paints[1]. The sample plates were approximately 25 mm square with 2 mm thickness of TPX, and are usually used as the front and back of the sample holders.

The paints are identified by colour, and are detailed in Table 6.1. The stickers were peel off ‘gold’ stickers obtained by mail order from *kookykards.com* [2], with the ‘rabbit’ coming from sheet number 1850 (“Farm Animals”) and the ‘foot’ coming from sheet number 1844 (“Francoise Read: Washing line and Feet”). Table 6.2 details the construction of each of the three phantoms. Photographs of the phantoms are shown in Figure 6.3, although note that the ‘THZ’ phantom is shown before the top coat was applied.

Colour	Description	Name
Blue	Homebase matt vinyl emulsion for interior walls and ceilings	“Fisherman’s Blue”
Red	Homebase “Sanctuary” matt emulsion	“Firedance”
Gold	Crown metallic effect	“Gold Palace”
Green	Crown shimmer coat	“Green Shimmer”

Table 6.1: Description of the paints used to create the phantoms.

Name	Description
Foot	A solid sticker of a foot (taken from inside the hollow outline).
Rabbit	A hollow sticker of a rabbit, having two large regions filled with red and green paint, and two smaller regions filled with yellow and blue paint.
THZ	A single coat of blue paint, followed by the letters ‘T’, ‘H’, and ‘Z’ painted in green, followed by a covering coat of blue paint.

Table 6.2: Description of the sticker and paint phantoms.

The sticker phantoms were designed to be imaged in transmission mode, however operational issues beyond our control prevented the acquisition of these data at Leeds,

---

[1] TPX, or poly 4 methyl pentene-1, is a low loss polymer that is almost transparent at terahertz frequencies [5].

[2] *http://www.kookykards.com/peel_off_stickers_images.htm* Last visited November 27, 2003.

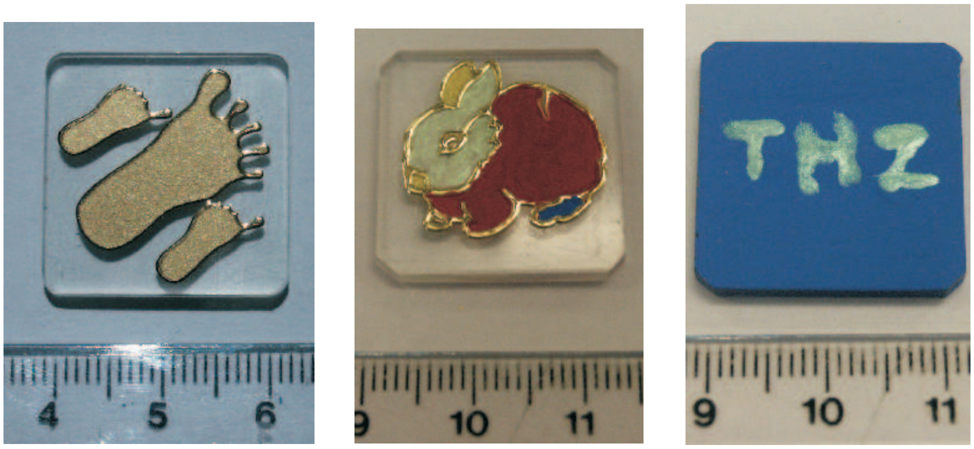

(a) (b) (c)

Figure 6.3: Photographs of the phantoms (a) ‘foot’, (b) ‘rabbit’, and (c) the partially completed ‘THZ’ phantom.

and instead they were imaged in reflection mode. They were imaged at Teraview Ltd., Cambridge, with the acquisition parameters shown in Table 6.3. The top part of the table lists the machine settings, in upper case, while the lower section lists some useful derived parameters. Note that the x direction in the Teraview machine is scanned continuously by sampling at constant time intervals as the x stage moves. In this way, approximately constant spatial x intervals between the pixels is achieved. This means that the number of pixels acquired in the x direction is dependent on the distance being scanned and the speed of the x stage, according to the formula $\text{total pixels} \approx \frac{\text{total x distance (mm)}}{\text{xspeed}} * \text{scanfreq}$. The scanning did vary however, and only 99 x pixels were obtained in the ‘rabbit’ image, whereas 101 x pixels were obtained in the ‘THZ’ image. The ‘foot’ had exactly 100 x pixels.

The pulses demonstrated a large d.c. offset, and this was removed using the fast Fourier transform before any clustering was undertaken. The phantoms were imaged face-up, with the bottom of the TPX sample plate resting on the imaging window. This arrangement results in a time-series consisting of two pulse responses, as shown in Figure 6.4. The first of these is the response from the quartz/TPX interface, and in a perfect system would be identical at every pixel. The second pulse is the response from the

TVL Version	3
DATASIZE	2048
APEAMP	800.0
PHASE	2
SCANFREQ	15.0
HIGHFILTER	16
AVERAGES	1
DATATYPE	raw
SCANCAL	-0.10043094726
LOWFILTER	500
SCANAMP	5.13325516467
XSPEED	3.0
YCOUNT	100
XRANGE	-10 – 10 mm
YRANGE	-10 – 10 mm
Time samples	512
Time step	100 fs
Approximate x count	100

Table 6.3: Teraview acquisition parameters.

phantom side of the TPX slide, consisting of some combination of TPX/air, TPX/sticker, TPX/paint, sticker/air, sticker/paint, and paint/air interfaces, depending on the phantom and pixel. Obviously there are no phantom or pixel specific data in the time-series until the terahertz radiation has had time to travel through the TPX, reflect off the interesting interfaces, and return back through the TPX. In practice, this ‘non-interesting’ section represented more than half of the time series, and so for the purposes of this study the first 50% of each pulse was windowed out before taking time or frequency domain data.

The ground truth for the two sticker based phantoms was calculated by scaling and registering the optical image of the phantom to a high frequency absorption image using the edges of the stickers as the surface to be matched. Segmentation was carried out manually by a single observer.

The manual segmentation of ‘foot’ resulted in 3 regions — TPX (5184 pixels), thick sticker (820 pixels), and thin sticker (3996 pixels), making a total of 10,000 pixels. The sticker part was split into thick and thin areas since the outline of the sticker consisted of

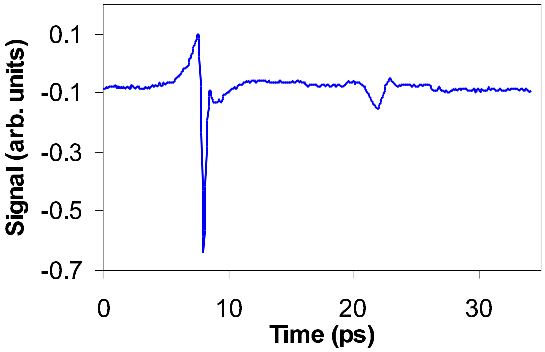

Figure 6.4: An example terahertz pulse acquired in reflection mode.

thicker plastic than the internal part. The manual segmentation of the ‘rabbit’ results in 6 regions — TPX (3083 pixels), sticker (2066 pixels), green paint (1371 pixels), gold paint (198 pixels), red paint (3058 pixels), and blue paint (124 pixels), making a total of 9900 pixels. Figure 6.5 shows the results of these manual segmentations.

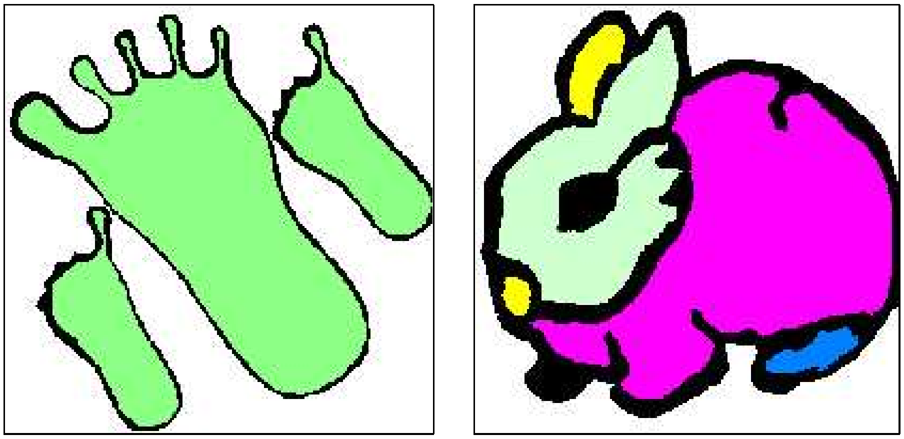

(a) (b)

Figure 6.5: The results of manual segmentation of the phantoms (a) ‘foot’, and (b) ‘rabbit’.

The ‘THZ’ phantom was highly speculative and intended as a qualitative test to see if the hidden paint regions could be distinguished by terahertz imaging. No ground truth image was therefore produced for this phantom.

6.2.2 Clustering

In this work I used standard *k*-means clustering with random initialisation and Euclidean distance metric. To compensate for the random initialisation the clustering was carried out between 20 and 30 times for each sample and feature vector, and the results are presented for the best outcome. The number of clusters to form was chosen based on *a priori* information about the constituency of each sample, as given by the number of regions in manual segmentation — three for the slices of tooth and ‘foot’, and five for ‘rabbit’.

Two approaches were used to form the feature vectors. Firstly, high dimensional vectors were formed based on the entire series from three different domains — the time domain, fast Fourier transform (FFT) domain, and the discrete wavelet transform (DWT) [40] domain. The amplitude of the FFT coefficients describe the spectral power content of the pulse, and the DWT coefficients provide a combined time and frequency distribution.

Secondly, low dimensional feature vectors were formed by combining parameters derived from the terahertz data which in themselves provide image contrast, such as phase shifts and the absorbances at various frequencies. For the analyses of the tooth data a 3-dimensional feature vector was built based on the integral phase shift between 0.5 and 1 THz and between 1 THz and 1.5 THz, and the absorbance at 1 THz. For the reflection data it was necessary to move to lower frequency bands as the frequency range of the Teraview scanner does not extend as high as that of the Leeds system. The integral phase shift between 0 THz and 0.5 THz and between 0.5 THz and 1.0 THz, and the absorbance at 0.8 THz was therefore used instead. The 3-dimensional vectors were then normalised to be uni-variate within a unit hypercube.

6.2.3 Evaluation

The segmentations were evaluated by counting the number of mis-classified pixels compared with the respective registered ground truths that were previously described. In each case the results of the best segmentation are presented.

The reflection data was marred by the occasional missed pixel, an artefact of continuous scanning. If the controlling computer became too busy the data buffer overran and any pixels scanned in during the busy period were lost. Those pixels were excluded from the clustering.

6.3 Results

6.3.1 Tooth Slices

A successful segmentation is one that results in three contiguous regions corresponding to the three regions present in the tooth. It was found that the success of the clustering was dependent on initialisation, with successful segmentations resulting around 70% of the time. It was also found that when the segmentations were successful the number of misclassified pixels varied by less than 1% of the total number of pixels.

	Time series	FFT coefficients	DWT coefficients	Feature vector
Synthetic Tooth	0%	0.04%	0%	1%
Slice of Real Tooth	19%	23%	19%	13%

Table 6.4: Number of mis-classified pixels when segmenting a synthetic and a real tooth using *k*-means clustering on a variety of vectors.

The standard *k*-means clustering segmented perfectly when using the time-series and DWT coefficients, with errors present when clustering on Fourier coefficients or the 3D feature vector (see Table 6.4). Figure 6.6(a-c) shows the best results of the *k*-means clustering using the various feature vectors. Figure 6.6(d) shows the same parameters and algorithm as (a) where the segmentation failed due to unfavourable initialisation. In these images, colour is used to show into which cluster the pixels were allocated, out of the three possibilities. The colours were manually assigned to the clusters to ease comparison.

The same segmentation techniques were applied to the real slice of tooth, and the

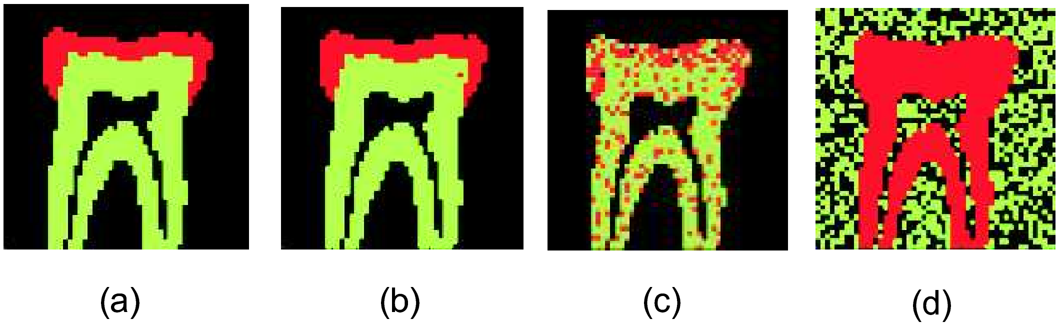

Figure 6.6: The clustering of the synthetic tooth images using *k*-means clustering on (a) Time series or DWT coefficients, (b) FFT coefficients, (c) 3D feature vector. (d) shows a failed clustering due to poor initialisation, in this instance on the time series data.

results shown in Figure 6.7. Figure 6.7(a) shows the relative amplitude of the pulse, which is an indicator of where the boundaries between the different areas lie. Figures 6.7(b)–(d) show the segmentation results as before, with the hand-segmented boundaries from the radiograph shown in white.

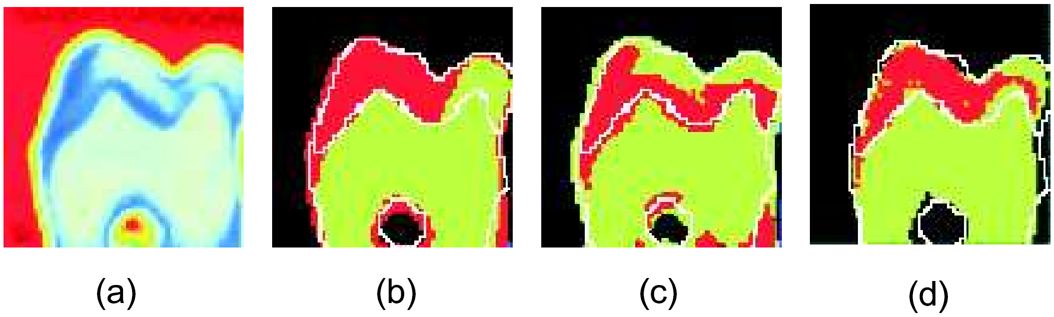

Figure 6.7: (a) The relative amplitude of the pulse through the real tooth, and the *k*-means clustering images on (b) Time series, (c) FFT coefficients, (d) 3D feature vector. The white lines show the boundaries extracted from the hand-segmented radiograph.

6.3.2 Phantoms

The phantoms imaged in reflection mode demonstrated less of a dependence on initialisation than the tooth, in that over 95% of the clusters resulted in groupings that approximated the ground truth. Successful runs again demonstrated less than 1% variation in the number of misclassified pixels.

It was also found that the frequency domain was subject to a lot of noise, below 0.234 THz and above 0.820 THz. To exclude this noise, I used a shortened frequency vector based on the coefficients only between these frequencies. However, it also found that the lower frequencies tended to demonstrate noise in the sense that while a given frequency might produce an acceptable image, the frequencies either side were subject to a lot of noise. The lower frequencies also had a blurring effect due to the longer wavelength. To compensate for these errors, the lower frequency limit was raised until no further improvement in the misclassification rate was obtained. This resulted in a frequency vector of the amplitude of the eight coefficients between 0.615 THz and 0.820 THz. This was the same for 'rabbit' and 'foot'.

Phantom	Time domain	Frequency domain	DWT domain	Feature vector
Foot	28.3%	29.6%	28.3%	39.8%
Rabbit	50.3%	56.9%	50.2%	64.4%

Table 6.5: Percentage of misclassified pixels for the phantoms imaged in reflection mode, using *k*-means clustering on three different vectors. In total there were 9,933 'foot' pixels and 9744 'rabbit' pixels classified.

The results for the 'foot' and 'rabbit' phantoms are shown in Table 6.5. As can be seen the percentage errors are far higher for these images than for the tooth, and this is due to the large influence of the boundaries, and the difficulty in resolving the fine structure of the stickers at terahertz frequencies. For instance, the smaller toes are less than 1 mm wide with around 1 mm separation, and yet radiation at 1 THz has a wavelength of 3 mm, leading inevitably to blurring and diffraction issues. The beam is focused to 0.5 mm at 1 THz which improves the sub-wavelength spatial resolution, however in the case of the toes the resolution is still not sufficient. Similarly fine structures of the rabbit, like its eye, nose, and feet, are too small to be resolved by the lower frequencies. Figures 6.8 and 6.9 show examples of the best attempt at segmenting the phantom images using the different vectors. Note that in these images, the excluded pixels are shown in red. The time domain image of the ‘foot’ particularly demonstrates the boundary effects — the fine detail of the toe region of the large foot has been entirely lost into one large merged area, as has the separation between the stickers. By careful selection of higher frequencies, it is possible to recover some of the toe detail, but the errors remain very large.

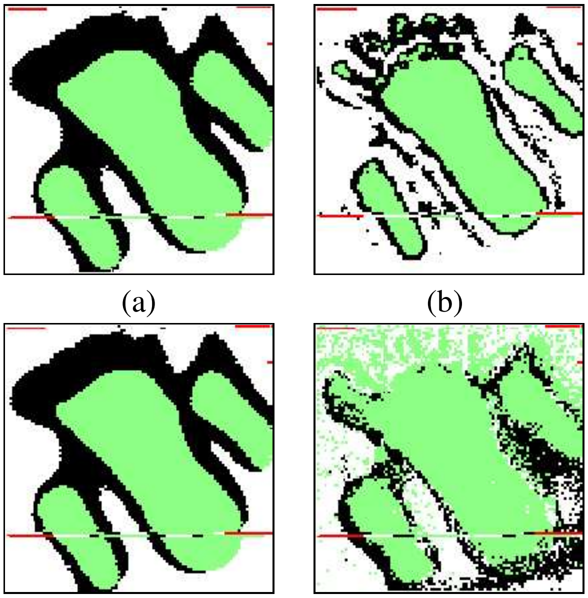
(a) (b)
(c) (d)

Figure 6.8: The results of segmenting the ‘foot’ phantom using (a) the time domain, (b) the frequency domain, (c) the DWT domain, and (d) the feature vector.

The challenge of the ‘THZ’ phantom proved too hard for this approach — the differences between the paints were too small and the lack of uniformity across the phantom too large for the regions to be distinguishable. An example output of the segmentation process on this phantom is shown in Figure 6.10, but that picture is typical of the output from any of the analysis techniques tried.

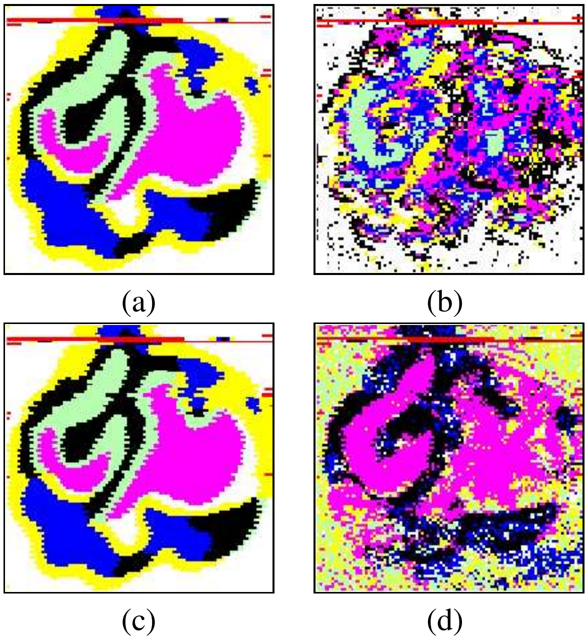

Figure 6.9: The results of segmenting the ‘rabbit’ phantom using (a) the time domain, (b) the frequency domain, (c) the DWT domain, and (d) the feature vector.

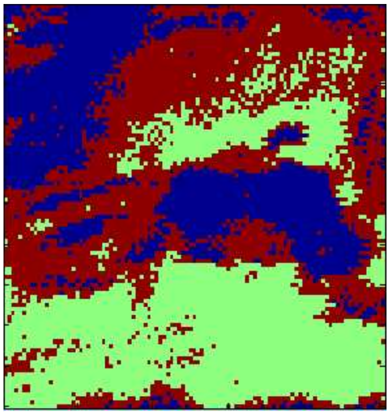

Figure 6.10: The results of segmenting the THZ phantom image using the frequency domain.

6.4 Discussion

Although speculative, this initial work is promising for transmission modality — k-means clustering segments the synthetic image without error when using a large feature vector.

The segmentation of the real tooth is also promising, although clearly nowhere near as clean as the synthetic image. This is to be expected, as extra factors such as boundary effects are now present. A physical examination of the slice of tooth showed that the top right hand corner of the tooth had been cut more thinly than the rest of the tooth, leading directly to the difficulties with segmentation in that region, so further quantification of the segmentation accuracy was not performed.

The poor performance of the phantoms imaged in reflection mode was partially expected. The phantoms were not designed for this modality — reflection mode imaging is best suited to thick, layered, samples where the interface between each of the layers produces a response. The boundary effects were more pronounced for the reflection images, leading to higher error rates than for the tooth. The buffer overruns for certain pixels led to entire rows of the image being misaligned, in each case contributing up to 100 pixels to the error. A further contribution to the error may have been mis-registration with the ground truth, owing to the difficulty in precisely locating edges. However, the clustering results were self-consistent, and in the case of the 'foot' the segmentation could easily distinguish between the broad areas of sticker, edge of sticker, and TPX. It was nevertheless disappointing that there were 50% or greater errors when classifying 'rabbit' pixels. This again was due to a combination of boundary effects, the inability of terahertz radiation to distinguish between the different paints, and inevitable physical variations in the application of paint, and the spatial resolution issues.

Although the 'THZ' phantom was not expected to be perfectly segmented, least of all in reflection mode, there was not even any evidence of the different paints used. This is due to a combination of factors, such as the two paints not being sufficiently distinguishable under this frequency of radiation, and varying thickness of paint tending to mask the effects of different paints. It would have been interesting to have imaged all the phantoms in transmission mode, but limited machine access did not allow this.

6.5 Conclusions

The study has demonstrated that a feature based clustering technique has potential value for the segmentation of transmission images, with the number of groups known *a priori* and a random initialisation strategy. Although not the focus of this exploration, existing initialisation strategies and methods of choosing the number of groups are expected to remove all the operator intervention, leading to truly unsupervised clustering.

The phantoms were deliberately designed to have a combination of macro- and micro-structures, in order to assess how well a clustering algorithm could cope with these challenges. The clustering performed well at the macro structure level, easily distinguishing between TPX and non-TPX regions in most cases, but the physical limitations of the phantoms, the technology, and use of reflection imaging meant that the micro structures could not reliably be extracted.

Chapter 7

Conclusions and future work

The work in this thesis was motivated by a desire to deepen our understanding of terahertz pulsed imaging and provide a range of analysis tools suitable for this domain.

7.1 Summary of Work

In Chapter 4 a model of the noise present in terahertz images was empirically built for the first time by fitting distributions from the stable family. The resulting models were statistically shown to accurately model the variation between terahertz pulses acquired through homogeneous material. Existing denoising techniques were applied in a new way with extreme denoising in order to quantify the errors introduced and their suitability for data compression applications.

Chapter 5 presented the novel application of time frequency analysis to terahertz data in order to estimate the optical properties of the sample. Both the short time Fourier transform and the wide band cross ambiguity functions were shown to be able to distinguish between nylon and resin, and the short time Fourier transform was additionally used to calculate the complex refractive index of a number of biomedical samples. A windowing technique was also used for the first time to demonstrate the impact that very short acquisition times would have on complex refractive index calculations.

The focus of Chapter 6 was on segmentation through the original application of an existing clustering technique to terahertz data by creating feature spaces drawn from the time domain, frequency domain, joint time frequency domains, and the domain of derived physical characteristics. The techniques were evaluated on both transmission and reflection mode images.

7.2 Discussion

This thesis set out to provide a number of novel time frequency analyses and analysis techniques suitable for terahertz pulsed imaging. In Section 1.2 some of the issues facing terahertz imaging were outlined, and these and other topics of interest are now discussed.

The *noise* present in terahertz data has been explored in two ways. Firstly it was established that the observed inter-pulse variation is not distributed according to a Gaussian distribution, or a distribution formed from a mixture of Gaussians, except under highly idealised conditions. This is an important finding, as the standard Gaussian model of noise is no longer appropriate. Instead, a distribution from the stable family (which includes Gaussians) must be employed, and the selection of parameters is non-trivial. Although this is a small step forward in our understanding of terahertz noise, it does pave the way for better models, better denoising techniques, and algorithms that are more robust against noise.

The extreme denoising has relevance to the *volume of data* typically acquired. It was demonstrated that only 20% of coefficients in the wavelet domain are needed for complex refractive index calculation. Of course, a truncation scheme such as this does not result in an 80% reduction in storage, since the positions of the untruncated coefficients still need to be stored. However combining a truncation scheme with a traditional lossless compression technique would help address this issue.

The issue of *long acquisition times* was partially dealt with through the examination of the effect that windowing the pulses has. This windowing is equivalent to the acquisition of fewer samples per pulse, which in turn directly corresponds to a shorter acquisition time. As expected, the shorter acquisition times did impact the calculation of complex refractive index, although the general trend was preserved. This means the acquisition time to be used depends on the application — if an image is being formed from relative transmission at one or two frequencies then a very short acquisition may be appropriate. On the other hand, if a detailed analysis of the complex refractive index at terahertz frequencies is required then longer acquisition times become necessary.

The rest of this work was focused on strategies to manage the *high dimensionality of data* in terahertz imaging. The short time Fourier transform looks to be a promising tool for calculating optical constants, especially for refractive index values of noisy data. In terahertz imaging one would not normally expect to have homogeneous samples of several different known thickness, and so in some regards the calculation of optical constants is an academic exercise. However, characterising these materials is useful for modelling their responses to terahertz radiation. The context for the extraction of optical constants was exploring contrast mechanisms, and if a technique such as the STFT can successfully be used to calculate optical constants, then it should also be able to highlight where materials of different optical properties are contained within an image. By the same token, the WBCAF shows potential as a tool to this end, although the inconsistency with traditional analyses does raise questions about its reliability.

Clustering mechanisms ultimately rely on the optical and physical properties of the sample being imaged of course, as it is these which cause differences to appear between the pulses. However, removing the need for *a priori* information and allowing the computer to group together pixels it considers to be similar greatly simplifies the problem. These techniques show great potential for transmission images, even if questions remain over the reflection images due to the sub-optimal design of phantom for this modality.

7.3 Future Work

The noise models created using distributions from the stable family form an extremely good fit to the observed noise. This begs the question of whether there is an underlying physical process that the stable family is modelling. The suggestion of a relationship between some of the stable parameters and thickness of material is an interesting avenue for further research, possibly in terahertz imaging methods such as dark field imaging [44]. These explorations would hopefully result in a better understanding of the noise processes present, so that better noise models could be built, enabling realistic noise to be added to synthetic images and more robust algorithms to be developed.

The denoising results suggest great potential for data compression. The basic truncation scheme used here could potentially be improved through the application of coefficient quantisation and perhaps some form of linear predictive encoding.

The limitations of an analysis based on the phase of the Fourier transform made assessing the accuracy of the short time Fourier transform difficult. Advances in either phase unwrapping techniques or the application of methods such as Monte Carlo modelling should provide the ground truth refractive index profile against which the results presented here could be evaluated. The wide band cross ambiguity analysis results were not consistent with the results from traditional analyses, and the reasons for this need further exploration. The obvious question is whether or not this matters if the technique allows different materials to be distinguished, and this would need further experimentation.

The time frequency techniques were explored with reflection geometry in mind, the mode in which the real advantage over traditional analyses would be found. The analyses presented would need refining before they are directly applicable to reflection mode images, however.

Clearly a feature based clustering technique has potential value for transmission images, as the real slice segmentations show. More work needs to be carried out on the parameters extracted to form the feature vectors. The absorbance should also be calculated for a frequency range to provide better robustness to noise. In this study the number of groups was known *a priori*, and the issue of initialisation was bypassed by using repeated random initialisations. In a complete system, the number of groups to be formed may not be known, and a better strategy would need to be adopted for initialisation. It is expected that principal components analysis to be an appropriate tool to try within the reduced feature spaces, using the Mahalanobis distance metric instead of Euclidean [10].

The development of a segmentation tool in the reflection mode would greatly benefit from the creation of phantoms better suited to this modality. The phantoms were originally designed to be imaged in transmission mode, but operational constraints with the terahertz scanner at Leeds led to them being imaged in reflection mode at Cambridge instead. Similarly, the features on the phantoms were at the limit of the spatial resolution of terahertz data, causing blurring. It is not known whether the boundary effects are more pronounced in reflection imaging, and this is a further area for future research.

The errors in the segmentation could be decreased through the application of a spatial reinforcement strategy, such as relaxation [32], which applies the principle that neighbouring pixels are likely to be of the same material. The application of an expectation maximisation technique with a hidden ‘thickness’ variable may also decrease the error of the segmentation of the slice of tooth.

Finally, k-means is just one of a large family of clustering tools, and with more data there would be increased scope for exploring the use of support vector machines, artificial neural networks, and other supervised and unsupervised techniques.

Bibliography

[1] S.J. Allen, K. Craig, B. Galkrikian, J.N. Heyman, J.P. Kaminski, J.S. Scott M.S. Sherwin, K. Unterrainer, M. Wanke, K. Campman, P.F Hopkins, A.C. Gossard, D.H. Chow, M. Lui, and T.K. Liu. Materials science in the far-IR with electrostatic based FELs. *Nuclear Instruments & Methods in Physics Research Section A: Accelerators Spectrometers Detectors and Associated Equipment*, 358:536–539, 1995.

[2] D.H. Austin and M.C. Nuss. Electro-optic generation and detection of femtosecond electrical transients. *IEEE Journal of Quantum Electronics*, 24(2):184–197, 1988.

[3] A.H. Barnett, J.P. Culver, A.G. Soresen, A. Dale, and D.A. Boas. Robust inference of baseline optical properties of the human head with three-dimensional segmentation from magnetic resonance imaging. *Applied Optics*, 42(16):3095–3108, 2003.

[4] E. Berry, A.J. Fitzgerald, N.N. Zinov'ev, G.C. Walker, S. Homer-Vanniasinkam, C.D. Sudworth, R.E. Miles, J.M. Chamberlain, and M.A. Smith. Optical properties of tissue measured using terahertz pulsed imaging. In *Medical Imaging*, volume 5030 of *Proceedings of SPIE*, pages 459–470. SPIE, 2003.

[5] J.R. Birch, J.D. Dromey, and J. Lesurf. The optical constants of some common low loss polymers between 4 and 40 cm^{-1}. Technical Report DES 69, NPL, February 1981.

[6] D. Boukerroui, A. Baskurt, J.A. Noble, and O. Basset. Segmentation of ultrasound images – multiresolution 2D and 3D algorithm based on global and local statistics. *Pattern Recognition Letters*, 24(4–5):779–790, 2003.

[7] J. Chambers, W. Cleveland, B. Kleiner, and P. Tukey. *Graphical Methods for Data Analysis*. Wadsworth, 1983.

[8] Q. Chen, Z.P. Jiang, and X.-C. Zhang. All-optical THz imaging. *Terahertz Spectroscopy and Applications*, 3617:98–105, 1999.

[9] L. Cohen. Time-frequency distributions — a review. *Proc. of the IEEE*, 77(7):941–981, 1989.

[10] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Training models of shape from sets of examples. In *Proc. British Machine Vision Conference*, pages 9–18. Springer–Verlag, 1992.

[11] D. Crawley, C. Longbottom, V.P. Wallace, B. Cole, D. Arnone, and M. Pepper. Three-dimensional terahertz pulse imaging of dental tissue. *Journal of Biomedical Optics*, 8(2):303–307, 2003.

[12] N Cristianini and J Shawe-Taylor. *Support Vector Machines and other kernel-based learning methods*. Cambridge University Press, 2000.

[13] R.E. Crochiere and L.R. Rabiner. Interpolation and decimation of digital signals - a tutorial review. *Proceedings of the IEEE*, 69(3):300–329, March 1981.

[14] J.R. Demers, T.M. Goyette, K.B. Ferrio, H.O. Everitt, B.D. Guenther, and F.C. De Lucia. Spectral purity and sources of noise in femtosecond-demodulation terahertz sources driven by ti:sapphire mode-locked lasers. *IEEE Journal of Selected Topics in Quantum Electronics*, 37(4):595–605, 2001.

[15] A. Demir and A. Sangiovanni-Vincentelli. *Analysis and Simulation of Noise in Nonlinear Electronic Circuits and Systems*. Kluwer Academic Publishers, MA, USA, 1998.

[16] A. Dempster, D. Rubin, and N. Laird. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B*, 39:1–38, 1977.

[17] D.L. Donoho. De-noising by soft thresholding. *IEEE Trans. on Information Theory*, 41(3):613–627, 1995.

[18] T. Dorney, R. Baraniuk, and D. Mittleman. Material parameter estimation with terahertz time-domain spectroscopy. *Journal of the Optical Society of America A*, 18(7):1562–1571, 2001.

[19] T. Dorney, W. Symes, R. Baraniuk, and D. Mittleman. Terahertz multistatic reflection imaging. *Journal of the Optical Society of America A*, 19:1432–1443, 2002.

[20] T.D. Dorney, J. Johnson, D.M. Mittleman, and R.G. Baraniuk. Imaging with terahertz pulses. In *Proc. SPIE. Application of Digital Image Processing XXIII (SPIE's 45th Annual Meeting)*, volume 4115, pages 688–99, 2000.

[21] L. Duvillaret, F. Garet, and J-L. Coutaz. A reliable method for extraction of material parameters in terahertz time-domain spectroscopy. *IEEE Journal of Selected Topics in Quantum Electronics*, 2(3):739–746, 1996.

[22] L. Fausett. *Fundamentals of Neural Networks*. Prentice Hall International, 1994.

[23] W. Feller. *Introduction to Probability and Its Applications, Vol 2*. John Wiley, 1966.

[24] B. Ferguson and D. Abbott. De-noising techniques for terahertz responses of biological samples. *Microelectronics Journal*, 32(12):959–967, 2001.

[25] B. Ferguson and D. Abbott. Wavelet de-noising of optical terahertz pulse imaging data. *Fluctuation and Noise Letters*, 1(17):L65–L70, 2001.

[26] B. Ferguson, S. Wang, D. Gray, D. Abbot, and X.-C. Zhang. Identification of biological tissue using chirped probe THz imaging. *Microelectronics Journal*, 33:1043–1051, 2002.

[27] A.J. Fitzgerald, E. Berry, N.N. Zinov'ev, S. Homer-Vanniaskinkam, R.E. Miles, J.M. Chamberlain, and M.A. Smith. Catalogue of human tissue optical properties at terahertz frequencies. *Journal of Biological Physics*, 29(2–3):123–128, 2003.

[28] G.P. Gallerano, A. Doria, E. Giovenale, and A. Renieri. Compact free electron lasers: From Cerenkov to waveguide free electron lasers. *Infrared Physics & Technology*, 40(3):161–174, 1999.

[29] J.C. Goswami and A.K. Chan. *Fundamentals of Wavelets: Theory, Algorithms, and Applications*. Wiley Series in Microwave and Optical Engineering. John Wiley & Sons, 1999.

[30] D. Grischkowsky, S. Keiding, M. van Exter, and C. Fattinger. Far-infrared time-domain spectroscopy with terahertz beams of dielectrics and semiconductors. *Journal of the Optical Society of America B: Optical Physics*, 7(10):2006–2015, 1990.

[31] S. Hadjiloucas, R. Galvão, and J. Bowen. Analysis of spectroscopic measurements of leaf water content at terahertz frequencies using linear transforms. *Journal of the Optical Society of America A*, 19:2495–2509, 2002.

[32] E.R. Hancock and J. Kittler. Discrete relaxation. *Pattern Recognition*, 23(7):711–733, 1990.

[33] J.A. Hartigan. *Clustering Algorithms*. John Wiley & Sons, New York, 1975.

[34] M. Herrmann, M. Tani, and K. Sakai. Display modes in time-resolved terahertz imaging. *Japanese Journal of Applied Physics*, 39(11):6254–6258, November 2000.

[35] M. Herrmann, M. Tani, M. Watanabe, and K. Sakai. Terahertz imaging of objects in powders. *IEE Proc. Optoelectron.*, 149(3):116–120, June 2002.

[36] B.B. Hu and M.C. Nuss. Imaging with terahertz waves. *Optics Letters*, 20:1716–1718, 1995.

[37] R.H. Jacobsen, D.M. Mittleman, and M.C. Nuss. Chemical recognition of gases and gas mixtures with terahertz waves. *Optics Letters*, 21(24):2011–2013, 1996.

[38] D.A. Jaroszynski, B. Ersfield, G. Giraud, S. Jamison, D.R. Jones, R.C. Issac, B.M.W. McNeil, A.D.R. Phelps, G.R.M. Robb, H. Sandison, G. Vieux, S.M. Wiggins, and K. Wynne. The Stratchclyde terahertz to optical pulse source (tops). *Nuclear Instruments & Methods in Physics Research Section A: Accelerators Spectrometers Detectors and Associated Equipment*, 445:317–319, 2000.

[39] N.L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 1 of *Wiley Series in Probability and Mathematical Statistics*. Wiley & Sons, 2 edition, 1994.

[40] G. Kaiser. *A Friendly Guide to Wavelets*. Birkhaüser, 1994.

[41] J.T. Kindt and C.A. Schmuttenmaer. Far-infrared dielectric properties of polar liquids probed by femtosecond THz pulse spectroscopy. *Journal of Physical Chemistry*, 100:10373–10379, 1996.

[42] T. Kleine-Ostmann, P. Knobloch, M. Koch, S. Hoffmann, M. Breede, M. Hofmann, G Hein, K. Pierz, M. Sperling, and K. Donhuijsen. Continuous-wave THz imaging. *Electronics Letters*, 37(24):1461–1463, 2001.

[43] K. Kobayashi. Self-organizing wavelet-based neural networks. In M. Akay, editor, *Time Frequency and Wavelets in Biomedical Signal Processing*, chapter 28, pages 685–702. IEEE Press, 1998.

[44] T. Löffler, T. Bauer, K.J. Siebert, H.G. Roskos, A.J. Fitzgerald, and S. Czasch. Terahertz dark-field imaging of biomedical tissue. *Opt. Express*, 9:616—621, 2001.

[45] T. Löffler, K.J. Siebert, T. Bauer, H.G. Roskos, A.J. Fitzgerald, and S. Czasch. Medical and non-medical case studies on THz image generation based on coherent time-domain measurements. In *First Int. Conf. on Biomedical Imaging and Sensing Applications of Terahertz Technology*, 2001.

[46] T. Löffler, K.J. Siebert, S. Czasch, T. Bauer, and H.G. Roskos. Visualization and classification in biomedical terahertz pulsed imaging. *Physics in Medicine and Biology*, 47(21):3847–4852, November 2002.

[47] A.G. Markelz, A. Roitberg, and E.J. Heilwiel. Pulsed terahertz spectroscopy of DNA, bovine serum albumin and collagen between 0.1 and 2.0 THz. *Chemical Physics Letters*, 320:42–48, 2000.

[48] S. Mickan, D. Abbott, J. Munch, X.-C. Zhang, and T. van Doorn. Analysis of system trade-offs for terahertz imaging. *Microelectronics Journal*, 31(7):503–514, July 2000.

[49] D.M. Mittleman, R.H. Jacobsen, and M.C. Nuss. T-ray imaging. *IEEE Journal of Selected Topics in Quantum Electronics*, 2(3):679–692, 1996.

[50] D.M. Mittleman, R. Neelamani, R.G. Baraniuk, and M.C. Nuss. Applications of terahertz imaging. In *Nonlinear Optics '98: Materials, Fundamentals and Applications Topical Meeting*, pages 294–296, 1998.

[51] J. Nolan. *STABLE Program for Windows*, 3.10 edition, 2002. http://academic2.american.edu/˜jpnolan/stable/stable.html - Last visited 25 November 2003.

[52] T. Nozokido, J. Bae, and M. Mizuno. Scanning near field millimeter-wave microscopy using a metal slit as a scanning probe. *IEEE Trans. Microwave Theory Tech.*, 49:491–499, 2001.

[53] T.G. Phillips and J. Keene. Submillimeter astronomy. *Proc. of the IEEE*, 90:1662–1678, 1992.

[54] A. Poppe, L. Xu, F. Krausz, and C. Spielmann. Noise characterisation of sub-10-fs Ti:Sapphire oscillators. *IEEE Journal of Selected Topics in Quantum Electronics*, 4:179–184, 1998.

[55] The Royal Society. *Proceedings of the Terahertz Gap: the Generation of far-infrared radiation and its applications*, June 2003. to appear.

[56] H. Rubens and E.F. Nichols. Heat rays of great wave length. *Physical Review*, 4(4):314–323, 1897.

[57] K.J. Siebert, H. Quast, R. Leonardt, T Löffler, M. Thomson, T. Bauer, and H.G. Roskos. Continuous-wave all-optoelectronic terahertz imaging. *Applied Physics Letters*, 80(16):3003–3005, April 2002.

[58] B. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 1986.

[59] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis and Machine Vision*. Brooks/Cole Publishing Company, Pacific Grove, second edition, 1999.

[60] M. van Exter, Ch. Fattinger, and D. Grischowsky. Terahertz time-domain spectroscopy of water vapor. *Optics Letters*, 14(20):1128–1130, 1989.

[61] M. van Exter and D. Grishchkowsky. Characterization of an optoelectronic terahertz beam system. *IEEE Trans. on Microwave Theory and Techniques*, 38:1684–1691, 1990.

[62] S. Wang, B. Ferguson, D. Abbot, and X.-C. Zhang. T-ray imaging and tomography. *Journal of Biological Physics*, 29(2–3):247–256, 2003.

[63] S. Webb, editor. *The Physics of Medical Imaging*. Bristol : Hilger, 1988.

[64] L.G. Weiss. Wavelets and wideband correlation processing. *IEEE Signal Processing Magazine*, 11(1):13–32, January 1994.

[65] R.M. Woodward, V.P. Wallace, D.D. Arnone, E.H. Linfield, and M. Pepper. Terahertz pulsed imaging of skin cancer in the time and frequency domain. *Journal of Biological Physics*, 29(2–3):257–261, 2003.

[66] R.M. Woodward, V.P. Wallace, R.J. Pye, B.E. Cole, D.D. Arnone, E.H. Linfield, and M. Pepper. Terahertz pulse imaging of ex vivo basal cell carcinoma. *Journal of Investigative Dermatology*, 120(1):73–78, 2003.

[67] K.C. Yao, M. Mignotte, C. Collet, P. Galerne, and G. Burel. Unsupervised segmentation using a self-organizing map and a noise model estimation in sonar imagery. *Pattern Recognition*, 33(9):1575–1584, 2000.

[68] R.K. Young. *Wavelet Theory and its Applications*. Kluwer, 1993.

[69] X.-C. Zhang. Three-dimensional terahertz wave imaging. In *Proceedings of the Terahertz Gap: the Generation of far-infrared radiation and its applications* [55]. to appear.

[70] G. Zhao, R.N. Schouten, N. van de Valk, W. Th. Wenckebach, and P.C.M. Planken. Design and performance of a THz emission and detection setup based on a semi-insulating GaAs emitter. *Review of Scientific Instruments*, 73(4):1715–1719, 2002.

Appendix A

Probability Plot Correlation Coefficient Critical Values

A.1 Critical Values

N	0.05	0.01	N	0.05	0.01	N	0.05	0.01
3	0.8687	0.879	36	0.954	0.9686	190	0.9897	0.9927
4	0.8234	0.8666	37	0.9551	0.9693	200	0.9903	0.993
5	0.824	0.8786	38	0.9555	0.97	210	0.9907	0.9933
6	0.8351	0.888	39	0.9568	0.9704	220	0.991	0.9936
7	0.8474	0.897	40	0.9576	0.9712	230	0.9914	0.9939
8	0.859	0.9043	41	0.9589	0.9719	240	0.9917	0.9941
9	0.8689	0.9115	42	0.9593	0.9723	250	0.9921	0.9943
10	0.8765	0.9173	43	0.9609	0.973	260	0.9924	0.9945
11	0.8838	0.9223	44	0.9611	0.9734	270	0.9926	0.9947
12	0.8918	0.9267	45	0.962	0.9739	280	0.9929	0.9949
13	0.8974	0.931	46	0.9629	0.9744	290	0.9931	0.9951
14	0.9029	0.9343	47	0.9637	0.9748	300	0.9933	0.9952
15	0.908	0.9376	48	0.964	0.9753	310	0.9936	0.9954
16	0.9121	0.9405	49	0.9643	0.9758	320	0.9937	0.9955
17	0.916	0.9433	50	0.9654	0.9761	330	0.9939	0.9956
18	0.9196	0.9452	55	0.9683	0.9781	340	0.9941	0.9957
19	0.923	0.9479	60	0.9706	0.9797	350	0.9942	0.9958
20	0.9256	0.9498	65	0.9723	0.9809	360	0.9944	0.9959
21	0.9285	0.9515	70	0.9742	0.9822	370	0.9945	0.996
22	0.9308	0.9535	75	0.9758	0.9831	380	0.9947	0.9961
23	0.9334	0.9548	80	0.9771	0.9841	390	0.9948	0.9962
24	0.9356	0.9564	85	0.9784	0.985	400	0.9949	0.9963
25	0.937	0.9575	90	0.9797	0.9857	450	0.9954	0.9967
26	0.9393	0.959	95	0.9804	0.9864	500	0.9959	0.997
27	0.9413	0.96	100	0.9814	0.9869	550	0.9963	0.9973
28	0.9428	0.9615	110	0.983	0.9881	600	0.9965	0.9975
29	0.9441	0.9622	120	0.9841	0.9889	650	0.9968	0.9977
30	0.9462	0.9634	130	0.9854	0.9897	700	0.997	0.9978
31	0.9476	0.9644	140	0.9865	0.9904	750	0.9972	0.998
32	0.949	0.9652	150	0.9871	0.9909	800	0.9974	0.9981
33	0.9505	0.9661	160	0.9879	0.9915	850	0.9975	0.9982
34	0.9521	0.9671	170	0.9887	0.9919	900	0.9977	0.9983
35	0.953	0.9678	180	0.9891	0.9923	1000	0.9979	0.9984

Table A.1: Critical Values for probability plot correlation coefficients.

Appendix B

Full results of STFT Analysis

B.1 Refractive Index Profiles

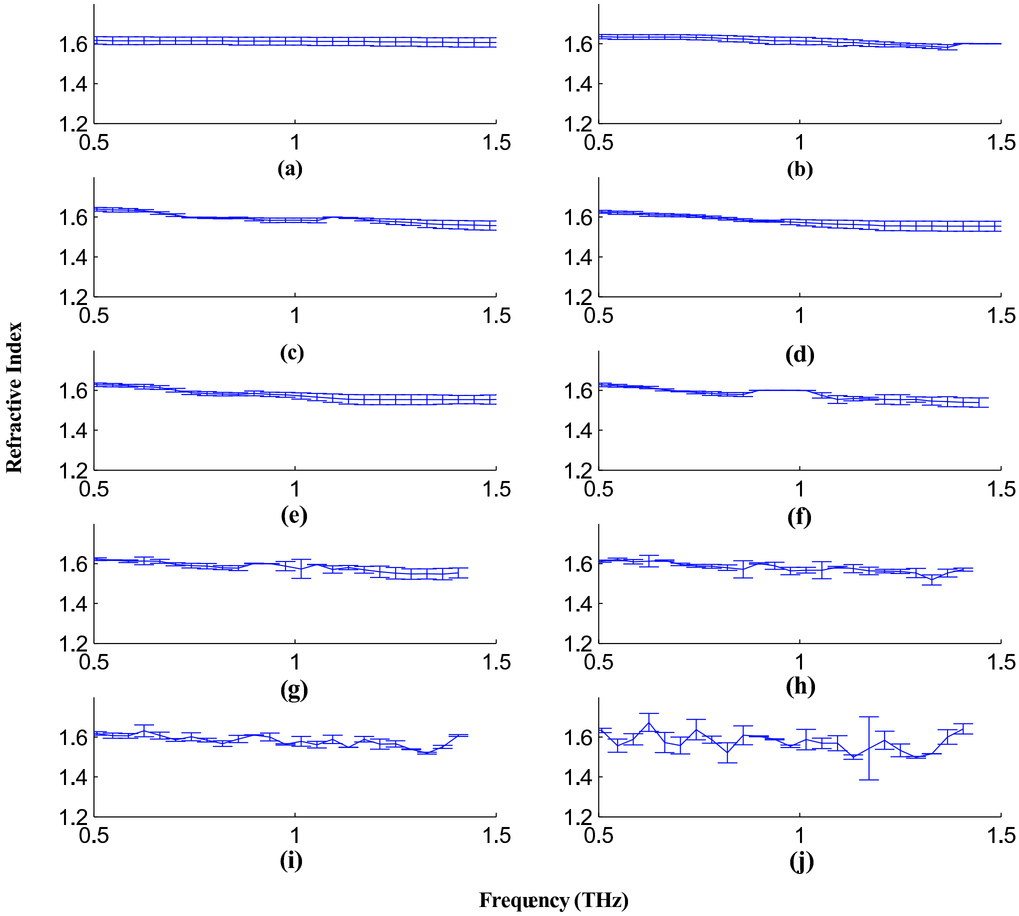

Figure B.1: Refractive index profiles for a resin step wedge using the Gaussian windowed STFT with widths (a) 10^{-1} (0.02 ps), (b) $10^{-0.5}$ (0.063 ps), (c) 10^{0} (0.2 ps), (d) $10^{0.5}$ (0.63 ps), (e) 10^{1} (2 ps), (f) $10^{1.5}$ (6.3 ps), (g) 10^{2} (20 ps), (h) $10^{2.5}$ (63 ps), (i) 10^{3} (200 ps), and (j) $10^{3.5}$ (632 ps).

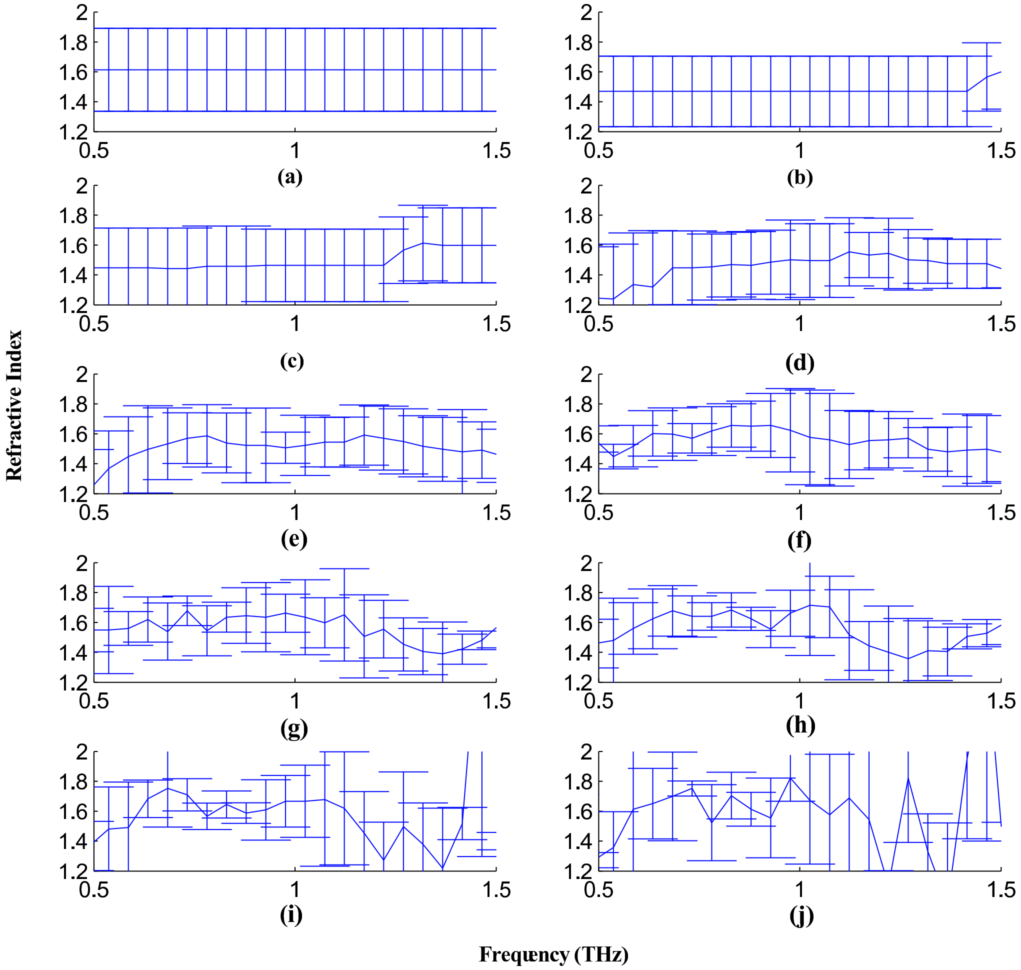

Figure B.2: Refractive index profiles for excised skin using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

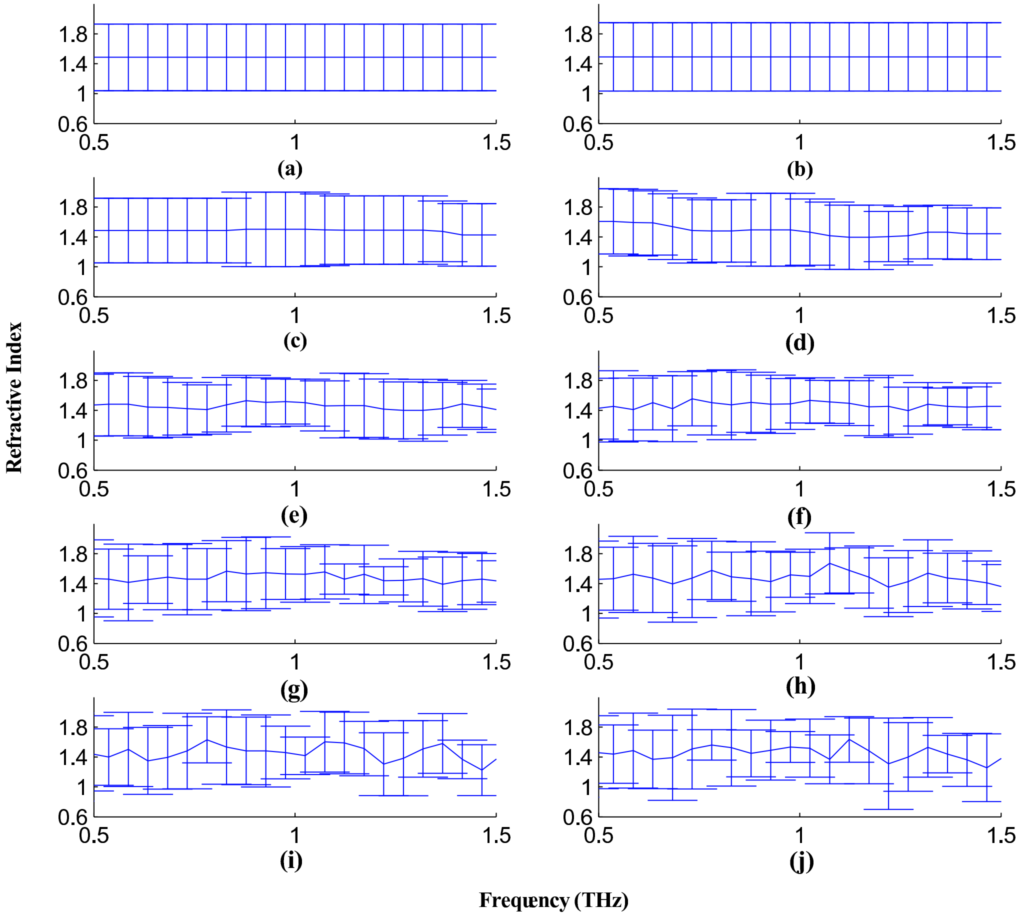

Figure B.3: Refractive index profiles for excised fat using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

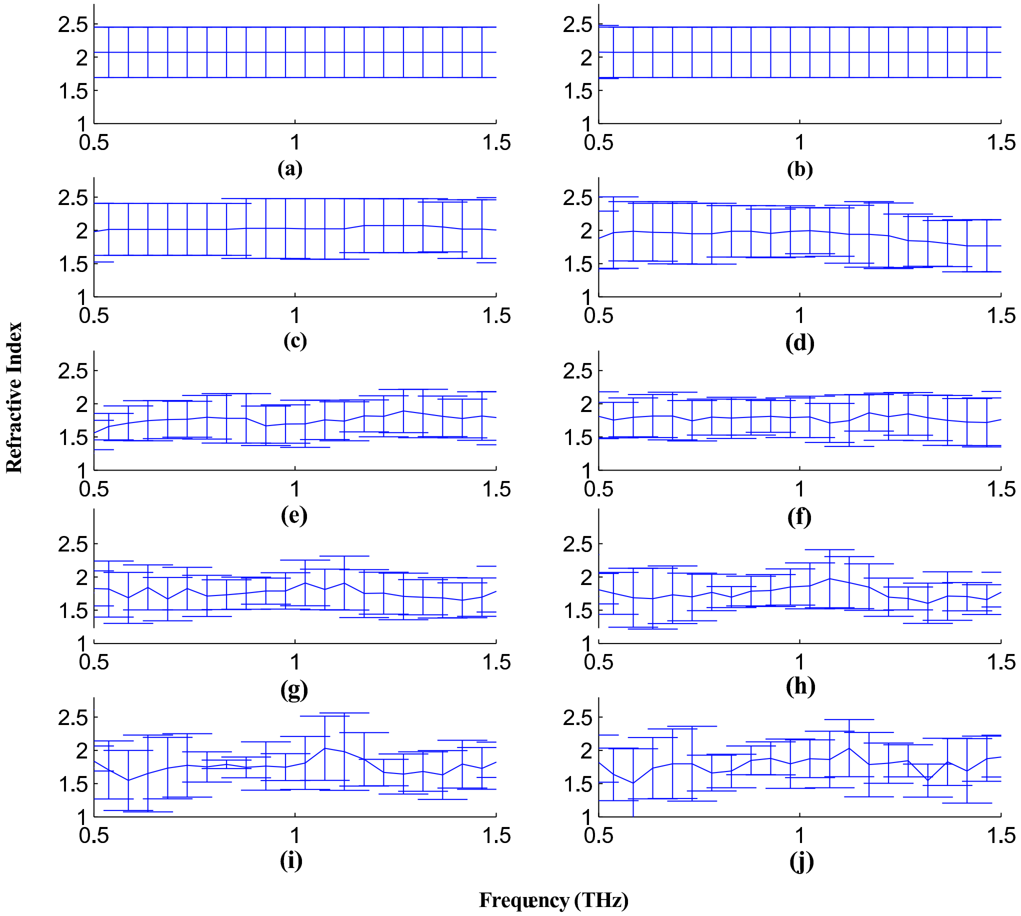

Figure B.4: Refractive index profiles for excised muscle using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

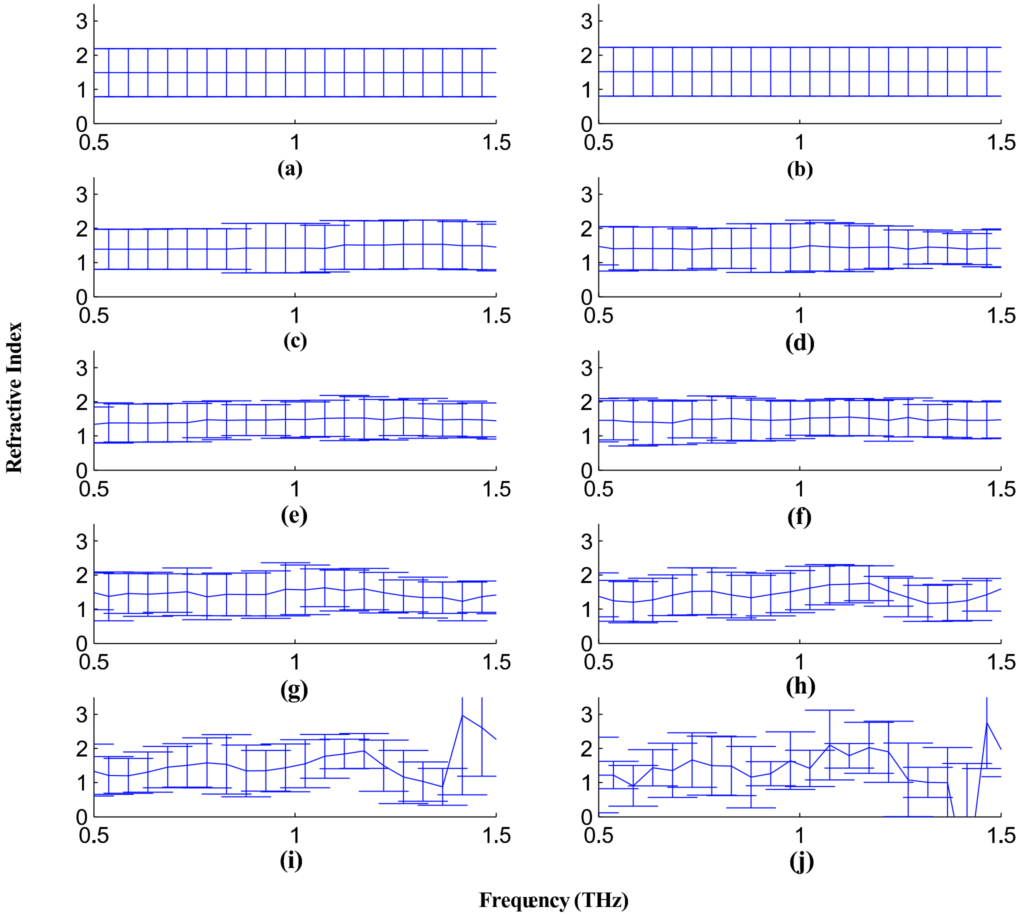

Figure B.5: Refractive index profiles for excised vein using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

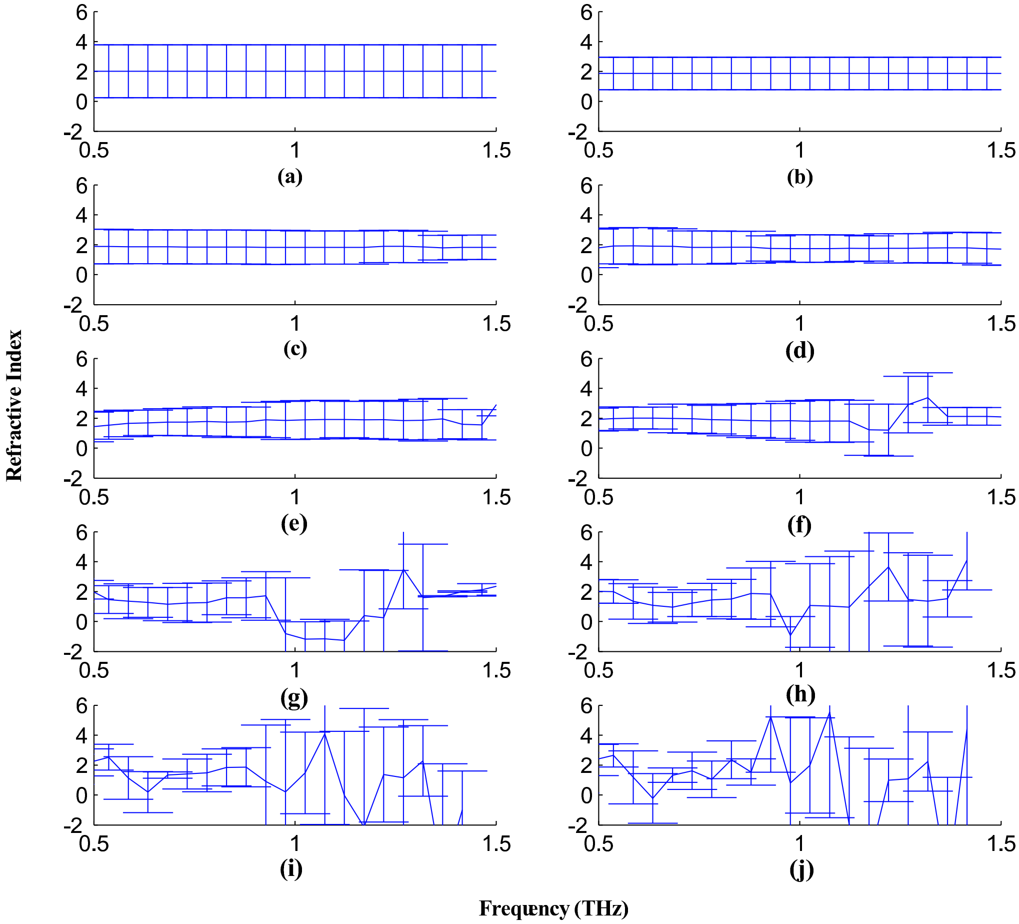

Figure B.6: Refractive index profiles for excised artery using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

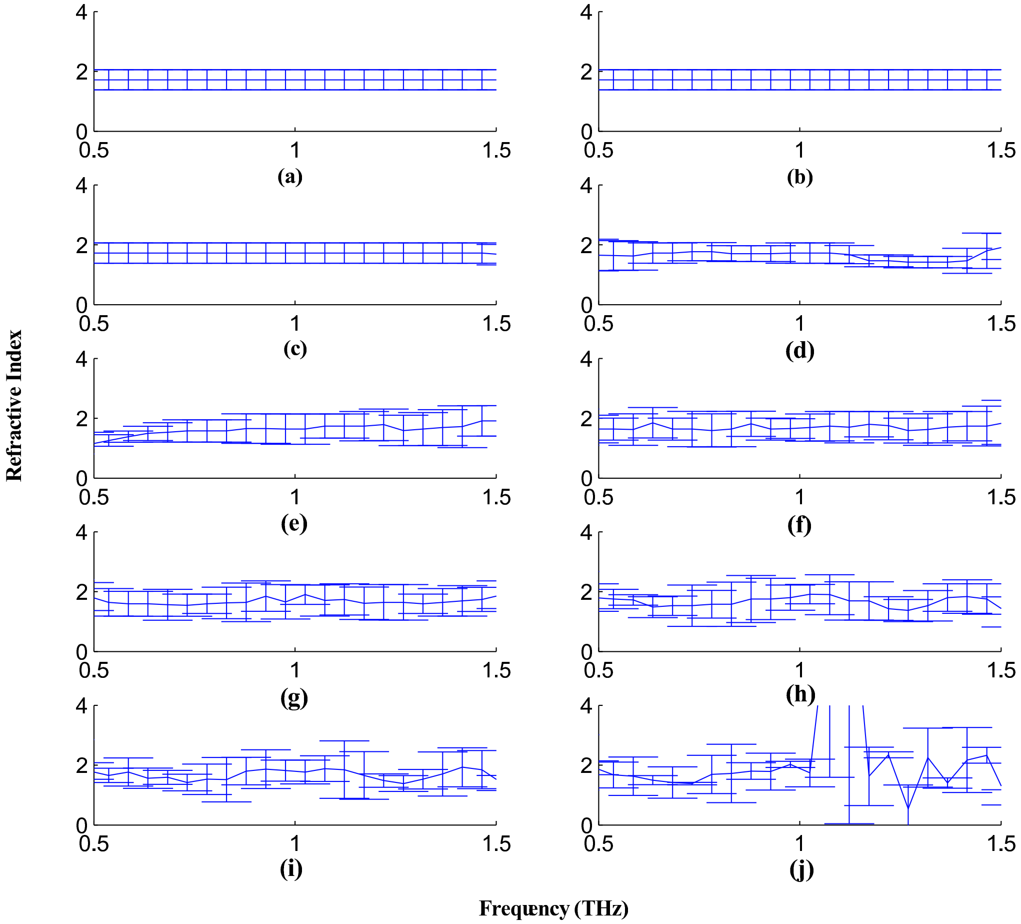

Figure B.7: Refractive index profiles for excised nerve using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

B.2 Absorption Coefficient Profiles

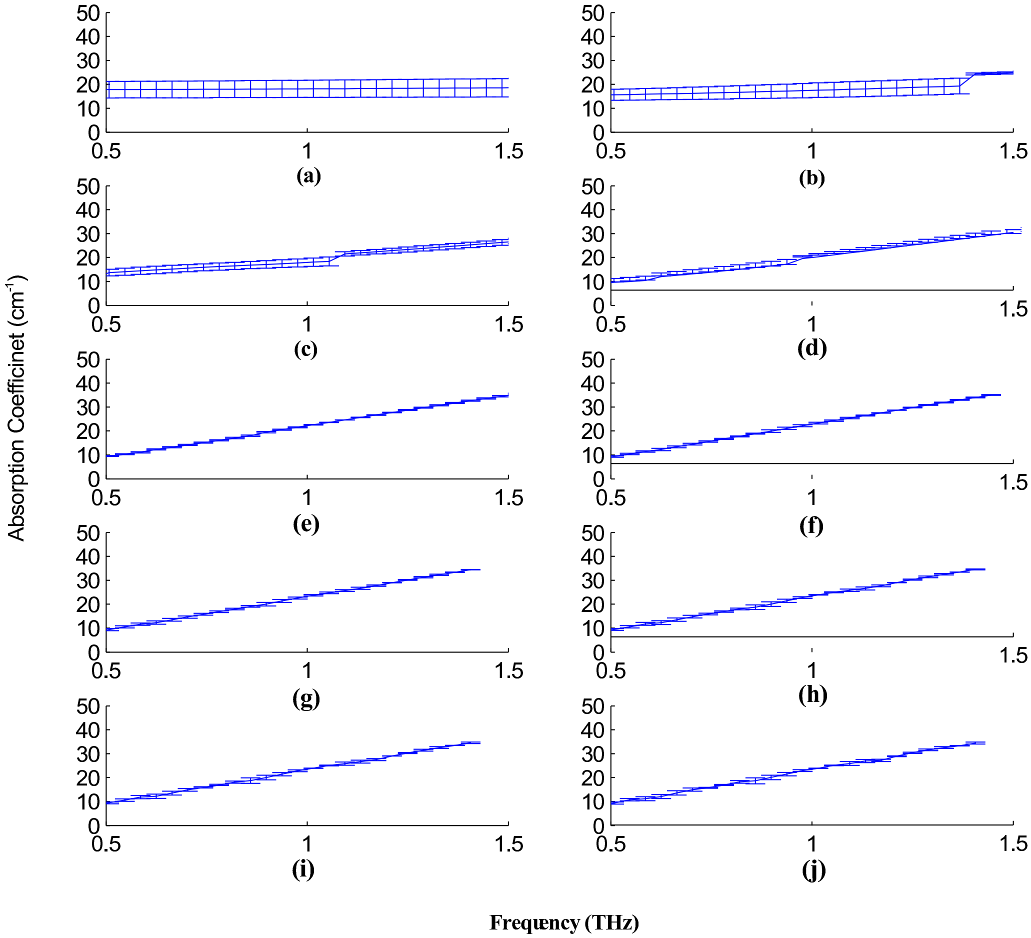

Figure B.8: Absorption coefficient profiles for a resin step wedge using the Gaussian windowed STFT with widths (a) 10^{-1} (0.02 ps), (b) $10^{-0.5}$ (0.063 ps), (c) 10^{0} (0.2 ps), (d) $10^{0.5}$ (0.63 ps), (e) 10^{1} (2 ps), (f) $10^{1.5}$ (6.3 ps), (g) 10^{2} (20 ps), (h) $10^{2.5}$ (63 ps), (i) 10^{3} (200 ps), and (j) $10^{3.5}$ (632 ps).

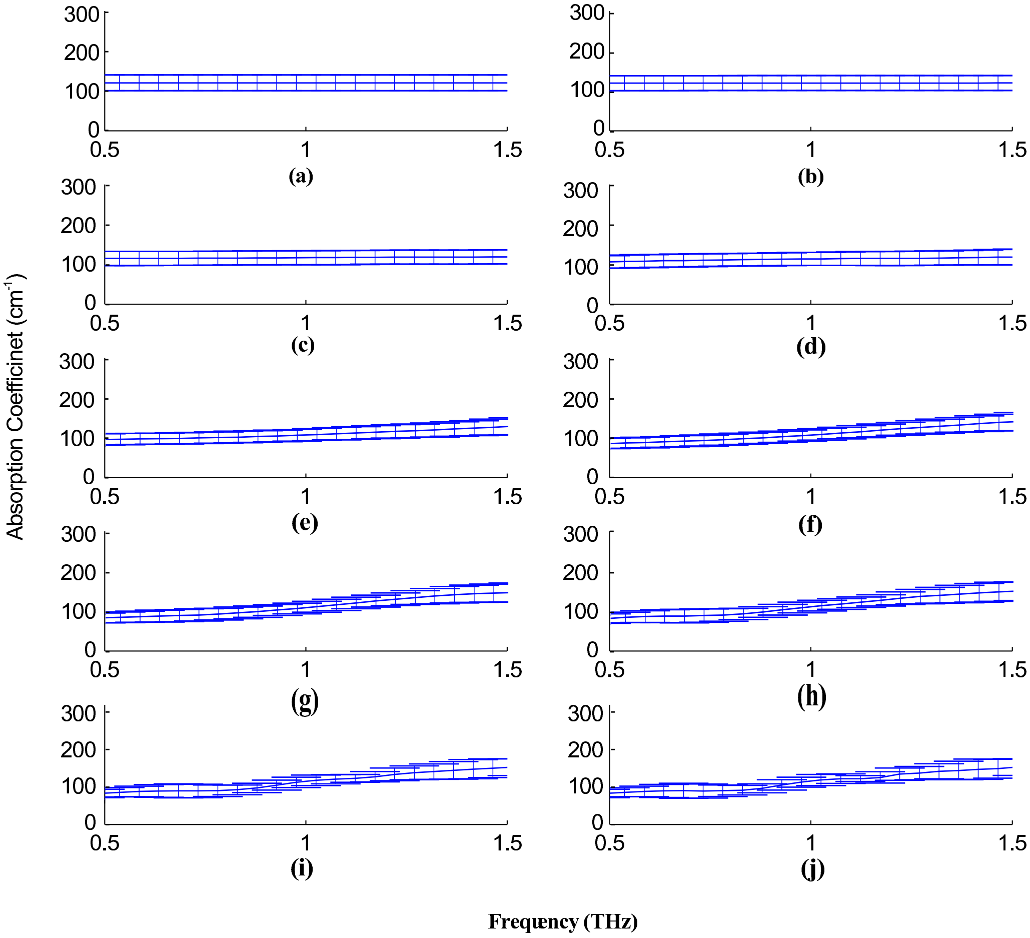

Figure B.9: Absorption coefficient profiles for excised skin using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

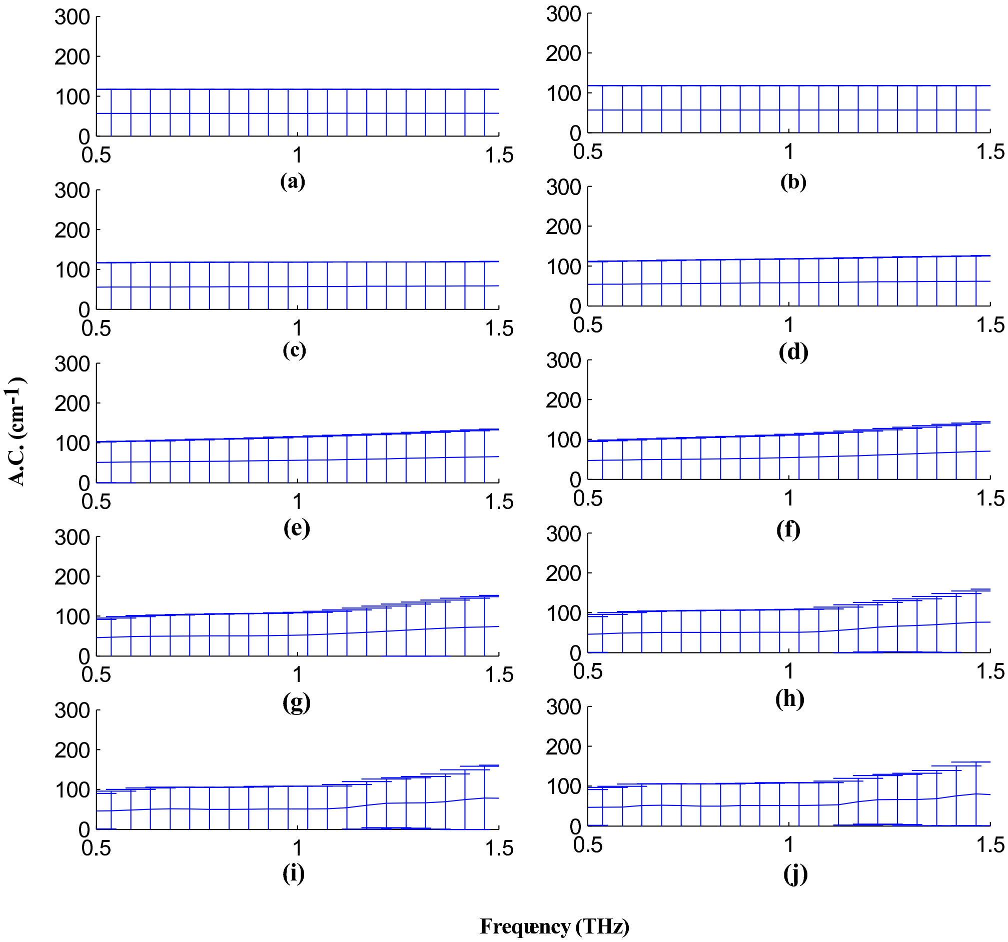

Figure B.10: Absorption coefficient profiles for excised fat using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

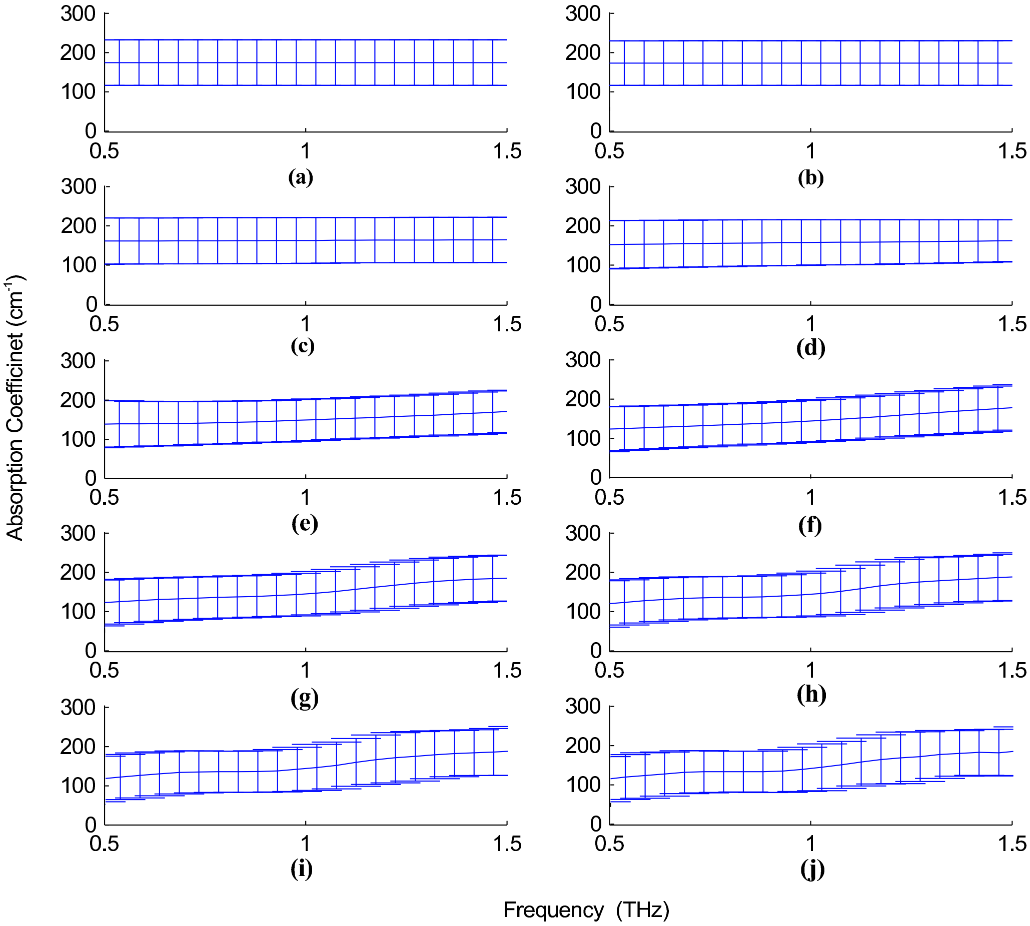

Figure B.11: Absorption coefficient profiles for excised muscle using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

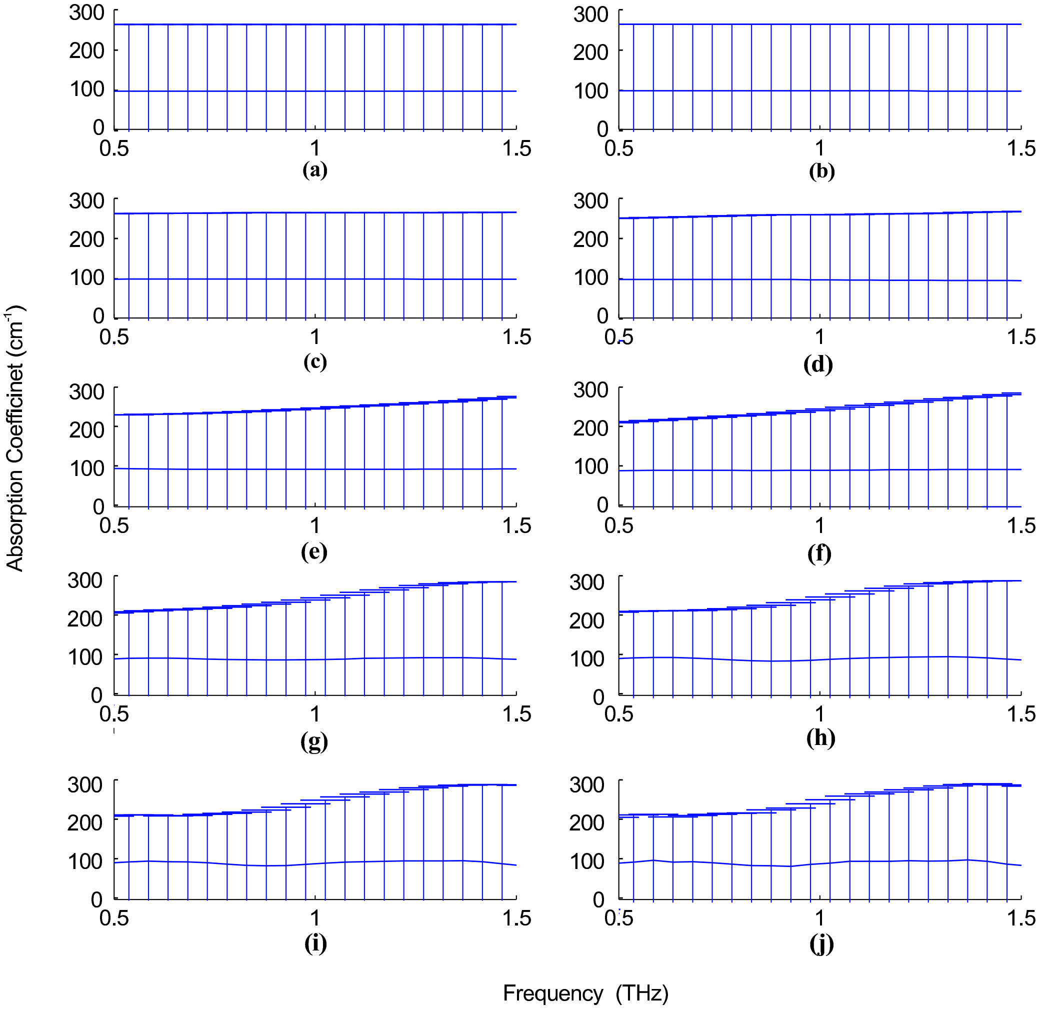

Figure B.12: Absorption coefficient profiles for excised vein using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

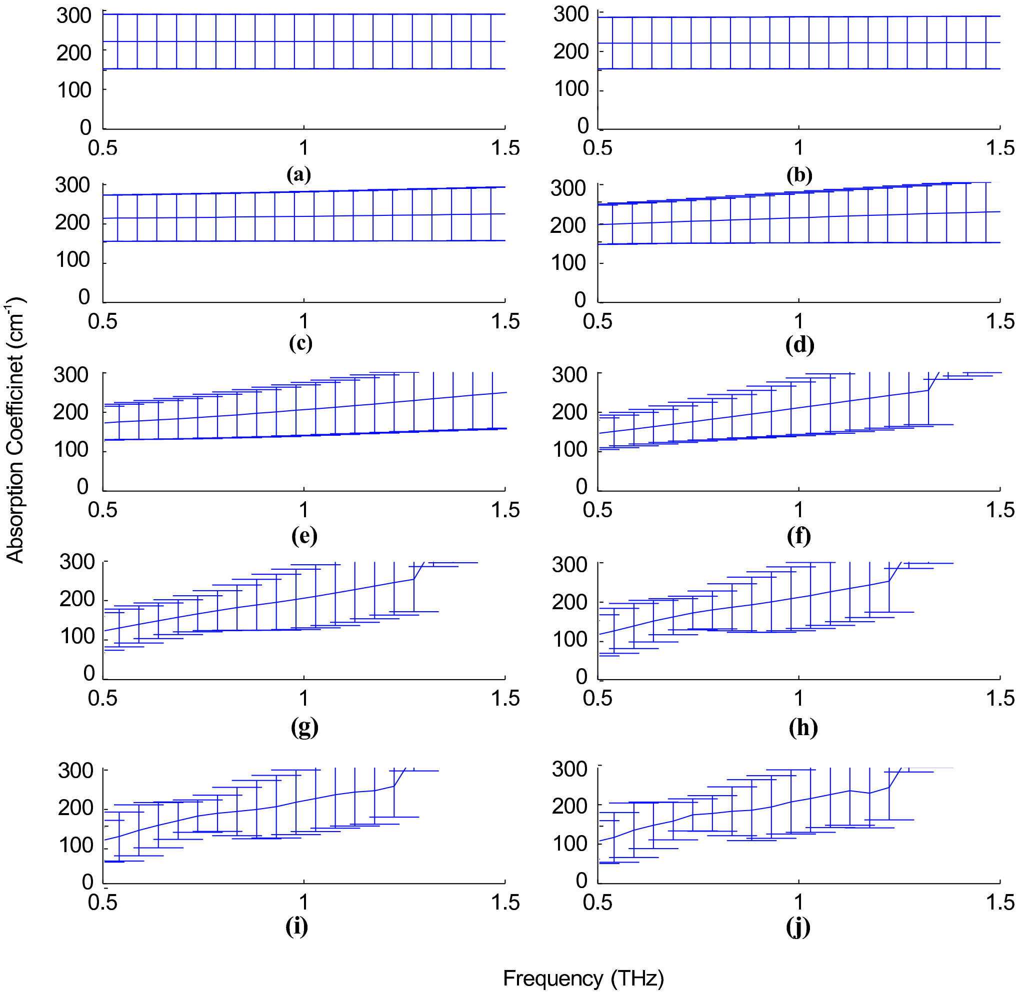

Figure B.13: Absorption coefficient profiles for excised artery using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).

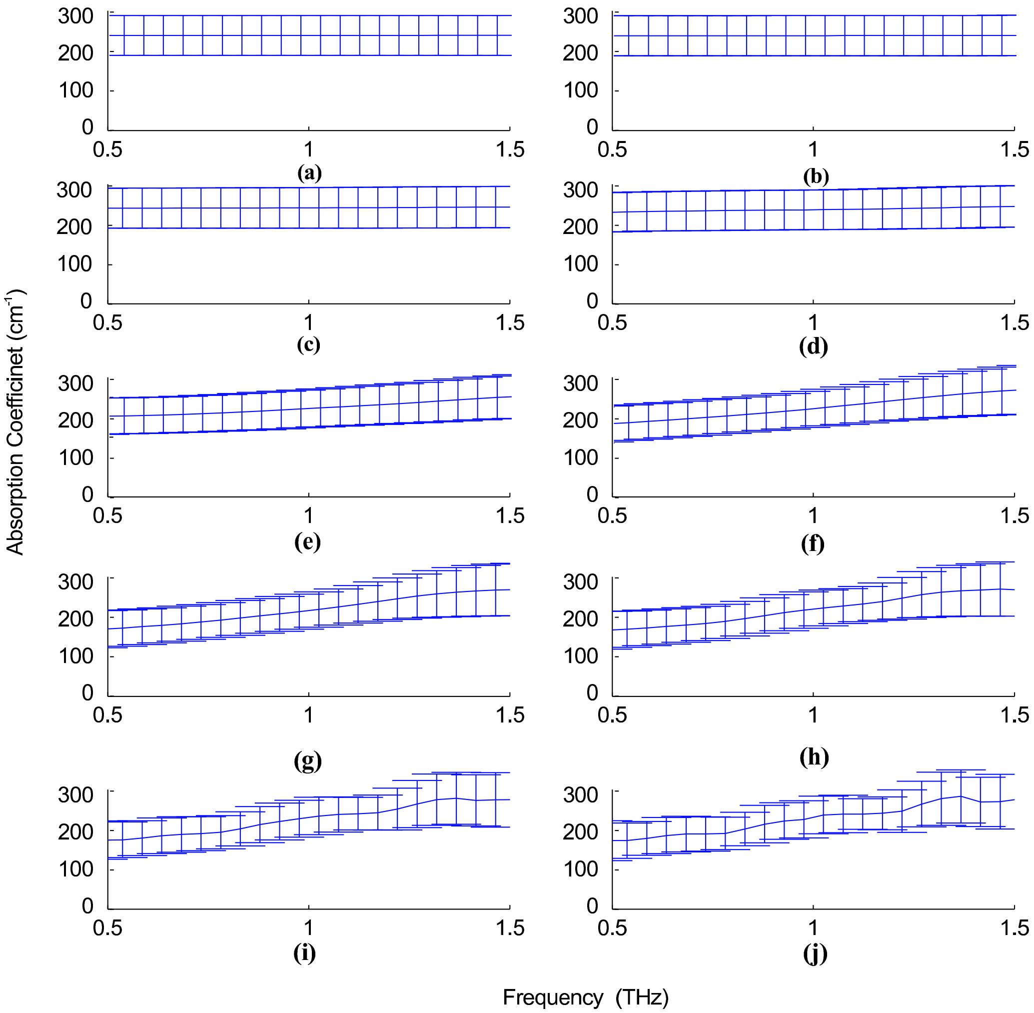

Figure B.14: Absorption coefficient profiles for excised nerve using the Gaussian windowed STFT with widths (a) 10^{-1} (0.008 ps), (b) $10^{-0.5}$ (0.025 ps), (c) 10^{0} (0.08 ps), (d) $10^{0.5}$ (0.25 ps), (e) 10^{1} (0.8 ps), (f) $10^{1.5}$ (2.5 ps), (g) 10^{2} (8 ps), (h) $10^{2.5}$ (25 ps), (i) 10^{3} (80 ps), and (j) $10^{3.5}$ (253 ps).