

Short Title: An Ultrasonic Image Model for Shape Trobaugh, D.Sc. 2000

WASHINGTON UNIVERSITY
SEVER INSTITUTE OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING

AN IMAGE MODEL FOR ULTRASOUND INCORPORATING
SURFACE SHAPE AND MICROSTRUCTURE AND
CHARACTERISTICS OF THE IMAGING SYSTEM

by

Jason W. Trobaugh, M.S.

Prepared under the direction of Professor R.M. Arthur

A dissertation presented to the Sever Institute of
Washington University in partial fulfillment
of the requirements for the degree of

Doctor of Science

August, 2000

Saint Louis, Missouri

WASHINGTON UNIVERSITY
SEVER INSTITUTE OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING

ABSTRACT

AN IMAGE MODEL FOR ULTRASOUND INCORPORATING
SURFACE SHAPE AND MICROSTRUCTURE AND
CHARACTERISTICS OF THE IMAGING SYSTEM

by Jason W. Trobaugh

ADVISOR: Professor R.M. Arthur

August, 2000

Saint Louis, Missouri

Model-based image analysis is a relatively new, but increasingly popular and widely successful, approach to analysis of object shape in images. The approach requires a physically-based model for the image data, a model that can be quite simple for MR and CT images. Existing applications for ultrasonic images employ data models using assumptions that significantly limit the general applicability of their results. Towards a data model that more fundamentally represents ultrasonic images in terms of shape, a probabilistic data model has been developed for ultrasonic images of rough surfaces combining surface shape and microstructure and the characteristics of the imaging system. The model is based on a physical model for image formation

using a linear model for the imaging system with a new, discrete-scatterer model for the tissue surface. The physical model provides the means for simulation of individual images and a mathematical representation for extending the model to a probabilistic form.

From the physical models for the imaging system and tissue surface, the random phasor sum has been used to characterize the amplitude at each pixel in terms of its mean, μ , variance, σ^2 , and SNR_0 , $\frac{\mu}{\sigma}$. Approximations of the amplitude statistics have been derived for surfaces in two ways: 1) using a planar approximation to the pixel's local surface geometry, and 2) using the original surface representation. The amplitude SNR_0 has been used to characterize each pixel as Rayleigh- or non-Rayleigh-distributed. This characterization forms the basis for constructing an image model where pixels are considered Rayleigh-distributed when applicable and Gaussian otherwise.

The image model was evaluated using *in vitro* images of a cadaveric vertebra. Simulated images were generated for visual comparison with actual images. Approximate mean and SNR_0 images were generated using the above techniques for comparison to sample mean and SNR_0 images generated from repeated simulations. The image model, as constructed from the image statistics, was evaluated quantitatively based on performance in inferring the pose of a vertebra from a small set of images.

copyright by
Jason W. Trobaugh
2000

Contents

List of Tables	viii
List of Figures	ix
Acknowledgments	xi
1 Introduction	1
1.1 Inference of Shape	1
1.2 Objectives and Contributions	4
1.3 Organization of Dissertation	5
2 An Application in Spinal Registration	6
2.1 Computer-Assisted Treatment of Spinal Disorders	6
2.2 Ultrasound-based Non-invasive Registration	10
2.3 Problem Definition: A Vertebra <i>In Vitro</i>	13
2.4 Methods: Data Collection	14
3 Pattern Theory and Model-Based Image Analysis	16
3.1 Pattern Theory	16
3.2 Relevant Constructs	18
3.2.1 Surfaces	18
3.2.2 Rigid transformations	20
3.2.3 Image Models	21
3.2.4 Inference and Nonlinear Optimization	22
4 Ultrasonic Imaging	27
4.1 A Linear Systems Model for Image Formation	28

4.1.1	Propagation: The Linear Wave Equation in a Homogeneous Medium	29
4.1.2	The Wave Equation with Sources	31
4.1.3	Low-Level Scattering Sources	34
4.1.4	Transducer sources	36
4.1.5	A Pulse-Echo Equation	37
4.2	Design and Analysis of Imaging Systems	41
4.2.1	Numerical Methods for PSF Modeling	43
4.2.2	Approximate Methods for PSF Modeling	44
4.3	Interpretation of Medical Ultrasonic Images	49
4.4	Tissue Models	51
4.5	Scattering from Rough Surfaces	52
4.6	Probabilistic Ultrasound Models and Statistical Tissue Characterization	52
4.7	Conclusions	57
5	A Discrete-Scatterer Model for Rough Surfaces	58
5.1	Image Formation	58
5.1.1	Imaging System Model	59
5.1.2	Surface Model	59
5.2	Methods	62
5.2.1	System Model Implementation	63
5.2.2	Surface Model Implementation	64
5.3	Results	65
5.4	Discussion	69
5.5	Conclusions	71
6	Amplitude Mean and Variance for the Random Phasor Sum	73
6.1	Image Formation as a Random Phasor Sum	73
6.2	The Gaussian Approximation to the Complex Sum	74
6.3	Computing the Complex Gaussian Parameters	76
6.4	Computing the Amplitude Mean and Variance	81
6.5	Relevant Issues for the Image Model	84
7	Image Model Statistics for Surfaces: Theory	86
7.1	The Random Phasor Sum and Tissue Surfaces	86

7.2	Computing Moments Directly from the Triangulation	88
7.3	Computing Moments Using a Locally Planar Approximation to the Surface	90
7.3.1	Area of intersection	91
7.3.2	Moments of functions of A_i, ϕ_i	92
7.4	Computing Moments with the Planar Approximation and $A_i \perp \phi_i$. .	94
7.4.1	Moments for scatterer amplitude	94
7.4.2	Axial extent of intersection	95
7.5	Surface Roughness	96
7.6	Locally Planar Approximations for Arbitrary Surfaces	96
8	Image Model Statistics for Surfaces: Results	99
8.1	General Results from the Planar Surface Approximation	99
8.1.1	Effects of PSF, surface and plane parameters	100
8.1.2	Effects of Surface Curvature	105
8.1.3	Discussion: Value of the Planar Approximation	111
8.2	Results for Arbitrary Surfaces	114
8.2.1	An Image Plane on the Transverse Process	114
8.2.2	Transverse process images with a reduced elevation width . . .	120
8.2.3	Sagittal plane images along the lamina and articular processes	123
8.3	Discussion	128
8.4	Conclusions	134
9	Inference of Vertebral Pose Using the Image Model	135
9.1	A Rayleigh/Gaussian Image Model	135
9.2	Inference	137
9.3	Results	139
9.3.1	Results - Actual Image	139
9.3.2	Results - Simulated Images	141
9.4	Implications for Inference	146
10	Conclusions and Future Directions	147
10.1	Conclusions	147
10.2	The Physical Models for Image Formation	148
10.3	The Image Model	149
10.4	Inference of Shape Using the Image Model	150

10.5 Adaptive Image Formation and Tissue Characterization	151
Appendix A Calculations involving triangles	152
Appendix B A Parametrization for the Intersection of a Plane and Ellipsoid	155
References	161
Vita	168

List of Tables

6.1	Table of moments for uniform phase	80
-----	--	----

List of Figures

1.1	Typical spinal ultrasonic image	2
2.1	Components of an image-guided surgery system.	8
2.2	Traditional open surgical approach to the spine.	9
2.3	Ultrasonic image 1 of the spine with corresponding CT image.	12
2.4	Ultrasonic image 2 of the spine with corresponding CT image.	12
2.5	Ultrasonic image 3 of the spine with corresponding CT image.	13
2.6	Sample <i>in vitro</i> image.	14
2.7	Phantom and rendered vertebral surface.	15
3.1	Sample triangulated surface rendering	20
4.1	Ultrasonic imaging systems	28
4.2	Typical axial pulse	49
4.3	Typical ultrasonic image	51
5.1	Sample actual and simulated images	66
5.2	Image plane and triangulated surface	66
5.3	Zoomed view of actual and simulated images.	67
5.4	Sensitivity of simulated images to gross shape	68
6.1	Plots of moments for uniform phase	80
6.2	Amplitude mean and variance vs. phase width	84
8.1	Effects of rotation about the y axis.	101
8.2	Effects of rotation about the x axis.	102
8.3	Mean and SNR_0 versus rotation and translation of plane.	103
8.4	Mean and SNR_0 versus rotation and concentration.	103
8.5	Mean and SNR_0 versus rotation and roughness.	104

8.6	Sample sphere images from curvature experiment.	107
8.7	Simulation and planar mean for 8 mm sphere.	109
8.8	SNR ₀ images for the sphere.	110
8.9	Effect of curvature for various PSF sizes.	110
8.10	Sample images for spherical and ellipsoidal PSFs.	111
8.11	Sample images generated from planar statistics for sphere images. . .	113
8.12	Image plane along the transverse process.	115
8.13	Sample simulated images for usual PSF	116
8.14	Simulation, planar and triangle means for transverse process images. .	118
8.15	Simulation, planar and triangle SNR ₀ images.	119
8.16	N,d images for planar approximation.	120
8.17	Sample simulated images for PSF with small elevation PSF width. . .	121
8.18	Mean images for reduced elevation width $\sigma_x = 0.5$ mm.	122
8.19	SNR ₀ images for reduced elevation width $\sigma_x = 0.5$ mm.	122
8.20	N and d images for small elevation width.	122
8.21	Views of image plane for lamina image	123
8.22	Sample simulated images for image of lamina and articular processes. .	124
8.23	Simulation and planar images for lamina.	125
8.24	Simulation and triangle mean images	126
8.25	Simulation and triangle (4x resampling) mean images.	127
8.26	Simulation and triangle SNR ₀ images.	127
8.27	Samples of image model for transverse process and typical PSF. . . .	131
8.28	Samples of image model for transverse process and narrower PSF. . .	132
8.29	Samples of image model for lamina image.	133
9.1	Sensitivity of the log likelihood to rotation and translation.	138
9.2	Results with an actual image of the lamina	141
9.3	Data set for tests using simulated images.	142
9.4	Results for a set of simulated images.	144
9.5	Scenes from an optimization from 4 mm, 4 degrees.	145
A.1	Equi-area resampling of a triangle.	153

Acknowledgments

Over the course of this work, I have received technical assistance from many collaborators. From Surgical Navigation Technologies, Paul Kessman and Troy Holsing assisted in areas from clinical software and phantom construction to various discussions. Sarang Joshi and Rob Teichman helped with the construction of triangulated surfaces from the spinal CT data set. Drs. Richard Bucholz and Jaimie Henderson at the Department of Neurosurgery in the Saint Louis University School of Medicine provided clinical resources. Dennis Dietz and the Tetrad Corporation assisted in development with their commercial ultrasound system. During his stay at Washington University in St. Louis, Alessandro Chiuso was of assistance in verifying and simplifying some mathematics, in addition to serving as a much-needed climbing partner.

For their support, I owe Surgical Navigation Technologies, now a subsidiary of Medtronic, Inc., a great deal of gratitude. They supported me financially for the past four years and supplied resources, from software to clinical collaborations, without which this work would not have been possible. Kurt Smith was largely responsible for this opportunity, and I thank him personally for his efforts.

I am sincerely grateful for the chance to be a part of the Electronic Signals and Systems Research Laboratory. With its faculty and students, it provides a rich environment for academic growth. The scientific, rigorous approach to imaging and modeling, has deeply affected my work. A special thanks to my advisor, Martin Arthur, for his willingness to work with me under unconventional circumstances, for his confidence and support, and for the many discussions on life, religion, existence and his birth in a supernova.

Finally, the continuing support of my family and friends has allowed me to retain some sanity. My thanks to all of them, especially Mom, Dad, Angela, Darin, Kait, Steve, Alessandro, Dave, Dawn and Abby.

Jason W. Trobaugh

*Washington University in Saint Louis
August 2000*

Chapter 1

Introduction

Despite its tendency to befuddle the average viewer, medical ultrasound remains in common practice because of its many advantages. Relative to x-ray computed tomography (CT) or magnetic resonance imaging (MR), it is inexpensive and extremely portable. Ultrasonic energy is non-ionizing, and, for most patients, an ultrasound exam is much more pleasant than the often-feared CT or MR scan. The biggest downside to ultrasound is the overwhelming amount of experience required for an operator to become skilled in the art of interpreting the images. This difficulty could be simplified greatly by a robust approach to computer-assisted image interpretation. If successful and accurate, an approach could expand the range of applications for ultrasound beyond common diagnostic usage. For instance, ultrasound has been investigated by many medical researchers as a means for tissue registration for image-guided surgery and delivery of radiosurgery [1, 2, 3].

1.1 Inference of Shape

Within image interpretation, or image analysis, this dissertation focuses on the inference of tissue shape. While shape analysis is difficult with any image modality, it is especially difficult in medical ultrasonic images. Consider the image in Figure 1.1, with the associated description of gross structure on the right. The region of skin at the top of the image is marked by a low-intensity, homogeneous speckle texture, the fascial layer by a slightly higher-intensity homogeneous texture. The muscle layer is identified by the presence of individual fibers within the layer, where the intensity varies along the fibers, depending on the orientation of the fiber with respect to the

axial direction (down in the image). Finally, the transverse process of a lumbar vertebra is marked by both a varying speckle texture, along the sides of the process, and some high-intensity, coherent echoes at the top of the process where the normal to the surface is aligned with the axial direction. Such variation is typical in all ultrasonic images, with image features representing a mix of the random speckle texture (as seen for the skin) and coherent echoes (as seen for some of the muscle fibers and for the top of the transverse process).

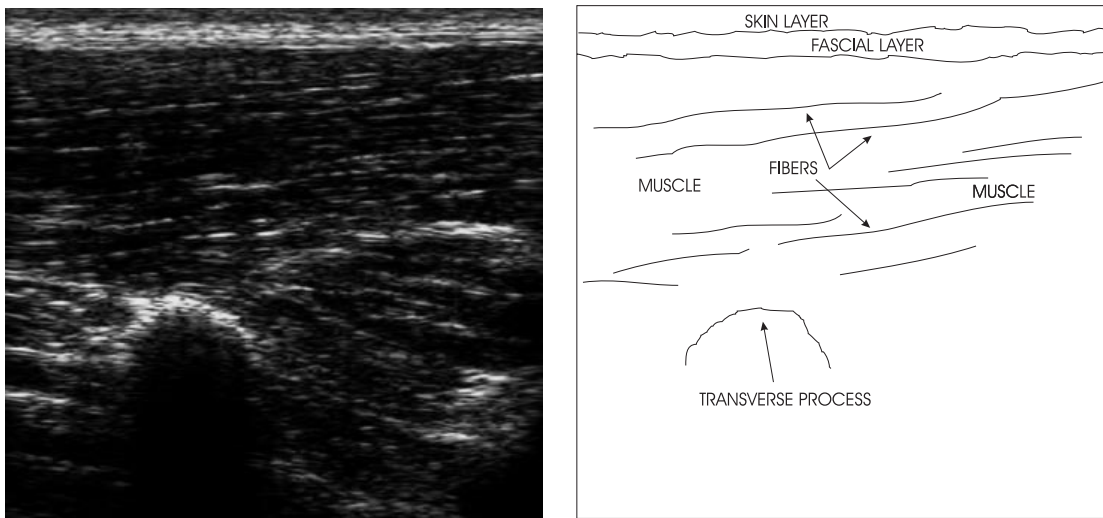


Figure 1.1: The spinal ultrasonic image on the left shows typical characteristics of several tissue regions, skin, fascia, muscle and bone. The drawing on the right shows the locations of selected tissue regions.

With such high variation in ultrasonic image intensities, one should not expect much success from traditional, image-processing approaches to pattern recognition. Approaches based on the simplest techniques, e.g., thresholding, are not even attempted. More sophisticated approaches with simple statistical models, e.g., constant Rayleigh statistics for a given region, have seen limited success [4, 5, 6, 7]. In typical approaches, features are extracted from the images and then composed in some way to form patterns. The features may be based on statistics at a pixel or texture in a region surrounding a pixel. Pixels are then classified and composed into patterns, where the structure of the patterns may be constrained according to reasonable assumptions about some basic properties, e.g., smoothness of a border. The simplicity of these approaches could make them valuable in some situations. In most ultrasonic images, however, the combination of varying speckle texture and intense, coherent echoes that are sensitive to the angle of insonification, will cause extracted features

and composed patterns to vary significantly regardless of the level of sophistication used to represent them. As a result, the traditional approaches are only moderately effective for such highly variable images.

In addition, the lack of a deep, underlying framework to approaches such as these can leave one with little understanding of the failures and successes of the approach or of the underlying problem. Modifications must be made according to intuition, are usually *ad hoc*, and must be judged purely by their effects on the solution, i.e., in a manner external to the problem. Assessment of the success of an approach is typically based on empirical data for lack of an underlying framework in which to couch the evaluation. Further, without a solid framework, the understanding afforded by such approaches has little carryover to similar problems.

Recently, Bayesian methods based on Grenander's pattern theory have demonstrated great potential in a variety of image analysis problems [8, 9, 10]. This approach combines explicit, detailed representations of structure shape with physically-based models for the image data. For a given structure, shape is represented with a template model and probabilistic models for its possible variations. Estimates of shape from a new observation employ an *a priori* probability density function, a probabilistic model for transformations of the shape, and a data likelihood, a conditional density representing the image data for a given shape (specified by a transformation from the template). Standard statistical estimation techniques, e.g., maximum likelihood, minimum-mean-squared-error, etc. can then be used for inference. Such a rigorous, probabilistic representation for shape often requires difficult mathematics as seen in the literature [8, 9, 10]. Given these results, though, the approach can be applied to any image analysis problem with appropriate shape constructs and an image model.

Application of these methods to ultrasonic images, thus, requires a probabilistic model describing the ultrasonic image data as a function of structure shape. For use in analysis of ultrasonic images, a probabilistic image model must accurately represent the variation of image intensity given only the gross surface shape. Since characteristics of the imaging system and surface microstructure both influence the image, the effects of these characteristics must be inherent within the model.

Many probabilistic models have already been developed for representing ultrasonic image data [11, 12, 13, 14, 15, 16, 17]. These models are based on a random walk representation of scattering that is intractable except under special circumstances. As a result, researchers have neglected the influence of such complicating factors as the system point-spread function (PSF) and the shape of underlying structure. For the

traditional purpose of characterizing the tissue microstructure, useful results may be achieved under these assumptions. Studies have shown, though, that data statistics can vary significantly from predicted distributions due to many factors, including the PSF [16, 17]. Regularity in tissue structure is also known to affect the statistics of scattering data [18]. A treatment incorporating the PSF has been developed for the one-dimensional (1D) case [17], although the exact solution is non-trivial, requiring intensive Monte Carlo simulations.

1.2 Objectives and Contributions

In the substantial literature on probabilistic ultrasound models, no comprehensive, pixel-based model exists that incorporates the system characteristics, tissue microstructure and gross shape. This dissertation comprises a significant step towards such a model, developed in detail for the surface aspect of gross shape. The objectives of this work were as follows.

1. To develop a physical model for image formation, including the effects of surface shape and microstructure, and characteristics of the imaging system. The physical model forms the basis for simulating individual images and the mathematical basis for the other objectives.
2. To develop methods for computing amplitude statistics at each pixel. Pixel-based statistics, the amplitude mean and variance, represent variation at each pixel and permit the construction of a likelihood model.
3. To develop a likelihood image model conditioned on shape. The likelihood model permits inference of the underlying shape using statistical estimation techniques.
4. To investigate the applicability of the image model to inference of shape. The data likelihood is used, in conjunction with nonlinear optimization approaches, for an algorithm designed to find the maximum likelihood estimate of vertebral pose for registration.

1.3 Organization of Dissertation

The previous objectives has been investigated in the context of ultrasound-based registration of the spine for image-guided surgery and delivery of radiosurgery. In Chapter 2, this application is dicussed, along with methods for data collection that will be used throughout the dissertation. Chapter 3 covers background on pattern theory and model-based image analysis. In Chapter 4, a physical model for ultrasonic imaging systems is derived, and relevant literature is reviewed. In chapter 5, a discrete-scatterer model for tissue is presented to complete the physical model for image formation. The basis for computing pixel-based amplitude statistics is developed for a 1D, axial image model in Chapter 6. In Chapter 7, methods for computing pixel-based statistics are developed that employ a locally planar approximation to the surface. This method, and another approximation based on direct calculations from the surface model, are investigated in Chapter 8. A data likelihood constructed from the pixel-based statistics is used in Chapter 9 in an investigation into the use of nonlinear optimization techniques for maximum likelihood estimation of vertebral pose. The dissertation concludes with Chapter 10, covering the image model and its possible impact on clinical and research directions within ultrasound.

Chapter 2

An Application in Spinal Registration

One motivation for this work comes from treatment guidance applications for the spine. The application serves as more than motivation, though. The treatment guidance environment also supplies nearly all the methods for data collection, including CT scans, cadaveric vertebrae, registration techniques, methods for constructing surfaces, and equipment and methods for tracking the ultrasonic image plane. The spinal vertebrae, the anatomical structures of interest in the application, provide a good test medium for the image model because of their extreme curvature. An image model that accurately represents the sensitivity to their intricate features should fare well in application to simpler structures.

2.1 Computer-Assisted Treatment of Spinal Disorders

In the human body, the spine plays the vital role of protecting the spinal cord, the center of control for nearly the entire body. Disorders of the spine, from degenerative disease to spinal trauma, can distort the normal spine and cause difficulties from mild discomfort to paralysis and complete loss of musculoskeletal function [19]. For many conditions, surgery is a primary treatment option in the form of stabilization of the spine, correction of deformity, or decompression of neural tissues. By its nature, surgery is invasive and includes risks of further damage to vital tissues. Exposure is necessary not only for access to the tissues but also for navigational purposes, allowing

the surgeon to see into the body. Improvements in surgical techniques and equipment often involve reducing invasiveness, increasing navigational capacity or both.

Technological advances have greatly improved the means for navigation and, at the same time, have allowed less invasive procedures to enable surgeons to produce better outcomes in less time with more accuracy [20]. The surgical microscope, popularized in the 1970s and widely used today, produces high-resolution video images of the anatomy that allow the surgeon to operate on a small scale. The surgical endoscope, an optical instrument that can be inserted into the body through a small incision, provides video of the internal anatomy for navigation. Medical imaging devices such as CT (Computed Tomography) and MRI (Magnetic Resonance Imaging) produce detailed images of the patient anatomy that allow the surgeon to visualize structures and relationships directly, e.g., the pedicle and adjacent bony structures in the vertebra, instead of having to rely on general knowledge about those relationships. These detailed images are also used in two of the latest fields of navigational technology, image-guided surgery and stereotactic radiotherapy [20, 21, 22].

In image-guided surgery, volumetric CT or MR images are interactively displayed to show positions of surgical instruments for navigational assistance. Instrument position is tracked using spatial localization technology and displayed on orthogonal and 3D images as shown on the left in Figure 2.1. The images are used in visualization of structures, planning of the procedure, and intra-operative guidance. On the right in Figure 2.1, a complete system for image-guided surgery is shown, including workstation, surgical instruments and a spatial localization system. In this system, an electro-optical camera array is used to track individual infra-red LEDs with sub-millimeter accuracy. By attaching LEDs to instruments, the positions and trajectories of those instruments can be monitored and displayed on images as in Figure 2.1. Use of these techniques is becoming common in cranial, spinal and orthopedic procedures and is expected to result in increased accuracy and reduced invasiveness [20, 23].

Similar technology is used in planning and guidance for stereotactic radiosurgery and radiotherapy [21, 22]. Radiation beams can be guided with accuracy of less than one millimeter to targets precisely identified on volumetric CT or MR images. Malignant tumor tissue is affected more deeply than normal healthy tissue, and successful treatment of disease can be achieved through a single, high-powered dose to a small target (radiosurgery) or through repeated, fractionated doses to a larger target (radiotherapy).

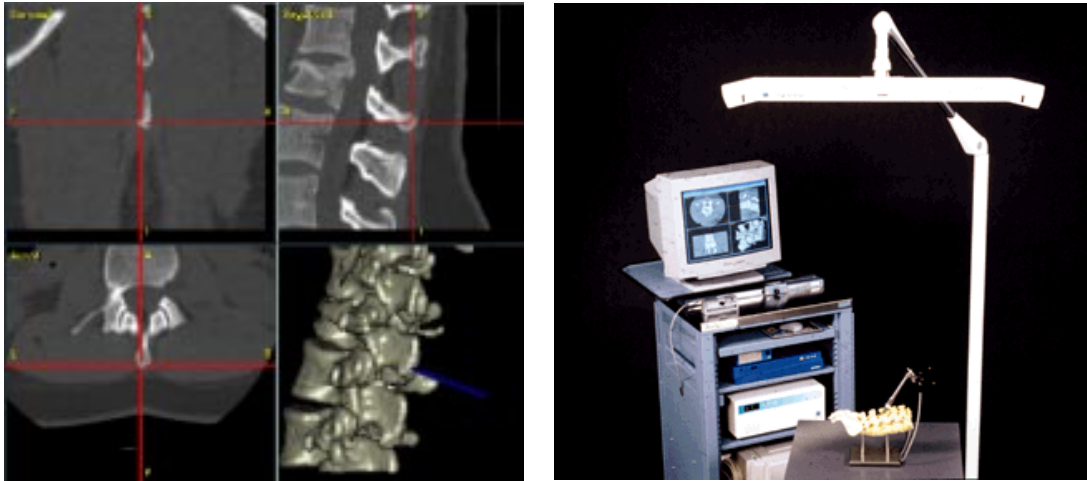


Figure 2.1: Components of an image-guided surgery system. On the left is a typical display of a spinal CT data set showing the position of an instrument on images in anatomical planes and on a 3D model. Clockwise from the lower left, the windows are the 1) axial, 2) coronal, 3) sagittal planes, and 4) a 3D rendering of the vertebrae. The position of the pointer in the 3D model is indicated by the crosshairs in the anatomical planes. The image on the right shows the physical components of the system: workstation, electro-optical camera array for instrument localization, modified surgical instrumentation (and lumbar phantom). Images courtesy of Surgical Navigation Technologies, a subsidiary of Medtronic, Inc.

A surgical example involving the surgical placement of screws into the vertebral pedicle illustrates some of the benefits of image guidance technology. Spinal instability is commonly corrected through a spinal fusion, in which two or more adjacent vertebrae are fused together. In the posterior surgical approach, i.e., approach from the posterior side, the vertebrae are exposed as shown on the left in Figure 2.2. The goal of the surgery is to place a bone graft between the vertebrae as shown in the center diagram of Figure 2.2. The graft grows over time to unite the vertebrae and stabilize the spine. Spinal implants, e.g., metal hardware consisting of screws and plates, are installed to immobilize the vertebrae during the healing process to increase the chances of a successful fusion. On the right of Figure 2.2, screws are shown installed in adjacent lumbar vertebrae through the pedicles and into the vertebral bodies with additional hardware securing the vertebrae. The implants can cause complications by interfering with neural tissues [19]. In a traditional open procedure, as in Figure 2.2, the surgeon must blindly determine the trajectory and distance for screw placement based on general knowledge of relationships between the pedicle and adjacent structures. Judgement errors within one or two millimeters can place

the screw in contact with sensitive nerves and arteries or possibly even fracture the vertebra [23, 24].




Figure 2.2: Open, posterior exposure of the spine (left), bone grafting (center), and placement of pedicle screws (right). All muscles on the posterior aspect of the vertebrae are scraped away from the bone and retracted to expose the vertebral column. Bone harvested from transverse processes is placed between the vertebral bodies and between spinous processes. Screws are shown inserted through pedicles and deep into the vertebral bodies for securing fusion by the bone graft. From Esses, *Textbook of Spinal Disorders*.

The surgical endoscope, the optical instrument that allows video of the internal anatomy through a small incision, has provided the initial means for enabling minimally invasive spinal fusion [25, 26]. Bone grafting in the intervertebral spaces can be accomplished without a complete posterior exposure using endoscopic guidance. Screw placement still presents a danger since the entire pedicle cannot be visualized with the endoscope alone. Image guidance technology allows the surgeon to precisely plan and guide screw placement through the pedicle and into the vertebral body using patient-specific anatomical data from the volumetric images. A posterior exposure is still required, however, for image registration, the process of spatially relating the images and the patient.

Image registration is required for both surgical guidance and delivery of radiation treatment. Registration is typically achieved by identifying corresponding landmarks in the images, via the workstation, and on the patient, using a probe tracked by the localization system [3, 23, 24]. In radiation treatments, this is accomplished with a stereotactic frame that is rigidly attached to the vertebrae. The CT scan is taken after the frame is attached, and landmarks on the frame are then easily identified in images and directly on the frame. In surgery, images are registered to a single vertebra using anatomical landmarks on the vertebra. Identification directly on the vertebra requires a traditional posterior exposure. The exposure process is lengthy

and significantly increases recovery time and pain for the patient beyond that necessary for hardware placement and bone grafting. A minimally-invasive registration procedure could improve surgical efficiency and reduce patient difficulties.

The accuracy of a registration is crucial for determining its applicability in a given procedure. For example, the number of complications occurring with screw placement increases with the uncertainty in registration. In this procedure, registration accuracy on the order of one to two millimeters, typically defined as the maximum translational error within some region of the target point, is generally considered necessary to justify the use of image guidance equipment [27]. Algorithms for registration based on landmarks are listed in [3]. Landmark-based techniques are limited in spinal registration because of ambiguity in identifying specific landmarks on the vertebra. Currently, a surface-based refinement of the landmark-based registration procedure is used that results in improved accuracy, but no reduction in exposure. Given an approximate landmark-based registration, a set of vertebral surface points (with no corresponding landmarks in the image set) are matched to a surface model of the vertebra to refine the registration [3, 28, 29]. The accuracy achieved by this method is the standard by which any minimally-invasive registration procedure must be judged. Any loss in accuracy must be weighed against the reduction in exposure that is achieved.

2.2 Ultrasound-based Non-invasive Registration

The use of ultrasound as a means for achieving non-invasive spinal registration was introduced in 1993 with results on a spinal phantom [1]. Later results include preliminary studies on clinical data [2]. Any localization system can be used to track an appropriately modified ultrasound probe and, thus, the associated image plane. By scanning the spine with such a probe, vertebral landmarks and a set of surface points can be identified on the patient via the ultrasonic images to achieve a completely non-invasive registration [1, 2, 27]. Other possibilities for non-invasive registration include the use of fluoroscopy data and are listed in [27]. The primary advantage of ultrasound over fluoroscopy is the lack of ionizing radiation. The use of fluoroscopy exposes the patient to additional risk and increases the time required for a procedure because the surgical staff must leave the room during the exam or be exposed to the radiation.

While ultrasound has been used successfully on phantoms and clinically to non-invasively register the spine, many issues must be resolved before such a registration can be safely used in guiding a surgical procedure. The accuracy of the technique is currently being investigated in clinical studies using manual interpretation of the ultrasonic images and currently available registration techniques. Even if accuracy requirements can be met, use of the technique will be limited by the experience required by the operator and the time required to interpret the images. Applicability would be greatly enhanced if the technique could be partially or fully automated, reducing the necessary time and user expertise.

Ultrasound images are inherently difficult to interpret due to low contrast resolution for many tissues, a relatively poor signal-to-noise ratio due to depth-dependent attenuation, and textural representations of tissue regions that may appear nonsensical to an inexperienced observer. Ultrasound images of the spine are no exception and may be arguably worse. The images exhibit a high degree of variability due to imaging mechanisms and anatomical variation. Images vary substantially with the scanning angle of the probe due to the angle-dependent intensity and texture of scattering from the vertebral surface. The signal-to-noise ratio varies with tissue composition due to depth- and tissue-dependent attenuation of the ultrasonic pulse used to interrogate the tissue. The interface between fat and muscle layers causes further image degradation due to phase aberration, distortion of the insonifying pulse, a problem that has been studied extensively in abdominal imaging [30]. Vertebral shape and the composition of fat and muscle tissue vary significantly both along the spine of an individual and from person to person at a given spinal level, increasing the variability beyond that attributable to modality-dependent factors. These problems of image variability combine to make both acquisition and interpretation of the images difficult without substantial operator expertise and, thus, complicate the task of registration with ultrasound.

A small sample of the image variability is shown in Figures 2.3, 2.4, and 2.5. The images shown were acquired from a spinal-fusion patient. The ultrasonic images are displayed with corresponding CT images of the same anatomical regions, which are constructed based on the results of a manual registration. The ultrasonic image in Figure 2.3 was acquired in the sagittal plane centered on the spinous process of L2, the second lumbar vertebra. In this image, scattering from the surface of the spinous process is relatively simple to identify, although identifying precisely which pixels actually lie exactly on the surface of the process is not trivial. Figure 2.4 is another

image acquired in the sagittal plane but at an oblique angle to the cross-section of the vertebral surface shown in the corresponding CT image. Scattering from the surface exists in this image but is undetectable even for an expert operator. In Figure 2.5, the image was again taken in the sagittal plane but approximately normal to the surface of the two articular processes that can be seen in both images.




Figure 2.3: The image on the right is an ultrasonic image of the spine in the sagittal plane with the spinous process centered. The image on the left is a CT image of the same anatomical region. Image height and depth are 5 cm each. The top of the image is posterior; the bottom is anterior.




Figure 2.4: The image on the right is an ultrasonic image of the spine in the sagittal plane with little or no vertebral surface scattering immediately apparent. The image on the left is a CT image of the same anatomical region. The image size and orientation are the same as those in the previous image.




Figure 2.5: The image on the right is an ultrasonic image of the spine in the sagittal plane at the level of the articular processes. The image on the left is a CT image of the same anatomical region. The image size and orientation are the same as those in the previous image.

2.3 Problem Definition: A Vertebra *In Vitro*

Clinical spinal registration is a worthy goal but a substantially difficult one. A complete development of algorithms for spinal registration is beyond the scope of this work, and the intent here is focused as much on an *understanding* of shape in ultrasound as the actual implementation. With this in mind, the application in this dissertation has been limited to registration of a single vertebra from images obtained *in vitro*. This much simpler problem isolated the fundamental issues regarding representation of surface shape by removing additional complexity associated with the *in vivo* environment. The image characteristics of the vertebral surface are quite similar for images obtained *in vivo* or *in vitro*, with the exception of the contrast to the surrounding medium. For example, Figure 2.6 shows a sample *in vitro* image with the associated image plane displayed on a rendering of the vertebra. The image contains the same mix of speckle texture and coherent echoes as clinical images, where echoes are coherent only for regions where the surface normal is orthogonal to the axial image direction.

In vitro vertebral images are used throughout the dissertation. The actual images are used primarily to verify the physical model of image formation (Chapters 4 and 5) through a visual comparison of actual images to images simulated using the model. Further development of the image model employs simulated images primarily. A set of simulated images is used for testing inference algorithms, although the potential for application to actual images is also investigated.




Figure 2.6: Sample *in vitro* image of a cadaveric vertebra. The speckle texture on the left is from the facet joint. Coherent echoes are from the lamina (center) and the inferior articular process (top of the peak on the right).

2.4 Methods: Data Collection

Data collection employed several tools from the image-guided surgery equipment. In addition to the following items, software from the StealthStationTM treatment-guidance platform from Surgical Navigation Technologies was used extensively for analyzing and visualizing clinical and experimental data. The associated treatment-guidance platform was also used for registration of phantom and images. An experimental phantom was constructed as shown on the left in Figure 2.7, containing a cadaveric L4 vertebra. The phantom was scanned with a CT imaging system to produce an image volume from which the vertebra could be segmented, allowing the construction of a triangulated surface. A rendering of that surface is shown on the right of Figure 2.7. The ten aluminum spheres mounted on the outside of the phantom allowed for registration of the physical phantom to the CT images and, thus, the surface.

Ultrasonic images were acquired using an imaging system from the Tetrad Corporation with a model 6C, 128-element, linear array transducer. Focus for the transducer is fixed in the elevation dimension and electronic in the lateral dimension. Relevant specifications for the transducer include a center frequency of 6.0 MHz and elevation focus at 33 mm. Based on information from the manufacturer, beam width is approximately 1.0 mm in the lateral direction throughout most of the image and approximately 3.0 mm in the elevation direction (at a depth of 40 mm). The imaging system used was chosen largely because of its availability through collaborators of the authors. From the standpoint of investigating the proposed models, however, the transducer design and operation are representative of conventional imaging systems.




Figure 2.7: Cadaveric vertebra phantom (left) and rendered surface (right). The phantom contains a cadaveric lumbar (L4) vertebra, constructed to allow registration of CT images, triangulated surface representation, and ultrasonic images. On the right is a rendering (side view) of the triangulated surface representing the vertebra in the phantom of the previous figure.

The ultrasound probe was modified to allow tracking with an optical localization system. The accuracy of using this method to track the position of objects identified in the ultrasonic image was measured to be approximately 2 mm, which is roughly in agreement with other published results [31]. By tracking the probe with the same system used to register the phantom and CT images, the ultrasonic images were acquired in known relation to the surface model, allowing for a direct comparison between simulated and actual images. A sample actual image is shown in the left of Figure 2.6.

Chapter 3

Pattern Theory and Model-Based Image Analysis

3.1 Pattern Theory

Pattern theory is “a way to approach patterns through a mathematical formalism, a way of *reasoning about patterns*” [32]. The work began in the late 1960’s by Grenander and continues today with additional researchers [8, 33]. From [33], the basic objectives of pattern theory are: 1) the creation of mathematical representations intended for representing and understanding patterns both natural and man-made, 2) mathematical analysis, including statistical inference, of the resulting representations, and 3) development of particular applications with implementation of the underlying structures. An underlying tenet is that truly complex structures require complex and detailed representation. This is in contrast to other theories, e.g., fractal theory, which build apparent complexity from simple structures.

Pattern theory is based on an algebra of patterns designed specifically for the aforementioned objectives. Consequently, the application of fields such as geometry, topology and probability theory to the pattern representations provide a deep, rich and robust theory for representation and inference on patterns in widely varying circumstances. A fundamental and powerful notion within pattern theory is the representation of patterns via *deformable templates*. Equivalence within a class of shapes is defined by a similarity group, a group of transformations which deforms any shape in the class to any other shape in the class. A set of equivalent shapes

can then be completely represented by any element of that set, a *template*, and the similarity group. This notion is made formal in [32, 33].

The deformable templates approach is powerful for two significant reasons. First, any knowledge about the class of shapes can be stored in the template and mapped to any other shape via a transformation. Second, inference on the shape of an object can be formulated as estimation of a transformation, a space of much smaller dimensionality than the shape itself for complex structures. Restrictions can be placed on the shapes more naturally through their variation than on the shape, including probabilistic variation. Successes have been made in many applications in diverse fields such as neuroanatomical variability, Automatic Target Recognition (ATR) and language modeling [8, 9, 34, 35, 36, 37, 38].

One very successful application of pattern theory has involved the development of a neuroanatomy atlas [8, 9, 37]. Transformations are constrained deterministically to preserve anatomical topology and probabilistically to prefer variation according to elastic and fluid models of movement. Using these constraints on variation, a template volume of Magnetic Resonance (MR) images is deformed to match an MR volume of an individual. In combination with the deformation, knowledge about anatomical structures in the template, e.g., identification of structures such as the ventricles and major nuclei, is inherent in the deformed volume. Detailed surface representations, which can be well-represented in both form and variation using pattern-theoretic constructs, can be identified automatically for an individual once constructed for the template volume [36]. By deforming several brains, researchers have been able to study the variability of the human and monkey brain and the associated neuroanatomical structures [36, 39].

In Automatic Target Recognition (ATR), researchers have used pattern theory to develop a Bayesian framework for inference on scenes involving variation in target type, pose, and number [8, 35, 40]. Templates consist of detailed models constructed for individual targets such as specific aircraft and tanks which could potentially be part of the scene. Individual target pose is represented by a rigid transformation in two or three dimensions depending on whether the target is land- or air-based. Fundamental within this framework are methods for searching the parameter space which incorporate the geometric properties of the Special Euclidean group, $SE(3)$, the group of rigid transformations [10, 35]. The Bayesian framework also allows observations from multiple sensors to be incorporated naturally.

In another biological application, a pattern-theoretic framework was used to identify mitochondria and membranes in electron micrograph images of cardiac muscle cells [8]. A scene representation similar to that used in ATR was used to represent the number and types of structures. In these images, the interiors of the mitochondria are distinguished by a texture that differs from the rest of the structures in the cells. Markov random fields, a common choice for texture representation, were used to represent the interior of the mitochondria. The Bayesian problem formulation again supports this probability structure naturally.

The relative complexity of representations in pattern theory are both the cost and value of the approach to inference. The advantage over traditional image processing approaches is the framework which combines representation of the underlying structure and observations in the form of images. In a specific application, an approach can be made simpler or more complex while remaining within the framework. Inference problems which are formulated in Bayesian terms can be extended to incorporate probabilistic sensor models, and variations in image representation such as texture and intensity can also be accommodated naturally within the Bayesian formulation. Through a fundamental understanding of the underlying structure and the mechanisms of image formation, the framework allows an approach to inference which can accommodate a broad range of complexity in solutions to a problem.

3.2 Relevant Constructs

3.2.1 Surfaces

The pattern-theoretic, or model-based, approach to recognizing surfaces in ultrasonic images requires two mathematical structures involving the template: 1) a representation for the surface, and 2) a class of transformations that can be applied to the surface. For the scope of this work, the surface representation requires local surface geometry and practical means for computational implementation, while transformations have been limited to rigid-body transformations since the image model is of primary interest. Ultimately, higher-dimensional transformations will be desired, and the triangulated-mesh representation used in pattern theory work [36] satisfies our surface needs without introducing unnecessary computational burden.

Observations of the template are described in the pattern-theoretic approach via a physically-based data model. The data model forms the data likelihood for use

in inference, where patterns are estimated as transformations acting on the template structures. This dissertation focusses primarily on a data model for ultrasonic images, but the context for that model is its application to inference of underlying gross structure. Performance in inference is, thus, the target for all developments and the basis for quantitative analysis.

Representation

Mappings, typically defined [41] as functions from \mathbb{R}^n to \mathbb{R}^m , form the basis for both the surface representation and the transformations that act on the surfaces. Surfaces will be required both for modeling the tissue surface and for modeling interactions with the system point-spread function (PSF). For the PSF, a simple representation will be sufficient where the entire surface can be parametrized by an analytic function, $\mathbf{r}(u, v)$. In this case, the mapping is of the form $\mathbf{r} : A \rightarrow \mathbb{R}^3$ where $(u, v) \in A \subset \mathbb{R}^2$. Tissue surfaces require a more arbitrary structure, and a similar but more general class of surfaces is required, with its definition based on *local* mappings, or patches, defined for all points in the surface.

For the arbitrary surface, the local mappings represent local variations in the shape properties, e.g., various forms of the surface curvature. Such properties are denoted shape because they represent the surface in such a way that it is invariant to rigid transformations. In this work, the important aspect of the surface representation is that it includes local surface characteristics and permits computation. In applying the image model to higher-dimensional transformations, it will be significant because it provides a basis for describing variation of the surface in the deformable templates approach [42]. For a thorough, mathematical description of this differential-geometric representation for surfaces, see, e.g., [36, 41, 43, 44].

Computing with the Triangulated Mesh

Following Joshi, et al, [36], the triangulated mesh has been chosen as the discrete surface representation to use for computing. The representation consists of a set of N_v vertices, $\{\mathbf{v}_i \in \mathbb{R}^3, i \in [1, \dots, N_v]\}$, which are simply points in \mathbb{R}^3 , and a set of N_p triangles, $\{p_i = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3), \mathbf{v}_j \in [1, \dots, N_v], i \in [1, \dots, N_p]\}$, where each triangle connects three vertices. In addition to vertices and triangles, the local surface geometry, up to a quadratic approximation, can be derived from the triangulated mesh [36, 45]. At any vertex, or over any triangle, the normal to the surface can be

estimated quite easily. In this work, the vertices and triangles have been used, as well as the normals. Future work could include application of the curvature, information that is already used in applications of pattern theory to shape analysis [36, 44].

In this work, triangulated mesh representations were generated from CT volumes using an implementation of the Marching Cubes algorithm [46], with the implementation developed by Sarang Joshi and the IntellX Corporation. Spinal CT volumes were segmented by hand, indicating the interior of the desired vertebra. Application of the Marching Cubes algorithm produced the vertices and triangles of the triangulated mesh. A sample rendering of a triangulated surface generated for a cadaveric vertebra is shown in Figure 3.1. Generated surfaces were verified visually using additional software from the IntellX corporation showing contours of the surface overlaid on the original CT images.



Figure 3.1: A rendering of a triangulated surface created using the Marching Cubes algorithm for a cadaveric vertebra.

3.2.2 Rigid transformations

From the classes of transformations that could be applied to the surface in inferring its shape, this investigation has been confined to the relatively simple rigid transformation since the image model is of primary interest. There are many ways to define rigid transformations, see, e.g., [47, 48], but essentially these transformations define distance-preserving motion of an object. For the purposes of this work, it will suffice to consider rigid transformations as the composition of a rotation and a translation. Formally, rigid transformations form the special Euclidean group, $SE(3)$, a mathematical structure offering advantages in certain instances. These advantages are more

relevant to other recognition work that is more theoretical and mathematical in nature. The interested reader is referred to [10, 35] for a detailed description and other references.

Consider first the group of rotation matrices, the set of matrices \mathbf{R} such that $\mathbf{R}^t\mathbf{R} = I$. Any rotation matrix can be conveniently constructed from rotations about the coordinate axes, $\mathbf{R}_x, \mathbf{R}_y, \mathbf{R}_z$, with rotations of θ, ϕ and ψ about the x, y and z axes, respectively,

$$\mathbf{R}_x = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) \\ 0 & -\sin(\theta) & \cos(\theta) \end{bmatrix}, \quad (3.1)$$

$$\mathbf{R}_y = \begin{bmatrix} \cos(\phi) & 0 & \sin(\phi) \\ 0 & 1 & 0 \\ -\sin(\phi) & 0 & \cos(\phi) \end{bmatrix}, \quad (3.2)$$

$$\mathbf{R}_z = \begin{bmatrix} \cos(\psi) & \sin(\psi) & 0 \\ -\sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (3.3)$$

Translations are simply vector addition, i.e., translation by $\mathbf{t} \in \mathbb{R}^3$ is defined as $\mathbf{y} = \mathbf{x} + \mathbf{t}$.

Any rigid transformation can then be represented in terms of a rotation, \mathbf{R} , followed by a translation, \mathbf{t} . The transformation $T = [\mathbf{R}, \mathbf{t}]$ is defined here as a mapping $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ acting on points $\mathbf{x} \in \mathbb{R}^3$, as

$$T(\mathbf{x}) = \mathbf{R}\mathbf{x} + \mathbf{t}. \quad (3.4)$$

Transformation of a *set* of points is defined as the set of transformed points. For example, for a rigid transformation, a transformed surface is defined by a set of vertices resulting from the action of the transformation on the original vertices.

3.2.3 Image Models

Inference with the pattern-theoretic approach requires a probabilistic model for observation data given a transformed version of the template. Observations typically comprise massive amounts of data, e.g., a set of images, thus success in inference depends on an efficient but accurate representation of the data and its dependence

on the transformed template. For any imaging system, probabilistic models allow a natural and convenient representation.

Probabilistic descriptions of image data are multi-dimensional stochastic processes, typically called *random fields* [49]. In this work, we are interested in the data likelihood, i.e., the joint probability density function, $p(x_1, x_2, \dots, x_N|h)$, conditioned on the template transformation, h , over all image samples, x_1, x_2, \dots, x_N . For the large amount of data forming any image, simplification of the joint density is required, thus, in any random field model, some assumptions must be made about neighbor independence. Neighbor properties of the one-dimensional Markov process have been extended to 2D random fields and have received considerable attention because of their relative simplicity in representing local texture [50]. Given the importance of speckle texture in ultrasonic images, Markov random fields could be a wise area to explore. Such application was beyond the scope of this work, however, and in the image model developed in later chapters, neighboring pixels will be assumed independent.

In the neuroanatomy textbook work [9, 37, 42], a simple additive Gaussian noise component was sufficient for modeling variation in MR images. Images of any brain were assumed to vary only by transformation of the underlying anatomy and Gaussian noise. As a result, the likelihood is simple to compute for any transformation of the anatomy. The details can be found in [37], but the relevant result is that these assumptions produce a log likelihood that can be reduced to the square of the difference between the observation image and the transformed template image. Unfortunately, the sensitivity of ultrasonic images to many underlying factors means that a much much more sophisticated image model is required.

3.2.4 Inference and Nonlinear Optimization

Representation of the image data with a probabilistic model permits couching inference of shape as an estimation theory problem. Any of the common approaches (Maximum Likelihood (ML), Maximum a Posteriori (MAP), Minimum Mean-Squared Error (MMSE)) can be taken to achieve various objectives. For instance, an *a priori* probability describing the relative likelihood of various transformations could be incorporated in the MAP approach. For likelihoods that are non-smooth, the MMSE estimate, or conditional mean, may give good estimates. A procedure for finding the MMSE estimate based on stochastic flows on $SE(3)$ is given in [10]. In this dissertation, the emphasis is on the image model rather than techniques for inference,

thus the ML estimate is used because of its simplicity, i.e., no prior is used as in the MAP estimate, and no integration is needed as in the MMSE estimate. The specific inference problem of interest here was to find an ML estimate of the pose given some close initial guess.

For the purposes of investigating the applicability of the image model to inference of shape, the mature field of nonlinear optimization [51, 52] was sufficient. Important issues in optimization involve the differentiability and convexity of the objective function (for an ML estimate, the log likelihood is the objective function). These properties are important because bounds and results can be given under circumstances where differentiability and convexity can be shown explicitly. The complexity of the computations for the image model preclude the development of any theoretical proofs regarding these properties, but practical and computational exploration of the smoothness and convexity have guided the development of the algorithms.

The optimization algorithms of interest are the well-known, derivative-based versions, such as the gradient ascent and quasi-Newton algorithms. In these approaches, an iterative two-step algorithm is followed, where the first step at each iteration involves determining a direction to search, and the second step involves a one-dimensional search along that direction to find a maximum.

In general, the algorithm proceeds as follows for objective function $f(x)$, $x \in \mathbb{R}^n$:

1. Choose an initial guess, x_0 .
2. For iteration k ,
 - (a) Find a search direction, p_k .
 - (b) Search along p_k to maximize $f(x_k + \alpha p_k)$,

$$\alpha_k = \max_{\alpha \in \mathbb{R}} f(x_k + \alpha p_k) \quad (3.5)$$

$$x_{k+1} = x_k + \alpha_k p_k. \quad (3.6)$$

3. Continue until convergence criteria are met.

Direction-finding

In derivative-based algorithms, first- and/or second-order derivative information is used to find a search direction. The algorithms are based on first- or second-order approximations to the objective function under a Taylor series expansion. The gradient-ascent, or steepest-ascent, algorithm uses the gradient as the direction, $p_k = \nabla f(x_k)$, because it is the steepest direction at the current estimate. The approach is relatively simple and straightforward to implement, but it tends to suffer in later stages of the algorithm due to zigzagging and other problems [51, 52]. To improve performance, Newton's algorithms use the Hessian matrix, $H(x_k)$, of second-order partial derivatives, $H_{i,j} = \frac{\partial^2 f(x_k)}{\partial x^{(i)} \partial x^{(j)}}$, to deflect the gradient direction, $p_k = H^{-1}(x_k) \nabla f(x_k)$, improving convergence rates significantly under appropriate conditions such as a quadratic objective function. For practical situations, several problems occur with using the actual Hessian matrix. For instance, inversion of the Hessian can be computationally difficult for problems with a large number of variables. In practical computations, the Hessian can become ill-conditioned, possibly allowing the deflected gradient to produce a decreasing, rather than increasing direction. Such difficulties are exacerbated in problems where significant computation and approximation limit accuracy in computing the objective function.

Several algorithms exist that use an approximation to the inverse of the Hessian, allowing for a computationally efficient approach that maintains the requirement of a positive-definite Hessian. These algorithms are termed quasi-Newton algorithms and build up second order information about the objective function using the gradient calculation at each step. The approximation is built to maintain positive-definiteness, and computational burden is eased since the matrix is built from gradients computed at each iteration. Among the many approximations that exist, one developed independently by Broyden, Fletcher, Goldfarb and Shanno (the BFGS approximation) has been shown in many circumstances, both theoretical and practical, to be the best general choice for nonlinear optimization [51, 52]. The equation for the update offers little insight into the algorithm, thus the interested reader is referred to [51, 52] for further information.

While the algorithms have a sound theoretical basis and have been very useful in many applications, implementation is usually difficult in engineering problems [52], especially those with objective functions requiring extensive computations. For cases when the gradient cannot be computed analytically and must be approximated using

a finite difference approach, special care must be taken in choosing the finite difference interval and form (forward, backward or central difference). Gill, Murray and Wright [52] suggest several practical methods for estimating the optimal interval size and finite difference form and strongly recommend re-evaluation of both at various stages of the algorithm. In the end, an optimization algorithm may require several internal checks and modifications to perform well consistently.

Line search

Given a search direction, a step size is determined in a one-dimensional optimization along the search direction. The accuracy and computational demands of the line-search method can greatly affect the accuracy and convergence of the overall optimization algorithm. Approximate line searches, e.g., Armijo's inexact line search [51, 52], provide nearly optimal solutions at low computational cost. Some algorithms require exact solutions to the line search for various reasons [51, 52], though, despite the increased computational cost. For a survey of line search algorithms see, e.g., [51, 52] with practical implementation covered more thoroughly in [52]. Of practical importance, choices such as a maximum and minimum step size can have a significant impact.

Convergence criteria

Criteria for stopping the algorithm and claiming success are also based in theory with some relatively *ad hoc* practical requirements. Theoretically, the following conditions should hold at a solution, \hat{x} :

1. The gradient is zero, $\nabla f(\hat{x}) = 0$.
2. Change in the objective function is very small, $f(x_k) - f(x_{k-1}) \approx 0$.
3. Step size is very small, $\alpha_k \approx 0$.

In practice, convergence criteria depend on several factors, and these three criteria must be approximations with limits adjusted for the specific problem. The interested reader is referred to [52] for insights into successful practical implementation.

In the end, optimization algorithms, while intuitively straightforward and theoretically attractive for certain cases, demand a moderate amount of fine-tuning for success in a practical setting. The amount of tuning can be expected to be especially high when the objective function requires relatively complicated computations.

Finally, this high sensitivity to details in computation of the objective function motivate a solid theoretical and computational basis for any developed image model to maximize the potential for successful inference.

Chapter 4

Ultrasonic Imaging

Ultrasonic imaging has been and continues to be an intensive area of research. While no single model previously existed for meeting the objectives of this dissertation, much of the research is relevant. In this chapter, relevant research is reviewed in the context of developing a physical model for image formation from the point of the imaging system. A tissue model for surfaces will be derived in the next chapter. The objective of the system model is a model that is based on fundamental principles and is general enough to describe the wide variety of commercial ultrasonic imaging systems. In addition, the model must be flexible enough to allow a range of sophistication in describing both the effects of system and tissue. The model should provide an intuitive, qualitative description of the system operations as well as a mathematical foundation for expansion to a probabilistic image model.

The system model is derived from basic physics in this chapter. A linear model for imaging is derived in Section 4.1. The associated point-spread function (PSF) description of the system is established in section 4.2. These sections are included as background for the model used in this dissertation and as the basis for work that could evolve from this dissertation. For readers uninterested in the mathematics and physics, the intent of the dissertation can still be understood without careful reading of these two sections.

In typical ultrasonic imaging systems, a focused, pulsed wavefield is transmitted from the transducer to the patient, and energy is reflected back to the transducer as shown in Figure 4.1. As ultrasound propagates through soft tissue, various acoustic discontinuities scatter small portions of the incident wavefield. The velocity of propagation ranges from 1450 to 1600 m/s in soft tissue and is slow enough to permit the scattered wavefield to be related to depth, or range, in the medium. For example,

assuming a speed of 1500 m/s, the time required for an ultrasonic pulse to travel from the transducer to a depth of 5 cm (typical for scanning the spine) and back to the transducer is approximately $67 \mu\text{s}$. By transmitting a short-duration ultrasonic pulse and focusing the energy along the direction of propagation, the recorded signal relates directly to a thin column of tissue in the propagation direction. Cross-sectional images are constructed by scanning the tissue region, either mechanically or electronically with an array. At $67 \mu\text{s}$ round-trip travel time, a pulse can be transmitted and received nearly fifteen thousand times in a second without interference between pulses. For an image generated from 128 individual scan lines, over 100 frames per second could be achieved. Typically, no more than 30 frames per second are required, and, in array-based systems, the resulting flexibility is used to improve focus in the image, a concept which will be addressed briefly in the following sections.




Figure 4.1: Commercial ultrasonic imaging systems operate by sending pulses of focused acoustic energy into the patient and recording the scattered energy received at the transducer.

4.1 A Linear Systems Model for Image Formation

The acoustic processes of interest in this work are the transmission, propagation, scattering and reception of acoustic energy. Refraction is also an issue but will not be addressed. The implications of this choice are discussed later in the chapter. Equations describing the phenomena of interest are well-known and have been thoroughly studied in many situations. In imaging, these processes are employed to form a focused image of the acoustic properties of a tissue region. In modern systems, focusing is achieved through both fixed (lenses) and electronic (array beamforming) means.

The system designer is concerned with forming images which optimize certain diagnostically useful characteristics while minimizing system cost and complexity. Precise modeling of the resulting system can be quite complex. Consequently, the imaging model must include tradeoffs between computational tractability and quantitative accuracy.

The physical processes involved can be combined fundamentally into a linear systems model for (most of) image formation, a common approach in both system design and image analysis. In the linear systems model, the system is characterized by its impulse response, or point-spread function (PSF), the system response to a point-like source. The PSF provides a basis for evaluating system performance, including important characteristics such as point resolution, signal-to-noise ratio (SNR), and contrast resolution [53]. The PSF is also essential for analysis of images produced by the system because it provides the basis for distinguishing characteristics of the medium from those of the system.

In both system design and image analysis, quantitative evaluation of the PSF is necessary. Ideally, a system would produce images for which the PSF is shift-invariant, i.e., constant throughout the image. In practice, this goal can only be approximately attained. Highly accurate estimates of the PSF can be found through numerical simulation and are invaluable for validating a design. Such techniques offer little intuition to the design engineer, though, and can be computationally intensive, especially when analyzing the PSF over the entire image. Analytical approximations at varying levels of simplicity are available for guidance in the design process or as a rough approximation in analysis, but these simplifications can be grossly inaccurate at some levels of detail. The value and limitations of these approaches to modeling the PSF are discussed in a later section.

The treatment of image formation given here is taken from a combination of several sources [53, 54, 55, 56, 57, 58]. The treatment includes descriptions of the various physical processes involved in ultrasonic imaging and concludes by combining them in a single equation describing pulse-echo imaging.

4.1.1 Propagation: The Linear Wave Equation in a Homogeneous Medium

Wave propagation in a fluid medium is well described by the linear acoustic wave equation, the second-order partial differential equation describing propagation of acoustic

energy in a medium described by its compressibility, κ , and density, ρ [55]. The equation is a result of combining linear (first-order) equations for conservation of mass (the Equation of Continuity), conservation of momentum (Euler's Equation) and the equation of state which relates changes in pressure and density. The linear wave equation has been found to be accurate in describing many of the fundamental phenomena involved in ultrasonic imaging and forms the basis of most imaging models. The case of a homogeneous medium is discussed first, then the equation is extended to include arbitrary acoustic sources. The two sources of interest are (1) small scattering sources, i.e., discontinuities in the acoustic properties of the medium, and (2) arbitrary transducer sources, which provide the insonifying wavefield.

In a homogeneous medium of compressibility, κ_0 , and density, ρ_0 , the wave equation describing the pressure, $p(\mathbf{r}, t)$, as a function of time and space takes the form [55],

$$\nabla^2 p(\mathbf{r}, t) - \rho_0 \kappa_0 \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} = 0. \quad (4.1)$$

Pressure fields of the form, $p(\mathbf{r}, t) = f(t - \boldsymbol{\alpha} \cdot \mathbf{r})$, where $f(\cdot)$ is any twice-differentiable function (or signal) and $\alpha^2 = \rho_0 \kappa_0$, are easily shown to satisfy the wave equation ($\nabla^2 p(\mathbf{r}, t) = \alpha^2 \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2}$). In this relation, $\boldsymbol{\alpha}$ is termed the slowness vector as in [59] and indicates the velocity of propagation, $c_0 = \frac{1}{|\boldsymbol{\alpha}|}$, and the direction of propagation, $\frac{\boldsymbol{\alpha}}{\alpha}$. Note that this relation expresses the intuitive notion that as time progresses, the signal represented by the wavefield moves, or propagates, in space. Note also that most one-dimensional signals of interest obey the twice-differentiable requirement and can propagate as acoustic energy in space and time as a plane wave (the term plane wave refers to a wavefield where planes in the direction of propagation have constant phase).

Further insight into the nature of propagating wavefields can be gained via the superposition principle. Any sum of solutions to the wave equation is also a solution due to the linearity of the differential operators in 4.1. A medium can, thus, support any sum of waves of the form $f(t - \boldsymbol{\alpha} \cdot \mathbf{r})$, and a wavefield can consist of an arbitrary number of plane waves traveling in different directions. This result can be expressed more simply by considering monochromatic plane waves as complex exponentials, $Ae^{j(\mathbf{k} \cdot \mathbf{r} - \omega t)}$, where ω is the temporal frequency and \mathbf{k} is the wavenumber vector (amplitude is spatial frequency and direction is the direction of propagation). In this form, ω and $k = |\mathbf{k}|$ are related to the velocity of propagation by the linear dispersion

relation $\frac{\omega}{k} = c_0$ (for a homogeneous, non-dispersive medium). The monochromatic plane waves are significant because they constitute the basis for the four-dimensional Fourier transform relationship between the space-time and wavenumber-frequency representations of the pressure wavefield,

$$p(\mathbf{r}, t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P(\mathbf{k}, \omega) e^{j(\mathbf{k} \cdot \mathbf{r} - \omega t)} d\mathbf{k} d\omega. \quad (4.2)$$

The linear dispersion relation confines the wavenumber-frequency spectrum to the cone defined by $\frac{\omega}{k} = c_0$ in agreement with propagation according to the wave equation. The Fourier transform relationship plays a significant role in the following development, allowing many expressions to be simplified by considering them in the frequency domain.

4.1.2 The Wave Equation with Sources

The wave equation for a homogeneous medium gives insight into the nature of propagating wavefields, but an acoustic source must be present to initiate the wavefield. In general, two acoustic sources are of interest in this case, 1) scattering inhomogeneities within the medium which scatter the incident wavefield. and 2) the transducer which is used to insonify the medium. Both situations can be approached with the Green's function method for partial differential equations [60].

Consider a partial differential equation,

$$Lu(x) = f(x), x \in D \subseteq \mathbb{R}^n \quad (4.3)$$

with boundary conditions,

$$Bu(x) = h(x), x \in \partial D \quad (4.4)$$

where L is a linear partial differential operator, e.g., $\nabla^2 - k^2$, defined on a domain, D , $f(x)$ is a driving function, and B is an expression (possibly differential) describing boundary conditions on the boundary of D , denoted ∂D . The Green's function method provides a way to turn the differential equation in $u(x)$ into an integral equation. The Green's function of the operator L is a function of the form, $g(x, y)$,

and is a solution to the equation

$$Lg(x, y) = \delta(x - y). \quad (4.5)$$

The fundamental solution, $g(x, y)$, is the kernel of an integral operator, L^{-1} , that inverts L . Assuming the operator L operates on x only and that the boundary conditions are satisfied, the integral equation for $u(x)$ takes the following convenient form [60]:

$$u(x) = L^{-1}f(x) = \int_D f(y)g(x, y)dy. \quad (4.6)$$

It is easy to see that $u(x)$ in this equation satisfies equation 4.3,

$$Lu(x) = L \int_D f(y)g(x, y)dy = \int_D f(y)Lg(x, y)dy = \int_D f(y)\delta(x - y)dy = f(y). \quad (4.7)$$

The value of the approach is that given the Green's function for a problem, the solution can be found for any driving function, $f(y)$. Note that the integral in equation 4.6 is a superposition integral with impulse response, $g(x, y)$.

The Green's function for a particular problem depends only on the operator L and the boundary conditions. The *general* Green's function, $G(x, y)$, is a sum of $g(x, y)$, the fundamental solution, and $\chi(x, y)$, any solution to the homogeneous equation,

$$G(x, y) = g(x, y) + \chi(x, y). \quad (4.8)$$

It also satisfies equation 4.5 and is chosen in a specific problem to satisfy the boundary conditions. In some cases, e.g., wave propagation, the complete integral equation solution includes other integral terms that depend on the boundary conditions.

The Green's function method is applied to the propagation of acoustic energy by extending the wave equation to include a general driving function, $f(\mathbf{r}, t)$, some distribution of acoustic sources, on the right-hand side of 4.1,

$$\nabla^2 p(\mathbf{r}, t) - \rho_0 \kappa_0 \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} = f(\mathbf{r}, t). \quad (4.9)$$

The solution is easily pursued in the temporal frequency domain. Taking the temporal Fourier transform of equation 4.9 yields a partial differential equation of the form in

equation 4.3 with $L = \nabla^2 - k^2$,

$$\nabla^2 p(\mathbf{r}, \omega) - k^2 p(\mathbf{r}, \omega) = f(\mathbf{r}, \omega). \quad (4.10)$$

The fundamental solution, $g_f(\mathbf{r}, \omega | \mathbf{r}_0)$, (known in acoustic texts as the free-space Green's function), is found by considering equation 4.10 with an impulsive source at \mathbf{r}_0 as the driving function,

$$\nabla^2 g_f(\mathbf{r}, \omega | \mathbf{r}_0) - k^2 g_f(\mathbf{r}, \omega | \mathbf{r}_0) = \delta(\mathbf{r} - \mathbf{r}_0). \quad (4.11)$$

The solution to 4.11 is a spherical pulse-wave traveling outward from \mathbf{r}_0 at velocity c , [55],

$$g_f(\mathbf{r}, \omega | \mathbf{r}_0) = \frac{e^{jkR}}{4\pi R}, \quad R = |\mathbf{r} - \mathbf{r}_0|. \quad (4.12)$$

Note that $g_f(\mathbf{r}, \omega | \mathbf{r}_0)$ is symmetric and spatially invariant, depending only on the distance, R ,

$$g_f(\mathbf{r}, \omega | \mathbf{r}_0) = g_f(\mathbf{r} - \mathbf{r}_0, \omega) = g_f(\mathbf{r}_0 - \mathbf{r}, \omega). \quad (4.13)$$

These properties will be useful in subsequent sections. Again, the general Green's function, $G(\mathbf{r}, \omega)$, is the sum of the free-space Green's function, $g_f(\mathbf{r}, \omega)$, with any other solution to the homogeneous wave equation.

The integral equation solution for the pressure wavefield, $p(\mathbf{r}, \omega)$ in the driven wave equation, 4.10, can be shown as on pages 320 and 321 of [55] to be

$$p(\mathbf{r}, \omega) = \iiint f(\mathbf{r}_0, \omega) G(\mathbf{r} - \mathbf{r}_0, \omega) dv_0 + \iint \left[G(\mathbf{r} - \mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} p(\mathbf{r}_0, \omega) + p(\mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} G(\mathbf{r} - \mathbf{r}_0, \omega) \right] dS_0 \quad (4.14)$$

where $G(\mathbf{r} - \mathbf{r}_0, \omega)$ is the Green's function chosen to suit the boundary conditions for the problem, and $\frac{\partial}{\partial n_0}$ is the partial derivative with respect to the direction normal to the surface. The first term is the superposition integral form of equation 4.6. The second is a surface integral of the boundary values of the pressure wavefield and its normal derivative. Note that it is two superposition integrals, first with respect to the Green's function and then with respect to the normal derivative of the Green's function. For the purposes here of acoustic sources and scattering sites, the

general Green's function will be seen to be shift-invariant, thus these integrals will be convolution integrals. The choice of Green's function and boundary conditions will be made to simplify the formulation in each case.

4.1.3 Low-Level Scattering Sources

Consider first the case of a medium with inhomogeneities in compressibility and density that become sources by scattering the incident wavefield. The medium is described by its density, $\rho_e(\mathbf{r})$, and compressibility, $\kappa_e(\mathbf{r})$, where $\delta\rho(\mathbf{r})$ and $\delta\kappa(\mathbf{r})$ are small perturbations (inhomogeneities) in density and compressibility, respectively,

$$\rho_e(\mathbf{r}) = \begin{cases} \rho_0 + \delta\rho(\mathbf{r}) & \text{in regions of inhomogeneity,} \\ \rho_0 & \text{else.} \end{cases} \quad (4.15)$$

$$\kappa_e(\mathbf{r}) = \begin{cases} \kappa_0 + \delta\kappa(\mathbf{r}) & \text{in regions of inhomogeneity,} \\ \kappa_0 & \text{else.} \end{cases} \quad (4.16)$$

In the presence of inhomogeneities, the density term in the wave equation remains inside the divergence operator,

$$\nabla \cdot \left(\frac{1}{\rho_e} \nabla p(\mathbf{r}, t) \right) - \kappa_e \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} = 0. \quad (4.17)$$

The inhomogeneous wave equation can be rewritten in the form of equation 4.9 by considering the inhomogeneities as variations about the mean values in the medium,

$$\begin{aligned} \gamma_\kappa(\mathbf{r}) &= (\kappa_e(\mathbf{r}) - \kappa_0)/\kappa_0, \\ \gamma_\rho(\mathbf{r}) &= (\rho_e(\mathbf{r}) - \rho_0)/\rho_0. \end{aligned} \quad (4.18)$$

Scattering from these inhomogeneities then becomes the driving function in 4.9 [55, 57],

$$f(\mathbf{r}, t) = \gamma_\kappa(\mathbf{r}) \frac{1}{c^2} \frac{\partial^2 p(\mathbf{r}, t)}{\partial t^2} + \nabla \cdot (\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, t)) \quad (4.19)$$

or in the frequency domain,

$$f(\mathbf{r}, \omega) = -k^2 \gamma_\kappa(\mathbf{r}) p(\mathbf{r}, \omega) + \nabla \cdot (\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega)) \quad (4.20)$$

where $\frac{\omega}{c}$ has been replaced by the wavenumber k .

Neglecting for now the original source of the pressure wavefield, the scattered wavefield can be calculated with the Green's function approach. The medium can be considered unbounded [55], and the integral equation 4.14 reduces to

$$p(\mathbf{r}, \omega) = \iiint f(\mathbf{r}_0, \omega) g_f(\mathbf{r} - \mathbf{r}_0, \omega) dv_0 \quad (4.21)$$

because the surface integral terms go to zero. The relation can be interpreted as the sum of the fields induced by the small regions of inhomogeneity, which act as scattering sources. Note again that the relation is a superposition integral with the free-space Green's function as impulse response.

The solution in 4.21 is intuitive but is computationally difficult for two reasons. First, the scattered pressure is dependent on the total pressure wavefield, consisting of the scattered and incident wavefields. This problem is resolved by making the Born approximation of weak scattering [55], in which the scattered pressure is assumed to be much smaller than the incident wavefield, allowing the total pressure in the right-hand side of 4.21 to be approximated by the incident pressure. The other difficulty in the source distribution is the density term, $\nabla \cdot (\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega))$. Many authors have simplified by neglecting the density terms, although others have found these terms to play a substantial role [57]. In [61], the term was shown to simplify to $-k^2 \gamma_\rho(\mathbf{r})$ for the assumptions of 1) a point-source insonification, 2) a distance from scatterer to sensor which is long relative to a wavelength ($r \gg \frac{1}{k}$), and 3) the backscatter assumption (source of incident wave and sensor of scattered wave are in the same location). Insana [57] adds some insight into the nature of these assumptions, showing that for the long-distance assumption and incident plane wave, the scattering contributions are monopole scattering (spherical waves) from inhomogeneities in density and dipole scattering ($\cos(\theta)$ dependence where θ is the angle between direction of incidence and direction of scattering) from inhomogeneities in compressibility. The backscatter assumption in [61] reduces the $\cos(\theta)$ dependence to minus one.

For a pressure wavefield equal to a superposition of point sources, and making the Born approximation, the backscatter assumption, and observation distance large

relative to a wavelength, the integral for the scattered pressure can be reduced to

$$p_s(\mathbf{r}, \omega) = \iiint \left[k^2 (\gamma_\kappa(\mathbf{r}_0) + \gamma_\rho(\mathbf{r}_0)) \right] p_i(\mathbf{r}_0, \omega) g_f(\mathbf{r} - \mathbf{r}_0, \omega) dv_0 \quad (4.22)$$

where $p_i(\mathbf{r}_0, \omega)$ is the incident pressure wavefield. The compressibility and density terms can be combined to characterize the medium with a single reflectivity function, $q(\mathbf{r}) = \gamma_\kappa(\mathbf{r}) + \gamma_\rho(\mathbf{r})$. Finally, the scattered wavefield due to inhomogeneities is given by

$$p_s(\mathbf{r}, \omega) = \iiint k^2 q(\mathbf{r}_0) p_i(\mathbf{r}_0, \omega) g_f(\mathbf{r} - \mathbf{r}_0, \omega) dv_0. \quad (4.23)$$

This is the direct result obtained from the Green's function approach with the simplified scattering source. Each scattering site acts as an acoustic source with strength determined by the incident field and the value of the inhomogeneity relative to the surrounding medium.

4.1.4 Transducer sources

In imaging, the incident pressure wavefield is generated by a piezoelectric transducer, literally an instrument that converts electrical energy to pressure and vice versa. The analysis of wavefields produced by these transducers borrows from work in both optics and acoustics [55, 58]. For a transducer, the source in 4.14 is distributed along the boundary surface, thus the volume integral vanishes and the pressure wavefield becomes an integral over an arbitrary surface which bounds the medium (the surface must bound the point \mathbf{r}),

$$p(\mathbf{r}, \omega) = \iint \left[G(\mathbf{r}|\mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} p(\mathbf{r}_0, \omega) + p(\mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} G(\mathbf{r}|\mathbf{r}_0, \omega) \right] dS_0. \quad (4.24)$$

Calculation of the pressure from a transducer source using 4.24 requires the knowledge of both the pressure and its normal derivative across the entire surface, a condition which can lead to inconsistencies [58]. Note that the general Green's function is used and is not necessarily assumed to be spatially invariant. By appropriate choice of the Green's function, however, equation 4.24 can be reduced to depend on either the pressure or its normal derivative [53, 58], resulting in the Rayleigh-Sommerfeld

equations, both of which can be formulated solely in terms of the free-space Green's function,

$$p(\mathbf{r}, \omega) = \iint 2g_f(\mathbf{r} - \mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} p(\mathbf{r}_0, \omega) dS_0, \quad (4.25)$$

$$p(\mathbf{r}, \omega) = \iint 2p(\mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} g_f(\mathbf{r} - \mathbf{r}_0, \omega) dS_0. \quad (4.26)$$

The factor of two is a result of the choice of Green's function in each case [58]. With either of the two equations, calculating the pressure based on the boundary conditions at the transducer surface amounts to making various assumptions about the behavior of the pressure and/or its normal derivative outside the transducer surface [62]. In general, either equation can be used with the values on a surface away from the transducer, obtained by theoretical or experimental means [53]. Equation 4.25 will be used here since it can be written as a convolution with the free-space Green's function,

$$p(\mathbf{r}, \omega) = 2 \frac{\partial}{\partial z} p_0(\mathbf{r}, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \quad (4.27)$$

where the value of the wavefield generated by the transducer is assumed to be known in the x, y plane and is represented by $p_0(\mathbf{r}, \omega)$, (dS_0 becomes $dx dy$ and permits writing the equation in terms of a convolution in x and y).

In summary, the processes of transmission, propagation and scattering can be written based on the free-space Green's function. Various assumptions are made in modeling the processes in this manner, most notably the Born approximation or weak scattering assumption, observation distance much larger than a wavelength, and the backscatter assumption. Attenuation has not been explicitly included, although the velocity term in the wave equation can be made complex to account for attenuation in tissue without changing the form of the equations [53].

4.1.5 A Pulse-Echo Equation

The results of the preceding sections can be combined to form a single equation describing the received signal for a pulsed transducer. The development here follows that of Wright [53] and is for a single focal point. In cases of interest, the incident wavefield can be modeled as in equation 4.27. In general, the transmitted pressure wavefield is produced by pulsing the transducer elements with electrical pulses that

can be varied in amplitude and time for focusing. The resulting wavefield can be represented as follows, with a transfer function, $H_t(\omega)$, representing the acoustic version of the the electrical pulse, and $w_t(x, y, \omega)$, an arbitrary weighting function representing spatial and temporal variation of the pulses over the transducer or transducer elements,

$$p_t(\mathbf{r}, \omega) = H_t(\omega)w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega). \quad (4.28)$$

The scattered wavefield is given by equation 4.23, expressed here as a convolution,

$$p_s(\mathbf{r}, \omega) = k^2q(\mathbf{r})p_i(\mathbf{r}, \omega) \underset{x,y,z}{*} g(\mathbf{r}, \omega). \quad (4.29)$$

In this equation, the incident wavefield, $p_i(\mathbf{r}, \omega)$, is generated by the transducer and given by 4.28, resulting in the following wavefield at the plane of the transducer ($z = 0$),

$$p(\mathbf{r}, \omega)|_{z=0} = \left[k^2q(\mathbf{r})p_t(\mathbf{r}, \omega) \underset{x,y,z}{*} g(\mathbf{r}, \omega) \right] \Big|_{z=0}. \quad (4.30)$$

The received signal, $R(\omega)$, is then generated from signals across the transducer elements with delays and amplitudes chosen similarly as in the transmit case to achieve a specific focus. This processing is represented using a receive transfer function, $H_r(\omega)$, for the conversion of pressure to electrical energy, and weighted as in the transmit case with a weighting function, $w_r(x, y, \omega)$,

$$R(\omega) = H_r(\omega) \iint dx dy w_r(x, y, \omega) p(\mathbf{r}, \omega)|_{z=0}. \quad (4.31)$$

The received signal equation can be simplified as follows. Substituting the previous expressions for pressure into the received signal equation,

$$R(\omega) = H_r(\omega) \iint dx dy w_r(x, y, \omega) \left[\left[k^2q(\mathbf{r})H_t(\omega)w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \underset{x,y,z}{*} g(\mathbf{r}, \omega) \right] \Big|_{z=0}. \quad (4.32)$$

Combining the receive and transmit transfer functions with the k^2 term as $H_{rt}(\omega) = k^2 H_r(\omega) H_t(\omega)$,

$$R(\omega) = H_{rt}(\omega) \iint dx dy w_r(x, y, \omega) \left[\left[q(\mathbf{r}) w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \underset{x,y,z}{*} g(\mathbf{r}, \omega) \right] \Big|_{z=0}. \quad (4.33)$$

The convolution in z and evaluation at $z = 0$ can be reduced to a single integral over z as follows. Operating on the bracketed expression in 4.33,

$$\begin{aligned} & \left[\left[q(\mathbf{r}) w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \underset{x,y,z}{*} g(\mathbf{r}, \omega) \right] \Big|_{z=0} \\ &= \int dz' \left[q(x, y, z') w_t(x, y, \omega) \underset{x,y}{*} g(x, y, z', \omega) \right] \underset{x,y}{*} g(x, y, -z', \omega) \end{aligned} \quad (4.34)$$

where z' is a dummy variable in the convolution. The free-space Green's function is symmetric in x, y , and z , thus $g(x, y, -z', \omega) = g(x, y, z', \omega)$, resulting in the following expression,

$$= \int dz \left[q(\mathbf{r}) w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \underset{x,y}{*} g(\mathbf{r}, \omega). \quad (4.35)$$

This expression can be substituted into equation 4.33, yielding

$$R(\omega) = H_{rt}(\omega) \int d\mathbf{r} w_r(x, y, \omega) \left[q(\mathbf{r}) w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \underset{x,y}{*} g(\mathbf{r}, \omega). \quad (4.36)$$

The receive weighting function, $w_r(x, y, \omega)$ can be grouped with the second convolution using the following relation, which holds for symmetric functions $h(x)$,

$$\int dx f(x) [g(x) * h(x)] = \int dx f(x) \int dx' g(x') h(x - x') \quad (4.37)$$

$$\stackrel{h(x)=h(-x)}{=} \int dx' g(x') \int dx f(x) h(x' - x) \quad (4.38)$$

$$= \int dx' g(x') [f(x') * h(x')]. \quad (4.39)$$

Using this result in equation 4.36,

$$R(\omega) = H_{rt}(\omega) \int d\mathbf{r} q(\mathbf{r}) \left[w_r(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \left[w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right]. \quad (4.40)$$

This equation can be stated more simply as a superposition integral,

$$R(\omega) = \int d\mathbf{r} q(\mathbf{r}) s(\mathbf{r}, \omega) \quad (4.41)$$

where $s(\mathbf{r}, \omega)$ is the system response to a spatial impulse at \mathbf{r} ,

$$s(\mathbf{r}, \omega) = H_{rt}(\omega) \left[w_t(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \left[w_r(x, y, \omega) \underset{x,y}{*} g(\mathbf{r}, \omega) \right] \quad (4.42)$$

and contains terms associated with the transmitted wavefield, the filtered version of the received wavefield and the roundtrip pulse.

Equation 4.40 is quite general and can be used to analyze many complex imaging situations. To illustrate how this equation can describe imaging, consider a simple example consisting of a point-source at the origin, a point-receiver at the origin, and a spatial impulse at \mathbf{r}_0 ,

$$\begin{aligned} H_{rt}(\omega) &= 1, \\ w_t(x, y, \omega) &= \delta(x, y), \\ w_r(x, y, \omega) &= \delta(x, y), \\ q(\mathbf{r}) &= A\delta(\mathbf{r} - \mathbf{r}_0). \end{aligned} \quad (4.43)$$

This gives a received signal,

$$\begin{aligned} R(\omega) &= \iiint A\delta(\mathbf{r} - \mathbf{r}_0) \frac{1}{4\pi r} e^{jkr} \frac{1}{4\pi r} e^{jkr} dx dy dz \\ &= \left(\frac{1}{4\pi r_0} \right)^2 A e^{j2kr_0} \\ r(t) &= \left(\frac{1}{4\pi r_0} \right)^2 A \delta(t - 2r_0/c). \end{aligned} \quad (4.44)$$

This is exactly the signal one should expect, an impulse delayed by the time to travel from the origin to \mathbf{r}_0 and back and scaled by the square of the amplitude spreading factor for a spherical wave.

Consider now an additional factor in the receive weighting function, $w_r(x, y, \omega)$, of $(4\pi r_0)^2 e^{-j2kr_0}$,

$$w_r(x, y, \omega) = (4\pi r_0)^2 e^{-j2kr_0} \delta(x, y). \quad (4.45)$$

When multiplied by the receive signal, this function eliminates the frequency dependence and the amplitude factor, leaving the scattering potential at \mathbf{r}_0 at time $t = 0$. This example illustrates a basic approach to imaging. The weighting function is designed to produce at $t = 0$ an estimate of the medium reflectivity, effectively eliminating time from the calculations, allowing a description of image formation dependent only on the spatial variables.

In any medium of interest, the reflectivity function will consist not of a single spatial impulse but of a complex and intricate structure. In this case, the transducer and weighting functions are designed to provide received signals that can be filtered to obtain estimates of the reflectivity function over a compact region. The general pulse-echo equation, 4.40, can be used to analyze many approaches to image formation, from transducer design to array-based beamforming. The separation of the system response into the transmitted wavefield, the received wavefield, and the round-trip pulse allows the role of each to be analyzed individually. In the next section, the use of this approach in the design and analysis of imaging systems is illustrated.

4.2 Design and Analysis of Imaging Systems

In an ultrasonic imaging system, “ideal” varies depending on the application, but most design issues are common among all systems. In the best case, an imaging system would produce an image which is spatially invariant, i.e., an object would appear the same regardless of where it is in the image. In practice, this goal can be met well locally but only to some degree globally. The system point-spread function (PSF), the response of the system to a point-like scatterer, can be used to characterize the system. Specifically, the PSF can be analyzed to determine the system point resolution, or the spacing between distinguishable targets, and its contrast resolution, or the difference in echogenicity required to distinguish adjacent regions. Another important measure of the system performance is the temporal resolution, or the rate at which the system forms images (independent images).

Many tradeoffs are made in designing a system to achieve the goals of spatial invariance and high point, contrast and temporal resolution. Focusing is typically achieved using both fixed lenses and electronic beamforming. For a given transducer geometry and beamforming process, the PSF can be simulated and measured experimentally. The system design engineer must use approximate but intuitive methods

based on geometric optics for initial design, then refine and test the design with accurate numerical methods and measurement.

In addition to characterizing system performance, the system PSF is used in tissue characterization work to provide a means for distinguishing system from tissue effects and for exploring the relation between them. In the pursuit of a suitable model for ultrasonic images based on shape, computational complexity must be balanced with accuracy. The same methods used for analyzing the PSF to determine system performance, e.g., measurement, approximation, and simulation, can be used to construct a PSF for use in modeling system response to tissue.

The model developed in the previous section is general enough to permit analysis of the PSF for imaging in a variety of scenarios from fixed focusing to electronic focusing with arrays. In beamforming, signals received by the array elements are weighted, delayed and summed to focus the beam. Many options are available for choosing these weights and delays. In [59], many general-purpose array beamforming techniques are described. A nice derivation of some options specifically for ultrasonic imaging can be found in [53].

In general, the weighting function for focus at \mathbf{r}_0 has the form of an amplitude term, $a(x - x_0, y - y_0, \omega)$, and a delay term, $e^{jk|\mathbf{r}-\mathbf{r}_0|}$,

$$w(x, y, \omega) = a(x - x_0, y - y_0, \omega)e^{jk|\mathbf{r}-\mathbf{r}_0|} \quad (4.46)$$

where the amplitude term compensates for loss such as propagation loss and the delay term compensates for the propagation time. An image pixel is then formed by sampling the received signal at $t = 0$ as in the earlier example. This process is equivalent to integrating over frequency, resulting in the following equation for $i_{rf}(\mathbf{r}_0)$, the RF image intensity at \mathbf{r}_0 ,

$$i_{rf}(\mathbf{r}_0) = \int d\mathbf{r} q(\mathbf{r}) \int d\omega H_{rt}(\omega) \left[w_r(x, y, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right] \left[w_t(x, y, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right] \quad (4.47)$$

where, again, $q(\mathbf{r})$ is the reflectivity at \mathbf{r} , $H_{rt}(\omega)$ is the round-trip transfer function, $w_r(x, y, \omega)$ is the receive weighting function, $w_t(x, y, \omega)$ is the transmit weighting function, and $g_f(\mathbf{r}, \omega)$ is the free-space Green's function. Written as a superposition

integral,

$$i_{rf}(\mathbf{r}_0) = \int d\mathbf{r} q(\mathbf{r})h(\mathbf{r}; \mathbf{r}_0) \quad (4.48)$$

where $h(\mathbf{r}; \mathbf{r}_0)$ is the system PSF for a focus at \mathbf{r}_0 ,

$$h(\mathbf{r}; \mathbf{r}_0) = \int d\omega H_{rt}(\omega) \left[w_r(x, y, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right] \left[w_t(x, y, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right]. \quad (4.49)$$

In modeling the PSF, the difficult calculation is for the transmitted and received field terms. Such modeling has been of interest for quite some time in both acoustics and optics, and many methods have been developed, ranging from numerical solutions [58, 62, 63, 64] to the commonly used Fresnel and Fraunhofer approximations [53, 58]. Fundamentally, the required calculation is the surface integral for transducer sources, equation 4.24, listed again here,

$$p(\mathbf{r}, \omega) = \iint \left[G(\mathbf{r}|\mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} p(\mathbf{r}_0, \omega) + p(\mathbf{r}_0, \omega) \frac{\partial}{\partial n_0} G(\mathbf{r}|\mathbf{r}_0, \omega) \right] dS_0 \quad (4.50)$$

where the wavefield is desired for a given transducer geometry, beamforming approach, and transducer vibration (driving function). Recall that either of the Rayleigh-Sommerfeld equations in 4.24 can be used to calculate this expression. In the formulation used here, i.e., in equation 4.49, the weighting function, $w(x, y, \omega)$, represents the normal derivative of the pressure at the surface of the transducer.

4.2.1 Numerical Methods for PSF Modeling

Liu and Waag [62] present a review of the numerical techniques available for modeling transmitted wavefields. The most widely-used methods are the impulse response method [63, 64] and the angular spectrum method [53, 58]. Each uses one of the Rayleigh-Sommerfeld versions of equation 4.24. The two differ in assumptions about the nature of the boundary conditions, specifically whether the pressure or its normal derivative are zero outside the transducer surface. The difference in results between the two is typically small, though [62], and negligible in this work.

These methods are well-developed and could be implemented if necessary. The main advantage of the numerical methods over the approximations of the next section

is the accuracy of the results. This accuracy is crucial for the field of quantitative ultrasonic imaging [62] but has limited applicability to this work at this stage. The disadvantages of the numerical methods are the computation required and the lack of intuition provided. Small changes to the imaging system, e.g., the beamforming steps that change several times throughout the image, require a completely new simulation, and results can be viewed but are not easily simplified. The approximate methods provide a quantitative model that is valid under certain assumptions that are rarely met completely. Their power, though, is the simplicity of the results, which provide useful relationships between the transducer geometry and the PSF.

4.2.2 Approximate Methods for PSF Modeling

Recall the relation for the system point-spread function for a focus at \mathbf{r}_0 ,

$$h(\mathbf{r}; \mathbf{r}_0) = \int d\omega H_{rt}(\omega) \left[w_r(x, y, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right] \left[w_t(x, y, \omega) \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right] \quad (4.51)$$

where $H_{rt}(\omega)$ is the round-trip transfer function, $w_r(x, y, \omega)$ is the receive weighting function, $w_t(x, y, \omega)$ is the transmit weighting function, and $g_f(\mathbf{r}, \omega)$ is the free-space Green's function. As stated previously, the general weighting function for focus at \mathbf{r}_0 has the form of an amplitude term, $a(x - x_0, y - y_0, \omega)$, and a delay term, $e^{jk|\mathbf{r}-\mathbf{r}_0|}$. With this substitution, the PSF becomes

$$h(\mathbf{r}; \mathbf{r}_0) = \int d\omega H_{rt}(\omega) \left[a_r(\mathbf{r} - \mathbf{r}_0, \omega) e^{jk|\mathbf{r}-\mathbf{r}_0|} \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right] \left[a_t(\mathbf{r} - \mathbf{r}_0, \omega) e^{jk|\mathbf{r}-\mathbf{r}_0|} \underset{x,y}{*} g_f(\mathbf{r}, \omega) \right]. \quad (4.52)$$

The bracketed expressions for the transmitted and received wavefields are, once again, the difficult calculation. Consider the expression for either field, writing out the convolution and substituting the expression for the free-space Green's function,

$$[\dots] = a(\mathbf{r} - \mathbf{r}_0, \omega) e^{jk|\mathbf{r}-\mathbf{r}_0|} \underset{x,y}{*} g_f(\mathbf{r}, \omega) \quad (4.53)$$

$$= \iint dx' dy' a(\mathbf{r}' - \mathbf{r}_0, \omega) e^{jk|\mathbf{r}'-\mathbf{r}_0|} \frac{e^{-jk|\mathbf{r}'-\mathbf{r}|}}{4\pi|\mathbf{r}' - \mathbf{r}|} \quad (4.54)$$

$$= \iint dx' dy' a(\mathbf{r}' - \mathbf{r}_0, \omega) \frac{e^{jk(|\mathbf{r}'-\mathbf{r}_0| - |\mathbf{r}'-\mathbf{r}|)}}{4\pi|\mathbf{r}' - \mathbf{r}|}. \quad (4.55)$$

This equation relates the pressure field to the amplitude and frequency weighting function, $a(\mathbf{r} - \mathbf{r}_0, \omega)$. With a simple approximation to the integral, a rough estimate of the field can be found easily for any transducer geometry. First, it will be helpful to make the substitution $\mathbf{r} = \mathbf{r}_0 - \delta$ where $\delta = (\delta_x, \delta_y, \delta_z)$ is much smaller than \mathbf{r}_0 .

$$[\dots] = \iint dx' dy' a(\mathbf{r}' - \mathbf{r}_0, \omega) \frac{e^{jk(|\mathbf{r}' - \mathbf{r}_0| - |\mathbf{r}' - (\mathbf{r}_0 - \delta)|)}}{4\pi|\mathbf{r}' - (\mathbf{r}_0 - \delta)|}. \quad (4.56)$$

Equation 4.56 is quite similar to those found in Goodman [58] and Macovski [54] for the transmitted field for an arbitrary planar aperture, with a slight difference due to the focusing delays included in this equation. The terms requiring approximation are the vector distances. A close look at the equation reveals that the result is much more sensitive to terms in the exponent than the denominator. The standard Fraunhofer approximation [54, 58] approximates the exponent using a binomial expansion; the same result can be achieved using a Taylor series expansion in δ [53].

The interested reader is referred to the texts [53, 54, 58] for details of the expansion in those cases. For the notation and formulation used here, the distance terms in the exponent, $|\mathbf{r}' - \mathbf{r}_0|$ and $|\mathbf{r}' - (\mathbf{r}_0 - \delta)|$, can be approximated as

$$|\mathbf{r}' - \mathbf{r}_0| - |\mathbf{r}' - (\mathbf{r}_0 - \delta)| \approx \frac{1}{z_0}(-\delta_x x' - \delta_y y' - \delta_z z_0). \quad (4.57)$$

This approximation can be found by a first-order Taylor series expansion in δ , followed by the assumption that $|\mathbf{r}_0| \approx z_0$ and $\frac{k_0}{z_0}(\delta_x x_0 + \delta_y y_0) < 1$ (resulting in a phase contribution of less than 1 radian). The same result can be found by a binomial expansion similar to that used in [54, 58], followed by removal of quadratic terms and the assumption again that $\frac{k_0}{z_0}(\delta_x x_0 + \delta_y y_0) < 1$. Note that both x_0 and y_0 can be chosen to be 0 according to choice of the coordinate system, rendering the assumption trivial.

Using the Fraunhofer approximation for the exponent and approximating $|\mathbf{r}' - (\mathbf{r}_0 - \delta)| \approx R_0 = |\mathbf{r}_0|$ in the denominator, the PSF takes the following form,

$$h(\mathbf{r}; \mathbf{r}_0) \approx \frac{1}{(4\pi R_0)^2} \int d\omega H_{rt}(\omega) e^{-j\omega(\frac{2\delta_z}{c_0})} \iint dx' dy' a_r(\mathbf{r}', \omega) e^{-j2\pi[x' \frac{\delta_x}{\lambda_{z_0}} + y' \frac{\delta_y}{\lambda_{z_0}}]} \iint dx'' dy'' a_t(\mathbf{r}'', \omega) e^{-j2\pi[x'' \frac{\delta_x}{\lambda_{z_0}} + y'' \frac{\delta_y}{\lambda_{z_0}}]} \quad (4.58)$$

where $\lambda = \frac{f}{c_0}$ is the wavelength. The weighting term has been represented as a function of \mathbf{r} instead of $\mathbf{r} - \mathbf{r}_0$ for convenience; recall that this relation and the resulting PSF are for a fixed focus at \mathbf{r}_0 . The main result requires the further assumption that the signal be narrowband, allowing the spatial integral to be written for a center frequency, f_0 , and separated from the temporal frequency integral. The frequency dependence of the amplitude weighting terms is assumed to be such that it can be absorbed into the round-trip transfer function. The relation can then be written as

$$h(\mathbf{r}; \mathbf{r}_0) \approx \frac{1}{(4\pi R_0)^2} \left[\int d\omega H_{rt}(\omega) e^{-j\omega(\frac{2\delta_z}{c_0})} \right] \left[\iint dx' dy' a_r(\mathbf{r}') e^{-j2\pi[x' \frac{\delta_x}{\lambda_0 z_0} + y' \frac{\delta_y}{\lambda_0 z_0}]} \right] \left[\iint dx'' dy'' a_t(\mathbf{r}'') e^{-j2\pi[x'' \frac{\delta_x}{\lambda_0 z_0} + y'' \frac{\delta_y}{\lambda_0 z_0}]} \right]. \quad (4.59)$$

The utility of this relation comes from the Fourier transform relations that result in each of the bracketed expressions. In this case, the PSF can be written as the product of the round-trip axial impulse response (in terms of $\frac{2\delta_z}{c_0}$, accounting for round-trip travel time), and transmit and receive fields expressed as Fourier transforms of the respective aperture functions. The following Fourier transform relationships express the correspondence between the beam response in the subsequent equation and the round-trip transfer function and aperture functions.

$$\begin{aligned} h_{rt}(t) &\xleftrightarrow{\mathcal{F}} H_{rt}(\omega) \\ a_r(x, y) &\xleftrightarrow{\mathcal{F}} A_r(u, v) \\ a_t(x, y) &\xleftrightarrow{\mathcal{F}} A_t(u, v) \end{aligned} \quad (4.60)$$

$$h(\mathbf{r}; \mathbf{r}_0) \approx h_{rt} \left(\frac{2\delta_z}{c_0} \right) A_r \left(\frac{\delta_x}{\lambda_0 z_0}, \frac{\delta_y}{\lambda_0 z_0} \right) A_t \left(\frac{\delta_x}{\lambda_0 z_0}, \frac{\delta_y}{\lambda_0 z_0} \right) \quad (4.61)$$

An important aspect of this result is the separability of the axial component of the PSF from the lateral and elevation components. The first term, the round-trip axial impulse response, is often referred to as the axial pulse. It is typically represented as an envelope modulating a carrier term at the center frequency. This representation will be useful in the subsequent statistical analysis.

The response in the lateral and elevation plane is the product of the receive and transmit responses, each of which is approximated as the Fourier transform of the aperture function. A useful result is that a basic relation can be determined for beam width in terms of the transducer geometry. For a separable aperture, the transform

is also separable. Consider a rectangular aperture of width (D_x, D_y) , the transform of which is well known as a 2D sinc function. The width between the first zeros of the sinc function, which is often used to approximate the width of the beam, is also well known [58] as $(\frac{\lambda_0 z_0}{D_x}, \frac{\lambda_0 z_0}{D_y})$.

Note that the beam width is inversely proportional to aperture width and proportional to the wavelength and range. In optics and ultrasonics, the f-number, defined as the ratio of range to aperture width, is often used to characterize a system. The utility of the f-number is that it gives a simple expression for the resolution as the product of f-number and wavelength. In array-based imaging systems, the goal of image uniformity is often sought by adjusting beamforming delays and amplitudes to achieve a nearly constant f-number throughout the image.

The Fraunhofer approximation, while useful for obtaining approximate results is rarely accurate for describing the exact shape of the PSF in ultrasonic imaging. This is in contrast to optics, where the narrowband assumption can be met quite well, along with the distance assumptions necessary for the approximations to be valid. In ultrasonic imaging, axial resolution is inversely proportional to the bandwidth, thus wideband signals are desirable. Approximations can be accurate at the center frequency, though, and the basic results for beam width are universally used as a first approximation in system design and analysis.

Other advantages are gained by making the Fraunhofer approximation and narrowband assumption in terms of achieving a simplified representation of the image formation process. Again, the approximations allow the axial component of the PSF to be separated from the lateral and elevation components. Because the axial component is a bandpass signal, it can be represented as the real component of an analytic signal [65], the product of a complex envelope term, $\mathbf{A}(\delta_z)$, and a complex sinusoid, e^{j2k_0z} , at the wavenumber, k_0 , corresponding to the center frequency, f_0 , of the transducer.

$$h_{rt}(\delta_z) = \text{Re}\{\mathbf{A}(\delta_z)e^{j2k_0z}\}. \quad (4.62)$$

The notation can be extended to the 3D PSF using a 3D envelope, $\mathbf{A}(\delta)$,

$$h(\mathbf{r}; \mathbf{r}_0) = h(\delta) = \text{Re}\{\mathbf{A}(\delta)e^{j2k_0\delta_z}\}. \quad (4.63)$$

A typical choice in modeling the point-spread function is to assume a real envelope with Gaussian curves in each dimension, where the standard deviations, $\sigma_x, \sigma_y, \sigma_z$,

represent the widths of the beam,

$$A(\delta) = e^{-\left(\frac{\delta_x^2}{\sigma_x^2} + \frac{\delta_y^2}{\sigma_y^2} + \frac{\delta_z^2}{\sigma_z^2}\right)}. \quad (4.64)$$

The analytic signal representation for the PSF allows the image formation equation, equation 4.48, also to be written in analytic signal notation. The image formation equation represents the RF image as an integral of the tissue reflectivity, $q(\mathbf{r})$ and the system PSF, $h(\mathbf{r}; \mathbf{r}_0)$, which can be rewritten in analytic signal notation,

$$\begin{aligned} i_{rf}(\mathbf{r}) &= \int d\mathbf{r}_0 q(\mathbf{r}_0) h(\mathbf{r}; \mathbf{r}_0) \\ &= \text{Re} \left\{ \left[\int d\mathbf{r}_0 q(\mathbf{r}_0) A(\delta) e^{-j2k_0 z_0} \right] e^{j2k_0 z} \right\} \\ &= \text{Re} \left\{ \mathbf{i}(\mathbf{r}) e^{j2k_0 z} \right\} \end{aligned}$$

where $\mathbf{i}(\mathbf{r}) = |\mathbf{i}(\mathbf{r})|e^{j\angle\mathbf{i}(\mathbf{r})}$ is a complex envelope representation for the RF image. The complex envelope can be considered as a spatially varying phasor with varying amplitude and phase. In conventional imaging systems, the envelope of the RF image is displayed, which effectively represents the amplitude of the complex envelope. This analytic signal expression is also significant for a second reason. In relating the image intensity to the underlying tissue structure in terms of the point-spread function, the integral representing the complex envelope can be considered as a sum of individual phasors. Each individual phasor can represent a scattering element with a given amplitude and phase. In probabilistic models, the amplitude and phase of each scatterer can be considered random, along with the number of scatterers affecting the intensity at any location \mathbf{r}_0 . The resulting random phasor sum forms the basis of many probabilistic models used in ultrasound and other scattering models.

The development in the preceding sections establishes the foundation for building the models that will be used in the rest of this dissertation. The development is general enough to describe the many possible approaches to system design and operation. Accuracy can be varied by removing or adding approximations in a trade-off of computational complexity for precision. Insights available from results of the Fraunhofer approximation will be used in the following section for a qualitative guide to the interpretation of typical medical ultrasonic images.

4.3 Interpretation of Medical Ultrasonic Images

The high variability of ultrasonic images is due to the anisotropic nature of scattering and variation of the acoustic properties of tissue at the microstructural level. The sensitivity of that relationship results in image features that perplex the untrained observer but can also be useful in visualizing soft tissue structure and detecting abnormalities. Expertise in interpreting the images requires substantial training, though. Each tissue region requires specific training regarding typical image characteristics unique to that anatomical structure. For an excellent and thorough introduction to qualitative interpretation of medical ultrasonic images, see the second chapter of [66]. The most basic image variation, and the aspect primarily used for interpretation, is the texture and relative intensities representing different regions. The relation of these basic image characteristics to the underlying tissue structure is understandable from the simple approximation to the PSF developed at the end of the previous section.

A simulated axial response representing that from a typical PSF is shown in Figure 4.2. Each scattering element in the tissue responds with a scaled and delayed version of that pulse. The separation distance of two elements determines the nature of the combined response. For a distance much smaller than a wavelength, the elements will add coherently, producing a strong echo at the transducer. As the distance increases, however, the pulses add incoherently, resulting in an incoherent response. This phenomenon is the source of image speckle, the texture commonly seen in ultrasonic images. Tissue is typically complex, with a highly varied acoustic structure at the sub-wavelength level, producing a multitude of possible image appearances.




Figure 4.2: A typical axial pulse consists of a few cycles of a sinusoid at a frequency between 1 and 10 MHz. For typical speeds of sound in tissue (and a center frequency of 6MHz), this corresponds to the indicated spatial scale.

Image intensity and texture, thus, vary significantly depending on the spacing and strength of scattering sources in the tissue. Variation in the regularity of spacing

produces echoes ranging from high and constant intensity at a smooth interface to speckled texture from a random scattering medium. Interface echoes range in intensity according to acoustic properties of the interface and the degree of variation in the structure. Speckled textures range in intensity according to the strength, spacing, and concentration of scattering elements in a region.

Consider again the spinal ultrasonic image of Chapter 1, shown again here in Figure 4.3. The diagram on the right indicates the location of basic tissue structures in the region. The skin and fascia are both distinguished by fairly homogeneous textures, with scattering from the fascial layer more intense than that from the skin. The muscle layer between the fascia and transverse process is represented primarily from scattering that comes from the fibers running parallel to the image plane. Scattering from the fibers varies in intensity depending on orientation to the beam and also includes some speckled texture. The bony transverse process of a lumbar vertebra is marked by intense scattering from the surface of the process. The scattering varies from a strong, coherent echo at the top of the process where the surface is normal to the beam, to a speckled texture along the sides of the process where microscopic roughness is distributed along the bone. The acoustic properties of bone differ significantly from that of surrounding tissue, producing high scattering and also strong attenuation of the beam. Note the shadow below the transverse process caused by nearly complete attenuation of the signal.

The variations in texture and intensity in Figure 4.3 are indicative of the range of characteristics typical of ultrasonic images. In general, though, characteristics are unique to each tissue region, with an underlying basis in the interaction between the system PSF and the acoustic properties of the tissue microstructure. The high sensitivity of this interaction results in substantial variation with changes in the system or the tissue microstructure.

A description of qualitative interpretation of ultrasonic images would be incomplete without mention of artifacts. Tissue can appear much different from behavior predicted by the models such as the one used here. The most significant example is attenuation in tissue, which is typically assumed to be exponential with depth, allowing compensation with a depth-dependent gain. When attenuation deviates from the assumed value, however, as with the bone in the previous figure, changes in image brightness occur that do not reflect structural variation but can be mistaken as such [66]. These variations are not addressed in this model and are outside the scope of this work.




Figure 4.3: The spinal ultrasonic image on the left shows typical characteristics of several tissue regions, skin, fascia, muscle and bone. The drawing on the right shows the locations of selected tissue regions.

4.4 Tissue Models

The extreme sensitivity to interactions between the tissue structure and imaging system is such that no single model exists for all tissue structure, and it appears unlikely that any such model will ever be developed [57]. Various tissue models have been proposed and evaluated, with results valid for only specific situations. Tissue representations are universally concerned with the microstructural level of detail and can generally be categorized according to a continuous or discrete representation of the underlying inhomogeneities [57]. The discrete representation is of considerable interest here because it permits a simplified probabilistic description of scattering in terms of a random phasor sum. In the discrete models, various assumptions are made regarding the tissue microstructure. These typically involve the strength and spacing of the scatterers, including various levels of regularity in the spacing [18]. While many options have been explored regarding the microstructural characteristics of tissue, none have considered directly the influence of gross, or large-scale, structure, even though such structures are inherently regular and have a significant influence on the image data.

4.5 Scattering from Rough Surfaces

Substantial work has also been done by many researchers specifically on modeling the scattering of waves, both electromagnetic and acoustic, from rough surfaces [67, 68]. Theoretical scattering models are based on the Helmholtz integral equation, providing a foundation for both analytical and numerical investigation. Statistical results can be derived directly for simple cases of monochromatic plane waves and planar rough surfaces with simple statistical properties [67, 68]. For more complex scenarios, statistical results can be generated empirically via simulation for independent realizations of a random rough surface [69, 70]. The derived results for monochromatic wavefields, however, do not extend directly to pulsed, or nonmonochromatic, wavefields as shown in [69]. Extension of the existing methods to model conventional pulse-echo imaging is possible but time-consuming and computationally expensive. In all of the existing methods, a continuous surface representation is used with roughness described in terms of random perturbations of the surface [68]. For any anatomical surface of interest, such a representation would require significant computational resources.

4.6 Probabilistic Ultrasound Models and Statistical Tissue Characterization

The high sensitivity to changes in system characteristics and acoustic microstructure have led many researchers to investigate the use of statistical methods in characterizing ultrasonic scattering. Of the existing methods, those of interest are based on a random phasor sum analysis of scattering, a highly developed area with many applications. The intensity of the scattered signal at any image pixel is considered as a sum of contributions from scatterers within the “resolution cell”, a volume around the pixel location in which scatterers can contribute to the intensity. The dimensions of the resolution cell are determined by the shape of the system PSF. Each scatterer can then be considered as a phasor contributing an amplitude and phase. Quantities in the sum that are considered random are the number of scatterers in the cell and the amplitude and phase of each scatterer. This will be made clear in the following development, based on the results from the previous section on the linear systems model of image formation. The goal is a probability density function describing the amplitude given densities for the number of scatterers and their amplitudes and phases.

The phasor sum relations are based on a discrete representation of the tissue in terms of individual scattering elements with given locations and strengths. In the notation of the previous sections, the tissue reflectivity, $q(\mathbf{r})$, is then composed of a sum of scaled and delayed 3D impulses, $q_i\delta(\mathbf{r} - \mathbf{r}_i)$, where q_i and \mathbf{r}_i are the scatterer strength and location, respectively,

$$q(\mathbf{r}_0) = \sum_{i=1}^N q_i \delta(\mathbf{r} - \mathbf{r}_i). \quad (4.65)$$

In equation 4.48 for image formation, the discrete representation for the tissue structure can be substituted, along with the analytic signal representation for the PSF of equation 4.63 in terms of an envelope, $A(\mathbf{r}; \mathbf{r}_i)$, and a complex sinusoid, $e^{j2k_0(z-z_i)}$.

$$i_{rf}(\mathbf{r}) = \int d\mathbf{r}_0 q(\mathbf{r}_0) h(\mathbf{r}; \mathbf{r}_0). \quad (4.66)$$

Each term in the resulting sum is a phasor associated with a particular scatterer,

$$i_{rf}(\mathbf{r}) = \text{Re} \left\{ \sum_{i=1}^N q_i A(\mathbf{r}; \mathbf{r}_i) e^{j2k_0(z-z_i)} \right\}. \quad (4.67)$$

The sum can be rewritten in analytic signal form by taking a complex sinusoid at the transducer center frequency outside the sum,

$$i_{rf}(\mathbf{r}) = \text{Re} \left\{ \left[\sum_{i=1}^{N_{\mathbf{r}_i}} q_i A_i e^{j\phi_i} \right] e^{j2k_0 z} \right\} \quad (4.68)$$

where the number of scatterers has been reduced to $N_{\mathbf{r}_i}$, those in the resolution cell of location \mathbf{r}_i , the amplitude of each scatterer is the product of the scatterer strength, q_i , and a location-dependent amplitude, $A_i = A(\mathbf{r}; \mathbf{r}_i)$, and the phase of each scatterer is $\phi_i = -2k_0 z_i$. Note that the system characteristics are inherent in the amplitude because of the location-dependent PSF envelope term and in the phase based on the transducer center frequency.

Again, the RF image can be interpreted as the real part of an analytic signal, where the complex envelope, $\mathbf{i}(\mathbf{r})$, is a phasor sum,

$$i_{rf}(\mathbf{r}) = \text{Re} \left\{ \mathbf{i}(\mathbf{r}) e^{j2k_0 z} \right\}. \quad (4.69)$$

The amplitude of the resulting phasor is the quantity displayed in conventional images and, thus, of significant interest,

$$\mathbf{i}(\mathbf{r}) = \sum_{i=1}^{N_r} q_i A_i e^{j\phi_i} \quad (4.70)$$

$$|\mathbf{i}(\mathbf{r})| = \left| \sum_{i=1}^{N_r} q_i A_i e^{j\phi_i} \right|. \quad (4.71)$$

While the phasor sum notation greatly simplifies the characterization of scattering, the sensitivity to system characteristics and tissue microstructure requires random models for the amplitudes, phases, and number of scatterers. An exact solution for the resulting sum with no approximations or assumptions is intractable. Typical methods for characterizing the sum include calculation of the moments for real and imaginary components and the derivation of exact forms for densities describing the amplitude and phase in certain special cases. Common assumptions involve, for example, the independence of scatterer amplitude and phase, the number of scatterers, and, perhaps most importantly, the phase density.

The most common approximation is to begin by assuming a large number of scatterers [65, 67]. The real and imaginary components can then be considered asymptotically Gaussian under the Central Limit Theorem assuming that densities of the amplitude and phase satisfy certain requirements [67, 65]. As a result, the sum can be modeled as a complex Gaussian random variable and described in terms of the means and variances for the real and imaginary components. For simplicity, consider the sum, $\mathbf{i}(\mathbf{r})$, as the complex quantity, $x + jy$, with real component, x , and imaginary component, y . In general, a complex Gaussian density has the form

$$p_{x,y}(x, y) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-r^2}} \exp \left\{ -\frac{1}{2(1-r^2)} \left[\frac{(x-\mu_x)^2}{\sigma_x^2} - 2r \frac{(x-\mu_x)(y-\mu_y)}{\sigma_x\sigma_y} + \frac{(y-\mu_y)^2}{\sigma_y^2} \right] \right\} \quad (4.72)$$

where μ_x and σ_x^2 are the mean and variance of the real component, μ_y and σ_y^2 are the mean and variance for the imaginary component, and $r = \frac{E(xy) - E(x)E(y)}{\sigma_x\sigma_y}$ is the correlation coefficient. If the phase is distributed symmetrically about a constant, the sum can be rotated in the complex plane to be distributed symmetrically about the

real axis. In this case, μ_y , the mean of the imaginary component and the correlation coefficient, r , are both zero. The density then reduces to the following form,

$$p_{\mathbf{i}}(\mathbf{i}) = \frac{1}{2\pi\sigma_x\sigma_y} e^{-\frac{(x-\mu_x)^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}}. \quad (4.73)$$

Because of its simplicity and mild assumptions, the complex Gaussian result has been employed in many areas [65, 67]. Means and variances can be computed exactly for many densities describing amplitude and phase, allowing for a variety of characterizations for scattering data [67].

The complex Gaussian density function is a simple relation for expressing the variation of the RF image intensity at a location. The difficulty arises after transforming from a real-imaginary representation to the amplitude-phase representation necessary for describing the amplitude of the sum. As shown in [67], a general, but computationally demanding, result can be derived for a normalized amplitude, $u = \frac{|\mathbf{i}(\mathbf{r})|}{\sqrt{\sigma_x^2 + \sigma_y^2}}$, in terms of, $K = \sqrt{\frac{\sigma_r^2}{\sigma_i^2}}$, a measure of the asymmetry between the real and imaginary variances, and $B^2 = \frac{\alpha^2}{\sigma_r^2 + \sigma_i^2}$, the ratio of power between the mean value and variances,

$$p_u(u) = \frac{K^2 + 1}{K} u \exp \left[-\frac{1 + K^2}{2} \left(B^2 + \frac{1 + K^2}{2K^2} u^2 \right) \right] \sum_{m=0}^{\infty} (-1)^m \epsilon_m I_m \left(\frac{K^4 - 1}{4K^2} u^2 \right) I_{2m}[B(1 + K^2)u] \quad (4.74)$$

where I_m is a modified Bessel function of order m .

The general density, 4.74, for the amplitude of the phasor sum is significantly more complicated than the associated density for the complex random variable. As a result, a few common approximations are typically made to produce much simpler densities. The most common is the Rayleigh density, which results from assuming independent amplitude and phase, identically-distributed amplitude (denoted r here), and phase identically-distributed with a uniform density over one wavelength,

$$p_r(r) = \frac{r}{\sigma_x^2} e^{-\frac{r^2}{2\sigma_x^2}}. \quad (4.75)$$

The assumptions are equivalent to assuming a zero-mean real component and equal real and imaginary variances. The result is a much simpler expression, of course. The Rayleigh density has been used extensively in attempts to characterize scattering from

tissue [11]. The assumptions have been shown to be too strict, however, to accurately model scattered ultrasonic data [16].

Another common approximation is the Rician Distribution. It results from the sum of a constant vector and a Rayleigh-distributed vector and has the following form,

$$p_r(r) = \frac{r}{\sigma_x^2} e^{-\frac{(\mu_r^2 + r^2)}{2\sigma_x^2}} I_0\left(\frac{r\mu_x}{2\sigma_x^2}\right). \quad (4.76)$$

Even with this simplification, the expression still contains a modified Bessel function. The Rician density has been used extensively in attempts at tissue characterization [11], however its utility is yet to be determined in those applications.

A recent solution with very few approximations comes from [71] but requires computationally intensive Monte Carlo simulations for evaluation. Effects of the axial PSF were studied thoroughly within the random phasor sum approach. The effects of amplitude-phase independence were studied thoroughly for typical versions of the axial pulse. The results showed that independence assumptions can be too strict and limit the applicability of the approximating densities that are commonly used. The direct applicability of these models to scattering from gross shapes has not been tested, but the importance of assuming amplitude-phase independence will be relevant in later chapters.

Other analytic forms for the amplitude distribution have been proposed and investigated by many researchers. For instance, the k-, generalized k- and Nakagami-m distributions have all been considered [12, 72, 73]. The k- and generalized k-distributions have the advantage of being derived directly from the random phasor sum without any approximations. Both assume a negative binomial distribution for the scatterer population with N tending toward ∞ . The generalized k-distribution was recently found to perform fairly well and, in fact, better than other typically used distributions, in describing scattering from breast and abdominal tissue under aberrating conditions [72]. The complexity of that specific situation suggests promise for the generalized k-distribution, although shortcomings were evident, e.g., the poor fit for a broadband pulse, i.e., a typical pulse used in conventional systems (good fits were obtained for harmonic amplitude distributions). All of these models require the calculation of sums of gamma and Bessel functions. Such calculations, combined with their questionable utility for data from actual systems, make them impractical for this work.

A fundamental limitation to all existing probabilistic models is the lack of an implicit representation for gross structure. Regularity in scatterer spacing is known to affect the result significantly [18], and gross structure is a component in virtually all tissues and possesses inherently regular spacing. Approaches to representing the effect of regular structure range from the Wold decomposition of a signal [15] to the generalized spectrum [74], all of which rely on the regular structure having a predictable contribution to the statistics of the scattered data. The contribution is then detected and removed. These approaches have not been shown to provide useful results in any clinical setting and also fail to incorporate gross structure directly into the model.

4.7 Conclusions

The existing models for ultrasonic imaging, in terms of system models, tissue models, and probabilistic scattering models, provide a rich background for the development of improved models. These models are incomplete, however, because of the lack of an appropriate characterization of the combined interactions of gross shape, tissue microstructure, and the system characteristics. The linear systems model for imaging systems forms the foundation upon which tissue models and probabilistic models are based. The associated random phasor sum representation for the RF intensity provides the basis for developing a single model representing image intensity with an appropriate characterization of the imaging system and underlying tissue structure.

Chapter 5

A Discrete-Scatterer Model for Rough Surfaces

In the previous chapter, a physical model was developed for the imaging system. In this chapter, a surface model completes the physical description for image formation. The objective of the complete physical was an accurate models with minimal computational requirements that incorporated system characteristics, surface shape and surface microstructure, and could also be extended naturally to a probabilistic model. For the imaging system, a linear systems model with a simple approximation to the 3D PSF was investigated. For the surface model, a representation consisting of discrete scatterers with positions based on the gross shape was investigated. The complete physical model was investigated via comparisons of simulated and actual images.

5.1 Image Formation

Linear systems image formation models have been used previously for simulation purposes in attempts to study statistical properties of scattering data [17, 18, 75]. In most work, the tissue model has been only 1D or 2D. In one exception [75], a 3D tissue model was created, but only a 2D slice was used in the simulation. In that work, the tissue structure was represented on a uniform grid with very fine spacing. This raises an issue common to all simulation studies. In a discrete-scatterer model, individual scatterers have positions on the continuum. In computing, positions are approximated by either storing the position to the precision of the data type used or by representing the entire medium by a grid with sufficiently fine spacing. For large

surfaces represented in 3D, positioning on the continuum can significantly reduce the computational resources required for storage while remaining consistent with the theoretical model. In an early study in this area, the surface was represented by a 3D grid with sub-wavelength (approximately $\lambda/5$) spacing [76]. In this chapter, the tissue model is still 3D but with scatterer placement on the continuum.

5.1.1 Imaging System Model

Recall from Chapter 4 that the RF image can be modeled as a linear system with PSF, $h(\mathbf{r}; \mathbf{r}_0)$, and tissue reflectivity, $q(\mathbf{r})$,

$$i_{rf}(\mathbf{r}) = \int d\mathbf{r}_0 q(\mathbf{r}_0)h(\mathbf{r}; \mathbf{r}_0). \quad (5.1)$$

For the experimental parts of this dissertation, the PSF is assumed to be shift-invariant, $h(\mathbf{r}; \mathbf{r}_0) = h(\mathbf{r} - \mathbf{r}_0)$, allowing the integral to be written as a convolution,

$$i_{rf}(\mathbf{r}) = h(\mathbf{r}) * q(\mathbf{r}). \quad (5.2)$$

The PSF is further simplified as in Chapter 4 by modeling with a 3D Gaussian envelope, $A(\mathbf{r})$, resulting in the following expression for the PSF, with wavenumber, k_0 ,

$$h(\mathbf{r}) = A(\mathbf{r}) \cos(2k_0z) \quad (5.3)$$

$$A(\mathbf{r}) = e^{-x^2/\sigma_x^2} e^{-y^2/\sigma_y^2} e^{-z^2/\sigma_z^2} \quad (5.4)$$

where the standard deviations, $\sigma_x, \sigma_y, \sigma_z$, denote the beam widths in each of the directions. The envelopes of the RF images were generated along the axial dimension. Envelope detection was accomplished by taking the magnitude of the complex signal, generated with the Hilbert transform.

5.1.2 Surface Model

The surface was modeled as a collection of discrete scatterers, a representation that has many advantages over the continuous representation used in existing models. These advantages stem from a simplified view of the combined effects of system and tissue characteristics. In this discrete-scatter representation, each scatterer represents a major scattering element of size smaller than a wavelength. The roughness of

surfaces is typically characterized in such a way that small (sub-wavelength) hills and valleys cover the surface. A hill or valley pointing in the direction of the transducer could be considered a major scattering element in this model. In this approach, a collection of small scattering elements comprised the entire acoustic representation of the surface.

For this work, the following advantages of the discrete representation make it an attractive choice:

- Computation is simplified because the small scattering elements enable a unification of structure and imaging system in a linear systems approach. For a continuous representation, the surface integral equation for the scattered wavefield would have to be computed numerically for each frequency of the incident wavefield then summed to calculate the image.
- A similar intuitive simplicity results from considering the surface as a collection of distinct elements, with their contributions adding constructively or destructively depending on phase separation, rather than considering the solution to a surface integral equation.
- The discrete-scatterer model allows use of the convenient random-phasor-sum analysis of scattering. Most probabilistic models are based on that model, providing an extensive background of previous work for the development of a new probabilistic image model.

Continuous surfaces contain structure inherently larger than a wavelength, violating the assumption that scattering regions are small relative to a wavelength. Insana [57] notes that studies have shown that the linear models appear to apply well to surfaces, anyway. The applicability to surfaces is still relatively untested, though. The further simplification of reducing an entire surface to a collection of its major scattering elements was, to my knowledge, untested before this work.

The gross surface was represented as a triangulated surface, i.e., a set of triangular elements defined by their vertices and edges. The triangulated surface representation can be easily produced from a volume segmentation using the Marching Cubes algorithm [36, 77]. While only the triangles were used, this representation inherently defines second-order (curvature) information, which could be used for a more nearly accurate description of the surface [36]. Furthermore, several computational algorithms exist in the computer graphics arena for the manipulation of these surfaces.

The discrete version of the surface permits a wide variety of medium parametrization. Scatterer spacing on the surface can be characterized in many ways, from completely regular spacing to random spacing. For example, models that fundamentally describe spacing regularity in 1D scattering arrangements are now being extended [18] to multiple dimensions. In addition to spacing, the surface can be characterized by the concentration of scatterers. Surface roughness can also be incorporated independently as a perturbation of each scatterer in the direction normal to the surface.

Any choice for distributing scatterers on the surface results in a discrete representation consisting of a collection $X = \{\mathbf{r}_i, i \in [1, \dots, N]\}$ of N scatterers, with scatterer i of amplitude A_i at position (\mathbf{r}_i) . The reflectivity function consists of a sum of appropriately scaled 3D delta functions, $\delta(\mathbf{r} - \mathbf{r}_i)$,

$$q(\mathbf{r}) = \sum_{i=1}^N A_i \delta(\mathbf{r} - \mathbf{r}_i). \quad (5.5)$$

The RF image is then a sum of scaled and delayed versions of the PSF,

$$i_{rf}(\mathbf{r}) = h(\mathbf{r}) * q(\mathbf{r}) \quad (5.6)$$

$$= \sum_{i=1}^N A_i h(\mathbf{r} - \mathbf{r}_i). \quad (5.7)$$

In the frequency domain, the convolution is a product, allowing the tissue response to be represented as a sum of scaled delay terms,

$$I_{rf}(\mathbf{u}) = H(\mathbf{u})Q(\mathbf{u}) \quad (5.8)$$

$$= H(\mathbf{u}) \sum_{i=1}^N A_i e^{-j(\mathbf{u}\mathbf{r}_i)} \quad (5.9)$$

where

$$i_{rf}(\mathbf{r}) \xleftrightarrow{\mathcal{F}_3} I_{rf}(\mathbf{u}) \quad (5.10)$$

where \mathcal{F}_3 denotes a 3D Fourier transform relationship.

The system response is assumed to be bandlimited, thus the image response can be computed exactly for any number of scatterers, without the limitation on scatterer positions imposed by a uniform grid as in [75, 76]. Computation can be performed in

the spatial domain as a sum of shifted versions of the PSF. Alternatively, a complete tissue response can be computed in the frequency domain, with the image computed using an inverse transform of the product of tissue and system responses.

5.2 Methods

Models were evaluated by comparing ultrasonic images of a cadaveric vertebra *in vitro* to simulated images. The vertebral surface provides a good medium for evaluation because of its intricate curvature and sub-wavelength roughness. The models could potentially be used to describe images of any rough surface, however. Imaging of the spine, the vertebrae in particular, is of interest in the area of treatment guidance (delivery of radiosurgery and guidance of traditional surgery) based on CT images of the spine as described in Chapter 2.

The weak scattering assumption inherent in the linear systems model is violated for the bone surfaces used here, although we are primarily interested in accurately representing scattering coming directly from the bone. Additional scattering that results from the high scattering strength, i.e., multiple reflections, is not necessary in the model at this point.

In the next section, simulated images are compared to actual images of the same vertebra. Registration between the actual and simulated images was achieved using methods and equipment from image-guided surgery, including optical localization for tracking the ultrasound probe and methods for registration of the images and vertebra. For a description of similar methods, see, e.g., [3].

A satisfactory quantitative measure for evaluating the model is elusive. Such a measure is desirable in many instances; first, to quantify the basic performance of the model, second, as a basis for improving the model and choosing parameter values, and finally, for use as a cost function for image analysis. Specifically, the difficulty is in quantitatively comparing the simulated and actual images. From the deterministic description of gross shape, even two simulated images would be expected to have significant variation at the pixel level because of the random description of the microstructure.

Simple quantitative measures, e.g., energy of a difference image, correlation coefficient, etc., were investigated briefly by attempting to “match” two images simulated from the same image plane over a 2 mm range of translation. While such measures were productive for images with coherent scattering (accurately identifying

the true translation), the measures were ambiguous around the correct translation for images with only texture from incoherent scattering. The failure of these simple comparisons makes sense. No theoretical basis exists for evaluating the similarity between two arbitrary random variables by taking the difference energy, correlation, etc. of two samples of that random variable.

Accuracy of registration between the actual images and the anatomical surface is a further limitation. The 2 mm error in registration is far greater than the difference in image registration that can be detected visually. Even if a meaningful quantitative measure did exist, evaluation would require an almost exact registration between the actual images and those simulated from the anatomical surface.

Finally, these problems further motivate the need for a pixel-based image model based on system characteristics and the gross and microstructural surface characteristics. In addition to satisfying the needs for model-based image analysis (and a means for more accurate registration), such a model would also provide the theoretical basis for development of a quantitative measure for model evaluation. The probabilistic model is the focus of Chapters 6, 7 and 8 of this dissertation. In this chapter, model validity has been assessed by visual comparison of the simulated and actual images. Attention has been directed to the location and extent of scattering, specifically the texture and relative intensities in the image and the sites of coherent and incoherent scattering.

Data for this chapter were collected using image-guided surgery equipment and methods as described in Chapter 2.

5.2.1 System Model Implementation

The system was modeled using the simplified PSF model of equation 5.3. The center frequency, $f_c = 6.0$ MHz, was given by the manufacturer. The PSF was assumed spatially invariant for this study, and the elevation and lateral PSF widths, $\sigma_y = 1.5$ mm and $\sigma_x = 0.5$ mm, were calculated from equations given by the manufacturer. The axial width, σ_z , was chosen to be 0.2 mm. It was selected from a range of approximately 0.15 to 0.4 mm based on visual comparison of actual images (from the Tetrad system) to images simulated with the various values for the width. In this study, exact relations between the PSF parameters and the actual system characteristics were not considered significant, and the approximate values used were deemed sufficient because of the quality (determined visually) of the results.

5.2.2 Surface Model Implementation

In images of bone surfaces, attenuation is a significant effect that warrants consideration. We make the assumption that bone is completely attenuating, or occluding. Before the discrete scatterers were generated, the surface was modified to account for this effect. A ray-tracing approach was used to determine visibility and was implemented using a modified version of the well-known computer graphics algorithm called Hidden Surface Removal [46]. In typical use of this algorithm for rendering a surface to a display, each triangle is projected onto the viewing plane, rasterized according to the display grid, and processed under lighting assumptions to generate an intensity value. During processing, the depth (Z) at each screen pixel is stored in a “Z buffer” so that only the closest triangle is displayed. In this case, the index of the closest triangle was stored instead of the rendering intensity, and, instead of displaying the Z Buffer, its contents were used to remove those triangles that were not visible to the transducer. A more nearly exact approach in this case would have been to “clip” those triangles which are partially occluded, but the computational intensity of such an approach makes it unsuitable at this time.

After accounting for occlusion, a collection of discrete scatterers was generated for the remaining triangles to form the acoustic model of the surface. We have parametrized the distribution of scatterers on the surface according to concentration (scatterers/area) and surface roughness. For each triangle, the number of scatterers was calculated as the product of the triangle area and scatterer concentration. The in-plane position of each scatterer was then generated from a 2D uniform distribution over the triangle. Scatterer position, $y \in \mathbb{R}^3$, was generated from two uniformly-distributed random variables, λ_1 and λ_2 for triangle vertices, $x_1, x_2, x_3 \in \mathbb{R}^3$ as follows:

1. Generate $\lambda_1, \lambda_2 \sim U[0, 1]$ until $\lambda_1 + \lambda_2 \leq 1$,
2. $y = x_3 + \lambda_1(x_1 - x_3) + \lambda_2(x_2 - x_3)$,

where the triangle borders and interior are represented by the combination of the point x_3 and the vectors $x_1 - x_3$ and $x_2 - x_3$ with $\lambda_1, \lambda_2 \in [0, 1]$. After the in-plane position was determined for each scatterer, it was perturbed in the direction normal to the surface to account for roughness. The perturbation was generated from a Gaussian-distributed random variable, with roughness characterized by the standard deviation of that perturbation.

In our investigation, we used surface scatterer concentration values between 50 and 150 scatterers/mm². For surface roughness, standard deviations between 0.001 and 0.1 mm were used. Changes in the concentration and roughness yielded modest changes in the images. The most visible change was a decrease in the coherent scattering sites when the roughness was nearly one wavelength. In the images that follow, scatterer concentration of 64 scatterers/mm² and surface roughness of 0.01 mm were used because the texture and intensity produced in the simulations were similar to those of the actual images.

Images were acquired and simulated for several image planes across the entire vertebra. Images in the results section are for one such image plane that contained a mix of coherent and incoherent scattering, as well as substantial out-of-plane effects. Performance of the simulation was similar for other images, when based on visual appearance.

5.3 Results

These methods were used to generate simulated images of a cadaveric vertebra. The images were compared to actual images of the same vertebra, with accuracy assessed via visual comparison (see Methods). The primary characteristics of interest were variations in texture and intensity, with attention to locations of coherent and incoherent scattering.

Figure 5.1 shows images typical of the results. The image plane is the sagittal plane, several views of which are shown in Figure 5.2 over the rendered surface. The major vertebral structures are, from left to right, the facet joint, the lamina and the inferior articular process. The image plane was oriented obliquely to the curvature of the facet joint, resulting in a broad texture of incoherent scattering. The lamina was oriented normally to the image plane, resulting in intense coherent scattering. The orientation of the inferior articular process varies along its surface, resulting in intense coherent scattering at the top of the process and various textures from incoherent scattering along the sides of the process. Note that some incoherent scattering texture, e.g., along the facet joint and articular process, exists in regions that are flat in the lateral dimension. This incoherent scattering is a result of the oblique orientation in the elevation dimension and is reproduced in the simulated image.




Figure 5.1: Actual and simulated images showing incoherent scattering from the inside of the facet joint (left), coherent scattering from the lamina (center), and incoherent and coherent scattering from the inferior articular process (right).






Figure 5.2: View of image plane overlaid on surface rendering

Figure 5.3 shows a closer view of the inferior articular process from the previous images, revealing substantial similarities. At the peak of the process, coherent scattering is evident in both images, with similar intensities relative to the surrounding texture. The surrounding texture on the sides of the process is of similar width and character as well. There are, however, differences between the two images. At the peak of the process, the coherent scattering in the simulated image is wider laterally and narrower axially than in the actual image. The incoherent texture is of slightly different character in the two images, perhaps a little sharper in the simulated image. As seen in the next figure, though, variation at this subtle level is well within the range of observed image variation within the accuracy in tracking the probe (around 2 mm).



Actual Image



Simulated Image

Figure 5.3: Close-up of actual and simulated images from previous figure showing only the inferior articular process of the cadaveric L4 vertebra.

The images in Figure 5.4 are all simulated, from image planes separated by 0.4 mm in the elevation dimension, from lateral (top) to medial (bottom). The six images cover the range of expected error in tracking the probe. Note the wide variation relative to the differences between simulated and actual images in the previous figures. Specifically, note on the left of the images the change in location and extent of scattering from the inside of the facet joint. In the top image, the image plane cuts through the upper point of the facet joint, where the surface of the joint is more oblique to the image plane. The result is a texture that is wider axially since scattering returns from a broader part of the surface. Similarly, the contributions from the curved aspect of the lamina (between the lamina and inferior articular process) in the top image do not appear in the other images. Finally, the curved region between the lamina and facet joint in the actual image of Figure 5.1 appears in various forms in some of the simulated images.

The high sensitivity of the images in Figure 5.4 to change in the image plane is due to translation only. Registration inaccuracy includes error in both translation and rotation. Rotation of the image plane, or equivalently the object, also causes changes




Figure 5.4: Six simulated images from planes separated by 0.4 mm in the elevation dimension (covering a total of 2 mm). Note the substantial variation in image characteristics due to the curvature of the vertebral surface.

in the image. Of course, the degree of change depends on the origin of rotation, thus effects of rotation depend on the choice of coordinate systems for representing the image plane and the object. The effects can be misleading because of the dependence on an origin and can be difficult to visualize. They are a part of the registration problem, however, and add another dimension to the problem of validating the proposed model.

The images of Figure 5.4 show that the model can produce visibly similar images well within the variation expected due to tracking error. The high sensitivity to image plane makes model validation a difficult one. Images could be simulated over a range of translation and rotation, with a best match chosen based on visual similarity. A range of two degrees in each rotation dimension and two mm in each translation direction would produce several hundred images, however, requiring significant computation time and hours of analysis for a result that would still be subjective.

5.4 Discussion

These results indicate that much of the observed variation in ultrasonic images of rough surfaces can be accounted for with a relatively simple model for the complete imaging process. The parametrized discrete-scatterer model for representing the acoustic properties of the surface is a simple approximation to the continuous description commonly used. The assumptions of separability and Gaussian beam widths for the PSF are simple approximations as well. For such simple models, however, the simulated images are quite remarkable in the amount of detail produced and the overall similarity to the actual images.

Much of the image variation, thus, can be obtained from only the gross surface shape with some assumptions about the basic characteristics of the surface roughness and the imaging system. Slight variation in the gross shape is apparent in the images, warranting the “highly-detailed” shape representation advocated by others for image analysis [8]. The linear systems imaging model used here contains basic characteristics that produce both coherent and incoherent scattering. The 3D PSF must be used, however, since out-of-plane aspects of the gross surface shape can induce a substantial contribution to the images. The simple microstructural model used here produced good visual similarity in the texture and intensity of the actual and simulated images. Limitations in the accuracy of registration of phantom and images, along with the absence of an appropriate quantitative measure of similarity between simulated and

actual images, prevent further evaluation without the image model of later chapters. Again, the physical models of this chapter provide the basis for that development.

Refinement of the model is a topic of relevance in each chapter of this dissertation. The models could be refined in many ways. For instance, the surface characterization could be modified with various parametrizations of the discrete-scatterer representation. The PSF could be refined at many levels, including numerical modeling, removal of the separability assumption, or modeling of the spatial variation. The Hilbert-transform-based envelope detection scheme used here is an ideal method. Many imaging systems use simpler schemes, however, and these could be included in the modeling. Refinements should be made based on an improvement in performance, and the relevant area of performance in this work is inference of shape. Under that premise, significant refinements or modifications were left for future work so that effort could be focused on developing the image model and investigating its potential in inference.

For the purposes of tissue characterization in rough surfaces, the discrete-scatterer model could potentially have advantages over continuous representations, even for characterization at the microstructural level. A typical approach to tissue characterization is to develop a parametrized model for the tissue microstructure, then estimate parameters of the model for classification. The discrete representation may make varying the parametrization of the surface easier. The extension of existing probabilistic models for tissue characterization may be simpler than for image analysis, with beneficial results without inclusion of the PSF, for instance. Also, investigation using the discrete representation is computationally simpler than the continuous representation, simplifying the required analysis.

For any application of the model to clinical images, a patient-specific model of the gross shape would be required. The images in Figure 5.4 show that ultrasonic images can be quite sensitive to the gross shape. An accurate model can be difficult to obtain, though, and it would have to be acquired in the clinic. In pattern theory, this difficult problem of inferring shape is approached with *deformable templates* [8, 9] (see Chapter 3). For a single category of shapes, e.g., lumbar vertebrae, a model for any new shape is found by estimating a transformation that relates the new structure to a similar template structure. For example, for lumbar vertebrae, a representative vertebral surface would be built very carefully from CT images of a “typical” patient. The surface of any new lumbar vertebra, either that of a new patient or another from the lumbar spine of the same patient, is then found as a transformation of

the template, instead of performing the difficult task of creating, effectively, a new template shape.

While the model has been tested only for images of vertebrae, it is applicable to the surface of any anatomical organ, e.g., arterial surfaces, the surface of the liver, etc., given only gross shape and parameters for the concentration and roughness. While images of these organs are typically used to characterize the interior of the structure, information about the surface can also be useful, both in tissue characterization and in identification of structure position. Evaluation with these soft-tissue structures using a similar approach to that used here would be more difficult due to increased difficulties in tissue registration. For example, isolation of the tissue surface for imaging *in vitro* with either CT or ultrasound would be difficult with soft tissue. For such structures, evaluation would probably have to be based on *in vivo* images. Tissue models would have to be modified to account for movement of the tissue between CT imaging and ultrasonic imaging. Modeling and compensation for movement between scans is an additional image analysis problem. With soft-tissue structures, the variation in surface shape is substantially more complicated, involving elastic deformation instead of the rigid motion incurred by a vertebra. Inclusion of the surrounding volumes of tissue in the model adds another level of complexity as well.

For the purposes of image analysis, the tissue surrounding a surface raises an interesting point. Some structures are more appropriately characterized by their interior than by their surface. The same model-based approach to image analysis is applicable, and a similar probabilistic data model representing image as a function of gross shape is required. The discrete-scatterer basis of the surface model used here could also provide a foundation for modeling volumes. The model could be evaluated and parametrized in similar fashion, and extension to a probabilistic model incorporating the PSF is expected to be similar for volumes.

5.5 Conclusions

A physical model was developed for formation of ultrasonic images of rough surfaces based on a discrete-scatterer model for the rough surface and a simple 3D PSF representing the imaging system. The representation of an entire surface by a collection of discrete scattering elements was untested before this work. The model was developed specifically for application to model-based image analysis, with a design that was both

computationally efficient and extendable to a probabilistic form. Simulated images were generated from the model which were visually similar to actual images of the same structure. The model provides the basis for the development of a probabilistic image model in later chapters.

Chapter 6

Amplitude Mean and Variance for the Random Phasor Sum

The tissue and system models of the previous chapter were developed explicitly for extension to a probabilistic model representing image intensities as a function of underlying structure shape. The models are relatively simple but produce substantial complexity when analyzed in terms of a random phasor sum. In this chapter, methods are explored for computing the amplitude mean and variance of the random phasor sum. The technical issues involve characterization of the amplitude in terms of probability densities for elements of the random phasor sum, i.e., N , the number of phasors, A , the phasor amplitude, and ϕ , the phasor phase. Implications for application to the shape-based image model will be mentioned along the way.

6.1 Image Formation as a Random Phasor Sum

Recall from previous chapters the physical model for image formation. Since the discrete-scatterer tissue model has been used, image formation can be modeled with a random phasor sum as in Chapter 4. At any pixel, with position \mathbf{r} , the quantity of interest is the amplitude of the phasor sum, $\mathbf{i}(\mathbf{r})$, which has the following form from Equation 4.70,

$$\mathbf{i}(\mathbf{r}) = \sum_{i=1}^{N_r} q_i A(\mathbf{r}; \mathbf{r}_i) e^{j\phi_i} \quad (6.1)$$

where $N_{\mathbf{r}}$ is the number of scatterers in the pixel's resolution cell, and, for each scatterer i , q_i is the reflectivity, $A(\mathbf{r}; \mathbf{r}_i)$ is the position-dependent amplitude of the PSF envelope, and $\phi_i = -2k_0 z_i$ is the position-dependent phase. For simplification, the scatterer strength and envelope amplitude will be combined as $A_i = q_i A(\mathbf{r}; \mathbf{r}_i)$ to denote the phasor amplitude for scatterer i .

For any tissue of interest, including the discrete-scatterer surface model of Chapter 5, components of the phasor sum are random. For a given pixel, the strengths, locations, and number of scatterers in its resolution cell are random and produce random amplitude, A_i , and phase, ϕ_i in ways that depend on the system PSF. Furthermore, for a given surface shape, these interactions change *at every pixel*, motivating the comprehensive model of this dissertation.

The image model requires a probability density function, $p_{|\mathbf{i}(\mathbf{r})|}(|\mathbf{i}(\mathbf{r})|)$, describing the echo amplitude at each image pixel. The effects of shape on local densities for N , A and ϕ , can significantly affect the amplitude density. General solutions for the amplitude density, however, can be quite complex as seen in Chapter 4. In this work, instead of focusing on an exact form for the amplitude density and dismissing local effects due to shape and system characteristics, the amplitude is characterized by its mean and variance at every pixel with comprehensive treatment of the shape, microstructure and system characteristics. The derivation begins by following [65, 67] in using the complex Gaussian approximation for the complex sum. The complex Gaussian density is then used as a basis for characterizing the *amplitude* of the sum by its mean and variance.

6.2 The Gaussian Approximation to the Complex Sum

This development follows that of Beckmann [67], primarily, with the exception that, here, N is treated as random. Typically, N is assumed to be large enough that the sum is complex Gaussian [65, 67]. Formally, this holds asymptotically and requires some independence assumptions. Practically, the sum can be considered complex Gaussian for $N \geq 25$ [67]. The sum can be written in terms of its real and imaginary

components, denoted $x(\mathbf{r})$ and $y(\mathbf{r})$, respectively, as

$$\mathbf{i}(\mathbf{r}) = \sum_{i=1}^{N_r} A_i e^{j\phi_i} \quad (6.2)$$

$$= \sum_{i=1}^{N_r} A_i \cos \phi_i + j \sum_{i=1}^{N_r} A_i \sin \phi_i \quad (6.3)$$

$$= x(\mathbf{r}) + jy(\mathbf{r}). \quad (6.4)$$

In general, for a complex random variable, $x + jy$, the complex Gaussian density is given as in Equation 4.72,

$$p_{x,y}(x, y) = \frac{1}{2\pi\sigma_x\sigma_y\sqrt{1-r^2}} \exp \left\{ -\frac{1}{2(1-r^2)} \left[\frac{(x-\mu_x)^2}{\sigma_x^2} - 2r \frac{(x-\mu_x)(y-\mu_y)}{\sigma_x\sigma_y} + \frac{(y-\mu_y)^2}{\sigma_y^2} \right] \right\} \quad (6.5)$$

where μ_x and σ_x^2 are the mean and variance of the real component, μ_y and σ_y^2 are the mean and variance for the imaginary component, and $r = \frac{E(xy) - E(x)E(y)}{\sigma_x\sigma_y}$ is the correlation coefficient. Several assumptions are usually made to simplify this form (see Chapter 4), but, since some of those assumptions will not always hold in this work, the development will proceed from this point. For any densities describing N, A_i and ϕ_i , the complex sum can be characterized by computing the means and variances of the real and imaginary components and the correlation coefficient.

The typical assumptions made in the phasor sum analysis involve the following [65, 67]:

1. Independence of individual phasor quantities, $A_i \perp A_j, \phi_i \perp \phi_j, i \neq j$,
2. Independence of the number of scatterers from other quantities,
3. Independence of amplitude and phase for each phasor, $A_i \perp \phi_i$,
4. Identically distributed amplitudes and identically distributed phases,
5. Symmetric phase distributions that allow rotating the phasor sum to align with the real axis (see Chapter 4).

Assumptions (3) and (5) will be too strong to hold at all times in this work, thus the development that follows assumes only (1), (2) and (4). Effects of assumptions (3) and (5) will be mentioned throughout and also used for illustration purposes.

6.3 Computing the Complex Gaussian Parameters

The \mathbf{r} dependence, indicating the pixel basis for the sum, will be dropped for the rest of this development, although the importance of the implied pixel-based model cannot be understated. The parameters of the complex Gaussian are given by the following well-known relations:

$$\mu_x = E(x), \quad (6.6)$$

$$\mu_y = E(y), \quad (6.7)$$

$$\sigma_x^2 = E[(x - E(x))^2] = E(x^2) - E^2(x), \quad (6.8)$$

$$\sigma_y^2 = E[(y - E(y))^2] = E(y^2) - E^2(y), \quad (6.9)$$

$$r = \frac{E(xy) - E(x)E(y)}{\sigma_x \sigma_y}. \quad (6.10)$$

Consider first the mean of the real component (the mean of the imaginary component will be analogous). The real component, x , of the sum is given by a sum of weighted cosines,

$$x = \sum_{i=1}^N A_i \cos \phi_i. \quad (6.11)$$

The mean can be computed via the expectation, $E(\cdot)$, as follows: using (1) the iterated expectation [78] to compute the mean via a conditional expectation for random N , (2) linearity of the expectation, and (3) the i.i.d. (independent and identically distributed) nature assumed of different scatterers. Expectations are with respect to the variables in each argument.

$$\mu_x = E(x) \quad (6.12)$$

$$= E\left(\sum_{i=1}^N A_i \cos \phi_i\right) \quad (6.13)$$

$$\stackrel{(1)}{=} E\left[E\left(\sum_{i=1}^N A_i \cos \phi_i | N\right)\right] \quad (6.14)$$

$$\stackrel{(2)}{=} E\left[\sum_{i=1}^N E(A_i \cos \phi_i)\right] \quad (6.15)$$

$$\stackrel{(3)}{=} E[NE(A_i \cos \phi_i)] \quad (6.16)$$

$$= E(N) E(A_i \cos \phi_i). \quad (6.17)$$

This result allows the mean of the real component to be computed in terms of the means of N and $A_i \cos \phi_i$. Given a shape model, these quantities will generally vary at the pixel level and must be computed from phase, amplitude and number densities derived specifically for that pixel.

The variance of the real component can be calculated similarly. First, the individual real components, $A_i \cos \phi_i$, of the sum will be denoted x_i for simplification,

$$x = \sum_{i=1}^N A_i \cos \phi_i = \sum_{i=1}^N x_i. \quad (6.18)$$

The variance calculation requires the second moment of x . It can be calculated as follows, using (1) the iterated expectation, (2) the linearity of the expectation, and (3) independence of different samples:

$$E(x^2) \stackrel{(1)}{=} E[E(x^2|N)] \quad (6.19)$$

$$= E\left[E\left(\sum_{i=1}^N x_i \sum_{j=1}^N x_j\right)\right] \quad (6.20)$$

$$\stackrel{(2)}{=} E\left[\sum_{i=1}^N \sum_{j=1}^N E(x_i x_j)\right] \quad (6.21)$$

$$E(x_i x_j) \stackrel{(3)}{=} \begin{cases} E(x_i^2) & i = j, \\ E^2(x_j) & i \neq j. \end{cases} \quad (6.22)$$

$$= E[N E(x_i^2) + N(N-1) E^2(x_i)] \quad (6.23)$$

$$= E(N) E(x_i^2) + E(N^2 - N) E^2(x_i) \quad (6.24)$$

$$= E(N) \sigma_{x_i}^2 + E(N^2) E^2(x_i). \quad (6.25)$$

The variance is then easily obtained by combining this expression with the previous result for the mean of the real component,

$$\sigma_x^2 = E(N) \sigma_{x_i}^2 + E(N^2) E^2(x_i) - E^2(N) E^2(x_i) \quad (6.26)$$

$$= E(N) \sigma_{x_i}^2 + \sigma_N^2 E^2(x_i). \quad (6.27)$$

This result is true, in general, for a sum of i.i.d. random variables [79]. For deterministic N , σ_N^2 is zero, and this expression simplifies to that in [67] and the well-known result that the variance of a sum of i.i.d. variables is the sum of the variances [78].

The additional term adds the variance in the number of scatterers scaled by the power in the individual real components.

The previous expression for the variance can be written out to show the dependence on the densities for N , A , and ϕ .

$$\sigma_x^2 = E(N) \left[E(A_i^2 \cos^2 \phi_i) - E^2(A_i \cos \phi_i) \right] + \sigma_N^2 E^2(A_i \cos \phi_i). \quad (6.28)$$

Note that the entire expression can be computed from the first and second moments for N and $A_i \cos \phi_i$.

The variance of the imaginary component is identical with the exception of a sin term instead of the cos term in equation 6.28,

$$\sigma_y^2 = E(N) \left[E(A_i^2 \sin^2 \phi_i) - E^2(A_i \sin \phi_i) \right] + \sigma_N^2 E^2(A_i \sin \phi_i). \quad (6.29)$$

For the correlation coefficient, the expectation, $E(xy)$, of the product of the real and imaginary components must be computed. The computations are similar to those for the variance.

$$E(xy) = E \left(\sum_i A_i \cos \phi_i \sum_j A_j \sin \phi_j \right) \quad (6.30)$$

$$= \sum_i \sum_j E(A_i \cos \phi_i A_j \sin \phi_j) \quad (6.31)$$

$$E(A_i \cos \phi_i A_j \sin \phi_j) = \begin{cases} E(A_i^2 \cos \phi_i \sin \phi_i) & i = j, \\ E(A_i \cos \phi_i) E(A_j \sin \phi_j) & i \neq j. \end{cases} \quad (6.32)$$

$$= E(N) E(A_i^2 \cos \phi_i \sin \phi_i) + E(N^2 - N) E(A_i \cos \phi_i) E(A_j \sin \phi_j). \quad (6.33)$$

Derivations of the mean and variance of the imaginary component are similar. In summary, the parameters are given by the following equations:

$$\mu_x = E(N)E(A_i \cos \phi_i), \quad (6.34)$$

$$\mu_y = E(N)E(A_i \sin \phi_i), \quad (6.35)$$

$$\sigma_x^2 = E(N)E(A_i^2 \cos^2 \phi_i) + \sigma_N^2 E^2(A_i \cos \phi_i), \quad (6.36)$$

$$\sigma_y^2 = E(N)E(A_i^2 \sin^2 \phi_i) + \sigma_N^2 E^2(A_i \sin \phi_i), \quad (6.37)$$

$$r = \frac{1}{\sigma_x \sigma_y} \left[E(N) [E(A_i^2 \cos \phi_i \sin \phi_i) - E(A_i \cos \phi_i)E(A_i \sin \phi_i)] + \sigma_N^2 E(A_i \cos \phi_i)E(A_i \sin \phi_i) \right]. \quad (6.38)$$

Again, the expressions depend only on the first and second moments of N , $A_i \cos \phi_i$ and $A_i \sin \phi_i$. In cases where no more assumptions are made, these results will be computed using numerical integration techniques.

The assumptions that are typically made in phasor sum analyses of scattering simplify the expressions a great deal. First, if amplitude and phase are considered independent for each scatterer ($A_i \perp \phi_i$), all expectations involving A_i and $\cos \phi_i$ or $\sin \phi_i$ can be computed as products of expectations, e.g., $E(A_i \cos \phi_i) \underset{A_i \perp \phi_i}{=} E(A_i)E(\cos \phi_i)$. If, in addition, the phase is assumed to be symmetric about zero, the expected value of $\sin \phi_i$ is zero, and the mean, μ_y , of the imaginary component becomes zero as well as the correlation coefficient, r .

Example: Sensitivity of the Complex Gaussian Parameters to the Phase Density

Much of the sensitivity to the phasor sum involves the phase density. A good choice for illustration purposes is a uniform density centered about 0 and having a width $2a$, i.e., uniform on the interval $[-a, +a]$. Specifically, let

$$p_\phi(\phi) = \begin{cases} \frac{1}{2a} & -a \leq \phi \leq a, \\ 0 & \text{else.} \end{cases} \quad (6.39)$$

where the subscript i has been dropped for convenience. Again, for simplicity, assume that amplitude and phase are independent and that phase is distributed symmetrically about zero. Then the first and second moments of the $\cos \phi$ and $\sin \phi$ terms needed for Equation 6.28 can be computed easily (also given in [67]). For deterministic

N and deterministic, unity A , these moments demonstrate the contribution of the phase density to the real and imaginary components. Analytic expressions for the moments are given in Table 6.3 and plotted in Figure 6.1 as a function of the density width, a . At $a = 0$, or constant phase, the mean value is one and the variances zero, as expected. For $a \in [0, \pi]$, as a approaches π , the mean value decreases and the variances increase to limits that correspond to the Rayleigh density. The Rayleigh result is produced by phase uniformly distributed on the interval $[-\pi, +\pi]$. Since any interval $[-n\pi, +n\pi]$, n an integer, is equivalent to $[-\pi, +\pi]$, the Rayleigh limit is also seen for any integer a . As n increases, the deviation from uniform phase decreases and the deviation from the Rayleigh limits for the mean and variances also decreases, as can be seen from the figure.

Component	Expression
$E(\cos\phi)$	$\text{sinc}a$
$E(\sin\phi)$	0
$E(\cos^2\phi)$	$\frac{1}{2}(1 + \text{sinc}2a)$
$E(\sin^2\phi)$	$\frac{1}{2}(1 - \text{sinc}2a)$
$\sigma_{\cos\phi}^2$	$\frac{1}{2}(1 + \text{sinc}2a - 2\text{sinc}^2a)$
$\sigma_{\sin\phi}^2$	$\frac{1}{2}(1 - \text{sinc}2a)$

Table 6.1: First and second moments of real and imaginary components for $\phi \sim U[-\pi, \pi]$ and unity A and N .

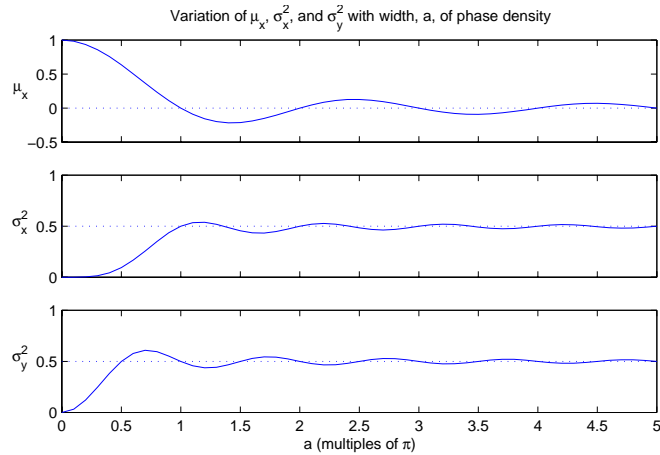


Figure 6.1: Plots of mean and variances with variation of phase density width, a (deterministic, unity N and A). The statistics are Rayleigh (mean and variances at dotted lines) at integral multiples of π and converge to the same as a increases.

When the phase density is uniform over an integral multiple of π , $E(\cos\phi)$ is zero, eliminating those terms including it. As a result, the mean of the real component is zero, and the variances of the imaginary and real components are simplified and equal and depend only on the expected value of N and the second moment of the amplitude,

$$\text{For } \phi \sim U[-\pi, \pi], \quad \sigma_x^2 = \sigma_y^2 = \frac{1}{2}E(N)E(A_i^2). \quad (6.40)$$

In other words, as the contribution of the phase density becomes negligible, the statistics become Rayleigh. Likewise, assumption of the Rayleigh density implies that the phase density does not contribute. The simplicity of the Rayleigh density makes it attractive for constructing an image model. In cases where the phase density is important, e.g., when echoes are coherent, the Rayleigh assumptions do not apply. Because the coherent echoes are an important part of the basic image features, there is a strong incentive for developing a procedure for classifying pixels as Rayleigh or non-Rayleigh.

6.4 Computing the Amplitude Mean and Variance

The amplitude of the complex sum is still the quantity of interest in characterizing commercial images. As discussed in Chapter 4, much research has focused on deriving probability densities for the amplitude. Many of those results require assumptions that will be too strict to meet the needs of the image model. For instance, the commonly used combination of Rayleigh and Rician densities [11] cannot even completely model the variation in the simple uniform phase distribution example of the previous section. Others, e.g., the k-distribution and generalized-k, although sufficiently general, would have impractical computational requirements for this image model and also produce a level of detail that is of little utility here. In this work, accuracy in the form of the pixel densities is weighed against accuracy in describing the *differences* in pixel densities. Of special importance in this work is how the densities vary when the shape changes, i.e., when the likelihood of a different transformation of the shape is being assessed. With this in mind, the mean and variance of the amplitude are reasonable choices for a first approximation.

The amplitude of the complex sum is given by $\rho = \sqrt{x^2 + y^2}$. The amplitude mean, μ_ρ , is computed as an expectation over the complex Gaussian density,

$$\mu_\rho = E\left(\sqrt{x^2 + y^2}\right) \quad (6.41)$$

$$= \iint \sqrt{x^2 + y^2} p_{x,y}(x, y) dx dy \quad (6.42)$$

$$= \iint \sqrt{x^2 + y^2} \frac{e^{-\frac{1}{2(1-r^2)} \left[\frac{(x-\mu_x)^2}{\sigma_x^2} - 2r \frac{(x-\mu_x)(y-\mu_y)}{\sigma_x \sigma_y} + \frac{(y-\mu_y)^2}{\sigma_y^2} \right]}}{2\pi \sigma_x \sigma_y \sqrt{1-r^2}} dx dy. \quad (6.43)$$

In general, no closed form solution exists for this integral. In computations for the image model, the statistics have been assumed Rayleigh whenever applicable, and, otherwise, Simpson's rule [80] has been used to compute the integral 6.43 numerically.

In contrast to the mean, the variance is quite simple to calculate. Starting with the standard expression for the variance, a relationship can be derived in terms of the second moments of the real and imaginary components and the amplitude mean,

$$\sigma_\rho^2 = E(\rho^2) - E^2(\rho) \quad (6.44)$$

$$= E(x^2) + E(y^2) - \mu_\rho^2. \quad (6.45)$$

The second moment of the real component, $E(x^2)$, was given in Equation 6.25. This quantity, and $E(y^2)$, the second moment of the imaginary component, are easy to compute given the necessary moments for N , $A_i \cos \phi_i$ and $A_i \sin \phi_i$.

For many cases, the statistics will satisfy the Rayleigh density, parameterized by a single constant α ,

$$p_\rho(\rho) = \frac{\rho}{\alpha^2} e^{-\frac{\rho^2}{2\alpha^2}}. \quad (6.46)$$

The mean of a Rayleigh density is well-known [78] and can easily be derived from this development [67],

$$\mu_\rho = \alpha \sqrt{\frac{\pi}{2}} = \frac{1}{2} \sqrt{\pi E(N) E(A_i^2)}. \quad (6.47)$$

The variance for the Rayleigh-distributed random variable is given by

$$\sigma_\rho^2 = \left(2 - \frac{\pi}{2}\right) \alpha^2 = \left(1 - \frac{\pi}{4}\right) E(N) E(A_i^2). \quad (6.48)$$

The simplicity of the relations between the Rayleigh mean and variance and the moments of N and A_i makes the Rayleigh density an attractive choice for the image model. The difficulty is in classifying a given pixel as Rayleigh or non-Rayleigh. In this regard, an important quantity for a Rayleigh-distributed random variable is the ratio of the mean to the standard deviation, often termed the SNR_0 [11], which has a constant value of 1.91,

$$\text{SNR}_0 = \frac{\mu_r}{\sigma_r} = 1.91. \quad (6.49)$$

Because this ratio is a constant, it can be used to classify a pixel as Rayleigh or non-Rayleigh. Specifically, if the amplitude mean and variance are computed (as in the expressions of the previous sections), the SNR_0 can be computed, and the result can be used in a decision to approximate a pixel as Rayleigh or not. If the result could be predicted using some more fundamental descriptor of the scattering elements, i.e., the axial extent of the distribution of scatterers for a given resolution cell, computation could be simplified since some parameters of the complex Gaussian could be neglected and the amplitude mean and variance could be computed from the previous expressions for the Rayleigh density. Such an approach will be considered in the next section and subsequent chapters, although it will be found to have limited application in some cases of surface geometry and system characteristics.

Example: Sensitivity of the Amplitude to the Phase Density

The previous example showed the sensitivity of the complex Gaussian parameters to the width of a uniform phase density. It is worth examining the effects on the amplitude statistics as well. In this example, the amplitudes A_i were deterministic with unity amplitude, and phases, ϕ_i , were uniformly distributed on $[-a, +a]$. Amplitude mean and variance were then computed for a varying phase-density width, a , for four different values of N , as shown in Figure 6.2.

The figure contains several important details that lend insight to the rest of this investigation. First, the mean decreases from a maximum of $E(N)$ at $a = 0$ to the Rayleigh limit at $a = \pi$, with deviation from the Rayleigh limit as seen previously for the mean and variances of the complex sum. Note, however, that as N is increased, deviation from the Rayleigh result increases, and the phase-density width required for convergence to the Rayleigh result also increases. This result is an indication of the high sensitivity of the amplitude to the phase interactions, and it is especially




Figure 6.2: Amplitude Mean, Variance and SNR_0 vs. width of uniform phase density with N deterministic at various values. Again, the statistics are Rayleigh (mean and variances at dotted lines) at integral multiples of π and converge to the same as a increases. For higher values of N , however, convergence is slower and variations from the Rayleigh result are larger.

important in the context of using the width of the phase density to predict whether or not a pixel can be modeled as Rayleigh-distributed.

6.5 Relevant Issues for the Image Model

The relations derived here provide methods for computing the amplitude mean and variance from pixel-based densities for N , A_i and ϕ_i , the elements of the random phasor sum. Extension to the surface description of Chapter 5 will entail computing these densities from local surface characteristics. Such computations will require approximations, and the assumptions that are made will be significant. Assumptions typically made in the random-phasor-sum characterization of scattering will prove to be too strict. The Rayleigh density and associated assumptions will play a special role because of their simplicity. The Rayleigh result holds when phase can be disregarded, but predicting when phase can be disregarded is not trivial. The SNR_0 ratio for a Rayleigh density will prove to be a useful indicator of when to classify a pixel as Rayleigh, but the ratio still requires substantial computation. In the end,

the assumptions and approximations will require a tradeoff between computational requirements and accuracy.

Chapter 7

Image Model Statistics for Surfaces: Theory

The previous chapter provided the basis for an image model based on the mean and variance of the amplitude. Extension of those results to surfaces requires techniques for computing the parameters of the complex Gaussian for arbitrary surfaces. For surfaces, the random phasor sum depends locally on the intersection between the surface and a 3D resolution cell representing the envelope of the PSF. Because of the complexity of those interactions, approximations must be used. In the approximations, accuracy and computational requirements are competing objectives. In this chapter, two approaches are outlined: one where the computations are made directly from the triangulated representation for the surface, and one where the surface is locally approximated as planar. The direct approach provides better accuracy, but the locally planar approximation is more efficient computationally and produces intuitive relationships between the characteristics of the surface and imaging system.

7.1 The Random Phasor Sum and Tissue Surfaces

Because the amplitude mean and variance can be computed from the parameters for the complex Gaussian, these parameters form the subject of interest in characterizing images of surfaces. Recall from Equations 6.34 to 6.38 that the parameters are found in terms of the local, pixel-based moments for N , the number of scatterers, and the products of amplitude, A_i , and phase, ϕ_i , $A_i \cos \phi_i$, and $A_i \sin \phi_i$. The equations are

listed again here for reference,

$$\mu_x = E(N)E(A_i \cos \phi_i), \quad (7.1)$$

$$\mu_y = E(N)E(A_i \sin \phi_i), \quad (7.2)$$

$$\sigma_x^2 = E(N)E(A_i^2 \cos^2 \phi_i) + \sigma_N^2 E^2(A_i \cos \phi_i), \quad (7.3)$$

$$\sigma_y^2 = E(N)E(A_i^2 \sin^2 \phi_i) + \sigma_N^2 E^2(A_i \sin \phi_i), \quad (7.4)$$

$$r = \frac{1}{\sigma_x \sigma_y} \left[E(N) \left[E(A_i^2 \cos \phi_i \sin \phi_i) - E(A_i \cos \phi_i) E(A_i \sin \phi_i) \right] + \sigma_N^2 E(A_i \cos \phi_i) E(A_i \sin \phi_i) \right]. \quad (7.5)$$

All of the moment computations require surface integrals, where the region of interest is the local intersection of the surface and the PSF resolution cell. The resolution cell is simply a region around the pixel location in which scatterer contributions are included. For the 3D Gaussian envelope used in this dissertation, the associated resolution cell is an ellipsoidal volume with the lengths of the principal axes defined by the widths of the Gaussian envelope. The resolution cell is described formally in a later section. The item of importance here is the intersection surface, S_\cap , defined as the intersection of the gross tissue surface and the local resolution cell. Note that technically the intersection of surface and volume could also be a point, but this case is not of interest and is not considered.

The moments of N and $A_i \cos \phi_i$ require different computational approaches. Depending on the model used for the number of scatterers, N , the moments for N will depend in some way on the area of the intersection surface, S_\cap ,

$$\text{Area}(S_\cap) = \iint_{S_\cap} dA. \quad (7.6)$$

For the rest of this dissertation, N is modeled as deterministic primarily for the sake of simplicity. With the surface microstructure parametrized in terms of the scatterer concentration, the moments for N will depend only on the product of the scatterer concentration and the intersection area. For any realistic model of scatterer distribution, the area of intersection would be an important quantity. For example, for N random with Poisson density, the density parameter would also be the product of the area of intersection and scatterer concentration [81]. The differences between N deterministic and N Poisson are in the variance and second moment as described in

the previous chapter. As shown in later sections, the area of intersection is computed in different ways for the different approximations.

The first and second moments for the necessary functions of A_i and ϕ_i require additional computations to the area of intersection. Recall that the scatterer amplitude and phase are both functions of position and that the position is assumed to be uniformly distributed on the surface. Expected values for functions of that position can, thus, be computed with respect to a density for scatterer position. For scatterers uniformly distributed over the surface, the density on the scatterer position, \mathbf{r} , is simply the reciprocal of the area of intersection,

$$p_{\mathbf{r}}(\mathbf{r}) = \frac{1}{\int_{S_{\cap}} dA} \quad \text{for } \mathbf{r} \in S_{\cap}. \quad (7.7)$$

The expectation of $A_i \cos \phi_i$ is then given by a surface integral defined in terms of position, \mathbf{r} ,

$$E(A_i \cos \phi_i) = \iint_{S_{\cap}} A_i(\mathbf{r}) \cos \phi_i(\mathbf{r}) p_{\mathbf{r}}(\mathbf{r}) dA \quad (7.8)$$

$$= \frac{1}{\text{Area}(S_{\cap})} \iint_{S_{\cap}} A_i(\mathbf{r}) \cos \phi_i(\mathbf{r}) dA. \quad (7.9)$$

Other moments for products of the amplitude and phase require similar surface integrals because the moments are all functions of position.

In general, computation of the five parameters of the complex Gaussian involves solving various forms of these two surface integrals. The two approaches given here for computing those parameters differ in the way the integrals are approximated and computed and, subsequently, in the accuracy that is achieved.

7.2 Computing Moments Directly from the Triangulation

Of the methods developed in this chapter for computing the required moments, the most straightforward and potentially most accurate employs approximations to the integrals based on the triangles of the surface representation. The area of intersection is approximated quite simply as a sum over those triangles in the resolution cell,

$$\text{Area}(S_{\cap}) = \iint_{S_{\cap}} dA \approx \sum_{\Delta_i \in S_{\cap}} \text{Area}(\Delta_i). \quad (7.10)$$

The expectation of the amplitude and cosine of the phase is given as follows with the associated numerical approximation,

$$E(A_i \cos \phi_i) = \iint_{S_\cap} A_i(\mathbf{r}) \cos \phi_i(\mathbf{r}) p_{\mathbf{r}}(\mathbf{r}) dA \quad (7.11)$$

$$\approx \frac{1}{\text{Area}(S_\cap)} \sum_{\Delta_i \in S_\cap} \iint_{\Delta_i} A_i(\mathbf{r}) \cos \phi_i(\mathbf{r}) dA. \quad (7.12)$$

If the amplitude is approximated as constant over the triangle with the value at the triangle midpoint, $A_i(\mathbf{r}_{\Delta_i})$, the computations are simplified further because the integral of the cosine of the phase can be calculated analytically,

$$E(A_i \cos \phi_i) \approx \frac{1}{\text{Area}(S_\cap)} \sum_{\Delta_i \in S_\cap} A_i(\mathbf{r}_{\Delta_i}) \iint_{\Delta_i} \cos \phi_i(\mathbf{r}) dA. \quad (7.13)$$

The integral of the cosine of the phase can be calculated analytically for the triangle as shown in Appendix A. The remaining expectations of functions of A_i and ϕ_i can also be computed with analytic results for integrals of the trigonometric functions as in Appendix A.

The phase contribution is the most sensitive part of the computations, but because it can be computed analytically, the accuracy in using the direct computation depends on approximating the amplitude as constant over a triangle and approximating the intersection surface by triangles from the original surface. As shown in Appendix A, triangles can be resampled to any desired area, thus the approximations involved, i.e., approximating the amplitude as constant over each triangle and approximating the intersection surface by triangles, can be made as accurate as desired at the expense of increased computation.

Given the parameters of the complex Gaussian, the means, variances and correlation coefficients, the amplitude mean can be computed by integrating numerically as in Chapter 6. The main advantages of this method over the those of the next two sections is that the full surface representation is used and that no assumptions are made regarding the independence of amplitude and phase. The disadvantages are the computation required to perform the integrals and the lack of intuition developed from inspection of the solutions.

7.3 Computing Moments Using a Locally Planar Approximation to the Surface

The remaining methods for computing the required moments are based on an approximation to the local surface geometry. By approximating the surface as locally planar, computation can be simplified and, more importantly, the results provide an intuitive, qualitative description of scattering, i.e., the angle between the axial direction and the local surface normal determines whether the displayed echoes are coherent or incoherent.

As in the previous section, the integrals of Equations 7.6 and 7.8 are required. In general, the integral of any function $f(\mathbf{r})$ defined over the surface with parametrization, $\mathbf{r}(u, v)$, can be computed as follows [80],

$$\iint_S f(\mathbf{r})dA = \iint_{R_0} f(\mathbf{r}(u, v)) |\mathbf{r}_u \times \mathbf{r}_v| dudv \quad (7.14)$$

where \times denotes the cross product, \mathbf{r}_u and \mathbf{r}_v denote the partial derivatives of the transformation with respect to u and v , and R_0 is the subset of \mathbb{R}^2 that is mapped onto the surface S , by $\mathbf{r}(u, v)$. See Chapter 3 for a review of parametric representations for surfaces. In this case, the parametrization will define the intersection surface in terms of the planar approximation and the resolution cell ellipsoid. The parametrization used in this chapter is derived in Appendix B.

The standard definition for the ellipsoid [80] surface uses the principal axes a, b and c to define the extent of the ellipsoid in the x, y and z dimensions, respectively,

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1. \quad (7.15)$$

The numbers a, b and c represent multiples of the PSF widths, σ_x, σ_y and σ_z in each of the dimensions. An equivalent matrix representation is particularly useful in the calculations of the next two sections. For more details on a linear algebraic approach to elementary geometry, see [82]. For the ellipsoid, the matrix representation is defined in terms of a matrix, \mathbf{M} , and vector, $\mathbf{x} \in \mathbb{R}^3$,

$$\mathbf{x}^t \mathbf{M}^2 \mathbf{x} = 1 \quad (7.16)$$

where

$$\mathbf{M} = \begin{bmatrix} \frac{1}{a} & 0 & 0 \\ 0 & \frac{1}{b} & 0 \\ 0 & 0 & \frac{1}{c} \end{bmatrix} \quad \text{and} \quad \mathbf{M}^{-1} = \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & c \end{bmatrix}. \quad (7.17)$$

In the following sections, the plane is described in matrix notation by a normal, $\mathbf{N} = [N_x \ N_y \ N_z]$, and distance, d , to the origin,

$$\mathbf{N}\mathbf{x} = d \quad (7.18)$$

where \mathbf{N} is assumed to be a unit normal, i.e., $|\mathbf{N}| = 1$. The description is the same as that in Appendix B.

7.3.1 Area of intersection

The area of intersection is computed as the integral over the intersection surface, S_\cap , with $f(\mathbf{r}(u, v)) = 1$. Using the parametrization of equation B.29, the region R_0 is the circle $u^2 + v^2 \leq 1 - d'^2$ for $d' < 1$, with area of the circle simply $\pi(1 - d'^2)$. For $d' \geq 1$, the intersection is a point with equality and empty otherwise. In the case of a circle intersection, the integral can be computed quite simply as the product of the cross product term and the area of the circle in the \mathbf{x}'' coordinate system,

$$\iint_S dA = \iint_{R_0} |\mathbf{r}_u \times \mathbf{r}_v| dA_0 \quad (7.19)$$

$$= |\mathbf{r}_u \times \mathbf{r}_v| \iint_{R_0} dA_0 \quad (7.20)$$

$$= \sqrt{b^2 c^2 N_x'^2 + a^2 c^2 N_y'^2 + a^2 b^2 N_z'^2} \pi(1 - d'^2). \quad (7.21)$$

Note that this result is in the parameters for the plane of the \mathbf{x}' coordinate system. Alternatively, the result can be expressed in the parameters of the original coordinate system after substitution and some simplification,

$$\iint_S dA = \pi \left(1 - \frac{d^2}{a^2 N_x^2 + b^2 N_y^2 + c^2 N_z^2} \right) \frac{abc}{\sqrt{a^2 N_x^2 + b^2 N_y^2 + c^2 N_z^2}}. \quad (7.22)$$

The area of intersection for any plane with any ellipsoid, thus, depends only on the parameters for the ellipsoid, a, b, c , and the plane, $[\mathbf{N}, d]$. As expected, this

result is symmetric in a, b, c and N_x, N_y, N_z , even though the transformation used to derive the result was not.

7.3.2 Moments of functions of A_i, ϕ_i

The other required expectations are not so simple to derive. Consider the product of A_i and $\cos \phi_i$. The required integral is that of Equation 7.8,

$$E(A_i \cos \phi_i) = \frac{1}{\text{Area}(S_\Omega)} \iint_{S_\Omega} A_i(\mathbf{r}) \cos \phi_i(\mathbf{r}) dA. \quad (7.23)$$

The amplitude $A_i(\mathbf{r})$ can be represented as an exponential with a matrix Σ of variances representing the widths of the PSF envelope,

$$A_i(\mathbf{r}) = e^{-\mathbf{r}^t \Sigma \mathbf{r}} \quad (7.24)$$

where

$$\Sigma = \begin{bmatrix} \frac{1}{\sigma_x^2} & 0 & 0 \\ 0 & \frac{1}{\sigma_y^2} & 0 \\ 0 & 0 & \frac{1}{\sigma_z^2} \end{bmatrix}. \quad (7.25)$$

The amplitude function can be written in terms of the matrix \mathbf{M} by choosing the ellipsoid axes to be constant multiples, α , of the PSF widths,

$$\mathbf{M} = \begin{bmatrix} \frac{1}{a} & 0 & 0 \\ 0 & \frac{1}{b} & 0 \\ 0 & 0 & \frac{1}{c} \end{bmatrix} = \frac{1}{\alpha} \Sigma^{\frac{1}{2}} = \begin{bmatrix} \frac{1}{a\sigma_x} & 0 & 0 \\ 0 & \frac{1}{b\sigma_y} & 0 \\ 0 & 0 & \frac{1}{c\sigma_z} \end{bmatrix}. \quad (7.26)$$

In terms of \mathbf{M} , the amplitude equation can then be written

$$A_i(\mathbf{r}(u, v)) = e^{-\alpha^2 \mathbf{r}^t \mathbf{M}^2 \mathbf{r}}. \quad (7.27)$$

When substituted into the amplitude equation, this expression simplifies the exponent, where $\mathbf{r}(u, v) = \mathbf{M}^{-1}\mathbf{R}\mathbf{u}$,

$$A(\mathbf{r}(u, v)) = e^{-\alpha^2 \mathbf{u}^t \mathbf{R}^t \mathbf{M}^{-1} \mathbf{M}^2 \mathbf{M}^{-1} \mathbf{R} \mathbf{u}} \quad (7.28)$$

$$= e^{-\alpha^2 \mathbf{u}^t \mathbf{R}^t \mathbf{R} \mathbf{u}} \quad (7.29)$$

$$= e^{-\alpha^2 (u^2 + v^2 + d'^2)}. \quad (7.30)$$

The integral can then be simplified as follows, where, from previous chapters, the phase is given by $\phi_i = 2k_0 z$,

$$E(A_i \cos \phi_i) = \frac{1}{\text{Area}(S_\cap)} \iint_{S_\cap} A_i(\mathbf{r}) \cos \phi_i(\mathbf{r}) dA \quad (7.31)$$

$$= \frac{1}{|\mathbf{r}_u \times \mathbf{r}_v| \pi (1 - d'^2)} \times \iint_{R_0} e^{-\alpha^2 (u^2 + v^2 + d'^2)} \cos(2k_0 z(u, v)) |\mathbf{r}_u \times \mathbf{r}_v| dA_0 \quad (7.32)$$

$$= \frac{1}{\pi (1 - d'^2)} \iint_{R_0} e^{-\alpha^2 (u^2 + v^2 + d'^2)} \cos(2k_0 z(u, v)) dA_0 \quad (7.33)$$

$$= \frac{e^{-\alpha^2 d'^2}}{\pi (1 - d'^2)} \iint_{R_0} e^{-\alpha^2 (u^2 + v^2)} \cos(2k_0 z(u, v)) dA_0 \quad (7.34)$$

$$(7.35)$$

where $z(u, v)$ is simply the z component of the parametrization, $\mathbf{r}(u, v)$. No simple analytical expression is known for this integral, thus it has been computed numerically for the results of the next chapter. The other required expectations are similar and have also been computed numerically.

The locally planar approximation provides simple results that can build intuition regarding the interaction of the PSF and tissue surface. The results of the next chapter show accuracy limitations to the locally planar approximation, specifically in cases where the surface curvature is high relative to a wavelength. In the meantime, the results of this section still require numerical integration. In the next section, the locally planar approximation is simplified further by removing assumptions of independent amplitude and phase. Analytical results are derived for the parameters of the complex Gaussian, providing a computationally efficient approach to finding the amplitude mean and variance.

7.4 Computing Moments with the Planar Approximation and $A_i \perp \phi_i$

With independent amplitude and phase, the expectation of the amplitude is separated from that of the phase, e.g., $E(A_i \cos \phi_i) \underset{A_i \perp \phi_i}{=} E(A_i)E(\cos \phi_i)$. The first and second moments for the amplitude will be derived for the Gaussian envelope. The phase will be assumed uniform, and the axial extent of the intersection surface will provide the width of the associated uniform phase density.

7.4.1 Moments for scatterer amplitude

The first moment for A_i can be derived as before for $A_i \cos \phi_i$, where the phase term is now absent,

$$E(A_i) = \frac{1}{\text{Area}(S_\Gamma)} \iint_{S_\Gamma} A_i(\mathbf{r}) dA \quad (7.36)$$

$$= \frac{1}{|\mathbf{r}_u \times \mathbf{r}_v| \pi (1 - d'^2)} \iint_{R_0} e^{-\alpha^2(u^2+v^2+d'^2)} |\mathbf{r}_u \times \mathbf{r}_v| dA_0 \quad (7.37)$$

$$= \frac{e^{-\alpha^2 d'^2}}{\pi (1 - d'^2)} \iint_{R_0} e^{-\alpha^2(u^2+v^2)} dA_0 \quad (7.38)$$

$$= \frac{e^{-\alpha^2 d'^2}}{\pi (1 - d'^2)} \left[2\pi \int_0^{\sqrt{1-d'^2}} r e^{-\alpha^2 r^2} dr \right] \quad (7.39)$$

$$= \frac{e^{-\alpha^2 d'^2}}{\pi (1 - d'^2)} \left[-\frac{\pi}{\alpha^2} e^{-\alpha^2 r^2} \Big|_0^{\sqrt{1-d'^2}} \right] \quad (7.40)$$

$$= \frac{e^{-\alpha^2 d'^2}}{\alpha^2 (1 - d'^2)} \left[1 - e^{-\alpha^2 (1-d'^2)} \right]. \quad (7.41)$$

The second moment for the individual phasor amplitude can be computed in a similar manner, (computations are exactly the same with α replaced by $\sqrt{2}\alpha$),

$$E(A^2) = \frac{e^{-2\alpha^2 d'^2}}{2\alpha^2 (1 - d'^2)} \left[1 - e^{-2\alpha^2 (1-d'^2)} \right]. \quad (7.42)$$

Note that both moments depend only on the scaling parameter, α , which determines the extent of the resolution cell as a multiple of the σ widths of the PSF, and d' , the distance of the plane to the origin adjusted in terms of the parameters of the ellipsoid as in B.9.

7.4.2 Axial extent of intersection

For this approximation assuming independent amplitude and phase, the axial extent of the intersection surface will give a distance that will be used to define a uniform phase density, allowing the simple results of Table 6.3 to be used in computing the required moments for $\cos \phi_i$ and $\sin \phi_i$. For the axial extent of intersection, the maximum and minimum z values of the intersection surface are desired. These points must lie on the boundary of the surface, thus the problem can be posed as follows, as an optimization over the circle $u^2 + v^2 = 1 - d'^2$ of the z coordinate of the surface parametrization in (u, v) :

$$\max(\min)_{u^2+v^2=1-d'^2} z(u, v) = \begin{bmatrix} -\frac{cN'_x N'_z}{\sqrt{N'_y{}^2 + N'_z{}^2}} & -\frac{cN'_y}{\sqrt{N'_y{}^2 + N'_z{}^2}} & cN'_z \end{bmatrix} \begin{bmatrix} u \\ v \\ d' \end{bmatrix}. \quad (7.43)$$

The problem can be solved by parametrizing the circle in t as

$$(u(t), v(t)) = \left(\sqrt{1 - d'^2} \cos t, \sqrt{1 - d'^2} \sin t \right), \quad (7.44)$$

differentiating $z(u, v) = z(t)$ with respect to t , setting the result to 0 and solving,

$$\frac{\partial z}{\partial t} = -c\sqrt{1 - d'^2} \left[-\frac{N'_x N'_z}{\sqrt{N'_y{}^2 + N'_z{}^2}} \sin t + \frac{N'_y}{\sqrt{N'_y{}^2 + N'_z{}^2}} \cos t \right] \quad (7.45)$$

$$\frac{\partial z}{\partial t} = 0 \Rightarrow -\frac{N'_x N'_z}{\sqrt{N'_y{}^2 + N'_z{}^2}} \sin t + \frac{N'_y}{\sqrt{N'_y{}^2 + N'_z{}^2}} \cos t = 0 \quad (7.46)$$

$$\Rightarrow \tan t = \frac{N'_y}{N'_x N'_z}. \quad (7.47)$$

These equations are satisfied for

$$(u, v) = (\cos t, \sin t) = \left(\pm \frac{N'_x N'_z}{\sqrt{N'_y{}^2 + N'_z{}^2 N'_x{}^2}}, \pm \frac{N'_y}{\sqrt{N'_y{}^2 + N'_z{}^2 N'_x{}^2}} \right) \quad (7.48)$$

with corresponding values of z ,

$$z = c \left(N'_z d' \pm \sqrt{\frac{(1 - d'^2)(N'_y{}^2 + N'_x{}^2 N'_z{}^2)}{(N'_y{}^2 + N'_z{}^2)}} \right). \quad (7.49)$$

Of course, the real value of interest is the axial extent, w_z , or the difference between the max and min values for z ,

$$w_z = c \sqrt{\frac{(1 - d'^2)(N_y'^2 + N_x'^2 N_z'^2)}{(N_y'^2 + N_z'^2)}}. \quad (7.50)$$

7.5 Surface Roughness

The methods of this chapter were developed without including the surface roughness parameter. This omission was made to speed up development and allow sufficient time and effort for an inference investigation. In the methods requiring numerical integration, the surface integrals would become volume integrals where the volume includes an extension of the surface along the surface normal. For the planar approximation where amplitude and phase are assumed independent, the effects of roughness can be modeled by simply adding the roughness width to the axial extent, w_z , of Equation 7.50. With this approach, the effects of roughness for planar surfaces will be investigated briefly in the next chapter. The roughness will only affect the coherent scattering results, and even that effect will be small for roughness values of interest, i.e., small relative to a wavelength. Recall that the discrete-scatterer model inherently implies a roughness to the surface. Modeling the physical perturbation, if small, may be unnecessary in many applications.

7.6 Locally Planar Approximations for Arbitrary Surfaces

For arbitrary surfaces, the planar approximation methods for computing the amplitude statistics require a preliminary step of constructing a planar approximation for each pixel. Ideally, a plane would be chosen to minimize some error in the amplitude statistics, but the associated expressions are too complex to permit an algorithm with practical computational requirements. To be practical, an approach must balance accuracy and computational efficiency. Recall that the triangulated surface consists of a set of vertices and triangles. At the least, the algorithm should compute planar approximations for all pixels in a single pass through the triangulation since computation would be prohibitive if, for each pixel, every triangle were visited.

For this work, a simple method has been developed. Each triangle inherently defines a normal, and a normal can easily be estimated at any vertex. From this information, an “average” tangent plane can be computed for any local part of the tissue surface, e.g., that section of the surface interior to the local resolution cell. The plane for any pixel i is approximated by computing, first, an amplitude-weighted average of the normal, \mathbf{n}_i , over the local surface region, and, second, an amplitude-weighted average distance, d_i , over the surface region assuming \mathbf{n}_i as the normal to the plane. The amplitude weighting reflects the influence of the Gaussian PSF envelope and smooths the resulting planar approximations over neighboring pixels.

The local surface is denoted S_{\cap_i} and is given by the intersection of the surface, S , and the volume, V_i , of the resolution cell at sample location $\mathbf{x}_i \in \mathbb{R}^3$. The amplitude-weighted average, \mathbf{n}_i , is given by

$$\mathbf{n}_i = \frac{1}{\iint_{S_{\cap_i}} A_i(\mathbf{r}) dA} \iint_{S_{\cap_i}} A_i(\mathbf{r}) \mathbf{n}(\mathbf{r}) dA \quad (7.51)$$

where the amplitude is given by the amplitude of the PSF envelope.

The approximating plane is completed with the distance, d_i , along the average normal, \mathbf{n}_i , from the center of the resolution cell to the plane. For any point, \mathbf{r} , the distance, $d_{\mathbf{r}}$, from the resolution cell center to a plane passing through that point and having normal \mathbf{n}_i is given by $\mathbf{n}_i \cdot \mathbf{r}$. Again using an amplitude-weighted average, d_i is given by

$$d_i = \frac{1}{\iint_{S_{\cap_i}} A_i(\mathbf{r}) dA} \iint_{S_{\cap_i}} A_i(\mathbf{r}) d_{\mathbf{r}} dA. \quad (7.52)$$

In computing these terms from a triangulated mesh for a Gaussian PSF envelope, the integration must be performed numerically. The local surface of intersection must be determined, and the integration must be computed over discrete portions of the surface. Local surfaces of intersection can be determined in terms of collections of triangles that are “in” the local resolution cell volume. In this work, triangles are included in computations if *any* vertex falls in the resolution cell.

The integrals have been computed numerically as follows as sums over triangles where \mathbf{r}_i is the midpoint of triangle i :

$$\mathbf{n}_i \approx \frac{1}{\sum_{\Delta_i \in S_{\cap_i}} \text{Area}(\Delta_i)} \sum_{\Delta_i \in S_{\cap_i}} A_i(\mathbf{r}_i) \mathbf{n}(\mathbf{r}_i) \text{Area}(\Delta_i) \quad (7.53)$$

$$d_i \approx \frac{1}{\sum_{\Delta_i \in S_{n_i}} \text{Area}(\Delta_i)} \sum_{\Delta_i \in S_{n_i}} A_i(\mathbf{r}_i) d_{\mathbf{r}_i} \text{Area}(\Delta_i). \quad (7.54)$$

Computation of the locally approximating plane will prove to be a difficulty for arbitrary surfaces in the next chapter. The utility of this simple method is that it has provided an initial approach for investigating the planar surface approximation in computing the amplitude mean and variance. Possible refinements of the method are suggested in the results chapter, but implementation has been left for future work.

Chapter 8

Image Model Statistics for Surfaces: Results

For this chapter, the methods of the previous section were investigated for their potential in computing pixel-based amplitude statistics for surfaces. Since the methods all employ approximations, evaluation is based on comparison of the computed statistics with statistics generated from repeated simulation. The first section comprises results obtained using the locally planar approximation to the surface. These results confirm the observation that statistics depend on the angle of incidence and expose some of the limitations of the planar approximation. The second section contains results comparing the locally planar approximation to the direct computation method for arbitrary image planes of the cadaveric vertebra. Throughout the chapter, the following labels are used for the associated results: *planar* denotes statistics generated using the locally planar approximation, *triangle* denotes statistics generated directly from the triangulation, and *simulation* denotes statistics generated from repeated simulations. The chapter concludes with a discussion on the implications for inference using an image model derived from the statistics.

8.1 General Results from the Planar Surface Approximation

The planar approximation methods permit efficient investigation of basic effects of the PSF, surface and plane parameters on the amplitude statistics. These effects are demonstrated in the next part of this section. The subsequent part of this section

shows limitations of using the planar approximation for a curved surface, a sphere, and introduces entire images of pixel-based statistics computed from a shape with associated microstructure and given system characteristics.

8.1.1 Effects of PSF, surface and plane parameters

Many factors of the physical model affect the amplitude statistics. In this section, results are shown that demonstrate effects of the following changes: 1) rotation of the plane about the y axis, 2) rotation of the plane about the x axis, 3) rotation about the y axis and translation away from the center of the resolution cell, 4) scatterer concentration for the planar surface, and 5) surface roughness.

Statistics were generated according to the planar approximation results of the previous chapter. The assumption of independent amplitude and phase was varied in the different investigations as indicated in each result. Simulation results were generated using the methods of Chapter 5 with scatterers distributed uniformly within the plane according to the chosen concentration and roughness. Unless otherwise specified, concentration was 64 scatterers/mm², and roughness was assumed uniform over a range of 0.1 wavelengths. The results for rotation are shown for positive and negative angles to indicate the symmetry of variation over this range, although values were computed only for the positive angles (the positive results are displayed for the negative angles as well). For the simulation results, 500 trials were averaged. In all trials, the PSF values of Chapter 5, $\sigma_x = 1.5$, $\sigma_y = 0.5$ and $\sigma_z = 0.2$ mm, were used unless otherwise specified. The scaling parameter, $\alpha = 2$, defined the size of the resolution cell ellipsoid in terms of the PSF widths and was chosen by hand based on plots of the exponential amplitude function. This choice was then justified by the success of the results.

In Figure 8.1, the planar and simulation mean, standard deviation, SNR_0 and the axial extent of intersection are shown versus rotation of the plane about the y axis. As expected, mean values are greatest near 0 degrees (plane oriented normal to axial direction) where the axial distribution of scatterers is small. For planar results without the $A \perp \phi$ assumption, and for the simulation results, the mean falls off smoothly with increase in angle from the maximum to a Rayleigh value ($\text{SNR}_0 = 1.91$). The transition between non-Rayleigh and Rayleigh regions occurs at an axial extent of approximately two wavelengths. Note that prior to that transition, in the non-Rayleigh region, the planar result assuming $A \perp \phi$ includes deviations from the

other results. These deviations are evidence of the importance and sensitivity of the phase contributions to observed features in ultrasonic images. In actual images, pixel intensities vary smoothly over coherent regions, as opposed to the oscillations generated with amplitude assumed independent of phase.




Figure 8.1: Effects of rotating the plane about the y axis for full range (left) and narrow range (right). Mean, standard deviation, and SNR_0 are shown versus rotation around the y axis. Axial extent of intersection is also shown for reference. Simulation and planar results are included, where the planar results were obtained with and without assuming amplitude and phase independent.

Figure 8.2 shows statistics versus rotation of the plane around the x axis. The main difference from the plots versus rotation around the y axis is the width of the coherent region, or the transition angle between Rayleigh and non-Rayleigh regions. For rotation around x , the transition occurs around 3 degrees, whereas it occurs around 7 degrees for rotation around y . The difference is due to the rate at which the axial extent of intersection changes with rotation around the two axes, which is affected by the relative sizes of the ellipsoid widths, or, equivalently, the PSF widths. As the plane is rotated around x , the change in intersection is associated with the PSF width in the y dimension. Similarly, as the plane is rotated around y , the change in intersection depends on interaction with the PSF width in the x dimension. Since the PSF is shorter in the y dimension (lateral) than the x dimension (elevation), the axial extent changes more quickly with rotation around x than y . This is apparent from the plots of axial extent; notice that the transition occurs at an axial extent of 2 wavelengths in both cases.

Figure 8.3 shows effects on the mean and SNR_0 caused by rotating around the y axis and translating the plane relative to the center of the resolution cell (changing




Figure 8.2: Effects of rotating the plane about the x axis for full range (left) and narrow range (right). Mean, standard deviation, SNR_0 and axial extent of intersection are shown versus rotation around the x axis. Simulation and planar results are included, where planar results were obtained with and without assuming amplitude and phase independent. Effects are similar to those for rotation around the y axis with the exception that that non-Rayleigh region is narrower due to interactions with the shorter PSF width in the y dimension.

the d plane parameter). In general, the mean falls off with distance according to the size of the PSF in the direction normal to the plane, and roughly with the Gaussian shape of the PSF envelope. The simulation and planar mesh plots are quite similar, with only a slight difference in the SNR_0 values near the ends of the translation range. In the SNR_0 plot, the SNR_0 in the Rayleigh region is approximately 1.91, as expected, and greater in the non-Rayleigh region. An interesting feature of the SNR_0 plots is that, for a given angle of rotation, the SNR_0 is constant with translation of the plane for both simulation and planar results. Translation of the plane, thus, appears to have little effect on whether or not the scattering follows a Rayleigh density.

Figure 8.4 shows the effects of varying the scatterer concentration. The mean and SNR_0 are shown for different concentrations at angles in both the Rayleigh and non-Rayleigh regions. The most significant change occurring with increase in concentration is an increase in the relative amplitude of non-Rayleigh means to Rayleigh mean. This relative amplitude of coherent and incoherent scattering could be an important parameter in inference if the concentration needs to be estimated from the data. The SNR_0 exhibits similar changes with concentration indicating that any change in the standard deviation (with concentration) does not appear to be significant.




Figure 8.3: Effects of rotating and translating the plane. Simulation (bottom) and planar (top) results are shown. The planar results did not assume independence of amplitude and phase in this case. Note that while the mean decreases rather quickly with translation, the SNR_0 value is constant for much of the translation range and varies mostly with rotation, where it follows the same variation from non-Rayleigh to Rayleigh statistics as in the previous figure.




Figure 8.4: Effects of changing the scatterer concentration. Simulation (bottom) and planar (top) mean and SNR_0 are shown. The most significant effect of increasing scatterer concentration is an increase in the relative amplitude of the non-Rayleigh and Rayleigh means and SNR_0 values.

Figure 8.5 shows the effects of varying the surface roughness parameter. Planar results were generated assuming independence of amplitude and phase, since that method was the only one of the previous chapter to incorporate roughness. The roughness was modeled as uniformly distributed in the direction normal to the surface. The primary effect of changing the roughness is similar to that of changing the concentration, i.e., to alter the relative amplitudes of the Rayleigh and non-Rayleigh means. The main difference is that some roughness values can eliminate the non-Rayleigh region entirely, forcing Rayleigh statistics at normal incidence. This result is as expected since a roughness of one or more wavelengths means that, even if the axial extent of intersection is zero, scatterers will be distributed over an entire wavelength. The result is evident in the plots, especially the simulation SNR_0 plot, which is constant at about 1.91 over all angles for a roughness parameter equal to one wavelength. Differences between the simulation and planar values are again associated with the oscillations in the non-Rayleigh region of the planar values when amplitude-phase independence is assumed.




Figure 8.5: Effects of changing the surface roughness parameter. Changes in roughness also change the relative amplitude of non-Rayleigh and Rayleigh means, with the additional change that for a roughness of at least one wavelength, statistics are Rayleigh for all rotation angles.

To summarize, the basic phenomena associated with planar surfaces are the following: 1) scattering can be separated into non-Rayleigh and Rayleigh regions, 2) the transition from Rayleigh to non-Rayleigh regions appears to be correlated with the axial extent of the intersection surface, which depends on factors such as the PSF widths, and the center spatial frequency, k_0 , 3) relative amplitudes of non-Rayleigh

and Rayleigh regions are dependent on the scatterer concentration and surface roughness, and 4) the amplitude-phase independence assumption produces discrepancies between planar and simulation statistics in the non-Rayleigh region, specifically with regard to the smoothness of variation with rotation. With this basic knowledge, the next experiments were intended as an investigation of the effects of surface curvature on the accuracy of computations based on the locally planar approximation.

8.1.2 Effects of Surface Curvature

Previous results have provided substantial evidence that the greatest sensitivity of the scattered signal to the underlying structure involves non-Rayleigh statistics that occur when the local axial distribution of scatterers is smaller than a wavelength. For curved surfaces approximated by a plane, one can expect that the greatest error will occur when the approximating plane occupies a smaller axial range than the curved surface. This effect can be examined directly and carefully for spherical surfaces because they are easily parametrized for varying the curvature and because they permit a simple approximation via the tangent plane.

Since the resolution cell has been modeled as an ellipsoid (a quadric surface), the surface of intersection with another quadric, e.g., the sphere, is non-trivial and beyond the scope of this work. To simplify the problem for this study with spheres, the majority of the results that follow employed a spherical resolution cell so that, for the spherical surface medium, the following advantages could be used: 1) the surface could be approximated locally (within the resolution cell) by the tangent plane at the surface point that was closest to the center of the resolution cell, and 2) the surface of intersection could be parametrized in case a numerical solution was desired.

Image statistics (pixel-based mean and variance over the entire image) were computed using the planar approximation methods for comparison with results generated from simulation. The images of this section are for the center plane of a single sphere (with a fixed radius) to show general behavior over the whole image (and the associated range of surface orientations with respect to the axial dimension). Results were then generated for key individual pixels to show the effects of varying the curvature (via the sphere radius). The single-pixel curvature experiments and findings were then extended to an ellipsoidal resolution cell representing the PSF used throughout most of this thesis. Application to arbitrary surfaces are addressed in the next section.

In generating the simulation results, many parameters could be varied. The sphere represented the gross shape, with the radius of the sphere variable to change the surface curvature. In accord with the discrete-scatterer surface model, the sphere was characterized by a scatterer concentration (scatterers/mm²) and surface roughness (width, in wavelengths, of the uniform distribution). The PSF was represented in terms of its widths, σ_x , σ_y and σ_z and the spatial wavelength, k_0 . Experiments were based on both image simulations and pixel simulations, each implemented with different code but executing the same linear system model for image formation.

In the simulations, the number of scatterers was determined deterministically based on the scatterer concentration and surface area. Note that this choice represents a minor inconsistency in the models because, over a local volume like the resolution cell, the number of scatterers will still be random. The effect is small, however, and the deterministic scatterer number has been kept because it seems to be a more appropriate model for a real surface. For the images, a sphere of fixed radius 8 mm has been used. For the single-pixel results with varying curvature, two different PSFs were used, the spherical PSF mentioned previously, with all widths 0.5 mm, and an ellipsoidal PSF with the widths used in previous chapters, $\sigma_x = 1.5$, $\sigma_y = 0.5$, and $\sigma_z = 0.2$. Where appropriate, PSF widths are listed in the figures.

Figure 8.6 shows sample images for the center plane of an 8 mm sphere centered in the image. Note the substantial difference in intensity and texture between the coherent regions at the top and bottom of the sphere relative to the speckle texture along the sides. These differences correspond exactly to the non-Rayleigh/Rayleigh effects observed with rotation of the plane in the previous results. Of course, the tangent plane to the sphere is normal to the axial direction at the top and bottom of the sphere and varies gradually from normal along the sides of the sphere. Note also the smoothness of the variation in intensity from the top of the sphere towards the sides. The spherical PSF was used in these images, causing the images to vary in appearance from typical images, e.g., the coherent regions at the top and bottom are approximately circular instead of flattened as in a typical image (the axial resolution is typically significantly better than the lateral resolution). For comparison, observe the two images in Figure 8.10 showing sample images of the sphere obtained with spherical and ellipsoidal PSFs.

Figure 8.7 shows the simulation and planar mean images for the 8 mm sphere with spherical PSF. The image computed without amplitude-phase independence is close to the simulation over the whole image, with some difference in amplitude at the

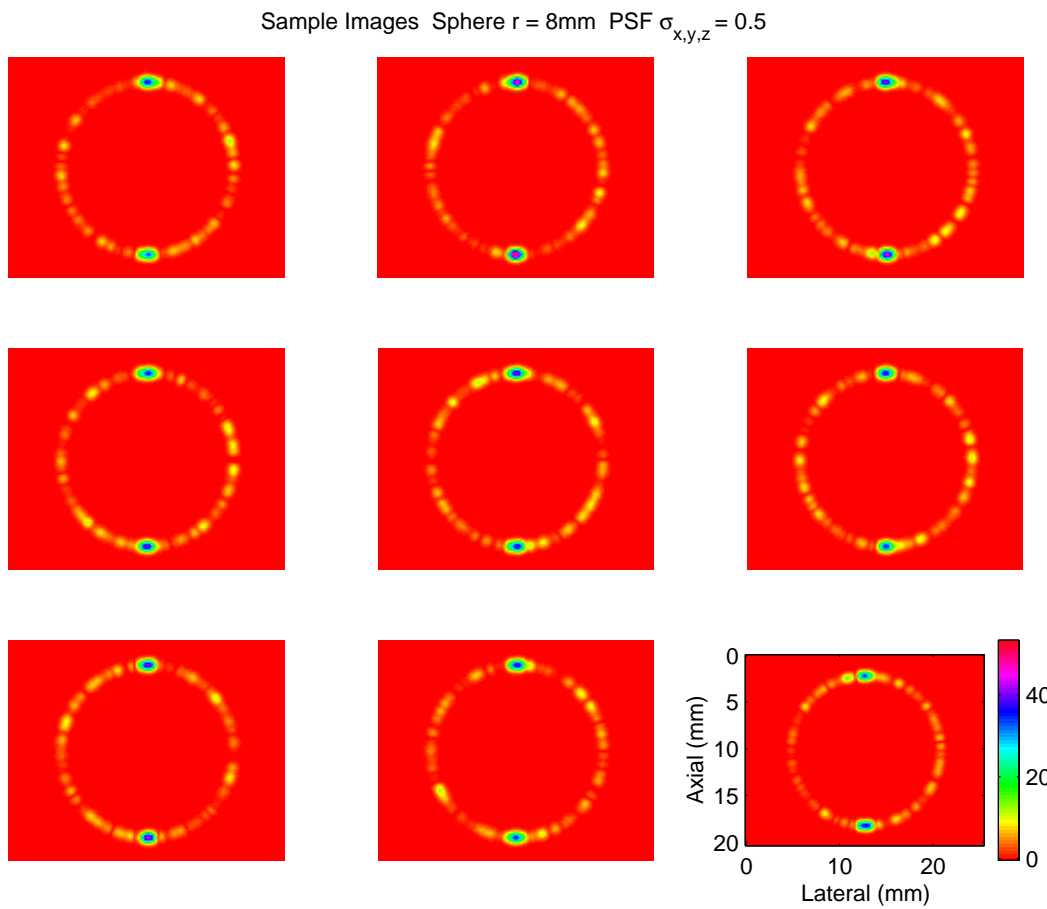


Figure 8.6: Sample images for an 8mm radius sphere with a spherical $\sigma = 0.5$ mm PSF.

top of the sphere where the curvature has the greatest effect on the phase. The main difference with amplitude-phase independence is again the presence of oscillations in the non-Rayleigh region that were seen in the previous results. The Rayleigh region around the upper left side of the sphere can be seen closely in the zoomed images on the right of the figure. The variation with distance from the sphere surface is shown in detail for that region and holds regardless of whether or not amplitude and phase are assumed independent.

SNR_0 images are shown in Figure 8.8 for the 8 mm sphere with spherical PSF. The distinction between Rayleigh and non-Rayleigh regions is clear here, where the Rayleigh region has a constant SNR_0 of 1.91. As before, the SNR_0 is constant with translation of the plane as seen previously for rotated and translated planes.

From these results, it appears that the plane approximation has substantial potential in representing the mean and variance for a curved surface. The greatest differences are in the non-Rayleigh region. To examine the effect of the surface curvature on the size of the discrepancy, curvature was varied while statistics were computed for two points, one at the top of the sphere and one on the left side of the sphere, for varying curvature. The top graph in Figure 8.9 shows the resulting variation in the mean, μ_ρ , for a spherical PSF. The results for the point at the top of the sphere show that the mean converges to the planar result as the radius of the sphere is increased. The results for the point on the side of the sphere, where the SNR_0 value indicates a Rayleigh distribution, show no effects of changing the curvature.

These results for a spherical PSF can be generalized to the ellipsoidal PSF that is more representative of typical imaging systems. For a qualitative comparison of differences between effects of the two PSFs, Figure 8.10 shows sample images of the sphere simulated with each PSF. The coherent region shrinks in the axial direction for the ellipsoidal PSF, as would be expected. Also, the nature of the texture in the Rayleigh region changes from a radially-directed texture for the spherical PSF to a laterally-directed texture for the ellipsoidal PSF. Quantitatively, the maximum error should still occur at the most coherent point on the object, or the peak of the sphere. The bottom graph in Figure 8.9 shows the effect of varying the surface curvature with the ellipsoidal PSF. In this case, the radius at which the simulation mean approaches the planar mean is much greater than that for the spherical PSF. This difference is a result of the increase in the elevation width, σ_x , from 0.5 to 1.5 mm from the spherical to ellipsoidal PSF. With the increased width, the surface of intersection is




Figure 8.7: Simulation (top) and planar (middle and bottom) mean images for the 8 mm sphere. The planar result agrees quite well with the simulation when amplitude-phase independence is not assumed (middle). Otherwise (bottom), discrepancies exist in the non-Rayleigh regions at the top and bottom of the sphere. The images on the right show a zoomed view of the upper-left region of the sphere. In this region, the assumption of independent amplitude and phase makes no noticeable difference.




Figure 8.8: Simulation (left) and planar (right) SNR_0 images. The images agree quite well and show the Rayleigh and non-Rayleigh regions distinctively.

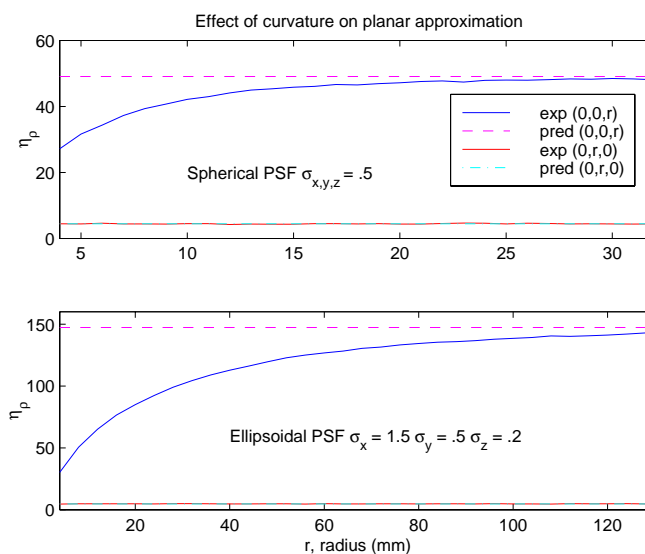


Figure 8.9: Effects of curvature on the planar approximation for the spherical PSF (top) and ellipsoidal PSF (bottom). Plots show the mean at the top-most point on the sphere location $(0,0,r)$, the non-Rayleigh location of maximum error, and the left-most point on the sphere, $(0,r,0)$, a location in a Rayleigh region. The curvature has negligible effect on the point in the Rayleigh region. At the non-Rayleigh point, the simulation results converge to the planar approximation as the radius is increased (as the surface becomes more planar). The ellipsoidal PSF requires a greater curvature for convergence due to the larger value for σ_x , the elevation width of the PSF.

larger, and the discrepancy between the actual axial extent of the sphere over that region and the extent generated by the approximating plane is greater.

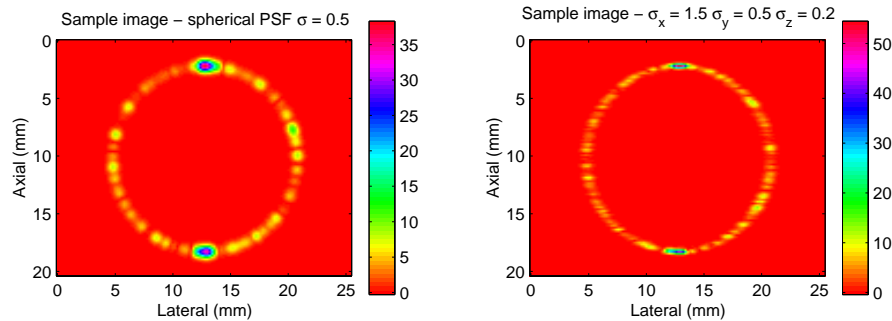


Figure 8.10: Sample simulated images of the sphere using the spherical and ellipsoidal PSFs. Several qualitative differences are evident. Expected differences include the axial width of the non-Rayleigh regions at top and bottom of the sphere. Less expected differences include the radially-oriented texture for the spherical PSF and the laterally-oriented texture for the ellipsoidal PSF.

8.1.3 Discussion: Value of the Planar Approximation

Overall, the potential accuracy of the local planar approximation model appears to be quite good judging from the successes in these basic experiments. While countless other details could also be examined, many basic phenomena have been investigated and now provide a base for further study when deemed necessary. As in previous chapters, any clear path towards improving the model must be based on a quantitative measure of performance in image analysis.

Basic technical issues

These results further strengthen speculation that images can be categorized as Rayleigh or non-Rayleigh as a first step in characterizing the statistics at a pixel. Such a categorization is important in trading computation for accuracy, because Rayleigh pixels appear to be insensitive to factors such as amplitude-phase independence, some general characteristics of the PSF, and the local curvature of the surface, as well as computation of the amplitude mean and variance as seen in Chapter 5. In contrast, the non-Rayleigh region can be especially sensitive to all of those features, requiring greater computational resources to achieve similar accuracy.

The classification into Rayleigh and non-Rayleigh regions could have impact in assessing the applicability of the planar model to various surfaces and systems. For instance, a given anatomical surface may be modeled easily with limitations on the curvature, e.g., the maximum curvature of a liver is probably much less than that on a vertebral surface. Given the curvature limitations, system requirements could be specified under which the planar approximation would be valid. The dimension of worst resolution is the limiting factor, and the accuracy of the planar approximation may be an argument for improving the elevation focus of some systems, e.g., for motivating the use of dynamic focus in elevation. Quantitative criteria for system and structure interaction such as the curvature and PSF width plots in Figure 8.9 could be useful in such a case.

Application of the model to inference

The ultimate test of the model remains its application to inference. The planar model is especially attractive because of its potentially low computational requirements. For fixed acoustic characteristics, statistics could be pre-computed for all rotations and translations of the plane and implemented via a lookup table or even fitted with simple analytic functions. Computation of the objective function (data likelihood) would essentially be limited by computing the planar approximation for each pixel, and issues such as amplitude-phase independence would be unimportant because those calculations could be made offline.

The images and results from this section immediately provide new insight into the applicability of the model to inference of shape. From the pixel-based mean and variance computed for the 8 mm sphere, sample images from the resulting Gaussian-distributed image likelihood can be generated simply. Figure 8.11 shows nine such images that can be compared directly to those of Figure 8.6. The images show visually the information available in such an image model regarding the ensemble of images produced from the linear systems model for the spherical shape. Of specific interest is that the information is entirely first-order, i.e., no neighbor interaction is considered. For images where texture is significant, the amount of similarity between the images indicates the degree to which the underlying shape is represented by the pixel-based statistics. This issue raises the question of the importance of representing second-order information, or neighbor interaction, in a model for inferring shape. The next chapter will serve as the beginning of an answer to that question.




Figure 8.11: Sample sphere images generated from a Gaussian image distribution with pixel-based statistics computed using the planar surface approximation. The images are a visual representation of the information contained within the model about each sphere image.

8.2 Results for Arbitrary Surfaces

The investigations of the previous sections provided intuitive insights into the interactions of a planar surface and the imaging system and into the limitations of using such an approach for arbitrary surfaces. The following results are intended to demonstrate the potential and shortcomings of the methods in computing statistics for images of an arbitrary surface. Simulation results are compared to planar and triangle results, where the simulation data is generated, again, via the model of Chapter 4, and the other images are computed using either the planar approximation (without A, ϕ independence) or the direct approach using the triangles.

In the following results, the scatterer concentration was 64 scatterers/mm² as in the previous chapters. Roughness was not modeled since it has not been incorporated into the methods used for this section. On this note, recall that in using the discrete-scatterer model for surfaces, a roughness is implied. A roughness value of zero, thus, approximates a roughness that is negligible relative to a wavelength. Considering the plots of the previous section for the planar surface, the mean varied little between roughness values of 0 and 0.1 wavelengths. Two PSFs were used in this section. The first was the same used throughout the dissertation, with widths $\sigma_x = 1.5, \sigma_y = 0.5$, and $\sigma_z = 0.2$ mm, and the second was the same except that the elevation width, σ_x , was changed to 0.5 mm to reduce the influence of out-of-plane surface features. The 6 MHz center frequency of previous chapters was used again. For the images, two different planes on the cadaveric vertebra were used. One image shows the transverse process, providing good detail over a small region. The other image is the lamina image of previous chapters and shows a larger region of the vertebra with a broader range of image features. Simulation statistics were generated from trials of 100 simulated images.

8.2.1 An Image Plane on the Transverse Process

The first images show an image plane in the sagittal plane of the transverse process of the cadaveric L4 lumbar vertebra of Figure 8.12. The approximate image plane is shown on the surface. The view on the right is a zoomed view of the surface region used in computations; it consists of those triangles within +/- 5 mm of the image plane and not occluded from the probe by other triangles. The plane corresponds approximately to the plane of the actual image in Figure 8.13. That actual image was acquired by imaging carefully along the transverse process and searching for the

brightest (most coherent) echo. The sensitivity to angle was such that when imaging by hand, even the slightest wobble of the probe changed the image substantially. As a result, this plane represents a challenging surface geometry for the planar approximation model since the curvature is substantial and varied relative to the PSF widths and wavelength.



Figure 8.12: Approximate image plane for the upcoming transverse process images. The view on the right is a zoomed version showing portions of the surface within ± 5 mm of the image plane and not occluded from view. The size of the image region shown is approximately 5 mm high by 15 mm wide.

In Figure 8.13, the actual image is shown along with six simulated images. Substantial variation exists among the simulated images and is representative of the variability to be expected with the discrete-scatterer model. The actual image appears quite similar to the simulated images in shape, but the actual image seems to have a wider region of coherent scattering along the top of the process with a greater relative amplitude to the incoherent scattering. Close inspection of the shape in the images shows that it is visibly rotated counter-clockwise in the simulated images. Given the sensitivity to the angle of insonification, such a change could easily account for the differences between the images. As mentioned in Chapter 6, these differences are unavoidable with the experimental methods used here since the tracking error is on the order of 2 mm.

Sample mean and variance were computed from simulated images such as those in Figure 8.13. The simulation mean image is shown at the top of Figure 8.14. A small region of coherent scattering is present near the middle of the process with incoherent scattering along the rest of the process. Referring to the zoomed view of




Figure 8.13: Actual and simulated images of the transverse process. PSF values are listed in the figure. The image region shown is from a larger image (in the actual image produced by the Tetrads system) approximately 50 mm x 50 mm. The scale for the actual image is simply the 256-level 8-bit scale of the image data. No attempt has been made to quantitatively relate this scale to the simulation scale.

the imaged surface in Figure 8.12, it is difficult to predict that such a region of coherent scattering would be produced from that surface region. The mean computed using the planar approximation is shown in the middle left of Figure 8.14 with the relative error shown in the bottom left. As may be expected, the planar approximation misses this obscure region of coherent scattering in the middle of the process. Note also the artifacts in the planar approximation image in the middle and bottom sides of the transverse process. These artifacts are due to the approximation algorithm and will be addressed momentarily.

In contrast, consider the images on the right showing the triangle mean image and relative error. The coherent scattering site is predicted rather well, and no significant artifacts are produced. These results were produced using the triangulated mesh originally generated with the Marching Cubes algorithm, i.e., with no resampling. The triangles had maximum widths between 0.2 and 1 mm, all large relative to a wavelength, presumably representing the major source of error here. In the subsequent lamina images, one iteration of resampling is shown to dramatically improve the already small error in the triangle results.

The Rayleigh and non-Rayleigh regions in this transverse process image are easily distinguished when observing the SNR_0 images of Figure 8.15. From the simulation image, most of the non-zero values are Rayleigh at approximately 1.91 with a small region in the center of non-Rayleigh scattering. Variation from the exact Rayleigh SNR_0 value in the Rayleigh region is small (less than 10 %) and could be reduced by increasing the number of simulation trials. Again, recalling the zoomed view of the surface in Figure 8.12, it would be difficult to predict the presence and location of this non-Rayleigh region by simply looking at the surface. The planar approximation result misses this small region, while the triangle result produces even this obscure feature quite well.

The artifacts in the planar images result from two cases, discontinuities in the planar approximation and partial overlaps of the surface and resolution cell. Three main regions of artifact are present on the planar approximation mean image of Figure 8.14, one in the center of the bone section and one on the bottom of each side.

Artifacts in the center are due to the discontinuity in approximating the normal. In this region, the intersection of surface and resolution cell changes abruptly in the axial direction because of the surface regions that are intersected by the associated resolution cells. Consider the images of Figure 8.16 showing the orientation angle (relative to the beam axis) and distance of the approximating plane. At approximately

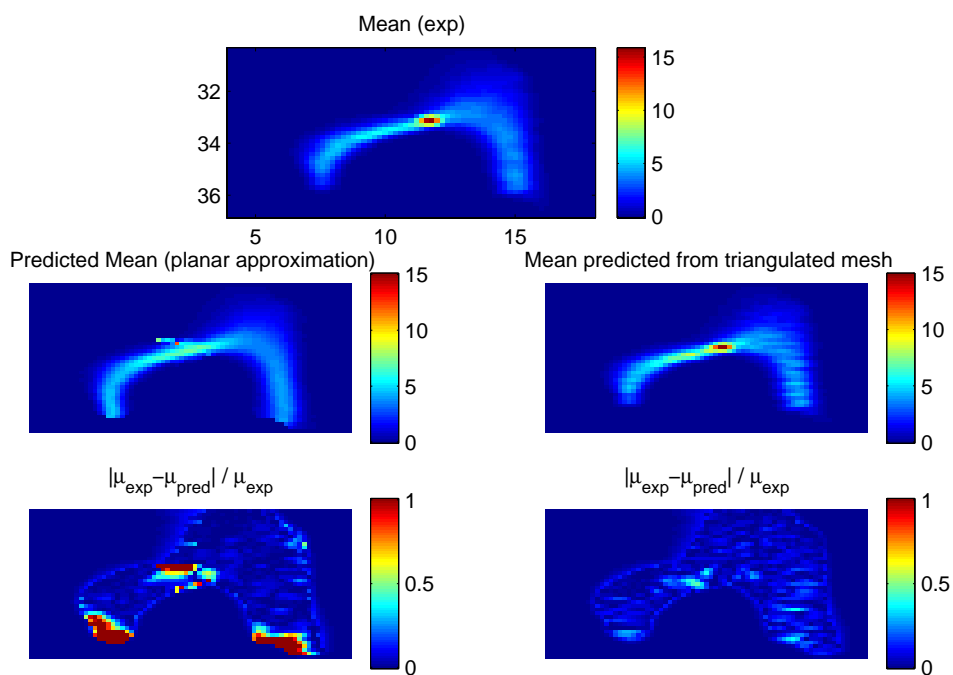


Figure 8.14: Simulation, planar and triangle means for the transverse process images. The simulation mean (top) is matched quite closely by the triangle mean (middle right with relative difference on bottom right), and matches in much of the region for the planar mean (middle and bottom left). See text for other details.

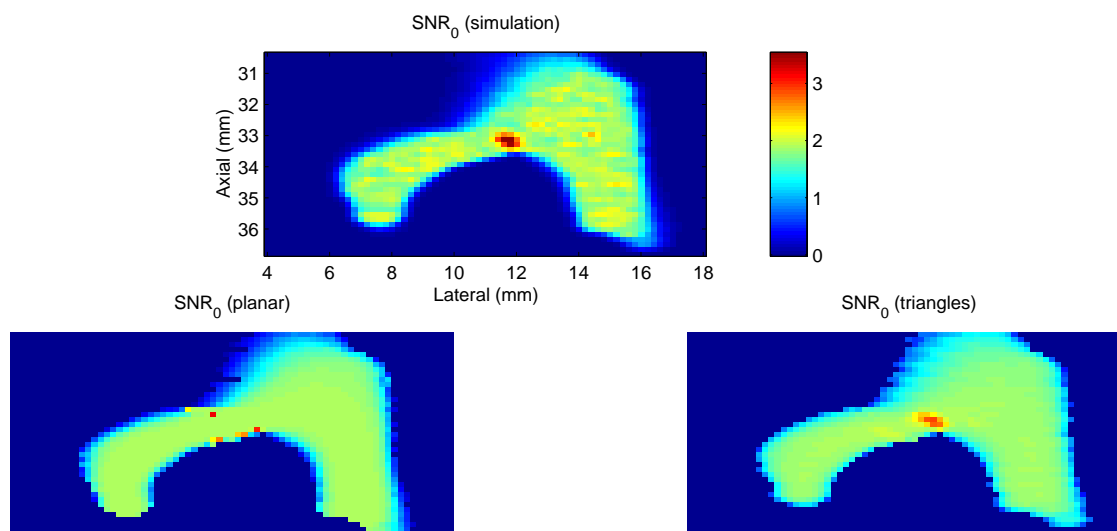


Figure 8.15: Simulation (top), planar (bottom left) and triangle (bottom right) SNR_0 images for the transverse process. The images agree over much of the image region where statistics are Rayleigh ($\text{SNR}_0 = 1.91$). The triangle image shows the small region with non-Rayleigh statistics where the planar image does not. Colormaps for the images are the same, with the associated colorbar shown next to the simulation image.

12 mm lateral and 33 mm axial, discontinuities exist in both the angle and distance of the approximating plane. Below the discontinuity, the approximating normal is affected mostly by the flat region of the surface, where the approximating normal is around 0 degrees. Above the point of discontinuity, the normal is affected by the out-of-plane curved region of the process that is oriented at about 60 degrees. The result is a discontinuity in the approximated mean, even though the actual change in mean amplitude is smooth. This problem with the planar approximation would be resolved if the surface were less curved relative to the width of the PSF.

The artifacts on the sides of the process are due to incomplete intersections of the surface and resolution cell, i.e., because of the removal of occluded sections of bone, open surface regions exist that only partially overlap the resolution cell at some pixels. The approximating plane overlaps the cell completely, resulting in a computed amplitude that is higher than the actual. These artifacts would not exist for non-occluding, closed surfaces like the soft tissue of the liver or brain, and they could potentially be compensated by comparing actual area (summed during the approximation) to the area computed using the planar approximation (which is already calculated during the approximation).




Figure 8.16: Images of the planar approximation for the transverse process image. Images show the angle (left) between the approximating normal, N , and the axial dimension and the distance (right) from the approximating plane to the resolution cell center. The discontinuity in both at around 11 mm lateral and 33 mm axial produces artifacts in the computed statistics.

8.2.2 Transverse process images with a reduced elevation width

Some of the planar approximation artifacts that result from the non-planar variations of the surface over the resolution cell are reduced by reducing the size of the PSF. In the following figures, images were generated using an elevation width, $\sigma_x = 0.5$ mm, reduced by a factor of 3 from the previous set of images. Other parameters were the same.

Sample images are shown in Figure 8.17 with the same actual image as in the previous analysis. Reducing the elevation width reduces the contribution of out-of-plane surface elements, which results in images that appear more like an in-plane cross section of the surface (aside from speckle-based textural variations, of course).

Simulation, planar and triangle mean images are shown in Figure 8.18. In this case, the non-Rayleigh scattering site has disappeared from the center of the image. The planar result agrees well over most of the image, except for the artifacts that result from partial overlap of the resolution cell and surface.

The SNR_0 images of Figure 8.19 show the improved agreement in the center of the image and show quantitatively that the scattered statistics are all Rayleigh in this image. The plane parameters shown in Figure 8.20 show none of the discontinuities of the previous case, indicating why the discontinuity artifacts are not present in this case.




Figure 8.17: Sample simulated images of the transverse process with a reduced elevation width in the PSF. PSF values are listed in the figure. Much of the coherent scattering evident with the wider elevation width has disappeared, and the maximum image value is much lower.




Figure 8.18: Mean images for a reduced elevation width, $\sigma_x = 0.5$ mm. The out-of-plane surface features have little influence here, and some artifacts of the planar approximations are eliminated as a result.




Figure 8.19: SNR_0 images for reduced elevation width $\sigma_x = 0.5$ mm. The images show quantitatively that scattering is all Rayleigh for this image. Variation from a Rayleigh SNR_0 is slight in the simulation image and, again, would decrease with more trials.




Figure 8.20: Images of the planar approximation for the reduced elevation width. The discontinuities disappear for this image plane since the surface does not curve faster than the resolution cell ellipsoid, i.e., the surface cannot intersect the ellipsoid in disconnected pieces.

8.2.3 Sagittal plane images along the lamina and articular processes

The transverse process images showed an obscure region of non-Rayleigh scattering and the elimination of that region with reduction of the PSF width. The sagittal image plane shown in Figure 8.21 represents a large region of varying surface curvature, featuring both Rayleigh and non-Rayleigh regions as well as substantial contributions from the out-of-plane surface in the facet joint (left side of the image). The following images give further examples of the image model applied to the vertebral surface, including improvement of the model with surface resampling in computing directly from the triangulated mesh.

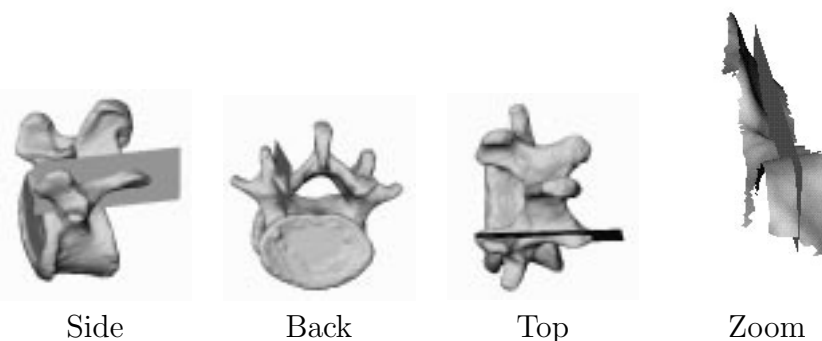


Figure 8.21: Views of image plane for lamina image overlaid on surface rendering. The zoomed image shows the extent of the surface (± 5 mm in the elevation dimension) that is considered in the computations. For the various image features, the angle of the surface relative to the image plane is primarily responsible for determining whether or not scattering amplitude obeys the Rayleigh distribution.

Figure 8.22 shows several sample simulated images of the lamina plane. The lamina and inferior articular process show relatively high-amplitude non-Rayleigh scattering in all of the images, with the lamina producing a higher amplitude more consistently. The facet joint is marked by Rayleigh scattering with slowly decreasing amplitude toward the top of the image. This slow decrease in amplitude results because the image plane is nearly parallel to the facet joint surface, with out-of-plane contributions decreasing slowly (due to increase in distance from the image plane) over a large axial range.

Planar mean images are shown in Figure 8.23 along with the simulation mean image. As for the transverse process images, the planar approximation suffers when the surface is nearly normal to the beam and includes regions with curvature greater




Figure 8.22: Sample simulated images of the lamina and articular processes along with actual image of approximately the same region. From left to right, anatomical structures are the facet joint on the left (Rayleigh scattering with wide axial extent), the lamina in the center (non-Rayleigh scattering with relatively high amplitude) and the inferior articular process on the right (mix of Rayleigh on the sides and non-Rayleigh at the peak).

than the resolution cell. These artifacts are easily seen along and around the lamina (center) and the peak of the inferior articular process. For a narrower resolution cell and/or a surface with reduced curvature, such artifacts would reduce as in previous cases. For this surface region, the surface extends over the entire image plane, thus no partial-overlap artifacts are present. Overall, the planar approximation works quite well in the Rayleigh region again.




Figure 8.23: Simulation and planar mean images for lamina image plane. The simulation image shows the sites of non-Rayleigh scattering on the lamina and peak of the articular process as well as the low-amplitude Rayleigh scattering along the facet joint.

The triangle mean image is shown in Figure 8.24. The artifacts of the planar approximation are absent, of course, and agreement is much better between the images. The relative error image shows that the disagreement is still quantitatively significant in some small regions, however, specifically along the lamina and peak of the articular process.

The triangles used for the results in Figure 8.24 were generated directly from the CT volume segmentation using the Marching Cubes algorithm and had a mean width of approximately 0.5 mm with a maximum width in the axial dimension reaching over 1 mm. These widths are much greater than the 0.2 mm width of the axial envelope for the PSF and could be expected to contribute significantly to the disagreement between the simulation and triangle image. In Figure 8.25, triangle mean images are shown using a set of triangles resampled by 4 times as in Appendix A with




Figure 8.24: Simulation and triangle mean images. The triangle mean shows substantial improvement over the planar mean image in terms of artifacts. Quantitatively, there are still substantial differences between the simulation and triangle means.

a mean width of around 0.2 mm and maximum width approximately 0.5 mm. The results show significant improvements in those small regions of disagreement. Further improvement could be expected from subsequent refinement of the mesh. This result, however, was considered sufficiently accurate for investigating inference. For one pixel of translation between the simulation and triangle images, the relative image error was much greater than the disagreement in either of the previous two figures, suggesting that the current result should have significant discriminating power over poses differing by the size of a pixel, or approximately 0.15 mm.

SNR_0 images are shown in Figure 8.26 for simulation and the resampled triangle method. Results agree quite well, with sites of non-Rayleigh scattering along the lamina and articular process as anticipated from the mean images. The greater SNR_0 along the lamina could be due to an angle of incidence closer to normal than that of the articular process, although it would be difficult to predict by looking at the surface because of the relative flatness of the articular process and its wider extent. This ambiguity is yet another example of the sensitivity of the interactions between surface and PSF that produce the highly varied images we regularly observe.

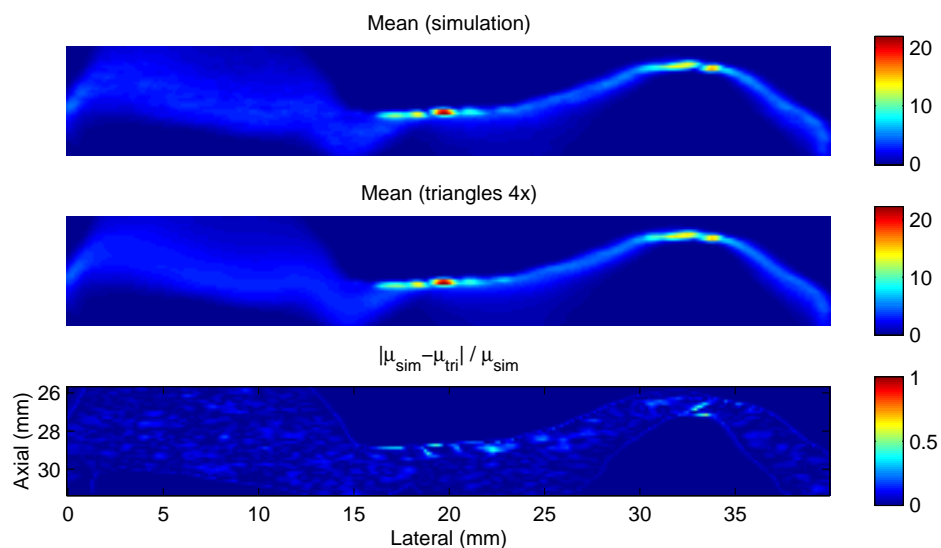


Figure 8.25: Simulation and triangle (from original triangles resampled by 4x) mean images. Computing from the resampled triangles reduced the error in the approximation substantially from that of the previous figure.




Figure 8.26: Simulation and triangle (from original triangles resampled by 4x) SNR₀ images. Regions of non-Rayleigh scattering are in agreement and are present where expected from the mean images (lamina and peak of the articular process), although the lamina represents a site of greater SNR₀ and higher relative amplitude of scattering.

8.3 Discussion

While images have been investigated for a limited selection of system PSF and tissue surface media, the Rayleigh/non-Rayleigh characterization of surface images does represent many fundamental features observed in both the actual and simulated images. Only two of an infinite number of possible image planes were used; only one surface (the surface of the cadaveric vertebra) was used; and the surface was used with fixed acoustic characteristics. The results generally behaved as expected, though, with Rayleigh and non-Rayleigh scattering occurring with a dependence roughly on the orientation of the local surface normal with respect to the axial image dimension.

Perhaps the most important results were those that provided further evidence of the extreme sensitivity of the interaction between the system PSF and the tissue surface. For instance, recall the high-amplitude coherent scattering that was evident in a case where it would have been difficult to predict from visual inspection alone. This sensitivity agrees with clinical imaging experience of both *in vivo* and *in vitro* spines in the sense that “finding” these locations of coherent scattering by hand is often quite difficult, even for the experienced operator. The relevance of this sensitivity to applications of the image model will again be dependent on the particular situation, i.e., on the characteristics of the imaging system and the surface geometry.

Image model accuracy

The planar approximation method performed as expected, with good results in Rayleigh regions and difficulties in non-Rayleigh, coherent scattering regions with high sensitivity to the exact surface structure. Artifacts could be reduced in most cases by additional computation, e.g., comparing computed area (computed from the planar approximation) and measured area (measured while computing the planar approximation). Artifacts resulting from a partial overlap of surface and resolution cell could probably be removed using such an approach, although these artifacts would not exist for closed surfaces, i.e., for soft tissue surfaces where occlusion is not considered. Other artifacts, such as those in the non-Rayleigh regions of Figures 8.14 and 8.23 are likely to be eliminated for the planar approximation only for regions where the surface does not curve significantly relative to the PSF.

Computation of statistics directly from the triangles was remarkably accurate in regions of coherent scattering, especially with resampling of the triangles. In Rayleigh regions, the accuracy suffers slightly with larger triangles, where the error

in the exact calculations is greater than that produced in approximating the surface as planar. A combination of the two models exploiting the advantages of each could potentially be developed for acceptable accuracy with minimal computational requirements.

Acceptable accuracy remains an undefined term without a quantitative measure. In the next chapter, the image model computed directly from the triangles will be used for inferring surface pose. Performance in inference will ultimately point to the development of a quantitative measure that could be used to evaluate model accuracy requirements and computational needs for specific applications.

Computational requirements

The results in this chapter were generated using algorithms coded for MATLABTM. Computation times were measured for the sagittal-plane lamina images. The region of the original surface used for the image (shown in the zoomed view in Figure 8.12) contained approximately 2500 triangles before resampling. The image region was 40x280 pixels. Calculation times listed are for a Silicon Graphics Indy workstation.

For the planar approximation, two steps were required, first, the calculation of the locally approximating plane at each pixel, and, second, calculation of the mean and variance from each locally approximating plane. For the lamina image, calculating the locally approximating planes for all image pixels required approximately 1.5 minutes. Using that information to compute the mean and variance over the image took 8.6 minutes. In the planar approximation, the mean and variance computations could easily be stored in a lookup table or fit with a polynomial, making the second set of computations negligible.

Computing directly from the triangles also required two steps. The first step consisted of computing the mean, variance and correlation coefficient parameters for the complex Gaussian at each image pixel. In the second step, the amplitude mean and variance were computed at each pixel from those parameters. For the triangulation at the original size, the first step required approximately 2 minutes, and the second step took 28 minutes. The second step is computationally intensive because at each pixel a 2D numerical integration over the complex Gaussian pdf is performed to estimate the mean (the variance calculation is relatively trivial). For the resampled triangles, the first step took 8 minutes (4 times the time for the original as should be expected with 4 times the triangles) and the second step was the same. As with the planar approximation, the second step can be implemented in a lookup table

relating the 5 complex Gaussian parameters to the amplitude mean and variance. Computation time then depends only on the number of triangles used.

The planar approximation, thus, represents the computationally attractive choice for application of the model, although the artifact issues would probably have to be resolved for most applications. Whether or not the artifacts present a problem depends on the impact on application performance. With a PSF of size small relative to the curvature of the surface, though, the planar approximation might be sufficient without change. For the rest of this dissertation, computation directly from the triangles will be used, leaving applications of the planar approximation for future work.

Computing directly from the triangles holds less promise for reducing the computation time, with the first step taking nearly 4 times as long as that for the planar approximation (assuming resampling by 4 times). The potential accuracy and lack of artifacts make it an attractive choice, though, either by itself or in combination with the planar approximation model. A hybrid model combining the simplicity of the planar approximation with the possibility of computing directly from the triangles could have potential benefit. For instance, an indicator of the planar approximation reliability could be constructed by comparing the area of intersection computed with the planar approximation to that computed from the triangles (could be computed at no cost while calculating the planar approximation). Such an indicator could prevent many of the planar approximation artifacts, especially if combined with the direct computation for artifact-likely situations.

Inference

As in the previous section, one can get a glimpse into the representation of shape provided by the image model by viewing sample images realized from the model. The model consists of pixels considered either Rayleigh or non-Rayleigh, with a single parameter defining each Rayleigh-distributed pixel and a mean and variance defining a Gaussian distribution for each non-Rayleigh-distributed pixel. The following sample images were generated from samples of the Rayleigh/Gaussian-distributed image models.

Figure 8.27 shows sample images for the transverse process image with the typical PSF ($\sigma_x = 1.5$ mm). The images were generated from the statistical model produced directly from the triangulated mesh. As samples of the imaging process, the images should be compared to the simulated images of Figure 8.13. While the

speckle texture of the simulated images is not present because of the independent pixel assumption, the images do reveal the same representation of the transverse process shape, with hints of the obscure region of coherent scattering at the center of the process, the same one that was seen in the simulations and mean images.

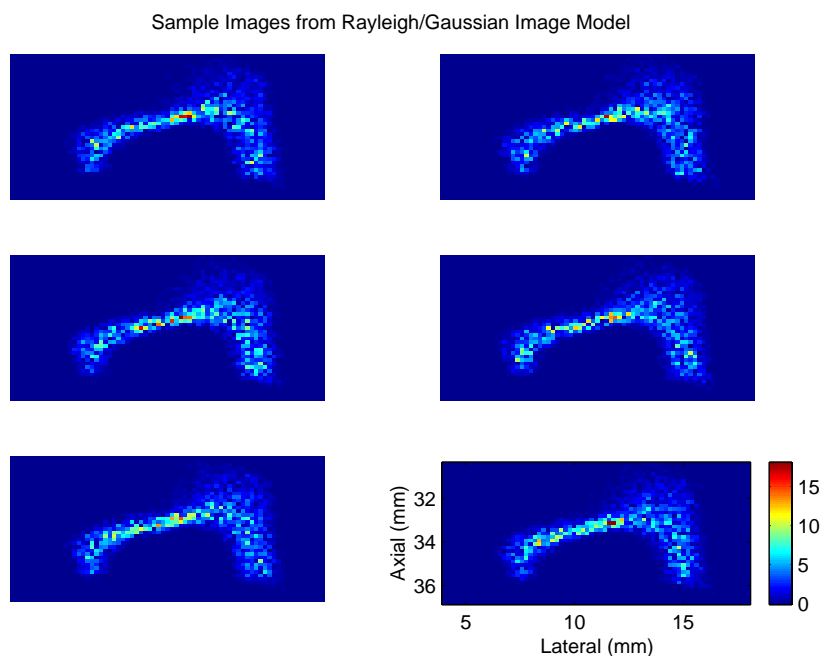


Figure 8.27: Samples of a Rayleigh/Gaussian image model for the transverse process image plane with typical PSF ($\sigma_x = 1.5$ mm).

Figure 8.28 shows samples of images generated from the planar approximation image model for the transverse process image plane imaged with the reduced-elevation-width PSF. Because the pixels are all Rayleigh-distributed, the images look less similar to their counterparts of Figure 8.17 than the images with non-Rayleigh scattering as well. This difference is due to the fact that much of the image appearance, the model representation of shape, is based on the relative amplitude of the pixels. The basic shape of the process, however, is still clearly represented.

Figure 8.29 shows sample images for the lamina image plane, with the samples generated from the image model that was computed directly from the resampled triangles. These images should be compared to the simulated images of Figure 8.22 and, because of the larger image region, give a broader indication of how the image model represents shape. Because of the non-Rayleigh coherent regions in the image,




Figure 8.28: Samples of a Rayleigh/Gaussian image model for the transverse process image plane with narrow PSF ($\sigma_x = 0.5$ mm).

the appearance is again closer to that of the simulated images than for the previous all-Rayleigh image.

In any image region, the connection between the presence of coherent scattering and information about the surface shape has not been investigated quantitatively. Qualitatively, the importance is clear, especially in clinical images, when contrast between adjacent tissues becomes the deciding factor in whether or not the shape can be recognized. Extending this importance to a quantitative assessment of shape is not trivial, but the subject will be addressed briefly in the next chapter, an investigation of performance in inferring vertebral pose using the image model.

From the visual comparisons, the Rayleigh/Gaussian images appear to capture significant information about the shape in the image. Without quantitative assessment, a conclusive argument on the validity of the Rayleigh/Gaussian characterization may be premature, but it does seem that little would be gained by further specificity in the amplitude distributions, i.e., it seems unlikely that the k - or generalized k - distributions used in other statistical characterizations of scattering would add significant value to the image model.




Figure 8.29: Samples of a Rayleigh/Gaussian image model for the lamina image plane. Statistics were computed directly from the resampled triangles.

8.4 Conclusions

Algorithms have been developed for generating a Rayleigh/Gaussian image model for an arbitrary surface. Statistical mean and SNR_0 images produced using the algorithms have matched simulation results quite well. How well remains an area of work, where a quantitative measure of performance is still required. The visual quality of the results suggests, though, that the model would perform well in inference, thus the next chapter provides the initial investigation in the area. In the limited context of inferring vertebral pose, the performance becomes the initial quantitative measure for evaluating the image model, providing the basis for application-specific direction of future improvements in the model.

Chapter 9

Inference of Vertebral Pose Using the Image Model

In the previous chapter, methods were developed for computing the mean and variance at any image pixel from the system and surface characteristics. From the pixel-based mean and variance, inference of the surface shape is framed in a probabilistic sense by constructing a data likelihood for the image data conditioned on the surface shape. In the pattern-theoretic representation, the surface shape consists of a template surface and a set of allowable transformations that act on the template. Transformations are limited here to rigid transformations, thus the data likelihood represents the probability of an observation conditioned on template pose.

9.1 A Rayleigh/Gaussian Image Model

The data likelihood characterizes all observation data, i.e., all image pixel measurements, with a single probability density function. Pixel intensities have been assumed independent, reducing the data likelihood, $p(x|h)$, to a product of probability density functions, $p_{x_i}(x_i)$, for N individual pixels,

$$p(x|h) = \prod_{i=1}^N p_{x_i|h}(x_i|h). \quad (9.1)$$

From the mean and variance at any pixel, the pixel is represented as either Rayleigh or Gaussian, depending on the value of $\text{SNR}_0 = \frac{\mu}{\sigma}$ at the pixel. For this work, any SNR_0 value less than 1.95 was assumed Rayleigh. The number is a little higher than the

theoretical value to allow for some computation error without losing the advantages of the Rayleigh characterization. The data likelihood is then a product of Rayleigh and Gaussian probability density functions with parameters derived from the system and surface characteristics,

$$p(x|h) = \prod_{x_i \text{ Rayleigh}} \frac{x_i}{\alpha_i^2} e^{-\frac{x_i^2}{2\alpha_i^2}} \prod_{x_j \text{ Gaussian}} \frac{1}{\sqrt{2\pi\sigma_j^2}} e^{-\frac{(x_j - \mu_j)^2}{\sigma_j^2}} \quad (9.2)$$

where the dependence on the pose, h , is implicit but not shown.

The likelihood serves the purpose of a cost function, or objective function, for estimating the template pose. The log likelihood, the logarithm of the likelihood, preserves the maxima and is simpler to compute in this case since the product becomes a sum.

$$\ln p(x|h) = \ln \left[\prod_{x_i \text{ Rayleigh}} \frac{x_i}{\alpha_i^2} e^{-\frac{x_i^2}{2\alpha_i^2}} \prod_{x_j \text{ Gaussian}} \frac{1}{\sqrt{2\pi\sigma_j^2}} e^{-\frac{(x_j - \mu_j)^2}{\sigma_j^2}} \right] \quad (9.3)$$

$$= \sum_{x_i \text{ Rayleigh}} \ln \left(\frac{x_i}{\alpha_i^2} \right) - \frac{x_i^2}{2\alpha_i^2} - \sum_{x_j \text{ Gaussian}} \frac{1}{2} \ln (2\pi\sigma_j^2) + \frac{(x_j - \mu_j)^2}{\sigma_j^2}. \quad (9.4)$$

The log likelihood function will be used in the remainder of the chapter.

To this point, noise in the image has not been discussed. In general, the issue of noise in an ultrasonic image is a complex one when representing shape. In a conventional image, additive system noise may not even be detectable for typical tissue structures and is, for applications of interest, of much less importance than random variation due only to scattering from the microstructure. For using the image model to infer shape in a clinical image, however, it may be desirable to model tissue surrounding the shape of interest as noise. Ideally, the surrounding tissue would be modeled as thoroughly as the shape of interest, but such an approach may not be efficient or feasible. The issue of structured noise (tissue noise) or unstructured (electronic system noise) is an interesting one but one left for future research since it is beyond the scope of this work.

For the results that follow, additive Rayleigh noise is assumed with some constant mean value. In the image model, the noise is added at the RF stage as a zero-mean, complex Gaussian with equal real and imaginary variances (the equivalent of the Rayleigh amplitude). In the simulations, a realization of the same is added to the RF image before envelope detection.

9.2 Inference

The availability of a log likelihood function allows the search for a maximum likelihood estimate over the parameter space. The techniques of nonlinear optimization have been applied to the search problem to assess the validity of the model for inference of shape. Note that these results are intended to show feasibility, i.e., that it can work, rather than to prove that a certain algorithm solves the problem of inferring shape, i.e., that it will always work.

Ideally, one would like some theoretical proof on the convexity of the objective function, the presence of local and global maxima, and other characteristics regarding optimization of the function. Unfortunately, the computations of the log likelihood used here do not permit such a characterization. For instance, convexity requires computation of the Hessian matrix, but since the gradient must be computed here by finite difference, a reliable estimate of the Hessian is difficult to obtain. The complex nature of the involved computations makes analytic solutions even for the gradient unlikely. Likewise, theoretical analysis of the function behavior over any range is not possible at this point.

From a practical standpoint, one can get a sense of the potential for an optimization algorithm by simply plotting the sensitivity of the log likelihood to the rotation and translation variables, or the variation of the log likelihood as the surface is rotated or translated. In Figure 9.1, this sensitivity is shown for a single image taken from the sagittal image of previous chapters. The plots on the left show the sensitivity when a simulated image was used as observation, and the plots on the right show sensitivity for the actual image. In both plots, the variation is smooth, and the curves have a single peak close to the expected maximum. Sensitivity to the different variables is substantially different. For instance, translation in the axial dimension (cyan curve) causes much greater variation in the log likelihood than rotation in any dimension (blue, red, or yellow curves). Because of this difference in sensitivity, one should expect inference results with better accuracy in translation than rotation. One should also expect some effects on the performance of the various optimization algorithms because of the associated differences in magnitudes of the gradient vector in different dimensions.

In the results of Figure 9.1 as well as those that follow, the theoretical log likelihood was modified to include a minimum log value of -10, corresponding to a probability of approximately $4.5 * 10^{-5}$. Without this modification, highly unlikely

pixel values drive the algorithm, i.e., search directions might be generated to change the log likelihood at one very unlikely pixel out of the thousands that make up the image. For instance, a pixel predicted to be coherent (Gaussian with a high mean and relatively low variance) could produce extreme variation in the likelihood as the shape is translated through a region of noise. Such high variation is much greater than the useful information, resulting in noisy variation of the objective function and local maxima. The noted change to place a minimum on the computable value for the log probability at any pixel improved convergence of the optimization algorithms dramatically. The quantity used (minimum log value = -10), was chosen based on observation and experimentation.




Figure 9.1: Variation of the log likelihood with change in the rotation and translation variables. Variation is for the sagittal plane lamina image with a simulated image (left) and an actual image (right). Scale is +/- two degrees (for rotation) or mm (for translation)

Choice of observation data also has a significant effect on performance in inferring pose. A small number of images is desirable, and coherent scattering sites will presumably provide more sensitivity to variation of pose than incoherent scattering textures. Intuitively, coverage of the surface of interest by the images is desirable and will affect performance. A thorough examination of the effects of these choices was not of great interest here, except to the extent that computation was affected, i.e., the smallest possible data set for consistent results was desired.

9.3 Results

After much experimentation with the details of the optimization algorithms, successful results were obtained. For direction-finding algorithms, the gradient and a BFGS quasi-Newton approximation were both tested extensively. Step size was based on an iterative quadratic fit along the search direction [51] for an “exact” line search. Gradients were calculated by finite difference with either forward or central differences using varying interval sizes. Convergence criteria included small change in the log likelihood and gradient near zero, although both required assessment for the specific data set since the magnitudes depend on the number of images and amount of data. Minor but important results were obtained using a single actual image, but much more significant and thorough testing was done for a data set of simulated images.

The algorithms for inference were written in C++, improving the computational demands significantly. Computation of the pixel-based amplitude mean from the five parameters of the complex Gaussian was implemented using interpolation from a lookup table. For the image of Chapter 8 that required 40 minutes in MATLABTM, approximately 4 seconds were required in this implementation. This speedup is remarkable and may even seem impossible. The computations required for the image model are rarely matrix computations, though. Several loops are required, and speedup of these is enormous when coded directly in comparison with MATLABTM.

9.3.1 Results - Actual Image

Results for the actual data are limited by several difficulties inherent to the problem. First, registration of the actual images includes error in the relative alignment of different image planes. The magnitude of this error is not well established, but the 2 mm range expected for general tracking errors is probably close. Accurate models for this error are difficult because of the complexity of tracking, and use of the model-based approach without an error model would be of limited value.

In addition to tracking error, potential scaling errors exist in the many stages of obtaining the experimental data. Because of the high sensitivity of the image model to very small (sub-millimeter) changes in shape, even small scaling errors in the CT scan, the segmentation of the CT images, the Marching Cubes construction of the surface, and calibrating the ultrasound probe, could significantly affect the performance of an optimization algorithm. Because of these potential inaccuracies, the simulation

environment provided a better testbed for inference, one with scaling known to precision of the computing environment. In addition, the purpose of this work was to investigate inference when shape is represented by the highly variable mechanism of ultrasonic scattering. Practical details such as the accuracy in representation of the surface from CT, registration issues, etc. are left for future work.

With these limitations in mind, experimentation using a single actual image did provide some insight into the potential for the algorithm with real data. For the actual images, a scale factor was required for matching the intensity range on the actual image with that of the image model. The scale factor was chosen based on visual comparison of the two images. The surface microstructure parameters were chosen as concentration of 64 scatterers/mm² and roughness of 0 (recall that roughness was not implemented for computation directly from the triangles). PSF parameters were assumed as before with $f_c = 6.0$ MHz, $\sigma_x = 1.5$ mm, $\sigma_y = 0.5$ mm and $\sigma_z = 0.2$ mm. Some variation of these parameters was investigated, with limited effects on the results.

Figure 9.2 shows the sagittal actual image, along with a good result from several inference trials. The images show improvement of the registration based on the alignment of the structures in the actual image and mean as seen in the difference images (difference between the actual image and mean). The images on the left represent an initial guess taken from a registration between the phantom and CT images followed by tracking of the ultrasound probe and subsequent images. The images on the right represent the pose with a likelihood much higher than the original guess. Clearly, the alignment is better in the images on the right, those at the higher likelihood.

While these results are quite limited, they still show strongly the potential of the image model to represent the relationship between the surface, its pose, and the image data. Despite the limitations of the data acquisition process, recall that these results were obtained for a 3D pose with just a single image. The sensitivity of Figure 9.1 is also evidence of the potential of the image model for representing shape. Further results with actual data would require a much stricter data collection process to truly evaluate the potential of the model and likelihood.

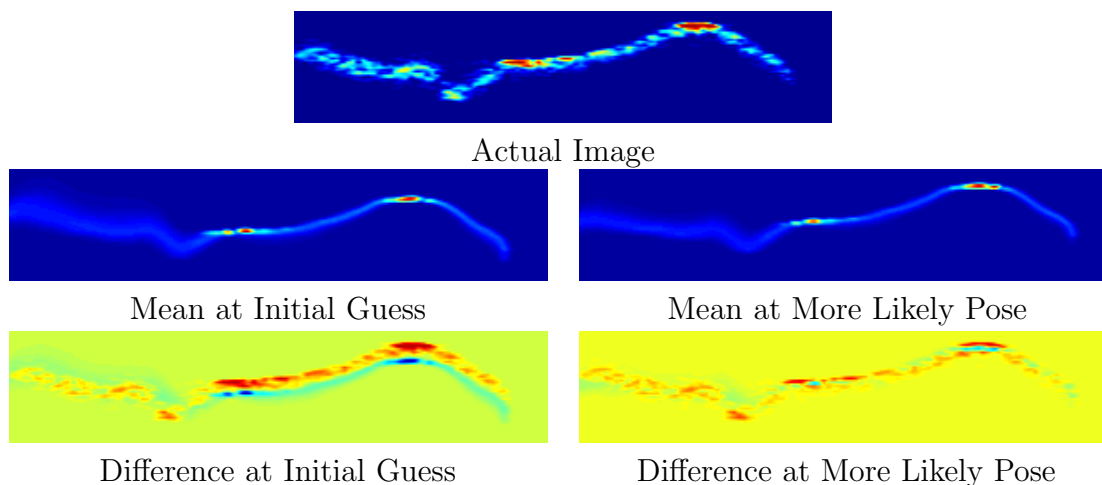



Figure 9.2: A good result from several trials with the single sagittal actual image. Difference images (difference between actual image and mean) show the improvement in alignment at the higher-likelihood pose.

9.3.2 Results - Simulated Images

Many trials and experiments were run with simulated images as observation data. The number of images was varied, along with the type of algorithm (gradient ascent, BFGS), form of finite difference (forward, central) and size of finite-difference interval. With such a variety of options and limited computational power (computational requirements increase quickly with the number of images and the central finite difference), no obvious theoretical basis existed for choosing a minimal set of options for success. As a result (as for many other practical optimization problems [52]), success required a mix of intuition and trial and error to find a data set and optimization options producing consistent and accurate results for a given experiment.

For the following results, a data set of three images was used. The image planes, the simulated images, mean images (at the known true pose), and difference images are shown in Figure 9.3. These simulated images include additive Rayleigh noise at a level of approximately 10% of the coherent signal strength in the image. The difference images give a visual indication of the alignment between images. Again, the difference image was not used for any quantitative assessment, only to indicate the alignment. The difference images will appear again in Figure 9.5 to show the improvement in alignment throughout an optimization trial.

Several trials of the experiment were run before a configuration of the optimization algorithm was found that worked well consistently. The greatest difficulties were due to noise in computing the log likelihood and to the unequal scaling in the



Orientation of Image Planes




Image 1




Image 2




Image 3

Figure 9.3: Data set for tests using simulated images. The rendered surface shows the relative orientations of the three image planes. Image 1 lies in the axial plane cutting through the spinous process. Image 2 is the familiar sagittal plane image. Image 3 is an image of the transverse process. For each image plane, the images shown are (from top to bottom) the simulated image, the mean image, and the difference image. The images also include added Rayleigh noise.

pose variables (as seen previously in Figures 9.1). The log likelihood was smooth at the resolution of the sensitivity figure (approximately 0.05 mm(deg) per step), but at some poses and with finite difference intervals smaller than 0.05 mm(deg), computation noise was sometimes larger than the actual increase of the function.

Noise in the objective function comes from many sources. Recall that several steps are taken with chances for discontinuous step changes. The triangulated surface is first modified (with Hidden Surface Removal) to account for occlusion of the beam by the bone surface. This occlusion has been modeled by making each triangle either visible or not, i.e., triangles are not clipped to produce smooth edges at the borders of occluded regions, thus variation with change in pose is not guaranteed to be absolutely smooth. This problem alone may explain the noise in the objective function. In addition, coherent scattering sites are sensitive enough to various conditions that variation of pose over a small distance can cause a noisy variation in the objective function due to changes in the sampling of the coherent region. Such a problem could potentially be solved by higher-resolution sampling or computation accomplishing the same. For the purposes of this study, however, a larger finite difference interval was acceptable, and much simpler.

Despite these difficulties, a configuration of the optimization routine was developed that produced good results most of the time. Figure 9.4 shows the magnitude of the rotation and translation error vectors over 20 trials with a BFGS quasi-Newton algorithm. The algorithm used a forward-difference gradient approximation with a difference interval of 0.05 mm (deg). Convergence was based on small change in the log likelihood and a small gradient (values were chosen based on plots of the sensitivity near the true pose). The algorithm was also designed to stop if a search direction was not an increasing direction. The initial guess was a random pose composed of a random rotation vector (each element taken from a uniform distribution then scaled for a magnitude of 2 degrees) and a random translation vector with magnitude of 2 mm.

The results are between good and great with good consistency. Excellent results were found on 15 of the 20 trials, with rotation error less than 0.4 degrees and translation less than 0.2 mm. On 4 of the 5 remaining trials, the rotation error is still below one degree, and the translation error still less than 0.3 mm. Only on one trial was performance poor, and on that trial a local noise maximum was found, i.e., on the line formed by the search direction, the pose was a local maximum. Since the local maximum was due to noise, the likelihood of such an occurrence could probably




Figure 9.4: Rotation (blue stars) and translation (red) errors over 20 trials using a BFGS quasi-Newton algorithm with the three images of the previous figure.

be reduced by, for instance, the commonly-used technique of inserting a random perturbation to assure optimality of the current estimate. Also of note is that the 5 trials that were less than excellent had log likelihood values that were less than those of the other 15, i.e., the error is due to an inadequacy of the optimization algorithm for finding the maximum rather than to poor representation of the shape by the log likelihood.

Figure 9.5 shows snapshots of the progress of the algorithm during a trial with an initial guess 4 mm and 4 deg away from the true pose. Each snapshot shows the surface relative to the image planes and the three difference images (difference between the simulated image and the mean image). The difficulty of the problem is evident in the first snapshot, where the mean images show cross-sections of the structure that are different than those in the the simulated images, i.e., the pose estimation problem is not just a translation in any of these images. In finding the pose, the largest movement occurs in the first 10 to 20 iterations, but even at 30 and 40 iterations, slight improvement is still occurring. This minimal movement near the convergence is on the order of 0.5 mm. At this level, the images show the detail that allows for such precise estimation of the pose, as seen in the plots of Figure 9.4.




Figure 9.5: Scenes from an optimization from 4 mm, 4 degrees

9.4 Implications for Inference

The conditions for these trials were relatively simple compared to possible conditions. Images were simulated directly from the model used to predict their variation. No surrounding tissue was present. The relations of the image planes to each other were known. At the same time, however, inferring pose this well in these conditions is quite a feat. Only three noisy images were used, where the images contained a mix of speckle texture and coherent echoes to represent the very complex geometric surface of a vertebra. The results are great testimony to the value of the image model in representing shape in ultrasound. As such, they motivate further investigation and development of the image model for shape in other tissue structures as well as for inference of other characteristics of the tissue and also the imaging system.

Chapter 10

Conclusions and Future Directions

10.1 Conclusions

This dissertation includes contributions that extend beyond its primary focus on inference of gross surface shape. The combination of system characteristics with gross tissue shape and microstructure into a comprehensive, pixel-based model has not been accomplished before. By starting at the physical level and extending the representation to a probabilistic model, many potentially new areas are opened.

The physical model of Chapters 4 and 5 is the basis for all of the models. Shape, microstructure and system characteristics have been included in the model in such a way that the level of detail in any component of the model can be adjusted to fit any application of interest. Shape representations other than surfaces, e.g., volumes, points and curves, could easily be investigated via the simulation methods. Similarly, other microstructural representations and system characteristics, including effects of tissue, could be investigated. The physical model was developed to provide a mathematical basis for the probabilistic, pixel-based models of Chapters 6, 7 and 8. These probabilistic models, too, could be modified to fit the application of interest. For example, if the specificity of the generalized k-distribution is necessary for an application, the random phasor sum basis of the model could be adjusted for the correct fit.

The simulation environment permitted by the physical model is truly more than a means for visually comparing results. For the complex interactions underlying ultrasonic imaging, simulation does not simply confirm intuition with a mathematical basis, it also provides a visual means for observing and understanding those interactions and their effects. As seen in Chapter 9, the simulations also provide an

environment isolated from some of the practical difficulties associated with actual data to allow for more precise investigation of practical algorithms for inference.

The comprehensive, probabilistic image models described for surfaces in Chapter 7 were presented here for the first time. With their basis in the physical model of image formation, the probabilistic descriptions could be extended in many ways. The pixel basis of these models and the introduction of the tissue shape and microstructure with system characteristics provides a new approach to inference of any of the included components. The results of Chapter 9 indicate that very good results for inference of pose can be achieved using the approach with well-known algorithms for optimization. For other characteristics of the tissue and system, similar applications could be developed.

For inference of shape and any other application, development of the model concurrently with the inference algorithms is imperative. Performance in inference then becomes the quantitative measure by which the model and algorithms are evaluated. Similarly, that performance can be used to direct improvement of the models and algorithms to achieve desired results. This relation between the basis of the model and its evaluation is a main advantage of model-based image analysis. Development under that paradigm has a good chance of not only leading to robust algorithms for inference but also a better understanding of underlying phenomena and the exploitation of that understanding in algorithm development.

10.2 The Physical Models for Image Formation

The physical model for image formation could be modified in endless ways. More or less sophisticated methods could be used for characterizing the system. Different models for tissue microstructure could be used. Other shape models could be included. The physical model is general enough that any of these modifications could be made. Which ones are made will depend on the application of interest.

Many effects could be included within the system model. Perhaps most interesting are the effects of tissue, e.g., attenuation and phase aberration, on the system description. In a simple description, attenuation could be modeled in the amplitude of the PSF. Phase aberration could similarly be defined as having an impact on the PSF amplitude and widths. Such simple models could be useful here because of the probabilistic, pixel-based nature of the model. Knowledge about the likely variation of the attenuation and phase aberration could be incorporated into the probabilistic

model as a prior. Inference of the effects could then be implemented with a maximum a posteriori (MAP) approach. On an even simpler level, the imaging system used to collect the data in this dissertation used a simpler envelope detection scheme than the Hilbert transform used in the physical model. The result is that the axial resolution in the actual images is a little worse than in the simulated images. By including a more nearly accurate description of this process in the physical model, a better match may be obtained between the simulated and actual images.

The tissue models also have plenty of room for modification. For realistic images of shape, the most crucial change would be the addition of adjacent tissue structures for a description of an entire tissue region. Adding different tissues is easy within the framework. The parameters for the complex Gaussian simply add because of the linearity of the equations for the RF image. The difficulty comes in modeling different tissues because of the complexity of the scattering interactions. For instance, models of the muscle layer would require models for individual fibers at some level. Again, changes such as these would require application-specific incentives to justify the level of effort.

In addition to the shape models, models for the microstructure could be changed in many ways. The discrete-scatterer representation is fundamental to the model, but many tissues can be represented in this way. Parametrization of the microstructure is probably the more interesting component. Random models for the number of scatterers in a resolution cell would have a significant effect on the amplitude statistics. Similarly, the regularity of scatterer spacing could produce significant changes. Microstructural differences such as these would probably have the greatest effect on characterization of the *local* microstructure. For example, for inference of shape, a microstructural model is absolutely necessary, but performance is probably more dependent on the shape characterization than the exact description of the microstructure.

10.3 The Image Model

From a given physical model, the image model is well-defined mathematically, but, as seen in Chapters 7 and 8, choices in computing the amplitude statistics will have a significant effect on the accuracy and computational requirements of the model. For any given application, the associated requirements and acceptable level of accuracy will determine the type of approximations that can be used. For inference of surface

shape, the curvature of the shape is of utmost importance. For a shape that is smooth relative to the size of the PSF, approximations with less accuracy could be used to achieve quicker computation and faster rates of convergence.

Another interesting area with respect to the image model is the issue of neighbor interactions in the probabilistic model. In this work, neighboring pixels were assumed independent to simplify computation. For inference of shape, additional modeling is not likely to be helpful since the shape is well-represented without the texture information. Items of interest such as the system characteristics, however, have a direct influence on the interdependence between neighboring pixels. Inference of these characteristics, thus, could be aided by extending the model to represent texture using, e.g., a Markov random field approach.

10.4 Inference of Shape Using the Image Model

From the results of Chapter 9, the image model appears to hold promise in inference of shape. In terms of tissue shape, the vertebra represents a very challenging case because of its high curvature. Tissue shapes with less curvature would have a better chance for success because the likelihood would vary more smoothly and predictably. Similarly, the flexibility of non-rigid transformations that allow for more variation in the shape would probably allow better performance. Perhaps the most important criteria for a good shape for inference, however, is the contrast between the shape and the surrounding tissue. In clinical spine images, the vertebrae can be very difficult to image because of the low contrast. As a result, spinal registration with ultrasound is less likely to be successful than registration or identification of soft tissue structures with good contrast. In light of this, the treatment guidance applications most likely to benefit from this work would probably be registration for the liver or prostate (because of the simplicity of their gross shapes) and, potentially, the assessment of brain shift (because of the potential for using a prior model to describe likely deformation of the tissue).

10.5 Adaptive Image Formation and Tissue Characterization

Because the image model represents the system characteristics and tissue, any component could potentially be inferred using an appropriate modification of the image model. One area of interest is the adaptive formation of images using information inferred about the tissue characteristics and their effects on the system PSF. By adaptively forming this image, the image itself would provide additional information about the underlying characteristics. Such an approach could potentially provide improved image quality by correcting for degrading effects such as attenuation and phase aberration, while forming estimates of the underlying tissue characteristics at the same time.

Again, any of these applications would require specific modifications to the image model to fit the purpose. The advantage of the image model is its generality. Any of these applications *could* be represented in this new way, combining any desired information about the tissue shape, the tissue microstructure or the system characteristics.

Appendix A

Calculations involving triangles

Triangle geometry

The triangle consists of three vertices $\{v_1, v_2, v_3\}, v_i \in \mathbb{R}^3$. The area of any triangle can be found as half the magnitude of the vector cross product formed by any two of the three vector edges, (recall that the magnitude of the vector cross product equals the area of the parallelogram formed by the vectors, half of which is the area of the triangle).

$$\text{Area}(\Delta) = \frac{1}{2}|(v_2 - v_1) \times (v_3 - v_1)| \quad (\text{A.1})$$

$$= \frac{1}{2}|(v_1 - v_2) \times (v_3 - v_2)| \quad (\text{A.2})$$

$$= \frac{1}{2}|(v_1 - v_3) \times (v_2 - v_3)|. \quad (\text{A.3})$$

The normal to the triangle is obtained by the vector cross product itself. Any triangle can be resampled to produce four equi-area triangles by splitting each edge in half and connecting as in Figure A.1.

Triangle parametrization and surface integrals

Surface integrals over a triangle surface will be computed in many instances. As a surface, the triangle can be parametrized in two dimensions, λ_1, λ_2 as

$$r = \lambda_1(v_2 - v_1) + \lambda_2(v_3 - v_1) + v_1 \quad (\text{A.4})$$




Figure A.1: Any triangle can be resampled into four triangle of equal area by splitting the edges and reconnecting as shown.

where $\lambda_1, \lambda_2 \in [0, 1]$ and $\lambda_1 + \lambda_2 = 1$. With a surface parametrization, $\mathbf{r}(\lambda_1, \lambda_2)$, the integral of any function $f(\mathbf{r})$ over the surface, can be computed as follows [80],

$$\iint_S f(\mathbf{r})dA = \iint_{R_0} f(\mathbf{r}(\lambda_1, \lambda_2)) |\mathbf{r}_{\lambda_1} \times \mathbf{r}_{\lambda_2}| d\lambda_1 d\lambda_2 \quad (\text{A.5})$$

where \times denotes the cross product, \mathbf{r}_{λ_1} and \mathbf{r}_{λ_2} denote the partial derivatives of the transformation with respect to λ_1 and λ_2 , and R_0 is the subset of \mathbb{R}^2 that is mapped onto the surface S , by $\mathbf{r}(\lambda_1, \lambda_2)$. For the triangle parametrization of equation A.4, the cross product, $|\mathbf{r}_{\lambda_1} \times \mathbf{r}_{\lambda_2}|$ is given by twice the area of the triangle,

$$|\mathbf{r}_{\lambda_1} \times \mathbf{r}_{\lambda_2}| = |(v_2 - v_1) \times (v_3 - v_1)| = 2\text{Area}(\Delta). \quad (\text{A.6})$$

Any integral over the triangle can be computed in the λ_1, λ_2 coordinates by multiplying by twice the triangle area,

$$\iint_{S_\Delta} f(\mathbf{r})dA = 2\text{Area}(\Delta) \int_0^1 \int_0^{1-\lambda_2} f(\mathbf{r}(\lambda_1, \lambda_2))d\lambda_1 d\lambda_2. \quad (\text{A.7})$$

Specific integrals for phasor sum calculations

In representing the phasor sum, certain integrals of trigonometric functions over the triangle surface will be useful in analytic form. The following functions are of interest,

$\cos(r)$, $\sin(r)$, $\cos^2(r)$, $\sin^2(r)$ and $\cos(r)\sin(r)$. The argument in each will depend only on the axial, z , component as $r = \omega z$. The integrals can be found through simple calculations, or with an application such as MathematicaTM. They are listed here for reference, where $z_{31} = (z_3 - z_1)$, $z_{12} = (z_1 - z_2)$, and $z_{23} = (z_2 - z_3)$,

$$\iint_{S_\Delta} \cos(\omega z) dA = 2\text{Area}(\Delta) \frac{z_{31} \cos(\omega z_2) + z_{12} \cos(\omega z_3) + z_{23} \cos(\omega z_1)}{\omega^2 z_{12} z_{23} z_{31}} \quad (\text{A.8})$$

$$\iint_{S_\Delta} \sin(\omega z) dA = 2\text{Area}(\Delta) \frac{z_{31} \sin(\omega z_2) + z_{12} \sin(\omega z_3) + z_{23} \sin(\omega z_1)}{\omega^2 z_{12} z_{23} z_{31}} \quad (\text{A.9})$$

$$\begin{aligned} \iint_{S_\Delta} \cos^2(\omega z) dA = \\ 2\text{Area}(\Delta) \frac{2\omega^2 z_{12} z_{23} z_{31} + z_{31} \cos(2\omega z_2) + z_{12} \cos(2\omega z_3) + z_{23} \cos(2\omega z_1)}{8\omega^2 z_{12} z_{23} z_{31}} \end{aligned} \quad (\text{A.10})$$

$$\begin{aligned} \iint_{S_\Delta} \sin^2(\omega z) dA = \\ 2\text{Area}(\Delta) \frac{2\omega^2 z_{12} z_{23} z_{31} + z_{31} \sin(2\omega z_2) + z_{12} \sin(2\omega z_3) + z_{23} \sin(2\omega z_1)}{8\omega^2 z_{12} z_{23} z_{31}} \end{aligned} \quad (\text{A.11})$$

$$\begin{aligned} \iint_{S_\Delta} \cos(\omega z) \sin(\omega z) dA = \\ 2\text{Area}(\Delta) \frac{z_{31} \sin(2\omega z_2) + z_{12} \sin(2\omega z_3) + z_{23} \sin(2\omega z_1)}{8\omega^2 z_{12} z_{23} z_{31}}. \end{aligned} \quad (\text{A.12})$$

Appendix B

A Parametrization for the Intersection of a Plane and Ellipsoid

For a given surface, a parametrization, $\mathbf{r}(u, v)$, was defined as in Chapter 3 as a mapping from a subset of \mathbb{R}^2 to \mathbb{R}^3 , $\mathbf{r} : A \rightarrow \mathbb{R}^3, (u, v) \in A \subset \mathbb{R}^2$,

$$\mathbf{r}(u, v) = \begin{bmatrix} x(u, v) \\ y(u, v) \\ z(u, v) \end{bmatrix}. \quad (\text{B.1})$$

The utility of the matrix expression for the ellipsoid is that it simplifies the task of finding that parametrization for the intersection of the ellipsoid with an arbitrary plane. The parametrization can be found in a manner similar to those in [82] by constructing a sequence of transformation that transforms the intersection surface into a circle in the x-y plane.

In matrix notation, the plane is described by a normal, $\mathbf{N} = [N_x \ N_y \ N_z]$, and distance, d , to the origin,

$$\mathbf{N}\mathbf{x} = d \quad (\text{B.2})$$

where \mathbf{N} is assumed to be a unit normal, i.e., $|\mathbf{N}| = 1$. With this notation, the intersection of plane and ellipsoid can be derived using the following transformations on the coordinates \mathbf{x} .

1. Scale so that the ellipsoid becomes a sphere. The plane is also transformed and then intersects the sphere in one of the following ways: not at all, a point, or a circle [82].
2. Rotate and translate to align the approximating plane with the xy ($z = 0$) plane. If the intersection exists, it is either a point (not of interest) or a circle.
3. If the intersection exists, the parametrization is then defined as a transformation of the intersection circle (the set, A , required for the parametrization), back to the original coordinate system.

The first coordinate transformation from \mathbf{x} to \mathbf{x}' is a scaling by the inverse of the ellipsoid matrix,

$$\mathbf{x} = \mathbf{M}^{-1}\mathbf{x}' \quad (\text{B.3})$$

to transform the ellipsoid into a sphere. Since $\mathbf{x}^t = \mathbf{x}'^t\mathbf{M}^{-1}$ and $\mathbf{M}^{-1}\mathbf{M} = \mathbf{I}$, The ellipsoid becomes a sphere with unit radius,

$$\mathbf{x}^t\mathbf{M}^2\mathbf{x} = 1 \quad (\text{B.4})$$

$$\mathbf{x}'^t\mathbf{M}^{-1}\mathbf{M}^2\mathbf{M}^{-1}\mathbf{x}' = 1 \quad (\text{B.5})$$

$$\mathbf{x}'^t\mathbf{x}' = 1. \quad (\text{B.6})$$

The plane, $[\mathbf{N}, d]$, becomes

$$\mathbf{N}\mathbf{M}^{-1}\mathbf{x}' = d \quad (\text{B.7})$$

which can be rewritten as the plane $[\mathbf{N}', d']$, where

$$\mathbf{N}' = \frac{\mathbf{N}\mathbf{M}^{-1}}{|\mathbf{N}\mathbf{M}^{-1}|} = \frac{[aN_x \quad bN_y \quad cN_z]}{\sqrt{a^2N_x^2 + b^2N_y^2 + c^2N_z^2}} \quad (\text{B.8})$$

and

$$d' = \frac{d}{|\mathbf{N}\mathbf{M}^{-1}|}. \quad (\text{B.9})$$

In the \mathbf{x}' coordinate system, the problem becomes one of finding the intersection of the unit sphere $\mathbf{x}'^t\mathbf{x}' = 1$ and a plane, $\mathbf{N}'\mathbf{x}' = d'$. By choosing a transformation

that aligns the plane with the $x'y'$ plane, the intersection (if it exists) will either be a circle (with known radius) or a point.

The $x'y'$ plane has normal $\mathbf{N}' = [0 \ 0 \ 1]$ and $d = 0$, and the desired transformation can be constructed by, first, rotating the arbitrary plane so that it is parallel to the $x'y'$ plane, ($\mathbf{N}'' = [0 \ 0 \ 1]$), then translating it to the the origin ($d'' = 0$). The transformation can be constructed as follows, in terms of a rotation matrix, \mathbf{R} , and translation vector, \mathbf{t} ,

$$\mathbf{x}' = \mathbf{R}(\mathbf{x}'' + \mathbf{t}) \quad (\text{B.10})$$

where \mathbf{R} and \mathbf{t} are chosen such that substitution into the plane equation $\mathbf{N}'\mathbf{x}' - d' = 0$ produces the equation

$$\mathbf{N}''\mathbf{x}'' = 0 \quad \text{with} \quad \mathbf{N}'' = [0 \ 0 \ 1] = \mathbf{N}'\mathbf{R}. \quad (\text{B.11})$$

It is easy to verify that the following matrices, R_y and R_x , are rotations around the y and x axes (see Chapter 3), thus the product is also a rotation matrix, \mathbf{R} , and it satisfies the relation $\mathbf{N}'\mathbf{R} = [0 \ 0 \ 1]$,

$$\mathbf{R} = \mathbf{R}_x\mathbf{R}_y \quad (\text{B.12})$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{N'_z}{\sqrt{N'_y'^2 + N'_z'^2}} & \frac{N'_y'}{\sqrt{N'_y'^2 + N'_z'^2}} \\ 0 & -\frac{N'_y'}{\sqrt{N'_y'^2 + N'_z'^2}} & \frac{N'_z'}{\sqrt{N'_y'^2 + N'_z'^2}} \end{bmatrix} \begin{bmatrix} \sqrt{N'_y'^2 + N'_z'^2} & 0 & N'_x \\ 0 & 1 & 0 \\ -N'_x & 0 & \sqrt{N'_y'^2 + N'_z'^2} \end{bmatrix}. \quad (\text{B.13})$$

Substituting these expressions back into the equation for the arbitrary plane allows a translation to be chosen to complete the transformation,

$$\mathbf{N}'\mathbf{x}' = d' \quad (\text{B.14})$$

$$\mathbf{N}'\mathbf{R}(\mathbf{x}'' + \mathbf{t}) = d' \quad (\text{B.15})$$

$$\mathbf{N}''\mathbf{x}'' = d' - \mathbf{N}''\mathbf{t} = d'' = 0. \quad (\text{B.16})$$

Since $\mathbf{N}'' = [0 \ 0 \ 1]$, a translation vector of $\mathbf{t} = [0 \ 0 \ d']^t$ will complete the transformation of the original plane to the $x''y''$ plane.

Combining the above transformations results in a single transformation mapping \mathbf{x} to \mathbf{x}'' or vice versa,

$$\mathbf{x} = \mathbf{M}^{-1}\mathbf{x}' \text{ and } \mathbf{x}' = \mathbf{R}_x\mathbf{R}_y(\mathbf{x}'' + \mathbf{t}) \quad (\text{B.17})$$

$$\mathbf{x} = \mathbf{M}^{-1}\mathbf{R}_x\mathbf{R}_y(\mathbf{x}'' + \mathbf{t}) \quad (\text{B.18})$$

$$\text{or} \quad (\text{B.19})$$

$$\mathbf{x}'' = \mathbf{R}_y^t\mathbf{R}_x^t\mathbf{M}\mathbf{x} - \mathbf{t}. \quad (\text{B.20})$$

It remains to transform the sphere by the second transformation so that the intersection in the $x''y''$ plane can be found. Recall that this intersection completes the definition of the parametrization of the surface of intersection (it will comprise the set $A \in \mathbb{R}^2$ that is transformed to the original coordinates). In the \mathbf{x}' coordinate system, the sphere satisfies the equation $\mathbf{x}'^t\mathbf{x}' = 1$. The following intermediate expression (without the scaling term) relates \mathbf{x}' and \mathbf{x}'' ,

$$\mathbf{x}' = \mathbf{R}(\mathbf{x}'' - \mathbf{t}). \quad (\text{B.21})$$

Substituting this expression into the unit sphere equation produces an equation for the sphere in the \mathbf{x}'' coordinate system,

$$(\mathbf{R}(\mathbf{x}'' + \mathbf{t}))^t(\mathbf{R}(\mathbf{x}'' + \mathbf{t})) = 1 \quad (\text{B.22})$$

$$(\mathbf{x}'' + \mathbf{t})^t\mathbf{R}^t\mathbf{R}(\mathbf{x}'' + \mathbf{t}) = 1 \quad (\text{B.23})$$

$$(\mathbf{x}'' + \mathbf{t})^t(\mathbf{x}'' + \mathbf{t}) = 1 \quad (\text{B.24})$$

$$\mathbf{x}''^t\mathbf{x}'' + 2\mathbf{t}^t\mathbf{x}'' + \mathbf{t}^t\mathbf{t} = 1 \quad (\text{B.25})$$

$$\mathbf{x}''^t\mathbf{x}'' + 2\mathbf{t}^t\mathbf{x}'' = 1 - d'^2. \quad (\text{B.26})$$

In the \mathbf{x}'' coordinate system, the plane of interest is the $x''y''$ ($z'' = 0$) plane, thus the intersection of the sphere and plane is obtained by setting $z'' = 0$ in equation B.26,

$$x''^2 + y''^2 = 1 - d'^2. \quad (\text{B.27})$$

Of course, this expression describes a circle for $1 - d'^2 > 0$ (or for $d' < 1$). The surface of intersection in the original coordinate system can now be parametrized as a transformation, $\mathbf{r}(u, v)$, or a mapping, $\mathbf{r} : A \rightarrow \mathbb{R}^3$ for $A = \{(u, v) : u^2 + v^2 \leq$

$$1 - d'^2\} \subset \mathbb{R}^2,$$

$$\mathbf{r}(u, v) = \mathbf{M}^{-1}\mathbf{R}(\mathbf{x}'' + \mathbf{t}) \quad \text{where} \quad \mathbf{x}'' = \begin{bmatrix} u & v & 0 \end{bmatrix}. \quad (\text{B.28})$$

The parametrization can be represented more simply by substituting the translation, $\mathbf{t} = \begin{bmatrix} 0 & 0 & d' \end{bmatrix}^t$, and employing a new vector, $\mathbf{u} = \begin{bmatrix} u & v & d' \end{bmatrix}^t$,

$$\mathbf{r}(u, v) = \mathbf{M}^{-1}\mathbf{R}\mathbf{u}. \quad (\text{B.29})$$

From Equation 7.14, the surface integrals of interest require the cross product of the partial derivatives, \mathbf{r}_u and \mathbf{r}_v , from the parametrization of Equation B.29. In terms of the quantities of the plane and PSF envelope, the parametrization is given by

$$\mathbf{r}(u, v) = \mathbf{M}^{-1}\mathbf{R} \begin{bmatrix} u \\ v \\ d' \end{bmatrix} \quad (\text{B.30})$$

$$= \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & c \end{bmatrix} \begin{bmatrix} \sqrt{N_y'^2 + N_z'^2} & 0 & N_x' \\ -\frac{N_x'N_y'}{\sqrt{N_y'^2 + N_z'^2}} & \frac{N_z'}{\sqrt{N_y'^2 + N_z'^2}} & N_y' \\ -\frac{N_x'N_z'}{\sqrt{N_y'^2 + N_z'^2}} & -\frac{N_y'}{\sqrt{N_y'^2 + N_z'^2}} & N_z' \end{bmatrix} \begin{bmatrix} u \\ v \\ d' \end{bmatrix} \quad (\text{B.31})$$

$$= \begin{bmatrix} a\sqrt{N_y'^2 + N_z'^2} & 0 & aN_x' \\ -\frac{bN_x'N_y'}{\sqrt{N_y'^2 + N_z'^2}} & \frac{bN_z'}{\sqrt{N_y'^2 + N_z'^2}} & bN_y' \\ -\frac{cN_x'N_z'}{\sqrt{N_y'^2 + N_z'^2}} & -\frac{cN_y'}{\sqrt{N_y'^2 + N_z'^2}} & cN_z' \end{bmatrix} \begin{bmatrix} u \\ v \\ d' \end{bmatrix}. \quad (\text{B.32})$$

The partial derivatives in u and v are easily found as the first and second columns, respectively, of the last matrix,

$$\mathbf{r}_u = \begin{bmatrix} a\sqrt{N_y'^2 + N_z'^2} \\ -\frac{bN_x'N_y'}{\sqrt{N_y'^2 + N_z'^2}} \\ -\frac{cN_x'N_z'}{\sqrt{N_y'^2 + N_z'^2}} \end{bmatrix} \quad \mathbf{r}_v = \begin{bmatrix} 0 \\ \frac{bN_z'}{\sqrt{N_y'^2 + N_z'^2}} \\ -\frac{cN_y'}{\sqrt{N_y'^2 + N_z'^2}} \end{bmatrix}. \quad (\text{B.33})$$

The determinant of the cross product gives the integration term, $|\mathbf{r}(u, v)|$,

$$|\mathbf{r}_u \times \mathbf{r}_v| = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ a\sqrt{N_y'^2 + N_z'^2} & -\frac{bN_x'N_y'}{\sqrt{N_y'^2 + N_z'^2}} & -\frac{cN_x'N_z'}{\sqrt{N_y'^2 + N_z'^2}} \\ 0 & \frac{bN_z'}{\sqrt{N_y'^2 + N_z'^2}} & -\frac{cN_y'}{\sqrt{N_y'^2 + N_z'^2}} \end{vmatrix} \quad (\text{B.34})$$

$$= |\mathbf{i}(bcN_x') + \mathbf{j}(acN_y') + \mathbf{k}(abN_z')| \quad (\text{B.35})$$

$$= \sqrt{b^2c^2N_x'^2 + a^2c^2N_y'^2 + a^2b^2N_z'^2}. \quad (\text{B.36})$$

References

- [1] C. Barbe, J. Troccaz, B. Mazier, and S. Lavalée. Using 2.5D echography in computer assisted spine surgery. *IEEE Engineering in Medicine and Biology Proceedings*, pages 160–61, 1993.
- [2] J.W. Trobaugh, P.J. Kessman, D.R. Dietz, and R.D. Bucholz. Ultrasound in image fusion: A framework and applications. *Proceedings of the IEEE Ultrasonics, Ferroelectrics and Frequency Control Symposium*, 1997.
- [3] S. Lavalée. Registration for computer-integrated surgery: Methodology, state of the art. In *Computer-Integrated Surgery*, chapter 5, pages 77–97. The MIT Press, Cambridge, Massachusetts, 1995.
- [4] E.A. Ashton. *Segmentation and Feature Extraction Techniques, with Applications to Biomedical Images*. PhD thesis, University of Rochester, 1995.
- [5] E.A. Ashton. Multiple resolution Bayesian segmentation of ultrasound images. *Ultrasonic Imaging*, 17:291–304, 1995.
- [6] J.M.B. Dias. Wall position and thickness estimation from sequences of echocardiographic images. *IEEE Transactions on Medical Imaging*, 15(1):25–38, 1996.
- [7] I.L. Herlin, D. Bereziat, G. Giraudon, C. Nguyen, and C. Graffigne. Segmentation of echocardiographic images with Markov random fields. *Lecture Notes in Computer Science*, 801:201–206, 1994.
- [8] U. Grenander and M.I. Miller. Representation of knowledge in complex systems. *J. R. Statist. Soc.*, 56(4):549–603, 1994.
- [9] M.I. Miller, G. Christensen, Y. Amit, and U. Grenander. Mathematical textbook of deformable neuroanatomies. *Proceedings of the National Academy of Science*, 90(24), December 1993.

- [10] A. Srivastava, M.I. Miller, and U. Grenander. Ergodic algorithms on Special Euclidean groups for ATR. In C.I. Byrnes, B.N. Datta, C.F. Martin, and D.S. Gilliam, editors, *Systems and Control in the Twenty-First Century*, pages 327–350. Birkhauser, 1997.
- [11] R.F. Wagner, S.W. Smith, J.M. Sandrik, and H. Lopez. Statistics of speckle in ultrasound B-scans. *IEEE Transactions on Sonics and Ultrasonics*, 30(3):156–163, 1983.
- [12] P.M. Shankar. A model for ultrasonic scattering from tissues based on the k distribution. *Physics in Medicine and Biology*, 40:1633–1649, 1995.
- [13] U.R. Abeyratne, A.P. Petropulu, and J.M. Reid. On modeling the tissue response from ultrasonic B-scan images. *IEEE Transactions on Medical Imaging*, 15(4):479–490, 1996.
- [14] F.L. Lizzi, M. Astor, E.J. Feleppa, M. Shao, and A. Kalisz. Statistical framework for ultrasonic spectral parameter imaging. *Ultrasound in Medicine and Biology*, 23(9):1371–1382, 1997.
- [15] F.S. Cohen, G. Georgiou, and E.J. Halpern. Wold decomposition of the backscatter echo in ultrasound images of soft tissue organs. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 44(2):460–472, 1997.
- [16] T.A. Tuthill, R.H. Sperry, and K.J. Parker. Deviations from Rayleigh statistics in ultrasonic speckle. *Ultrasonic Imaging*, 10:81–89, 1988.
- [17] L.C. Gilman. First-order statistics of pulsed-sinusoid backscatter from random media: Basic elements of an exact treatment. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 44(4):798–804, 1997.
- [18] R.M. Cramblitt and M.R. Bell. Marked regularity models. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 46(1):24–34, 1999.
- [19] S.I. Esses. *Textbook of Spinal Disorders*. J.P. Lippincott Company, U.S.A., 1995.
- [20] M.P. Heilbrun. Foreword. *Neurosurgery Clinics of North America: Clinical Frontiers of Interactive Image-guided Neurosurgery*, 7(2):xv–xvi, 1996.
- [21] Friedman W.A. and Bova F.J. The University of Florida radiosurgery system. *Surg-Neurol*, 32(5):334–42, 1989.

- [22] Friedman W.A., Bova F.J., and Mendenhall W.M. Stereotactic radiosurgery. *Medical Progress through Technology*, 18(4):239–51, 1992-93.
- [23] K.T. Foley and M.M. Smith. Image-guided spine surgery. *Neurosurgery Clinics of North America*, 7(2):171–187, 1996.
- [24] K.T. Foley, K.R. Smith, and R.D. Bucholz. Image-guided intraoperative spinal localization. In *Intraoperative Neuroprotection*, pages 325–340. Williams and Wilkins, 1996.
- [25] F.L. Hansjoerg and R.K. Hauser. Percutaneous endoscopic lumbar spine fusion. *Neurosurgery Clinics of North America: Percutaneous Spine Techniques*, 7(1):107–118, 1996.
- [26] H.H. Mathews and B.E. Mathern. *Percutaneous Procedures in the Lumbar Spine*, pages 731–745. Williams and Wilkins, 1998.
- [27] S. Lavallee, J. Troccaz, P. Sautot, B. Mazier, P. Cinquin, P. Merloz, and J.P. Chirossel. Computer-assisted spinal surgery using anatomy-based registration. In *Computer-Integrated Surgery*, chapter 32, pages 425–449. The MIT Press, Cambridge, Massachusetts, 1995.
- [28] D.A. Simon, M. Hebert, and T. Kanade. Techniques for fast and accurate intra-surgical registration. *Journal of Image-Guided Surgery*, 1:17–29, 1995.
- [29] D.A. Simon. *Fast and Accurate Shape-Based Registration*. PhD thesis, Carnegie Mellon University, 1996.
- [30] T.D. Mast, L.M. Hinkelman, M.J. Orr, V.W. Sparrow, and R.C. Waag. Simulation of ultrasonic pulse propagation through the abdominal wall. *Journal of the Acoustical Society of America*, 102(2):1177–1190, 1997.
- [31] R.W. Prager, R.N. Rohling, A.H. Gee, and L. Berman. Rapid calibration for 3-D freehand ultrasound. *Ultrasound in Medicine and Biology*, 24(6):855–869, 1998.
- [32] U. Grenander. *Elements of Pattern Theory*. Johns Hopkins University Press, Baltimore; London, 1996.
- [33] U. Grenander. *General Pattern Theory*. Oxford University Press, Inc., New York, 1993.

- [34] U. Grenander, Y. Chow, and D. Keenan. *HANDS: A Pattern Theoretic Study of Biological Shapes*. Springer-Verlag, New York, 1990.
- [35] A. Srivastava. *Inferences on Transformation Groups Generating Patterns on Rigid Motions*. PhD thesis, Washington University in St. Louis, 1996.
- [36] S.C. Joshi, M.I. Miller, and U. Grenander. On the geometry and shape of brain sub-manifolds. *International Journal of Pattern Analysis and Artificial Intelligence. Special Issue on processing MR Images.*, 1997.
- [37] G. Christensen. *Deformable Shape Models for Anatomy*. PhD thesis, Washington University in St. Louis, 1994.
- [38] K.E. Mark. *Markov Random field Models for Natural Language*. PhD thesis, Washington University in St. Louis, 1997.
- [39] J.W. Haller, G.E. Christensen, S.C. Joshi, J.W. Newcomer, M.I. Miller, J.C. Csernansky, and M.W. Vannier. Hippocampal MR imaging morphometry by means of general pattern matching. *Radiology*, 199(3):787–791, June 1996.
- [40] A. Srivastava. Automated tracking-recognition. Master’s thesis, Washington University in St. Louis, 1993.
- [41] B.O. O’Neill. *Elementary Differential Geometry*. Academic Press, San Diego, California, 1966.
- [42] S.C. Joshi. *Large Deformation Diffeomorphisms and Gaussian Random Fields for Statistical Characterization of Brain Sub-Manifolds*. PhD thesis, Washington University in St. Louis, 1998.
- [43] W.M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry*. Academic Press, Orlando, Florida, 1975.
- [44] N. Khaneja. Geometry of cortical features. Master’s thesis, Washington University in St. Louis, 1996.
- [45] B. Hamann. Curvature approximation for triangulated surfaces. *Supplementum to Computing*, 8:139–153, 1993.
- [46] A. Watt. *3D Computer Graphics*. Addison-Wesley Publishing Company, Wokingham, England, 1993.

- [47] J.M. McCarthy. *Introduction to Theoretical Kinematics*. MIT Press, USA, 1990.
- [48] O. Faugeras. *Three-Dimensional Computer Vision*. MIT Press, USA, 1993.
- [49] E. Vanmarcke. *Random Fields*. The MIT Press, Cambridge, Massachusetts, 1983.
- [50] G.R. Cross and A.K. Jain. Markov random field texture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5(1):25–39, 1983.
- [51] M.S. Bazaraa, H.D. Sherali, and C.M. Shetty. *Nonlinear Programming: Theory and Algorithms*. John Wiley and Sons, United States of America, 1993.
- [52] P.E. Gill, W. Murray, and M.H. Wright. *Practical Optimization*. Academic Press, London, 1981.
- [53] J.N. Wright. Image formation in diagnostic ultrasound. Short Course, IEEE International Ultrasonics Symposium, 1997.
- [54] A. Macovski. *Medical Imaging Systems*. Prentice-Hall, Inc., U.S.A., 1983.
- [55] P.M. Morse and K.U. Ingard. *Theoretical Acoustics*. McGraw-Hill, Inc., U.S.A., 1968. Ch. 7, 8.
- [56] R.M. Arthur and S.R. Broadstone. Imaging via inversion of ellipsoidal projections of solutions to the linear acoustic wave equation. *IEEE Transactions on Medical Imaging*, 8(1):89–95, 1989.
- [57] M.F. Insana and D.G. Brown. Acoustics scattering theory applied to soft biological tissues. In K.K. Shung and G.A. Thieme, editors, *Ultrasonic Scattering in Biological Tissues*, pages 75–124. CRC Press, U.S.A., 1993.
- [58] J.W. Goodman. *Introduction to Fourier Optics*. The McGraw-Hill Companies, Inc., U.S.A., 2nd edition, 1988.
- [59] D.H. Johnson and D.E. Dudgeon. *Array Signal Processing*. P T R Prentice-Hall, Inc., U.S.A., 1993.
- [60] G.F. Roach. *Green's Functions*. Cambridge University Press, Cambridge, 2nd edition, 1982.

- [61] S.J. Norton and M. Linzer. Ultrasonic reflectivity imaging in three dimensions: Exact inverse scattering solutions for plane, cylindrical, and spherical apertures. *IEEE Transactions on Biomedical Engineering*, BME-28(3):202–220, 1981.
- [62] D.L. Liu and R.C. Waag. Propagation and backpropagation for ultrasonic wavefront design. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 44(1):1–13, 1997.
- [63] J.C. Lockwood and J.G. Willette. High-speed method for computing the exact solution for the pressure variations in the nearfield of a baffled piston. *Journal of the Acoustical Society of America*, 53(3):735–741, 1973.
- [64] P.R. Stepanishen. Transient radiation from pistons in an infinite planar baffle. *Journal of the Acoustical Society of America*, 49(5):1629–1638, 1970.
- [65] J.W. Goodman. *Statistical Optics*. John Wiley and Sons, U.S.A., 1985.
- [66] K.K. Shung and G.A. Thieme, editors. *Ultrasonic Scattering in Biological Tissues*. CRC Press, U.S.A., 1993.
- [67] P. Beckmann and A. Spizzichino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. The MacMillan Company, New York, 1963.
- [68] J.A. Ogilvy. *Theory of wave scattering from random rough surfaces*. Institute of Physics Pub., Philadelphia, 1991.
- [69] J.A. Ogilvy. Computer simulation of acoustic wave scattering from rough surfaces. *Journal of Physics D: Applied Physics*, 21:260–277, 1988.
- [70] B.J. Dean. *Angular Spectrum-Based Statistical Formulation of Ultrasound Scattering by Rough Surfaces*. PhD thesis, Worcester Polytechnic Institute, 1997.
- [71] L. Clifford. *On the First-Order Amplitude Statistics of Myocardial Ultrasonic Backscatter*. PhD thesis, Dartmouth College, 1995.
- [72] D.L. Liu and R.C. Waag. Harmonic amplitude distribution in a wideband ultrasonic wavefront after propagation through human abdominal wall and breast specimens. *Journal of the Acoustical Society of America*, 101(2):1172–1183, 1997.

- [73] M. Nakagami. The m-distribution - a general formula of intensity distribution of rapid fading. In W.C. Hoffman, editor, *Statistical Methods in Radio Propagation*, pages 3–36. Pergamon, 1960.
- [74] K.D. Donohue, F. Forsberg, C.W. Piccoli, and B.B. Goldberg. Analysis and classification of tissue with scatterer structure templates. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 46(2):300–309, 1999.
- [75] D. Phillips and K.J. Parker. Evaluation of a three-dimensional fractal model for scattering in the liver. In *22nd International Ultrasonics Symposium*, 1997.
- [76] J.W. Trobaugh and R.M. Arthur. A simulation study of variability in ultrasonic images of vertebrae. In *23rd International Ultrasonics Symposium*, 1998.
- [77] T.T. Elvins. A survey of algorithms for volume visualization. *Computer Graphics*, 26(3):194–201, 1992.
- [78] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw Hill, U.S.A., 1991.
- [79] E. Parzen. *Stochastic Processes*. Holden-Day, Inc., USA, 1962.
- [80] E. Kreyzig. *Advanced Engineering Mathematics*. John Wiley and Sons, U.S.A., 1993.
- [81] D.L. Snyder and M.I. Miller. *Random Point Processes in Time and Space*. Springer-Verlag New York, Inc., U.S.A., 1991.
- [82] J. Roe. *Elementary Geometry*. Oxford University Press, Great Britain, 1993.

Vita

Jason W. Trobaugh

- Date of Birth** May 1, 1969
- Place of Birth** Fort Leavenworth, Kansas
- Degrees** B.S. Cum Laude, Electrical Engineering, May 1991
M.S. Electrical Engineering, May 1993
D.Sc. Electrical Engineering, August 2000
- Professional Societies** IEEE
Eta Kappa Nu
Tau Beta Pi
- Publications** J.W. Trobaugh and R.M. Arthur. A Discrete-Scatterer Model for Ultrasonic Images of Rough Surfaces. Accepted for publication in *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*.
- J.W. Trobaugh and R.M. Arthur. Simulation of Ultrasonic Images of Rough Surfaces Using a Parametrized Discrete-Scatterer Model. In *24th International Ultrasonics Symposium*, 1999.
- J.W. Trobaugh and R.M. Arthur. A Simulation Study of Variability in Ultrasonic Images of Vertebrae. In *23rd International Ultrasonics Symposium*, 1998.
- J.W. Trobaugh, P.J. Kessman, D.R. Dietz and R.D. Bucholz. Ultrasound in Image Fusion: A Framework and Applications. In *Proceedings of the IEEE Ultrasonics, Ferroelectrics and Frequency Control Symposium*, 1997.
- R.D. Bucholz, D.D. Yeh, J.W. Trobaugh et al. The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. *CVRMed-MRCAS '97 Proceedings*, March 1997, 459-66.

J.W. Trobaugh, D.J. Trobaugh, and W.D. Richard. Three-Dimensional Imaging with Stereotactic Ultrasonography. *Computerized Medical Imaging and Graphics*, 18:5, 315-323, 1994.

J.W. Trobaugh, W.D. Richard, K.R. Smith and R.D. Bucholz. Frameless Stereotactic Ultrasonography - Method and Applications. *Computerized Medical Imaging and Graphics*, 18:4, 235-246, 1994.

J.W. Trobaugh. *Frameless Stereotactic Ultrasonography in Neurosurgery*. Master's thesis, Washington University in St. Louis, May 1993.

R.D. Bucholz, J.W. Trobaugh and W.D. Richard. Three-Dimensional Intraoperative Ultrasonography Using an Optical Digitizer. *Proceedings of the 61st Annual Meeting of the American Association of Neurosurgeons*, 61:149; 1993.

August 2000