A MODALITY INDEPENDENT APPROACH TO ELASTICITY IMAGING

By

Chad Wayne Washington

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

Biomedical Engineering

August, 2003

Nashville, Tennessee

Approved by:

Michael I. Miga

Robert J. Roselli

ACKNOWLEDGEMENTS

I consider myself to be extemely fortunate to have had the opportunity to work in the Biomedical Modeling Lab. Dr. Michael Miga has been a great help through his advisement in both academics and research. I extend a tremendous amount appreciation to him for this. Thanks also to Dr. Robert Roselli for his assessment of this work.

Along with Dr. Miga, the other members our lab, Dave Cash, Prashanth Dumpuri, and Colin Caughran, have made this a rewarding experience. Last, but not least, well he is actually smaller than everyone else, Tuhin Sinha. Without his help in linux and coding in general, this work would have not been possible. I really do owe him a steak.

Two people outside of the lab, but still partakers of lab lunches, that have supported me with entertainment and a little MATLAB help are Dr. Mark Bray and Steve Gebhart.

Also, my family has been a tremendous help with their support. Whether it was listening to me complain or providing much needed home cooking, they were there for whatever I needed. I pray that they will never have a need for my research.

TABLE OF CONTENTS

Page

LIST OF FIGURES

CHAPTER I

BACKGROUND AND LITERATURE REVIEW

I.1 Breast Cancer

Cancer is a term that has unfortunately become an ever-present part of our society's vocabulary. Its origin has been credited to the early Greek physician Hippocrates (460-370 B.C.), who used the terms *carcinos* and *carcinoma* to describe different types of tumors. The reference to a crab is believed to have most likely resulted from the tumors' finger-like spreading projections which seemed to almost mimic the shape of a crab. While cancer has become synonymous with a broad range of today's most grievous of diseases, the fact is that many of these illnesses are quite different and vary greatly with respect to their origins, symptoms, and plans of treatment. However, it is generally accepted that the commonality found among all "cancers", is the uncontrollable proliferation of abnormal cells. Today, discoveries are being made on an almost continual basis into the secrets of such pathologies. Yet, in spite of this, cancer remains one of the leading causes of death in the United States and the world [1].

One of the most prolific forms of cancer found today was also the first to be formally diagnosed, breast cancer. Early Egyptian writings (1600 B.C.) describe treating ulcers of the breast with cauterization and they go further to say that "There is no treatment" for the disease [1]. Fortunately, for breast cancer, this idea of "no treatment" has since been replaced by a strong sense of hope for recovery. This hope is greatly due to an increased awareness of the disease which has led to improvements in screening programs and significant advances in methods of detection and treatment. With this awareness there has also come the realization of women's high susceptibility to this disease. One of breast cancer's most formidable characteristics, is its disregard for ethnicity, socio-economic status, or country of origin. Hence, women now live with the anxiety of knowing that their chances for acquiring this disease only increase with age [2].

Technically, breast cancer is the unconstrained growth of cells within breast tissue. This tissue

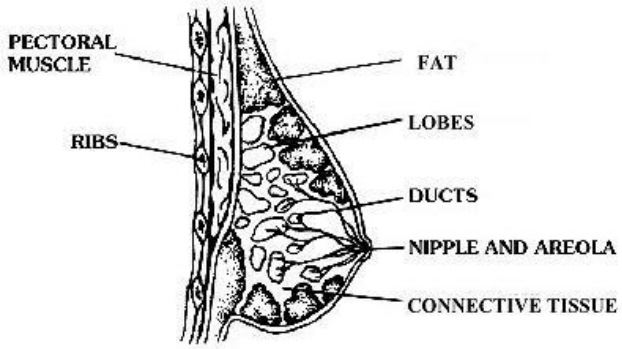

Figure I.1: Anatomical view of female human breast.

consists of milk producing glands, known as *lobules*, and a system of *ducts*. These ducts connect the lobules to the nipple. Fatty, connective, and lymphatic tissues make up the remaining portions of the breast. A simplified anatomical view of the breast can be seen in Figure I.1 [3]. Generally, breast cancer makes itself apparent through the formation of tumors. However, not all tumors are cancerous. In fact, most are not. These tumors are considered to be *benign* ; that is, while they are abnormal they do not grow and spread like cancer. Of those tumors which are cancerous there are two main types, *in situ* and *invasive*. *In situ* describes breast cancers which are contained within the ducts or lobules, known as *ductal carcinoma in situ* and *lobular carcinoma in situ*, respectively. These types most often times are fully treatable. Invasive tumors, however, are considered to be a more serious type of cancer. These masses have broken through the ductal or lobular walls and have infiltrated the surrounding tissue. The seriousness of invasive breast cancer is described by the stages - local, regional, or distant. These refer to the extent to which the cancer has spread. The local stage describes cancers confined to the breast, regional those which have spread to the lymph nodes, and distant where the tumors have metastasized [4].

In a report released by the American Cancer Society, a woman in the United States has a 1 in 8 lifetime risk of developing breast cancer, and these odds could increase depending on factors such as age, family history, and genetic predisposition. It also presented statistics illustrating the strong connection between survival rate and the stage of cancer development at diagnosis(local-

ized (96.4%), regional (77.7%), and distant (21.1%))[1]. Along with the obvious benefits in lifetime expectancy, early detection also provides for significant economic savings as well. Johnston, reports a possible 20.9% decrease in health related costs through appropriate screenings [5]. This is a substantial gain considering that in 1996 an estimated 5.4 billion dollars was spent on medical treatment of breast cancer in the United States alone. At 13.1% of all monies, this was the largest expenditure for one individual type of cancer [6]. With such evidence suggesting the benefits of early breast cancer detection, there has been a concerted effort in the medical community to develop imaging methods and modalities that not only have a high sensitivity and specificity to breast lesions but ones that are also cost-effective.

I.2 Elastography

One such method that has shown great promise in fulfilling these characteristics is elastography. Elastography (the imaging of tissue mechanical properties) is based on the premise that tissue pathology can be directly assessed through the quantification of tissue mechanical properties. There has been a longstanding precedence among physicians that stiffness can be used as an indicator of possible cancerous lesions. In fact, Krouskop et al. [7] reported data showing significant differences between the elastic modulus, ($E = \frac{\sigma}{\varepsilon}$, which is the ratio of stress, σ, to strain, *strain*, and can be considered as a measure of "stiffness") of healthy tissue, benign, *in situ*, and invasive breast tumors. Following this, a report was released by the Harvard BioRobotics Laboratory stating that infiltrating ductal cancer can be as much as 70 times stiffer than normal fat tissue and 5 times as stiff compared to healthy glandular tissue [8]. Information such as this suggests that having the ability to accurately quantify these values, non-invasively, offers the possibility of not only detecting but also diagnosing breast abnormalities.

An extensive body of literature regarding ultrasound (US) research for assessing tissue stiffness was seen in the late 1980's. It was during this time investigators began to notice a distinct correlation between US images and pathology of tissue. These studies focused on non-invasively characterizing tissue through the analysis of tissue motion [9, 10, 11, 12, 13, 14, 15, 16]. The basis

[1] These statistics represent the 5-year survival rate

for a new imaging modality was formed from this work and made its first appearance in the pages of *Ultrasonic Imaging* in April of 1991 [17]. The term *elastography* was first keyed by Ophir et al. in this paper and demonstrated a connection between strain profiles and elastic modulus. Since then, there has been a proliferation of studies into elasticity imaging. Thus far the two main areas of investigation have focused on ultrasound and magnetic resonance elastography (USE and MRE, respectively).

I.2.1 Ultrasound Elastography (USE)

Ophir et al. [17] describe how a cross-correlation analysis of pre- and post-compression A-line US pairs could be used to compute a strain profile of the perturbed tissue, and from this, deduce the elastic modulus distribution within the tissue. From here, this group continued to research and expand on this idea by looking into various applications and the clinical feasibility of such techniques [18, 19, 20]. These studies strongly suggested that elastography could be developed into an invaluable tool for soft-tissue characterization. However, along with the attractiveness of its applicability, they also noted that a number of significant barriers must be overcome in order to take full advantage of elastography. The first was the high sensitivity of the elastographic signal-to-noise ratio (SNR) to the intrinsic characteristics of US signal. Secondly, decorrelation of the pre- and post compressed US signals results at modest levels of strain. Along with these, additional corruption, i.e. numerical round-off and etc., of the data is caused by the signal processing techniques. Subsequently, a consequence of these factors is that the implementation of this method is extremely difficult.

In an attempt to address these matters, O'Donnell et al. presented a method to help decrease the presence of signal decorrelation at large strains by shortening the correlation window [21]. Skovoroda et al., in order to take advantage of the increased SNR between displacement and strain at large deformations, implemented an elastographic technique that used a non-linear displacement-strain relationship [22]. Others are continuing in this vein, and are trying to account for these SNR and decorrelation artifacts through the introduction of analytical techniques in the estimation of strain [23, 24].

Due to difficulty in implementation many have shifted away from the correlation technique. Skovoroda et al. developed a reconstruction method which makes use of the US images, rather than the signals, to estimate internal tissue displacements and strains, and from these calculate a modulus distribution [25]. Yet others are attempting to use techniques such as a least squares approach to strain estimation [26] and spectral strain estimators [27], which are notably less precise, but have the advantage of being more robust than cross-correlation.

All previously mentioned methods have been dependent upon a "static" deformation being applied externally to the tissue of interest. Also, their performance is a function of accurate displacement measurements. Hence, this "manual" application of a deformation adds an additional level of complexity to such techniques by increasing their chances for error, i.e. patient movement, incidental high strain, along with numerous other factors. To remedy this a variety of techniques have been proposed. Of these perhaps two of the most exciting are those using focused ultrasound and/or low-frequency shear waves [28, 29, 30, 31, 32].

Sarvazyan et al. make use of both of these tools in their implementation of "Shear Wave Elasticity Imaging" (SWEI) [33]. They use the radiation force of focused ultrasound to remotely induce low frequency shear waves. The internal motion resulting from these waves is then measured allowing for strain estimation. Compared to other methods of elasticity imaging this technique allows for the highly localized application of tissue deformation, and thus allowing for evaluation of very precise regions of interest. A second technique that is showing great promise is "Sonoelastography" presented by Taylor et.al [34]. This method propagates low-amplitude, low-frequency shear waves through internal organs. Using Doppler imaging techniques, hard inhomogeneities can be detected due to the fact that the amplitude of the shear wave will decrease at the location of the tumor. Also, by taking advantage of image segmentation techniques, a full 3D tumor volume can be visualized, which helps to account for the 2D limitations of ultrasound.

I.2.2 Magnetic Resonance Elastography (MRE)

While USE has shown great promise, it is in reality a two-dimensional imaging modality facing a three-dimensional problem, and there is a degree of error associated with this fact. From this, one

can easily recognize that acquisition of 3D data would be extremely beneficial. Realization of such a technique was first introduced in a 1995 issue of *Science* by Muthupillai et al [35]. This group took advantage of MRI's ability to detect phase shifts when tissue is oscillated within a magnetic field gradient. The resultant phase shift allows for the calculation of the internal deformation field. The significance of these measurements is that they are three-dimensional. Using these displacements allowed for the accurate calculation of mechanical tissue properties and "Magnetic Resonance Elastography" (MRE) was created.

In addition to this, many teams are working to establish new approaches to MRE. For example, Plewes et al. use a deformation method similar to Muthupillai in conjunction with phase-contrast enhanced MR imaging. Chenevert et al., in order to simplify the physics behind the elastic model, introduced a method using MRI to detect the internal displacements of tissue under "quasi-static" deformation [36]. Extensions of this work using stimulated echo NMR [37] have followed. Then others, such as Van Houten et al. have introduced a zonal approach to elastography [38] within the context of harmonic shear wave elastography and have taken advantage of MRI's ability to produce localized data. This allows for the property reconstruction of a given region of interest.

MR Elastography has yet to be implemented consistently in a clinical setting. There have been, however, numerous investigations into its detection and diagnosis capabilities along with the feasibility of its application. Thus far, most have found it to be extremely promising [39], especially for property determination in specific regions [40]. Recently, McKnight et al. [41], demonstrated MRE's ability to quantitatively delineate breast tissue properties. Using MRE they were able to detect that the shear stiffness of a breast carcinoma was 418% higher than the surrounding tissues. Another promising study was released by Lorenzen et al. [42]. In this study they gave statistics indicating MRE's ability to distinguish between healthy and diseased tissue. They also found that there can be overlap in stiffness properties between malignant and benign tumors. Others have investigated its application into areas such as the assessment of thermal tissue ablation [43]. While all studies found magnetic resonance elastography to be of great promise, they also indicated some shortcomings, such as specificity and difficulty in clinical application, that must be overcome.

I.2.3 Inverse Elastography Techniques

Thus far, what has been presented are techniques dealing mainly with the application and determination of tissue deformations. In order for these measurements to be of any value, methods must be developed which are capable of deriving "elastograms" (stiffness images) from this data. There is a large body of literature dealing with this topic, but the ones most relevant to this work deal with solving the inverse problem. In practice, an inverse problem is one in which a constitutive behavior is assumed and a set of *a priori* data are given. Then, through methods of optimization a parameter set that characterizes the constitutive behavior is derived such that if used will produce the *a priori* data.

One such method was given by Kallel and Bertrand [44]. In their proposed technique, they use a Newton-Raphson based optimization and a linear elastic finite element model (FEM) to construct a set of shear modulus parameters. In their reconstruction scheme, they are attempting to match a set of axial tissue displacement fields acquired with US. With this method they found that they were able to reduce artifacts found in more traditional strain images. They also discuss the ill-conditioned aspects of the problem, requiring the introduction of some type of regularization mechanism. A similar method was described by Doyley et al. [45], with added spatial filtering which helped to improve the stability of the reconstruction.

Another approach was introduced by Manduca et al. [46]. Given a set low-frequency shear wave deformations acquired through MR Elastography, they were able to calculate the complete measurements of wave displacements through the direct solution of the wave equation using Green's function integrals. Using these displacements they were then able to iteratively reconstruct an elasticity map of the perturbed tissue. They found that this inverse technique, similar to the method above, was able to remove many of the resolution issues and artifacts associated with the more common forward approach.

Further researchers have expanded on these initial investigations and have implemented a wide variety of techniques in order to more accurately and efficiently reconstruct elasticity parameters. Some of these include the use of maximum-likelihood estimators [47], three-dimensional FEM's

[48], and regional constraints [49]. Regardless of the technique each of the above methods are dependent upon their *a priori* information. They all need accurate displacements, and unfortunately, internal tissue deformations continue to be extremely difficult to measure, especially within the clinical setting.

I.3 Similarity Based Elastography

As mentioned above, perhaps the greatest limiting factor of elastography has been the requirement of accurately measured internal displacements. This necessity has significantly hindered elastography's implementation into the clinical setting. From this, a logical step would be to develop an approach which is independent of this need. Such a method was presented by Miga in a 2003 issue of *Physics in Medicine and Biology*, which described a mutual information/finite element based form of elastography (the details of mutual information along with other similarity metrics will be discussed within the methods section) [50]. This method represents a shift from those presented above in that it removes the necessity of having to accurately measure internal tissue deformations.

This movement away from a displacement driven reconstruction is not unfounded. Quite the contrary, Sarvazyan et al. noted that the information needed for tissue properties was contained within the images themselves, independent of any modality [51]. In their method they considered the relative motion of tissue elements, and based their reconstruction on "motion pattern". A similar idea was described by Fowlkes et al. [52]. They also discuss "relative motion" and note its independence from any imaging modality. However, these should still be considered displacement based techniques, as they used MR tagging [53], which uses special imaging sequences to calculate displacements. Also, image similarity has long been used in the realm of image registration, which is essentially the estimation of movement based on images. A very relevant application of this technique was presented by Rueckert et al. in their implementation of a non-rigid registration approach driven by mutual information [54]. Their algorithm makes use of free-form deformations to account for tissue motion between subsequent imaging scans of breast tissue.

I.4 Purpose

The purpose of this work is to present preliminary results of image similarity/finite element based elastography. This new approach to elastography does not focus on refining techniques to measure displacements but rather uses the similarity, or likeness, between images to drive the reconstruction. The benefits of such a method are its simplicity in clinical implementation (i.e. no special imaging sequences), low-cost, and its independence of any particular imaging modality. The uniqueness of this approach is that the stiffness properties are considered to be functions of image similarity, and thus allows for the problem to be cast into the realm of non-rigid image registration.

CHAPTER II

A MODALITIY INDEPENDENT APPROACH TO ELASTICITY IMAGING

II.1 Introduction

In a 2001 report released by the American Cancer Society, a woman in the United States had a 1 in 8 lifetime risk of developing breast cancer, and these odds increased depending on factors such as age, family history, ethnicity, and genetic predisposition. It also presented statistics illustrating the strong connection between survival rate and stage of cancer development at diagnosis (localized (96.4%), regional (77.7%), and distant (21.1%))[1] [4]. Along with the obvious benefits in lifetime expectancy, early detection also provides for significant economic savings as well. Johnston, reports a possible 20.9% decrease in health related costs through appropriate screenings [5]. This is a substantial gain considering that in 1996 an estimated 5.4 billion dollars was spent on medical treatment in the United States alone [6]. With such evidence suggesting the benefits of early breast cancer detection, there has been a concerted effort in the medical community to develop imaging methods and modalities that not only have a high sensitivity and specificity to breast lesions but ones that are also cost-effective.

One such method that has shown great promise is elastography. Elastography (the imaging of tissue mechanical properties) is based on the premise that tissue pathology can be directly assessed through the quantification of tissue mechanical properties. There has been a longstanding precedence among physicians that stiffness can be used as an indicator of possible cancerous lesions. In fact, recent reports have shown significant differences between the modulus of healthy tissue, benign, *in situ*, and invasive breast tumors [7]. Having the ability to accurately quantify these values, non-invasively, offers the possibility of not only detecting but also diagnosing breast abnormalities.

Since initial work presented by Ophir et al. [17], there has been a proliferation of studies into elastic imaging. Thus far the two main areas of investigation are focused on ultrasound and mag-

[1] These statistics represent the 5-year survival rate

netic resonance elastography (USE and MRE, respectively). Within USE, substantial emphasis has been placed on relating strain profiles to the distribution of elastic parameters within tissue. Ophir et al. has produced an extensive amount of work using cross-correlation analysis of pre- and post-compression A-line pairs [17], least squares strain estimation [26], and spectral strain estimators [27]. Another significant advancement came from Muthupillai et al. with their publication in a 1995 issue of *Science* [35]. It introduced a MR method that allowed for the measurement of three-dimensional strain effects created by propagating harmonic mechanical waves. Others have expanded on these initial studies, and a more complete review of current USE and MRE techniques can be found in the June, 2003 issue of *Physics in Medicine and Biology* [47]

The purpose of this paper is to present preliminary results of the similarity-based elastography approach presented by Miga [50]. This method represents a shift from the those presented above in that it removes the necessity of having to accurately measured internal tissue deformations. Rather, it uses the information shared between two images at the anatomic structural level to drive its reconstruction (as opposed to speckle tracking in USE). While iterative schemes have been proposed [44, 45], they require very accurate measurements of the internal deformations. This has shown to be a difficult task, particularly in a clinical setting, and to date is one of the major obstacles that must be overcome before current methods are a feasible option. This movement away from a displacement driven reconstruction is perhaps one of the most attractive characteristics of this technique, and these results help to show its viability as a future clinical screening method.

II.2 Methods

Similarity-based elastography considers the task of determining mechanical tissue properties within the scope of image registration. The essential premise to this technique is that regional image similarity can be used as the driving force in calculating appropriate updates for these properties within a FEM elastic image reconstruction framework. By viewing these properties as solutions to an inverse problem an elegant reconstruction algorithm has been developed. The entire process can be divided into two sections. The first is the acquisition of an appropriate image set and generation of a FE model. The second is the actual optimization procedure used to calculate tissue properties.

The sections that follow discuss these in detail along with other issues dealing with this method, such as algorithm implementation, phantom generation, and software development.

II.2.1 Image Acquisition and FE Model Generation

Prior to an image reconstruction, a two image series is acquired in any traditional imaging modality (i.e. MR, CT, etc.). The first set represents a baseline image of the tissue of interest in its undeformed state. The second set is acquired after a user-prescribed displacement. An example of this image series can be seen in Figure II.1.

These images are then used to construct a FEM mesh along with appropriate boundary conditions. The mesh is generated using the AIM software, which is based on the Algebraic Integer Method, along with supporting software for bandwidth reduction of the model [55]. The mesh is based on a user-defined contour constructed manually from the pre-deformed image. Then Dirichlet and Neumann boundary conditions are determined visually from the images through the use of MATLAB based software developed in the Biomedical Modeling Laboratory (BML) at Vanderbilt University. This is accomplished by first assuming that all boundaries are stress free, i.e. a Neumann condition with a value of zero for each boundary node. The post-compressed image is then used to determine which nodes require Dirichlet boundary conditions. These boundary conditions are calculated based on distance between the mesh boundary and the edge of the imaged tissue. It is important to note that while these are currently defined manually, through the use of some simple image processing techniques, it is plausible that the FE model generation procedures could be automated.

II.2.2 Elastic Parameter Optimization

The general basis of the inverse problem is that a computer model can be used to non-rigidly register the image acquired in the undeformed state to the image acquired in the deformed state by systematically varying the model's material properties (Young's modulus), E. The reconstruction is complete when the acquired deformed image and the model generated deformed image become adequately "similar". The image comparison basis is a regional image similarity measure. Similarity refers to the quantitative measure of the likeness between two images based on pixel

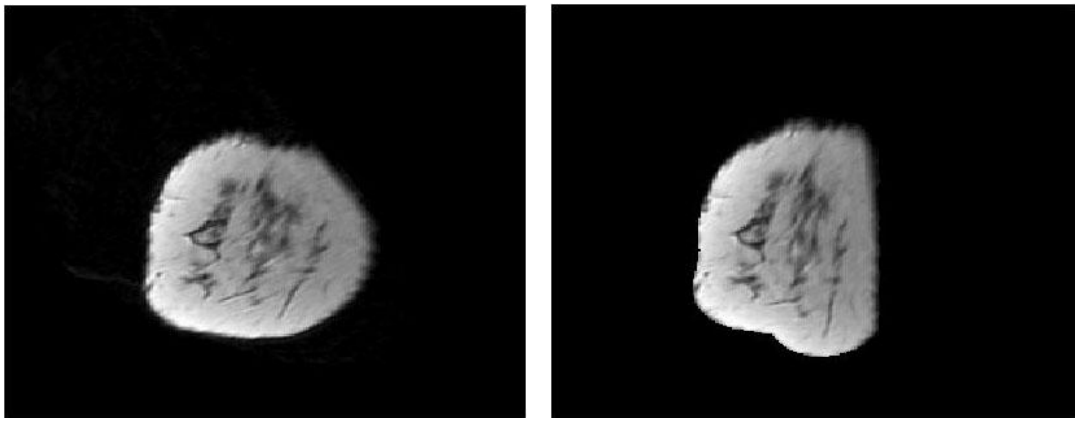

(a) (b)

Figure II.1: Acquired image series (a) MR slice of breast pre-compression, (b) MR slice of breast post-compression

intensity values. Using this idea of regional similarity and casting the problem as a least squares minimization we define our objective function, Φ, as (II.1)

$$\Phi(\vec{E}) = ||\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_E)||^2 \tag{II.1}$$

where $\mathrm{S}(\vec{E}_T)$ is the similarity between regions given the true modulus parameters, $\vec{E}_T$, and $\mathrm{S}(\vec{E}_E)$ is the similarity between corresponding regions given the current estimated modulus values, $\vec{E}_E$. $\mathrm{S}(\vec{E}_T)$ is calculated by comparing the post-compression image to itself. $\mathrm{S}(\vec{E}_E)$ is the comparison of the post-compressed image to the model-generated image. Both $\vec{E}_T$ and $\vec{E}_E$ are $Nx1$ vectors representing the N material properties being reconstructed. It is important to note that (II.1) is a function of the image similarity profile, S, rather than $\vec{u}$, the displacements, which are normally used by other methodologies.

In order to optimize S, we must minimize Φ. To do this we have chosen a Newton-Raphson approach, which is a gradient based optimization method. Therefore, we take the derivative of

(II.1) and set it equal to zero arriving at Equation (II.2)

$$\frac{\partial \Phi}{\partial \vec{E}_E} = (\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_E))\frac{\partial S(\vec{E}_E)}{\partial \vec{E}_E} = 0 \tag{II.2}$$

or in matrix notation it becomes Equation (II.3)

$$\mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_E)) = 0, \tag{II.3}$$

with the *MxN* Jacobian matrix, $\mathrm{J} = \frac{\partial \mathrm{S}(\vec{E})}{\partial \vec{E}}$, where M is the number of similarity measurements.

Using (II.3) we can develop a scheme to iteratively update our material properties. This is done by expanding Equation (II.3) about the true set of properties in Taylor series and specifing $\vec{E}_T \equiv \vec{E}_0 + \Delta\vec{E}$ which gives us (II.4).

$$\mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_T)) = 0 = \mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_0)) + \frac{\partial(\mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_0)))}{\partial \vec{E}}|_{E_0}\Delta\vec{E} + \cdots \tag{II.4}$$

Rewriting this and neglecting higher order terms we arrive at Equation (II.5).

$$0 = \mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_0)) - \mathrm{J}^T\mathrm{J}\Delta\vec{E} \tag{II.5}$$

$$\vec{E}_{k+1} = \vec{E}_k - \frac{\mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_E))}{-\mathrm{J}^T\mathrm{J}}. \tag{II.6}$$

A common nuance of inverse problems is that the matrix $\mathrm{J}^T\mathrm{J}$, also known as the Hessian matrix H, tends to be ill-conditioned. To account for this we have introduced the α term found in the standard Levenberg-Marquardt approach [56]. Our property updates are then calculated from the expression (II.7) seen here written in the matrix form $\mathrm{A}\vec{x} = \vec{b}$ as

$$(\mathrm{H} + \alpha\mathrm{I}) = \mathrm{J}^T(\mathrm{S}(\vec{E}_T) - \mathrm{S}(\vec{E}_E)) \tag{II.7}$$

where α is a regularization factor and I is the identity matrix. The term α is calculated using an

approach similar to that described by Joachimowicz et al. [57] and is given by the expression

$$\alpha = \gamma \left[\frac{1}{N} \sum_{i=1}^{N} H(i,i) \right] (error)^2, \tag{II.8}$$

where γ represents the degree of regularization, this parameter is empirically derived, and *error* is the summed difference between similarities of consecutive iterations. Scaling of the Hessian along with other techniques are used for additional regularization.

II.2.3 Reconstruction Algorithm

The preceding section describes the theoretical basis of our optimization, yet the implementation of such a technique may not be completely intuitive. A diagram of the algorithm can be seen in Figure II.2. Within the optimization, there are five major steps which are involved,

- construction of the Jacobian matrix
- calculating new material properties
- solving the forward problem model
- creation of a model deformed image
- evaluating the resultant image similarity.

Both solving the forward problem and evaluating image similarity will be discussed within their own sections.

Construction of the Jacobian

Constructing the Jacobian matrix (Figure II.3) is an extremely expensive computational process. In doing this, each material property, E_i, must be perturbed. Then the forward problem solved for the corresponding displacements. Using these displacements a model-deformed image is generated. The similarity between this image and the post-compression image is measured, which is then used in the determination of the derivative $\frac{\partial S(E_i)}{\partial E_i}$. A simple backward-difference, first order

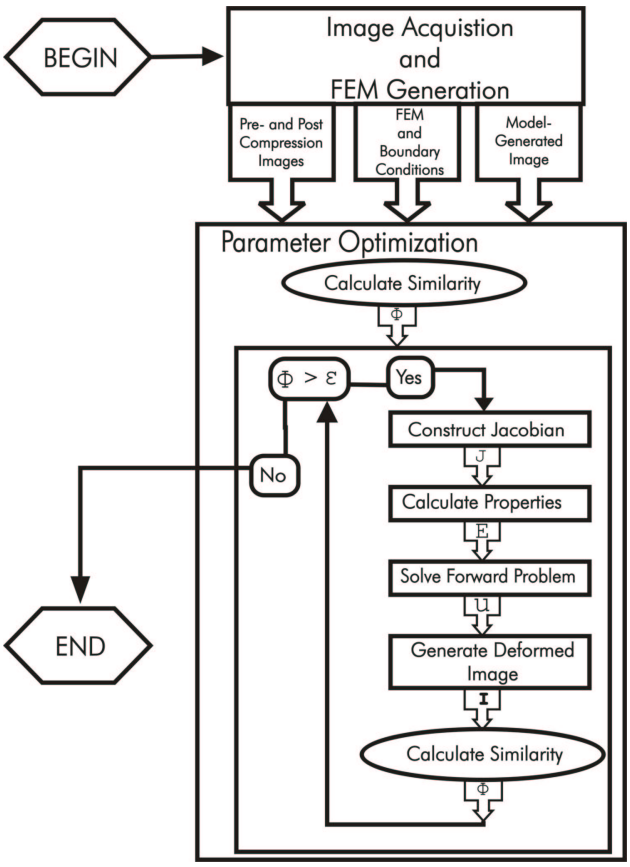

Figure II.2: Reconstruction algorithm.

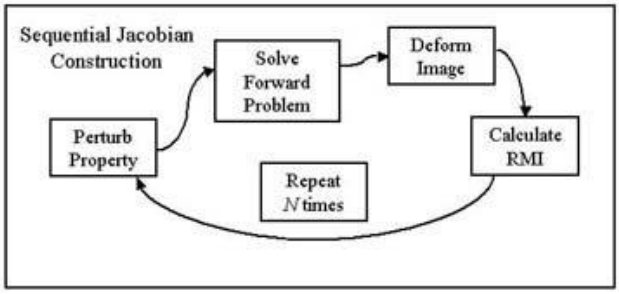

Figure II.3: Sequential Jacobian construction algorithm

approximation is used, given by Equation (II.9).

$$\frac{\partial \mathrm{S}(E_i)}{\partial E_i} = \frac{\mathrm{S}(E_i) - \mathrm{S}(E_{i-1})}{E_i - E_{i-1}} \tag{II.9}$$

The implications of moving to a difference approximation of the derivative is a reduction in the converstion rate of the traditional Newton Raphson method to slightly less than $O(h^2)$ - often called the secant method.

Calculating Material Property Updates

We have already discussed the method for determining our material property updates in the section dealing with the optimization of material properties. Unfortunately, due to the fact that inverse problems are inherently noisy, additional steps need to be taken in order to ensure a well-behaved reconstruction. Therefore, in calculating our material properties we have added a smoothing factor, θ, and a relaxation parameter, δ. For our smoothing factor we have implemented a spatial averaging technique, which is given in Equation (II.10), and a visual representation of this seen in Figure II.4. We define our material properties nodally, therefore in Equation (II.10), n_i represents, given node i, the number of its connected nodes. Relaxation is accomplished by a weighted sum between the current parameter value and the new updated value and is given by Equation (II.11). Both of these techniques combine to help improve numerical stability.

$$\hat{E}_i^k = \theta E_i^k + (1-\theta)\frac{\sum_{j=1}^{n_i} E_j^k}{n_i} \tag{II.10}$$

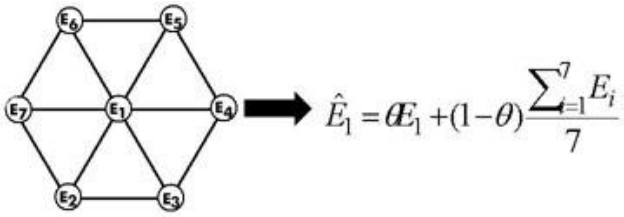

Figure II.4: Spatial averaging technique

$$\hat{E}_i^k = \delta E_i^k + (1-\delta)E_i^{k-1} \tag{II.11}$$

Image Deformation

Image deformation is another computationally intensive process. In order to ensure that the deformed images are consistent with the displacements generated by the model, a reverse mapping technique must be implemented. Therefore, after the FEM model has been updated with the current set of displacements, for each pixel, the element in which it is contained must be found. Generally this process would be an $O(n^2)$ operation. We have improved this significantly by using an ANN k-d search tree [58]. This allows for quick look-up of a pixels associated elements. Once the element is found, a deformation is calculated for the pixel, based on the elemental basis functions. The pixel is then back-projected into undeformed image space and an appropriate intensity value is determined. Figure II.5 represents this process.

II.2.4 Forward Problem

For each inverse problem there is also an associated forward problem (FP). In similarity based elastography, the FP is the prediction of internal displacement fields given a set of boundary conditions and elasticity parameters. The implementation of the FP, as in the case of the optimization algorithm, must be carefully considered. If it is incapable of accurately representing the physics involved in the deformation, then these inaccuracies cascade into the remaining portions of the algorithm. Knowing this, the two most significant factors are the model used to represent the system, and its implementation. We have selected Plane Strain Linear Elasticity as our model, and use the

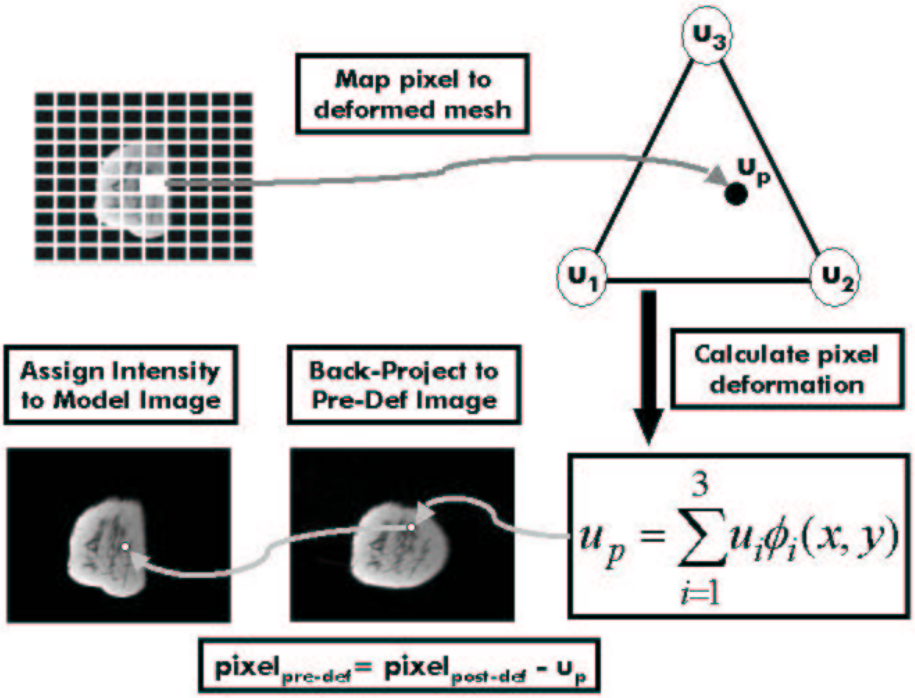

Figure II.5: Image deformation algorithm

finite element method as the mechanism by which the model is evaluated.

Plane Strain Linear Elasticity

The choice to use plain strain theory as the governing model has been made by numerous other investigators in their elastographic imaging attempts, and has been found to produce reasonable approximations of the displacements given small strains [44, 59, 45, 60]. The roots of this theory lie in Hooke's Law represented by Equation (II.12).

$$F = k \cdot u \tag{II.12}$$

This simply states that force, F, and displacements, u, are directly related through some proportionality constant, k. This was later generalized by Cauchy to the form given in Equation (II.13)

$$\sigma = c \cdot \varepsilon, \tag{II.13}$$

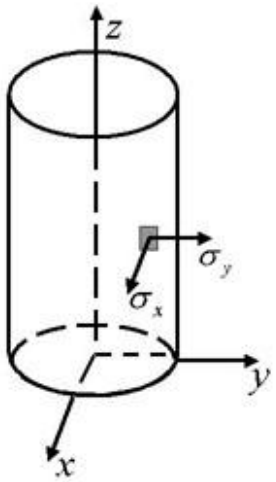

Figure II.6: Plain strain describes the cross-sectional loading of a thick object

which relates stress, σ (the ratio between force and area), to strain, ε (the ratio between Δl and l), by the coefficients c. This fully describes the three-dimensional nature of elastic deformations, and consequently results in c being a set of thirty-six constants, also known as stiffnesses. However, with assumptions of symmetry and isotropy, plain strain linear elasticity simplifies the problem to one with only two constants.

Plane strain is used to describe the cross-sectional loading behavior of a thick object(i.e. Figure II.6) allowing for the two-dimensional analysis of a three-dimensional problem [61]. The five key assumptions being made are:

- mechanical equilibrium
- a symmetric material
- an isotropic material
- strain in z direction is zero
- displacements in (x,y) are planar.

Using these we arrive at

$$\nabla \cdot \sigma = 0, \tag{II.14}$$

where σ is the stress tensor

$$\sigma = \begin{bmatrix} \sigma_x & \tau_{xy} \\ \tau_{xy} & \sigma_y \end{bmatrix} \tag{II.15}$$

and the constitutive relationships seen in Equation (II.16)

$$\begin{Bmatrix} \sigma_x \\ \sigma_y \\ \tau \end{Bmatrix} = \frac{E(1-\nu)}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1 & \frac{\nu}{1-\nu} & 0 \\ \frac{\nu}{1-\nu} & 1 & 0 \\ 0 & 0 & \frac{1-2\nu}{2(1-\nu)} \end{bmatrix} \begin{Bmatrix} \frac{\partial u}{\partial x} \\ \frac{\partial v}{\partial y} \\ \frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} \end{Bmatrix}. \tag{II.16}$$

In (II.16), ν is Poisson's ratio, and E is Young's modulus. Poisson's ratio is a measure of compressibility and for tissue is assumed to be very close 0.5. Young's Modulus is defined as the ratio between stress and strain and is considered to be a measure of stiffness. This is the parameter that we are reconstructing.

Finite Element Methods

Finite element analysis allows for the approximation of solutions to partial differential equations (PDE's) through the discretization of a continuous domain. Subdividing complex systems into well understood "elements" has become a standard technique in various fields for discretely solving what is often times an unsolvable continuous problem. The following is a discussion of plain strain elasticity within the realm of finite elements and FEM issues specific to similarity based elastography. For a more general and complete description see *The Finite Element Method* by Zienkiewicz and Taylor [62].

In considering Equation (II.14) within the finite element framework, we begin by applying what is known as the weighted residual method to (II.14). This technique allows for the approximation of the PDE by multiplying it by a spatially continuous weighting function and then integrating over the problem domain. This results in Equation (II.17)

$$\langle \phi_i \nabla \cdot \sigma \rangle = 0, \tag{II.17}$$

where $\langle\cdot\rangle$ represents integration and ϕ_i is the ith member of a complete set of standard finite element C^0 local Lagrange polynomial interpolants. Treating (II.17) with integration by parts results in the weak form (II.18)

$$-\langle\sigma\cdot\nabla\phi_i\rangle = -\oint\sigma\cdot\hat{n}\phi_i ds, \tag{II.18}$$

where the symbol, $\oint$, indicates a boundary integral and $\hat{n}$ is the outward-pointing normal direction to the boundary. The problem is now spatially discretized following the standard Galerkin formulation. This is accomplished by replacing the stress tensor, σ, with the constitutive relationships (II.16). Then rewriting the displacements u as a sum of known coefficients multiplied by known functions of position, $u(x,y) = \sum_j u_j\phi_j(x,y)$. Substituting this into (II.18) and viewing it at the element level gives

$$\begin{bmatrix} \left\langle \frac{\partial\phi_j}{\partial x}\frac{\partial\phi_i}{\partial x} + \frac{1-2\nu}{2(1-\nu)}\frac{\partial\phi_j}{\partial y}\frac{\partial\phi_i}{\partial y}\right\rangle & \left\langle \frac{\nu}{1-\nu}\frac{\partial\phi_j}{\partial y}\frac{\partial\phi_i}{\partial x} + \frac{1-2\nu}{2(1-\nu)}\frac{\partial\phi_j}{\partial x}\frac{\partial\phi_i}{\partial y}\right\rangle \\ \left\langle \frac{\nu}{1-\nu}\frac{\partial\phi_j}{\partial x}\frac{\partial\phi_i}{\partial y} + \frac{1-2\nu}{2(1-\nu)}\frac{\partial\phi_j}{\partial y}\frac{\partial\phi_i}{\partial x}\right\rangle & \left\langle \frac{\partial\phi_j}{\partial y}\frac{\partial\phi_i}{\partial y} + \frac{1-2\nu}{2(1-\nu)}\frac{\partial\phi_j}{\partial x}\frac{\partial\phi_i}{\partial x}\right\rangle \end{bmatrix} \begin{Bmatrix} u_{xj} \\ u_{yj} \end{Bmatrix} = \begin{Bmatrix} \hat{x}\cdot(-\oint\sigma\cdot\hat{n}\phi_i ds) \\ \hat{y}\cdot(-\oint\sigma\cdot\hat{n}\phi_i ds) \end{Bmatrix} \tag{II.19}$$

Integrating over the entire domain (II.19) may be written in matrix notation as $\mathrm{K}\vec{u} = \vec{b}$, with K being the stiffness matrix, $\vec{u}$ the nodal displacements, and $\vec{b}$ the right hand side.

Generally, the forward problem operates on a single domain, but due to the nature of this problem there is a significant difference between the resolution needed to resolve the physics and the resolution at which parameters can be reconstructed. In response to this we have implemented a dual-mesh scheme similar to that of Paulsen et al. [63]. The elastic properties are defined at the nodes of a coarse "parent-mesh". The resolution of the parent mesh is such that the algorithm is capable of reconstructing its parameters. Unfortunately, this mesh does not allow us to accurately model the physics of the problem. Therefore, we sub-divide each parent element, creating a "sub-mesh". This is the mesh used in the linear elastic model which gives us a better approximation to the displacements. However, a second problem arises. The model requires the stiffness properties to be defined at the sub-element level. To accommodate this, we use the basis functions associated with the parent elements, and interpolate its nodal properties to the mid-point of its corresponding

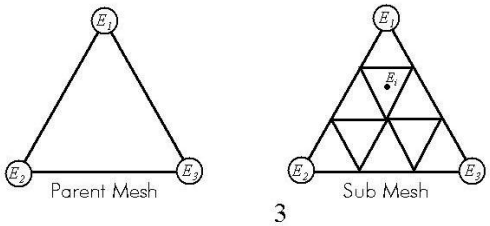

$$E_i = E(x_i, y_i) = \sum_{k=1}^{3} \phi_k(x_i, y_i) E_k$$

Figure II.7: Sub-element properties are determined through interpolation of the Parent-element's nodal properties.

sub-elements. A visual representation of this can be seen in Figure (II.7).

II.2.5 Similarity Metrics

Similarity metrics have developed from a vast body of research of which a significant portion has come from the area of image-processing, particularly medical image-processing. These metrics have formed the basis for what is known as intensity-based registration methods, where intensity refers to the scalar values in image pixels. Rather than using points or surface features, intensity-based methods attempt to optimize the "similarity" between images. Determining how to measure this "similarity" has been and continues to be the focus of numerous studies, and development of such methods is beyond the scope of this work. Our purpose here is to evaluate the effectiveness of such measures in their ability to drive the elastographic reconstruction process. For a more complete understanding of image registration and the role of similarity measures in this process see the *Medical Imaging Handbook* [64]. We have elected to examine four such similarity metrics.

Of the similarity metrics used today, one that is sparking a tremendous amount of interest is mutual information (MI). Its power is derived from its ability to effectively register cross-modality images (i.e. CT to MR, MR to PET, etc.). The origins of MI are found within the field of information theory. Shannon presented a measure, H, known as entropy, which calculates a value for the information contained within a set of parameters $\vec{s}$ based on $P(\vec{s})$, the probabilities of $\vec{s}$ [65]. In a

discrete sense this value is given by the Equation (II.20)

$$H = -\sum_i P(s_i) * log(P(s_i)) \tag{II.20}$$

This was then expanded upon by two groups independently [66, 67] to what is now referred to as MI. The idea demonstrated is a registration technique that operates by minimizing the joint entropy shared between two images I_1 and I_2. This translates into Equation (II.21)

$$MI(I_1, I_2) = H(I_1) + H(I_2) - H(I_1, I_2), \tag{II.21}$$

or in its normalized form

$$MI(I_1, I_2) = \frac{H(I_1) + H(I_2)}{H(I_1, I_2)}. \tag{II.22}$$

The individual entropies and the joint entropy may be calculated through the joint probability distribution function of the two images given the Equations (II.23), (II.24), and (II.25).

$$H(I_1, I_2) = -\sum_{j,k} PDF[j,k] log PDF[j,k] \tag{II.23}$$

$$H(I_1) = \sum_i \left(\sum_k PDF[j,k] log \sum_l PDF[j,l] \right) \tag{II.24}$$

$$H(I_2) = \sum_k \left(\sum_i PDF[i,k] log \sum_j PDF[j,k] \right) \tag{II.25}$$

In theory the model generated image and the post-compression data image should be identical given the correct set of tissue properties and model. Therefore, the difference between the two images will be zero. Working under this assumption the similarity metric known as SSD (sum of squares of differences) would be a reasonable choice for driving the reconstruction. In fact it has been proven that SSD is the optimal measure when two images differ only by Gaussian noise. Unfortunately, most noise within medical imaging is not Gaussian [64]; however, the simplicity of

this measure warrants its investigation. SSD is defined as

$$SSD = \frac{1}{N}\sum_{i}^{N} |I_1(i) - I_2(i)|^2 \quad \forall i \in I_1 \cap I_2 \tag{II.26}$$

Another metric that is elegantly simple is the correlation coefficient (CC) measure. It assumes that the intensity from I_1 are linearly related to those of I_2, and when this is true CC has been shown to be the ideal similarity measure. Also, this is not an unreasonable assumption given intra-modality images. CC can be calculated with the Equation (II.27)

$$CC = \frac{\sum_i (I_1(i) - \bar{I_1})(I_2(i) - \bar{I_2})}{\sqrt{\sum_i (I_1(i) - \bar{I_1})^2 \sum_i (I_2(i) - \bar{I_2})^2}} \quad \forall i \in I_1 \cap I_2, \tag{II.27}$$

where $\bar{I_1}$ and $\bar{I_2}$ are the mean pixel values of images I_1 and I_2.

The fourth and last metric to be evaluated is gradient correlation (GC) [68]. GC differs from previous measures most significantly in the fact that it operates on the gradient of the two images. The advantage of this method is that the process of calculating the gradient images removes the low spatial frequency differences between the two images. These differences are often times a result of soft tissue. Secondly, GC is essentially operating on the edges found within the two images, which seems plausible considering we are trying to determine structural movement. Calculation of GC is performed by first filtering the two images with the Sobel template in the x and y directions. This produces gradient images with respect to x, $\frac{\partial I_1}{\partial x}$ and $\frac{\partial I_2}{\partial x}$, and gradient images with respect to y, $\frac{\partial I_1}{\partial y}$ and $\frac{\partial I_2}{\partial y}$. The CC is then calculated using (II.27) between images $\frac{\partial I_1}{\partial x}$ and $\frac{\partial I_2}{\partial x}$ and images $\frac{\partial I_1}{\partial y}$ and $\frac{\partial I_2}{\partial y}$. GC is the average of these two values. Figures II.8(a) and II.8(b) represent gradient images $\frac{\partial I}{\partial x}$ and $\frac{\partial I}{\partial y}$, respectively.

Generally, image similarity is considered as a global metric. We have taken a different approach in measuring similarity. Our inverse problem needs more than one measurement, and therefore, we have developed a regional approach to calculating similarity. This is represented in Figure II.9. The process is initiated by first defining a bounding region about the tissue in the deformed image. This is done because similarity regions outside of the mesh are not influenced by tissue parameters, and

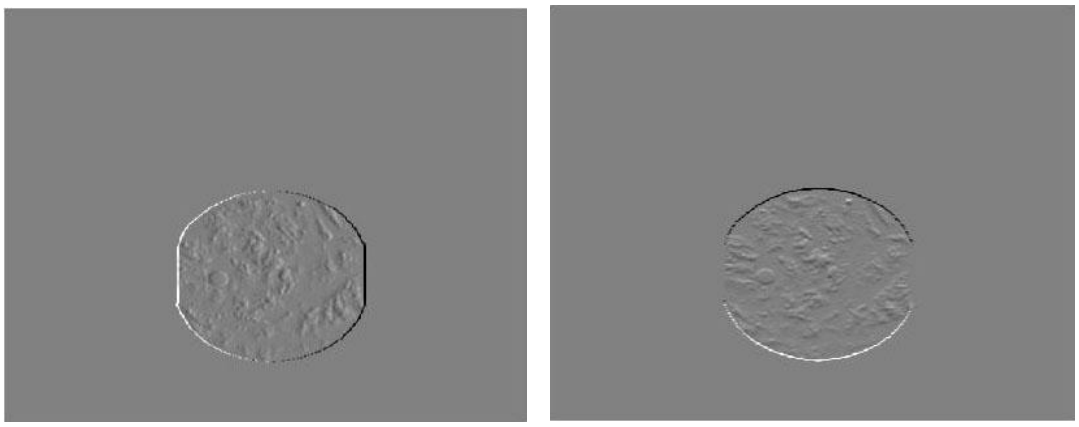

(a) (b)

Figure II.8: Gradients of CT image set (a) Gradient image $\frac{\partial I}{\partial x}$ (b) Gradient image $\frac{\partial I}{\partial y}$

thus gives no useful information. After the bounding box is defined, it is discretized into the regions that will be used in calculating similarity. This discretization is determined by a radius defining the size of the regions. There is one more consideration when determining the areas for calculating similarity. It is whether to include the regions located on the boundary of the mesh, or to only include those which are completely contained within the mesh. This idea of inclusion/exclusion is seen in Figure II.10. The effects of this will be considered within the experiments. Once the regions are defined, similarity is calculated using a chosen metric between corresponding regions of the post-deformed data image and the model-generated image.

II.2.6 Phantom Generation and Deformation

Creating an adequate imaging phantom has shown to be a very difficult process requiring several perturbations in construction procedures. The difficulty arises in trying to recreate the intricate intensity heterogeneity pattern seen in breast images while still maintaining consistent material properties throughout the phantom. Notice in Figure II.1(a) this is the objective of our phantom development. Figures II.11(a) through II.11(d) recount the progression of our attempts to generate a suitable phantom. In addition to the imaging phantom, a compression device compatible with MR/CT and ultrasound imaging modalities has been developed and is shown in Figure II.12.

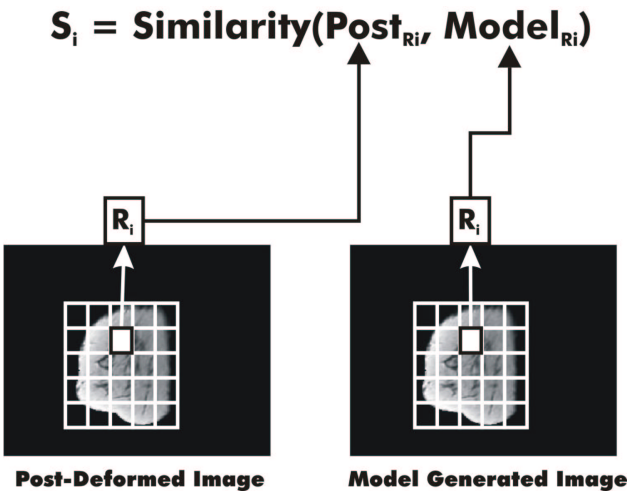

Figure II.9: Calculation of regional image similarity.

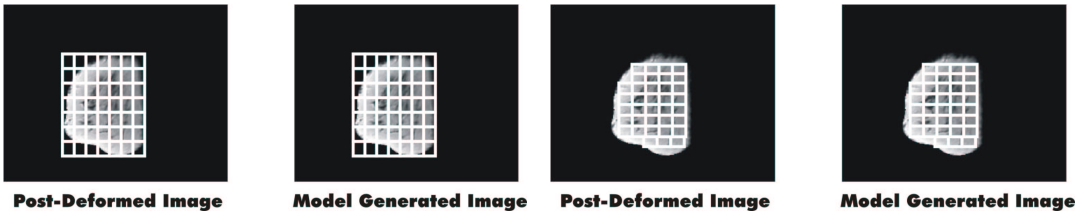

(a) (b)

Figure II.10: Calculation of image similarity (a) with boundary data (b) without boundary data

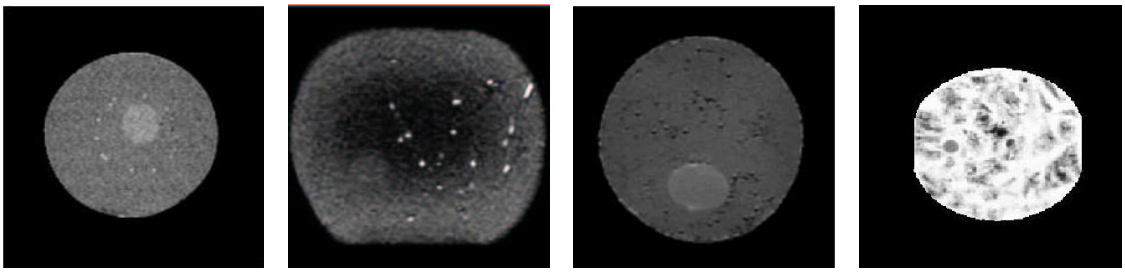

(a) (b) (c) (d)

Figure II.11: Progression in imaging phantom: (a) 1^{st} phantom (cotton twine and gel), (b) 2^{nd} phantom (cotton twine and gel), (c) 3^{rd} phantom (synthetic fibers and glass beads), (d) 4^{th} phantom (synthetic sponge material contained in silicon with wooden dowel rod inclusion). Note, a phantom with a characteristic image similar to II.1(a) is the objective.

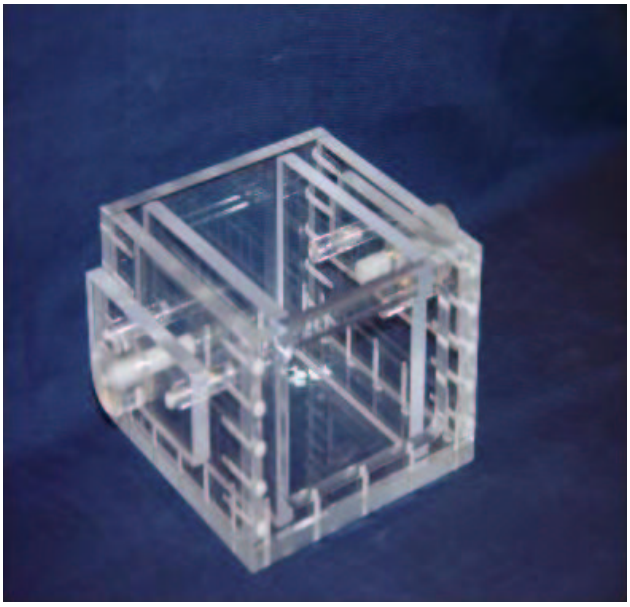

Figure II.12: Compression device compatible in MRI, CT, and ultrasound

II.2.7 Software Development

Before the implementation of this technique could be achieved a framework for code development had to be adopted. The decision was made to work within The Insight Toolkit (ITK) [69]. ITK is an open-source project funded by the National Library of Medicine within the National Institute of Health. The purpose of this program is to support the Visible Human Project, which is attempting to create a complete, anatomically detailed, three-dimensional representation of the human male and female bodies. Specifically, ITK is a toolkit which provides a wide range of segmentation and registration techniques for medical imaging applications. It is implemented in C++ with a generic, object-oriented philosophy (i.e. it operates on an "image" independent of any storage scheme). We maintained this philosophy in the implementation of our algorithm. Also, additional features were added to ITK, such as FEM methods, and integration of the numerical solving library LAPACK [70].

A second issue that arose during the planning phase, was the computational expense of such a method. As previously discussed, construction of the Jacobian matrix and image deformation are very time consuming operations. Therefore, the code was developed such that a parallel implementation could be easily integrated into the pre-existing code when computing resources became

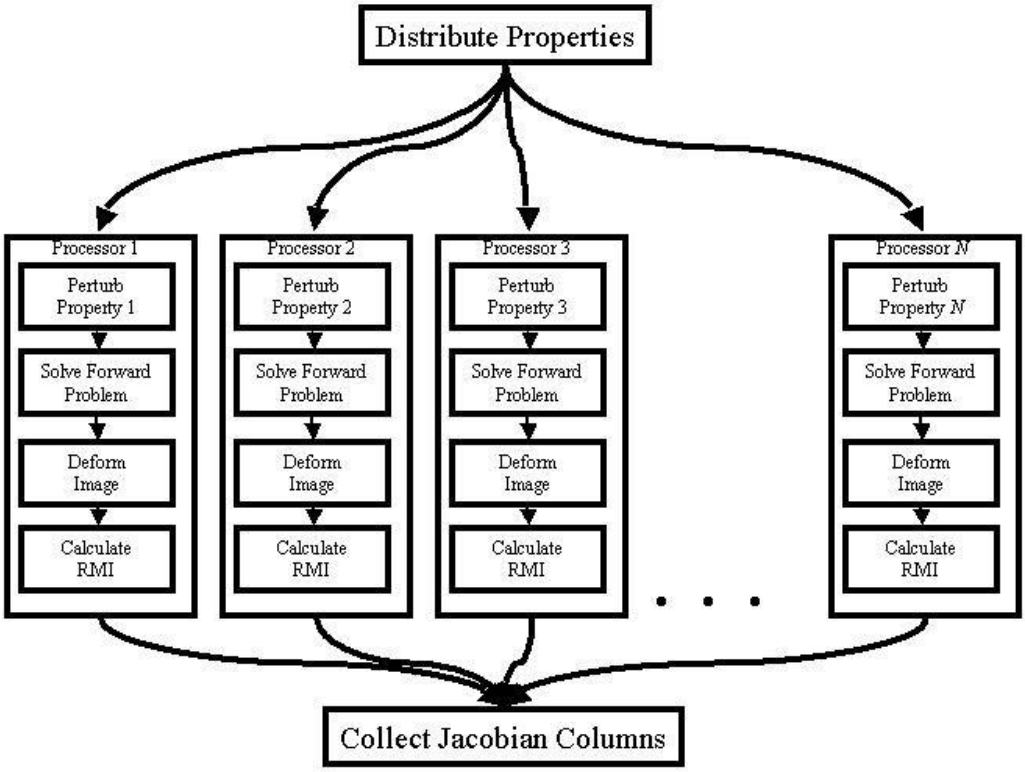

Figure II.13: Parallel algorithm for Jacobian construction

available. The most sensible place to paralyze the code is in the creation of the Jacobian. This is due to the fact that each iteration in this process is independent of all other iterations. From this a new Jacobian algorithm was implemented, seen in Figure II.13. Through the use of MPICH, an open source version of the Message Passing Interface, and the Beowulf computing cluster in the BML this was accomplished [71].

II.2.8 Computational Experiments

The following sections describe the set of experiments performed to assess the effectiveness of the proposed method, and to also more clearly understand the fundamental limits of such a technique. Each experiment falls into one of two categories, simulation or acquired data. Simulation refers to the experiments which use a specified modulus distribution and the finite element model to generate a post-compression image for the reconstruction. This will allow for us to better ascertain

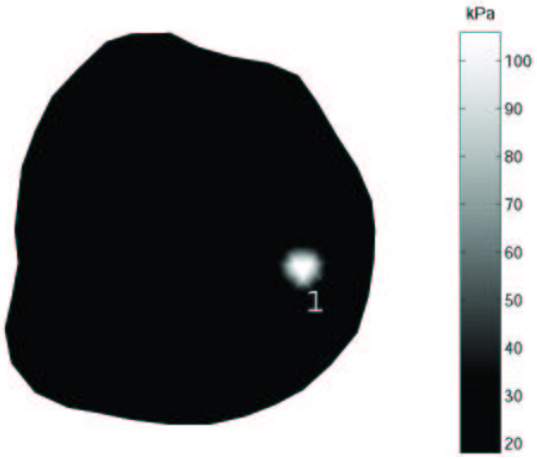

Figure II.14: Modulus distribution used in creating simulated post-compression images.

what factors play important roles in algorithm performance by removing inaccuracies of the model and noise acquired during imaging. Acquired data represents the experiments that use an actual image set taken in either CT or MRI. These tests will give a more realistic determination of the methods applicability into a clinical setting, along with helping to find an optimal set of procedures needed for implementing similarity-based elastography.

Simulation 1: Evaluation of Similarity Metrics

The purpose of this set of experiments is to detect and compare differences between the similarity metrics. A reconstruction was run using each of the individual measures. A pre- and post-deformed image set was simulated by providing the elastic model with a modulus distribution defining a simulated tumor (i.e. stiff region) with a modulus value of 106 *kPa* (non-tumor tissue is assigned a value of 18 *kPa* in all simulations) and a 7*mm* radius. This can be seen in Figure II.14 and the image set in Figure II.1. All other parameters remained constant except for the α regularization parameter. It must be adjusted due to the similarity metrics operating on different ranges of values. For example, normalized mutual information ranges from 1 to 2 while SSD may be on the order of hundreds. This affects the scale of the Hessian matrix, and thus requires adjustments in the regularization.

Simulation 2: Variations in Methods for Evalutating Similarity

One of the key elements of the reconstruction algorithm is the evaluation of image similarity. Within this there are two important variables, the metric used, which was addressed in Experiment 1, and the procedure used when taking the measurements. The latter is the focus of this experiment. While the general method for measuring similarity remained constant (seen in Figure II.9), the parameters defining the location and size of the similarity regions were varied. The radius of the regions was perturbed, which directly affects the number of measurements taken, and for each radius considered two reconstructions were run, one including the mesh boundary and the other without the boundary. This idea of inclusion/exclusion of the boundary can be seen in Figure II.10 and is the second parameter considered. Effects of these variations will be evaluated for gradient correlation and mutual information. All other parameters remained constant. The pre- and post-compression images were created as described above and are represented in Figure II.1.

Simulation 3: Number of Tumors Reconstructed

These computational experiments were used to examine the ability of similarity-based elastography to detect multiple tumors. Again using the MRI breast slice as our base image, simulated images were constructed each using a different modulus distribution, where the distributions differ by the number of tumors simulated. The number of tumors ranged from 1 to 4 and each tumor was specified to be of modulus 106 kPa and a radius of 7mm. A modulus distribution was then generated using the simulated images. The similarity measure used was gradient correlation.

Simulation 4: Reconstruction of Near Field/Far Field Tumors

An important characteristic of the algorithm is its ability to detect tumors considered to be in the far field with respect to the externally supplied deformation. Two post-compression images were simulated, one using a modulus distribution with the tumor specified to be in the near field with respect to the deformation, and the second distribution with the tumor located in the far field. Figure II.47 represents the locations of the simulated tumors both with radius 7mm each with a modulus value of 106 kPa, and the MRI breast slice was the base line image. The simulated images were each used in a reconstruction with gradient correlation as the similarity metric.

Simulation 5: Variations in Tumor Radius

One of the defining qualities of any imaging modality is the resolution at which it can produce quality images. It is highly desirable for methods of elastography to be able to detect very small regions of stiffness. These experiments look to determine the sensitivity of this technique to tumor size. With the MRI breast slice as the pre-compression image, modulus distributions with tumor sizes of 1, 3, 5, and 10*mm* and modulus 106 *kPa* were used to generate simulated post-compression images. Figure II.53 shows the location of the tumors' centers. Reconstructions using each of the simulated image sets were considered.

Simulation 6: Determination of Stiffness Resolution

The major difference between this method and others is its use of image similarity to drive the reconstruction rather than measured displacements. As a result, it is dependent upon the stiffness of tissue changing the image signal between uncompressed and compressed states. The purpose of this experiment is to determine the level of contrast between tumor and healthy tissue that can be detected. To evaluate this a simulated image with a homogeneous modulus, specified as 18 *kPa*, distribution was created. Then the stiffness of a given tumor region was repeatedly increased from a ratio of 1:1 up to 55:1 and used to create a second simulated image. The similarity between this image and the initial homogeneous image was calculated. The tumor region is the same as that defined by the tumor in Figure II.14. Gradient correlation was the similarity measure.

Simulation 7: Effects of Inaccurate Poisson's Ratio

Plain strain elasticity is based on two parameters, Poisson's ratio and Young's modulus, however, our method is only reconstructing the latter of these parameters while holding the other constant at some assumed value. The purpose here is to analyze the effects of using a Poisson's ratio that has some associated degree of error from that of the true value. To do this a model generated image was constructed using the modulus distribution found in Figure II.14 along with the MRI breast slice. This initial image was created with a Poisson's ratio of 0.495 and tumor modulus of 106 *kPa*. Reconstructions were then carried out holding all parameters equal except for Poisson's ratio, which was differed by 0%, 1%, 3%, and 7% from the true value. Gradient correlation was used.

Simulation 8: Three-Dimensional Effects

A major assumption that is being made is that the deformations found within the tissue can be represented by a two-dimensional model. However, preliminary studies indicated that three-dimensional effects could play a pivotal role in the effectiveness of the similarity-based technique. This experiment uses a three-dimensional linear elastic model to generate a post-compression image. A tumor was simulated in the region of the dowel rod inclusion and given a modulus value of 1.8 *MPa*. The 3D mesh used to create this image can be seen in Figure II.69. The base line image can be seen in Figure II.15(a). Then two reconstructions were carried out with one using mutual information and the second using gradient correlation.

Simulation 9: Evaluation of a Parallel Implementation

In order to determine the effectiveness of our parallel implementation of the algorithm, a number of reconstructions using the same input parameters were run. The differences between runs were the number of processors used in the reconstructions. The number of processors ranged from 1 to 16.

CT Data 1: Effects of Image Blurring

The image set for the CT experiments are represented by Figures II.15(a) and II.15(b). The darker gray area, positioned center-left is the stiff inclusion, a wooden dowel rod. Also, reconstructions used a Poisson's ratio of 0.4635, which was determined through a thorough, but not exhaustive, parameter search. This can only be validated through material testing. This has not been done as of yet.

Often times in image registration image blurring is used to smooth the function space of the given similarity metric. The degree of blurring is defined by the size of the kernel used in the convolution (i.e. a image blurred with a kernel 3x3x3 will be blurred less than an image blurred with kernel 5x5x3). Using a low-pass filter, images from the CT image were blurred with kernel sizes 0x0x0, 3x3x1, 3x3x3, 5x5x1, 5x5x3, 7x7x1, 7x7x3. Figure II.16 is an example of a blurred image. It represents a 5x5x3 sized kernel being applied to II.15(a). The effects of these blurrings are then seen through the evaluation of their respective reconstructions. Gradient correlation was

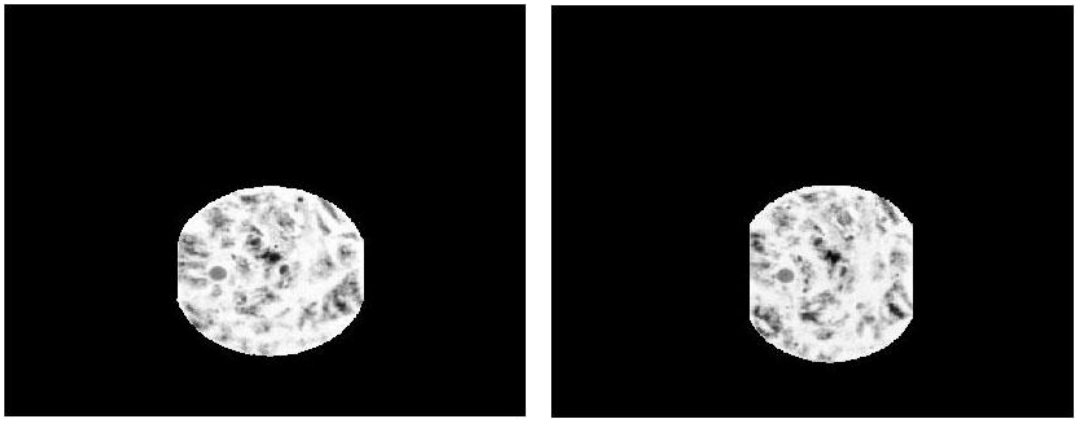

(a) (b)

Figure II.15: Acquired image series (a) CT slice of silicon phantom pre-compression, (b) CT slice of silicon phantom post-compression

again used.

CT Data 2: Reconstructing at Various Depths

As briefly discussed above, preliminary results have indicated that model inaccuracies due to a two-dimensional representation of a three-dimensional problem may play a part in reconstruction performance. Perhaps, one way to help reconcile this is to consider image slices located at lower depths of the phantom, because the lower depths have the additional constraint of the compression-device's bottom panel. To test the hypothesis a series of reconstructions using gradient correlation were carried out where the depth of the image slices were varied. The depths ranged from slice 9 of 30 to slice 19. Figure II.85 is a visual representation of the slice locations in order to provide a proper orientation perspective. The images were blurred with a 5x5x3 low-pass kernel.

CT Data 3: Variations in Parameter Space Resolution

One of the unique aspects of our method is the use of a dual finite element scheme. Described earlier in the methods, the parent mesh is the coarsely resolved mesh defining the property reconstructed domain. The purpose of these trials is to determine an upper and lower threshold for the resolution at which a reconstruction can be effective. Modulus distributions for slices from the CT

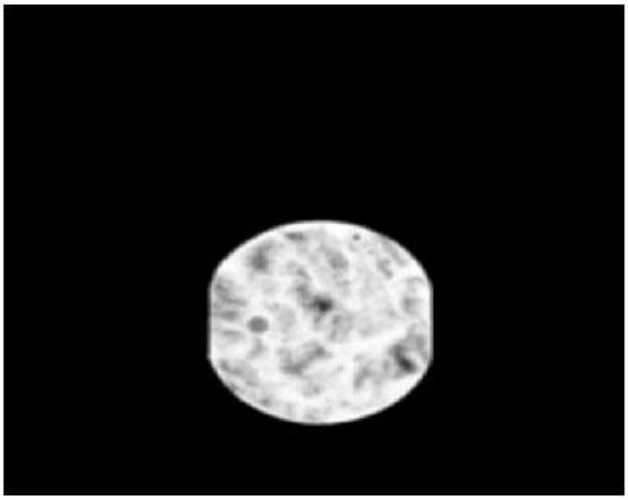

Figure II.16: CT image low-passed filtered with a 5x5x3 kernel.

image set were constructed using parent meshes with 379, 152, and 84 tissue properties. Again, image blurring was implemented and GC the metric.

MRI Data: MRI Parameter Search

This test is to help determine the realism of casting this technique as modality independent. It uses the image set acquired in the 3T MRI. A representative image set can be seen in Figures II.17(a) and II.17(b). The black area lower-center is the wooden dowel rod inclusion.

Optimal values for parameters such as Poisson's ratio and the α regularization term appear to have some dependence upon the image set being used in reconstructions. Therefore, the intent for this experiment is to determine these values for the MRI data set. An image set blurred with a 5x5x3 low-pass kernel was used. Reconstructions using gradient correlation were produced which varied both Poisson's ratio and α.

II.3 Results

II.3.1 Simulation Results

Simulation 1: Evalutation of Similarity Metrics

Figure II.18 shows the error profile for the different similarity metrics used (Equations (II.22) thru (II.27) and Gradient Correlation). Due to the differences in similarity ranges all error values have been normalized between 0 and 1 by subtracting the minium error and dividing by

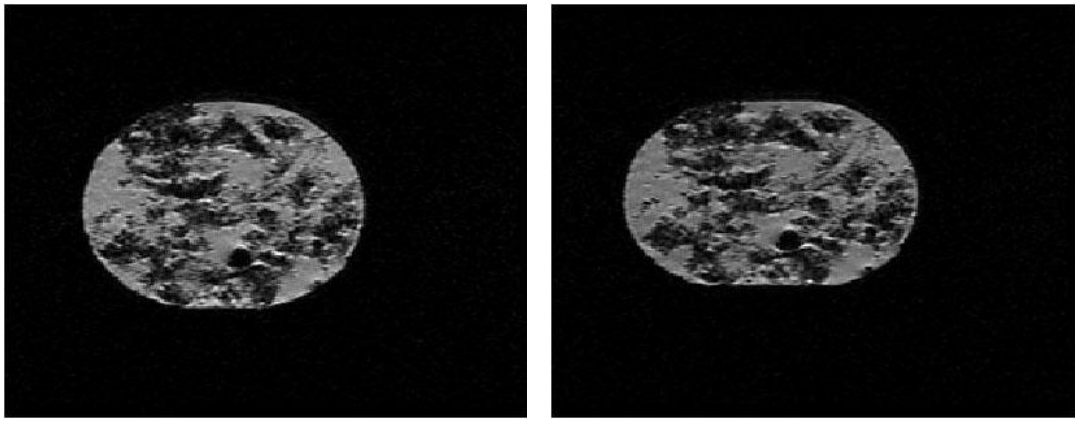

(a) (b)

Figure II.17: Acquired image series (a) MRI slice of silicon phantom pre-compression, (b) MRI slice of silicon phantom post-compression

$(error_{maximum} - error_{minimum})$. Therefore, error values of zero do not necessarily represent a perfect reconstruction. Also, the plots are meant to represent the behavior of a reconstruction with respect to stability, not the overall "goodness" or "badness" of a modulus reconstruction. As of now, to know this requires the analysis of the error plot in conjunction with the reconstructed paramter distribution.

With this in mind we see that all metrics behaved smoothly. Mutual information was somewhat sporadic, but overall the reconstruction still maintained a general decline in error. After considering the performance of a reconstruction's error, we look to calculated tissue parameters in order to better qualify it. Figure II.19 represents the reconstructed modulus values for those elements which were defined as being tumor in the simulated modulus distribution (see Figure II.14). Notice, that again the profiles for the modulus increase are smooth, except for that of MI. However, MI calculated values are approximately 39% of the original value while the other methods averaged only 19% of the initial 106 *kPa*. Figures II.20 thru II.23 demonstrated the spatial distribution of calculated properties. All did a good job of localizing the tumor. SSD performed the best spatially with a very tight localization while CC produced a very blurred distribution. Also, GC had good localization, but introduced small pockets of increased modulus. A concern though is that the

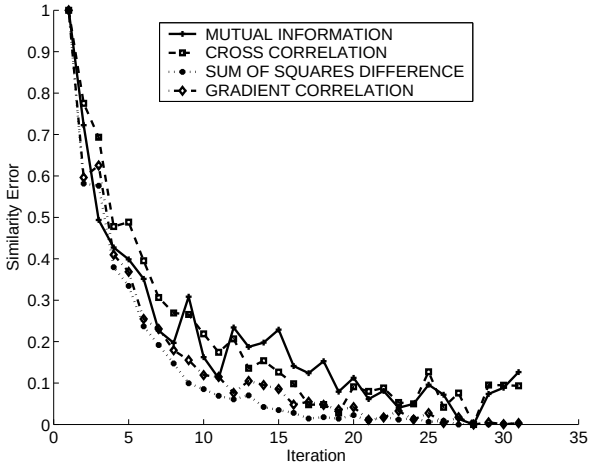

Figure II.18: Similarity error profile for various metrics.

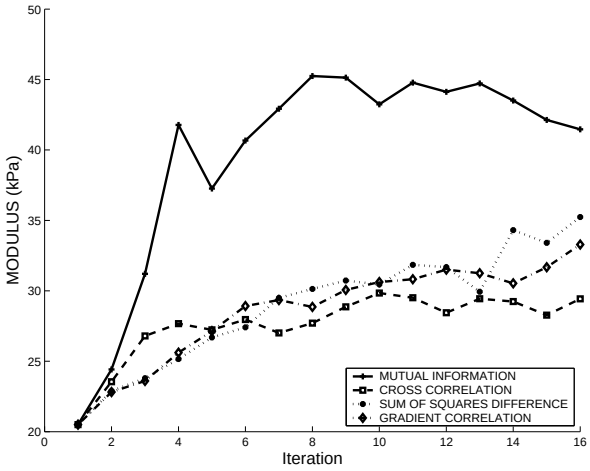

Figure II.19: Average modulus values of elements that define tumor for each of the similarity metrics.

reconstructed modulus values are 2-4 times lower than the actual value.

It is important to note that the stiff inclusion was only represented in the model. The tumor had the same intensity pattern as the surrounding tissue, and hence within the image there were no defining boundaries indicating where the stiff region was located. It is not surprising that SSD performed well in the simulation. Given the correct properties the post-compression image and the model image would be exact, and as mentioned earlier SSD is the optimal metric in this case. However, it is naive to believe this to be the case in actual data. What is surprising is the higher contrast produced by MI while having the poorest error profile. Lastly, it is encouraging that GC was able to perform so well. It primarily operates on edge maps which were not distinct between tumor and tissue.

Simulation 2: Variations in Methods for Evalutating Similarity

Figures II.24 and II.26 show the error profiles of GC and MI using similarity measurements excluding the image at the mesh boundary. Figures II.25 and II.27 give profiles with relatively equivalent measurement numbers but with the inclusion of the boundary. As can readily be seen GC peformed much better without the boundary when the number of measurements was below 598. While MI

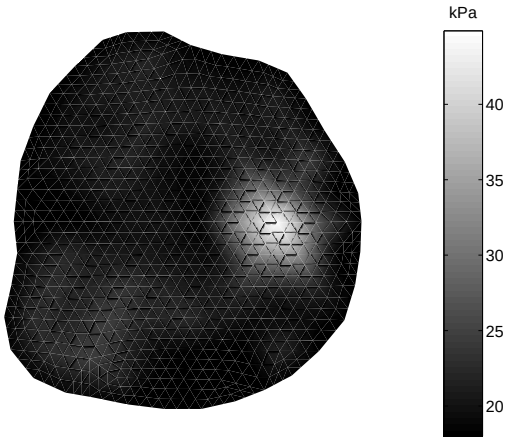

Figure II.20: Modulus profile reconstructed using Mutual Information.

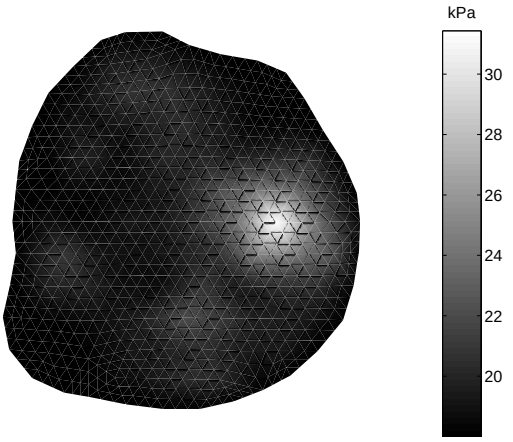

Figure II.21: Modulus profile reconstructed using Cross Correlation.

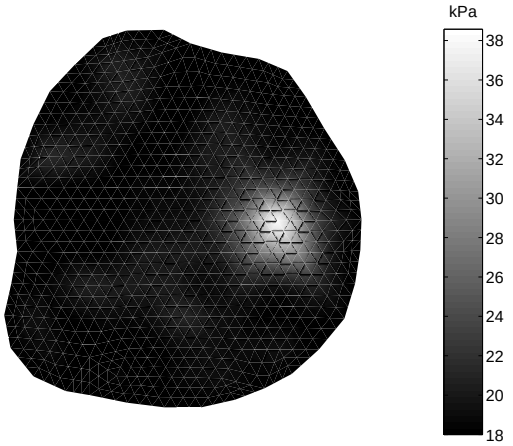

Figure II.22: Modulus profile reconstructed using Sum of Squares Difference.

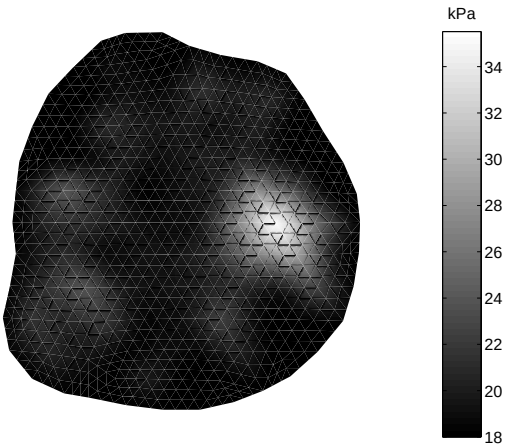

Figure II.23: Modulus profile reconstructed using Gradient Correlation.

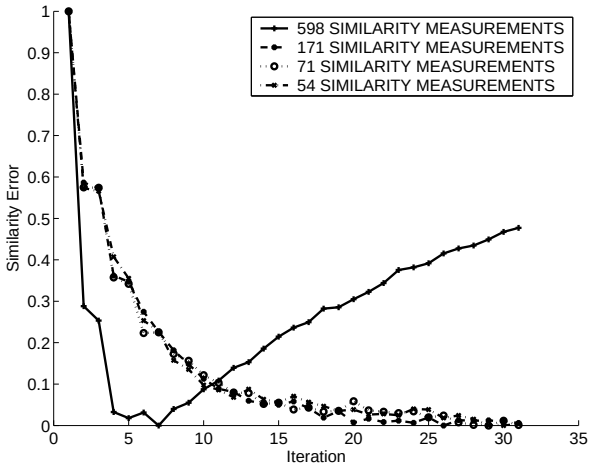

Figure II.24: Similarity error profile for GC without boundary measurements.

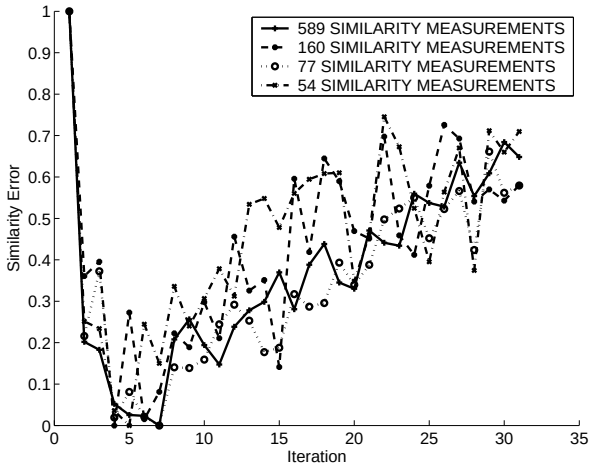

Figure II.25: Similarity error profile for GC with boundary measurements.

only reconstructed smoothly when using measurements numbering 171 or more, and the boundary was not considered. The reconstructed modulus distributions can be seen in Figures II.28 thru II.31, which were all taken at iteration 30. Both MI and GC calculated significantly more accurate and localized values without the boundary, even though it consisted of fewer measurements.

One of the most significant factors in the algorithms performance is a Jacobian matrix that accurately represents the change in similarity with respect to tissue properties. Each row of the Jacobian represents the change of a specific similarity region with respect to the modulus parameters. The mean of these values for each region is shown in Figures II.32 thru II.35. These are all the initial Jacobians created for a given reconstruction. Notice, that for both GC and MI, there is enough noise to produce values outside of the mesh region which will have an influence on the Jacobian. Also notice that for GC there is an order of magnitude difference between the non-boundary and boundary values. This directly effects the scale of the Hessian matrix.

Figure II.36 and Figure II.37 indicate that progression of regional similarity during a reconstruction when using and not using the boundary for GC. While the similarity generally improved for both methods. The technique of not using the boundary seemed to have the most significant regions localized about the tumor, and with the boundary, similarity was much more dispersed.

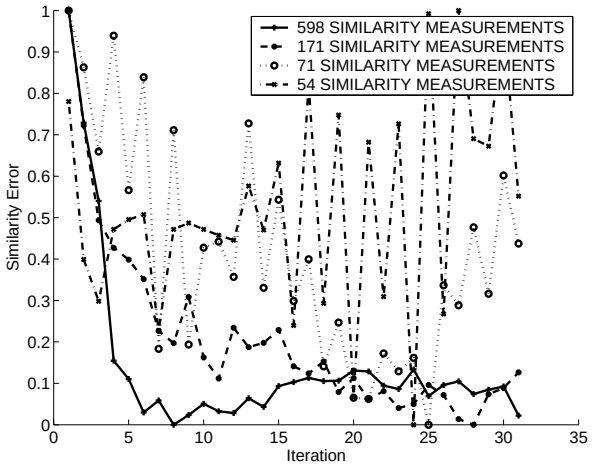

Figure II.26: Similarity error profile for MI without boundary measurements.

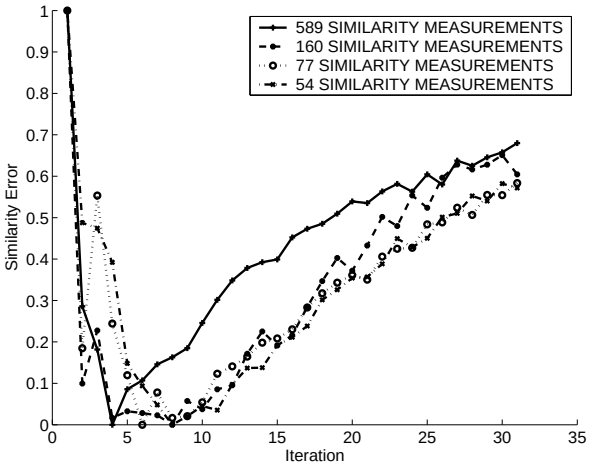

Figure II.27: Similarity error profile for MI with boundary measurements.

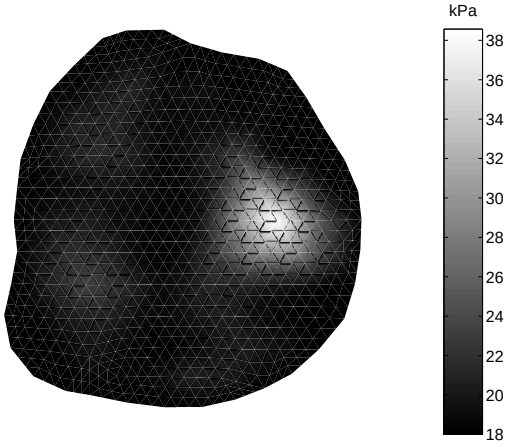

Figure II.28: Reconstructed modulus properties using GC without boundary measurements (71 measurements).

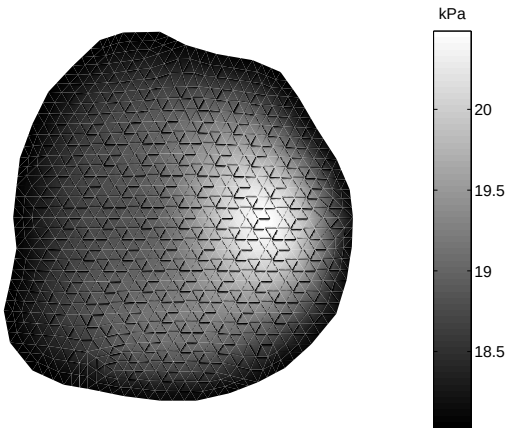

Figure II.29: Reconstructed modulus properties using GC with boundary measurements (77 measurements).

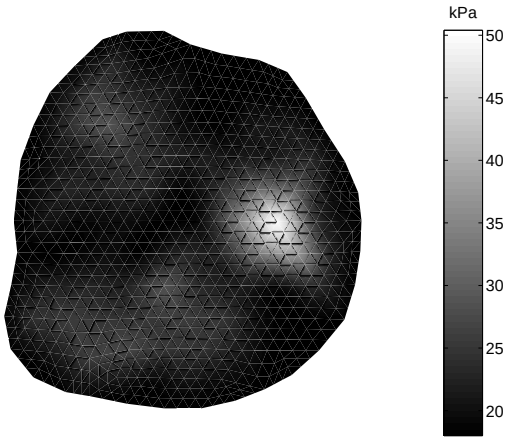

Figure II.30: Reconstructed modulus properties using MI without boundary measurements (71 measurements).

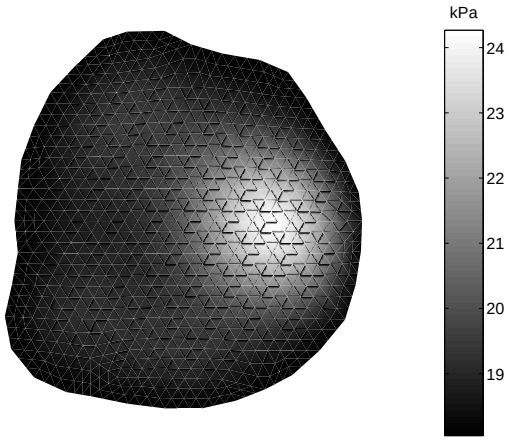

Figure II.31: Reconstructed modulus properties using MI with boundary measurements (77 measurements).

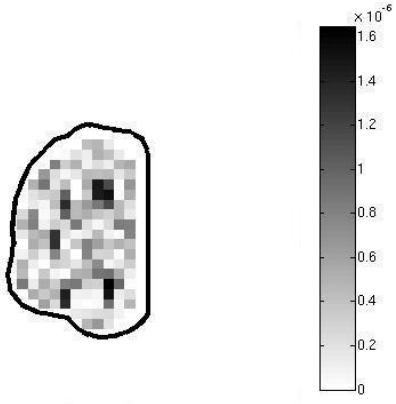

Figure II.32: The mean value of the Jacobian at the respective similarity measurement indices for GC without the boundary.

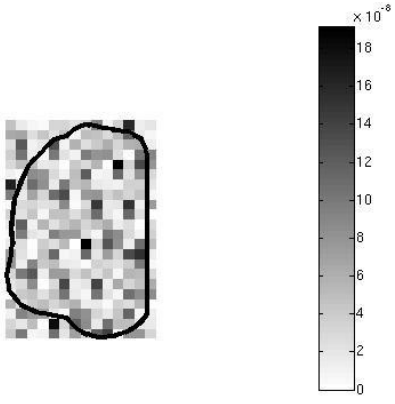

Figure II.33: The mean value of the Jacobian at the respective similarity measurement indices for GC with the boundary.

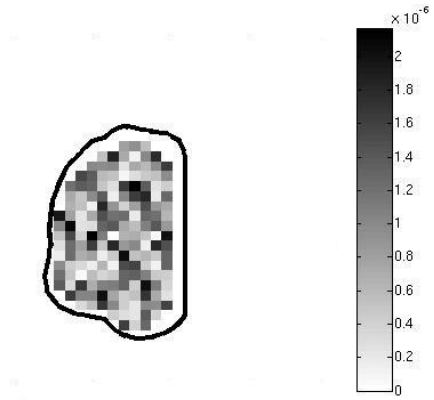

Figure II.34: The mean value of the Jacobian at the respective similarity measurement indices for MI without the boundary.

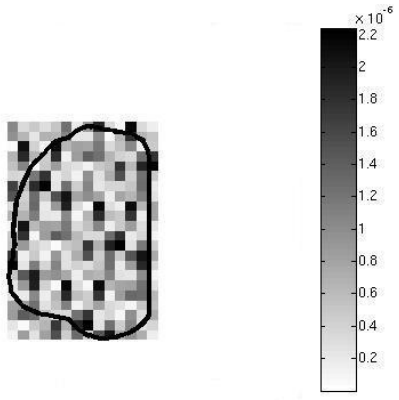

Figure II.35: The mean value of the Jacobian at the respective similarity measurement indices for MI with the boundary.

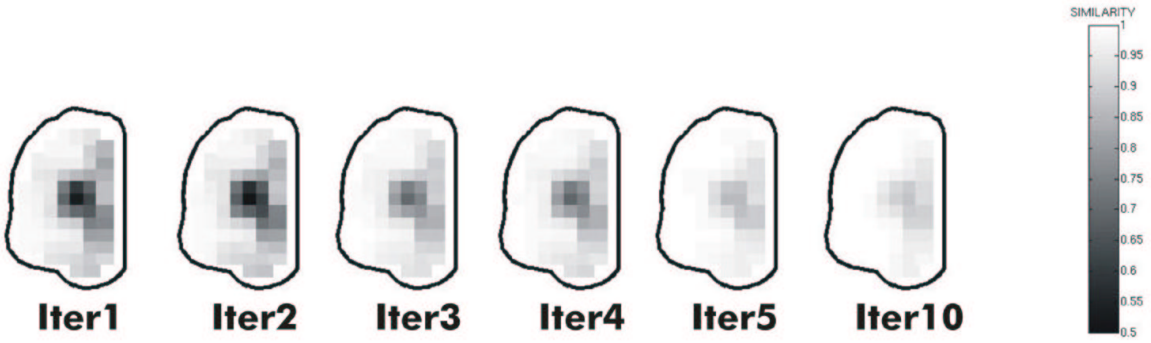

Figure II.36: Progression of regional similarity values for GC without boundary.

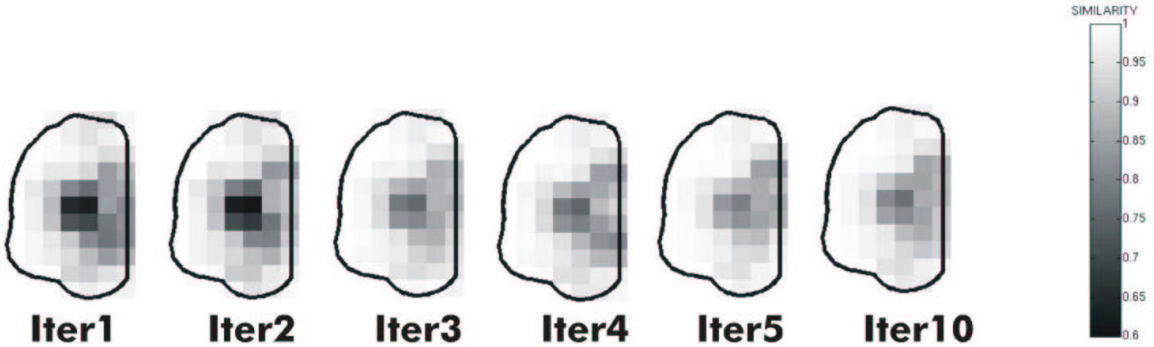

Figure II.37: Progression of regional similarity values for GC with boundary.

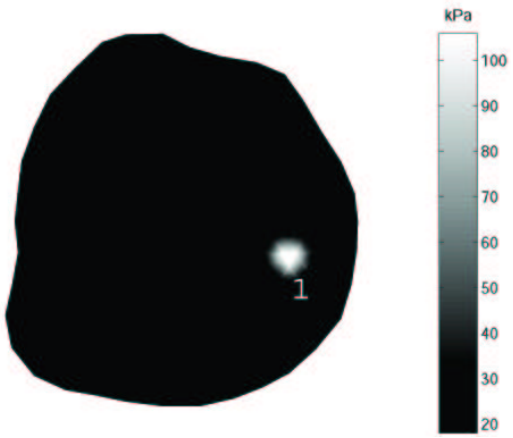

Figure II.38: Simulated modulus profile with 1 tumor (106 *kPa*).

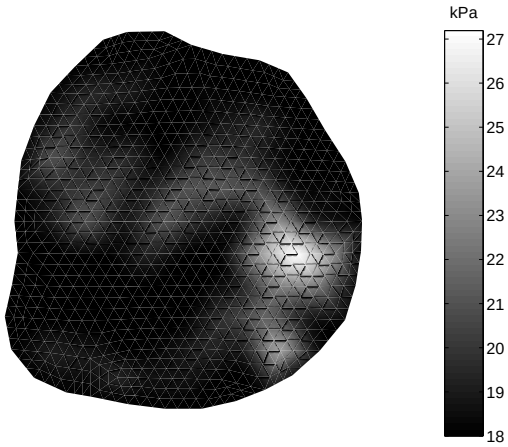

Figure II.39: Reconstructed modulus profile with 1 tumor.

Simulation 3: Number of Tumors Reconstructed

The ability of the algorithm to reconstruct properties in tissue with multiple tumors specified can be seen in Figures II.38 thru II.45, where the left image is the simulated modulus distribution and the right the reconstruction. Also the error performance can be seen in II.46. All behave smoothly. What is unusual is that the worst reconstruction is actually the simulation with only one tumor. Also interesting is the fact that tumor 3 is prominent in the reconstruction using three simulated tumors, but this is not true in the reconstruction with 4 tumors. We can also see in the fourth reconstruction that there is significant blurring between the tumor boundaries of tumors 1 and 4.

Simulation 4: Reconstruction of Near Field/Far Field Tumors

To better understand the role of tumor location in the parameter reconstruction, Figure II.47 was used in the generation of two simulated post-compression images. One with the tumor in the far field and the second using the near field tumor. Near and far taken in reference to the applied deformation which was simulated from the right pushing to the left. The error profile characterizing the optimization can be seen in Figure II.48. As expected the tumor in the near field produced the smoother plot. However, the far field tumor actually produced a better reconstruction. A possible explanation for this may be seen in Figures II.51 and II.52. These represent the difference in the

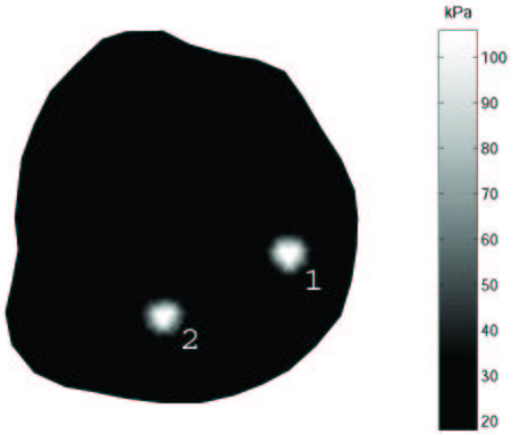

Figure II.40: Simulated modulus profile with 2 tumors (106 *kPa*).

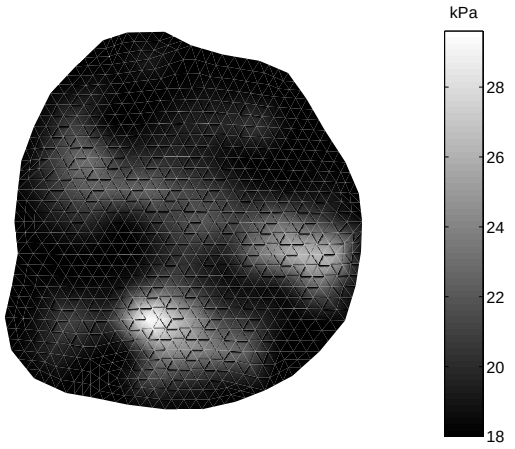

Figure II.41: Reconstructed modulus profile with 2 tumors.

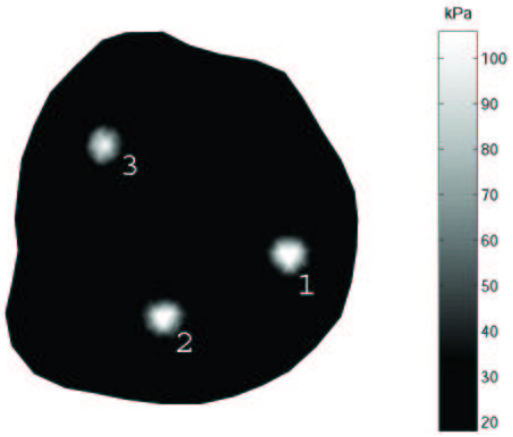

Figure II.42: Simulated modulus profile with 3 tumors (106 *kPa*).

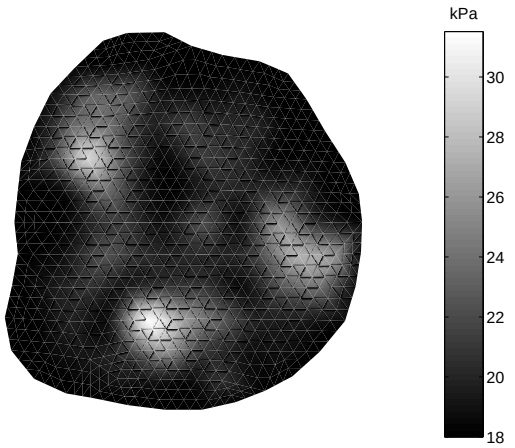

Figure II.43: Reconstructed modulus profile with 3 tumors.

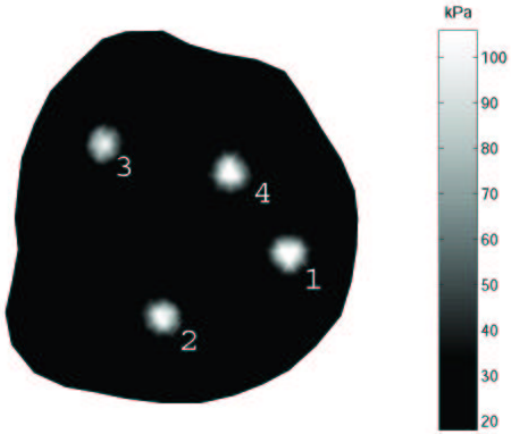

Figure II.44: Simulated modulus profile with 4 tumors (106 *kPa*).

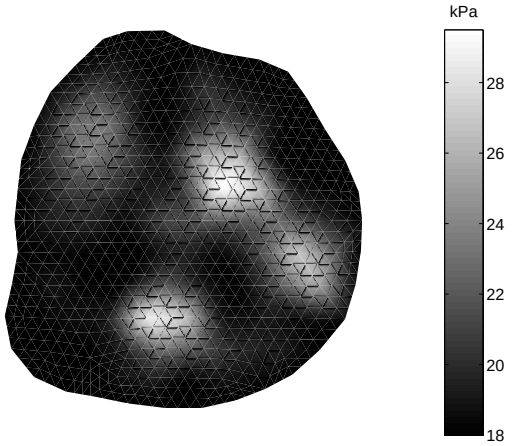

Figure II.45: Reconstructed modulus profile with 4 tumors.

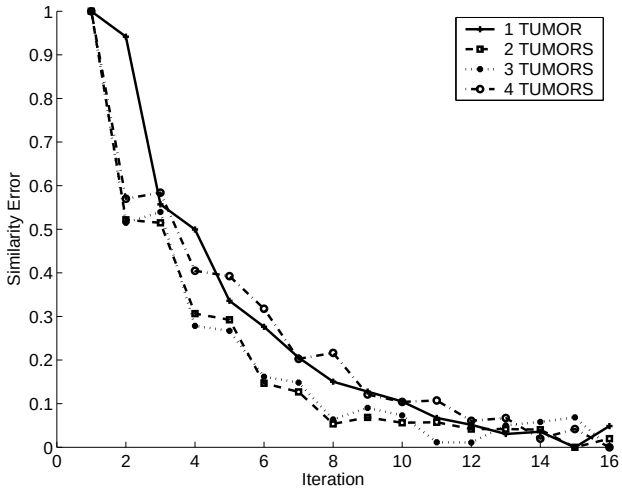

Figure II.46: Error profile of reconstructions with various number tumors.

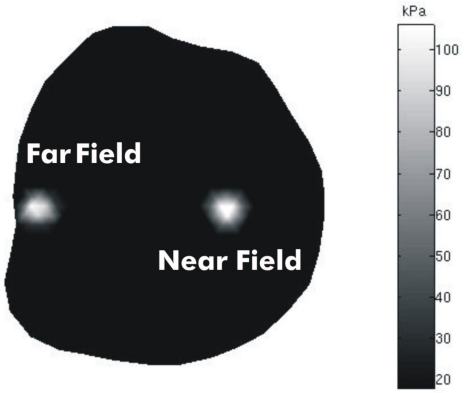

Figure II.47: Location of simulated tumors (106 *kPa*). One reconstruction using "Far Field" tumor and one using "Near Field" tumor.

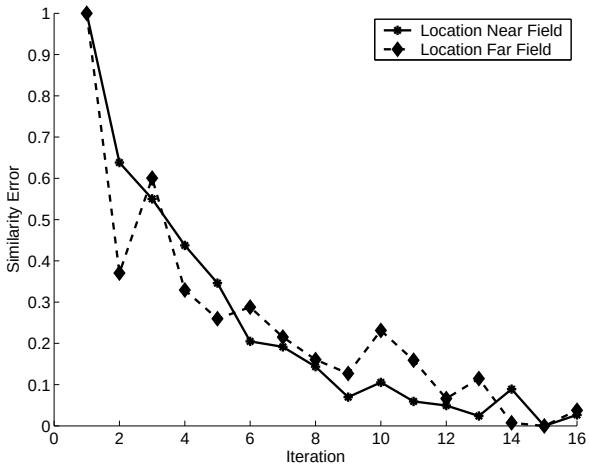

Figure II.48: Error profile for far and near field reconstructions.

deformation field produced by a homogeneous distribution and the simulated tumor distributions. Notice that the difference in the near field deformation is distributed much more throughout the tissue and is an order of magnitude greater than the far field deformations. This fact indicates that the near field tumor will produce a post-compressed image more different than that produced by the far field tumor. A result of this would perhaps be that the near field reconstruction starts farther from the local minimum than the far field, and thus introducing greater chances for error. Also, an extremely impressive result is that even with deformation differences being below the resolution of a pixel the far field tumor is capable of being detected.

Simulation 5: Variations in Tumor Radius

The error plot of reconstructions with different sized tumors can be seen in II.54. All tumors were centrally located at the location seen in Figure II.53. Their associated reconstructed parameter profiles are given in Figures II.55 thru II.58. Earlier it was stated that the error profile does not always indicate a successful reconstruction. This is the case for the simulation using a 1*mm* sized tumor. While its error is somewhat erratic it still performs smoothly enough to be considered an error decreasing reconstruction. Yet, its modulus profile is completely wrong (Figure II.55. It appears that 3*mm* is the smallest size that can be indicated, and at least 5*mm* to achieve any significant contrast.

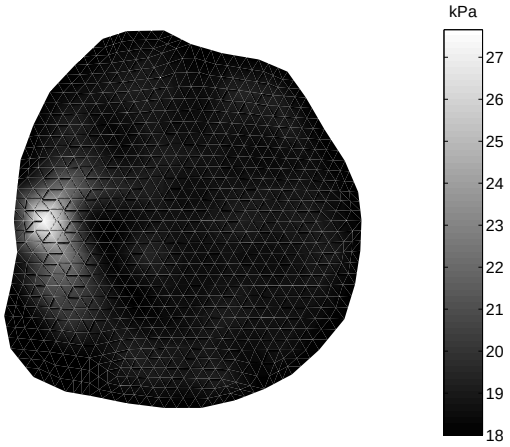

Figure II.49: Reconstructed modulus with simulated far field tumor.

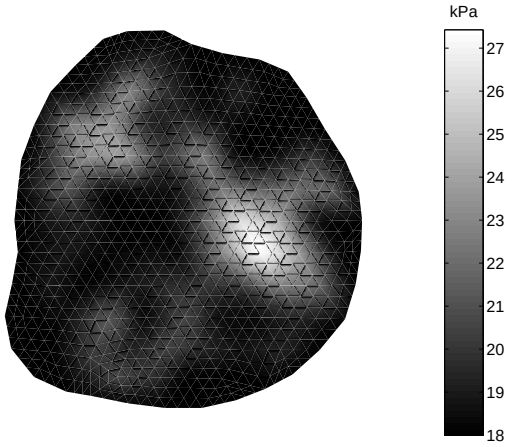

Figure II.50: Reconstructed modulus with simulated near field tumor.

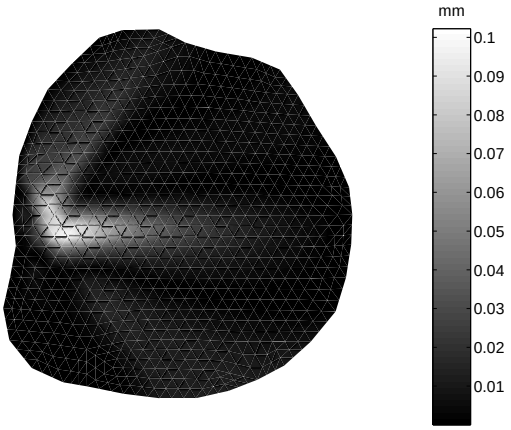

Figure II.51: Difference in deformation field due to far field tumor.

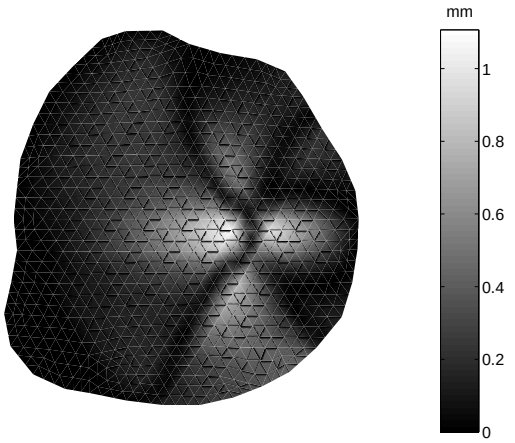

Figure II.52: Difference in deformatin field due to near field tumor.

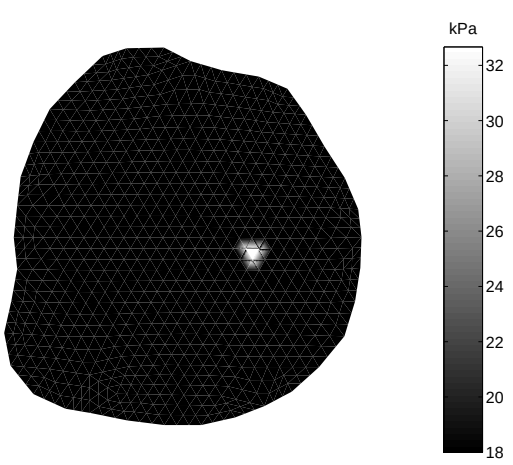

Figure II.53: Center point for simulated tumors. All specified to have modulus value of 106 *kPa*.

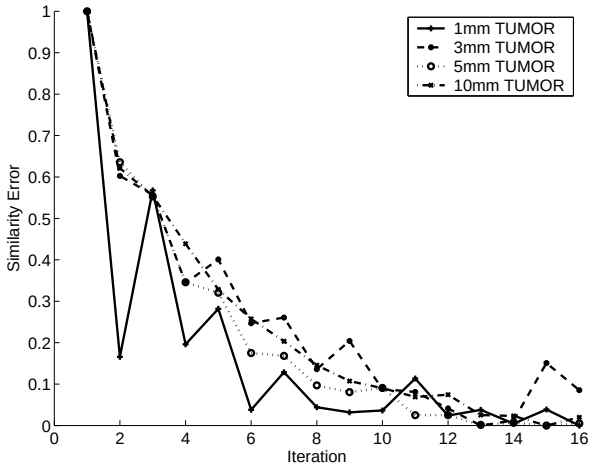

Figure II.54: Error profile for various tumor sizes.

Not astonishing is the fact that the 10*mm* tumor produced the best reconstruction. Its size is on the order of a parent element (1*cm* side lengths).

Another interesting note is the similarity between the deformation field in II.59 and II.51. While both are on the same order of deformation, evidence of a tumor seems to be more apparent in II.59 (maximum deforamation of 0.18*mm* compared to 0.10*mm*). However, it results in a set of totally inaccurate tissue properties, whereas the far-field reconstruction (Figure II.49) was able to achieve good localization of the simulated tumor.

Simulation 6: Determination of Stiffness Resolution

The change in image similarity with respect to tumor stiffness is given on a large scale in Figure II.61 and on a much smaller range in Figure II.62 (similarity here given on a log scale). Notice how II.61 decreases exponentially while II.62 appears to be almost linear. Also, important is the asymptotic nature of II.61, i.e. there is negligible difference between the 25:1 similarity and 55:1. The behavior of II.62 is also interesting in how it appears to change in very discrete steps of similarity.

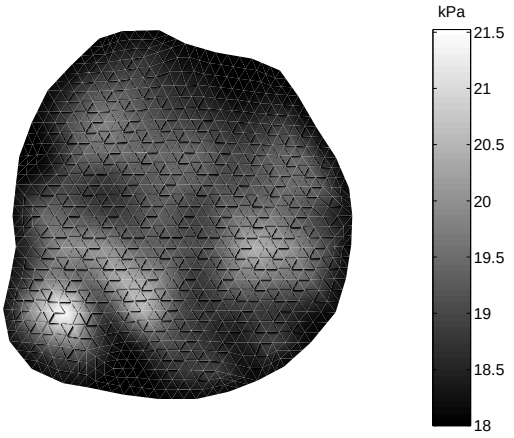

Figure II.55: Reconstruction with 1*mm* simulated tumor.

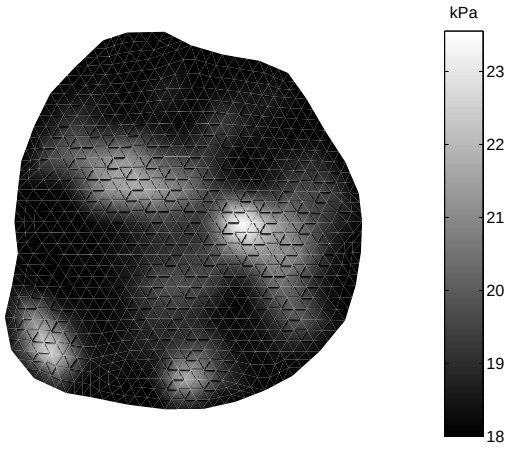

Figure II.56: Reconstruction with 3*mm* simulated tumor.

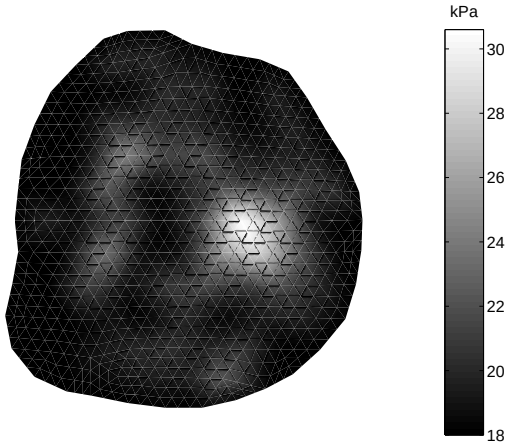

Figure II.57: Reconstruction with 5*mm* simulated tumor.

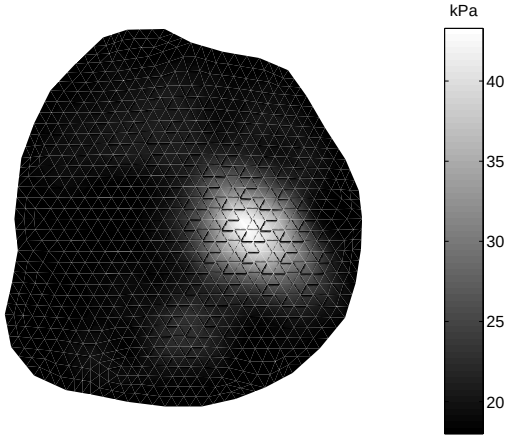

Figure II.58: Reconstruction with 10*mm* simulated tumor.

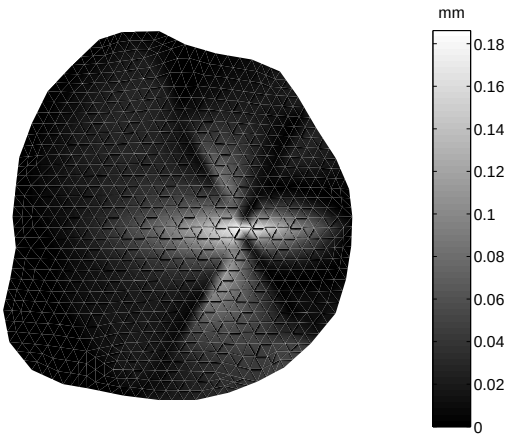

Figure II.59: Difference in deformation with a 1*mm* simulated tumor.

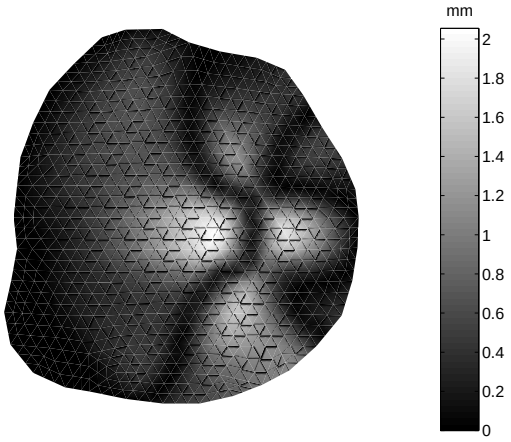

Figure II.60: Difference in deformation with a 10*mm* simulated tumor.

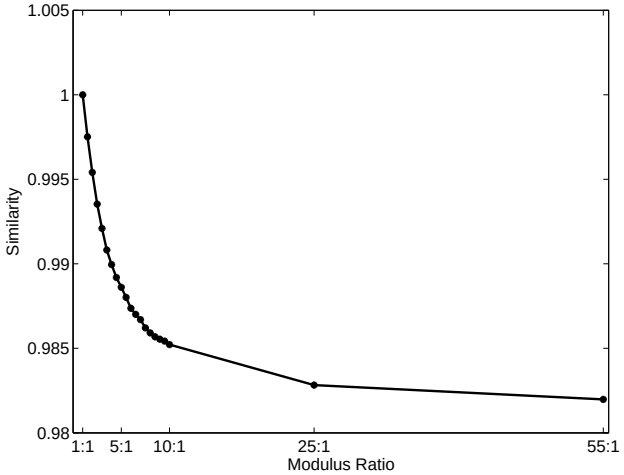

Figure II.61: Mean similarity for tumors of different stiffnesses.

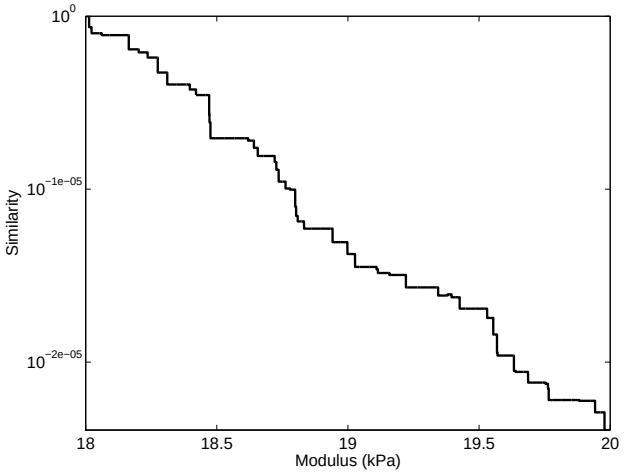

Figure II.62: Closer view of the effect of stiffness change on similarity.

Simulation 7: Effects of Inaccurate Poisson's Ratio

While Poisson's ratio is not being reconstructed, it plays an important part in the performance of the reconstruction which is evident in the error behavior given in II.63 and the modulus reconstructions (II.64-II.67). Figure II.64 represents the ideal case where the true Poisson's ratio and that assumed by the model are equal. The sporadic behavior seen toward the end iterations of 3% and 7%difference of II.63 gives clues to their poor reconstructions. Then inspection of II.66 and II.67 are definitive illustrations of this fact. Interestingly though, is that while both the 3% and 7% reconstructions (Figures II.66 and II.67) introduced artificial areas of increased modulus, they were still able to reconstruct values comparable to the 0% reconstruction within the region of the tumor. This is seen in Figure II.68, which represents the mean of the reconstructed modulus values of those elements specified as tumor in the simulated distribution.

Simulation 8: Three-Dimensional Effects

We have postulated that the model does not accurately represent the deformations being produced in real-world data. To test this hypothesis a 3D elastic model was used to create a set of simulated data. Figure II.71 is the difference between the 3D-model generated image and an actual post-compression data image. Compare this to Figure II.70, which is the difference between a 2D-model generated image and the data image. The significant fact seen here is how the 3D-model more

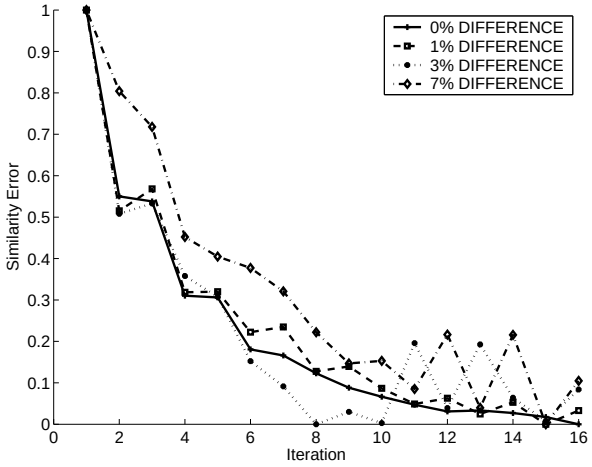

Figure II.63: Error profile for inaccurate Poisson's ratios.

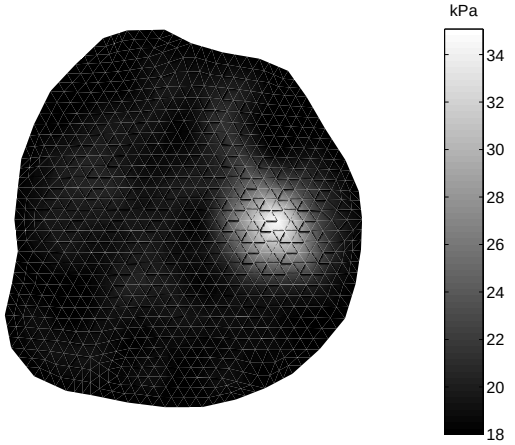

Figure II.64: Reconstruction with 0% difference in Poisson's ratios.

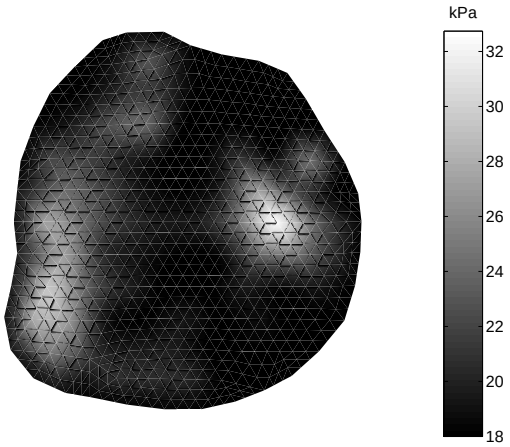

Figure II.65: Reconstruction with 1% difference in Poisson's ratios.

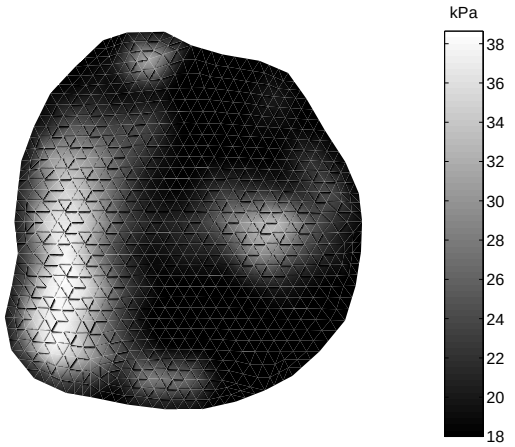

Figure II.66: Reconstruction with 3% difference in Poisson's ratios.

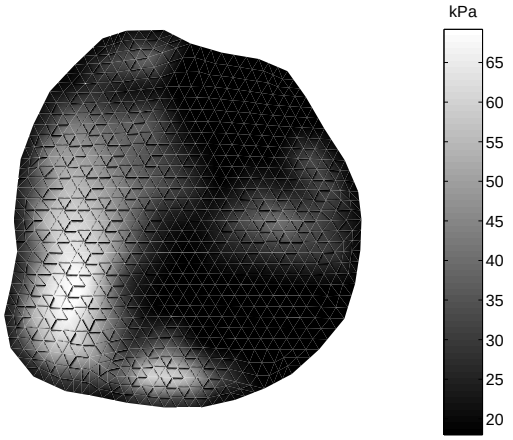

Figure II.67: Reconstruction with 7% difference in Poisson's ratios.

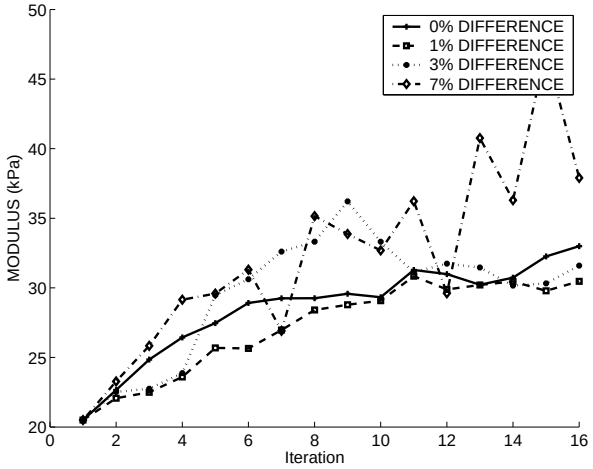

Figure II.68: Mean modulus values of elements defining tumor produced using inaccurate Poisson's ratios.

closely matches the boundary of the data image. Also, the reconstructions using the 3D simulated data, represented by II.73 and II.74, indicate that a three-dimensional model may more closely characterize the true deformations due to the fact that these are very similar to reconstructions seen in actual data (Figure II.75).

Simulation 9: Evaluation of a Parallel Implementation

Significant computational time was saved through the implementation of parallel techniques. This is evident in the Figure II.76 and Figure II.77. We can see that there is a limit to the improvement in reconstruction time. However, we do closely follow the theoretical limit. While the specifics to why we diverge from this limit are somewhat unclear at this time, a way to view this in a global sense is seen in II.77 which represents the factor of improvement for a given number of processors. Theoretically this is linear, i.e. 16 processors improve computation time by a factor of 16. Following Amdahl's Law [72], the speedup of a process can be represented by

$$S(p) = \frac{T_0}{(1-r)T_0 + \frac{rT_0}{p}}, \tag{II.28}$$

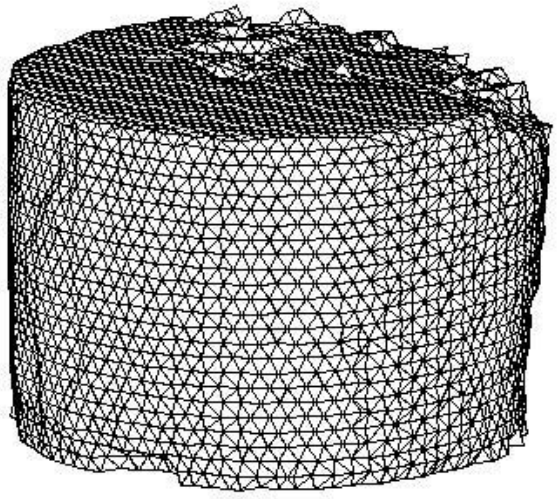

Figure II.69: Three-dimensional mesh used in generating simulated data.

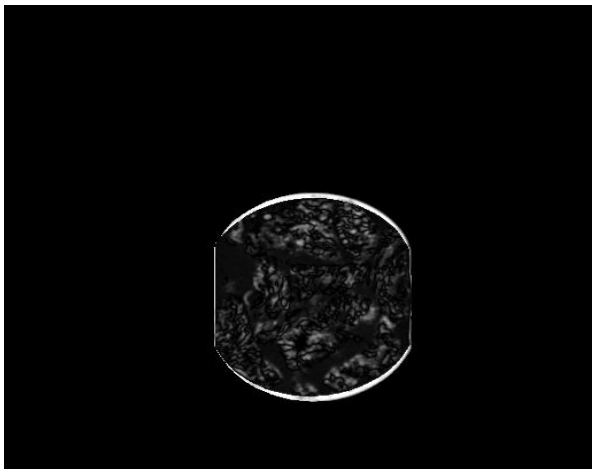

Figure II.70: Difference image between actual data and 2D-model generated image.

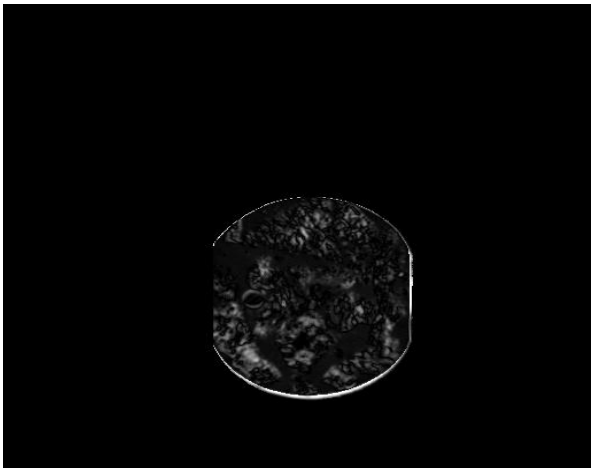

Figure II.71: Difference image between actual data and 3D-model generated image.

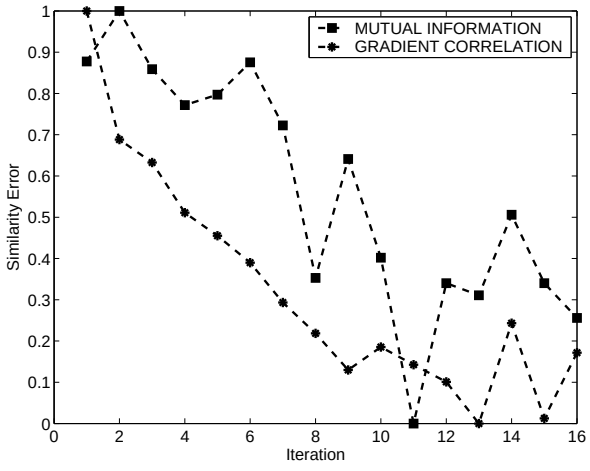

Figure II.72: Error profile of GC and MI using 3D simulated data.

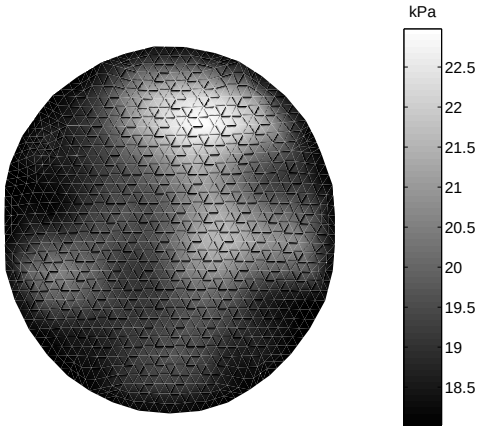

Figure II.73: Modulus profile generated using MI on 3D simulated data.

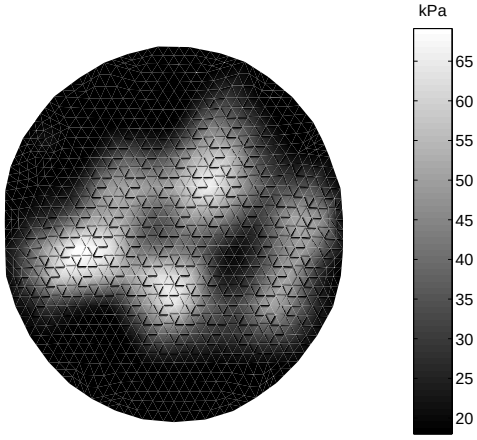

Figure II.74: Modulus profile generated using GC on 3D simulated data.

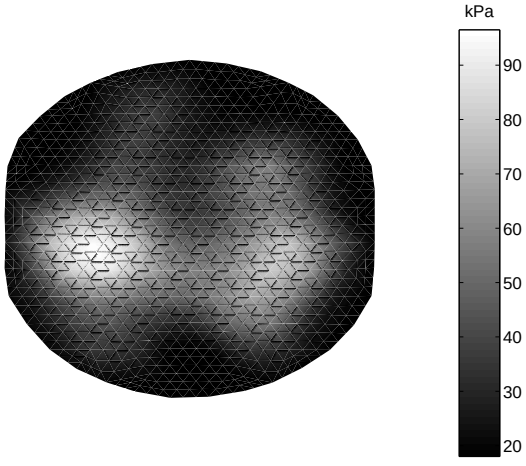

Figure II.75: Reconstructed modulus parameters generated using GC and CT data.

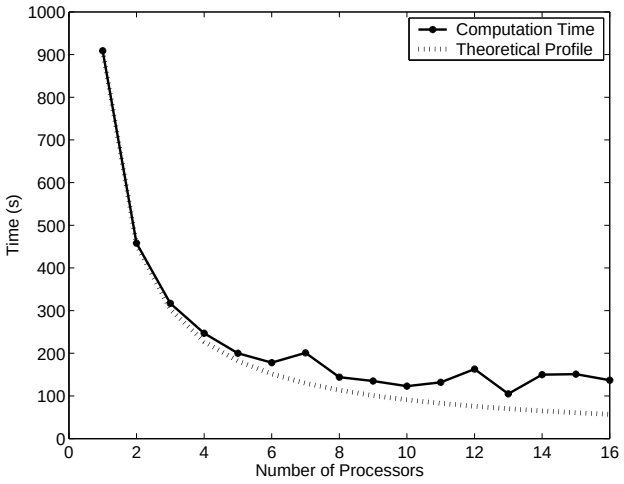

Figure II.76: Time of reconstruction vs. number of processors used.

where p is the number of processors, T_0 is the time taken if executed on one processor, and r is a measure of parallelization. A program is said to be "perfectly parallel" when $r = 1$. Using a least squares fit, a curve was fitted to the speedup data and an estimated value for r is 0.9263 for our current code.

II.3.2 CT Reconstructions

CT Data 1: Effects of Image Blurring

Figure II.78 compares the effects of varying the kernel size when blurring. Here we see that blurring does improve the smoothness of the error curve in later iterations, and also that the larger kernel size, 7x7x3, appears to perform worse than the smaller 3x3x3. However, while the overall size of the kernel does show to be of importance, what seems to be most significant is the blurring across image slices. Evidence of this can be seen in Figure II.79. The modulus profiles seen in II.81 and II.82 also demonstrate this fact. Another phenomenon that has been seen before is also exhibited in II.83. An optimization whose error profile is not the most well behaved, produces a more accurate set of modulus parameters.

It is important to note here that the true modulus values of the phantom are not known. The intensities of these values are given in order to provide a sense of the contrast being detected. Also, many of the reconstructions indicate increased modulus in areas other than the wooden dowel rod.

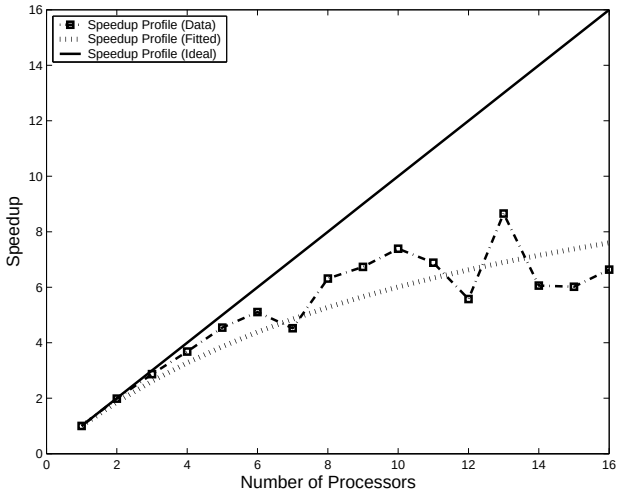

Figure II.77: Speedup factor obtained vs. the number of processors used.

With close inspection these regions appear to correspond to the sponge material that was used to produce image heterogeneity. Without proper material testing it is not known how these regions actually affect the modulus distribution.

CT Data 2: Reconstructing at Various Depths

Reconstructing image slices which correspond to different depths of the phantom is an effort to help reduce model inaccuracies. The idea being that lower depths will have fewer 3D-effects. Figure II.86 does not support this, and comparing Figure II.87 to II.88 actually negates the assumption. While the reconstruction using slice 15 has the most irregular error plot, again it produces the best reconstruction.

CT Data 3: Variations in Parameter Space Resolution

Figure II.90 is the error curves for reconstructions using different parameter space resolutions. It demonstrates that the rate of convergence increases as the number of reconstruction parameters decreases. This is also evident in the stiffness plots given in Figures II.91, II.92, II.93, and II.94. The highest resolution appears to be reconstructing in the correct location, however it has not obtained the contrast seen in II.92. Along with this, II.93 has similar contrast and spatial distribution to that of II.92 at iteration 23 of 30. Yet by iteration 30 it has begun to diverge from reconstructing in the

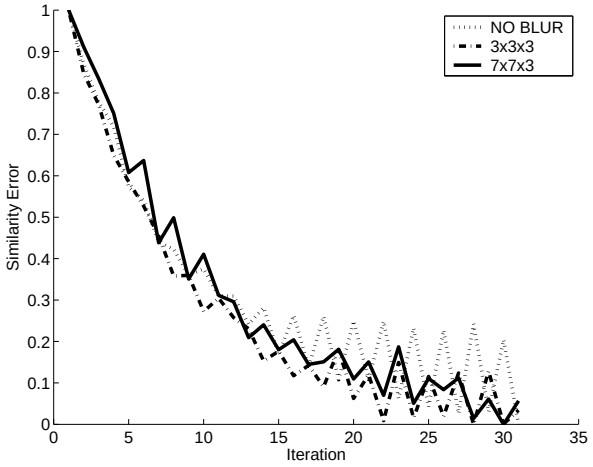

Figure II.78: Error profile of reconstructions using various blurring kernel sizes.

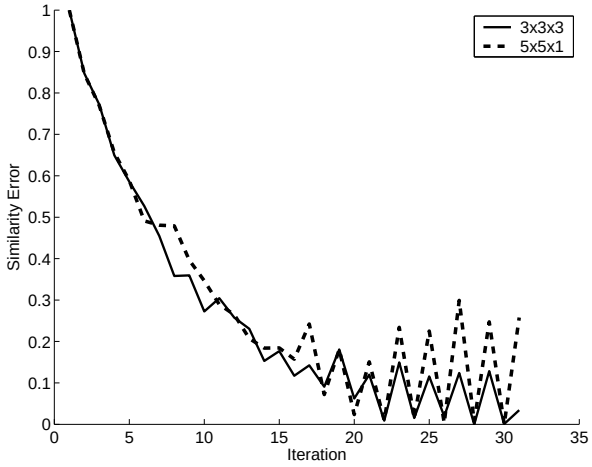

Figure II.79: Error profile comparing blurring within slice plane to blurring across slices.

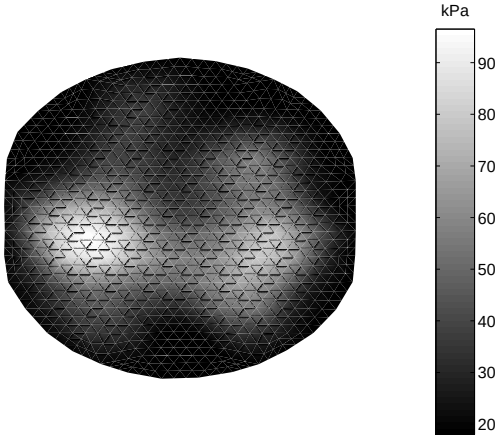

Figure II.80: Reconstructed modulus parameters generated using no image blurring.

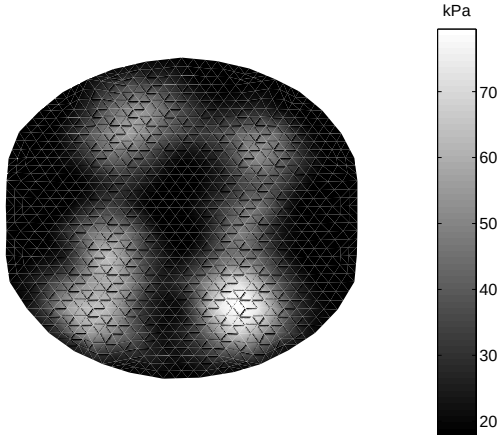

Figure II.81: Reconstructed modulus parameters generated using a 5x5x1 image blurring kernel.

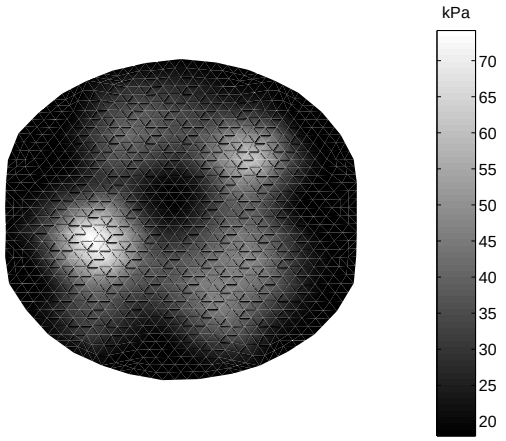

Figure II.82: Reconstructed modulus parameters generated using a 3x3x3 image blurring kernel.

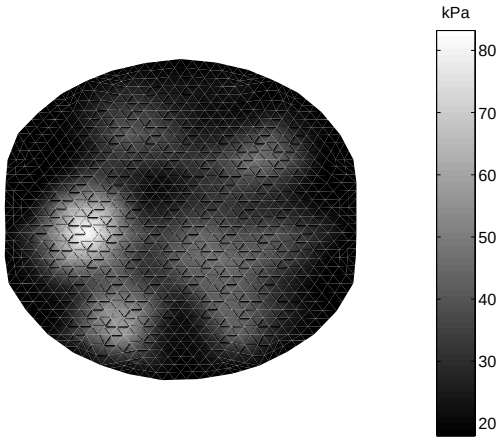

Figure II.83: Reconstructed modulus parameters generated using a 7x7x3 image blurring kernel.

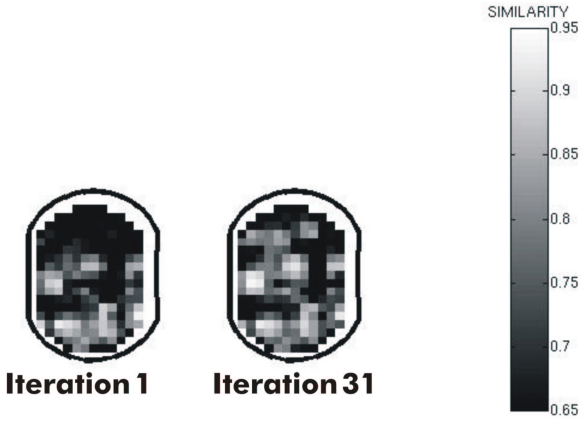

Figure II.84: Image similarity at first and last iteration using a 5x5x3 blurring kernel.

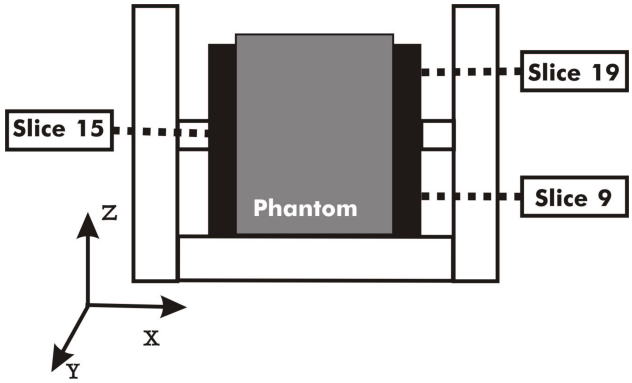

Figure II.85: Figure showing orientation for slice depth experiment.

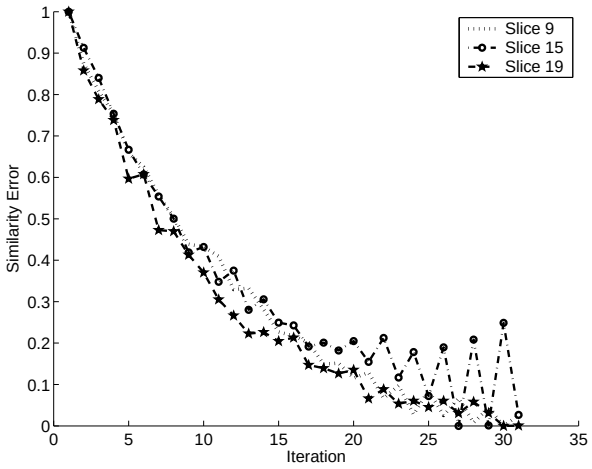

Figure II.86: Error profile using slices at different depths within the phantom.

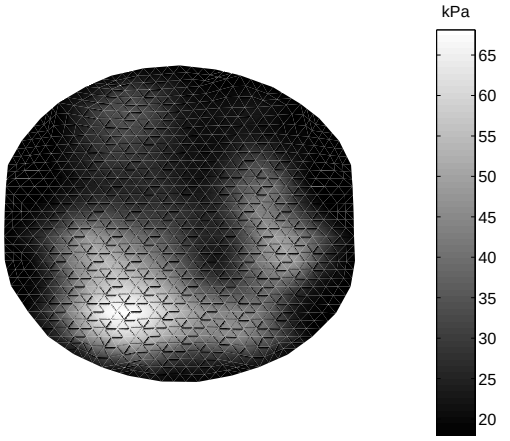

Figure II.87: Reconstruction using low-depth slice 9.

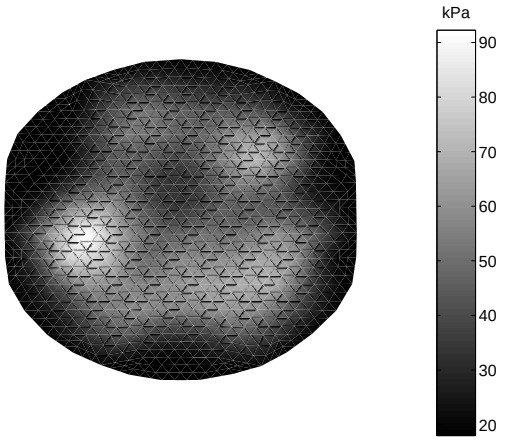

Figure II.88: Reconstruction using mid-depth slice 15.

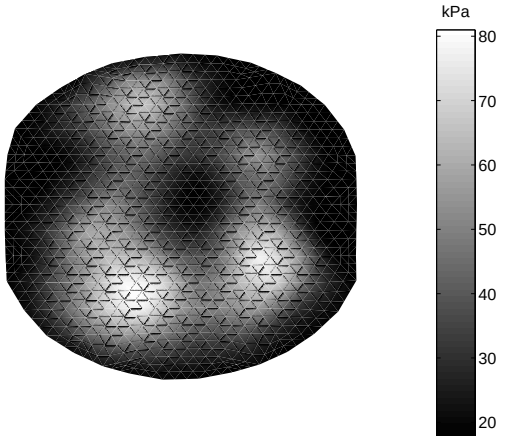

Figure II.89: Reconstruction using high-depth slice 19.

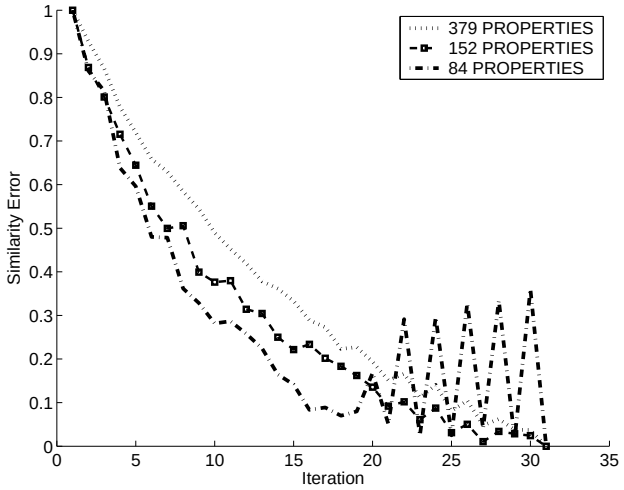

Figure II.90: Similarity error profile from reconstructing in parameter spaces of different resolutions.

appropriate locations. This coincides with the large jumps in error seen in II.90.

II.3.3 MRI Reconstructions

MRI Data: MRI Parameter Search

Perhaps the most exciting of all the results of this study can be seen in Figure II.96. Seen here is evidence strongly supporting the claim that this reconstruction scheme is independent from any particular modality. The modulus profile shown here is the parameter set at the eighth iteration, which corresponds to the lowest error value given in the plot II.95. After this the reconstruction begins to diverge, and Figure II.97 shows this. It is important to note, however, that while this last iteration is obviously not ideal, it did not completely deviate from the correct location. Comparing these results to the ones seen in the CT experiments the localization of the tumor is equivalent. However, the contrast in modulus values is not. It is not clear why this is true.

Another important point is that the parameter search found the optimal Poisson's ratio to be 0.464 compared to 0.4635 in CT, and the range of α regularization parameters that produced valid reconstructions was equivalent to that in CT. These are both significant in showing that these parameters appear to be phantom specific not modality.

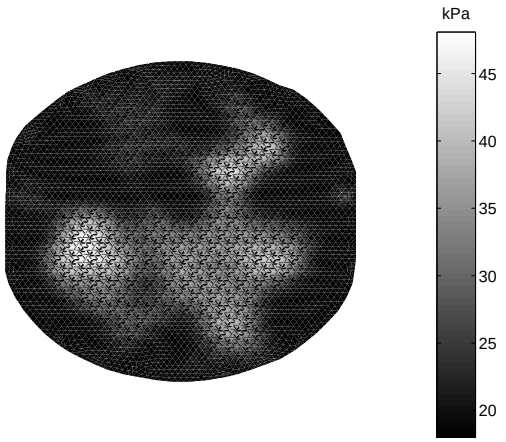

Figure II.91: Reconstruction using 379 parameters.

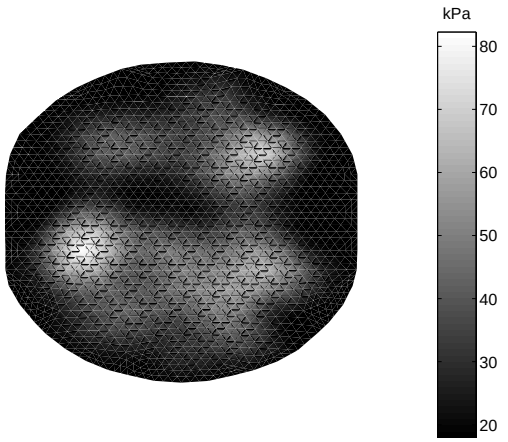

Figure II.92: Reconstruction using 152 parameters.

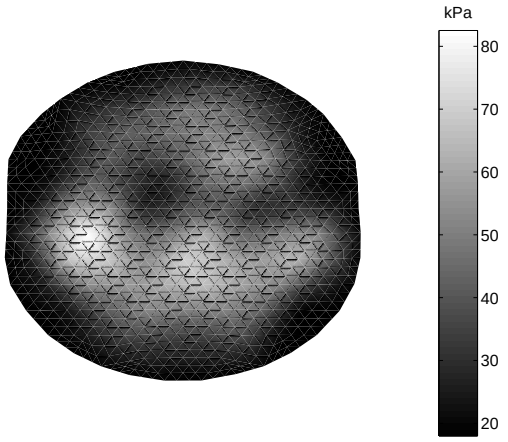

Figure II.93: Modulus profile using 84 parameters at iteration 23 of 30.

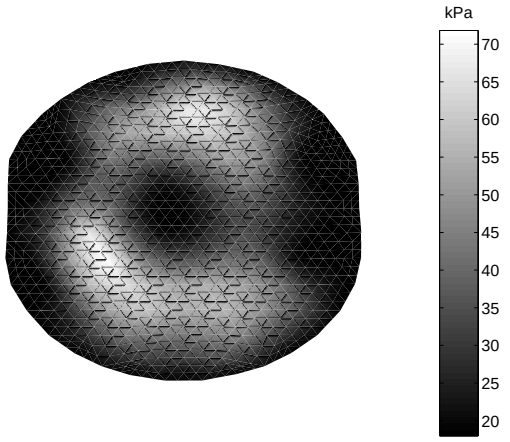

Figure II.94: Modulus profile using 84 parameters at iteration 30 of 30.

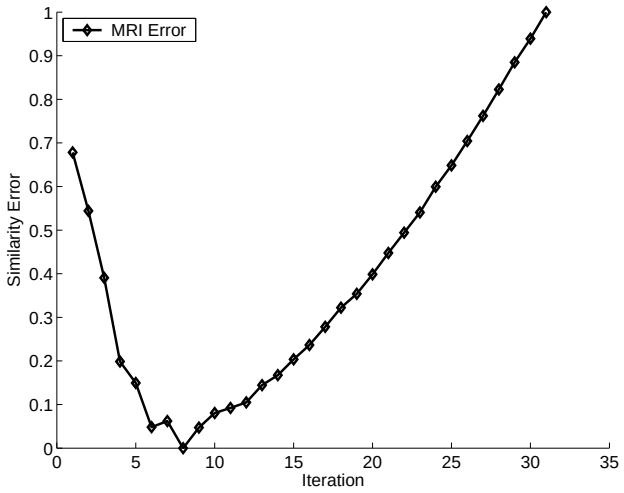

Figure II.95: Similarity error profile from MRI phantom reconstruction.

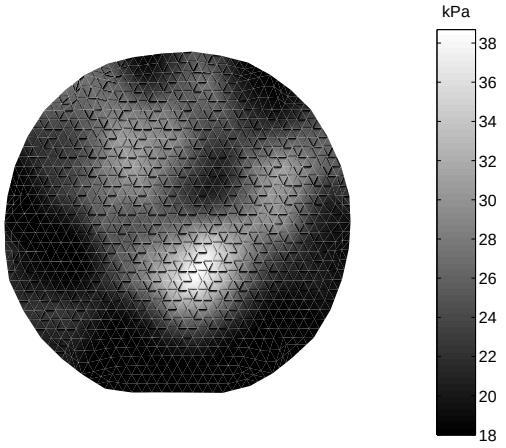

Figure II.96: Modulus values from MRI reconstruction at iteration 8.

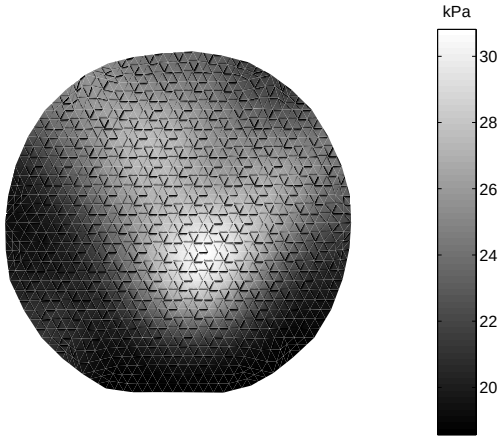

Figure II.97: Modulus values from MRI reconstruction at iteration 30.

II.4 Discussion

In this work we have presented an approach to elastography which is driven by the regional calculation of image similarity. Both the elegance and novelty of this technique lie in the fact that tissue properties are being reconstructed based on the information contained within a pre- and post-compressed image set. Previous techniques, which may be considered similar to this, have considered relative movement within images, but were confined to MRI due to the necessity of special imaging sequences. This method, however, is in essence modality independent. Its only requirement from the acquired data is an appropriate level of image heterogeneity, which is needed in order to detect changes in image pattern. By casting elastography in such a way we have significantly simplified many of the most difficult problems, such as data acquisition, associated with the more traditional methods.

The fundamental assumption we have made in order to make this shift away from displacement-based reconstructions is that the information needed to determine mechanical properties is contained within a set of pre- and post-deformed images. It is clear from the results seen in Figures II.18 thru II.23 that this is a reasonable assumption. What makes this point even clearer is the fact that the simulated tumor had no defining characteristics within the images. The similarity metrics were all operating on subtle changes in the image intensity patterns between pre and post. They were not simply honing in on areas that were clearly distinguishable within the images. Also impressive is that this change in pattern was detectable by all four similarity metrics used. It is not apparently clear at this time why MI was able to achieve a higher contrast ratio than the other metrics. One explanation maybe that MI is the most appropriate metric for detecting small changes between images, while a metric such as GC, which operates on edges, may be able to capture more gross movements in anatomical structures. Hence, each method may be best suited for a specific modality or particular image pattern.

While MI did reconstruct more accurate properties than those generated by other metrics, it was still over a 50% difference between its reconstruction and the original value of $106kPa$. This is a concern that was seen across all simulations with the best reconstruction being seen in Figure

II.30. It is believed that this is an indication of the methods dependence upon the resolution of the parameter space, i.e. the parent mesh resolution. The simulations were generated with tumors not being specified directly around a reconstruction property. That is, they were defined within and across parent elements. Further studies not presented here have shown that the contrast significantly improves in simulation if the tumor is defined around a parent node. This suggests that it may be necessary to implement a multi-resolution scheme or to optimize on a more refined domain.

Along with the issue of resolution, the procedure for calculating similarity has a significant effect on the performance of the reconstruction. This was seen in Figures II.24 thru II.36. From these we can see that inclusion of the tissue boundary in calculating similarity greatly hinders the ability of the reconstruction. The need to exclude the boundary in order to improve performance is not surprising. Within image registration it is generally accepted that similarity metrics should be used only on the intersection of the two images [64], else the chances for falling into local minimum are greatly increased. In our case this was true for both GC and MI.

With respect to the second characteristic used in calculating similarity, region size, GC and MI appear to behave opposite of each other (Figures II.24 and II.26). Surprisingly GC was able to produce a stable reconstruction with a number of measurements (71) below that of the unknown parameters (74). Where as MI is consistent with what has been historically true, that inverse problems need a ratio of roughly 2:1 between measurements and unknowns. Again the reason for this is not readily known. This does raise a concern that parameters of the algorithm are very much dependent upon the type of similarity metric being used. This is not necessarily bad, but it does require further investigation into the different function spaces of the particular metrics.

The computational experiments presented in Figures II.38 thru II.46 are intended to show the ability of the algorithm to detect multiple tumors. Figure II.45 fully demonstrates this capability, but it also introduces two important issues. The first is the idea of shading effects. As already mentioned tumor 3 was distinguishable in the distribution seen in II.43, yet in II.45 it could easily be mistaken for noise. This may be an issue of the deformations being dampened out by tumor 4. A result of this would be an insufficient change in image pattern in the area of tumor 3. A way to

test for this would be the introduction of multiple compressions where each is being applied from a different direction.

The second issue that arises from these results is blurring of tumor boundaries. In Figure II.41, tumor 1 and tumor2 are disjoint and easily distinguished as two tumors. Yet, in II.45 tumor 1 and tumor 4 appear to be almost joined, but it is apparent in II.44 that this is not the case. In this situation it is believed that both tumors 1 and 4 are influencing a common material property located between the two. This again raises the the issue of parent mesh resolution, and how a lack of appropriate resolution can result in reconstruction errors.

Similar to the multiple tumor test were the experiments dealing with variations in tumor location. These results are given in Figures II.48 thru II.52. The most encouraging points of these results are seen in Figures II.49 and II.51. The ability to reconstruct with such small differences in deformation again supports the idea that we are optimizing using subtle differences in image patterns. The tumor to tissue ratio is still seen to be low here. This may be due to spatial averaging with boundary nodes (boundary nodes are pinned to be $18kPa$).

We have already discussed issues dealing with parameter space resolution. Associated with this is the idea of tumor size. The results from the tests seen in Figures II.54 thru II.58 indicate that roughly $3mm$ is the lower limit in tumor size for the method. While this is important to know, more significant are the implications that can be derived from comparing Figure II.59 and Figure II.51. In these we see that a far field tumor of size $7mm$ produces deformations on the same scale as that of the $1mm$ tumor. Yet the $1mm$ could not be reconstructed (Figure II.55). This is the clearest example showing the importance of the parent mesh's resolution. We also believe that the location of the tumor relative to the reconstruction parameters plays an important role in detection capabilities.

Already we have seen that there several factors affecting the reconstructed modulus values. Primarily, these have dealt with issues of similarity and mesh attributes. Also important are the limitations associated with the plain-strain model. In our attempts to uncover these, we have focused on three aspects - tumor stiffness resolution, Poisson's ratio, and three-dimensional effects.

The results from the stiffness resolution experiments (II.61-II.62) introduce an interesting point that is not completely obvious; there is both an upper and lower threshold for relative differences in modulus at which similarity can be detected. This by itself is not necessarily provoking. What is, though, is that these thresholds are a result of both machine tolerance and image noise. The first becomes an issue when changes in similarity are below that which can be represented by the computer. This value can be determined. Image noise comes into play when changes in similarity are below the level of noise between the two images. A way of determining this may not be readily available. Therefore, this is a limiting factor. However, the literature has shown that differences in modulus values between healthy and cancerous tissue should be sufficiently above this level of noise.

The simulations which varied Poisson's ratio from its true value (see Figures II.63-II.68) indicate that reconstructions are sensitive to this value. Previous studies have shown that Poisson's ratio has a large influence on how well the boundary matches between the data and model generated image. Large differences obviously produce larger mismatches. This is significant due to the evidence shown by the Figures II.28 and II.29, which demonstrate how boundary information can have a significant impact on the reconstruction. It appears that differences on the order of 0.01, and greater, significantly hinder parameter estimation. Such high sensitivity will no doubt be an issue especially when considering clinical application. Exactly how we should approach this problem is not yet certain.

Knowing the importance that the boundary can play in the reconstruction, the results seen in Figures II.70 and II.71 give convincing evidence that three-dimensional effects are signficiant. First we notice that the 3D-generated image more closely matches the data image than the 2D image. Also, given the similarity between II.74 and II.75 it is very reasonable that the 3D model is more closely representing the deformations being produced. This strongly suggests that perhaps a three-dimensional elastic model should be implemented.

Without the speed improvements seen in Figures II.76 II.77 the degree to which this method has been evaluated would not have been possible. However, it can be clearly seen that we are quickly

deviating from the theoretical improvements. To improve this, restructuring of the code will be necessary. Currently, slave processors are only used in the construction of the Jacobian matrix. This creates a significant amount of idle time on these processors. With proper implementation this time could be used for other tasks, such as inversion of the Hessian matrix.

Considering the results from the CT data, Figures II.78 thru II.94, we see that there are a number of factors which can improve or worsen the reconstruction. The procedure that appears to have given the most improvement is image blurring. The best of these recontructions is seen in Figure II.83. This implies that the function space is extremely noisy. Such an issue is common among all elastography techniques. Along with this, we see in II.79, II.81, and II.82 that across slice averaging seems to have the greatest effect. An explanation for this is that across slice smoothing maybe helping to account for 3D-effects.

Within the data, three-dimensional inaccuracies appear to be very significant. To help account for this issue we proposed that reconstructions using slices from lower depths would be less influenced by the problem. However, we found that the middle slice, slice number 15, actually produced the most accurate results (II.88). This may be the optimial location due to three-dimensional effects actually being at a minimum at this position. A possible explanation for this is that strains in the z and $-z$ directions are canceling each other.

We believe that parent mesh resolution plays an important part in determining the ability of the algorithm to reconstruct appropriate modulus values. It was initially assumed that this technique would be similar to other inverse problems concerning resolution. Generally, inverse problems are required to operate in parameter spaces that are resolution limited by the number of measurements. Yet, amazingly, we were able produce a reasonable reconstruction using 379 properties, Figure II.91. We do not fully understand the implications of making internal measurements using similarity. Most inverse problems are primarily based on boundary data (which is the primary origin of the 2:1 ratio of data to the number of properties).

An important issue that became apparent, but was not a specific focus of any particular set of experiments, is divergence of the reconstruction. An example of this can be seen in the MRI

reconstruction (II.95 and II.96). An explanation for why the algorithm appears to be optimizing is that we are using Newton-Raphson. This is a root finding technique. We are not guaranteed to be reconstructing toward a minimum, much less the global minimum. Our assumption is that the initial guess is sufficiently close to the minimum. This is not necessarily true and techniques such as simulated annealing may need to be implemented.

Out of all experiments presented, the results seen in Figures II.95 and II.96 most clearly demonstrate the ability of similarity-based elastography to be modality independent. While the modulus values are not on the same scale as CT, the comparable localization is very exciting. Also promising, is the fact that the Poisson's ratio and the α regularization parameter are relatively close to those values used in CT, which show a dependence to the phantom rather than the modality. From this it is reasonable to assume that a general set of parameters can be found for clinical application.

CHAPTER III

CONCLUSIONS and FUTURE WORK

The results of this thesis show Similarity-Based Elastography to be a realistic goal for the clinical setting. While there are a number of issues that must first be overcome, these preliminary results clearly indicate that tissue parameters can be encoded within a set of images. Along with this is the idea that the encoding can be considered independent of the imaging modality. The reconstructions both in CT and MRI (Figures II.83 and II.96) strongly supports this claim.

Even though the results presented here are very promising, they have also indicated several shortcomings. These problems will be the focus of future work. For example, we know the model is inappropriate and does not fully represent the physics involved. Unfortunately, in order to provide adquate differences between pre- and post-compression images, this technique requires deformations that violate 2D small strain theory. To account for this a pseudo 3D model is being developed. The computational cost of the current method makes a full three-dimensional problem unfeasible. Therefore, rather than doing this, we are developing a model that is able to integrate 3D displacements into a 2D reconstruction technique.

The results also indicated that the method is very much dependent on an accurate value for Poisson's ratio. This has also been evident in the parameter searches that we have done. Not all have been presented. This is perhaps an unnecessary expense. A way to circumvent this tedious, time consuming task is to add Poisson's ratio as a parameter to be reconstructed. A question that will need to be answered when doing this is whether it should vary spatially or be considered a global parameter.

Speaking of computational cost, the parallelization of the code is a very initial attempt. Much more elegant techniques can be used to improve the speed of the program. For example PETSc [73], a parallel linear solving package, could be implemented to enhance performance of matrix inversion. Also, while significant efforts have been made to optimize the sequential portions of the

program, there is, without doubt, portions that can be further improved.

One issue that is secondary to the algorithm but essential for the continuation of this work, is the ability to do material testing of the phantoms. The goals for similarity-based elastography are both detection and diagnosis. Diagnosis is dependent upon at least being able to detect accurate contrasts within tissue if not true modulus values. The modulus reconstructions from the phantom data presented here provide the values of the parameters only to demonstrate the contrast of the image. We do not know the true values.

The image similarity based technique presented in this paper illustrates an alternative to displacement-based methods of elastography. While the dynamics between the distribution of tissue properties and the encoding of this information within the images is not yet fully understood, the results shown here indicate that image similarity can be a viable option for elasticity imaging. Given that this method relies solely on image intensity heterogeneity, the procedure has been cast as a modality independent elastography method. This fact makes this technique a truly novel approach to tissue property imaging. The potential for such a technique provides many exciting possibilities for future work and warrants more extensive research in model-based elasticity image reconstruction techniques.

REFERENCES

[1] The History of Cancer. Internet, 25 March 2002. http://www.cancer.org/docroot/CRI/content/CRI-2_6x_the_history_of_cancer_72.aspsitearea=CRI.

[2] C. K. Cross, J. Harris, and A. Recht. Race, Socioeconomic Status, and Breast Carcinoma in the U.S. *Cancer*, 95(9), 1 November 2002.

[3] New York Beth Israel Health Care System. Anatomy of the Breast. Internet, 2003. http://mammary.nih.gov/reviews/development/Human-breast001/.

[4] A. G. Thomas, A. Jemal, and M. J. Thun. Breast Cancer Facts & Figures 2001-2002. American Cancer Society, 2001.

[5] K. Johnston. Modelling the future costs of breast screening. *European Journal of Cancer*, 37:1752–1758, 2001.

[6] Costs of Cancer Care. Internet, 24 February 2003. http://progressreport.cancer.gov/doc.asp?pid=1&did=21&chid=13&coid=33&mid=vpeo.

[7] T. A. Krouskop, T. M. Wheeler, F. Kallel, B. S. Garra, and T. Hall. Elastic moduli of breast and prostate tissues under compression. *Ultrason Imaging*, 20(4):260–74., 1998.

[8] P. S. Wellman, R. D. Howe, E. Dalton, and K. A. Kern. Breast Tissue Stiffness in Compression is Correlated to Histological Diagnosis. Technical report, Harvard BioRobotics Laboratory, 1999.

[9] S. F. Levinson. Ultrasound propagation in anisotropic soft-tissues - the application of linear elastic theory. *Journal of Biomechanics*, 20(3), 1987.

[10] S. F. Levinson and V. L. Newhouse. The phase response of ultrasound to vibration - a method of measuring tissue elasticity. *Ultrasonic Imaging*, 10(1):74–75, 1988.

[11] M. Tristam, D. C. Barbosa, D. O. Cosgrove, J. C. Bamber, and C. R. Hill. Application of fourier-analysis to clinical-study of patterns of tissue movement. *Ultrasound in Medicine and Biology*, 14(8):695–707, 1988.

[12] E. Ueno, E. Tohno, S. Soeda, Y. Asaoka, K. Itoh, J. C. Bamber, M. Blaszczyk, J. Davey, and J. A. McKinna. Dynamic tests in real-time breast echography. *Ultrasound in Medicine and Biology*, 14:53–57, 1988.

[13] J. C. Bamber, L. Degonzalez, D. O. Cosgrove, and P. Simmons. Quantitative-evaluation of real-time ultrasound features of the breast. *Ultrasound in Medicine and Biology*, 14:81–87, 1988.

[14] R. J. Massay, R. B. Logansinclair, J. C. Bamber, and D. G. Gibson. Quantitative effects of speckle reduction on cross-sectional echocardiographic images. *British Heart Journal*, 62(4):298–304, 1989.

[15] C. R. Hill, J. C. Bamber, and D. O. Cosgrove. Performance criteria for quantitative ultrasonology and image parameterization. *Clinical Physics and Physiological Measurement*, 11:57–73, 1990.

[16] S. F. Levinson. Phase detection in acoustic microscopy. *Ultrasonic Imaging*, 12(4):292–308, 1990.

[17] J. Ophir, I. Cespedes, H. Ponnekanti, Y. Yazdi, and X. Li. Elastography - a quantitative method for imaging the elasticity of biological tissues. *Ultrasonic Imaging*, 13(2):111–134, 1991.

[18] I. Cespedes, J. Ophir, H. Ponnekanti, and N. Maklad. Elastography - elasticity imaging using ultrasound with application to muscle and breast in-vivo. *Ultrasonic Imaging*, 15(2):73–88, 1993.

[19] J. Ophir, F. Kallel, T. Varghese, M. Bertrand, I. Cespedes, and H. Ponnekanti. Elastography: A systems approach. *International Journal of Imaging Systems and Technology*, 8(1):89–103, 1997.

[20] B. S. Garra, E. I. Cespedes, J. Ophir, S. R. Spratt, R. A. Zuurbier, C. M. Magnant, and M. F. Pennanen. Elastography of breast lesions: Initial clinical results. *Radiology*, 202(1):79–86, 1997.

[21] M. Odonnell, A. R. Skovoroda, B. M. Shapo, and S. Y. Emelianov. Internal displacement and strain imaging using ultrasonic speckle tracking. *Ieee Transactions on Ultrasonics Ferroelectrics and Frequency Control*, 41(3):314–325, 1994.

[22] A. R. Skovoroda, M. A. Lubinski, S. Y. Emelianov, and M. O'Donnell. Reconstructive elasticity imaging for large deformations. *Ieee Transactions on Ultrasonics Ferroelectrics and Frequency Control*, 46(3):523–535, 1999.

[23] T. Varghese and J. Ophir. Experimental Corroboration of the Nonstationary Strain Estimation Errors in Elastography. *Ultrasound in Medicine and Biology*, 27(12):1677–1682, 2001.

[24] S. Srinivasan, F. Kallel, and J. Ophir. Estimating the Elastographic Signal-to-Noise Ratio Using Correlation Coefficients. *Ultrasound in Medicine and Biology*, 28(3):359–368, 2002.

[25] A. R. Skovoroda, S. Y. Emelianov, and M. O'Donnell. Tissue Elasticity Reconstruction Based on Ultrasonic Displacement and Strain Images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4), July 1995.

[26] F. Kallel and J. Ophir. A least-squares strain estimator for elastography. *Ultrasonic Imaging*, 19(3):195–208, 1997.

[27] E. E. Konofagou, T. Varghese, J. Ophir, and S. K. Alam. Power spectral strain estimators in elastography. *Ultrasound in Medicine and Biology*, 25(7):1115–1129, 1999.

[28] V. Dutt, R. R. Kinnick, R. Muthupillai, T. E. Oliphant, R. L. Ehman, and J. F. Greenleaf. Acoustic shear-wave imaging using echo ultrasound compared to magnetic resonance elastography. *Ultrasound in Medicine and Biology*, 26(3):397–403, 2000.

[29] A. Pesavento, A. Lorenz, S. Siebers, and H. Ermert. New real-time strain imaging concepts using diagnostic ultrasound. *Physics in Medicine and Biology*, 45(6):1423–1435, 2000.

[30] W. F. Walker, F. J. Fernandez, and L. A. Negron. A method of imaging viscoelastic parameters with acoustic radiation force. *Physics in Medicine and Biology*, 45:1437–1447, 2000.

[31] M. Fatemi and J. F. Greenleaf. Probing the dynamics of tissue at low frequencies with the radiation force of ultrasound. *Physics in Medicine and Biology*, 45(6):1449–1464, 2000.

[32] K. Nightingale, M. S. Soo, R. Nightingale, and G. Trahey. Acoustic Radiation Force Impulse Imaging: In Vivo Demonstration of Clinical Feasibility. *Ultrasound in Medicine and Biology*, 28(2):227–235, 2002.

[33] A. P. Sarvazyan, O. V. Rudenko, S. D. Swanson, J. B. Fowlkes, and S. Y. Emelianov. Shear wave elasticity imaging: A new ultrasonic technology of medical diagnostics. *Ultrasound in Medicine and Biology*, 24(9):1419–1435, 1998.

[34] L. S. Taylor, B. C. Porter, D. J. Rubens, and K. J. Parker. Three-dimensional sonoelastography: principles and practices. *Physics in Medicine and Biology*, 45(6):1477–1494, 2000.

[35] R. Muthupillai, D. J. Lomas, P. J. Rossman, J. F. Greenleaf, A. Manduca, and R. L. Ehman. Magnetic-resonance elastography by direct visualization of propagating acoustic strain waves. *Science*, 269(5232):1854–1857, 1995.

[36] T. L. Chenevert, A. R. Skovoroda, M. O'Donnell, and S. Y. Emelianov. Elasticity reconstructive imaging by means of stimulated echo mri. *Magnetic Resonance in Medicine*, 39(3):482–490, 1998.

[37] D. D. Steele, T. L. Chenevert, A. R. Skovoroda, and S. Y. Emelianov. Three-dimensional static displacement, stimulated echo nmr elasticity imaging. *Physics in Medicine and Biology*, 45(6):1633–1648, 2000.

[38] E. E. W. Van Houten, K. D. Paulsen, M. I. Miga, F. E. Kennedy, and J. B. Weaver. An overlapping subzone technique for mr-based elastic property reconstruction. *Magnetic Resonance in Medicine*, 42(4):779–786, 1999.

[39] R. Sinkus, J. Lorenzen, D. Schrader, M. Lorenzen, M. Dargatz, and D. Holz. High-resolution tensor mr elastography for breast tumour detection. *Physics in Medicine and Biology*, 45(6):1649–1664, 2000.

[40] D. B. Plewes, J. Bishop, A. Samani, and J. Sciarretta. *Physics in Medicine and Biology*, 45(6):1591–1610, 2000.

[41] A. L. McKnight, J. L. Kugel, P. J. Rossman, A. Manduca, L. C. Hartmann, and R. L. Ehman. MR Elastography of Breast Cancer. *American Journal of Radiology*, 178, June 2002.

[42] J. Lorenzen, R. Sinkus, M. Lorenzen, M. Dargatz, C. Leussler, P. Roschmann, and G. Adam.

[43] T. Wu, J. P. Felmlee, J. F. Greenleaf, S. J. Riederer, and R. L. Ehman. Assessment of thermal tissue ablation with mr elastography. *Magnetic Resonance in Medicine*, 45(1):80–87, 2001.

[44] F. Kallel, M. Bertrand, and J. Ophir. Fundamental limitations on the contrast-transfer efficiency in elastography: An analytic study. *Ultrasound in Medicine and Biology*, 22(4):463–470, 1996.

[45] M. M. Doyley, P. M. Meaney, and J. C. Bamber. Evaluation of an iterative reconstruction method for quantitative elastography. *Physics in Medicine and Biology*, 45(6):1521–1540, 2000.

[46] A. Manduca, V. Dutt, D. T. Borup, R. Muthupillai, R. L. Ehman, and J. F. Greenleaf. Reconstruction of elasticity and attenuation maps in shear wave imaging: An inverse approach. In *Medical Image Computing and Computer-Assisted Intervention - Miccai'98*, volume 1496 of *LECTURE NOTES IN COMPUTER SCIENCE*, pages 606–613. 1998.

[47] M. F. Insana, L. T. Cook, M. Bilgen, P. Chaturvedi, and Y. Zhu. Maximum-likelihood approach to strain imaging using ultrasound. *Journal of the Acoustical Society of America*, 107(3):1421–1434, 2000.

[48] M. Yamakawa and T. Shiina. Tissue Elasticity Reconstruction Based on 3-Dimensional Finite-Element Model. *Japanese Journal of Applied Physics*, 38, 1999.

[49] A. Samani, J. Bishop, and D. B. Plewes. A constrained modulus reconstruction technique for breast cancer assessment. *Ieee Transactions on Medical Imaging*, 20(9):877–885, 2001. English Article SEP IEEE TRANS MED IMAGING.

[50] M. Miga. A New Approach to Elastography Using Mutual Information and Finite Elements. *Physics in Medicine and Biology*, 48:467–480, 2003.

[51] A. P. Sarvazyan, A. R. Skovoroda, S. Y. Emelianov, J. B. Fowlkes, J. G. Pipe, R. S. Adler, R.B. Buxton, and P. L. Carson. Biophysical bases for elasticity imaging. *Acoustical Imaging*, 21:223–240, 1995.

[52] J. B. Fowlkes, S. Y. Emelianov, J. G. Pipe, A. R. Skovoroda, P. L. Carson, R. S. Adler, and A. P. Sarvazyan. Magnetic-resonance-imaging techniques for detection of elasticity variation. *Medical Physics*, 22(11):1771–1778, 1995.

[53] E. A. Zerhouni, D. M. Parish, W. J. Rogers, A. Yang, and E. P. Shapiro. Human-heart - tagging with mr imaging - a method for noninvasive assessment of myocardial motion. *Radiology*, 169(1):59–63, 1988.

[54] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: Application to breast mr images. *Ieee Transactions on Medical Imaging*, 18(8):712–721, 1999.

[55] J. M. Jr. Sullivan. *Integration of Automatic Mesh Generation For Simulation of Contaminants in Soils and GroundWater*, June 1994.

[56] D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal of Applied Mathematics*, 11:431–441, 1963.

[57] N. Joachimowicz, C. Pichot, and J. Hugonin. Inverse Scattering: An Iterative Numerical Method for Electromagnetic Imaging. *IEEE Transactions on Antennas and Propogation*, 39(2), December 1991.

[58] D. M. Mount. *ANN Programming Manaul*. Department of Computer Science and Institute for Advanced Computer Studies University of Maryland, College Park, Maryland, 0.1 edition, 1998.

[59] J. Bishop, A. Samani, J. Sciarretta, and D. B. Plewes. Two-dimensional mr elastrography with linear inversion reconstruction: methodology and noise analysis. *Physics in Medicine and Biology*, 45(8):2081–2091, 2000.

[60] A. Samani, J. Bishop, M. J. Yaffe, and D. B. Plewes. Biomechanical 3-d finite element modeling of the human breast using mri data. *Ieee Transactions on Medical Imaging*, 20(4):271–279, 2001. English Article APR IEEE TRANS MED IMAGING.

[61] A. P. Boresi. *Elasticity in Engineering Mechanics*. Prentice-Hall, Englewood Cliffs, NJ, 1965.

[62] O. C. Zienkiewicz and R. L. Taylor. *The Finite Element Method*. McGraw-Hill Book Company, London, 4 edition, 1989.

[63] K. D. Paulsen, P. M. Meaney, M. J. Moskowitz, and J. M. Sullivan. A dual mesh scheme for finite-element based reconstruction algorithms. *Ieee Transactions on Medical Imaging*, 14(3):504–514, 1995. SEP IEEE TRANS MED IMAGING.

[64] J. M. Fitzpatrick, D. L. G. Hill, and C. R. Maurer. Image registration. In M. Sonka and J. M. Fitzpatrick, editors, *Handbook of Medical Imaging*, volume 2, pages 447–513. SPIE Press, Bellingham, 2000.

[65] C. E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 623–656, 1948.

[66] P. Viola and W. M. III Wells. Alignment by Maximization of Mutual Information. *International Journal of Computer Vision*, 24(2):137–154, 1997.

[67] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *Ieee Transactions on Medical Imaging*, 16(2):187–198, 1997.

[68] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes. A Comparison of Similarity Measures for Use in 2-D-3-D Medical Image Registration. *IEEE Transactions on Medical Imaging*, 17(4):586–595, August 1998.

[69] L. Ibanez and W. Schroeder. *The ITK Software Guide*. Insight Consortium, 11 January 2003.

[70] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. SIAM, third edition, 1999.

[71] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra. *MPI- The Complete Reference*, volume 1. The MIT Press, Cambridge, MA, 2 edition, 1998.

[72] P. S. Pacheco. *Parallel Programming with MPI*. Morgan Kaufmann Publishers, Inc., San Francisco, CA, 1997.

[73] S. Balay, K. Buschelman, W. Gropp, D. Kaushik, M. Knepley, L. C. McInnes, B. Smith, and H. Zhang. *PETSc Users Manual*. Mathematics and Computer Science Division Argonne National Laboratory, http://www.mcs.anl.gov/petsc, 2.1.5 edition, 27 January 2003.