

Simulation-Based Joint Estimation of Body Deformation and Elasticity Parameters for Medical Image Analysis

Huai-Ping Lee

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
2012

Approved by:

Ming C. Lin

Mark Foskey

Marc Niethammer

Ed Chaney

Ron Alterovitz

© 2012
Huai-Ping Lee
ALL RIGHTS RESERVED

Abstract

**HUAI-PING LEE: Simulation-Based Joint Estimation of Body
Deformation and Elasticity Parameters for Medical Image Analysis.
(Under the direction of Ming C. Lin.)**

Elasticity parameter estimation is essential for generating accurate and controlled simulation results for computer animation and medical image analysis. However, finding the optimal parameters for a particular simulation often requires iterations of simulation, assessment, and adjustment and can become a tedious process. Elasticity values are especially important in medical image analysis, since cancerous tissues tend to be stiffer. Elastography is a popular type of method for finding stiffness values by reconstructing a dense displacement field from medical images taken during the application of forces or vibrations. These methods, however, are limited by the imaging modality and the force exertion or vibration actuation mechanisms, which can be complicated for deep-seated organs.

In this thesis, I present a novel method for reconstructing elasticity parameters without requiring a dense displacement field or a force exertion device. The method makes use of natural deformations within the patient and relies on surface information from segmented images taken on different days. The elasticity value of the target organ and boundary forces acting on surrounding organs are optimized with an iterative optimizer, within which the deformation is always generated by a physically-based simulator. Experimental results on real patient data are presented to show the positive correlation between recovered elasticity values and clinical prostate cancer stages.

Furthermore, to resolve the performance issue arising from the high dimensionality of boundary forces, I propose to use a reduced finite element model to improve the convergence of the optimizer. To find the set of bases to represent the dimensions for forces, a statistical training based on real patient data is performed. I demonstrate the trade-off between accuracy and performance by using different numbers of bases in the optimization using synthetic data. A speedup of more than an order of magnitude is observed without sacrificing too much accuracy in recovered elasticity.

Dedication

To my parents and Chen-Yi.

Acknowledgments

The past six years at UNC-Chapel Hill has been an exciting journey, and I was very lucky to be able to work with the best minds in the world. First and foremost, I would like to thank my advisor, Prof. Ming Lin, who guided and encouraged me through the PhD process. I also thank my committee members; Prof. Mark Foskey, who led me into the world of medical imaging and paper writing, and had been co-advising me until he took his talent to Morphormics (now part of Accuray); Prof. Marc Niethammer, who kept my math and logic as clean as possible; Prof. Ed Chaney, who gave me a lot of insights in cancer treatment and staging; and Prof. Ron Alterovitz, who provided discussions and comments about medical simulations.

I thank previous and present members of GAMMA group and MIDAG for discussion and feedback on my work, in particular, Dr. Vivek Kwatra, Rahul Narain, Yero Yeh, Prof. Stephen Pizer, Prof. Steve Marron, Josh Levy, Rohit Saboo, and Chen-Rui Chou. My presentations would not have the same quality without the feedback and criticism from group meetings and discussions, especially from Sean Curtis. I also thank people at Department of Radiation Oncology, especially Sue Xu, Gregg Tracton, Graham Gash, and Dr. Ron Chen, for providing prostate patient data and technical support.

It would be impossible for me to study abroad without support from my parents, who put my education above everything else. I am also grateful to my friends in Taiwan, in Chapel Hill, and around the world, for sharing laughs and keeping me motivated throughout the years.

Finally, my life would not be as enjoyable without my wife, Chen-Yi, who accompanies and supports me on the road to the PhD and beyond.

Table of Contents

List of Tables	x
List of Figures	xii
1 Introduction	1
1.1 Elastography	4
1.2 Deformable Registration	5
1.3 Parameter Estimation in Computer Graphics	6
1.4 Joint Estimation of Deformation and Elasticity Parameters	6
1.5 Thesis Statement	8
1.6 Main Results	9
1.6.1 Elasticity Estimation for Noninvasive Cancer Stage Assessment	9
1.6.2 Applications in Medical Image Analysis for Radiotherapy	11
1.6.3 Application in Physically-Based Animation	13
1.6.4 Acceleration of Optimization Framework Using Reduced-Dimension Finite Element Modeling	14
1.7 Organization	16
2 Previous Work	18
2.1 Elastography	18
2.1.1 Displacement Field Estimation	19
2.1.2 Inverse Modeling of Elastic Deformation	20
2.1.3 Other Elasticity Estimation Methods	20

2.2	Deformable Image Registration	21
2.2.1	Image-Based Methods	22
2.2.2	Landmark-Based Methods	23
2.2.3	Simulation-Based Methods	23
2.3	Parameter Estimation in Computer Graphics	24
2.3.1	Directing Simulations	24
2.3.2	Estimating Material Properties	25
3	Simulation-Based Joint Estimation of Deformation and Elasticity Pa- rameters	26
3.1	Introduction	27
3.2	Method	30
3.2.1	Linear Elasticity Model and Finite Element Modeling	30
3.2.2	Sensitivity Study	32
3.2.3	Distance-Based Objective Function	34
3.2.4	Numerical Optimization	36
3.2.5	Initial Guess of Parameters	37
3.3	Experiments	38
3.3.1	Synthetic Scene with Multiple Organs	40
3.3.2	Noninvasive Assessment of Prostate Cancer Stage	42
3.3.3	Study: Inhomogeneous Materials	48
3.3.4	Application: Registration of Segmented CT Images	51
3.3.5	Extension: Nonlinear FEM	52
3.4	Conclusion and Future Work	54
4	Treatment Image Estimation Based on Marker Locations for Dose Calculation	58
4.1	Introduction	59

4.2	Deformation Estimation Based on M-reps	62
4.2.1	M-reps	62
4.2.2	Rotational Flows and Deformation Field Generation	64
4.3	Simulation-Based Estimation of Deformation	65
4.3.1	Finite Element Modeling	65
4.3.2	Optimization	66
4.4	Dose Calculation	67
4.5	Experiments	68
4.5.1	Comparison of Organ Surfaces	68
4.5.2	Comparison of Calculated Dose	69
4.6	Summary	71
5	Fast Optimization-Based Elasticity Parameter Estimation Using Reduced Models	76
5.1	Introduction	77
5.2	Previous Work	79
5.3	Method	83
5.3.1	Optimization-Based Framework with Finite Element Modeling	84
5.3.2	Acceleration using Reduced-Dimension Modeling	85
5.4	Implementation and Results	88
5.4.1	Physically-Based 2D Shape Deformation	88
5.4.2	3D Elastography using Reduced Modeling	89
5.5	Conclusion	95
6	Discussion	97
6.1	Summary of Results	97
6.2	Future Work	99

Bibliography 101

List of Tables

3.1	Error in recovered modulus of elasticity in two synthetic multi-organ scenes; note that the error becomes much larger for elasticity values greater than 150 kPa.	41
3.2	Average and standard deviation of elasticity values for the prostate recovered from the patient data sets; the last column is the clinical cancer staging for the tumor for each patient.	45
3.3	Definition of clinical T-stages for prostate cancer	46
3.4	The recovered elasticity values for the prostate as a homogeneous material, when the organ contains a synthetic tumor of different sizes and a normal tissue; elasticity values are set to 100 kPa for the tumor and 50 kPa for normal prostate tissue.	48
3.5	Average error in landmark positions (distance in cm) inside the prostate, computed with the Demons method and my method; t-tests show that my method performs statistically significantly better in three of the five data sets. (If Bonferroni correction is used, my method is significantly better only for Patient 3.)	52
3.6	Error in recovered nonlinear modulus of elasticity in two synthetic multi-organ scenes.	53
3.7	Error in recovered nonlinear modulus of elasticity in two synthetic multi-organ scenes where the elasticity of surrounding tissue is doubled (20 kPa) when generating the synthetic data. The surrounding tissue elasticity is still set to 10 kPa in the optimization process, and I expect to see recovered values for the prostate to be half of the true values.	54
3.8	Average and standard deviation of elasticity values for the prostate recovered from the patient data sets using nonlinear FEM; the last column is the clinical cancer staging for the tumor for each patient.	55
4.1	Average errors in prostate surface and volume for the four patient data sets. The surface error is measured with average distance between two sets of sample points (D), and the volume overlap is measured with the ratio of intersection and union of the two volumes ($\frac{\text{Intersection}}{\text{Union}}$).	69

4.2	Comparison of equivalent uniform dose of the dose-volume histograms for the prostate for the four patient data sets; the errors resulted from estimated images (using m-rep and FE simulation) are relative to the values given by the IGRT (using real treatment images). Each patient is experimented on using five target volume margins around the prostate. .	73
4.3	Comparison of equivalent uniform dose of the dose-volume histograms for the anterior rectal wall for the four patient data sets; the errors resulted from estimated images (using m-rep and FE simulation) are relative to the values given by the IGRT (using real treatment images). Each patient is experimented on using five target volume margins around the prostate.	74
5.1	Error in recovered modulus of elasticity for two synthetic multi-organ scenes using a reduced-dimension FEM; the 18 to 62 times speedup is compared against an optimization using a full-dimension FEM.	95

List of Figures

1.1	Digital rectal exam (DRE); the doctor inserts a gloved, lubricated finger into the rectum to feel and check the prostate. Courtesy of National Cancer Institute.	2
1.2	An example elastography result for the prostate using a rectal probe. Left: overlaid colormap represents recovered elasticity. Right: corresponding ultrasound image. Courtesy of Hitachi Medical Systems.	3
1.3	Example 2D shape deformations showing different results with different Poisson’s ratios ν ; the boundary conditions are the same for both sets of objects (the gravity).	7
1.4	Overview of the optimization framework; the simulation parameters are updated according to the current error in deformed surface and fed into the simulator for the next iteration.	7
1.5	A typical scene of the male pelvis region: the movement of the bladder and rectum deforms the prostate. Left: a 3D rendering of the organ surfaces; the white surface represents the target surface of the prostate. Right: a slice of a CT image of the area; the red contour shows the segmentations of the reference image, and the blue contour shows the prostate in the target image.	9
1.6	Box plot showing the range of recovered elasticity values (Y axis) for patients within each cancer stage (X axis), where the top and bottom of the boxes show the 25th and 75th quartiles, and top and bottom lines show the maximum and minimum values; the data shows significant positive correlation between the two values (the estimated p-value is 0.024 using Spearman’s rank correlation).	10
1.7	A typical setting of the Calypso system. Left: three transponders are implanted into the prostate to provide localization; courtesy of Calypso. Right: a CT image slice showing the prostate and two of the transponders.	12
1.8	The initial shape (left), target shape (middle), and the plot of error versus Poisson’s ratio (right), where the error is based on the distance between corresponding vertices; the recovered value is 0.48 (ground truth is 0.49).	14
1.9	Visualization of the first two principal components (reduced bases) found by the statistical training of surface deformation.	15

1.10	The trade-off between speed and accuracy is achieved by using a different number of bases for representing nodal displacements in the FE model.	16
3.1	A sliced view of the synthetic scene, which consists of two concentric spheres; the inner (red) and outer (green) regions have different stiffness values (blue triangles represent outer surface, which is considered part of the green region).	32
3.2	The plots of the radius of the inner sphere (in cm) after deformation: (a) inner radius versus elasticity value (in kPa) of the inner region; (b) inner radius versus magnitude of forces (in N) acting on the outer surface; (c) inner radius (z -coordinate) versus elasticity and force magnitude with isocontours of inner radius on xy -plane; (d) outer radius versus elasticity; (e) outer radius versus magnitude of forces; (f) outer radius (z -coordinate) versus elasticity and force magnitude with isocontours of outer radius on xy -plane. The radii before deformation are 3 cm and 3.75 cm for two spheres, respectively, and the elasticity for the outer region is 10 kPa. The Poisson's ratios are fixed to 0.40 and 0.35 for the two regions, respectively.	33
3.3	Input to my algorithm: (a) a sliced view of the tetrahedral model of the moving image (light-blue triangles represent surfaces, not FEM regions; bladder and rectum are hollow); (b) a slice of the distance map of the prostate surface in the reference image.	36
3.4	Plot of Φ and \mathbf{E} (in kPa) with several sample values for finding an initial guess of elasticity value in a synthetic multi-organ scene. The plot suggests that the best initial guess is 50 kPa.	38
3.5	The moving surfaces and ground-truth boundary conditions in the two synthetic multi-organ scenes: the arrows shows Dirichlet boundary conditions applied to surfaces of bladder and rectum; the scaling of arrows are according to the magnitude of displacements.	41
3.6	Plots of relative errors in recovered elasticity vs. different Poisson's ratios for the prostate (0.3, 0.35, 0.4, 0.45, and 0.49) used for generating the synthetic surface data; each plot shows the result from one test scene, and each curve represents a true elasticity value (50, 100, and 150 kPa) used in the synthetic case. During the optimization process, the assumed Poisson's ratio is always fixed to 0.4.	43
3.7	Histogram of distances between the pairs of corresponding reference and target prostate meshes used in my experiments on segmented CT images.	44

3.8	Convergence graphs (plot of Φ and $\ \nabla\Phi\ $ versus iteration number) for a pair of CT image data: (left) convergence of the external forces; (right) convergence of the elasticity.	45
3.9	Box plot of average recovered elasticity value and cancer T stage for each patient data set shown in Table 3.2.	47
3.10	A sliced view of the tetrahedral mesh with a tumor (yellow) embedded in the prostate (red); the mesh is used to generate the synthetic target surface, while the prostate is still considered homogeneous in the optimization process.	49
3.11	Close-up view of the surfaces before (left) and after (right) deformation; the transparent white surface shown is the target surface of the prostate. Notice how the prostate surface move towards the white surface. Bladder and rectum surfaces are those with external forces applied.	49
3.12	Registration results for a pair of test images: (a) axial and sagittal views of the moving image, and a 4x4 checkerboard comparison with the planning image, before registration; (b) the two views of the registered image, along with a checkerboard comparison with the planning image; superimposed by segmentations of the reference image, shown in red, and the segmentation of the prostate in the daily image, shown in blue; notice that the image deforms towards the red contours.	50
3.13	Box plot of average recovered nonlinear elasticity value and cancer T stage for each patient data set shown in Table 3.8.	55
4.1	The Calypso System; three electromagnetic transponders are implanted in the prostate (left) prior to the treatment planning and are tracked in real time during each treatment fraction to provide localization (right). Courtesy of Calypso.	60
4.2	A simple 3×4 m-rep. A single shape is represented by the locations of the hubs and the lengths and directions of the spokes. The boundary of the object (not shown) passes through the ends of the spokes.	63
4.3	3D rotational flow between two oriented points $(\mathbf{x}_0, \mathbf{E}_0)$ and $(\mathbf{x}_1, \mathbf{E}_1)$ lies on a helix whose axis is the rotation axis between \mathbf{E}_0 and \mathbf{E}_1 . Courtesy of Levy et al. [2008].	64
4.4	A finite element mesh used for estimating the deformation using the simulation-based approach. The red elements represent the prostate.	66

4.5	Two examples of nodal displacement field resulted from the simulation-based optimization framework; the white transparent surface is the prostate.	67
4.6	Axial (left), coronal (center), and sagittal (right) views of an example m-rep comparison; blue contour: m-rep fitted to Calypso marker locations during the treatment fraction; green contour: surface deformed by the simulation-based scheme; red contour: m-rep created with real treatment image; the green crosshairs show the Calypso origin, which also serves as the isocenter of the treatment plan; (a, c) CT images estimated using fitted m-reps and the simulation-based scheme, respectively; (b, d) 4×4 checkerboard images comparing (a) and (c) to the real treatment image, respectively.	70
4.7	Differential (a) and cumulative (b) dose-volume histograms for the prostate (left) and anterior rectal wall (right); blue lines represent results with real treatment images, green lines represent results with images estimated using m-rep fitting, and red lines represent results with images estimated using FE simulation.	72
5.1	Flow chart of the optimization loop; the displacement field generated by the simulator is used in the objective function for the optimizer to update the parameters; updated parameters are fed back into the simulator, and so on.	83
5.2	The process of generating the the node correspondence and example deformations for one source surface; this procedure is repeated for multiple source meshes, and the resulting reduced basis can be applied to other FE meshes using the same node correspondence procedure.	88
5.3	Example 2D shape deformations showing different results with different Poisson's ratios ν (each column results from the same first-type boundary conditions); top: $\nu = 0.1$; middle: $\nu = 0.49$; bottom: comparison by overlaying the two results.	89
5.4	Flow chart of Poisson's ratio estimation given source and target 2D shapes; an FE model is built according to the source shape, and the distance between the deformed source shape and the target shape is minimized. The red cross marks the nodes with boundary conditions applied.	90
5.5	Test of accuracy of the recovered Poisson's ratios ν ; the four target shapes are generated using four different ν values; for each subfigure, on the left is the target shape, and on the right is the plot of surface errors versus ν values.	91

5.6	The tetrahedral mesh of the bladder-prostate-rectum area (left) and a close-up view of the three organs (right), where the white surface is the target prostate surface.	92
5.7	The first two principal components (visualized as nodal displacements on bladder and rectum surfaces) resulted from a PCA using 108 example deformations.	93
5.8	The number of reduced basis vectors provides a means of trade-off between optimizer performance and accuracy of recovered elasticity values; the error (blue) decreases and the number of optimization iterations (green) increases with the number of basis vectors used. The slight increase in error with 139 basis vectors is likely the result of optimizer sticking to a local minimum.	94

Chapter 1

Introduction

Screening and treatment of cancer are an important topic, since cancer is the second most common cause of death in the US [American Cancer Society, 2012]. In particular, prostate cancer is the second leading cause of cancer-related death in men, and there is no standard or routine screening procedure yet. A popular method for prostate cancer screening is the digital rectal exam (DRE), which involves inserting a gloved, lubricated finger into the patient's rectum to palpate and check for any abnormality in the prostate, as shown in Fig. 1.1. The palpation is an example of cancer detection methods based on elasticity, assuming that cancerous tissues are stiffer than normal ones. A palpation, however, provides only qualitative assessment, and the small size or deep location can make it impossible to detect the difference in stiffness by palpation. Instead, automatic approaches based on medical image analysis can be used to estimate internal tissue elasticity quantitatively. For example, elasticity reconstruction, or elastography [Ophir et al., 1999], is a non-invasive method for estimating elastic tissue properties, based on comparison of pre- and post-compression images and measured external forces. Elastography is an inverse problem to elastic body simulation: the elasticity is recovered from displacement and external forces. A rectal ultrasound probe is a popular device for elastography of prostate, where the external forces are provided by the probe, and the speckles in the ultrasound image provides an estimation of pixel-wise displacement field. Fig. 1.2 shows an example result of ultrasound elastography for the prostate, where the




Figure 1.1: Digital rectal exam (DRE); the doctor inserts a gloved, lubricated finger into the rectum to feel and check the prostate. Courtesy of National Cancer Institute.

color for each pixel represents elasticity values recovered from the displacement and the forces.

Elastic body simulation is also useful in cancer treatment methods such as image-guided radiotherapy (IGRT) and image-guided surgery. In an IGRT, the treatment plan is made according to the planning image, and a treatment image is taken before each treatment fraction to align the patient with the plan. Therefore, there is a need for finding the rigid and non-rigid transformation from the treatment image to the planning image to perform the alignment. An image registration method aims to find such a pixel-wise correspondence between the two images. Besides methods based on directly minimizing intensity errors and segmented surface distances, a physically-based simulation can also be used to find such a correspondence. The boundary conditions (prescribed nodal displacement or forces) need to be chosen carefully to match the deformed surfaces or image intensities, and the material properties are usually chosen according to *ex vivo* experiments on human tissues [Krouskop et al., 1998; Umut Ozcan et al., 2011]. In image guided surgery, it is also important to register the preoperative image to intra-operative




Figure 1.2: An example elastography result for the prostate using a rectal probe. Left: overlaid colormap represents recovered elasticity. Right: corresponding ultrasound image. Courtesy of Hitachi Medical Systems.

anatomy (landmarks) in order to compensate for the deformation during the surgery. The finite element method (FEM) [Zienkiewicz and Taylor, 2005] has been applied to model the biomechanics of the organ of interest, and the boundary conditions based on landmark matching is the key to a successful registration [Cash et al., 2005].

Material properties are also essential in surgical simulations targeted at a specific patient, since tissue properties differ from patient to patient. The patient-specific modeling of a surgical simulation is especially important for treatment planning. For example, a simulator can be useful for planning of a brachytherapy [Chentanez et al., 2009] for treating prostate cancer, and different prostate stiffness can produce different simulation results and thus change the optimal way of inserting the needles [Salcudean et al., 2006]. Since *in vivo* experiments on tissue properties are usually unavailable before treatment planning, the elastography becomes the main method in modeling surgical simulations.

1.1 Elastography

The elastography can be illustrated using a linear static finite element method. The basic idea of the FEM is to discretize a continuous domain with discrete elements and nodes, and only the values at the nodes are considered in the discretized equations. The linear static isotropic elasticity model and FEM gives the linear equation

$$\mathbf{K}\mathbf{u} = \mathbf{f},$$

where \mathbf{u} is the vector of nodal displacements, \mathbf{f} is the vector of external forces, and \mathbf{K} is the stiffness matrix determined by node connectivities and material properties. A forward FEM-based simulation uses a volumetric mesh, a set of material properties, and boundary conditions to build the matrix \mathbf{K} and vector \mathbf{f} and solve for the displacement \mathbf{u} . An elastography is the inverse problem of the forward simulation and aims to solve for the elasticity by finding nodal displacements and external forces.

Existing elastography methods usually requires a dense (pixel-wise) displacement field and a force exertion/measurement device. The displacements and external forces are considered as known values, and the elasticity can be computed directly as an inverse problem of the linear FEM [Zhu et al., 2003] or by iteratively minimizing errors in the displacement field [Kallel and Bertrand, 1996]. Since the effect of a static compression is usually limited to superficial tissues, more recent methods use vibrations instead to cause the displacement [Chopra et al., 2009]. In order to reach deep-seated organs such as the prostate, the vibration actuator needs to be inserted into the rectum or urethra. For other internal organs such as the liver, however, the ability to apply vibration could be limited by the patient's fat tissues. Moreover, the modality of the images is also limited: in images where the intensity is almost constant, the dense displacement field can only be approximated, and these elastography methods may not be as reliable.

1.2 Deformable Registration

The result of a deformable registration is usually represented as a pixel-wise displacement field. A popular registration method is to minimize pixel-wise intensity errors by treating the displacement vectors as the parameters to the optimization [Thirion, 1998]. The image-based method is likely to stick to a local minimum because the displacement for each vector is decided using only local intensity information. A regularization term is therefore needed to resolve ambiguities for uniform regions and to enforce some quality measures of the displacement field, such as the smoothness. Moreover, the intensity values are ambiguous — a particular value could represent different tissues in an image. If landmarks (local feature points) can be found in the two images, the matching of corresponding landmarks can provide more robust registration than matching intensities [Shen and Davatzikos, 2002]. However, it is often impossible to find a very dense set of landmarks, and displacement vectors for most pixels are affected by both the landmark matching and the interpolation method.

Physically-based simulations, such as the FEM-based approach described in Section 1.1, provide displacement fields of high quality and physical accuracy. And since internal nodes are also simulated in the FEM, the dependence of the registration on the interpolation is minimized. However, the material properties (needed for stiffness matrix \mathbf{K}) and boundary conditions (for force vector \mathbf{f}) are unknown. A surface matching can provide estimated boundary conditions [Ferrant et al., 2000], but geometrically deformed surfaces may have the problem of vertex redistribution. On the other hand, to decide the material properties, cycles of simulation-evaluation-adjustment may be needed. Notice that in some of the simulation-based registration methods, the material properties are less important than the deformation, since when there are boundary conditions applied to all the surface nodes, a good registration can always be achieved, if the materials are soft enough: one can apply larger forces on a stiffer material to achieve the same deformation. Also notice that the term “elastic registration” is often used to describe

methods using regularization terms based on elastic energy (computed with derivatives of the deformation field), and these “elastic” methods does not utilize simulations of elastic bodies.

1.3 Parameter Estimation in Computer Graphics

In computer animation, physically-based simulations can help generate realistic effects such as smoke, liquid, and elastic deformation, without low-level manipulations. In many cases, however, the animator would like high-level control of the physical effect through adjusting simulation parameters. For example, particular shapes of fluid can be achieved by adding artificial forces while keeping physical plausibility [Treuille et al., 2003; McNamara et al., 2004]. 2D animations can also be controlled by simulation parameters. Fig. 1.3 shows 2D shapes resulted from simulations using different Poisson’s ratios (compressibility). The look and feel of simulated cloth and soft objects, on the other hand, depend on the physical properties rather than the exerted forces. In order to achieve a certain effect, the material properties need to be estimated through tedious cycles of adjustment, simulation, and evaluation. Optimization-based methods have been proposed to automatically estimate material properties based on captured images/videos and measured forces [Bhat et al., 2003; Syllebranque and Boivin, 2008; Bickel et al., 2009]. Most of these methods require special image/video capturing systems and force measuring devices, which may limit their applicability.

1.4 Joint Estimation of Deformation and Elasticity Parameters

In this thesis, I aim to simultaneously solve the two main problems, elastography and deformable registration, by presenting a general deformation and elasticity parameter

Poisson's Ratio:
0.49 (dark colors) vs. 0.10 (light colors)



Figure 1.3: Example 2D shape deformations showing different results with different Poisson's ratios ν ; the boundary conditions are the same for both sets of objects (the gravity).

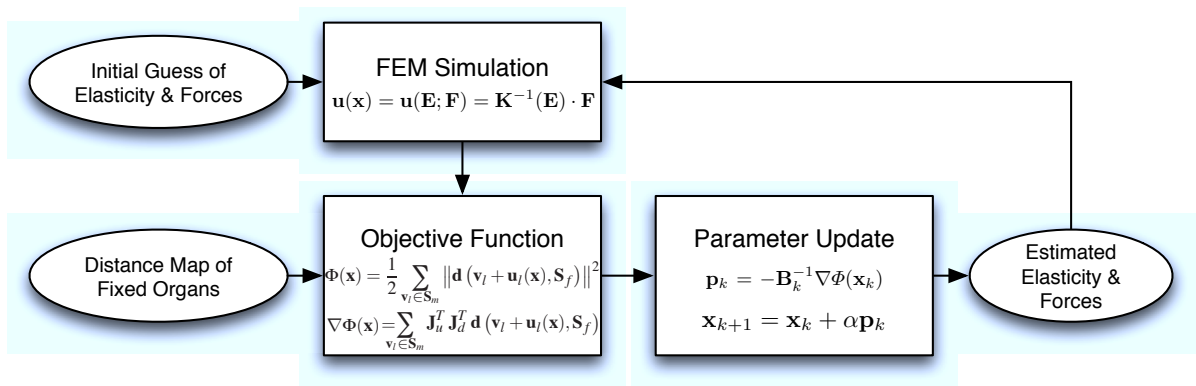


Figure 1.4: Overview of the optimization framework; the simulation parameters are updated according to the current error in deformed surface and fed into the simulator for the next iteration.

estimation method that only depends on surface information from multiple images, and the movement is provided by natural organ movements. The deformation is not known and is computed as an end result of the algorithm. The framework is built upon a physically-based simulator that requires some simulation parameters (such as material properties and boundary conditions) and generates nodal displacements, which are used to deform the surfaces. The distance between the deformed and the target surfaces are minimized in an optimization loop: in each iteration, an updated set of simulation parameters is computed according to the distance-based objective function, and the parameters are fed back to the simulator for the next iteration, and so on, as shown in Fig. 1.4.

My method has two main advantage over previous elastography methods. Firstly, since only surface information is used, any image that can be segmented can be used in my framework. This also means that the resolution of the resulting elastogram is limited to organ boundaries, but I argue that the recovered “average” stiffness of an organ also reflects the combination of tissues of different elasticities in it. Second, the problem of force exertion or vibration actuation is avoided since I do not actively generate the displacements but only observe the movements from the images. Although the absolute elasticity values cannot be recovered due to the lack of force measurements, I am able to find the ratio of elasticity values when there is more than one material.

1.5 Thesis Statement

Physically-based modeling and numerical optimization can be applied to simultaneously estimate deformation and material properties efficiently from deformed surfaces for elastography and deformable registration.

To support this thesis, I present a general framework for estimating elasticity parameters and boundary forces based on surface information without knowing the deformation field or external forces. The feasibility of the framework is examined by showing a significant positive correlation between recovered elasticity values and clinical prostate cancer stages. To show potential applications of the framework to assessment of a radiotherapy, I apply the optimization scheme to the deformable image registration problem, as well as the estimation of treatment images for dose calculation in the setting of a marker tracking system. The application in physically-based animation is shown with an example parameter estimation for 2D elastic body animation, where the Poisson’s ratio is recovered from the source and target shapes. Finally, to address the performance issue due to high dimensionality of external forces, I introduce an acceleration method using reduced-dimension finite element modeling, where the reduced basis is trained




Figure 1.5: A typical scene of the male pelvis region: the movement of the bladder and rectum deforms the prostate. Left: a 3D rendering of the organ surfaces; the white surface represents the target surface of the prostate. Right: a slice of a CT image of the area; the red contour shows the segmentations of the reference image, and the blue contour shows the prostate in the target image.

statistically across multiple patient data sets.

1.6 Main Results

1.6.1 Elasticity Estimation for Noninvasive Cancer Stage Assessment

In Chapter 3, I apply my optimization framework to estimate the stiffness of the prostate in real patient data and show that the recovered elasticity values correlate to prostate cancer staging. Fig. 1.5 shows a typical scene of the male pelvis region that I focus on in the experiments. In this area, the prostate is usually pushed around by the bladder and rectum, which deforms a lot due to different amounts of contents from time to time.

The clinical application I propose is complementary to existing elastography methods. When a number of images from the same patient is available, my method provides an additional indication of cancer staging. The experimental results show a significant positive correlation between recovered stiffness values and the cancer stages, as shown in

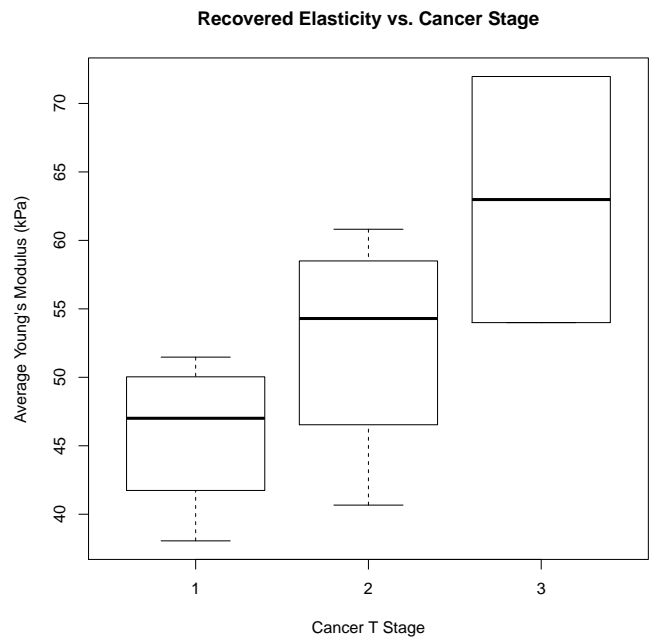


Figure 1.6: Box plot showing the range of recovered elasticity values (Y axis) for patients within each cancer stage (X axis), where the top and bottom of the boxes show the 25th and 75th quartiles, and top and bottom lines show the maximum and minimum values; the data shows significant positive correlation between the two values (the estimated p-value is 0.024 using Spearman's rank correlation).

Fig. 1.6. I also analyze the effect of inaccurate Poisson’s ratio and different combination of normal and cancerous tissues within the prostate. Furthermore, a nonlinear FEM can be integrated with my framework. These results were published in [Lee et al., 2012] and are presented in Chapter 3.

1.6.2 Applications in Medical Image Analysis for Radiotherapy

In an image-guided radiotherapy of prostate cancer, a CT image is taken before the treatment, and the treatment plan is made according to this CT image (*planning image*). On a predetermined subset of treatment fractions, in order to align the patient with the plan, another CT image (*treatment image*) is taken, and the patient is moved accordingly. A treatment CT images also provides an estimation of the radiation dose delivered to each part of the tissue, since the X-ray intensity represents the absorbed radiation. If the delivered dose deviates too much from the treatment plan, the plan can be modified accordingly. Therefore, finding the transformation (rigid and non-rigid) from the treatment space to the planning space is essential in evaluating delivered dose in a radiotherapy. The process of finding a transformation that matches two images is called image registration. Details are given in Chapter 2.

1.6.2.1 Physically-Based Image Registration

While the rigid registration is relatively easy, nonrigid image registration remains difficult since the movement of internal organs is complicated. A physically-based registration approach uses a simulator to generate a deformation field that matches the images, and the key to a good matching are the material properties and boundary conditions. However, most existing methods use surface matching to provide boundary conditions and use hand-picked material properties. My optimization framework provides a means of automatically finding the optimal elasticity values and boundary conditions. In contrast to registration methods that seek to maximize image similarities, my method ensures

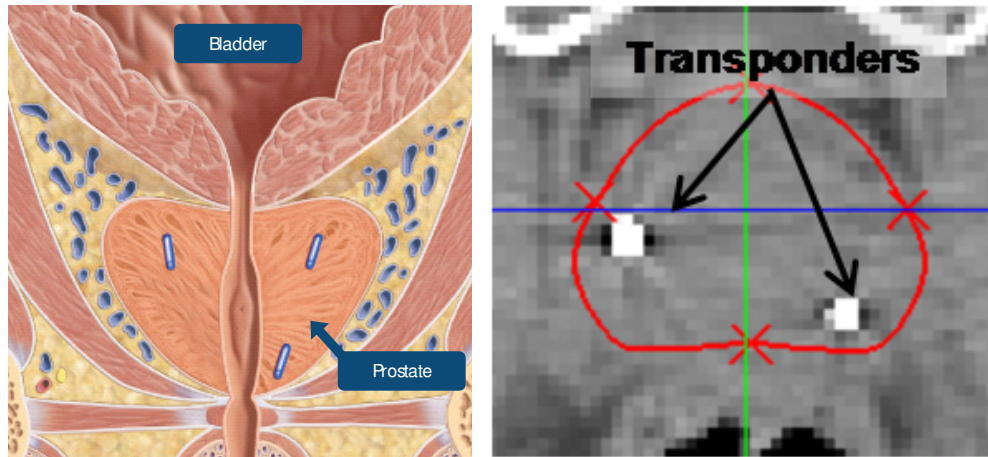


Figure 1.7: A typical setting of the Calypso system. Left: three transponders are implanted into the prostate to provide localization; courtesy of Calypso. Right: a CT image slice showing the prostate and two of the transponders.

quality of deformation since it is always generated by an FE simulator, assuming the FE model is suitable for the patient. In contrast to existing FEM-based methods, mine does not require hand picking the material properties. I experimented with CT images of the male pelvis region and compared the results with a popular image-based method. The results were published in [Lee et al., 2010b] and are presented in Chapter 3.

1.6.2.2 Treatment Image Estimation by Matching Implanted Markers

Instead of treatment images, a GPS-like system [Langen et al., 2008; Balter et al., 2005] can also be used to track a few markers implanted into the target organ. Fig. 1.7 shows a typical setting of such a system, where three electromagnetic transponders are implanted into the prostate. With such a tracking system, the target organ can be located without taking a treatment image. Such a system not only provide a more accurate localization than a CT image but also provide localization *during* the treatment fraction. However, since the treatment images are missing, there is no way to estimate the delivered dose.

I propose two methods to estimate the treatment image according to the planning image and the marker locations. One is based on *m-reps* [Pizer et al., 2003] fitted to the marker locations, providing displacements at the sample points on the surface. The

m-reps provide a coordinate system that represents any location within or near the surface of the object, including the marker locations. By fitting the planning m-rep to the marker locations in the treatment space, displacements on the surface sample points can be estimated. The dense deformation field is then interpolated for the entire region of interest, and the planning image warped using the resulting deformation serves as the estimated treatment image.

The other proposed method is based on the simulation-based optimization framework: I find the optimal external forces that minimize the error in marker locations. In this case, the parameters to the optimizer are the external forces acting on the surface nodes of the prostate, and the objective function is based on distance between deformed marker locations and the target locations. The optimal external forces gives the deformation field that matches the marker locations in the planning and treatment spaces and is used to generate the estimated treatment image. The simulation-based approach results in more physically accurate deformations for pixels far from sample points (in this work, the organ surface), since the entire 3D domain is considered in the FEM. On the other hand, the deformation generated by m-reps depends on the interpolation algorithm because only the deformation at sample points is given by the m-reps correspondence.

I assess the effectiveness of the estimated treatment image by comparing the deformed planning models (fitted to treatment marker locations) with the segmentation from real treatment image. Dose calculation results using the real and estimated treatment images are used to show the feasibility of the image estimation scheme. These results were published in [Lee et al., 2010a] and described in detail in Chapter 4.

1.6.3 Application in Physically-Based Animation

The selection of simulation parameters has been a difficulty in applying a physically-based simulation for generating animations. In 2D animations, for example, the notion of “stiffness” and “compressibility” has been studied in the literature, but heuristics are

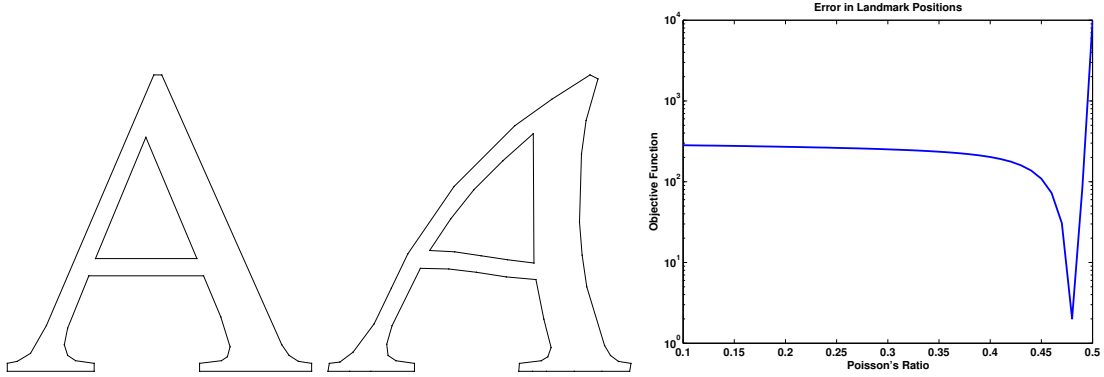


Figure 1.8: The initial shape (left), target shape (middle), and the plot of error versus Poisson’s ratio (right), where the error is based on the distance between corresponding vertices; the recovered value is 0.48 (ground truth is 0.49).

still much more popular than simulating animated objects as elastic bodies. I apply my optimization framework to 2D shape animation using the FEM. The simulation parameters to be optimized are boundary conditions of a few points and the Poisson’s ratio of the material (for a single object, Young’s modulus is irrelevant). Fig. 1.8 shows an example of recovering the Poisson’s ratio given the initial and target shapes. These results were published in [Lee and Lin, 2012] and presented in Chapter 5.

1.6.4 Acceleration of Optimization Framework Using Reduced-Dimension Finite Element Modeling

Reduced-dimension finite element modeling [Krysl et al., 2001] has been used to reduce the computational complexity of the nonlinear FEM. A linear basis representing the nodal displacements is computed by training a statistical model using sample displacements generated by example simulations [Krysl et al., 2001] or user-provided keyframes [Barbič et al., 2009]. Each of the bases can be viewed as nodal displacements, and the displacements solved by the FEM is always a linear combination of the bases used. Fig. 1.9 shows the first two bases for the surface displacements of the bladder and the rectum.

I apply the idea to my FEM-based elasticity optimization, since the high dimen-




Figure 1.9: Visualization of the first two principal components (reduced bases) found by the statistical training of surface deformation.

sionality of external forces hurts the performance of the optimizer. In my case, the example displacements for the statistical training are computed by matching surfaces from segmented images. The main difference between my method and previous applications of the reduced dimension modeling is in the cross-model training: I train for a reduced set of basis elements to be used for multiple FE models. The cross-model training also means that example displacements generated using different models can all be used in the statistical training and thus the reduced basis captures the statistics of multiple patients. The reduced basis can also be applied to any new patient data without re-training.

In order to achieve the basis sharing, all patient-specific FE models must be constructed in a way that they share the same set of boundary nodes (which have force applied), and only the dimensions for these nodes are reduced. I use a set of *atlas surfaces* consisting of the boundary nodes to build each specific FE model: the atlas surfaces are iteratively deformed to fit the specific surfaces, and the FE model is built on the fitted surfaces. As a result, the nodes on the atlas surface is shared among all specific FE models, and the reduced basis can be shared among these FE models.

Although the simulator itself does not benefit a lot from this partial reduction scheme, the speed of the optimizer is improved by more than an order of magnitude due to the

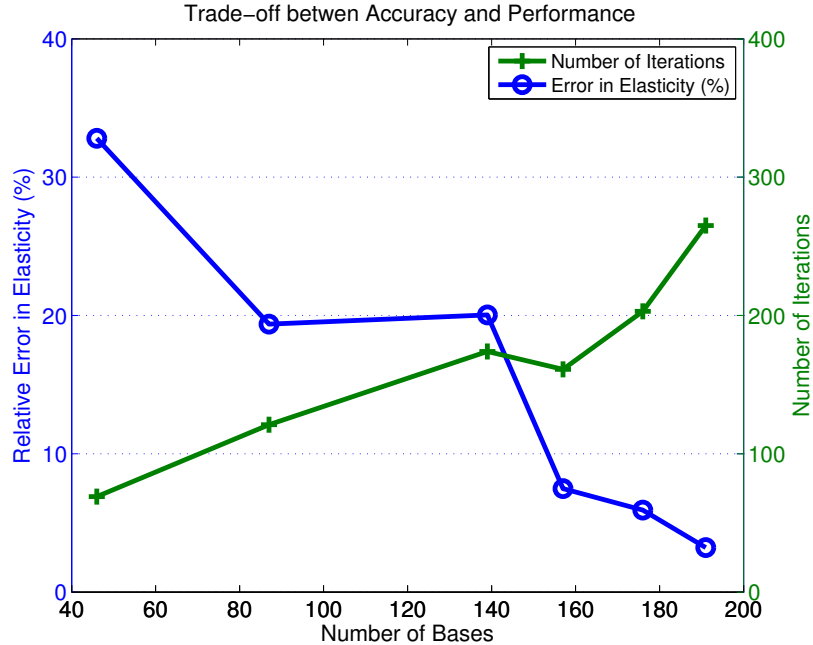


Figure 1.10: The trade-off between speed and accuracy is achieved by using a different number of bases for representing nodal displacements in the FE model.

reduced dimension of its parameter space (boundary forces). I present experiments on synthetic patient data to show the trade-off between the speed (number of basis vectors used) and accuracy in recovered elasticity, as shown in Fig. 1.10. These results were published in [Lee and Lin, 2012] and presented in Chapter 5.

1.7 Organization

The remainder of this dissertation is organized as follows.

Chapter 2 reviews previous work and basic concepts related to the thesis work, including medical image analysis and elasticity reconstruction.

Chapter 3 presents the core framework of the thesis for estimating deformation and elasticity parameters. The sensitivity of deformation to the elasticity is studied, and experiments on real patient data is presented. A clinical application complementary to traditional elastography is also presented, followed by examples in physically-based

medical image registration.

Chapter 4 discusses the problem of estimating treatment images for dose calculation when marker locations instead of CT images are used to localize the patient during treatment. A method based on fitted m-reps as well as the simulation-based approach is proposed and experimental results in dose calculation are presented.

Chapter 5 introduces the acceleration method for the optimizers using reduced-dimension FE modeling. An application of the optimization framework on physically-based 2D animation is also presented.

Finally, Chapter 6 concludes this thesis and discuss future research directions.

Chapter 2

Previous Work

In this chapter, I survey some related work on elasticity and deformation estimation in the fields of medical image analysis (elastography and image registration) and computer graphics (physically-based animation).

2.1 Elastography

Material properties are important in medical applications such as surgical simulation, motion compensation, and cancer detection. The most direct and accurate approach of material property estimation is measuring them *ex vivo*, where the tissue is extracted from the patient or the animal, and the deformation and forces are measured in a controlled environment [Krouskop et al., 1998; Zhang et al., 2008; Umut Ozcan et al., 2011]. In many cases, however, *ex vivo* experiments with the tissue are not feasible for live patients, and medical images have been used for noninvasive elasticity reconstruction.

Since many diseases are related to stiffer tissues, palpation has been an important means of diagnosis. However, the result of a palpation is subject to the person who performed the examination, and the application is limited to areas close to the skin. To overcome these limitations, a quantitative analysis of medical images can be used to evaluate tissue stiffness. For example, elastography (or elasticity reconstruction), is a noninvasive approach that makes use of displacements evaluated from medical images

(usually ultrasound images) and the measured forces.

2.1.1 Displacement Field Estimation

Based on how the external forces are applied, elastography methods can be roughly divided into two groups: one using quasi-static compression (also called *static elastography*), and the other using low-frequency vibrations (also called *dynamic elastography*). A common component in most methods of both groups is the use of a dense displacement field for estimation elasticity parameters. Once the displacement field is known, the elasticity can be computed using a direct or iterative optimization method.

2.1.1.1 Static Elastography

This type of methods estimate the dense displacement field using two images, one taken at the rest state, and the other taken when the tissue is compressed. Most of the methods use ultrasound images because the displacement field can be computed from the time delays between pre- and post-compression echo signals, assuming that speckle motion represents the underlying tissue movement for small, uniaxial compressions [Ophir et al., 1999]. Ultrasound elastography is usually only two dimensional, since the image resolution is highest on the plane perpendicular to the transducer face, and therefore it requires a sweeping of the transducer to produce a 3D elastogram [Lindop et al., 2006]. Another disadvantage is that the amount of compression is still limited to regions close to the skin, unless some probe is inserted into the body [Egorov et al., 2006].

2.1.1.2 Dynamic Elastography

Instead of static compressions, the external forces can be exerted through low-frequency vibrations. The advantage over the static approach is that it can penetrate into deeper tissues, although the effectiveness may be lowered for obese patients. An MRI or ultrasound machine in tune with the vibration can be used to find the displacement field

[Manduca et al., 2001; Salcudean et al., 2006]. When the target tissue is deep-seated, the vibration actuation mechanism may become sophisticated. For example, for an elastography of the prostate, the vibrator can be inserted into the rectum or the urethra [Chopra et al., 2009].

2.1.2 Inverse Modeling of Elastic Deformation

Once the displacement field is known, the elasticity can be computed by solving a least squares problem [Zhu et al., 2003; Becker and Teschner, 2007; Eskandari et al., 2011], if the algebraic equation resulting from the numerical solver is linear. Notice that the elasticity can be computed only if the force is also measured during the static compression or vibration actuation, and otherwise only the relative values can be found.

Alternatively, the error in the displacement field can be minimized by updating the elasticities iteratively [Kallel and Bertrand, 1996]. This type of methods use numerical optimization algorithms and are usually slower than direct methods. However, the optimization scheme does not depend on the linearity of the deformation model and is more general.

2.1.3 Other Elasticity Estimation Methods

The modality-independent elastography (MIE) [Washington and Miga, 2004] uses 3D medical images such as the CT or MR images and maximizes image similarity measures instead of finding the displacement field. The modality does not depend on the force exertion scheme, but a number of landmarks within the tissue is required. Therefore, the method is not suitable for organs with near-constant image intensities, such as the prostate.

Risholm et al. [2010] proposed a Bayesian image registration framework, where the physically-based registration problem is modeled with posterior probabilities, and a Markov Chain Monte Carlo sampling technique is used to characterize the posterior

distributions over deformation and elastic parameters. However, the high dimensionality requires a large number of samples to characterize the distribution and therefore limits the number of degrees of freedom of the problem.

In cardiac function estimation, sequential data assimilation [Sermesant et al., 2006; Moireau et al., 2008] has been applied to estimate the simulation parameters and predict the state (displacements) jointly. The parameters and states are modeled with probability density functions conditioned on the observed states, and a filtering procedure is applied over time to estimate the states. The accurate representation of the probability density function also suffers from high computational complexity, and therefore the dimension of the dynamic system is limited in practice.

2.2 Deformable Image Registration

The goal of an image registration is to find a voxel-to-voxel correspondence between two images (2- or 3-dimensional). Mathematically, an image can be defined as an N -dimensional real-valued function which maps a point \mathbf{x} in the N -dimensional space to the intensity value. The correspondence between two images I_f (fixed or *target* image) and I_m (moving or *source* image) can be represented as a function mapping a point to another point, and the goal of an image registration is to find the transformation T such that $I_f(\mathbf{x})$ is as close to $I_m(T(\mathbf{x}))$ as possible.

A simple example of such a transformation is a rigid transformation, which is essential during a radiation therapy in order to position the patient correctly with respect to the radiation beam, since the treatment table may be located differently on each treatment day. For most medical applications, however, a nonrigid registration is also desired since human organs are always deforming. In this thesis, I assume that a rigid registration is done as a preprocessing step [Foskey et al., 2005].

Image registration methods can be categorized based on the two main parts of the

algorithm: the form of the transformation T and the way of measuring image similarities of $I_f(\mathbf{x})$ and $I_m(T(\mathbf{x}))$. I discuss two image similarity measures in Sections 2.2.1 and 2.2.2. Section 2.2.3 introduces the type of methods using a physically-based simulation to generate the non-rigid transformation. See [Sotiras et al., 2012] for a survey on non-rigid image registration.

2.2.1 Image-Based Methods

Thirion [1998] proposed using a simple optical flow method with intensity difference as the similarity measure. The registration process is considered as an iterative process of moving each voxel in the reference image I_f , until the voxel intensities of the deformed reference image and moving image agree. Beginning with the initial I_f , in each iteration t , a voxel \mathbf{x} in I_f^t is moved in the direction of the image gradient $\nabla I_f^t(\mathbf{x})$ with the distance decided by the intensity difference

$$v^t(\mathbf{x}) = (I_m(\mathbf{x}) - I_f^t(\mathbf{x})) \frac{\nabla I_f^t(\mathbf{x})}{\|\nabla I_f^t(\mathbf{x})\|}. \quad (2.1)$$

Although $v^t(\mathbf{x})$ is computed by moving pixels in the reference image, the resulting non-rigid transformation is applied to the moving image implicitly. Notice that Eq. 2.1 is equivalent to a gradient descent optimization scheme using the sum of squared intensity differences as the objective function. Even with matched intensity, however, the resulting displacement field may not be smooth since each voxel is moved independently. To cope with the high dimensionality of the parameter space (displacement field), regularization terms such as image smoothness can be added in the process of minimizing the intensity difference. More sophisticated regularization terms based on compressible fluid [Christensen et al., 1996; Foskey et al., 2005] and linear elasticity model [Bajcsy et al., 1983; Holden, 2008] have also been used in the literature.

2.2.2 Landmark-Based Methods

A major drawback of image-based methods is that the displacement field for regions with nearly constant intensities (very low gradient) is inconclusive. It therefore can be desirable to make use of salient feature points (*landmarks*) in the images to compute displacements and interpolate the displacement field for the regions without features using some radial basis functions such as the thin plate spline [Rohr et al., 2001]. For organs with complex structures, landmarks may be found inside the organ, and a coarse-to-fine approach based on the number of landmarks can be used [Shen and Davatzikos, 2002]. For most organs, however, only smooth edges can be found, and the the landmarks could be vertices in the tessellation of surfaces of organs [Kaus et al., 2007], but the distribution of vertices may be different for the two surfaces. Another solution is to make use of sample points on the two m-reps [Pizer et al., 2003] from the two images. The sample points with the same *figural coordinates* on the two m-reps provides a natural correspondence and can be utilized for computing displacements [Levy et al., 2008].

2.2.3 Simulation-Based Methods

The displacement field can also be generated with a physically-based simulation to enforce the physical constraints exactly and take different material properties into account. The isotropic linear elasticity model with the finite element method has been a common simulation method for image registration [Ferrant et al., 1999, 2000; Liang and Yan, 2003; Hensel et al., 2007]. Boundary conditions are usually obtained from the gradient of the image [Ferrant et al., 1999] or from surface matching methods [Ferrant et al., 2000; Liang and Yan, 2003; Hensel et al., 2007], but they may induce uncertainty to the simulation since the surface matching is not unique. Material properties, on the other hand, are adjusted manually. Recently, an optimization-based 2D registration algorithm has been proposed [Alterovitz et al., 2006], where the boundary conditions and the material properties are optimized with a gradient descent method, but it has been applied

only to 2D images with a low-resolution triangular mesh. My method aims to solve the parameter estimation problem automatically using an FEM simulator combined within a numerical optimization method.

2.3 Parameter Estimation in Computer Graphics

Physically-based simulation is a popular method of generating realistic animations of natural phenomena. In order to achieve artistic control of these phenomena, artificial forces can be added into a simulated domain, or material properties can be adjusted to achieve a certain look and feel.

2.3.1 Directing Simulations

Recent advances in physically-based fluid [Stam, 1999; Fedkiw et al., 2001; Foster and Fedkiw, 2001; Enright et al., 2002] and deformable [Nealen et al., 2006; Teschner et al., 2005] modeling have greatly enhanced physical realism in computer animations. However, in many cases, artistic control is also desired to achieve a certain effect. For example, the artist may want to control the shape of the smoke while retaining the physical realism. Such an effect could be achieved by controlling the wind forces, but the process of adjusting parameters, simulating, and assessing the result is prohibitively tedious. Treuille et al. [2003] proposed a nonlinear optimization scheme to minimize the difference in density fields in order to achieve a keyframe-based control of the shape of smoke. The work is further improved with the adjoint method and applied to liquids as well [McNamara et al., 2004]. My method, on the other hand, solves an even harder problem because material properties are also taken into account. And since both the moving and reference models are acquired from real medical data, the accuracy requirement is higher than that in computer animations.

Another approach to address the direction of simulation is by adding physical details

to a rough animation. For example, wrinkles can be added to the skin of an animated character with physically-based simulations. Bergou et al. [2007] used constrained Lagrangian mechanics approach to simulate thin shells to add details to an animated object. Their work is different from mine in that I match simulated results rather than *adding* simulated details to an existing animation.

2.3.2 Estimating Material Properties

Estimation of material properties is also of interest for animation of deformable objects such as cloth and elastic bodies. An optimization scheme has been proposed to estimate cloth simulation parameters [Bhat et al., 2003]. The cloth model has stiffness and damping coefficients in an in-plane stretch term, an in-plane shear term, and an out-of-plane bending term, giving a total of six parameters [Baraff and Witkin, 1998]. The authors compared video of real fabric patches and simulated images to compute the error metric based on the orientation of each edge pixel, and the error is minimized with the continuous simulated annealing method [Press, 2007]. Syllebranque and Boivin [2008] used a similar optimization method with a force capture device, so that the boundary forces are known, to estimate the mechanical properties of deformable solids. They used video-based metrics to optimize for Poisson’s ratio and used the errors in computed boundary forces to optimize for Young’s modulus. A trinocular video system can be used to further improve the video tracking. Bickel et al. [2009] used such a 3D vision system with markers placed on the object, along with force measuring devices, to estimate the nonlinear behavior of the elastic object. Their system is further applied to computer-assisted material design [Bickel et al., 2010]. While these methods depend on the rendering and computer vision algorithms, my technique directly uses the surfaces of the deformed bodies to compute the error metric. In addition, these methods require a highly controlled setting for capturing the videos, which may not always be feasible.

Chapter 3

Simulation-Based Joint Estimation of Deformation and Elasticity Parameters

Estimation of tissue stiffness is an important means of noninvasive cancer detection. Existing elasticity reconstruction methods usually depend on a dense displacement field (inferred from ultrasound or MR images) and known external forces. Many imaging modalities, however, cannot provide details within an organ and therefore cannot provide such a displacement field. Furthermore, force exertion and measurement on the boundary of the target region can be difficult for some internal organs, making boundary forces another missing parameter. I propose a general method for estimating elasticity and boundary forces automatically using an iterative optimization framework, given the desired (target) output surface. During the optimization, the input model is deformed by the simulator, and an objective function based on the distance between the deformed surface and the target surface is minimized numerically. The optimization framework does not depend on a particular simulation method and is therefore suitable for different physical models. I show a positive correlation between clinical prostate cancer stage (a clinical measure of severity) and the recovered elasticity of the organ. Since the surface correspondence is established, my method also provides a non-rigid image registration, where the quality of the deformation fields is ensured, as they are computed using a physics-based simulation.

3.1 Introduction

Material property estimation has been an important topic in noninvasive cancer diagnosis, since cancerous tissues tend to be stiffer than normal tissues. Traditional physical examination methods, such as palpation, are limited to detecting lesions close to the skin, and reproducible measurements are hard to achieve. With the advance of medical imaging technologies, it becomes possible to quantitatively study the material properties using noninvasive procedures.

Computer vision methods in combination with force or pressure sensing devices have been proposed to find material properties of tissues [Kauer et al., 2002; Syllebranque and Boivin, 2008]. These methods require a controlled environment in order to capture the video and force (pressure), and therefore the experiments are usually done *ex vivo*. Kauer et al. [2002] combined the video and pressure capturing components into a single device to simplify the measurement process, so that it can be performed *in vivo* during a surgical intervention. However, the device still needs to be in direct contact with the tissue, and only a small portion of the tissue can be measured due to the size of the device.

Elasticity reconstruction, or elastography, is a noninvasive method for acquiring strain or stiffness images using known external forces and a known displacement field [Ophir et al., 1999; Manduca et al., 2001]. The reconstruction is usually formulated as an inverse problem of a physically-based simulation of elastic bodies, and a popular choice of the simulator is based on a linear elasticity model solved with the finite element method (FEM) [Zienkiewicz and Taylor, 2005], where the domain of the image is subdivided into tetrahedrons or hexahedrons called *elements*, with vertices known as *nodes*. Boundary conditions (displacement vectors or forces) on some of the nodes must be given to drive the simulation. Under this framework, nodal displacement vectors need to be computed based on a pair of images, and the force exertion mechanism needs to be controlled so that external forces can be measured. Otherwise, without measured forces,

only relative elasticity values can be recovered. Most existing elastography methods rely on pixel-wise correspondence or a dense set of image features, along with known boundary conditions, for reconstructing the elasticity. See Section 2.1 for a survey. While these methods are instrumental in their respective fields of interest, they are less well suited for a more general, multi-organ case where the image intensity may be almost constant within an organ, such as the prostate, and the lack of image details within the object makes it impossible to rely on pixel-wise correspondence. Moreover, the force exertion or vibration actuation mechanism can become complicated when the target tissues are deep inside the body.

I propose an entirely passive analysis of a pair of images that only uses information about the boundaries of corresponding internal objects. I assume the images have already been segmented, that is, the organ boundaries have been found. Since I do not assume a good correspondence for pixels inside an object, the resolution of the resulting elastogram is limited to the object boundaries. Namely, I assume that the elasticity is fixed within each object whose boundary can be identified. Natural movements inside the body provide the deformation of the organs, and I do not need an additional force exertion or vibration actuating mechanism. I minimize the distance between the deformed reference surface and the target surface and optimize for the elasticities and boundary forces. Currently, as a simplification, I consider only Young's modulus (which measures the stiffness or elasticity of the material). It is the simplest parameter to work with, and it is also important in noninvasive cancer detection techniques. The general optimization framework extends naturally to the inclusion of other parameters such as Poisson's ratio (which measures compressibility of the material), and in fact is suitable for a variety of physical models. In my experiments, the images are obtained from a prostate radiotherapy case, where there is one reference (planning) CT image and multiple target (daily) images for each patient, and the Young's moduli of the prostate recovered from the pairs of images are averaged. My initial investigation involving 10 patient data sets shows

that the recovered elasticity values positively correlate with the clinical tumor stages, which demonstrates its potential as a means of cancer stage assessment complementary to existing elastography methods. Furthermore, compared broadly to other work on simulation parameter estimation, my method does not require the inclusion of forces as part of the input and can therefore avoid the process of measuring the external forces (at the cost of only providing relative force information in my results).

My method also produces an image registration [Maintz and Viergever, 1998; Holden, 2008] (pixel-wise correspondence between images) since the distance between the pair of surfaces (segmentations) is minimized. The FEM has been applied to image registration, given that the images are segmented [Ferrant et al., 1999, 2000; Bharatha et al., 2001; Cash et al., 2005; Hensel et al., 2007; Wittek et al., 2007; Crouch et al., 2007]. Material properties, however, are not trivial to find from the images, and most authors use *ex vivo* experimental results to set up the materials. Moreover, due to the patient-to-patient differences, these material properties sometimes need hand adjustments. Alterovitz et al. [2006] incorporated an optimization of Young’s modulus and Poisson’s ratio into an FEM-based registration, but the method has only been implemented for coarse 2D meshes. As a non-rigid image registration method, mine improves over previous simulation-based methods by providing an automatic means of finding the parameters that are missing in the images. My current implementation uses both standard linear and nonlinear material models, but the optimization framework should be applicable to tissues with more advanced and complex physical models.

I explain the elastic model and the optimization scheme in Section 3.2, followed in Section 3.3 by experimental results using two synthetic scenes and 10 sets of real CT images to demonstrate the feasibility of my method. I conclude with a summary and discussion of future work.

3.2 Method

The idea of the algorithm is to optimize a function based on the separation between corresponding organ boundaries. In each iteration, the objective function is computed by first simulating and deforming the surface using the current set of parameters, and then computing surface distances. I consider only the elasticity value (Young’s modulus), with Poisson’s ratios being chosen according to previous work on simulation-based medical image registration [Hensel et al., 2007].

The inputs to the correspondence problem are two segmented images: a fixed (target) image with segmentation \mathbf{S}_f and a moving (source) image with segmentation \mathbf{S}_m . The bones are already aligned using a rigid registration method described in [Foskey et al., 2005]. Each segmentation is represented as a set of closed triangulated surfaces, one for each segmented object. I construct a tetrahedralization of the moving volume such that each face of \mathbf{S}_m is a face in the tetrahedralization, so that \mathbf{S}_m is characterized entirely by its set of nodes. My optimization framework is built on a physically-based simulator that generates deformation fields with n unknown parameters $\mathbf{x} = [x_1, \dots, x_n]^T$, and a numerical optimizer to minimize an objective function $\Phi(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by the deformation and surface matching metrics. During the optimization process, the physical model is refined in terms of more accurate parameters and converges to a model describing the deformation needed for the particular surface matching problem. Here I use the linear FEM to illustrate the optimization scheme, although the framework can also be incorporated with a nonlinear FEM. A flow chart of my algorithm is shown in Fig. 1.4 and will be explained in detail in this section.

3.2.1 Linear Elasticity Model and Finite Element Modeling

In the optimization loop, the displacement field $\mathbf{u} = [u, v, w]^T$ is always generated by a physically-based simulation, where the FEM is used to solve the constitutive equations

of the linear elasticity model. Assuming isotropic linear elasticity, I can write $\sigma = \mathbf{D}\varepsilon$, where σ is the stress vector induced by the *surface forces*, ε is the strain vector defined by the spatial derivatives of the displacement \mathbf{u} , and \mathbf{D} is a matrix defined by the material properties (assuming an isotropic material, the properties are Young's modulus E and Poisson's ratio ν). To solve the equations numerically, I approximate the derivatives of the deformation with the FEM, where the domain is subdivided into a finite set of elements, and each element consists of several nodes. Fig. 3.3a shows the finite element model used in one of my experiments, where four-node tetrahedral elements are used. The deformation field \mathbf{u}^{el} for any point \mathbf{p} within an element is approximated with a piecewise linear function $\hat{\mathbf{u}}^{el}(\mathbf{p}) = \sum_{j=1}^4 \mathbf{u}_j^{el} N_j^{el}(\mathbf{p})$, where \mathbf{u}_j^{el} is the deformation of the j -th node of the element, and $N_j^{el}(\mathbf{p})$ is the linear shape function that has value one at node j and is zero at all other nodes and outside of the element. After combining the approximated piecewise linear equation for each element, the resulting linear system is

$$\mathbf{K}\mathbf{u} = \mathbf{f}, \quad (3.1)$$

where \mathbf{K} is called the stiffness matrix, which depends on the material properties (Young's modulus and Poisson's ratio) and the geometry of the elements; \mathbf{f} is a vector of external forces. For a 3D domain with N_n nodes, \mathbf{K} is a $3N_n \times 3N_n$ matrix. Notice that since both \mathbf{K} and \mathbf{f} are unknown, they can be scaled by the same factor without changing the output deformation field. Therefore, unless I know the exact values of the forces, only the relative values of the material properties can be recovered.

To make the nodes deform, some boundary conditions need to be enforced, either by assigning displacement values or by assigning forces to some nodes. If all the surface nodes, including boundaries between two materials, are assigned displacement values, then the simulation is essentially an interpolation of the displacement field from surface matching results. This means that the elasticity values only affect internal nodes, for




Figure 3.1: A sliced view of the synthetic scene, which consists of two concentric spheres; the inner (red) and outer (green) regions have different stiffness values (blue triangles represent outer surface, which is considered part of the green region).

which I do not know the target positions. Therefore the elasticity cannot be recovered. Instead, I only assign boundary conditions to a part of the surface nodes, and other surface nodes without boundary conditions will be affected by the relative elasticities. For example, in a simulation of the male pelvis region, the bladder and the rectum are usually the organs that drive the deformation of the prostate, while the pelvic bone is considered static. An intuitive choice is to apply boundary conditions on boundary nodes of the bladder, the rectum, and the pelvic bone, and set all other entries in the force vector to zero (no external forces), as proposed in [Hensel et al., 2007].

3.2.2 Sensitivity Study

Since my method is based on the assumption that the deformed surface depends on both the elasticity and the external forces, I first conduct an experiment of forward simulations using different parameter values to see how sensitive the surface is to these parameters. The synthetic scene consists of two concentric spheres that form two regions, one inside the inner sphere, and the other between the two spheres, as shown in Fig. 3.1.

I fix the elasticity of the outer region and alter the elasticity of the inner sphere, as only the ratio of the two elasticity values matters. A force with a specified magnitude

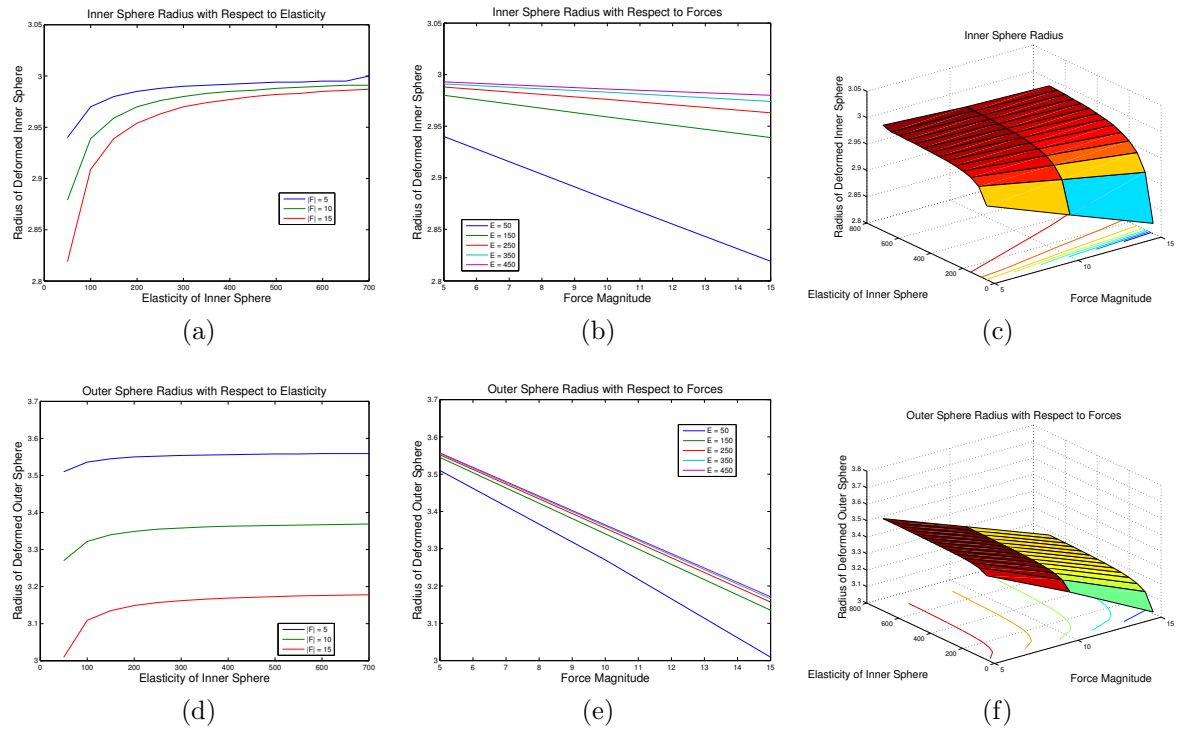


Figure 3.2: The plots of the radius of the inner sphere (in cm) after deformation: (a) inner radius versus elasticity value (in kPa) of the inner region; (b) inner radius versus magnitude of forces (in N) acting on the outer surface; (c) inner radius (z -coordinate) versus elasticity and force magnitude with isocontours of inner radius on xy -plane; (d) outer radius versus elasticity; (e) outer radius versus magnitude of forces; (f) outer radius (z -coordinate) versus elasticity and force magnitude with isocontours of outer radius on xy -plane. The radii before deformation are 3 cm and 3.75 cm for two spheres, respectively, and the elasticity for the outer region is 10 kPa. The Poisson's ratios are fixed to 0.40 and 0.35 for the two regions, respectively.

pointing towards the center of the spheres is applied on each node of the outer surface, and no external forces are applied on the inner surface. Several simulations using different elasticities of the inner region and force magnitudes were performed, and the plots of the sphere radius versus the elasticity value and versus force magnitude are shown in Fig. 3.2.

Notice that in these plots, the slope is much higher when the elasticity is low for each curve, which indicates that the shapes of both spheres are much more sensitive to the elasticity when the elasticity value is lower. These results suggest that my ability to recover the parameters is limited by how stiff the object is. When an object has a very

high stiffness, its shape becomes insensitive to the parameters. In this case, the shape can still be recovered, but the resulting parameters may not be accurate. Notice that the problem of solving for elasticity and for boundary forces is ill-posed with a single object. For example, drawing horizontal lines at some inner radius value in the plots in Fig. 3.2 would give multiple combinations of elasticity and external forces. However, when both the inner and outer surfaces are taken into account, the problem becomes well-posed: in the two-dimensional space formed by elasticity value and force magnitude, there is one curve that implies some radius of the inner sphere (an isocontour on the xy -plane in Fig. 3.2c) and another curve that results in some radius of the outer sphere (an isocontour on the xy -plane in Fig. 3.2f). The solution is at one of the intersections of the two curves, and I can eliminate unwanted solutions by limiting the range of elasticity and force magnitude according to experimental results on the specific materials.

3.2.3 Distance-Based Objective Function

The parameters needed in the simulator are $\mathbf{x} = [\mathbf{E}; \mathbf{F}]$, where \mathbf{E} consists of the material properties (in my case, the Young's moduli), and \mathbf{F} is the vector of external forces on boundary nodes. The objective function to be minimized is defined as the difference between the segmentations in the moving and target images,

$$\Phi(\mathbf{x}) = \frac{1}{2} \sum_{\mathbf{v}_l \in \mathbf{S}_m} \|\mathbf{d}(\mathbf{v}_l + \mathbf{u}_l(\mathbf{x}), \mathbf{S}_f)\|^2. \quad (3.2)$$

Here $\mathbf{u}(\mathbf{x})$ is the deformation field computed by the simulator with parameters \mathbf{x} , interpreted as a displacement vector for each surface node \mathbf{v}_l in the tetrahedralization. The notation $\mathbf{d}(\mathbf{v}, \mathbf{S})$ denotes the shortest distance vector from the surface \mathbf{S} to the node \mathbf{v} , and the sum is taken over all nodes of the moving surface.

The gradient of the objective function, which is needed in the iterative optimization,

is given by the chain rule,

$$\begin{aligned}\nabla\Phi(\mathbf{x}) &= \sum_{\mathbf{v}_l \in \mathbf{S}_m} \left[\frac{\partial \mathbf{u}_l}{\partial \mathbf{x}} \right] \left[\frac{\partial \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l, \mathbf{S}_f)}{\partial \mathbf{u}_l} \right] \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l, \mathbf{S}_f) \\ &= \sum_{\mathbf{v}_l \in \mathbf{S}_m} \mathbf{J}_u^T \mathbf{J}_d^T \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l, \mathbf{S}_f),\end{aligned}\tag{3.3}$$

where $\mathbf{J}_u = \left[\frac{\partial u_i}{\partial x_j} \right]$ is the Jacobian matrix of $\mathbf{u}(\mathbf{x})$ with respect to the parameters, and $\mathbf{J}_d = \left[\frac{\partial d_i}{\partial u_j} \right]$ is the Jacobian matrix of \mathbf{d} with respect to the deformation vector. Here I use the bracket $[\cdot]$ to represent a matrix and the curly braces $\{\cdot\}$ to denote a vector. Each column of \mathbf{J}_d , namely $\left\{ \frac{\partial \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l, \mathbf{S}_f)}{\partial u_j} \right\}$, is the derivative of $\mathbf{d}(\mathbf{v}_l + \mathbf{u}_l, \mathbf{S}_f)$ with respect to the j -th spatial coordinate ($j = 1, 2, 3$). The derivatives of \mathbf{u} with respect to the material properties are computed by differentiating both sides of (3.1),

$$\left[\frac{\partial \mathbf{K}}{\partial E_j} \right] \mathbf{u} + \mathbf{K} \left\{ \frac{\partial \mathbf{u}}{\partial E_j} \right\} = 0,\tag{3.4}$$

Therefore I have $\left\{ \frac{\partial \mathbf{u}}{\partial E_j} \right\} = -\mathbf{K}^{-1} \left[\frac{\partial \mathbf{K}}{\partial E_j} \right] \mathbf{u}$. The Jacobian matrix can then be computed by solving for each column of \mathbf{J}_u . The derivatives with respect to the boundary forces are computed in the same manner; by taking derivatives of both sides of (3.1), I have $\left[\frac{\partial \mathbf{K}}{\partial F_j} \right] \mathbf{u} + \mathbf{K} \left\{ \frac{\partial \mathbf{u}}{\partial F_j} \right\} = \mathbf{e}_j$, where \mathbf{e}_j is the j -th coordinate vector. On the right hand side, only the j -th entry is nonzero since $\frac{dF_i}{dF_j} = 0$ when $i \neq j$. And since \mathbf{K} is independent of F_j , $\frac{\partial \mathbf{K}}{\partial F_j} = 0$. Therefore I can solve for each column of the Jacobian with the equation $\mathbf{K} \left\{ \frac{\partial \mathbf{u}}{\partial F_j} \right\} = \mathbf{e}_j$. In practice, $\mathbf{d}(\mathbf{v}_l + \mathbf{u}_l(\mathbf{x}), \mathbf{S}_f)$ can be looked up in the precomputed vector distance map of the fixed organ, \mathbf{S}_f , and the derivatives $\partial \mathbf{d} / \partial u_j$ can be approximated with a centered finite difference operator applied on the map. Fig. 3.3b shows one of the distance maps used in my experiments. Notice that the physical model can be different, as long as the derivatives $\partial u_i / \partial x_j$ can be computed.

In my experiments, however, I observed that the magnitudes of gradients with respect to the material properties, $\|\partial \Phi / \partial \mathbf{E}\|$, are about 1000 times smaller than that with respect




Figure 3.3: Input to my algorithm: (a) a sliced view of the tetrahedral model of the moving image (light-blue triangles represent surfaces, not FEM regions; bladder and rectum are hollow); (b) a slice of the distance map of the prostate surface in the reference image.

to the forces, $\|\partial\Phi/\partial\mathbf{F}\|$, which caused the material properties to converge very slowly. To obtain a faster convergence of \mathbf{E} , I embed the optimization of the forces into the objective function evaluation at each \mathbf{E} value. That is, every time $\Phi(\mathbf{E})$ is evaluated, a full optimization of \mathbf{F} is performed with the fixed value of \mathbf{E} .

3.2.4 Numerical Optimization

I use a line search scheme for optimization: in each iteration k , I find a descent direction \mathbf{p}_k , find an optimal step size α in that direction with a line search algorithm, and then update the parameters with $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha\mathbf{p}_k$. The descent direction can be computed by using Newton's method to solve the equation $\nabla\Phi = 0$: $\mathbf{p}_k = -\mathbf{B}_k^{-1}\nabla\Phi(\mathbf{x}_k)$, where \mathbf{B} is the Hessian matrix, $\left[\frac{\partial^2\Phi}{\partial x_i\partial x_j}\right]$. A modified Newton's method has been used in elasticity reconstruction [Kallel and Bertrand, 1996], but the Hessian matrices can only be approximated and are usually ill-conditioned. Alternatively, I can use a Quasi-Newton method such as the BFGS formula to avoid computing the Hessian [Nocedal and Wright, 1999].

Quasi-Newton methods can reduce the computation yet still retain a super-linear

convergence rate. A line search enforcing the curvature condition ($\mathbf{s}_k^T \mathbf{y}_k > 0$) needs to be performed to keep the approximate Hessian positive definite. In my case, the number of parameters can be in the thousands, and therefore I adopt a limited-memory quasi-Newton method known as the L-BFGS method [Nocedal and Wright, 1999].

3.2.5 Initial Guess of Parameters

A good initial guess can prevent the optimizer from getting stuck in a local minimum. My initial guess for the forces is based on the distance field of the target surface: each node requiring a boundary condition is moved according to the distance field to compute a Dirichlet boundary condition. A forward simulation is performed using the set of boundary conditions and the initial guess of elasticities, and the output deformation is used, via (3.1), to compute the corresponding forces, which become the initial guess for the forces.

In the case of medical image registration, the initial guess of the elasticity is chosen based on knowledge of the simulated organs. My example images involve two materials: the prostate and the surrounding tissue. There have been *ex vivo* experiments on the prostate using different elasticity models. Krouskop et al. [Krouskop et al., 1998] reported an elastic modulus of 40-80 kPa for normal prostate tissue, 28-52 kPa for BPH tissue, and 80-260 kPa for cancerous tissue when receiving 4% compression. They also reported 10-30 kPa for breast fat tissue. Based on these numbers for fat tissue, I chose an elasticity value of 10 kPa for the tissue surrounding the prostate. This value is fixed, in my calculations, since only the ratio of the elasticity values matters.

The initial guess of elasticity for the prostate is chosen by a parameter search: I perform force optimizations with several elasticity values between 30 kPa and 200 kPa and choose the elasticity with the lowest objective function value after the force optimization. An example result of the parameter search is shown in Fig. 3.4, where the target surfaces are generated by artificially deforming a set of organ boundaries, so that I know




Figure 3.4: Plot of Φ and \mathbf{E} (in kPa) with several sample values for finding an initial guess of elasticity value in a synthetic multi-organ scene. The plot suggests that the best initial guess is 50 kPa.

the true elasticity value. The plot shows that the parameter search successfully located the global minimum in the synthetic case with multiple organs. In my experiments using synthetic and real organ boundaries of the male pelvis region, I have observed similar curves with a single minimum. If more than one local minimum is observed, an optimization can be performed using each of these values as the initial guess. To reduce the computation time, I use a lower-resolution mesh for the parameter search, and the resulting optimal forces are used as the initial guess when using a higher-resolution mesh for elasticity optimization.

3.3 Experiments

I used the male pelvis region as the test scene. To build the reference surfaces, I obtained segmentations of a 3D CT image of the male pelvis region, including the surfaces of the

bladder, prostate, rectum, and bones. A tetrahedral finite element mesh is constructed from a set of reference surfaces, as shown in Fig. 3.3a. The corresponding target surfaces are used to compute the distance map, as shown in Fig. 3.3b. In the tetrahedral mesh, the bladder and the rectum are made hollow to reflect the actual structure, and the bones are fixed during the simulations. Since the prostate is the main organ of interest, I apply forces only on the boundaries of the bladder and the rectum to reduce the uncertainty on the prostate, which will be moved by surrounding tissues. The setting also reflects the fact that the bladder and the rectum are the organs that have larger deformations due to different amount of fluid and gas, and the prostate is usually deformed by their movement.

During the iterative optimization, the objective function is evaluated over the surfaces of the bladder, rectum, and prostate. The Poisson’s ratios are fixed (0.40 for the prostate and 0.35 for surrounding tissues, chosen based on literature in image registration [Bharatha et al., 2001; Tanner et al., 2006; Hensel et al., 2007]), and I optimize for the elasticity values because of its importance in noninvasive cancer detection. Since only the relative values of material properties can be recovered, I fix the Young’s modulus of the surrounding tissues (the region outside all organs and bones) to 10 kPa and optimize that of the prostate.

I tested my algorithm on two types of surface data. First, I tested the accuracy of the optimization scheme using synthetic target surfaces generated by forward simulations, so that I know the true elasticity values. I then applied the technique to prostate cancer stage assessment based on multiple segmented target images of the same patient to show applicability to real data. Since the distances between reference and target surfaces are minimized, I also compare the visual result (the warped image) with that of an image-based image registration method.

The reference and target organ surfaces are obtained from real 3D CT images of the male pelvis region using the software MxAnatomy (Morphormics, Durham, NC), and

the bones are segmented using ITK-SNAP [Yushkevich et al., 2006]. Given the moving surfaces in the form of triangle meshes, the tetrahedral model for the entire domain is built with the software TetGen [Si, 2009], and the library ITK [Yoo et al., 2002] is used to compute the vector distance maps of the target surface. The FEM simulator uses the linear algebra library PETSc [Balay et al., 1997].

Mesh generation The image segmentation was done with an early semi-automatic version of software MxAnatomy. For the prostate, the user typically needed to specify 15-20 initial boundary points on five image slices, and it usually took 20 minutes to segment the three main organs (prostate, bladder, and rectum) in a CT image. The semi-automatic segmentation of bones (ITK-SNAP) requires some initial pixels (specified with a few spheres) that are roughly in line with the bones, and the algorithm iteratively grows or shrinks from these initial pixels until an optimal binary image of the bone is achieved. It usually takes 15 minutes for the segmentation of bones. Once the surface meshes are generated, the tetrahedralization takes a few seconds using the software TetGen.

3.3.1 Synthetic Scene with Multiple Organs

To test how well my algorithm recovers elasticity values, I use synthetic target surfaces generated with known elasticity and boundary conditions. The target surfaces are generated by a forward simulation with Dirichlet boundary conditions acquired from a real pair of segmented images applied to the bladder and rectum surfaces. The moving surfaces and boundary conditions in the two synthetic scenes are shown in Fig. 3.5, where the boundary conditions are shown with scaled 3D arrows. The elasticity value of the prostate is controlled, and I can therefore compare the value recovered by my method to the ground truth. I tested my algorithm with three elasticity values, and the results are shown in Table 3.1. The optimization process is terminated when $\|\partial\Phi/\partial\mathbf{E}\| < 10^{-7} \|\mathbf{E}\|$ and $\|\partial\Phi/\partial\mathbf{F}\| < 10^{-4} \|\mathbf{F}\|$, or when the optimizer cannot find a direction in the parame-




Figure 3.5: The moving surfaces and ground-truth boundary conditions in the two synthetic multi-organ scenes: the arrows shows Dirichlet boundary conditions applied to surfaces of bladder and rectum; the scaling of arrows are according to the magnitude of displacements.

Table 3.1: Error in recovered modulus of elasticity in two synthetic multi-organ scenes; note that the error becomes much larger for elasticity values greater than 150 kPa.

True Elasticity (kPa)		50	100	150	200	250	300	350
Scene 1	Recovered Value	49	101.18	158.79	141.57	136.65	204.45	176
	Relative Error	2%	1.2%	5.9%	29.2%	45.3%	31.9%	49.7%
Scene 2	Recovered Value	51.33	102.90	167.5	225.0	222.91	275.0	277.97
	Relative Error	2.7%	2.9%	11.7%	12.5%	10.8%	8.3%	20.6%

ter space that reduces the value of the objective function. The relative error is less than 12% in the cases where the elasticity values do not exceed 150 kPa, which corresponds to an elasticity ratio of 15 between the prostate and the surrounding tissue. Notice that according to the literature [Krouskop et al., 1998], the ratio is already beyond the range for normal tissues and is within the range for cancerous tissues. Therefore I expect to see worse accuracy in the case of stiffer cancerous tissues.

Effect of inaccurate Poisson’s ratios In order to show the effect of selecting different Poisson’s ratios, I repeat the experiments using synthetic target surfaces generated with five different Poisson’s ratios for the prostate (0.3, 0.35, 0.4, 0.45, and 0.49), while the assumed value is fixed to 0.4 during the optimization process. (Most previous work on image registration or elastography assumes values between 0.3 and 0.49). As shown

in Fig. 3.6, the relative errors in recovered elasticity increase with larger deviation of the Poisson’s ratio, and the effect is especially prominent in the cases with lower elasticity values (soft) and low Poisson’s Ratios (compressible). I observe errors of 45–60% with an elasticity of 50 kPa and a Poisson’s ratio of 0.3. The errors are generally below 13.3% in cases with Poisson’s ratios of 0.40–0.45 and can be as high as 20% for a ratio of 0.49 (nearly incompressible). These results show that my method is robust to inaccurate Poisson’s ratios in most cases.

3.3.2 Noninvasive Assessment of Prostate Cancer Stage

To show the effectiveness of my method applied to prostate cancer assessment, I repeated the experiments on the multi-organ settings, but with both the deformed and target surfaces taken from segmented 3D CT images of the male pelvis region. I consider 10 patient data sets (a total of 112 target images) taken throughout courses of radiotherapy for prostate cancer. Each patient data set consists of a set of reference surfaces (bladder, prostate, rectum, and bones), which is from the CT image (reference image) taken before the radiotherapy, and multiple sets of target surfaces, each of them representing the internal structures in one daily CT image during the therapy. The reference image is taken about a week before the first treatment, and treatment (target) images are typically taken twice a week. For each patient data set, I repeated the process of deforming the reference surfaces toward a set of target surfaces with my method, so that one elasticity value of the prostate is recovered for each daily image. Fig. 3.7 shows the histogram of surface-surface distance between pairs of reference and target prostate surfaces used in the experiments. The surface-surface distance is defined as the maximum of node-surface distance,

$$\max_{\mathbf{v}_l \in \mathbf{S}_m} \mathbf{d}(\mathbf{v}_l, \mathbf{S}_f), \quad (3.5)$$

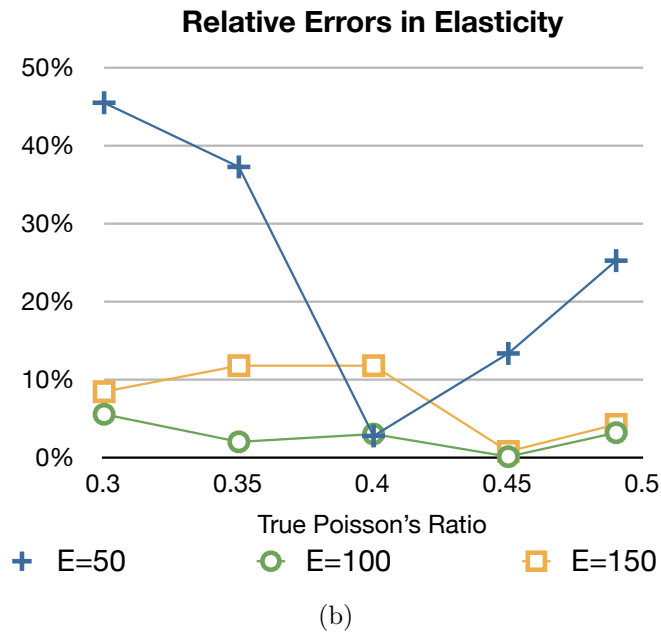
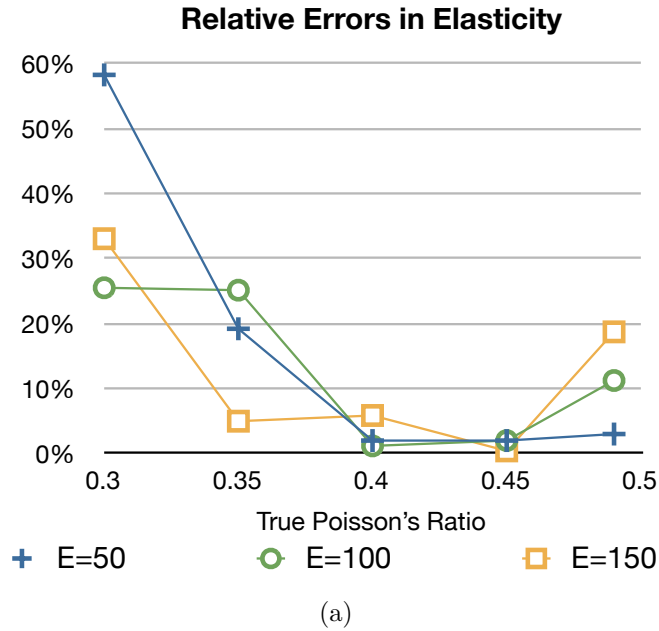


Figure 3.6: Plots of relative errors in recovered elasticity vs. different Poisson's ratios for the prostate (0.3, 0.35, 0.4, 0.45, and 0.49) used for generating the synthetic surface data; each plot shows the result from one test scene, and each curve represents a true elasticity value (50, 100, and 150 kPa) used in the synthetic case. During the optimization process, the assumed Poisson's ratio is always fixed to 0.4.

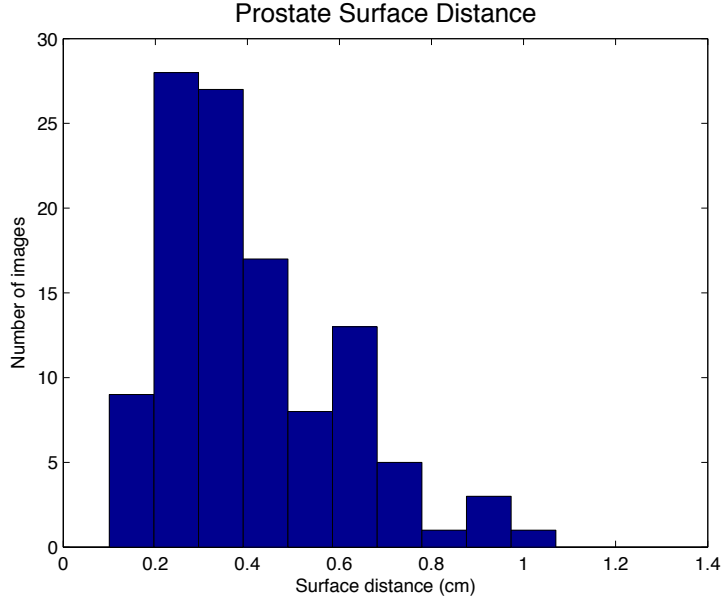


Figure 3.7: Histogram of distances between the pairs of corresponding reference and target prostate meshes used in my experiments on segmented CT images.

where \mathbf{S}_m is the reference (moving) prostate surface, and \mathbf{S}_f is the target surface. The average surface distance for the prostate among the 112 pairs of images is 0.41 cm, which is less than 10% of the diameter of a typical prostate (around 4–5 cm).

The convergence graphs (plots of Φ and $\|\nabla\Phi\|$ versus iteration number) and for boundary forces and for the material property from a typical image pair are shown in Fig. 3.8 (convergence graphs for other experiments are similar). Note that the optimization of forces was done in batches (in each evaluation of $\Phi(\mathbf{E})$), and the convergence graph for force optimization is the result of concatenating the steps for optimizing \mathbf{F} . With my current code, each iteration for the force optimizer takes about 19 seconds for a mesh with 34,705 tetrahedral elements and 6,119 nodes on a Xeon X3440 CPU, and the total number of iterations is around 1,700 (the total time is about nine hours), which means that my current implementation is only suitable for off-line processes. Note that I have not utilized any parallelism in the FEM computation. In the future, I plan to explore faster implementations of the FEM, such as those utilizing a many-core processor and reduced-dimension models.

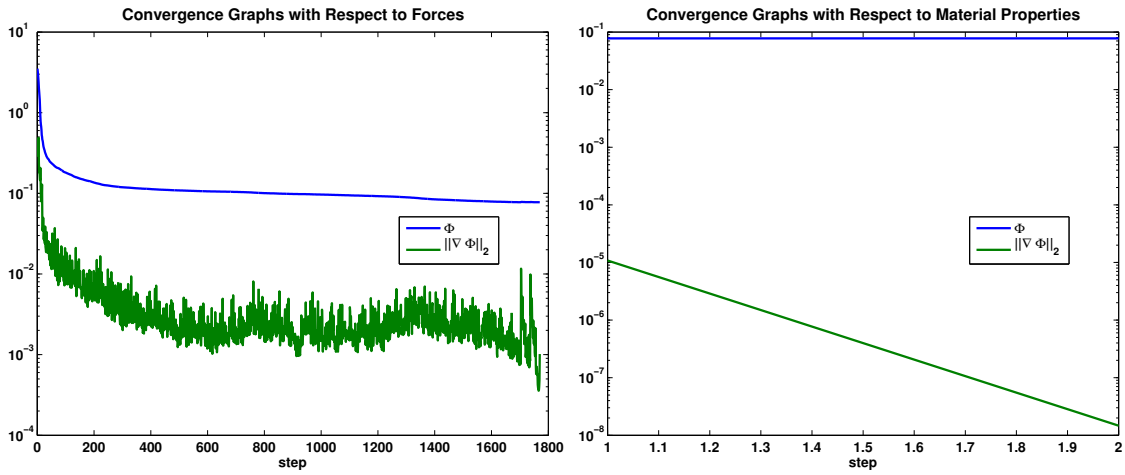


Figure 3.8: Convergence graphs (plot of Φ and $\|\nabla\Phi\|$ versus iteration number) for a pair of CT image data: (left) convergence of the external forces; (right) convergence of the elasticity.

Table 3.2: Average and standard deviation of elasticity values for the prostate recovered from the patient data sets; the last column is the clinical cancer staging for the tumor for each patient.

	Number of Targets	Average Young's Modulus (kPa)	Std. Deviation	Clinical T-Stage
Patient 1	8	48.60	2.41	T1
Patient 2	6	53.99	10.28	T3
Patient 3	7	71.97	4.35	T3
Patient 4	6	60.81	1.25	T2
Patient 5	16	38.06	13.91	T1
Patient 6	16	45.42	10.26	T1
Patient 7	17	40.67	16.34	T2
Patient 8	15	52.40	7.72	T2
Patient 9	9	51.47	7.50	T1
Patient 10	12	56.19	7.95	T2

Table 3.3: Definition of clinical T-stages for prostate cancer

Stage	Definition
TX	Primary tumor cannot be assessed
T0	No evidence of primary tumor
T1	Clinically inapparent tumor neither palpable nor visible by imaging
T2	Tumor confined within prostate
T3	Tumor extends through the prostate capsule
T4	Tumor is fixed or invades structures other than seminal vesicles

Each of the 10 patient data sets tested include 6 to 17 sets of target surfaces (daily images), namely 112 target images in total, and the recovered elasticity values of the prostate for each patient are shown in Table 3.2. Notice that the recovered values from all image pairs are within the range suggested in the literature [Krouskop et al., 1998], and the result shows consistency within each patient.

The aim of this study is to assess the relation between the recovered elasticity value and the cancer stage of each patient, under the assumption that prostates with more advanced tumors have higher stiffness. A common cancer staging system is the TNM (Tumor, lymph Nodes, Metastasis) system, where the clinical T-stage describes the size and extent of the primary tumor [Sobin, 2009]. The definitions of T-stages are shown in Table 3.3. I focus on the T-stage because of its relevance to the stiffness of the prostate. The clinical T-stages for the patients are shown in the last column of Table 3.2. In order to analyze the data statistically, I treat the average recovered elasticities and tumor stages as two random variables and use numbers 1, 2, and 3 to represent T-stages T1, T2, and T3, respectively (T0 and T4 are not present in my data sets), and I test if the recovered elasticity values and the T-stages are positively correlated. The resulting Pearson (linear) correlation coefficient is 0.662, and the p-value for the two-sided correlation test is 0.037, which indicates a significant positive correlation between the recovered elasticity values and the T-stages, based on a p-value threshold of 0.05. Since the tumor stage values are discrete and might be nonlinear with respect to the elasticity, a rank correlation coefficient, such as the Spearman’s rank correlation ρ , may

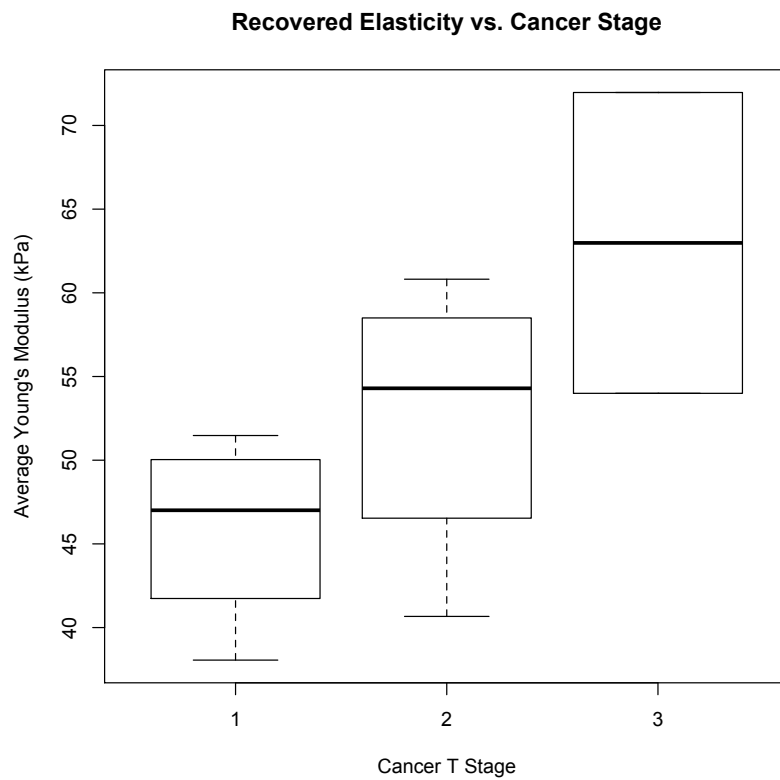


Figure 3.9: Box plot of average recovered elasticity value and cancer T stage for each patient data set shown in Table 3.2.

Table 3.4: The recovered elasticity values for the prostate as a homogeneous material, when the organ contains a synthetic tumor of different sizes and a normal tissue; elasticity values are set to 100 kPa for the tumor and 50 kPa for normal prostate tissue.

	tumor size / prostate size (%)			
	10%	25%	50%	75%
Scene 1	51.24	54.98	62.15	63.44
Scene 2	53.55	56.90	69.62	70

be more suited for the test. From the samples I have Spearman’s $\rho = 0.701$ and an estimated p-value of 0.024, which shows again a significant positive correlation. The box plot of the elasticity values and cancer stages is shown in Fig. 3.9.

3.3.3 Study: Inhomogeneous Materials

I assume a constant material property within an organ due to the limitation of the image modality, where the intensity is almost constant within the prostate, so that it is impossible to segment the tumor. The elasticity values recovered by my method are therefore “average” values in some sense, and a higher recovered elasticity indicates either a stiffer tumor or a larger tumor. Since the clinical T-stage for prostate cancer depends on the extent of the tumor, I conducted a study to show the correlation between the tumor size and the recovered elasticity value. Based on the settings in the synthetic multi-organ experiments in Section 3.3.1, I embedded an additional tumor inside the prostate for generating synthetic target surfaces, as shown in Fig. 3.10. The elasticity values for the tumor and the normal prostate tissues are set to 100 and 50 kPa, respectively. Notice that in the elasticity recovery process, I do not know the extent of the tumor due to the imaging limitation, and I only recover one value for the prostate. Table 3.4 shows the recovered elasticity values with different tumor sizes relative to the entire prostate. The results show increasing elasticity values with increasing tumor sizes in both scenes. Even though I assume homogeneous materials, the recovered values can still be used as an indicator of the extent of the tumor and are therefore correlated to cancer stages.




Figure 3.10: A sliced view of the tetrahedral mesh with a tumor (yellow) embedded in the prostate (red); the mesh is used to generate the synthetic target surface, while the prostate is still considered homogeneous in the optimization process.

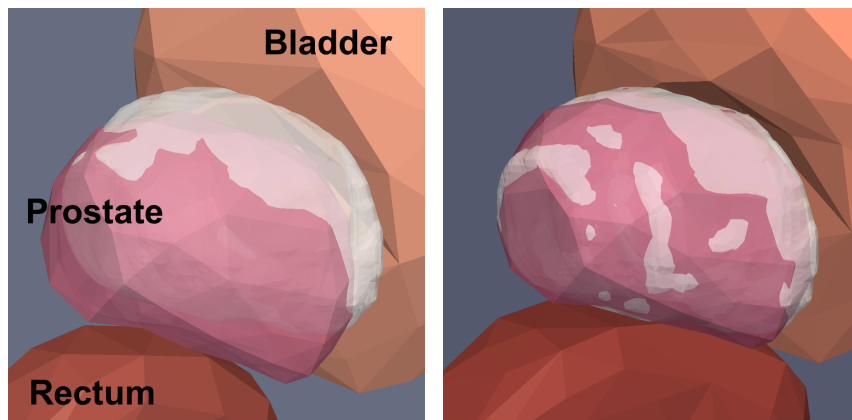




Figure 3.11: Close-up view of the surfaces before (left) and after (right) deformation; the transparent white surface shown is the target surface of the prostate. Notice how the prostate surface move towards the white surface. Bladder and rectum surfaces are those with external forces applied.



(a)



(b)

Figure 3.12: Registration results for a pair of test images: (a) axial and sagittal views of the moving image, and a 4x4 checkerboard comparison with the planning image, before registration; (b) the two views of the registered image, along with a checkerboard comparison with the planning image; superimposed by segmentations of the reference image, shown in red, and the segmentation of the prostate in the daily image, shown in blue; notice that the image deforms towards the red contours.

3.3.4 Application: Registration of Segmented CT Images

Since the distance between the fixed and moving surfaces is minimized during the optimization process, I also have an image registration as a result of optimizing for forces and elasticities. In my experiments, the final average value of the objective function is 0.09, corresponding to an RMS error of 0.01 cm, and a maximum of 0.22 cm, which are within the image resolution, $0.1 \times 0.1 \times 0.3$ cm. The deformed images of a typical image pair before and after registration are shown in Fig. 3.12, with the segmentations of the reference image (red) and the prostate in the moving image (blue) superimposed on the image. Notice how the prostate in the images moves from the blue contour to the red contour. Fig. 3.11 shows a 3D close-up view of the deforming surfaces from another image pair, where the surfaces of the bladder and the rectum are those with external forces applied, and the target surface of the prostate is shown in white.

I also compared my registration results with a popular image-based approach, the Demons method [Thirion, 1998], by looking at some landmarks inside the prostate. In most cases, the image intensity is almost constant inside an organ, but five of the patient data sets (a total of 32 image pairs) I experimented on have three “seeds” implanted in the prostate for location tracking during each treatment fraction, resulting in bright spots that can be observed in the CT image. The distance between the target and the deformed landmark positions from the two methods are shown in Table 3.5, and the two-tail t-tests for paired samples (distances) show that my method produces statistically significantly better results in three out of five patient data sets, with a p-value threshold of 0.05 (if Bonferroni correction is used, my method performs significantly better in only one data set). For regions with nearly uniform intensity, the deformation computed by the Demons method is entirely governed by the registration regularization terms, which do not need to be physically meaningful for the image-based method. my method enforces physically-based constraints and results in errors within the diameter of the spot. Notice that for the Demons method, I replaced the voxel values inside the prostate

Table 3.5: Average error in landmark positions (distance in cm) inside the prostate, computed with the Demons method and my method; t-tests show that my method performs statistically significantly better in three of the five data sets. (If Bonferroni correction is used, my method is significantly better only for Patient 3.)

	# Target Images	Demons		My Method		Paired t-test p-value
		Avg. (cm)	Std. dev.	Avg. (cm)	Std. dev.	
Patient 1	8	0.27	0.17	0.26	0.18	0.07
Patient 2	6	0.21	0.11	0.16	0.11	0.03
Patient 3	7	0.18	0.06	0.10	0.04	1.5e-4
Patient 4	6	0.17	0.07	0.13	0.03	0.02
Patient 5	5	0.21	0.15	0.20	0.08	0.86

with the average intensity within the organ, since the intensity and gradient information from the landmarks could also be utilized in the image-based registration, giving it an additional advantage, while my method is based purely on the physics-based simulation and does not take advantage of the landmarks.

3.3.5 Extension: Nonlinear FEM

To demonstrate that my optimization framework can also be applicable to nonlinear models, I incorporated a geometrically nonlinear FEM and the neo-Hookean model with the elasticity optimization scheme. The linearized equilibrium formulation of the nonlinear FEM is

$$\mathbf{K}(\mathbf{u}^n) \cdot \Delta \mathbf{u} = \mathbf{f}^{\text{ext}} - \mathbf{f}^{\text{int}}, \quad (3.6)$$

where \mathbf{f}^{ext} and \mathbf{f}^{int} are the external and internal force vectors, $\mathbf{K}(\mathbf{u}^n)$ is the stiffness matrix that depends on the current displacement vector \mathbf{u}^n , and $\Delta \mathbf{u}$ is used to update the vector \mathbf{u}^n in a Newton iteration ($\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta \mathbf{u}$). The Jacobian matrix $\mathbf{J}_u = \left[\frac{\partial u_i}{\partial E_j} \right]$ (derivative of displacements \mathbf{u} with respect to the elasticity parameter E_j) for the elasticity optimizer is approximated using the finite difference method due to the complexity of differentiating the internal forces with respect to the elasticity. Notice that I have not implemented force optimization for the nonlinear model, and boundary

Table 3.6: Error in recovered nonlinear modulus of elasticity in two synthetic multi-organ scenes.

True Elasticity (kPa)		50	100	150	200
Scene	Recovered Value	50.18	105.64	159	174
1	Relative Error	+0.35%	+5.6%	+6%	-13%
Scene	Recovered Value	45	105	141.5	197
2	Relative Error	-10%	+5%	-5.67%	-1.5%

conditions given by a surface matching is always used in the simulation.

3.3.5.1 Synthetic scene with multiple organs

I used the same multi-organ scenes in Section 3.3.1 and deformed them using the nonlinear FEM to generate the synthetic target surfaces. That is, the nonlinear FEM is used in both generating synthetic cases and in the optimization scheme. The resulting recovered elasticity values are shown in Table 3.6. The errors are within 13% for the range I tested (50-200 kPa).

Effect of inaccurate elasticity values for surrounding tissues The elasticity value for tissue surrounding the prostate is fixed to 10 kPa in my experiments. While only the ratio between two elasticity values can be recovered with a linear model without knowing true force values (as discussed in Section 3.2.1), the surrounding tissue elasticity could have a different effect on nonlinear models. However, with the small amount of displacement I have observed, I expect the surrounding tissue elasticity to have a similar effect as in the linear model. For example, if the true elasticity values for the prostate and surrounding tissue are 100 kPa and 20 kPa, respectively, I expect to recover the value 50 kPa for the prostate since the surrounding tissue elasticity is fixed to 10 kPa in the optimization. Table 3.7 shows the results using the nonlinear FEM where the true elasticity is twice the value used in the optimization process. The recovered elasticities for the prostate are very close to what I expect, with relative errors below 13%.

Table 3.7: Error in recovered nonlinear modulus of elasticity in two synthetic multi-organ scenes where the elasticity of surrounding tissue is doubled (20 kPa) when generating the synthetic data. The surrounding tissue elasticity is still set to 10 kPa in the optimization process, and I expect to see recovered values for the prostate to be half of the true values.

Elasticity for Surroundings		20 kPa			
Elasticity for Prostate (kPa)		50	100	150	200
Expected Value for Prostate		25	50	75	100
Scene 1	Recovered Value	28.05	50.17	72.98	106.17
	Relative Error	+12.2%	+0.34%	-2.7%	+6.17%
Scene 2	Recovered Value	25	45	76.52	108.16
	Relative Error	0%	-10%	+2.02%	+8.16%

3.3.5.2 Assessment of prostate cancer stage

I repeated the experiments in Section 3.3.2 using the nonlinear FEM. The recovered elasticity values for the 10 patient data sets are shown in Table 3.8, and the box plot of average recovered elasticity and clinical T-stage is shown in Fig. 3.13. The Pearson (linear) correlation coefficient for recovered elasticity values and T-stages is 0.704 with a p-value of 0.023, and the Spearman’s rank correlation ρ is 0.636 with a p-value of 0.048, which again shows a significant positive correlation between the stiffness value and the cancer stage for this group of patients. However, the recovered values are less consistent than those from the linear FEM implementation. I conjecture that the implementation using nonlinear FEM is more sensitive to the material properties and boundary conditions, and therefore the recovered values vary more than those using the linear FEM.

3.4 Conclusion and Future Work

I have presented a novel physically-based method for simultaneously estimating the 3D deformation of soft bodies and determining the unknown material properties and boundary conditions. Previous elastography methods are limited by imaging modalities and

Table 3.8: Average and standard deviation of elasticity values for the prostate recovered from the patient data sets using nonlinear FEM; the last column is the clinical cancer staging for the tumor for each patient.

	Average Young's Modulus (kPa)	Std. Deviation	Clinical T-Stage
Patient 1	47.29	3.25	T1
Patient 2	69.28	8.09	T3
Patient 3	78.91	4.81	T3
Patient 4	63.62	2.92	T2
Patient 5	47.45	16.62	T1
Patient 6	59.85	18.37	T1
Patient 7	62.73	18.34	T2
Patient 8	60.23	11.93	T2
Patient 9	69.74	11.46	T1
Patient 10	69.25	17.64	T2




Figure 3.13: Box plot of average recovered nonlinear elasticity value and cancer T stage for each patient data set shown in Table 3.8.

force measurement schemes, and I overcome these limitations by utilizing the surface information extracted from 3D images. Although the resolution of the resulting elastogram is limited to the object boundaries, I showed that the recovered value reflects the distribution of materials within the object, and the recovered elasticity values have a significant positive correlation with clinical prostate cancer staging in small-scale experiments. Therefore, my method has the potential to become a means of noninvasive cancer detection.

As a non-rigid image registration method, mine automatically determines the patient-specific material properties during the registration. The resulting deformation field is enforced to be physically plausible, since it is computed by the 3D FEM simulator with appropriate contact constraints among organs. The observed error on the boundary is within the resolution of the segmented images, and the error on the internal bright spots as landmarks in the prostate is comparable to the diameter of the spots.

The optimization framework for joint estimation of both 3D deformation and material parameters is generalizable. It is not limited to elasticity reconstruction and could be used for more sophisticated physiological models than the basic linear and nonlinear elasticity models I chose for simplicity in my current implementation. As an image registration technique, my method is reliable in terms of the registration error; as a parameter estimation method, my system can save an enormous amount of efforts adjusting the simulation parameters manually by automatically extracting patient-specific tissue properties. Furthermore, since only the 3D surfaces are used in my algorithm, applications other than medical image analysis could also adopt the method.

My current implementation assumes that the Poisson's ratio can be treated as known, which is also the case in most elastography studies, since the Young's modulus has more clinical significance in cancer detection. However, it has been reported that the Poisson's ratio plays a more important role than the elasticity in modeling deformation of breasts [Tanner et al., 2006], where the optimal Poisson's ratio also depends on the boundary

conditions — lower values could improve the results when volume changes need to be modeled. Therefore, a study of how different Poisson’s ratios affect elastography results could be a topic for future investigation.

In the near future, I plan to accelerate performance of the iterative scheme by many-core computing and model reduction. The resulting implementation can then be applied to more complicated physical and geometric models, such as situations with complex material property distributions, surface sliding, and large deformations. I would also like to explore the possibility of clinical trials of my method to noninvasive cancer staging based on the stiffness value. Virtual surgery and material engineering are some example application domains that would benefit from an automatic estimation of material properties and they can also directly benefit from this framework, worthy of further exploration.

Chapter 4

Treatment Image Estimation Based on Marker Locations for Dose Calculation

Tracking implanted seeds in the prostate during each radiation treatment delivery provides a high-frequency and accurate approximation of the organ locations, which enables the use of higher daily doses with tighter margins of the treatment beams and thus improves the efficiency of the radiotherapy. However, the lack of 3D image data with such a technique has prevented the use of dose calculation for assessing treatment results. I propose to use a reference statistical shape model generated from the planning image and a deformed version of the reference model fitted to the implanted marker locations during treatment to estimate a regionally dense deformation from the planning space to the treatment space. I also apply the optimization framework introduced in Chapter 3 to estimate the deformation by minimizing the distance between deformed and target marker locations.

My methods provide a means of estimating the treatment image by mapping planning image data to treatment space via the deformation field and therefore enables the calculation of dose distributions with marker tracking techniques during each treatment delivery.

4.1 Introduction

In order to deliver a high dose of radiation to cancerous tissue while sparing nearby normal tissues, modern radiotherapy protocols create steep dose gradients near the boundary of the target volume (I consider the prostate in this chapter). These techniques, however, are very sensitive to treatment uncertainties such as day-to-day changes in the geometry of the internal organs, because they have only a narrow margin of high dose around the target volume. To cope with this difficulty, image-guided radiotherapy (IGRT) [Wong et al., 2005] uses a CT image taken at planning time (the *planning image*) and a CT image taken immediately before a dose fraction (a *treatment image*) to detect potential positioning errors and changes in anatomic geometry relative to the planning image. Patients are realigned and beam apertures are reshaped to correct for these errors prior to treatment. Treatment images can also be used to calculate and assess the delivered dose in the context of adaptive radiation therapy, a form of treatment that compensates for differences between planned and delivered doses. The delivered dose is calculated in the treatment space given the known treatment parameters. A non-rigid transformation from the treatment space to the planning space is established by mapping the treatment image to the planning image using an image registration method [Maintz and Viergever, 1998; Foskey et al., 2005]. The transformation for each treatment day is then used to deform the radiation dose calculated for each treatment day to the planning space, where it can be added to dose delivered on other days. The result can then be used if necessary to modify the treatment plan to compensate for any discrepancy between the planned and actual cumulative delivered doses [Yan et al., 2000]. A problem with this approach is that the patient geometry is not sampled during actual treatment delivery [Litzenberg et al., 2007; Noel et al., 2009].

An alternative method to monitor the prostate motion is by tracking several (usually three) markers that are implanted in the prostate. For example, in the Calypso system, the markers take the form of electromagnetic transponders that can be tracked to sub-




Figure 4.1: The Calypso System; three electromagnetic transponders are implanted in the prostate (left) prior to the treatment planning and are tracked in real time during each treatment fraction to provide localization (right). Courtesy of Calypso.

millimeter accuracy at a sample rate of 10 Hz, using a GPS-like system [Langen et al., 2008; Balter et al., 2005]. The locations of the markers are used to position the patient accurately in the treatment machine, thus eliminating the need for image guidance or external marks on the patient’s skin or fixation device. Although the tracking information is accurate and takes intra-treatment motion into account, the image data needed for dose calculation is missing, eliminating the promise of calculating and accumulating delivered dose. In this chapter, I demonstrate that the non-rigid mapping from the planning space to the treatment space can be inferred in a neighborhood of the prostate from the three marker locations and used to estimate the missing treatment image data from the planning image.

My problem is different from the image registration problem in that the image data on the treatment day is missing, therefore the methods based on voxel-scale intensity matching [Foskey et al., 2005] or based on surface matching [Kaus et al., 2007] cannot be applied. Landmark-based registration methods [Rohr et al., 2001] require much more than three marker locations to interpolate the deformation field and are not directly applicable to my problem.

My first approach is to use a statistical shape model for the prostate so that the most likely prostate shape can be estimated given the measured marker locations. As

my shape model I choose to use the *m-rep* (medial representation) [Pizer et al., 2003] because it provides a coordinate system that can represent the interior and the nearby exterior of the object, so that the markers in the prostate can be also represented relative to an m-rep model of the prostate. I create a reference m-rep from the planning image and one for each treatment day, based on the marker locations. The deformation between the reference m-rep and estimated treatment m-rep provides a transformation between the planning space and the treatment space that can be used to estimate the treatment image to calculate delivered dose.

I also propose to use the simulation-based optimization framework presented in Chapter 3 to estimate the deformation around the target organ by minimizing the distance between the reference and treatment marker locations. A finite element (FE) model [Zienkiewicz and Taylor, 2005] for the reference organ and a bounding box is built and deformed during the optimization process. The external forces acting on the organ surface is optimized, and the resulting deformation field can be used to estimate the treatment image. The simulation-based approach can provide more reliable deformation at locations far from the sample points on the surface since the entire 3D domain is simulated, while with m-reps correspondence, the deformation also depends on the interpolation method.

I compare the estimated images with the real treatment images to demonstrate the ability to reconstruct treatment images, and I compare the dose histograms computed using the estimated treatment images to the histograms generated by an IGRT procedure (using actual CT images) to establish the feasibility of dose calculation using estimated treatment images.

For the rest of the chapter, I first present the deformation estimation method based on m-rep correspondence in Section 4.2 and the simulation-based approach in Section 4.3. The aspect of dose calculation for radiotherapy is reviewed in Section 4.4, followed by experimental results presented in Section 4.5.

4.2 Deformation Estimation Based on M-reps

Given the planning image, a reference m-rep of the prostate is created based on the gray-scale image and several user-specified landmarks on the organ boundary [Pizer et al., 2005]. Calypso markers are ~ 8.5 mm in length and ~ 1.5 mm in diameter that appear as bright spots in the image data. The location of each marker is computed as the center of mass of the bright spot, which closely corresponds to the origin of signals emitted by the miniature electromagnetic coil housed in the marker. For each treatment fraction, the three marker locations are given by the Calypso system, and the reference m-rep is then fitted to those locations by an optimization procedure. This procedure is similar to organ segmentation except that the seed locations, rather than image intensities, are what drive deformation within the trained shape space describing day-to-day shape variations of the prostate [Pizer et al., 2005]. The diffeomorphism implied by the two m-reps is then computed with a shape interpolation method called *rotational flows* [Levy et al., 2008]. The deformation is used to transform the planning image to generate an estimated treatment image for dose calculation.

4.2.1 M-reps

The discrete m-rep [Pizer et al., 2003] can be thought of as a discrete generalization of the medial axis [Blum and Wathen-Dunn, 1967]. A simple m-rep model consists of a discretely sampled medial surface. Sample points are referred to as *hubs*, and associated to each hub are two vectors (known as *spokes*) extending, on opposite sides of the medial sheet, from the hub to the boundary of the modeled object. The model is trained from a population of shapes by adjusting the parameters (the locations of the hubs and the directions and lengths of the spokes) to make the model fit each shape. From this data, a mean shape and its modes of variation can be determined [Merck et al., 2008]. This statistical representation can be used to reduce the effective dimension of the

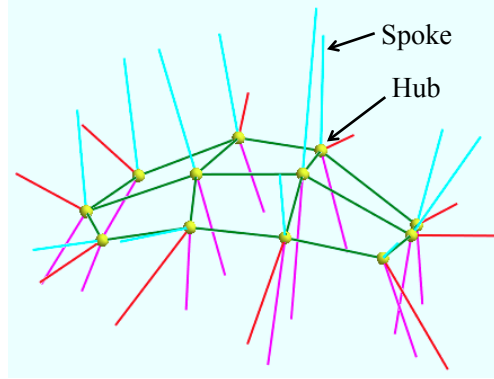


Figure 4.2: A simple 3×4 m-rep. A single shape is represented by the locations of the hubs and the lengths and directions of the spokes. The boundary of the object (not shown) passes through the ends of the spokes.

shape space, so that the limited information available from the three marker locations can be applied most effectively. For instance, changes orthogonal to the plane of the three points may be implied by their relative motions, given the statistical behavior of the shape space, even though that orthogonal direction is not explicitly sampled. At planning time, a patient-specific m-rep is created from the planning image for each patient. For each treatment fraction, the patient-specific m-rep is deformed using the modes of variation to fit the treatment image (in an IGRT) or the marker locations measured for each treatment fraction (in the case of using markers instead of treatment images).

An object-relative coordinate representation (*figural coordinates*) of the interior and nearby exterior of the organ is used to interpolate the locations of unsampled spokes. Two coordinates, u and v , reflect the arrangement of the hubs on the medial surface. Each hub, along with its associated spokes determines integer values for u and v , and points between them have non-integer values. An additional coordinate, ϕ , is equal to 1 along the entire top surface (bounded by the upper spokes extending from the outermost hubs) and -1 along the bottom surface, varying smoothly between those values over the intermediate portion of the boundary. A fourth coordinate, τ , gives the fraction of the distance from the medial sheet to the boundary of the organ. This fraction may be

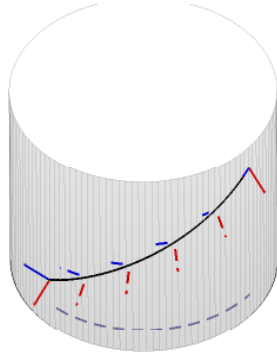


Figure 4.3: 3D rotational flow between two oriented points $(\mathbf{x}_0, \mathbf{E}_0)$ and $(\mathbf{x}_1, \mathbf{E}_1)$ lies on a helix whose axis is the rotation axis between \mathbf{E}_0 and \mathbf{E}_1 . Courtesy of Levy et al. [2008].

greater than 1 for points outside the organ.

4.2.2 Rotational Flows and Deformation Field Generation

Given the reference m-rep m_0 and the target m-rep m_1 created by fitting m_0 to the marker locations, the deformation from m_0 into m_1 is computed by interpolating between the sets of boundary points b_0 and b_1 , which are sampled using the same set of figural coordinates $\{(u_i, v_i, \phi_i) | i = 1, \dots, n\}$, on the boundary of m_0 and m_1 , respectively. Thus for each point \mathbf{x}_0 in the sets b_0 , there is a point \mathbf{x}_1 in b_1 that has the correspondence with \mathbf{x}_0 given by the m-reps and the figural coordinates. The goal of the rotational flows method is to interpolate between each pair of corresponding oriented vertices in \mathbb{R}^3 , $(\mathbf{x}_0, \mathbf{E}_0)$ and $(\mathbf{x}_1, \mathbf{E}_1)$, where $\mathbf{E}_i = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is the set of orthonormal bases representing the orientation for the point \mathbf{x}_i . In 3D cases, each oriented vertex moves along a helical path along the axis of rotation, as shown in Fig. 4.3. It can be shown that the interpolation is shape-maintaining if the two objects are similar and is size-maintaining if the two models are congruent [Levy et al., 2008].

Once the curved path for each pair of corresponding points is computed, the deformation for each voxel is computed by numerically integrating the subdivided displacement along each curved path. At each integration step, a radial basis function (RBF) inter-

polution is used to compute the deformation field in a rectangular box containing the prostate with a small margin, using the collection of landmarks. Some static landmarks are added on the boundary of the box to make the transition across the box boundary smoother. The deformation outside of the rectangular box is assumed to be zero. Due to the physics of dose deposition at high energies, this assumption has negligible effect on the accuracy of dose calculation.

4.3 Simulation-Based Estimation of Deformation

Given a reference surface model of the main organ (prostate) and three marker locations at the planning and treatment time, I create a finite element (FE) simulation that deforms to match the marker locations. A tetrahedral FE mesh is built from the segmentation of the planning image and a bounding box (region of interest). The deformation is modeled as an elastostatic process, where I assume different material properties for the main organ and for the surrounding tissue, and boundary forces are applied to the surface nodes of the main organ. During the optimization process, the external forces are updated to minimize the error in marker locations. A dense deformation field is then interpolated from the nodal displacements to generate the estimated treatment image.

4.3.1 Finite Element Modeling

The finite element method is based on a volumetric mesh representing the simulated domain. See Section 3.2.1 for more details. In this chapter, I consider a domain limited by a bounding box that roughly aligns with the surrounding bones (pelvic bones and the spine) where the main organ (prostate) resides. A tetrahedral mesh is built from the surface mesh of the main organ extracted from the planning image and the bounding box, as shown in Fig. 4.4. The two regions in the domain (organ and surrounding tissue) are each assigned a material property chosen based on previous work on FEM-based image




Figure 4.4: A finite element mesh used for estimating the deformation using the simulation-based approach. The red elements represent the prostate.

registration [Hensel et al., 2007; Lee et al., 2010b], and boundary conditions (prescribed nodal displacement or forces) are assigned to the surface nodes to drive the simulation. The bounding box is assumed to be static, and I optimize the nodal boundary forces at the organ surface.

4.3.2 Optimization

The objective function to be minimized is the sum of squared difference between deformed and treatment marker locations

$$\Phi(\mathbf{f}) = \sum_i \|\hat{\mathbf{p}}_i - \mathbf{p}_i\|^2,$$

where $\hat{\mathbf{p}}_i$ is the marker location in the deformed model, \mathbf{p}_i is the corresponding marker location in the treatment space, and \mathbf{f} is the external force vector, which is the parameters to be updated in the optimization loop. In practice, the marker locations in the planning space is modeled as nodes in the FE mesh, and the distances are computed by comparing the deformed node positions and the treatment marker locations. I employ the L-BFGS method [Nocedal and Wright, 1999] for the optimization. Once the optimal forces are computed, the displacements given by the FE simulation can be interpolated

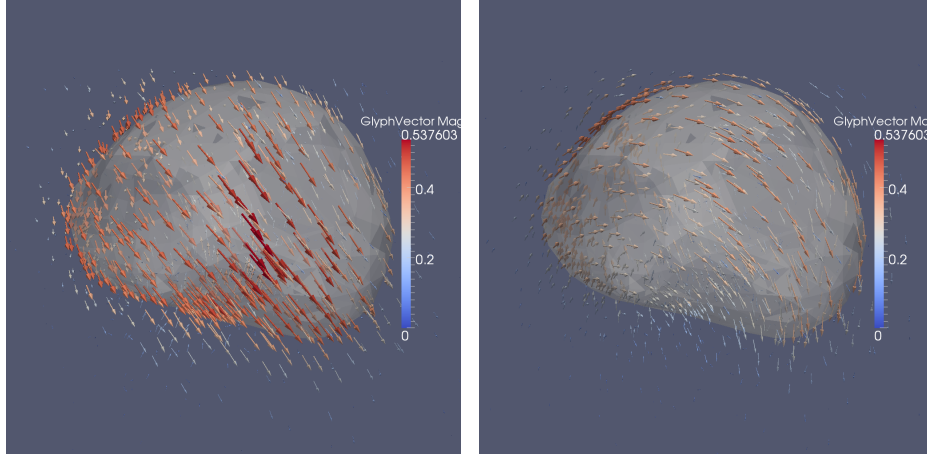


Figure 4.5: Two examples of nodal displacement field resulted from the simulation-based optimization framework; the white transparent surface is the prostate.

and used for estimating the treatment image. Fig. 4.5 shows two example visualizations of nodal displacement around the prostate resulted from FE simulations with optimal boundary forces. An advantage of the simulation-based approach over the m-reps-based method is that the entire 3D domain is considered in the FEM, and therefore the resulting deformation at internal pixels is more accurate physically. The fitted m-reps provide displacements only at sample points on the surface, and the deformation at internal pixels is a result of both the point correspondence and the interpolation method, which may not be as reliable at locations far from the sample points.

4.4 Dose Calculation

Calculation of delivered dose is a well-established part of radiation therapy planning. CT intensities are based on the absorptivity of tissue to radiation, so it is possible to calculate the amount of energy a given beam deposits in each portion of tissue it passes through or near, taking into account the attenuation of the beam and the scattering of radiation into neighboring tissue. I use the treatment planning system known as PlanUNC [Schreiber et al., 2006] to calculate dose. Dose calculations for treatment planning are expected to be $\pm 3\%$ compared to measurements.

These calculations are most commonly used in treatment planning. But when images are available for treatment fractions, dose can be calculated for those days too. Then, if a mapping can be constructed between the image of each treatment fraction and a common reference frame, the dose can be accumulated to determine the total amount of radiation received by each portion of the tissue. My approach is a reversal of the most common methods of dose accumulation. Taking the planning image as given, I derive a deformation first and derive the treatment image from it, rather than acquiring a treatment image first and then deriving the deformation.

4.5 Experiments

I evaluated my method on four patient data sets, which contain eight, six, seven, and six treatment fractions, respectively. The data for each patient consist of a planning image, treatment images, and Calypso marker locations during treatment. Note that for each treatment fraction, the patient is positioned in the treatment machine according to the Calypso origin, which also serves as the isocenter of the treatment plan (i.e., the point in the space where the central beam of radiation passes), and treatment images are taken only for comparison with estimated images. Since the treatment images are taken before the patient is positioned for treatment, they are visually aligned to the Calypso marker locations before the comparison with estimated images and the non-rigid image registration for dose calculation.

4.5.1 Comparison of Organ Surfaces

In order to assess the similarity between the estimated image and real treatment image, I create an m-rep of the prostate for each treatment image based on the gray-scale image and user-specified landmarks. The estimated m-rep (m_{marker} , reference m-rep fitted to the marker locations) is then compared to the treatment m-rep (m_{treat} , m-rep created

Table 4.1: Average errors in prostate surface and volume for the four patient data sets. The surface error is measured with average distance between two sets of sample points (D), and the volume overlap is measured with the ratio of intersection and union of the two volumes ($\frac{\text{Intersection}}{\text{Union}}$).


Patient	M-reps-based		FEM-based	
	D (cm)	$\frac{\text{Intersection}}{\text{Union}}$	D (cm)	$\frac{\text{Intersection}}{\text{Union}}$
1	0.08	0.90	0.07	0.91
2	0.04	0.93	0.07	0.90
3	0.10	0.85	0.06	0.90
4	0.04	0.95	0.05	0.94

with the treatment image) by computing the distance between sample points on the surfaces. For each sample point p on the surface of m_{treat} , its distance to the m-rep m_{marker} is approximated as $D(p, B_{\text{marker}}) = \min \{d(p, q) | q \in B_{\text{marker}}\}$, where B_{marker} is the set of sample points on the surface of m_{marker} , and $d(p, q)$ is the Euclidean distance between the two points p and q . For the surface deformed by the FE simulation, I use the surface nodes of the prostate in the volumetric mesh as the sample points. The quality of object matching can also be measured in terms of the ratio of intersection and union of the two volumes, which takes the value in $[0, 1]$, with the value one representing a perfect match, and zero meaning no overlap. Fig. 4.1 shows the average surface distances and volume overlap of the treatment prostate estimated using the two methods.


Note that these average distances are within the image resolution of the four data sets; two of them are $0.12 \times 0.12 \times 0.3$ cm, and the other two are $0.12 \times 0.12 \times 0.3$ cm. Fig. 4.6 shows an example visual comparison of the estimated image (using m-reps and the simulation-based scheme) and corresponding treatment image.

4.5.2 Comparison of Calculated Dose


The dose calculation using the estimated image is compared against the dose calculated with real treatment images to evaluate the accuracy of my method. The non-rigid transformation with the treatment images is computed in the same manner as with




(a) Image estimated using fitted m-reps



(b) Comparison of (a) and real treatment image



(c) Image estimated using simulation-based optimization scheme



(d) Comparison of (c) and real treatment image


Figure 4.6: Axial (left), coronal (center), and sagittal (right) views of an example m-rep comparison; blue contour: m-rep fitted to Calypso marker locations during the treatment fraction; green contour: surface deformed by the simulation-based scheme; red contour: m-rep created with real treatment image; the green crosshairs show the Calypso origin, which also serves as the isocenter of the treatment plan; (a, c) CT images estimated using fitted m-reps and the simulation-based scheme, respectively; (b, d) 4×4 checkerboard images comparing (a) and (c) to the real treatment image, respectively.

the image estimated using m-reps, as described in Section 4.2.2. The dose calculation is based on radiotherapy plans with the radiation beam targeting at volumes with different margins around the prostate (5 mm, 7 mm, 9 mm, 11 mm, and 13 mm). I consider the dose in the prostate and in the anterior rectal wall (the part of the rectum right next to the prostate). Fig. 4.7 shows the differential and cumulative dose-volume histograms (DVH) of one of the data sets using the estimated images and real treatment images (with 5 mm margin). Besides visual similarity, I also numerically compare the DVHs by considering the equivalent uniform dose (EUD) [Wu et al., 2002] of the differential DVHs, as shown in Tables 4.2 and 4.3, along with the errors of the values given by my methods relative to those given by the real images. There are no generally accepted standards for errors in calculating delivered dose, but $\pm 5\%$ is a reasonable goal. I observed relative errors of less than 3.2% in EUD for the four patient data sets I experimented on.


4.6 Summary

I presented two methods for estimating the treatment image using the planning image and locations of implanted markers during the treatment. The simulation-based approach has the advantage that the entire 3D domain is simulated, while the fitted m-reps only provide displacements at sample points on the surface, and the resulting deformation field depends more heavily on the interpolation method than in the simulation-based framework. I demonstrated the feasibility of image estimation by comparing the estimated images to real treatment images, and I also showed that the calculated dose histograms using the estimated images are close to those using real treatment images. My methods complement the Calypso system, where the prostate motion can be tracked accurately at a high frequency during the treatment but the image data is missing, so that the delivered dose distribution can be calculated and a safe delivery can be ensured.

In the future, I will experiment on more patient data. I will also investigate the



(a)



(b)

Figure 4.7: Differential (a) and cumulative (b) dose-volume histograms for the prostate (left) and anterior rectal wall (right); blue lines represent results with real treatment images, green lines represent results with images estimated using m-rep fitting, and red lines represent results with images estimated using FE simulation.

Table 4.2: Comparison of equivalent uniform dose of the dose-volume histograms for the prostate for the four patient data sets; the errors resulted from estimated images (using m-rep and FE simulation) are relative to the values given by the IGRT (using real treatment images). Each patient is experimented on using five target volume margins around the prostate.

Patient	Target Margin	IGRT EUD (cGy)	M-reps Est. Error	FE Est. Error
1	5 mm	4680.41	-0.10%	-0.17%
	7 mm	4665.60	0.04%	-0.27%
	9 mm	4665.88	0.13%	0.11%
	11 mm	4668.44	0.11%	0.05%
	13 mm	4669.89	-0.10%	0.01%
2	5 mm	4,647.33	1.67%	0.96%
	7 mm	4,634.99	1.71%	1.28%
	9 mm	4,627.81	1.72%	1.47%
	11 mm	4,479.13	1.73%	1.59%
	13 mm	4,629.33	1.52%	1.39%
3	5 mm	4,651.24	-0.23%	-0.33%
	7 mm	4,617.89	-0.11%	-0.20%
	9 mm	4,617.67	-0.04%	-0.09%
	11 mm	4,599.79	-0.03%	-0.03%
	13 mm	4,591.71	0.04%	0.01%
4	5 mm	4,477.87	0.35%	0.29%
	7 mm	4,467.08	0.37%	0.45%
	9 mm	4,459.99	0.36%	0.52%
	11 mm	4,460.13	0.32%	0.28%
	13 mm	4,458.89	0.56%	0.55%

Table 4.3: Comparison of equivalent uniform dose of the dose-volume histograms for the anterior rectal wall for the four patient data sets; the errors resulted from estimated images (using m-rep and FE simulation) are relative to the values given by the IGRT (using real treatment images). Each patient is experimented on using five target volume margins around the prostate.

Patient	Target Margin	IGRT EUD (cGy)	M-reps Est. Error	FE Est. Error
1	5 mm	4,348.68	0.27%	0.39%
	7 mm	4,417.91	-0.00%	0.08%
	9 mm	4,533.19	-0.22%	-0.25%
	11 mm	4,528.29	-0.19%	-0.21%
	13 mm	4,811.30	-0.63%	-0.71%
2	5 mm	4,463.73	1.19%	1.52%
	7 mm	4,479.61	1.31%	1.55%
	9 mm	4,504.62	1.42%	1.47%
	11 mm	4,333.20	1.48%	1.65%
	13 mm	4,543.98	1.61%	1.58%
3	5 mm	4,108.74	-2.12%	-3.12%
	7 mm	4,059.30	-2.10%	-3.15%
	9 mm	4,110.98	-1.95%	-2.94%
	11 mm	4,154.21	-1.73%	-2.63%
	13 mm	4,158.04	-1.80%	-2.70%
4	5 mm	3,757.31	1.15%	0.54%
	7 mm	3,916.76	0.97%	0.47%
	9 mm	4,043.46	0.83%	0.46%
	11 mm	4,182.75	0.70%	0.46%
	13 mm	4,255.77	0.64%	0.45%

feasibility of adjusting the treatment machine to compensate for the intra-treatment motion observed.

Chapter 5

Fast Optimization-Based Elasticity Parameter Estimation Using Reduced Models

Elasticity parameters are central to physically-based animation and medical image analysis. I present an accelerated method to automatically estimate these parameters for a deformation simulator using an iterative optimization framework, given the desired (target) output surface/shape. During the optimization, the input model is deformed by the simulator, and the distance between the deformed surface and the target surface is minimized numerically, as discussed in Chapter 3. To accelerate the optimization process, I introduce a dimension reduction technique to allow a trade-off between the computational efficiency and desired accuracy. The reduced model is constructed using statistical training with a set of example deformations. To demonstrate this approach, I apply the computational framework to 2D animations of elastic bodies simulated with a linear finite element method. I also present a 3D elastography example, which is simulated with a reduced-dimension finite element model to improve the performance of the optimizer.

5.1 Introduction

Physically-based simulations can help generate realistic scenes or animations without low-level control of a 3D model [Nealen et al., 2006]. To achieve a particular appearance, however, it frequently requires many iterations of adjusting simulation parameters, simulating, and assessing the results. For a large number of parameters and a complex, high-cost simulation, such an iterative process becomes extremely tedious and time-consuming, making parameter estimation a topic of significant interest in computer graphics [Treuille et al., 2003; McNamara et al., 2004; Bergou et al., 2007]. For example, estimation of material properties are essential in simulating the appearance of a particular type of cloth [Bhat et al., 2003], as well as in designing and fabricating materials for a certain deformation behavior [Bickel et al., 2010]. These material property estimation methods focus on materials that can be put into a specialized video capturing system to measure displacements, and a force measuring device is needed in the case of elasticity parameters [Bickel et al., 2009].

In 2D shape deformations, the notions of “rigidity,” “stiffness,” and “compressibility” are also conveyed in order to improve the physical plausibility of the animation [Igarashi et al., 2005; Weng et al., 2006; Yang et al., 2008]. However, physically-based methods have seldom been applied partly due to the difficulty in tuning simulation parameters.

Elasticity estimation is also of interest in non-invasive cancer detection, since human tissues are not always easy to obtain, and it is sometimes impossible to measure the actual parameters of a live patient. 2D Elastography [Ophir et al., 1999; Zhu et al., 2003; Kallel and Bertrand, 1996] is a method for estimating the elasticity value for each pixel in medical images, and most existing methods are based on a dense displacement field established by pixel-wise correspondence between pre- and post-compression images. However, in some imaging modalities, the intensity inside an organ is almost constant; only the displacements at the organ boundaries can be *approximated* using a surface matching or an intensity-gradient-based method, and it is difficult to find a reliable

dense displacement field.

In this chapter I introduce a fast elasticity parameter estimation method suitable both for 2D shape deformation and for 3D elastography. I extend the simulation-based framework introduced in Chapter 3, where two segmented images are registered using physics-based deformation. I minimize an (error) objective function that is based on the distance between the deformed surface and the target surface, with the elasticities and boundary forces as the parameters to the iterative optimizer. The framework does not require a force measuring device for finding elasticities (Young’s modulus) in a multi-material system, instead only the ratios between the elasticity values are computed. However, in applications where the dimensionality of the boundary forces is very high, the optimizer may become slow to converge and is more likely to get stuck in a local minimum. For such applications with example deformations available, I propose a novel method for statistically training a set of basis to represent the degrees of freedom for boundary forces, thereby improving the performance of the optimizer by more than an order of magnitude. I firstly acquire example deformations by matching the moving surface to several example target surfaces. A principal component analysis (PCA) is then performed with these example deformations to find the linear basis for model reduction.

For the rest of the chapter, I first review related work in Section 5.2. I give an overview of the optimization framework and the reduced-dimension finite element modeling in Section 5.3. I demonstrate the effectiveness of my approaches by showing experimental results in 2D shape deformation and 3D elastography in Section 5.4 and analyze the accuracy of the recovered elasticity values. I conclude with a summary and discussion of future work.

5.2 Previous Work

Physically-based deformable models have been applied to computer graphics for more than two decades, and there have been significant advances in many subareas such as numerical partial differential equations, multi-resolution modeling, modal analysis, and collision detection [Nealen et al., 2006; Teschner et al., 2005; Müller and Gross, 2004]. Artistic control of physically-based simulation has been an important topic in computer animation. For example, the visual shape of simulated fluid can be controlled with external forces [Treuille et al., 2003; McNamara et al., 2004]. By minimizing an objective function that measures the difference between the simulated density fields and the keyframes provided by the user, one can achieve the desired shape or look of the simulated fluid. Example-based methods has also been applied to art-directed elastic bodies [Martin et al., 2011], where a space of deformation is formed with example strains (keyframes). During a simulation, the configuration (deformation) is projected to the space of deformation. Additional external forces are then applied in order to match the projected and simulated configurations. My method, on the other hand, not only matches the deformed shape but also estimates the optimal elasticity values to achieve such a behavior. Optimizing elasticity and forces jointly is a much harder problem because the error function has different sensitivities to forces and to elasticity parameters. The problem of directing a simulation can alternatively be solved by adding physical details to an object (surface) animated by hand using the constrained Lagrangian mechanics approach [Bergou et al., 2007]. Their work is different from mine in that I am matching simulated results rather than *adding* simulated details to an existing animation.

Some work in computer graphics has explored parameter estimation for deformable model simulation. For example, Becker and Teschner [2007] proposed a framework using quadratic programming to find linear elastic parameters and analyzed the effects of noise in measurements. Pai et al. [2001] combined a trinocular stereo system and a

force measurement device to model deformable objects. The linear relationship between the tractions and displacements is estimated using a least squares formulation. [Bhat et al. \[2003\]](#) estimated cloth simulation parameters by comparing video of real fabric patches and simulated images and computing the error metric based on the orientation of each edge pixel, and they minimized the error using the continuous simulated annealing method [[Press, 2007](#)]. [Syllebranque and Boivin \[2008\]](#) used a similar optimization method with a force capture device, so that the boundary forces are known, to estimate the mechanical properties of deformable solids. They used video-based metrics to optimize for Poisson’s ratio and used the errors in computed boundary forces to optimize for Young’s modulus. [Bickel et al. \[2009\]](#) estimated non-linear heterogeneous material properties using a force sensor and a trinocular stereo vision system, so that the displacements of the vertices on the surface and the applied forces are known. They estimated the material properties by minimizing the error in vertex locations. While these methods depend on rendering and/or computer vision algorithms, my technique directly uses the surfaces of the deformed bodies to compute the error metric. In addition, the boundary conditions are unknown in my problem.

In 2D shape deformation, there have been attempts to optimize some quantities such as the area of the deformed mesh in order to create physically plausible results. [Igarashi et al. \[2005\]](#) presented an interactive shape manipulation method using only a few control points. They optimize for the rotation and the area of each triangle in the 2D mesh to create plausible deformations, and linearized optimizers are used due to the real-time constraint. [Weng et al. \[2006\]](#) formulated the problem using a non-linear optimization minimizing the Laplacian coordinates of the boundary curve as well as the local area of the interior. [Yang et al. \[2008\]](#) further improved the framework to make the stiffness tunable by allowing different amount of mesh distortion. However, these methods are not physically-based. My method can estimate the material properties and boundary conditions needed to achieve the desired shape, and the physical plausibility

is guaranteed by the simulator used in the optimization loop.

Estimation of material properties of human tissues is also important in the area of medical image analysis for detecting cancerous tissues, since cancerous tissues tend to be stiffer. Elasticity reconstruction, or 2D elastography, is a non-invasive method to acquire strain or elasticity images of soft tissues [Skovoroda and Emelianov, 1995]. Elastography is usually done by first estimating the optimal deformation field that relates two ultrasound images, one taken at the rest state, and the other taken when a known force is applied to the skin [Ophir et al., 1999; Rivaz et al., 2008]. Alternatively, the displacement field can be found with a modified MRI machine in tune with a mechanical vibration of tissues [Muthupillai and Ehman, 1996; Fu et al., 2000]. Once the deformation field and external forces are known, the material properties can be found by solving a least-squares problem [Zhu et al., 2003], assuming that the physical model is linear, or by using iterative optimization algorithms to minimize the error in the deformation field [Kallel and Bertrand, 1996; Balocco et al., 2008; Schnur and Zabararas, 1992]. Although slower than directly solving the inverse problem, these iterative methods do not require linearity of the underlying model and are therefore suitable for any physical model. Another kind of 2D elastography seeks to maximize image similarity while treating the displacement field as unknown [Washington and Miga, 2004]. However, the method depends on salient features within the object, which may not be present in human organs.

Surface capturing and registration are essential to my parameter estimation method. Computer vision methods such as trinocular stereo [Pai et al., 2001; Bickel et al., 2009] and 2D video tracking [Bhat et al., 2003; Syllebranque and Boivin, 2008] has been used to track the surfaces or landmarks. For medical images, statistical model-based segmentation methods such as those based on active shape models [Cootes et al., 1995] and m-reps [Pizer et al., 2003] can be used to extract smooth surfaces of organs. The extracted surfaces from different medical images needs to be registered in order to es-

timate nodal displacements. A typical polygonal surface registration method usually involves minimizing a surface distance error and a regularization term for maintaining vertex distribution [Kaus et al., 2007].

Finite element methods [Zienkiewicz and Taylor, 2005] (FEM) have become a standard simulation approach for elastic continua. In many engineering and medical applications, such as design analysis and image-guided surgery, some forms of simplified elasticity models and numerical methods are usually adopted in order to meet the performance requirements. For example, dimensional model reduction [Krysl et al., 2001] reduces the dimension of the finite element domain by finding a set of linear basis, or *reduced basis*, to transform the nodal displacement and boundary conditions into a lower-dimensional space. Subsequent solutions or integrations are done in the reduced space. Since high-frequency modes of the displacements are also removed in the process of dimension reduction, the reduced model saves computation through both the reduced matrix size and increased critical time step (for explicit integration), at the cost of reduced details in displacement. The set of reduced basis, however, is not trivial to find. Krysl et al. [2001] suggested the use of example displacements as a training set, where each sample is a snapshot of the displacement vectors when a certain load is applied to the full-dimension FE model. A PCA is then performed to find the most significant, important dimensions. Taylor et al. [2010] used the same training scheme with multiple load cases and implemented the simulator on the GPU. They applied the method to simulate brain shift during a synthetic surgery scene and reported a speedup of 9.6 times when using just five basis vectors. Barbič et al. [2009] used reduced-dimension control for keyframe animation. Animator-specified keyframes, along with some FEM-generated deformations filling the gaps in between and some natural vibration modes, are used as training samples for the PCA.




Figure 5.1: Flow chart of the optimization loop; the displacement field generated by the simulator is used in the objective function for the optimizer to update the parameters; updated parameters are fed back into the simulator, and so on.

5.3 Method

My algorithm minimizes an (error) objective function based on the separation between corresponding shapes. In each iteration, the objective function is computed by first simulating and deforming the surface using the current set of parameters, and then computing surface distances. My current implementation of the simulator uses the isotropic linear elasticity model, because it is widely used in real-time graphics with stiffness warping to reduce effects of non-linearity [Müller and Gross, 2004]. Other simulators can also be integrated with this framework.

The inputs to the problem are two triangulated 2D or 3D surfaces: the fixed (target) surface \mathbf{S}_f and the moving (source) surface \mathbf{S}_m . For 3D objects, I construct a tetrahedralization of the moving volume such that each face of \mathbf{S}_m is a face in the tetrahedralization, so that \mathbf{S}_m is characterized entirely by its set of nodes. The framework is built on a physically-based simulator that generates deformation fields with N_p unknown parameters $\mathbf{x} = [x_1, \dots, x_{N_p}]^T$, and a numerical optimizer to minimize an objective function $\Phi(\mathbf{x}) : \mathbb{R}^{N_p} \rightarrow \mathbb{R}$ defined by nodal displacements and surface matching metrics. During the optimization process, the physical model is refined in terms of more accurate parameters and converges to the model that can describe the deformation needed for the particular surface correspondence problem. The flow chart of my algorithm is shown in Fig. 5.1 and explained in detail in this section. I briefly review the optimization framework in Section 5.3.1 and present the acceleration method using

reduced-dimension modeling in Section 5.3.2.

5.3.1 Optimization-Based Framework with Finite Element Modeling

In the optimization loop, each nodal displacement $\mathbf{u}_l = [u_l, v_l, w_l]^T$ is always generated by a physically-based simulation, where any simulation method can be used to solve the elastostatic problem. Using a linear FEM, I have

$$\mathbf{K}\mathbf{u} = \mathbf{f}, \tag{5.1}$$

where \mathbf{u} is the displacement vector that consists of \mathbf{u}_l , \mathbf{K} is the stiffness matrix constructed using some material properties, and \mathbf{f} is the vector of external forces. The resulting deformed surface is compared against the target surface to evaluate the correspondence error (objective function). The optimizer then finds a descent direction to update the parameters, and the new parameters are fed back to the simulator for the next iteration. See Section 3.2 for detailed formulation of the FEM and the distance-based objective function.

5.3.1.1 Initial Guess of Parameters

The initial guess of external forces is important in avoiding local minima. I use a simple iterative surface matching to roughly match the boundary nodes, and the matching provides estimated displacements of the boundary nodes, which are substituted into the FEM (Equation 5.1) to find an initial guess of forces. For the material properties, I evaluate the values of $\Phi(\mathbf{E})$ using several \mathbf{E} values, and use the one that results in the smallest error as the initial guess for \mathbf{E} .

The iterative surface matching is based on the gradient of the distance map and vertex redistribution: in each iteration, each vertex on the source surface is moved along

the direction of the distance map of the target surface in order to minimize the surface distance. To avoid degenerate polygons, imaginary spring forces are applied to each edge to maintain the vertex distribution. Formally, in iteration $n + 1$, each node \mathbf{v}_l^n (from previous iteration) is pulled by an imaginary force

$$\mathbf{f}^{n+1} = -k_D \mathbf{d}(\mathbf{v}_l^n, \mathbf{S}_f) + \sum_{\mathbf{v}_N^n \in \mathcal{N}(\mathbf{v}_l^n)} k_N (\mathbf{v}_N^n - \mathbf{v}_l^n),$$

where k_D and k_N are constants controlling magnitudes of imaginary forces, and $\mathcal{N}(\mathbf{v}_l^n)$ is the set of neighboring nodes of \mathbf{v}_l^n . In the first term, $-\mathbf{d}(\mathbf{v}_l^n, \mathbf{S}_f)$ pulls the node towards the surface \mathbf{S}_f , and in the second term, $(\mathbf{v}_N^n - \mathbf{v}_l^n)$ maintains the relative spacing between the nodes so that degenerate triangles are avoided.

5.3.2 Acceleration using Reduced-Dimension Modeling

The iterative optimization scheme requires a large number of solutions of the global linear system provided by the FEM. While the performance is acceptable for simple models with low degrees of freedom, it may take more than eight hours for some complex 3D models with many thousands of degrees of freedom for boundary forces. Here I present an acceleration method that reduces the dimensionality of the parameter space (the boundary forces) and thus improves the convergence of the optimizer.

5.3.2.1 Reduced-Dimension FEM

Krysl et al. [2001] suggested that the FE model could be simplified with a set of basis for the displacement field. The reduced-dimension displacement vector $\hat{\mathbf{u}}$ is computed with

$$\mathbf{u} = \mathbf{W}\hat{\mathbf{u}}, \tag{5.2}$$

where \mathbf{W} is the column matrix of orthonormal basis called the reduced basis. Substituting Equation 5.2 into Equation 5.1 and left-multiply both sides with \mathbf{W}^T , I have the

reduced-dimension linear system

$$\begin{aligned}\mathbf{W}^T \mathbf{K} \mathbf{W} \hat{\mathbf{u}} &= \mathbf{W}^T \mathbf{f} \\ \hat{\mathbf{K}} \hat{\mathbf{u}} &= \hat{\mathbf{f}},\end{aligned}\tag{5.3}$$

where $\hat{\mathbf{K}} = \mathbf{W}^T \mathbf{K} \mathbf{W}$ and $\hat{\mathbf{f}} = \mathbf{W}^T \mathbf{f}$.

5.3.2.2 Training of Reduced Basis

In order to obtain \mathbf{W} , a set of example displacements is used in a PCA, and some of the resulting principal components are chosen as the reduced basis. However, the generation of the example displacements is nontrivial and depends on the specifics of the application. In keyframe-based animation, for instance, the keyframes can serve as the examples [Barbič et al., 2009]. In the case where there is one source surface and multiple target surfaces, example displacements are generated by matching the source surface to the target surfaces (using the method described in Section 5.3.1.1).

Notice that only the degrees of freedom (DOFs) of the boundary nodes (where boundary forces are applied) are reduced, whereas in [Krysl et al., 2001], all the DOFs are reduced. My approach does not provide as much acceleration in the FEM itself, but the dimensionality of the parameter is greatly reduced, and thus the performance of the optimization is improved. Moreover, example displacements are much easier to obtain for the boundary nodes than for all the internal nodes.

5.3.2.3 Training for Multiple Source Surfaces

I consider a more general setting where there are multiple source surfaces that are similar, and I want to find a common set of reduced basis for each FE model built from a similar source surface. For example, in medical image analysis, there can be multiple 3D images of each patient taken on different days. Once the organs in the images are

delineated, they can provide examples of organ deformations. Further, if images from different patients can all be used as example deformations, the resulting basis can more effectively represent the general set of displacements, even for surfaces that are not used in the training process. In order to achieve this, a one-to-one correspondence must be established between the boundary nodes in all the FE models, and only the degrees of freedom for these boundary nodes are reduced (and other DOFs are not affected by the basis matrix \mathbf{W}).

Correspondence of boundary nodes. I first choose an *atlas* surface mesh \mathbf{S}_a consisting of only the boundary nodes. The surface \mathbf{S}_a is deformed through an iterative surface matching (using the method described in Section 5.3.1.1) to fit each specific source surface \mathbf{S}_m . The atlas surfaces fitted to \mathbf{S}_m serve as the *standardized* source surface \mathbf{S}'_m , from which the FE mesh for that source surface is created. Namely, the boundary nodes in each specific FE mesh are all from the deformed atlas surfaces (fitted to \mathbf{S}_m), while other surface nodes that has no boundary forces applied are directly from \mathbf{S}_m . As a result, a one-to-one correspondence is established between the boundary nodes from the FE meshes that are created through the same process. Example deformations are then generated by matching each \mathbf{S}'_m to the corresponding target surfaces. The process of generating the node correspondence and example deformations is shown in Fig. 5.2. Because all these FE meshes share the same set of boundary nodes, the PCA training can make use of all these example displacements, and the resulting reduced basis can be used for any of these FE meshes. Although the nodal correspondence between each pair of \mathbf{S}'_m may not be accurate, I show in the experimental results that the cross-patient PCA training works well in practice.

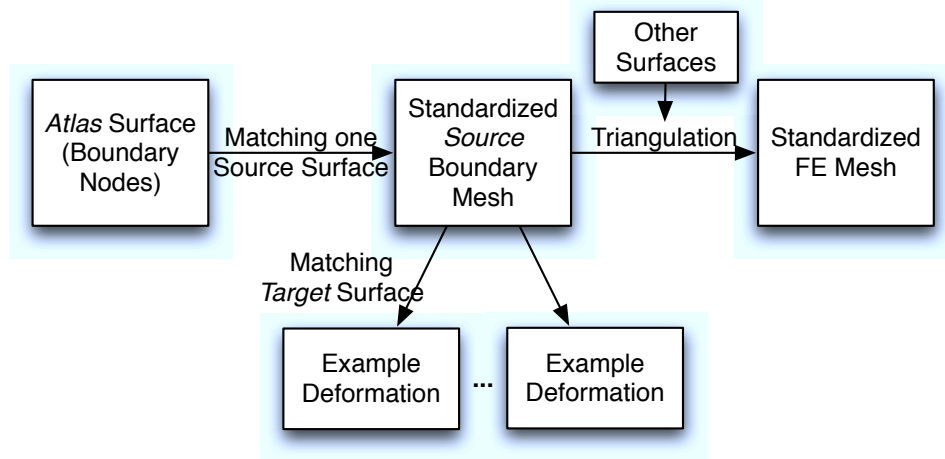


Figure 5.2: The process of generating the the node correspondence and example deformations for one source surface; this procedure is repeated for multiple source meshes, and the resulting reduced basis can be applied to other FE meshes using the same node correspondence procedure.

5.4 Implementation and Results

I have applied my method to 2D shape deformations and a 3D elastography example using the reduced-dimension modeling.

5.4.1 Physically-Based 2D Shape Deformation

In 2D shape deformation with a single material, the elasticity (Young’s modulus) cannot be estimated, since I do not know the true force applied. On the other hand, the compressibility (Poisson’s ratio) is the main source of difference in deformed shape, if the same first-type boundary conditions are applied. A perfectly incompressible material has the value 0.5 (but in the isotropic linear elasticity model, the value cannot be exactly 0.5). Fig. 5.3 shows the difference caused by two different Poisson’s ratios.

Given the source piecewise linear (represented by vertices and edges) shape, I build the triangular finite element mesh that fills the shape. The target shape is also given in the piecewise linear format, and the distance function between the two shapes is based on distance between corresponding vertices. The boundary forces to be optimized are



Figure 5.3: Example 2D shape deformations showing different results with different Poisson's ratios ν (each column results from the same first-type boundary conditions); top: $\nu = 0.1$; middle: $\nu = 0.49$; bottom: comparison by overlaying the two results.

applied to a subset of nodes, which can be chosen by the user. Fig. 5.4 shows the flow chart of the estimation process. The FE mesh and the recovered Poisson's ratio can be used to generate other deformations of the same object.

To test the accuracy of the estimated values, I perform a series of experiments where each target shape is generated with FEM using a Poisson's ratio ν , so that I know the ground truth. The dimension reduction approach is not applied to the 2D problem, since the number of DOFs is already small. Fig. 5.5 shows the target shape and a plot of surface errors for each ground truth ν value. The error in the ν value is about 10% in the four cases. In these examples, the nodes at the bottom of the alphabet 'A' are fixed, and I optimize for the boundary forces at the top of the 'A.' The results also show that the shape is more sensitive to Poisson's ratio when the value is close to 0.5 (nearly incompressible). The entire optimization process takes about four seconds for each case. The experiments show that my framework can be easily applied to 2D shape deformations with very high accuracy and efficiency.

5.4.2 3D Elastography using Reduced Modeling

I apply my framework to 3D elastography of the prostate. The domain considered consists of the bladder, the prostate, the rectum, and the bones. The space is filled

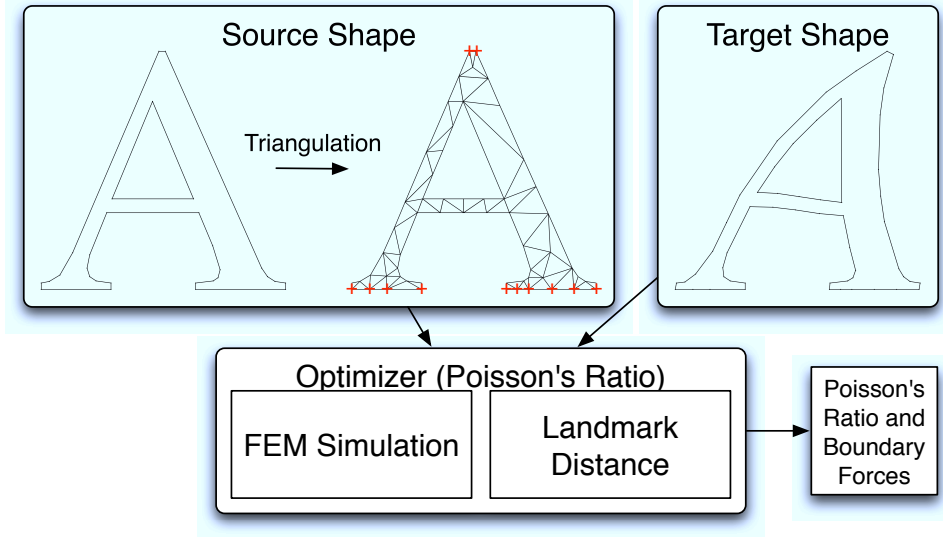
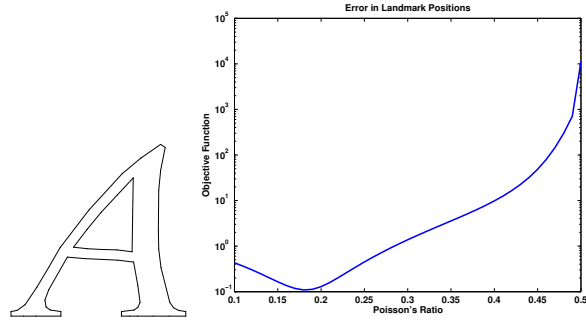



Figure 5.4: Flow chart of Poisson's ratio estimation given source and target 2D shapes; an FE model is built according to the source shape, and the distance between the deformed source shape and the target shape is minimized. The red cross marks the nodes with boundary conditions applied.


with tetrahedral elements, and the bladder and the rectum are hollow to reflect their structure, as shown in Fig. 5.6. There are two material properties, one assigned to the prostate (red elements in Fig. 5.6), and the other assigned to the elements between organs (shown in green in Fig. 5.6). The goal of the elastography is to find the ratio of Young's moduli of the prostate to that of the other elements. (The bone is always fixed, so its material property is irrelevant.) The Young's modulus is the material property to be optimized here because of its significance in non-invasive cancer detection. Poisson's ratios are set to 0.4 and 0.35 for the two materials, based on common values used in medical simulations [Hensel et al., 2007]. Boundary forces are applied to the surface nodes of the bladder and the rectum. The parameters to be optimized are the Young's modulus of the prostate and the boundary forces applied to the bladder and the rectum. I have not yet been able to estimate the Poisson's ratio along with the Young's modulus because the surface distance is much less sensitive to the Poisson's ratio.




(a) $\nu = 0.2$



(b) $\nu = 0.3$



(c) $\nu = 0.4$



(d) $\nu = 0.49$

Figure 5.5: Test of accuracy of the recovered Poisson's ratios ν ; the four target shapes are generated using four different ν values; for each subfigure, on the left is the target shape, and on the right is the plot of surface errors versus ν values.




Figure 5.6: The tetrahedral mesh of the bladder-prostate-rectum area (left) and a close-up view of the three organs (right), where the white surface is the target prostate surface.

Reduced basis training. The degrees of freedom for the boundary forces is around 4,900, which makes the optimizer slow to converge (taking thousands of iterations). Therefore, a dimension reduction is performed with 108 example deformations from CT images. These images are from nine patient data sets, each has one source image and several target images. The first step of generating examples is to create the node-to-node correspondence between source surfaces from different (patient) data sets. I use a surface matching approach to establish the correspondence: an atlas surface is iteratively deformed to match each patient-specific source surface, and therefore each patient-specific surface has the same set of nodes. The patient-specific source surface is then matched with corresponding target surfaces to generate the example deformations for the organ. Fig. 5.7 shows visualizations (as nodal displacements) of the first two principal components (basis vectors) resulted from the PCA. I choose to use 92 basis vectors for the bladder and 100 for the rectum due to the high accuracy requirement in elasticity parameter estimation. Fig. 5.8 shows the relationship between the number of basis vectors used, the accuracy of recovered elasticity, and the performance of the optimizer (in terms of number of iterations needed). A reduced accuracy of recovered elasticity is expected with a smaller number of basis vectors, since the accuracy in the simulated displacement is reduced due to the model simplification. I observe better




Figure 5.7: The first two principal components (visualized as nodal displacements on bladder and rectum surfaces) resulted from a PCA using 108 example deformations.

accuracy and slower convergence with a larger number of basis vectors used. Namely, the reduced basis provide a trade-off between the accuracy and the performance. Also notice that even with 191 basis vectors, the performance is still 32 times of that using a full-dimension model (which has 4,932 DOFs). The slight increase in error (see Fig. 5.8) when the number of basis vectors increased to 139 is likely the result of optimizer sticking to a local minimum.

I then test the accuracy of estimated elasticity values using two patient-specific FE models that are not used in the training process. The models are deformed using three different elasticity values to generate the target surfaces, so that I know the ground truth for the optimization. The recovered elasticity values in the six test cases are shown in Table 5.1. The relative errors in Young’s modulus range from 0.1% up to 8.1% and are generally less than 10% for the set of basis chosen. The average surface error (distance) is 0.08 cm, down from 1.81 cm with only rigid alignment. The accuracy is higher around the prostate (where the deformation is more predictable); the average error is 0.01 cm, down from 0.44 cm before deformation. The error in forces is not considered because force measurements are not present in my framework – only the surface data are required, and




Figure 5.8: The number of reduced basis vectors provides a means of trade-off between optimizer performance and accuracy of recovered elasticity values; the error (blue) decreases and the number of optimization iterations (green) increases with the number of basis vectors used. The slight increase in error with 139 basis vectors is likely the result of optimizer sticking to a local minimum.

Table 5.1: Error in recovered modulus of elasticity for two synthetic multi-organ scenes using a reduced-dimension FEM; the 18 to 62 times speedup is compared against an optimization using a full-dimension FEM.

True Elasticity (kPa)		25	50	75
Scene 1	Recovered	25.80	54.07	75.60
	Error (%)	3.2%	8.1%	0.8%
	Speedup	32x	62x	30x
Scene 2	Recovered	24.00	53.38	74.95
	Error (%)	4.0%	6.8%	0.1%
	Speedup	60x	19x	18x

the synthetic data sets are deformed using first-type boundary conditions (prescribed displacements at boundary nodes). The optimizer takes around 50-250 iterations to converge – in the amount of time that is about 1/60 to 1/20 of what would otherwise be needed when using a full FE model. The average time needed for each iteration of force optimization is about 9 seconds. These results show that the optimization framework using the dimension reduction approach can achieve an accuracy comparable to that using a full FEM, while taking substantially less time.

5.5 Conclusion

I presented an accelerated method based on a general framework for estimating deformable body simulation parameters for achieving a certain target shape. The simulation parameters can be a combination of Young’s modulus, the Poisson’s ratio, and the external forces, depending on the application. The error metric can be based on surface distance map or landmarks, depending on the features that one can find on the surfaces. The main contribution of my work is the acceleration of the optimization framework using reduced-dimension modeling: a set of example deformations are used for the statistical training for the basis, which are used to reduce the dimensionality of the parameter space. Furthermore, my acceleration method provides a trade-off between performance and accuracy by controlling the number of basis vectors used. However,

the accuracy of the reduced model is limited by the quantity of example deformations, and therefore my acceleration method is not suitable for applications lacking examples. I demonstrated this optimization-based framework on an implementation using a linear FEM and reduced-dimension modeling to accelerate the overall performance by more than an order of magnitude. More complex models, such as the non-linear FEM, can also be used. I have also shown that this method can be easily applied to 2D shape deformation and 3D elastography for the prostate.

In the future, I would like to experiment with different elasticity models, so that large deformations in character animations can also be simulated. Since the convergence of this approach depends on the number of basis vectors needed, the reduced-dimension acceleration technique could benefit from a better statistical training method, such as the non-linear PCA. Modal analysis of the FE model could also further reduce to the computation. I would also like to integrate the material property estimation with patient-specific virtual surgery, in order to provide a more accurate, patient-specific model.

Chapter 6

Discussion

In this dissertation, I have presented a general simulation-based optimization scheme for estimating simulation parameters with surface information. The method does not depend on known displacement field or external forces and is therefore suitable for any imaging modality where the objects can be segmented. I demonstrate the feasibility of using the optimization framework for medical image analysis by showing examples in prostate cancer staging, nonrigid image registration, and treatment image estimation using marker locations. The performance of the optimization scheme is improved by reducing the number of degrees of freedom in the FEM, thus providing a trade-off between speed and accuracy.

6.1 Summary of Results

My method is applied to prostate cancer stage assessment since the stiffness of the prostate is related to tumors inside the organ. I show experimental results with ten patient data sets, and the recovered elasticity values for the prostate are positively correlated to clinical cancer stages. Since the resolution of my elastogram is limited to organ boundaries (assuming constant material for each region), I present synthetic experimental results showing how the overall elasticity value for the prostate reflects the size of the tumor. The optimization framework is also integrated with a nonlinear

FEM when only the elasticity is optimized, and the experiments with real patient data also show significant positive correlations with cancer stages. With a linear FEM, only the ratio of elasticity values can be recovered, and the nonlinear FEM exhibits results (recovered ratios) that are very similar to the linear case under small deformation.

Since the optimization framework minimizes the surface distance, the resulting simulation also provides an image registration. Compared to image-based approaches, the deformation is generated by a simulator and therefore the quality of the displacement field is guaranteed by the simulation. Compared to landmark- or model-based approaches, my method generates displacements not only for the sparse set of points but also for nodes in the volumetric mesh. The resulting displacement field, therefore, is much less affected by the interpolation method employed. I also show experimental results showing image matching quality that is comparable to an image-based approach.

My optimization framework can also be based on landmark locations. For example, in radiotherapy with marker localization, the three marker locations instead of the surface data can provide information for generating a dense displacement field, which is used for estimating the treatment CT image for dose calculation. Compared to a model-based method such as using the m-reps fitted to the markers, the simulation-based approach generates more consistent deformation field since all the nodes in the mesh are simulated.

Material properties other than the elasticity can also be recovered using my method. For example, when only one object is present, the Poisson's ratio, rather than the elasticity, is the key to achieve a certain shape. I demonstrate the optimization of Poisson's ratio for 2D shape animations using landmarks in the error metric.

Finally, to address the efficiency issue of the force optimization, I propose to use statistical training of example surface displacements to reduce the degrees of freedom in the FEM, and therefore the dimension of the optimized parameters (forces). The examples are built from surface matching results of real patient data, and the FE mesh for each patient is built so that a common set of reduced basis can be used for all these

meshes. Synthetic experiments show the accuracy of the reduced-dimension method, and the speed up over the full-dimension method is more than an magnitude. The trade-off between speed and accuracy is achieved by choosing different numbers of basis vectors.

6.2 Future Work

There are several possible research directions extending this dissertation. Firstly, the accuracy of my method is currently verified using synthetic data with known parameters. Further verifications could involve building phantoms with known materials and taking 3D images of the phantom under different amount of compressions.

Since the error function in the optimization loop depends on surface distance, the accuracy will inevitably be affected by the quality of the segmentations, which depends on the algorithm and the image modality used. For example, segmentations for ultrasound images are usually less reliable due to the speckles, and the smoothness of the resulting surface can depend on the image resolution. Therefore, a more thorough analysis on errors in the segmentation and in recovered parameters can be conducted to clarify how much effect the segmentation method has.

In terms of the clinical application in cancer detection, a larger scale clinical trial of cancer stage assessment could provide more real patient data and better statistics. Along with other indicators of cancer (such as the antigen level), I could build a classifier using machine learning approaches. The accelerated optimization framework using reduced-dimension models can also benefit from more example displacements, since more sample in the PCA results in more principal components. Besides increasing the number of samples, a better statistical training method could also improve the simulation and optimization. A more thorough comparison between the reduced and full dimensional models can be conducted using more real patient data sets.

Application in patient-specific modeling. My method can also be applied to model patient-specific tissue parameters. For example, an image-guided surgery usually involves deformable tissues, and the deformation during the surgery makes the image guidance inaccurate. In order to compensate for the movements, physically-based modeling has been applied to approximate the deformation. Since each patient exhibits different tissue parameters, the simulation can be further improved if one knows the material properties of each patient. Since a planning image is already taken for the surgery, an additional image under different pose could be useful for estimating the tissue parameter, if taking such an additional image is possible. The resulting patient-specific simulation can also be applied to surgical training in order to prepare the surgeons for handling different material properties from different patients.

Application in computer graphics. Since simulation parameter estimation has been a difficulty in physically-based modeling, my approach can improve applications using physically-based simulations. For example, I have shown applications in 2D shape animation, and I can further extend the method to handle 3D animations. However, the surface matching (for providing initial guess of the force) can become much more difficult in general cases other than medical applications, where the difference between surfaces is usually small. Thus, a robust surface matching method can be crucial for a successful optimization.

Bibliography

- Alterovitz, R., Goldberg, K., Pouliot, J., Hsu, I., Kim, Y., Noworolski, S., and Kurhanewicz, J. (2006). Registration of MR prostate images with biomechanical modeling and nonlinear parameter estimation. *Medical Physics*, 33(2):446–454. 23, 29
- American Cancer Society (2012). *Cancer Facts & Figures 2012*. American Cancer Society, Atlanta. 1
- Bajcsy, R., Lieberman, R., and Reivich, M. (1983). A computerized system for the elastic matching of deformed radiographic images to idealized atlas images. *J Comput Assist Tomogr*, 7(4):618–625. 22
- Balay, S., Groppe, W. D., McInnes, L. C., and Smith, B. F. (1997). Efficient management of parallelism in object-oriented numerical software libraries. *Modern Software Tools in Scientific Computing*, pages 163–202. 40
- Balocco, S., Camara, O., and Frangi, A. F. (2008). Towards regional elastography of intracranial aneurysms. In *Medical Image Computing and Computer-Assisted Intervention*, volume 11, pages 131–8. PMID: 18982598. 81
- Balter, J. M., Wright, J. N., Newell, L. J., Friemel, B., Dimmer, S., Cheng, Y., Wong, J., Vertatschitsch, E., and Mate, T. P. (2005). Accuracy of a wireless localization system for radiotherapy. *International Journal of Radiation Oncology, Biology, Physics*, 61(3):933–937. PMID: 15708277. 12, 60
- Baraff, D. and Witkin, A. (1998). Large steps in cloth simulation. *Computer Graphics*, 32(Proceedings of SIGGRAPH 98):54, 43. 25
- Barbič, J., da Silva, M., and Popović, J. (2009). Deformable object animation using reduced optimal control. *ACM Transactions on Graphics - Proceedings of ACM SIGGRAPH 2009*, 28(3):1. 14, 82, 86
- Becker, M. and Teschner, M. (2007). Robust and efficient estimation of elasticity parameters using the linear finite element method. *Proc. of Simulation and Visualization*, pages 15–28. 20, 79
- Bergou, M., Mathur, S., Wardetzky, M., and Grinspun, E. (2007). TRACKS: toward directable thin shells. In *ACM SIGGRAPH 2007 papers*, page 50, San Diego, California. ACM. 25, 77, 79
- Bharatha, A., Hirose, M., Hata, N., Warfield, S. K., Ferrant, M., Zou, K. H., Suarez-Santana, E., Ruiz-Alzola, J., D’Amico, A., Cormack, R. A., Kikinis, R., Jolesz, F. A., and Tempany, C. M. C. (2001). Evaluation of three-dimensional finite element-based deformable registration of pre- and intraoperative prostate imaging. *Medical Physics*, 28(12):2551. 29, 39

- Bhat, K. S., Twigg, C. D., Hodgins, J. K., Khosla, P. K., Popovic, Z., and Seitz, S. M. (2003). Estimating cloth simulation parameters from video. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 37–51. Eurographics Association Aire-la-Ville, Switzerland. 6, 25, 77, 80, 81
- Bickel, B., Bächer, M., Otaduy, M. A., Lee, H. R., Pfister, H., Gross, M., and Matusik, W. (2010). Design and fabrication of materials with desired deformation behavior. In *ACM SIGGRAPH 2010 papers*, pages 1–10, Los Angeles, California. ACM. 25, 77
- Bickel, B., Bächer, M., Otaduy, M. A., Matusik, W., Pfister, H., and Gross, M. (2009). Capture and modeling of non-linear heterogeneous soft tissue. In *ACM SIGGRAPH 2009 papers*, pages 1–9, New Orleans, Louisiana. ACM. 6, 25, 77, 80, 81
- Blum, H. and Wathen-Dunn, W. (1967). A transformation for extracting new descriptors of shape. *Models for the Perception of Speech and Visual Form*, pages 380, 362. 62
- Cash, D., Miga, M., Sinha, T., Galloway, R., and Chapman, W. (2005). Compensating for intraoperative soft-tissue deformations using incomplete surface data and finite elements. *Medical Imaging, IEEE Transactions on*, 24(11):1479–1491. 3, 29
- Chentanez, N., Alterovitz, R., Ritchie, D., Cho, L., Hauser, K. K., Goldberg, K., Shewchuk, J. R., and O’Brien, J. F. (2009). Interactive simulation of surgical needle insertion and steering. In *ACM SIGGRAPH 2009 papers*, pages 1–10, New Orleans, Louisiana. ACM. 3
- Chopra, R., Arani, A., Huang, Y., Musquera, M., Wachsmuth, J., Bronskill, M., and Plewes, D. (2009). In vivo MR elastography of the prostate gland using a transurethral actuator. *Magnetic Resonance in Medicine*, 62(3):665–671. 4, 20
- Christensen, G., Rabbitt, R. D., and Miller, M. I. (1996). Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.*, 5(10). 22
- Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. (1995). Active shape models — their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59. 81
- Courtis, P. and Samani, A. (2007). Detecting mechanical abnormalities in prostate tissue using FE-Based image registration. In Ayache, N., Ourselin, S., and Maeder, A., editors, *MICCAI 2007*, volume 4792 of *Lecture Notes in Computer Science*, pages 244–251. Springer Berlin / Heidelberg.
- Crouch, J., Pizer, S., Chaney, E., Hu, Y.-C., Mageras, G., and Zaider, M. (2007). Automated finite-element analysis for deformable registration of prostate images. *Medical Imaging, IEEE Transactions on*, 26(10):1379–1390. 29
- Egorov, V., Ayrapetyan, S., and Sarvazyan, A. (2006). Prostate mechanical imaging: 3-d image composition and feature calculations. *Medical Imaging, IEEE Transactions on*, 25(10):1329–1340. 19

- Enright, D., Marschner, S., and Fedkiw, R. (2002). Animation and rendering of complex water surfaces. In *Proceedings of SIGGRAPH 02*, pages 736–744, San Antonio, Texas. ACM. 24
- Eskandari, H., Salcudean, S. E., Rohling, R., and Bell, I. (2011). Real-time solution of the finite element inverse problem of viscoelasticity. *Inverse Problems*, 27(8):085002. 20
- Fedkiw, R., Stam, J., and Jensen, H. W. (2001). Visual simulation of smoke. In *Proceedings of SIGGRAPH 01*, pages 15–22. ACM. 24
- Ferrant, M., Warfield, S. K., Guttman, C. R. G., Mulkern, R. V., Jolesz, F. A., and Kikinis, R. (1999). 3D image matching using a finite element based elastic deformation model. In *MICCAI*, pages 202–209. 23, 29
- Ferrant, M., Warfield, S. K., Nabavi, A., Jolesz, F. A., and Kikinis, R. (2000). Registration of 3D intraoperative MR images of the brain using a finite element biomechanical model. In *MICCAI*, pages 19–28. 5, 23, 29
- Foskey, M., Davis, B., Goyal, L., Chang, S., Chaney, E., Strehl, N., Tomei, S., Rosenman, J., and Joshi, S. (2005). Large deformation three-dimensional image registration in image-guided radiation therapy. *Physics in Medicine and Biology*, 50:5869–5892. 21, 22, 30, 59, 60
- Foster, N. and Fedkiw, R. (2001). Practical animation of liquids. In *Proceedings of SIGGRAPH 01*, pages 23–30. ACM. 24
- Fu, D., Levinson, S., Gracewski, S., and Parker, K. (2000). Non-invasive quantitative reconstruction of tissue elasticity using an iterative forward approach. *PHYSICS IN MEDICINE AND BIOLOGY*, 45(6):1495–1510. 81
- Han, L., Hipwell, J. H., Tanner, C., Taylor, Z., Mertzaniidou, T., Cardoso, J., Ourselin, S., and Hawkes, D. J. (2012). Development of patient-specific biomechanical models for predicting large breast deformation. *Physics in Medicine and Biology*, 57(2):455–472.
- Hensel, J. M., Menard, C., Chung, P. W. M., Milosevic, M. F., Kirilova, A., Moseley, J. L., Haider, M. A., and Brock, K. K. (2007). Development of multiorgan finite element-based prostate deformation model enabling registration of endorectal coil magnetic resonance imaging for radiotherapy planning. *Int. J. Radiation Oncology Bio. Phys.*, 68(5):1522–1528. 23, 29, 30, 32, 39, 66, 90
- Holden, M. (2008). A review of geometric transformations for nonrigid body registration. *Medical Imaging, IEEE Transactions on*, 27(1):111–128. 22, 29
- Igarashi, T., Moscovich, T., and Hughes, J. F. (2005). As-rigid-as-possible shape manipulation. *ACM Trans. Graph.*, 24(3):1134–1141. 77, 80

- Kallel, F. and Bertrand, M. (1996). Tissue elasticity reconstruction using linear perturbation method. *Medical Imaging, IEEE Transactions on*, 15(3):299–313. [4](#), [20](#), [36](#), [77](#), [81](#)
- Kauer, M., Vuskovic, V., Dual, J., Szekely, G., and Bajka, M. (2002). Inverse finite element characterization of soft tissues. *Medical Image Analysis*, 6(3):275–287. [27](#)
- Kaus, M. R., Brock, K. K., Pekar, V., Dawson, L. A., Nichol, A. M., and Jaffray, D. A. (2007). Assessment of a model-based deformable image registration approach for radiation therapy planning. *Int. J. Radiation Oncology Bio. Phys.*, 68(2):572–580. [23](#), [60](#), [82](#)
- Krouskop, T. A., Wheeler, T. M., Kallel, F., Garra, B. S., and Hall, T. (1998). Elastic moduli of breast and prostate tissues under compression. *Ultrasonic imaging(Print)*, 20(4):260–274. [2](#), [18](#), [37](#), [41](#), [46](#)
- Krysl, P., Lall, S., and Marsden, J. E. (2001). Dimensional model reduction in non-linear finite element dynamics of solids and structures. *International Journal for Numerical Methods in Engineering*, 51(4):479–504. [14](#), [82](#), [85](#), [86](#)
- Langen, K. M., Willoughby, T. R., Meeks, S. L., Santhanam, A., Cunningham, A., Levine, L., and Kupelian, P. A. (2008). Observations on real-time prostate gland motion using electromagnetic tracking. *International Journal of Radiation Oncology, Biology, Physics*, 71(4):1084–1090. PMID: 18280057. [12](#), [60](#)
- Lee, H., Foskey, M., Levy, J., Saboo, R., and Chaney, E. (2010a). Image estimation from marker locations for dose calculation in prostate radiation therapy. *Med Image Comput Comput Assist Interv*, 13(Pt 3):335–342. PMID: 20879417. [13](#)
- Lee, H.-P., Foskey, M., Niethammer, M., Krajcevski, P., and Lin, M. C. (2012). Simulation-based joint estimation of body deformation and elasticity parameters for medical image analysis. *IEEE Transactions on Medical Imaging*. [11](#)
- Lee, H.-P., Foskey, M., Niethammer, M., and Lin, M. (2010b). Physically-based deformable image registration with material property and boundary condition estimation. In *2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 532–535. IEEE. [12](#), [66](#)
- Lee, H.-P. and Lin, M. C. (2012). Fast optimization-based elasticity parameter estimation using reduced models. *The Visual Computer*, 28(6-8):553–562. [14](#), [16](#)
- Levy, J. H., Foskey, M., and Pizer, S. M. (2008). Rotational flows for interpolation between sampled surfaces. In *Computer Vision and Pattern Recognition Workshop*, volume 0, pages 1–8, Los Alamitos, CA, USA. IEEE Computer Society. [xiv](#), [23](#), [62](#), [64](#)
- Liang, J. and Yan, D. (2003). Reducing uncertainties in volumetric image based deformable organ registration. *Medical Physics*, 30:2116–2122. [23](#)

- Lindop, J. E., Treece, G. M., Gee, A. H., and Prager, R. W. (2006). 3D elastography using freehand ultrasound. *Ultrasound in Medicine & Biology*, 32(4):529–545. 19
- Litzenberg, D. W., Hadley, S. W., Tyagi, N., Balter, J. M., Ten Haken, R. K., and Chetty, I. J. (2007). Synchronized dynamic dose reconstruction. *Medical Physics*, 34(1):91–102. 59
- Maintz, J. and Viergever, M. A. (1998). A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36. 29, 59
- Manduca, A., Oliphant, T. E., Dresner, M. A., Mahowald, J. L., Kruse, S. A., Amromin, E., Felmlee, J. P., Greenleaf, J. F., and Ehman, R. L. (2001). Magnetic resonance elastography: Non-invasive mapping of tissue elasticity. *Medical Image Analysis*, 5(4):237–254. 20, 27
- Martin, S., Thomaszewski, B., Grinspun, E., and Gross, M. (2011). Example-based elastic materials. In *Proceedings of SIGGRAPH 2011*, volume 30, page 72:1–72:8. 79
- McNamara, A., Treuille, A., Popović, Z., and Stam, J. (2004). Fluid control using the adjoint method. *ACM Trans. Graph.*, 23(3):449–456. 6, 24, 77, 79
- Merck, D., Tracton, G., Saboo, R., Levy, J., Chaney, E., Pizer, S., and Joshi, S. (2008). Training models of anatomic shape variability. *Medical Physics*, 35(8):3584–3596. 62
- Moireau, P., Chapelle, D., and Tallec, P. L. (2008). Joint state and parameter estimation for distributed mechanical systems. *Computer Methods in Applied Mechanics and Engineering*, 197(6-8):659–677. 21
- Müller, M. and Gross, M. (2004). Interactive virtual materials. In *Proceedings of Graphics Interface 2004*, GI '04, page 239–246, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. Canadian Human-Computer Communications Society. 79, 83
- Muthupillai, R. and Ehman, R. L. (1996). Magnetic resonance elastography. *Nat Med*, 2(5):601–603. 81
- Nealen, A., Muller, M., Keiser, R., Boxerman, E., and Carlson, M. (2006). Physically based deformable models in computer graphics. *Computer Graphics Forum*, 25:809–836. 24, 77, 79
- Nocedal, J. and Wright, S. (1999). *Numerical Optimization*. Springer. 36, 37, 66
- Noel, C., Parikh, P. J., Roy, M., Kupelian, P., Mahadevan, A., Weinstein, G., Enke, C., Flores, N., Beyers, D., and Levine, L. (2009). Prediction of intrafraction prostate motion: Accuracy of pre- and Post-Treatment imaging and intermittent imaging. *International Journal of Radiation OncologyBiologyPhysics*, 73(3):692–698. 59

- Ophir, J., Alam, S., Garra, B., Kallel, F., Konofagou, E., Krouskop, T., and Varghese, T. (1999). Elastography: ultrasonic estimation and imaging of the elastic properties of tissues. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 213(3):203–233. **1, 19, 27, 77, 81**
- Pai, D. K., Doel, K. v. d., James, D. L., Lang, J., Lloyd, J. E., Richmond, J. L., and Yau, S. H. (2001). Scanning physical interaction behavior of 3D objects. In *Proceedings of SIGGRAPH 2001*, SIGGRAPH '01, page 87–96, New York, NY, USA. ACM. **79, 81**
- Pizer, S. M., Fletcher, P. T., Joshi, S., Gash, A. G., Stough, J., Thall, A., Tracton, G., and Chaney, E. L. (2005). A method and software for segmentation of anatomic object ensembles by deformable m-reps. *Medical Physics*, 32(5):1335–1345. **62**
- Pizer, S. M., Fletcher, P. T., Joshi, S., Thall, A., Chen, J. Z., Fridman, Y., Fritsch, D. S., Gash, A. G., Glotzer, J. M., Jiroutek, M. R., Lu, C., Muller, K. E., Tracton, G., Yushkevich, P., and Chaney, E. L. (2003). Deformable M-Reps for 3D medical image segmentation. *Int. J. Comput. Vision*, 55(2-3):85–106. **12, 23, 61, 62, 81**
- Press, W. H. (2007). *Numerical recipes*. Cambridge University Press. **25, 80**
- Risholm, P., Samset, E., and Wells, W. (2010). Bayesian estimation of deformation and elastic parameters in non-rigid registration. *Biomedical Image Registration*, page 104–115. **20**
- Rivaz, H., Boctor, E., Foroughi, P., Zellars, R., Fichtinger, G., and Hager, G. (2008). Ultrasound elastography: A dynamic programming approach. *Medical Imaging, IEEE Transactions on*, 27(10):1373–1377. **81**
- Rohr, K., Stiehl, H. S., Sprengel, R., Buzug, T. M., Weese, J., and Kuhn, M. H. (2001). Landmark-based elastic registration using approximating thin-plate splines. *Medical Imaging, IEEE Transactions on*, 20(6):526–534. **23, 60**
- Salcudean, S. E., French, D., Bachmann, S., Zahiri-Azar, R., Wen, X., and Morris, W. J. (2006). Viscoelasticity modeling of the prostate region using vibro-elastography. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, volume 4190, pages 389–396. Springer Berlin Heidelberg. **3, 20**
- Schnur, D. S. and Zabaras, N. (1992). An inverse method for determining elastic material properties and a material interface. *International Journal for Numerical Methods in Engineering*, 33(10):2039–2057. **81**
- Schreiber, E., Xu, Z., Lorenzen, A., Foskey, M., Cullip, T., Tracton, G., and Chaney, E. (2006). SU-FF-T-362: PPlanUNC as an Open-Source radiotherapy planning system for research and education. *Medical Physics*, 33(6):2129. **67**
- Sermesant, M., Moireau, P., Camara, O., Sainte-Marie, J., Andriantsimiavona, R., Cimirman, R., Hill, D. L. G., Chapelle, D., and Razavi, R. (2006). Cardiac function estimation from MRI using a heart model and data assimilation: Advances and difficulties. *Medical Image Analysis*, 10(4):642–656. **21**

- Shen, D. and Davatzikos, C. (2002). HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Transactions on Medical Imaging*, 21(11):1421–1439. 5, 23
- Shewchuk, J. R. (1996). Triangle: Engineering a 2D quality mesh generator and delaunay triangulator. In *Applied Computational Geometry Towards Geometric Engineering*, volume 1148, pages 203–222. Springer-Verlag, Berlin/Heidelberg.
- Si, H. (2009). TetGen: a quality tetrahedral mesh generator and three-dimensional delaunay triangulator. 40
- Skovoroda, A. and Emelianov, S. (1995). Tissue elasticity reconstruction based on ultrasonic displacement and strain images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):141. 81
- Sobin, L. H. (2009). *TNM Classification of Malignant Tumours*. John Wiley and Sons. 46
- Sotiras, A., Christos, D., and Paragios, N. (2012). Deformable medical image registration: A survey. Rapport de recherche RR-7919, INRIA. 22
- Stam, J. (1999). Stable fluids. In *Proceedings of SIGGRAPH 99*, pages 121–128. ACM Press/Addison-Wesley Publishing Co. 24
- Syllebranque, C. and Boivin, S. (2008). Estimation of mechanical parameters of deformable solids from videos. *The Visual Computer*, 24(11):963–972. 6, 25, 27, 80, 81
- Tanner, C., Schnabel, J. A., Hill, D. L. G., Hawkes, D. J., Leach, M. O., and Hose, D. R. (2006). Factors influencing the accuracy of biomechanical breast models. *Medical Physics*, 33(6):1758. 39, 56
- Taylor, Z. A., Crozier, S., and Ourselin, S. (2010). Real-time surgical simulation using reduced order finite element analysis. *MICCAI 2010*, 13(Pt 2):388–395. PMID: 20879339. 82
- Teschner, M., Heidelberger, B., Manocha, D., Govindaraju, N., Zachmann, G., Kimmerle, S., Mezger, J., and Fuhrmann, A. (2005). Collision handling in dynamic simulation environments: The evolution of graphics: Where to next? In *Eurographics 2005 Tutorial*. 24, 79
- Thirion, J.-P. (1998). Image matching as a diffusion process: an analogy with maxwell’s demons. *Medical Image Analysis*, 2(3):243–260. 5, 22, 51
- Treuille, A., McNamara, A., Popović, Z., and Stam, J. (2003). Keyframe control of smoke simulations. In *ACM SIGGRAPH 2003 Papers*, pages 716–723, San Diego, California. ACM. 6, 24, 77, 79

- Umut Ozcan, M., Ocal, S., Basdogan, C., Dogusoy, G., and Tokat, Y. (2011). Characterization of frequency-dependent material properties of human liver and its pathologies using an impact hammer. *Medical Image Analysis*, 15(1):45–52. **2, 18**
- Washington, C. W. and Miga, M. I. (2004). Modality independent elastography (MIE): a new approach to elasticity imaging. *IEEE Trans. on Medical Image*, 23(9):1117–1128. **20, 81**
- Weng, Y., Xu, W., Wu, Y., Zhou, K., and Guo, B. (2006). 2D shape deformation using nonlinear least squares optimization. *The Visual Computer*, 22(9-11):653–660. **77, 80**
- Wittek, A., Miller, K., Kikinis, R., and Warfield, S. K. (2007). Patient-specific model of brain deformation: Application to medical image registration. *Journal of Biomechanics*, 40(4):919–929. **29**
- Wong, J. R., Grimm, L., Uematsu, M., Oren, R., Cheng, C. W., Merrick, S., and Schiff, P. (2005). Image-guided radiotherapy for prostate cancer by CT-linear accelerator combination: Prostate movements and dosimetric considerations. *International Journal of Radiation Oncology*Biology*Physics*, 61(2):561–569. **59**
- Wu, Q., Mohan, R., Niemierko, A., and Schmidt-Ullrich, R. (2002). Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *International Journal of Radiation Oncology, Biology, Physics*, 52(1):224–235. PMID: 11777642. **71**
- Yan, D., Lockman, D., Brabbins, D., Tyburski, L., and Martinez, A. (2000). An off-line strategy for constructing a patient-specific planning target volume in adaptive treatment process for prostate cancer. *International Journal of Radiation Oncology*Biology*Physics*, 48(1):289–302. **59**
- Yang, W., Feng, J., and Jin, X. (2008). Shape deformation with tunable stiffness. *The Visual Computer*, 24(7-9):495–503. **77, 80**
- Yoo, T. S., Ackerman, M. J., Lorensen, W. E., Schroeder, W., Chalana, V., Aylward, S., Metaxas, D., and Whitaker, R. (2002). Engineering and algorithm design for an image processing api: a technical report on ITK—the insight toolkit. *Studies in health technology and informatics*, 85:586–592. PMID: 15458157. **40**
- Yushkevich, P. A., Piven, J., Hazlett, H. C., Smith, R. G., Ho, S., Gee, J. C., and Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *NeuroImage*, 31(3):1116–1128. **40**
- Zhai, L., Polascik, T. J., Foo, W.-C., Rosenzweig, S., Palmeri, M. L., Madden, J., and Nightingale, K. R. (2012). Acoustic radiation force impulse imaging of human prostates: Initial in vivo demonstration. *Ultrasound in Medicine & Biology*, 38(1):50–61.

- Zhang, M., Nigwekar, P., Castaneda, B., Hoyt, K., Joseph, J. V., Sant'Agnes, A. d., Messing, E. M., Strang, J. G., Rubens, D. J., and Parker, K. J. (2008). Quantitative characterization of viscoelastic properties of human prostate correlated with histology. *Ultrasound in Medicine & Biology*. 18
- Zhu, Y., Hall, T., and Jiang, J. (2003). A finite-element approach for young's modulus reconstruction. *Medical Imaging, IEEE Transactions on*, 22(7):890–901. 4, 20, 77, 81
- Zienkiewicz, O. C. and Taylor, R. L. (2005). *The Finite Element Method*. Butterworth-Heinemann, Oxford, 6 edition. 3, 27, 61, 82