

Simulation-Based Joint Estimation of Body Deformation and Elasticity Parameters

Huai-Ping Lee¹, Mark Foskey^{1,2}, Marc Niethammer¹, and Ming C. Lin¹

¹Department of Computer Science,

²Department of Radiation Oncology, University of North Carolina at Chapel Hill

Abstract

Physically-based elastic body simulation has been applied to medical image registration to enforce physical constraints, since each organ in the image can be viewed as an elastic body. However, it requires hand adjustment for the material properties to achieve the desired output shape, and the boundary conditions (forces or displacements on boundary nodes) cannot be acquired directly from the images. We present a general method for estimating these parameters for a simulator 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. Besides matching the surface boundaries, our method guarantees the quality of the deformation fields, as they are computed using a physics-based simulation. Compared to other FEM-based methods, our technique can determine the elastic properties automatically and does not rely on explicit surface matching for generating boundary conditions.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically-Based Modeling I.3.8 [Computer Graphics]: Applications—Image Registration

1. Introduction

Physically-based simulations can help generate realistic scenes or animations without requiring low-level control of a 3D model. To achieve a particular appearance, however, requires many iterations of adjusting simulation parameters, simulating, and assessing the results. For a large number of parameters and a complex, high-cost simulation, this adjustment process becomes very tedious, making parameter estimation a topic of interest in the field of computer graphics [BTH*03, SB08]. Physically-based simulations can also be used for 3D *image registration*, which is the process of finding correspondence between two images. In medical imaging applications, the images are taken using modalities such as computed tomography (CT) or magnetic resonance imaging (MRI), and human organs are the objects of interest. An application of particular interest arises in radiation therapy for cancer, in which treatment is delivered daily over a series of weeks. With the newest technology, patients can be imaged periodically in the treatment position. If the resulting

images can be brought into point-by-point correspondence, the radiation dose delivered to different parts of the tissue can be brought into a common reference frame and accumulated. Because human organs consist of elastically deforming tissues, a nonrigid registration is needed in addition to a rigid registration (which finds a transformation matrix to align fixed parts such as the bones).

When registering multiple images from the same patient, it makes sense to use an elastic model reflecting actual physical properties of the tissue. Provided with some information about correspondence, such as the locations of selected organ boundaries (an *image segmentation*), such a physical simulation can deform the moving image toward the fixed image, which constitutes a nonrigid registration through the implied correspondence [FWG*99, FWN*00, LY03, HMC*07]. However, there is no direct way to measure the material properties (the *modulus of elasticity*, or *Young's modulus*, which is a measure of stiffness of the material, and *Poisson's ratio*, which is a measure of compressibility)

and boundary forces acting on the organs within the patient. Although some reference values of material properties have been reported in the literature [KWK*98, ZNC*08], the values may vary significantly between patients. Surface matching algorithms can provide approximate boundary conditions in place of boundary forces [FWN*00, HMC*07], but we do not know if the surface matching satisfies the physical constraints, since they are not used in the surface matching algorithms and the solution to the surface matching problem is not unique.

In this paper we describe a physically-based registration method that estimates these material properties as part of the registration process. We minimize an objective function 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. Our method avoids the tedious manual parameter adjustment process and reduces the implicit error induced by surface matching methods which do not respect physical constraints. The general optimization framework is suitable for different physical models. Our technique is based on the surfaces of the models at the rest state and the deformed state, and image intensity information is not required after the segmentation boundaries are acquired. Therefore the method can be applied to areas other than the image registration problem.

Our method improves over previous simulation-based registration methods by providing an automatic means of finding the parameters that are missing in the images. Furthermore, compared to other work on simulation parameter estimation, our method does not require boundary conditions and can therefore avoid the process of measuring the external forces.

We discuss related work in Section 2 and explain the elastic model and the optimization scheme in Section 3, followed by experimental results using a synthetic scene and two pairs of real CT images in Section 4 to demonstrate the feasibility of our method. We conclude with a summary and discussion of future work.

2. Related 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 [NMK*06, THM*05]. Some work in the computer graphics community has explored parameter estimation for deformable model simulation. For example, an optimization scheme has been proposed to estimate cloth simulation parameters [BTH*03]. 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. 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 [Pre07]. Syllebranque and Boivin [SB08] 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. While these methods depend on the rendering and computer vision algorithms, our technique directly uses the surfaces of the deformed bodies to compute the error metric. In addition, the boundary conditions are unknown in our problem.

Image registration is an integral part of image-guided radiotherapy. In order to assess the progress of treatment over multiple daily radiation fractions, it is important to be able to trace the motion of points in the tissue with respect to the radiation beams. Given a 3D moving image, generally taken at the time of each treatment, and a 3D reference image, typically taken initially to plan the treatment course, the goal is to find the optimal correspondence from the moving image to the reference image. A traditional nonrigid image registration method treats the deformation field as parameters and minimizes some image-based metric [Thi98, SD02]. Because the optimization occurs in a high-dimensional parameter space, a physically-based energy function, such as the compressible fluid model [CRM96, FDG*05] or the linear elasticity model [RRM83], can be used in the minimization process to impose regularity on the deformation field. Basis functions such as radial basis functions or B-splines can also be used with methods that minimize the error in landmark positions [RSS*01] or in surface tessellation [KBP*07] to reduce the dimensionality. In these optimization-based methods, the objective function to be minimized consists of terms measuring image similarity and terms measuring some physical energy, and a coefficient needs to be chosen for each term carefully to balance their effects. The balancing is essentially a trade-off between image similarity and physical constraints. A survey of these methods can be found in [Hol08]. Unlike the methods using physically-based regularizations, simulation-based methods [FWG*99, FWN*00, LY03, HMC*07] enforce the physical constraints exactly and take different material properties into account. Boundary conditions are usually obtained from the gradient of the image [FWG*99] or from surface matching methods [FWN*00, LY03, HMC*07], but they may induce uncertainty to the simulation since the surface matching is not unique. Material properties, on the other hand, are adjusted by hand. Recently, an optimization-based 2D registration algorithm has been proposed [AGP*06], where the boundary conditions and the material properties are optimized with a gradient descent method, but it is applied only to 2D images with a low-resolution triangle mesh. Unlike previous methods, we simultaneously estimate body deformation and elasticity parameters based on a full 3D lin-

ear elasticity simulator. Physical constraints are enforced with simulations and the material properties are recovered through optimization. Our method uses surface matching results only as a means of generating the initial guess of forces and does not fix any boundary node of elastic bodies during the optimization.

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 elastography, is a non-invasive method to acquire strain or elasticity images of soft tissues [SE95]. 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. This estimate may be performed using cross-correlation analysis [OAG*99] or using dynamic programming to minimize similarity of echo amplitude and displacement discontinuity [RBF*08]. Alternatively, a modified MRI machine in tune with a mechanical vibration of tissues can be used to find the displacement field [ME96]. Once the deformation field is known, the material properties can be found by solving a least-squares problem [ZHJ03, BT07], if the physical model is linear. Another type of methods use iterative numerical optimization algorithms to minimize the error in the deformation field [KB96, BCF08]. Although these iterative methods are slower than directly solving the inverse problem, they do not require linearity of the underlying model and are therefore suitable for any physical model. Our method also uses an optimization scheme to find out the elasticity of an organ, but our problem is different from the elastography in several ways. Firstly, image registration is done for a multi-organ system, assuming the material property is constant within an organ, while elastography is used to find out different material properties within a tissue. Secondly, the deformation field and external forces are considered known in elastography, while in image registration, deformation cannot be measured directly and is therefore inferred from image information such as edges or corners.

3. 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.

Our current implementation of the simulator uses the isotropic linear elasticity model because it is widely used in the medical image registration of the prostate. We consider just the elasticity value (Young's modulus) in this paper due to the simplicity in computation and its importance in non-invasive cancer detection techniques.

The inputs to the registration problem are two segmented images: the fixed image with segmentation \mathbf{S}_f and the

moving image with segmentation \mathbf{S}_m . The bones are already aligned using a rigid registration method described in [FDG*05], and we focus on the nonrigid registration problem. Each segmentation is represented as a set of closed triangulated surfaces, one for each segmented object. We 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. Our registration algorithm 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 fields and image 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 image registration problem. We adopt the linear elasticity model with a finite-element model. The flow chart of our algorithm is shown in Fig. 1 and will be explained in detail in this section.

3.1. Linear Elasticity Model and Finite Element Modeling

In the optimization loop, the deformation field $\mathbf{u} = [u, v, w]^T$ is always generated by a physically-based simulation, where the finite element method (FEM) is used to solve the constitutive equations of the linear elasticity model,

$$\begin{aligned} \frac{\partial \sigma_x}{\partial x} + \frac{\partial \tau_{xy}}{\partial y} + \frac{\partial \tau_{xz}}{\partial z} + f_x &= 0 \\ \frac{\partial \tau_{xy}}{\partial x} + \frac{\partial \sigma_y}{\partial y} + \frac{\partial \tau_{yz}}{\partial z} + f_y &= 0 \\ \frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{yz}}{\partial y} + \frac{\partial \sigma_z}{\partial z} + f_z &= 0, \end{aligned} \quad (1)$$

where $\boldsymbol{\sigma} = [\sigma_x, \sigma_y, \sigma_z, \tau_{xy}, \tau_{yz}, \tau_{xz}]^T$ is the stress vector, which is a linear transformation of the strain vector $\boldsymbol{\varepsilon}$,

$$\boldsymbol{\varepsilon} = \left[\frac{\partial u}{\partial x}, \frac{\partial v}{\partial y}, \frac{\partial w}{\partial z}, \frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}, \frac{\partial v}{\partial z} + \frac{\partial w}{\partial y}, \frac{\partial u}{\partial z} + \frac{\partial w}{\partial x} \right]^T = \mathbf{L}\mathbf{u}, \quad (2)$$

where \mathbf{L} is the 6×3 matrix consists of the partial differential operators acting on the deformation \mathbf{u} . Assuming linear elasticity, we can write $\boldsymbol{\sigma} = \mathbf{D}\boldsymbol{\varepsilon}$, where \mathbf{D} is a matrix defined by the material properties (assuming isotropic material, the properties are Young's modulus and Poisson's ratio) of the body. To solve Eq. 1 numerically, we approximate the derivatives of the deformation with the finite element method (FEM) [ZT05], where the domain is subdivided into a set of elements, and each element consists of several nodes. Fig 5a shows the finite element model used in one of our 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

$$\hat{\mathbf{u}}^{el}(\mathbf{p}) = \sum_{j=1}^4 \mathbf{u}_j^{el} N_j^{el}(\mathbf{p}), \quad (3)$$

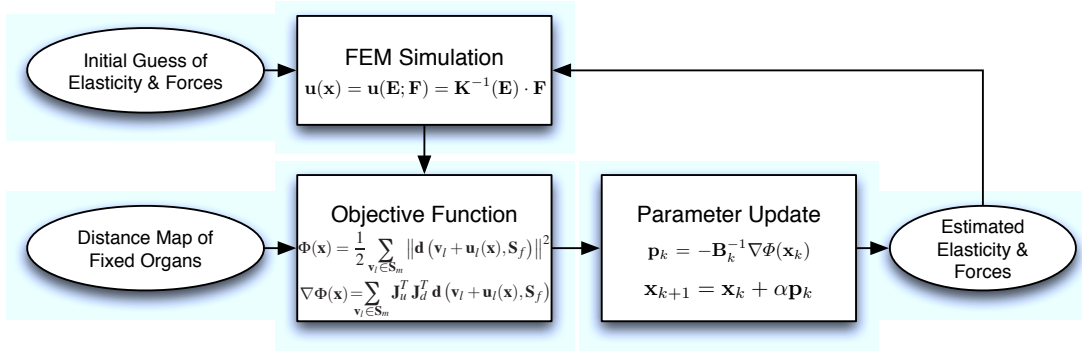


Figure 1: Flow chart of the optimization loop; the deformation field generated by the simulator is used in the objective function to update the parameters, which are fed back into the simulator, and so on.

where \mathbf{u}_j^{el} is the deformation of the j -th node of the element, and N_j^{el} 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

$$\mathbf{K}\mathbf{u} = \mathbf{F}, \quad (4)$$

where \mathbf{K} is called the stiffness matrix, which depends on the material properties \mathbf{E} 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 we 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 boundary nodes, including boundaries between two materials, are assigned displacement values, the simulation is essentially an interpolation of displacement field from surface matching results, and the elasticity values only affects internal nodes, and we do not know their target positions. Therefore the elasticity cannot be recovered. Instead, we only assign boundary conditions to a part of the boundary nodes, and other surface nodes without boundary conditions will be affected by the relative elasticities. For example, in a simulation of the male pelvis area, 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 Dirichlet 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, as proposed in [HMC*07].

3.2. Sensitivity Study

Since our method is based on the assumption that the deformed surface depends on both the elasticity and the exter-




Figure 2: A sliced view of the synthetic scene, which consists of two concentric spheres; the inner and outer regions have different stiffness values.

nal forces, we 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. 2. We fix the elasticity of the outer region and alter the elasticity of the inner sphere, since only the ratio of the two elasticity values matters. A force with a specified magnitude pointing towards the center of the spheres is applied on each node of the outer surface, and no boundary conditions (zero external forces) are applied on the inner surface. Several simulations using different elasticities of the inner region and force magnitudes are performed, and the plots of the inner sphere radius versus the elasticity value and versus force magnitude are shown in Fig. 3.

In Fig. 3a, the slope is much higher when the elasticity is low for each curve, which indicates that the shape of the inner sphere is much more sensitive to the elasticity when the elasticity value is low. In Fig. 3b, the magnitude of the slope

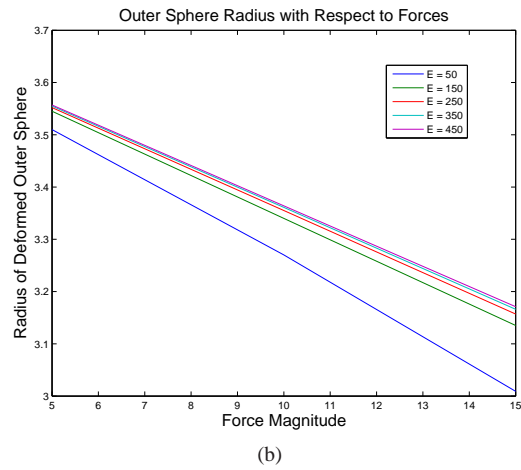
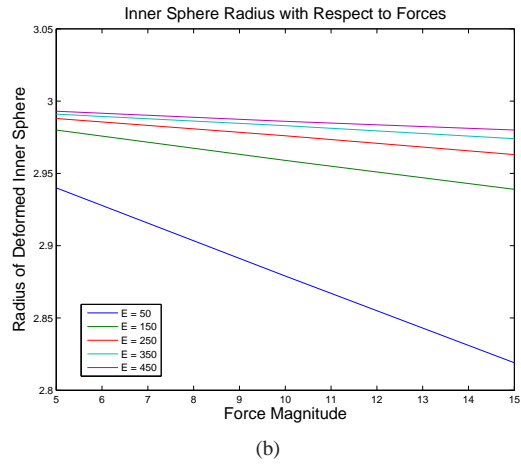
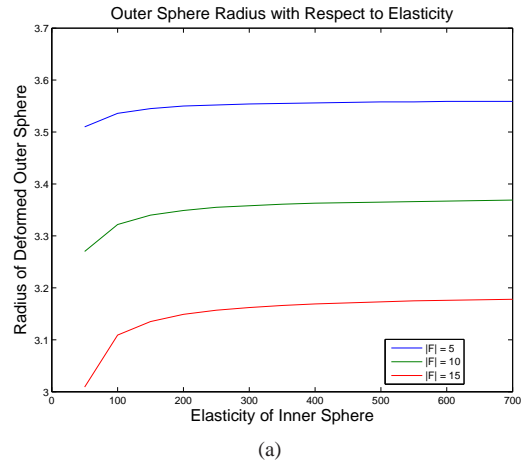
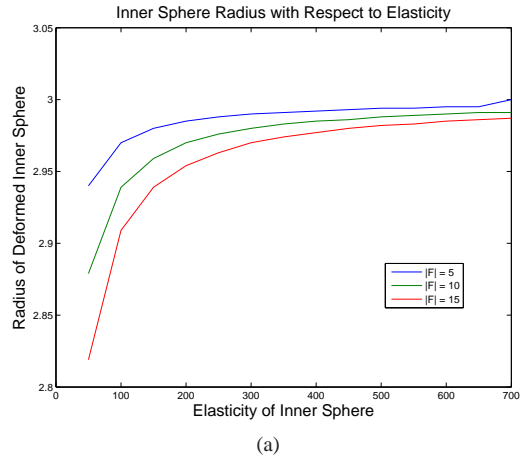


Figure 3: The plots of the radius of the inner sphere after deformation: (a) radius versus the elasticity value of the inner region; (b) radius versus the magnitude of forces acting on the outer surface. The radius before deformation is 3, and the elasticity ratio between the inner and outer region is 10. The Poisson's ratios are fixed to 0.40 and 0.35 for the two regions, respectively.

is higher for the lines with lower elasticity values, which indicates that the shape is more sensitive to the forces when the elasticity is low. Similar plots for the radius of the outer sphere are shown in Fig. 4, which also indicates that the outer radius is more sensitive to the parameters when the elasticity of the inner region is lower.

These results suggest that our ability to recover the parameters is limited by how stiff the object is. With a very high stiffness, the shape of the object becomes insensitive to the parameters. In this case, the shape can still be recovered, but the resulting parameters may not be accurate.

Figure 4: The plots of the radius of the outer sphere after deformation: (a) radius versus the elasticity value of the inner region; (b) radius versus the magnitude of forces acting on the outer surface. The radius before deformation is 3.75.

3.3. 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, and we consider only Young's moduli in this paper, 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 reference images,

$$\Phi(\mathbf{x}) = \frac{1}{2} \sum_{\mathbf{v}_l \in \mathbf{S}_m} \|\mathbf{d}(\mathbf{v}_l + \mathbf{u}_l(\mathbf{x}), \mathbf{S}_f)\|^2. \quad (5)$$

Here $\mathbf{u}(\mathbf{x})$ is the deformation field computed by the simulator with parameters \mathbf{x} , interpreted as a displacement vector for each 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 chain rule,

$$\begin{aligned}\nabla\Phi(\mathbf{x}) &= \sum_{\mathbf{v}_l \in \mathcal{S}_m} \left[\frac{\partial \mathbf{u}_l}{\partial \mathbf{x}} \right] \left[\frac{\partial \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l, \mathbf{S}_f)}{\partial \mathbf{u}_l} \right] \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l(\mathbf{x}), \mathbf{S}_f) \\ &= \sum_{\mathbf{v}_l \in \mathcal{S}_m} \mathbf{J}_u^T \mathbf{J}_d^T \mathbf{d}(\mathbf{v}_l + \mathbf{u}_l(\mathbf{x}), \mathbf{S}_f),\end{aligned}\quad (6)$$

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 we 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 essentially 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 Eq. 4,

$$\left[\frac{\partial \mathbf{K}}{\partial E_j} \right] \mathbf{u} + \mathbf{K} \left\{ \frac{\partial \mathbf{u}}{\partial E_j} \right\} = 0, \quad (7)$$

Therefore we 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 Eq. 4, we have

$$\left[\frac{\partial \mathbf{K}}{\partial F_j} \right] \mathbf{u} + \mathbf{K} \left\{ \frac{\partial \mathbf{u}}{\partial F_j} \right\} = \mathbf{e}_j, \quad (8)$$

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 we 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. 5b shows one of the distance maps used in our experiments. Notice that the physical model can be different, as long as the derivatives $\partial u_i / \partial x_j$ can be computed.

In the experiments, however, we 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 to the forces, $\|\partial\Phi/\partial\mathbf{F}\|$, and it caused the material property to converge very slowly. To obtain a faster convergence of \mathbf{E} , we use an alternating approach: first fix \mathbf{E} and optimize \mathbf{F} for a few steps (five steps in our experiments), then fix \mathbf{F} to optimize \mathbf{E} for a few steps, and then return to optimizing \mathbf{F} with the new \mathbf{E} value fixed, and so on.

3.4. Numerical Optimization

We use a line search scheme for optimization: in each iteration k , first find a descent direction \mathbf{p}_k , find an optimal step




Figure 5: Input to our algorithm: (a) a sliced view of the tetrahedral model of the moving image; (b) a slice of the vector distance map of the prostate in the reference image.

size α in the 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 with the Newton's method solving 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 [KB96], but the Hessian matrices can only be approximated and are usually ill-conditioned. Alternatively, the Hessians can be approximated with the BFGS formula [NW99],

$$\mathbf{B}_{k+1} = \mathbf{B}_k + \frac{\mathbf{y}_k \mathbf{y}_k^T}{\mathbf{y}_k^T \mathbf{s}_k} - \frac{\mathbf{B}_k \mathbf{s}_k (\mathbf{B}_k \mathbf{s}_k)^T}{\mathbf{s}_k^T \mathbf{B}_k \mathbf{s}_k}, \quad (9)$$

where $\mathbf{y}_k = \nabla\Phi(\mathbf{x}_{k+1}) - \nabla\Phi(\mathbf{x}_k)$ and $\mathbf{s}_k = \mathbf{x}_{k+1} - \mathbf{x}_k$. The main idea of the BFGS formula is to approximate the Hessian with the secant equation, $\mathbf{B}_{k+1} \mathbf{s}_k = \mathbf{y}_k$, while keeping \mathbf{B}_{k+1} symmetric and close to \mathbf{B}_k . 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 our case, the number of parameters can be up to thousands, therefore we use a limited-memory quasi-Newton method known as the L-BFGS method [NW99], combined with the line search algorithm proposed in [MT94]. Instead of storing the entire $3N_n \times 3N_n$ Hessian matrix, the L-BFGS method saves only m most recent vectors pairs $\{\mathbf{s}_i, \mathbf{y}_i\}$, $i = k-m, \dots, k-1$ and construct the approximate \mathbf{B}_{k+1} by applying Eq. 9 m times. Intuitively, the L-BFGS method reduces the memory usage by storing curvature information of the objective function from only the most recent steps.

3.5. Initial Guess of Parameters

A good initial guess can prevent the optimizer from getting stuck in a local minimum. In the case of medical image registration, the initial guess of the elasticity is chosen based on the knowledge of the simulated organs. Our example images involve two materials: the prostate and the surrounding tissue. There have been *in vitro* experiments on the prostate using different elasticity models. Krouskop et al. [KWK*98] reported the elastic modulus of 40-80 kPa for

normal prostate tissue and 80-270 kPa for cancerous tissue. They also reported 10-30 kPa for the breast fat tissue. Based on these numbers, we chose the values 100 for the prostate and 10 for the surrounding tissue as the initial guess of elasticity, since all the patients for our example application have prostate cancer, and only the ratio between the two elasticity values matters.

Initial guess of forces are 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 to compute the corresponding external forces, which is used as the initial guess, using Eq. 4.

To further reduce the risk of wrong initial guess of the elasticity, we employ a two-trial scheme: we run the optimization once, and use the resulting elasticity values as the initial guess for the second trial of optimization, while the initial guess of forces is generated in the same way as in the first trial. This scheme has helped to reduce errors in elasticity in our experiments.

4. Experiments

We tested our algorithm on two sets of surface data. Firstly, we tested the accuracy of the optimization scheme using synthetic target surfaces generated by forward simulations. We then applied the technique to a medical image registration problem to show the physical accuracy with real data. Given the moving surfaces in the form of triangle meshes, the tetrahedral model for the entire domain is built with the software TetGen [Si], and the library ITK [Ins] is used to compute the vector distance maps of the target surface. The FEM simulator uses the linear algebra library PETSc [Mat].

4.1. Synthetic Target Surface

To verify the accuracy of our method, we test the algorithm with synthetic target surfaces generated by a forward simulation, so that we know the true elasticity values. We use the same base model as used in the sensitivity study (Fig. 2) and apply inward Dirichlet boundary conditions of magnitude 0.75 cm on the outer sphere. The initial guesses and resulting errors for each elasticity value, along with the total number of optimization iterations, are shown in Table 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 parameter space that reduces the value of the objective function. The results showed that our method recovers the elasticity value to under 7% of error when the value is under 400, which corresponds to a ratio of 40 between elasticities of the inner and outer region, and the relative error is within 3.5% for elasticity values below 300. The result of higher errors for higher elasticity values

is consistent with the sensitivity study presented in Section 3.2, where we showed that the sensitivity of the radii of the spheres with respect to the elasticity is considerably lower when the elasticity of the inner region is high.

4.2. Segmented CT Images

We experimented on two pairs of 3D CT images of the male pelvis area with the main structures (bladder, prostate, rectum, and bones) segmented manually, and each image pair is from one patient on different days. The bladder and the rectum are made hollow to reflect the actual structure, and the bones are fixed during the simulations, as shown in Fig. 5a. Since only the relative values of material properties can be recovered, we fix the Young's modulus of the surrounding tissues (the region outside all organs and bones) to 10 and optimize that of the prostate. Since the prostate is the main organ of interest, we 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.

The optimization process is terminated when $\|\partial\Phi/\partial\mathbf{E}\| < 10^{-7} \|\mathbf{E}\|$ and $\|\partial\Phi/\partial\mathbf{F}\| < 10^{-3} \|\mathbf{F}\|$, or when the optimizer cannot find a direction in the parameter space that reduces the value of the objective function. The convergence graphs (plots of Φ and $\|\nabla\Phi\|$ versus iteration number) for the material property and for boundary forces from the first trial of the first patient are shown in Fig. 6 (convergence graphs for other trials are similar). Note that the optimization is done in the alternating way, and the convergence graphs shown are the result of concatenating the steps for optimizing \mathbf{F} and for optimizing \mathbf{E} separately.

The final elasticity ratio between the prostate and the surrounding tissue is 12.08 for the first patient and 35.63 for the second patient. From the elasticity values of the prostate and breast fat tissues reported in the literature [KWK*98], we have a speculated range of elasticity ratio, and the recovered value for first patient is within the range while that of the second patient is above the range. The final values of the objective function are 1.53 and 1.12, corresponding to an RMS error of 0.076 cm and 0.065 cm, respectively, for the two patients. The RMS errors are within the image resolution, $0.1 \times 0.1 \times 0.3$ cm. The moving images of the first patient before and after registration are shown in Fig. 8, 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. 7 shows a 3D close-up view of the deforming surfaces for the second patient, where the yellow surfaces denote the boundaries with external forces applied, and the target and moving surfaces of the prostate are shown in red and blue, respectively. We also compared our

True Values of Elasticity	50	100	150	200	250	300	350	400	450
Initial Guess of Elasticity	100	200	200	300	300	400	400	500	500
Recovered Elasticity	51.76	98.89	146.98	202.14	243.16	292.07	330.81	373.47	369.07
Relative Error of Elasticity	+3.5%	-1.1%	-2.0%	+1.1%	-2.7%	-2.6%	-5.5%	-6.6%	-18.0%
Number of Optimization Iterations	100	33	11	8	5	6	4	5	4

Table 1: Error in recovered modulus of elasticity in the synthetic scene




Figure 7: Close-up view of the deforming surfaces; the red surface shown in wireframe is the target surface of the prostate. Notice how the blue surface move towards the red surface. Yellow surfaces are those with external forces applied.

registration results with a popular image-based approach, the Demons method [Thi98], by looking at some landmarks inside the prostate. In most cases, the image intensity is almost constant inside an organ, but a small number of patients have some calcium accumulated in the prostate, resulting in bright spots of diameters about 0.2 – 0.6 cm 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 2. The estimated diameters of the spots are 0.25 cm in the first patient and 0.6 cm in the second patient. 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. Our method enforces physically-based constraints and results in errors within the diameter of the spot. Notice that we did not remove the bright spots in the images for the Demons method, so the intensity and gradient information from the landmarks is also utilized in the image-based registration, while our method is based purely on the physics-based simulation and does not take advantages of the landmarks, which are not always present in most of the other applications. For example, the spot in the second pair of CT images is considerably larger and brighter, therefore the intensity and gradient information is much stronger, enabling the Demons method to register well very quickly. Without the bright spots, the Demons method cannot guarantee the integrity of the deformation field.

5. Conclusion and Discussion

We have presented a novel physically-based method for simultaneously estimating the deformation of soft bodies and determining the unknown material properties and boundary

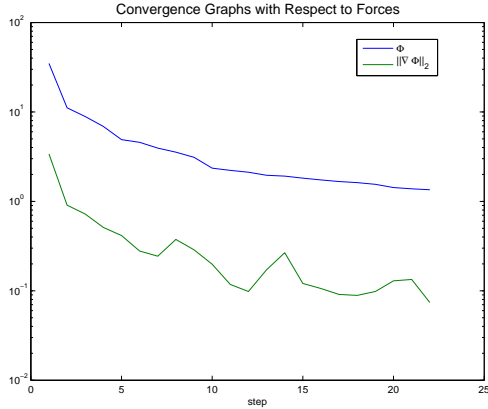
conditions. Our method is especially well suited for 3D medical image registration using a pair of scanned images to automatically determine the patient-specific material properties during the registration. The resulting deformation field is enforced to be physically-plausible, since it is computed by the 3D FEM simulator with appropriate contact constraints among organs. The observed error on the boundary due to our physically-based image registration method is within the resolution of segmented images, i.e. no larger than the segmentation error, and the error on the internal bright spots as landmarks in the prostate is comparable to the diameter of the spots.

The optimization framework is general. It is not limited to deformable image registration and could be used for more sophisticated physiological models than the linear elasticity model we chose for simplicity in our current approach. As an image registration technique, our method is reliable in terms of the registration error; as a parameter estimation method, our system can save the enormous amount of work adjusting the simulation parameters by hand. Furthermore, since only the 3D surfaces are used in our algorithm, applications other than image registration could also adopt our algorithm.

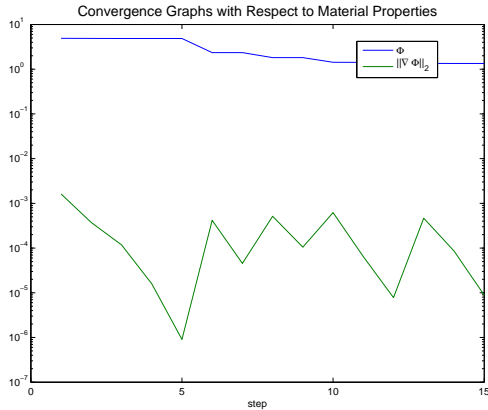
In the future, we would like to further improve the performance of the iterative scheme and apply it to more complicated physical and geometric models, such as the situations with complex material property distributions [NKJF09]. For the application in image registration, we would like to define the objective function directly with the information in the images, such as the intensities and local features, thereby further reducing the amount of work and the human error induced by hand segmentation in the image registration process. Virtual surgery [FLA*05, SHG*06, WBG07, MTG03] is another area that needs the estimation of material proper-

	CT Images from Patient 1					CT Images from Patient 2
	Landmark 1	Landmark 2	Landmark 3	Landmark 4	Average	Landmark 1
Rigid Registration	0.5594	0.5580	0.5180	0.3585	0.4985	0.3336
Demons	0.2793	0.2791	0.2036	0.2033	0.2413	0.0807
Our Method	0.1971	0.2830	0.2294	0.0513	0.1902	0.2865

Table 2: Error of landmark positions (distance in cm) inside the prostate



(a)




(b)

Figure 6: Convergence graphs (plot of Φ and $\|\nabla\Phi\|$ versus iteration number) and optimization history of the elasticity for the first CT image data: (a) convergence of the external forces; (b) convergence of the elasticity.


ties and may be combined with our framework. We would also like to explore the possibility of applying our technique to control the animation of general deformable bodies in graphics applications.

References

[AGP*06] ALTEROVITZ R., GOLDBERG K., POULIOT J., HSU I., KIM Y., NOWOROLSKI S., KURHANOWICZ J.: Registration of MR prostate images with biomechanical modeling and nonlinear parameter estimation. *Medical Physics* 33 (2006), 446.
 [BCF08] BALOCCO S., CAMARA O., FRANGI A. F.: Towards



(a)



(b)

Figure 8: Registration results for the first pair of test images: (a) axial and sagittal views of the moving image before registration; (b) the two views of the registered image; superimposed by segmentations of the reference image, shown in red, and the segmentation of the prostate in the moving image, shown in blue; notice that the image deforms towards the red contours.

regional elastography of intracranial aneurysms. In *Medical Image Computing and Computer-Assisted Intervention* (2008), vol. 11, pp. 131–8. PMID: 18982598.

[BT07] BECKER M., TESCHNER M.: Robust and efficient estimation of elasticity parameters using the linear finite element method. *Proc. of Simulation and Visualization* (2007), 15–28.
 [BTH*03] BHAT K. S., TWIGG C. D., HODGINS J. K., KHOSLA P. K., POPOVIC Z., SEITZ S. M.: Estimating cloth simulation parameters from video. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2003), Eurographics Association Aire-la-Ville, Switzerland, Switzerland, pp. 37–51.
 [CRM96] CHRISTENSEN G., RABBITT R. D., MILLER M. I.: Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.* 5, 10 (1996).
 [FDG*05] FOSKEY M., DAVIS B., GOYAL L., CHANG S., CHANEY E., STREHL N., TOMEI S., ROSENMAN J., JOSHI S.: Large deformation three-dimensional image registration in

- image-guided radiation therapy. *Physics in Medicine and Biology* 50 (2005), 5869–5892.
- [FLA*05] FRANCE L., LENOIR J., ANGELIDIS A., MESEURE P., CANI M. P., FAURE F., CHAILLOU C.: A layered model of a virtual human intestine for surgery simulation. *Medical image analysis* 9, 2 (2005), 123–132.
- [FWG*99] FERRANT M., WARFIELD S. K., GUTTMANN C. R. G., MULKERN R. V., JOLESZ F. A., KIKINIS R.: 3d image matching using a finite element based elastic deformation model. In *MICCAI* (1999), pp. 202–209.
- [FWN*00] FERRANT M., WARFIELD S. K., NABAVI A., JOLESZ F. A., KIKINIS R.: Registration of 3d intraoperative MR images of the brain using a finite element biomechanical model. In *MICCAI* (2000), pp. 19–28.
- [HMC*07] HENSEL J. M., MENARD C., CHUNG P. W. M., MILOSEVIC M. F., KIRILOVA A., MOSELEY J. L., HAIDER M. A., BROCK K. K.: 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 (2007), 1522–1528.
- [Hol08] HOLDEN M.: A review of geometric transformations for nonrigid body registration. *Medical Imaging, IEEE Transactions on* 27, 1 (2008), 111–128.
- [Ins] INSIGHT SOFTWARE CONSORTIUM: Insight segmentation and registration toolkit (itk). <http://www.itk.org/>.
- [KB96] KALLEL F., BERTRAND M.: Tissue elasticity reconstruction using linear perturbation method. *Medical Imaging, IEEE Transactions on* 15, 3 (1996), 299–313.
- [KBP*07] KAUS M. R., BROCK K. K., PEKAR V., DAWSON L. A., NICHOL A. M., JAFFRAY D. A.: Assessment of a model-based deformable image registration approach for radiation therapy planning. *Int. J. Radiation Oncology Bio. Phys.* 68, 2 (2007), 572–580.
- [KWK*98] KROUSKOP T. A., WHEELER T. M., KALLEL F., GARRA B. S., HALL T.: Elastic moduli of breast and prostate tissues under compression. *Ultrasonic imaging(Print)* 20, 4 (1998), 260–274.
- [LY03] LIANG J., YAN D.: Reducing uncertainties in volumetric image based deformable organ registration. *Medical Physics* 30 (2003), 2116–2122.
- [Mat] MATHEMATICS AND COMPUTER SCIENCE DIVISION, ARGONNE NATIONAL LABORATORY: Portable, extensible toolkit for scientific computation. <http://acts.nersc.gov/petsc/>.
- [ME96] MUTHUPILLAI R., EHMAN R. L.: Magnetic resonance elastography. *Nat Med* 2, 5 (May 1996), 601–603.
- [MT94] MOR'E J. J., THUENTE D. J.: Line search algorithms with guaranteed sufficient decrease. *ACM Trans. Math. Softw.* 20, 3 (1994), 286–307.
- [MTG03] MEEHAN M., TESCHNER M., GIROD S.: Three-dimensional simulation and prediction of craniofacial surgery. *Orthodontics & Craniofacial Research* 6, s1 (2003), 102–107.
- [NKJF09] NESME M., KRY P., JERABKOVA L., FAURE F.: Preserving topology and elasticity for embedded deformable models. *ACM Transaction on Graphics (proc. of SIGGRAPH 2009)* (2009).
- [NMK*06] NEALEN A., MULLER M., KEISER R., BOXERMAN E., CARLSON M.: Physically based deformable models in computer graphics. In *Computer Graphics Forum* (2006), vol. 25, Citeseer, pp. 809–836.
- [NW99] NOCEDAL J., WRIGHT S.: *Numerical Optimization*. Springer, 1999.
- [OAG*99] OPHIR J., ALAM S., GARRA B., KALLEL F., KONOFAGOU E., KROUSKOP T., VARGHESE T.: 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 (1999), 203–233.
- [Pre07] PRESS W. H.: *Numerical recipes*. Cambridge University Press, 2007.
- [RBF*08] RIVAZ H., BOCTOR E., FOROUGHI P., ZELLARS R., FICHTINGER G., HAGER G.: Ultrasound elastography: A dynamic programming approach. *Medical Imaging, IEEE Transactions on* 27, 10 (2008), 1373–1377.
- [RRM83] R B., R L., M R.: A computerized system for the elastic matching of deformed radiographic images to idealized atlas images. *J Comput Assist Tomogr* 7, 4 (August 1983), 618–625.
- [RSS*01] ROHR K., STIEHL H., SPRENGEL R., BUZUG T., WEESE J., KUHN M.: Landmark-based elastic registration using approximating thin-plate splines. *Medical Imaging, IEEE Transactions on* 20, 6 (June 2001), 526–534.
- [SB08] SYLLEBRANQUE C., BOIVIN S.: Estimation of mechanical parameters of deformable solids from videos. *The Visual Computer* 24, 11 (2008), 963–972.
- [SD02] SHEN D., DAVATZIKOS C.: Hammer: Hierarchical attribute matching mechanism for elastic registration. *IEEE Transactions on Medical Imaging* 21, 11 (2002), 1421–1439.
- [SE95] SKOVORODA A., EMELIANOV S.: Tissue elasticity reconstruction based on ultrasonic displacement and strain images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 42, 4 (1995), 141.
- [SHG*06] STEINEMANN D., HARDERS M., GROSS M., SZEKELY G., ZURICH E. T. H.: Hybrid cutting of deformable solids. In *Virtual Reality Conference, 2006* (2006), pp. 35–42.
- [Si] SI H.: Tetgen: A quality tetrahedral mesh generator and three-dimensional delaunay triangulator. <http://tetgen.berlios.de/>.
- [Thi98] THIRION J.-P.: Image matching as a diffusion process: an analogy with maxwell's demons. *Medical Image Analysis* 2, 3 (1998), 243–260.
- [THM*05] TESCHNER M., HEIDELBERGER B., MANOCHA D., GOVINDARAJU N., ZACHMANN G., KIMMERLE S., MEZGER J., FUHRMANN A.: Collision handling in dynamic simulation environments: The evolution of graphics: Where to next? In *Eurographics 2005 Tutorial* (2005).
- [WBG07] WICKE M., BOTSCH M., GROSS M.: A finite element method on convex polyhedra. In *Proceedings of Eurographics '07* (2007).
- [ZHO03] ZHU Y., HALL T., JIANG J.: A finite-element approach for young's modulus reconstruction. *Medical Imaging, IEEE Transactions on* 22, 7 (2003), 890–901.
- [ZNC*08] ZHANG M., NIGWEKAR P., CASTANEDA B., HOYT K., JOSEPH J. V., DI SANT'AGNESE A., MESSING E. M., STRANG J. G., RUBENS D. J., PARKER K. J.: Quantitative characterization of viscoelastic properties of human prostate correlated with histology. *Ultrasound in Medicine & Biology* (February 2008).
- [ZT05] ZIENKIEWICZ O. C., TAYLOR R. L.: *The Finite Element Method*, 6 ed. Butterworth-Heinemann, 2005.