

Bayesian Estimation of Deformation and Elastic Parameters in Non-rigid Registration

Petter Risholm^{1,2}, Eigil Samset², and William Wells III¹

¹ Harvard Medical School, Brigham and Women's Hospital

² Center of Mathematics for Applications, University of Oslo
pettri@ifi.uio.no

Abstract. Elastic deformation models are frequently used when solving non-rigid registration problems that are associated with neurosurgical image guidance, however, establishing precise values for the material parameters of brain tissue remains challenging. In this work we include elastography in the registration process by formulating these parameters as unknown random variables with associated priors that may be broad or sharp, depending on the situation.

A Bayesian registration model is introduced where the deformation probability is formulated by way of Boltzmann's equation and the linear elastic deformation and similarity energies. The full joint posterior on deformation and elastic random variables is characterized with a Markov Chain Monte Carlo method and can provide useful information beyond the usual "point estimates"; e.g. deformation uncertainty. Hard deformation constraints are easily accommodated in this framework which allows us to constrain the deformation of the brain to the intra-cranial space.

We describe preliminary experiments with synthetic 3D brain images for which ground truth is known for the elastic and deformation parameters. We compare a model with separate elastic parameters for three compartments (white matter, gray matter, and CSF), to a single compartment model, and show convergence, improved deformation estimates for the three compartment model and that plausible posteriors on the elastic parameters are obtained from the elastography process.

1 Introduction

Many surgical procedures induce tissue deformations such that image information acquired before surgery might not match the anatomy seen during surgery. In this paper we focus on deformations occurring during neurosurgery, however, the applicability of the framework we present is not restricted to the neurosurgical case. Opening of the dura leads to a gravitational shift of the brain tissue mainly due to the disappearance of pressure forces at the brain and ventricular boundaries[1]. This shift is commonly called brain-shift and has been reported to be ranging up to 24mm[2]. The "brain" consists of four main compartments, or tissue materials; skull (SK), gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF). These compartments will deform differently during

surgery because different tissues have different material properties. In particular, we may see a partial collapse of the ventricular areas due to CSF leakage[3]. A substantial uncertainty is involved in determining exact values for the material properties of brain tissue, and consequently the reported estimates of material parameters are quite divergent[4]. In addition, medication, radiation, or other factors related to the surgical procedure may significantly change the material parameters.

Many non-physical models have been proposed to recover tissue deformation, e.g. B-splines[5] and the Demons method[6]. However, these methods are not directly connected to the actual physical behavior of the tissue and do not explicitly model any material parameters. The most successful intra-operative registration methods use biomechanical models[7,8] which accommodates specifying separate elastic parameter values for different tissue types, but because the literature reports divergent values for the material parameters, most authors avoid the issue of varying tissue parameters and use fixed uniform material parameters for the whole registered anatomy. In [9], Ou et. al. introduced an interesting approach to elastography using a biomechanical model to estimate tissue parameters from two images that are deformed versions of each other. A disadvantage to their method is that the initial boundary conditions, i.e. the movement of boundary nodes in the biomechanical model, are not estimated in the method but needs to be specified.

In contrast to [9], we propose to integrate an elastography process in a Bayesian registration framework and thereby estimate both the boundary conditions as well as the elastic parameters simultaneously. A Markov Chain Monte Carlo sampling technique is applied to characterize the posterior distributions over deformation and elastic parameters. The sampling approach also effectively facilitates adding hard deformation constraints, for instance to constrain the deformation to the intra-cranial space.

2 Elastic Image Registration

Let $\mathbf{f}(\mathbf{x})$ and $\mathbf{m}(\mathbf{x})$, $\mathbf{x} \in \Omega$ be a pre- and intra-operative image of the brain respectively. They are defined over the d -dimensional image region $\Omega \subset \mathbb{R}^d$. We assume there exists a segmentation of \mathbf{f} into disjoint anatomical regions $\Omega = \cup_{l \in \{\text{CSF,GM,WM}\}} \Omega_l$, and require a segmentation of the skull region Ω_{SK} .

We assume there exists a displacement field $\mathbf{u}(\mathbf{x})$, $\mathbf{x} \in \Omega$ such that $\mathbf{m}(\mathbf{u}(\mathbf{x}) + \mathbf{x})$ is similar to \mathbf{f} . We measure the similarity between two images using an energy (similarity) function $E_s(\mathbf{u}; \mathbf{f}, \mathbf{m})$. Many types of similarity measures have been proposed in the literature. Any choice of similarity measure would be applicable in the proposed framework. However, in this work we restrict ourselves to the popular sum of squared differences

$$E_s(\mathbf{u}; \mathbf{f}, \mathbf{m}) = \int_{\Omega} (\mathbf{f}(\mathbf{x}) - \mathbf{m}(\mathbf{u}(\mathbf{x}) + \mathbf{x}))^2 dx . \quad (1)$$

It is common to restrict the deformations to model some physical meaningful process. Perhaps the most common model used in intra-subject registration is

to model brain tissue as a linear elastic material. This mechanism, which may be viewed as a regularizer, is defined in terms of an energy function

$$E_r(\mathbf{u}; \boldsymbol{\mu}, \boldsymbol{\lambda}) = \int_{\Omega} \frac{\boldsymbol{\mu}(\mathbf{x})}{4} \sum_{j,k=1}^d (\partial_{x_j} u_k + \partial_{x_k} u_j)^2 + \frac{\boldsymbol{\lambda}(\mathbf{x})}{2} (\operatorname{div} \mathbf{u}(\mathbf{x}))^2 d\mathbf{x}, \quad (2)$$

where $\mu, \lambda \in \mathbb{R}^+$ are the Lamé parameters. A subscript on \mathbf{x} and \mathbf{u} denotes a specific component of the 3D vector. Notice that we model the elastic parameters as functions of \mathbf{x} . In a traditional energy minimization registration method[10], we generally solve: $\operatorname{argmin}_{\mathbf{u}} E(\mathbf{u}; \mathbf{f}, \mathbf{m}, \boldsymbol{\mu}, \boldsymbol{\lambda}) = E_s(\mathbf{u}; \mathbf{f}, \mathbf{m}) + \alpha E_r(\mathbf{u}; \boldsymbol{\lambda}, \boldsymbol{\mu})$, where α is a weighting parameter.

2.1 Linear Elastic Finite Element Model

The Finite Element (FE) method is a powerful and versatile computational framework that is frequently used in the context of linear elastic problems over non-uniform domains. It provides a fast way of computing both E_r and E_s . FE-calculations are performed on a mesh that covers the region of interest. In this work we used a non-uniform tetrahedral mesh where the tetrahedral elements coarsely conform to the tissue boundaries delineated by Ω (GM, WM and CSF). The number of nodes and elements in the mesh is denoted N_e and N_n respectively.

In a FE setting we often work with the strain energy instead of the elastic potential in Eq. (2) (we refer to [11] for a detailed description of linear elastic FE-methods):

$$E_r = \frac{1}{2} \int_{\Omega} \boldsymbol{\epsilon}^T \boldsymbol{\sigma} d\mathbf{x}. \quad (3)$$

Let the displacement vector be $\mathbf{u} = [u, v, w]^T$, the strain vector is then defined as:

$$\boldsymbol{\epsilon} = \left[\frac{\partial u}{\partial x}, \frac{\partial v}{\partial y}, \frac{\partial w}{\partial z}, \frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}, \frac{\partial u}{\partial z} + \frac{\partial w}{\partial x}, \frac{\partial v}{\partial z} + \frac{\partial w}{\partial y} \right]^T. \quad (4)$$

We can rewrite this as $\boldsymbol{\epsilon} = \mathbf{B}\mathbf{u}$ where:

$$\mathbf{B} = \begin{bmatrix} \frac{\partial}{\partial x} & 0 & 0 & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} & 0 \\ 0 & \frac{\partial}{\partial y} & 0 & \frac{\partial}{\partial x} & 0 & \frac{\partial}{\partial z} \\ 0 & 0 & \frac{\partial}{\partial z} & 0 & \frac{\partial}{\partial x} & \frac{\partial}{\partial y} \end{bmatrix}^T.$$

The stress vector $\boldsymbol{\sigma}$ is related to the strain vector through Hooke's law, $\boldsymbol{\sigma} = \mathbf{C}\boldsymbol{\epsilon}$, where \mathbf{C} is the material matrix. For an isotropic material this matrix is defined by the two Lamé material constants λ and μ :

$$\mathbf{C} = \begin{bmatrix} \lambda + 2\mu & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda + 2\mu & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & \lambda + 2\mu & 0 & 0 & 0 \\ 0 & 0 & 0 & \mu & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix}. \quad (5)$$

Using these relations, adding work done by internal forces, \mathbf{f} , and minimizing with regards to \mathbf{u} , we end up with a linear matrix equation for each element:

$$\mathbf{K}_e \mathbf{u}_e = V_e \mathbf{B}_e^T \mathbf{C}_e \mathbf{B}_e \mathbf{u}_e = \mathbf{f}_e , \quad (6)$$

where \mathbf{K}_e is called the element stiffness matrix and V_e is the volume of the element. By assembling the contributions for each element into a square global stiffness matrix, \mathbf{K} , of size dN_n , we get a large sparse linear system:

$$\mathbf{K} \mathbf{u} = \mathbf{f} . \quad (7)$$

The linear system is under-determined – the solution can be seen as a mesh floating freely in space. Consequently, the positions of at least three nodes need to be specified to constrain the computations appropriately.

For a given deformation \mathbf{u} , the elastic energy in Eq. (3) can be computed efficiently using the FE-model. To compute the integral in Eq. (1), we use a 4-point Gaussian quadrature[11] and sum up the contributions from each element.

3 A Bayesian Framework for Elastic Image Registration

Our framework is based on a generative Bayesian probability model where the node deformations as well as the elastic material parameters can vary over the image domain.

3.1 Generative Model

We assume that the intra-operative image \mathbf{m} , the deformation \mathbf{u} , and the elastic parameters $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$ are random variables generated by the hierarchical model in Fig. 1. Notice that the pre-operative image, \mathbf{f} , is a model parameter. This leads to the following joint probability model:

$$p(\mathbf{m}, \mathbf{u}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = p(\mathbf{m}|\mathbf{u})p(\mathbf{u}|\boldsymbol{\lambda}, \boldsymbol{\mu})p(\boldsymbol{\lambda})p(\boldsymbol{\mu}) . \quad (8)$$

According to the theorem of conditional probability we can write the posterior as

$$p(\mathbf{u}, \boldsymbol{\lambda}, \boldsymbol{\mu}|\mathbf{m}) = \frac{p(\mathbf{m}, \mathbf{u}, \boldsymbol{\lambda}, \boldsymbol{\mu})}{p(\mathbf{m})} = \frac{p(\mathbf{m}|\mathbf{u})p(\mathbf{u}|\boldsymbol{\lambda}, \boldsymbol{\mu})p(\boldsymbol{\lambda})p(\boldsymbol{\mu})}{p(\mathbf{m})} . \quad (9)$$

Since we are mainly interested in characterizing the deformation, the Lamé parameters can be seen as nuisance parameters and marginalized out:

$$p(\mathbf{u}|\mathbf{m}) = \int_{\Omega} \int_{\Omega} \frac{p(\mathbf{m}|\mathbf{u})p(\mathbf{u}|\boldsymbol{\lambda}, \boldsymbol{\mu})p(\boldsymbol{\lambda})p(\boldsymbol{\mu})}{p(\mathbf{m})} d\boldsymbol{\lambda} d\boldsymbol{\mu} . \quad (10)$$

We use the Boltzmann distribution to define the likelihood, often called the similarity term

$$p(\mathbf{m}|\mathbf{u}) = \frac{1}{Z_s} \exp \left(-\frac{E_s(\mathbf{u}; \mathbf{m}, \mathbf{f})}{T_s} \right) , \quad (11)$$

where Z_s is a normalizing constant and T_s is the temperature parameter for the distribution. The Boltzmann distribution is similarly used to model the prior on the transformation which acts as a regularizer

$$p(\mathbf{u}|\boldsymbol{\lambda}, \boldsymbol{\mu}) = \frac{1}{Z_r} \exp\left(-\frac{E_r(\mathbf{u}; \boldsymbol{\lambda}, \boldsymbol{\mu})}{T_r}\right). \quad (12)$$

The Lamé parameters, $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$, should be non-negative to be compatible with human tissue. There are many suitable priors that can be used to model the elastic parameters depending on the prior information that is available. We model the elastic parameter priors, $p(\boldsymbol{\lambda})$ and $p(\boldsymbol{\mu})$, by Beta distributions scaled to the region $[r_1, r_2]$ and assume that the parameters are independent: $p(\boldsymbol{\mu}; A_\mu, B_\mu, r_1, r_2) = \prod_{\mu_i \in \boldsymbol{\mu}} \text{Beta}(\mu_i; A_\mu, B_\mu, r_1, r_2)$. The prior on $\boldsymbol{\lambda}$ is defined similarly. We have flexibility in modeling the prior – we can use separate priors for different brain compartments, or alternatively use a global prior.




Fig. 1. A generative model of the registration problem. The deformation random variable, \mathbf{u} , depends on the elastic random variables, $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$. The deformation governs the image, \mathbf{m} , we will observe given the image \mathbf{f} . $A_\lambda, A_\mu, B_\lambda$ and B_μ are parameters controlling the prior (beta) distributions for the elastic random variables. The two parameters, T_s and T_r , control the temperature of the likelihood distribution and the prior distribution on the deformation respectively. The plate parameter, N , is established by the resolution we sample the elastic parameters at. It can be uniform over the image domain, per tissue compartment or per element in the FE-model.

3.2 MCMC-Optimization

We wish to characterize the posterior in Eq. (10), given images \mathbf{f} and \mathbf{m} . Unfortunately, it is not possible to analytically compute the posterior nor feasible to draw samples directly from it. A common approach to generate samples from complicated posteriors is by way of the Metropolis-Hastings (MH) algorithm[12]. In our setting we will draw samples from $p(\mathbf{u}, \boldsymbol{\lambda}, \boldsymbol{\mu}|\mathbf{m})$ and “discard” the values for $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ to characterize Eq. (10). However, notice that the “discarded” values can be used for characterizing the marginal posterior distributions for the elastic parameters. While deterministic optimization procedures generally requires the Euler-Lagrange equations, MCMC-optimization stochastically optimizes the posterior distribution purely based on the probability distributions defined in the previous section and thereby eliminates the need for computing any gradients of the energy equations. Algorithm 1 provides an overview of the method.

Metropolis-Hastings. In the MH-algorithm, candidate samples are drawn from a proposal distribution $q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^n)$, with $\boldsymbol{\theta} = (\mathbf{u}, \boldsymbol{\lambda}, \boldsymbol{\mu})$, from which it should be easy to draw samples. For a particular step n of the algorithm, where the current state is $\boldsymbol{\theta}^n$, we draw a sample $\boldsymbol{\theta}^*$ from $q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^n)$ and accept it with probability

$$A(\boldsymbol{\theta}^*, \boldsymbol{\theta}^n) = \min \left(1, \frac{p(\boldsymbol{\theta}^*|\mathbf{m})q(\boldsymbol{\theta}^n|\boldsymbol{\theta}^*)}{p(\boldsymbol{\theta}^n|\mathbf{m})q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^n)} \right). \quad (13)$$

If the sample is accepted, $\boldsymbol{\theta}^{n+1} = \boldsymbol{\theta}^*$, otherwise $\boldsymbol{\theta}^{n+1} = \boldsymbol{\theta}^n$. Notice that the evaluation of the acceptance criterion does not require knowledge of the normalizing constants Z_s , Z_r and $p(\mathbf{m})$. For symmetric proposal distributions, such as the Normal distribution, the MH algorithm reduces to the standard Metropolis criterion where the ratio of the proposal distributions equals one.

Proposal Distributions. We model the proposal distribution with three Normal *jumping kernels*[12] and assume that the parameters are independent:

$$\boldsymbol{\theta}^* \sim q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^n) = N_{\mathbf{u}}(\mathbf{u}^*|\mathbf{u}^n, \sigma_u) N_{\boldsymbol{\lambda}}(\boldsymbol{\lambda}^*|\boldsymbol{\lambda}^n, \sigma_{\boldsymbol{\lambda}}) N_{\boldsymbol{\mu}}(\boldsymbol{\mu}^*|\boldsymbol{\mu}^n, \sigma_{\boldsymbol{\mu}}). \quad (14)$$

By using a rejection sampler, we restrict the elastic parameters to $\boldsymbol{\mu}, \boldsymbol{\lambda} \in \mathbb{R}^+$ and the node positions from moving into the skull region $(\mathbf{p} + \mathbf{u}) \notin \Omega_{\text{SK}}$, where \mathbf{p} are the initial nodal positions. Because the proposal samplers are Normal distributions, the proposal ratio $q(\boldsymbol{\theta}^n|\boldsymbol{\theta}^*)/q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^n)$ equals one.

Sampling Convergence. A problem in iterative simulations is to assess whether or not we are generating samples from the target distribution and whether we have generated enough samples to characterize it properly. A common approach of monitoring convergence is to compute the *scale reduction* for each scalar estimand using parallel Markov Chains[12]. We use three parallel Markov Chains and assume the posterior distribution has been adequately characterized when the scale reduction is less than 1.2 for all estimands.

3.3 Sampling Strategies

Deformation Sampling. We can sample deformations involving from three to all the nodes in the mesh. Calculation of Eq. (7) provides us with the nodal deformations for the nodes that are not sampled. In this work we sampled deformations for all nodes on the brain boundary, while the deformations for the inner nodes were found through the FE-calculations.

Deformation Constraints. Restricting the sampler from accepting nodal positions in Ω_{SK} avoids impossible nodal configurations. At the same time it allows the brain to slide along the skull boundary to find the minimum energy configuration.

Elasticity Sampling. We have the option of sampling the elastic parameters per element, per tissue label or uniformly over the image domain. For the present experiments we assume that the intra-tissue elastic variation is negligible and consequently sample the elastic parameters on a per-tissue basis. Hence, we sampled different λ and μ for Ω_{GM} , Ω_{WM} and Ω_{CSF} so that $\boldsymbol{\lambda} = (\lambda_{\text{CSF}}, \lambda_{\text{GM}}, \lambda_{\text{WM}})$ and $\boldsymbol{\mu} = (\mu_{\text{CSF}}, \mu_{\text{GM}}, \mu_{\text{WM}})$.

Because no internal forces are applied in the FE-calculations, $\mathbf{f} = \mathbf{0}$, we can see from inspecting Eq. (6) that the ratios of the elastic parameters controls the tissue elasticity while the scale of $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$ will have no affect on the calculations. Hence, to avoid drifting of the parameters towards zero, we fix one of the elastic parameters and sample the rest of them. It is also worth mentioning that since nodal positions are only sampled for nodes on the boundary of the mesh, it is the sampling of the elastic parameters that effectively explores the deformation space for the internal mesh nodes.

Algorithm 1. Bayesian estimation of deformation and elastic parameters.

repeat

 Sample $\mathbf{u}^* \sim N_{\mathbf{u}}(\mathbf{u}^n, \sigma_{\mathbf{u}})$, $(\mathbf{p} + \mathbf{u}^*) \in \Omega$.

 Sample $\boldsymbol{\lambda}^* \sim N_{\boldsymbol{\lambda}}(\boldsymbol{\lambda}^* | \boldsymbol{\lambda}^n, \sigma_{\boldsymbol{\lambda}})$, $\boldsymbol{\lambda}^* \in \mathbb{R}^+$

 Sample $\boldsymbol{\mu}^* \sim N_{\boldsymbol{\mu}}(\boldsymbol{\mu}^* | \boldsymbol{\mu}^n, \sigma_{\boldsymbol{\mu}})$, $\boldsymbol{\mu}^* \in \mathbb{R}^+$

 Compute the energies in Eqs. (1) and (3)

 Compute the MH criteria in Eq. (13).

if Accept **then** $\boldsymbol{\theta}^{n+1} = \boldsymbol{\theta}^*$

else $\boldsymbol{\theta}^{n+1} = \boldsymbol{\theta}^n$

end if

until Convergence

4 Results

4.1 Dataset

Preliminary validation of the framework was performed on a synthetic dataset where ground truth was established. We acquired an anatomic T1 weighted MR image together with a label map from the BrainWeb[13] database. A label-dependent tetrahedral mesh consisting of $N_e = 3807$ elements and $N_n = 801$ nodes was generated. We fixed 20 boundary mesh nodes with a deformation of 15 mm in the direction of the center of gravity of the brain and used the sampler to find the configuration of the remaining boundary nodes that minimized the elastic energy of the mesh. The literature[4] indicates that the elastic ratio of brain tissue should be approximately $\lambda/\mu \approx 24$. For generation of the dataset we fixed the elastic parameters to $\mu_{\text{CSF}} = 10^{-6}$, $\mu_{\text{GM}} = 1$, $\mu_{\text{WM}} = 1$, $\lambda_{\text{CSF}} = 10^{-6}$, $\lambda_{\text{GM}} = 10$, $\lambda_{\text{WM}} = 30$. The resulting dataset can be seen in Fig. 2. Even though CSF is clearly not an elastic material, we model it with small elastic values compared to GM and WM to approximate the fluid nature of CSF.

4.2 Experiments

Two separate experiments were carried out. In both experiments we sampled deformations for the 349 boundary nodes, while the deformations of the internal nodes were computed by solving Eq. (7). The elastic parameters were sampled in two different ways. In the first experiment the elastic parameters were uniformly sampled for all tissue types such that $\mu_{\text{CSF}} = \mu_{\text{GM}} = \mu_{\text{WM}}$ and $\lambda_{\text{CSF}} = \lambda_{\text{GM}} = \lambda_{\text{WM}}$. In the second experiment the parameters were sampled independently. We assumed no prior information on the elastic parameters and therefore used a uniform prior on them ($A_\lambda = A_\mu = B_\lambda = B_\mu = 1, r_1 = 0, r_2 = \infty$), identical deformation sampling kernels ($\sigma_u = 0.008$) and the same temperatures ($T_s = 0.05, T_r = 10\ 000$) for the two experiments. These parameters resulted in a MCMC acceptance rate of approximately 25%. We used a thinning factor of 10 and generated samples until the reduction scale was less than 1.2 for all estimands. A total of 250 000 samples were generated for each of the three parallel samplers, but, after thinning, only 25 000 samples were used to characterize

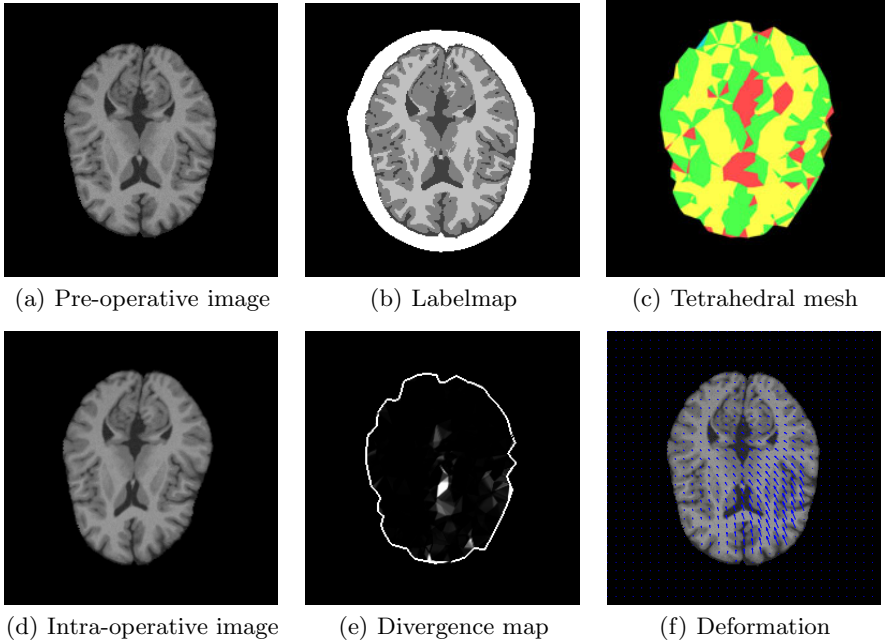


Fig. 2. The synthetic dataset. In (c) we show a cut through the mesh we used for the computations. CSF is colored in red, GM in green and WM in yellow. The poor correspondence between labels in (b) and (c) is a result of the coarse mesh. In (e) we show the divergence, which is governed by the λ parameter (see Eq. (2)), of the deformation field. Notice the bright areas in the divergence map, they are signs of the collapse of the ventricular areas. Also notice the skull label in (b), the sampler restricts movements of nodes into this area. Figure (f) shows the x - and y -components of the ground truth deformation field.




Fig. 3. The two images on the left shows the norm of the difference between the ground truth deformation and the estimated deformation field with varying elastic parameters (a) and with uniform elastic parameters (b). Comparing the deformation errors, it is clear that using varying elastic parameters provides a better estimate of the deformation in the interior of the brain. The collapse of the ventricles (see the divergence map in Fig. 2(e)) is better accommodated with varying elastic parameters than in the uniform case (compare the two rightmost images where (c) stems from using varying elastic parameters while (d) from uniform parameters).

Table 1. This table reports statistics on absolute error between the posterior means of the boundary node deformations and the ground truth deformations for these nodes. We also report statistics on the distance between the initial configuration of the nodes and the ground truth. From this and Tbl. 2 we can draw the conclusion that the sampler provides comparable estimates of the movements of the nodes on the boundary, while sampling different elasticity parameters provides us with a better estimate of the deformation in the interior of the brain. All values are in millimeters.

	Median	95th quantile	Max
Initial Configuration	2.56	8.24	15.08
Different Elastic Parameters	0.16	0.72	2.43
Uniform Elastic Parameters	0.18	0.83	2.40

the posterior distributions (a total of 75 000 samples). Each parallel sampler generated 5 samples per second which resulted in a total computation time of approximately 14 hours. Notice that we used the same sampling framework for generating the synthetic dataset as we used to estimate the elastic and deformation parameters. Hence, large discrepancies between the ground truth and the estimated parameters is not expected.

MCMC methods should use starting points that are crude estimates of the mode of the posterior distribution[12] that we are interested in. A simpler and faster registration method, such as the Demons method[6], can be used to find an initial configuration of the deformation parameters. We used initial deformation parameters for the boundary nodes that were in between an identity deformation and the ground truth deformation. In Tbl. 1 we report statistics on the distance between the initial configuration and the ground truth, as well as the distance between the estimated posterior means of the node deformations and the ground truth.

Table 2. In this table we report the Mean Squared Error (MSE) between the computed deformation and the ground truth deformation over different tissue types. The upper row contains results from varying the elastic parameters for each tissue type, while the results in the lower row stems from using uniform elastic parameters. It is evident that varying the elastic parameters provides a better estimate of the ground truth. All values are in millimeters squared.

	Ω	Ω_{CSF}	Ω_{GM}	Ω_{WM}
Different Elastic Parameters	0.040	0.049	0.049	0.023
Uniform Elastic Parameters	0.065	0.087	0.067	0.048




Fig. 4. Marginal posteriors of the elastic ratios for each tissue compartment. The ground truth values are, $\lambda_{CSF}/\mu_{CSF} = 1$, $\lambda_{GM}/\mu_{GM} = 10$ and $\lambda_{WM}/\mu_{WM} = 30$. It is evident that the method is able to recover elastic parameters that distinguishes between different tissue types, but not necessarily capable of recovering the exact underlying values. It is also evident that the elastic ratios are biased towards zero. This may be explained by the drifting mentioned in Sec. 3.3 and the Boltzmann distribution of the elastic energy. The lower the temperature is set, the more regularization we apply to the registration and the more the elastic parameters will tend towards zero. WM volume is also larger than for GM, hence the ratio for WM tends more towards zero than for GM.




Fig. 5. Translations were sampled for each node on the boundary of the mesh. This figure shows the marginal posterior distributions for each translational component for one of the boundary nodes. The posterior distribution of nodes moving along the skull-brain boundary gets the truncated shape we see in (a) and (b).

Figure 3 provides qualitative results from the registrations while Tbl. 2 reports quantitative Mean Squared Error (MSE) results from comparing the final deformation field with the ground truth deformation. It is evident that sampling

the elastic parameters per compartment provides better results in the brain interior, while the estimated deformations on the brain boundary (see Tbl. 1) are comparable. In Fig. 4 we plot posterior distributions of the elastic ratios. The experiments show that we are able to distinguish between tissue types and that we recover plausible elastic ratios, however the ratios are biased towards zero. Figure 5 shows the posterior distribution of one of the boundary nodes. Notice how the rejection of samples that fall in the skull region have truncated the posterior distributions for this particular node.

5 Discussion

In this paper we presented a Bayesian framework for simultaneous estimation of deformation and material parameters. One advantage of our approach is the characterization of marginal posteriors instead of point estimates for the registration parameters – not only can we find the most likely configuration of the deformation field given two images, but from the posterior distribution we can also quantify deformation uncertainty. Another novelty of the method is the ability for elastography without first determining the initial boundary conditions. Prior information on the material parameters is seamlessly included in the framework through the prior distributions on them. The MCMC sampler facilitates inclusion of hard deformation constraints; for example if a segmentation of the skull is available, brain movements can be restricted to the intra-cranial space.

Unfortunately, high-dimensional problems solved with MCMC methods require a large number of samples to adequately characterize the posterior distribution. Hence, it is a computationally intensive method that might not be suitable for intra-operative use in the near term. However, the use of parallel samplers on large computer clusters can reduce the computational time considerably. Because the method is highly computationally intensive, a coarse brain mesh was applied for the experiments to make it computationally tractable. Future work should study the sensitivity of the method to the tissue segmentation and the coarseness of the mesh. Notice that the need for a tissue segmentation can be eliminated by sampling the elastic parameters per element instead of per tissue compartment.

Three main conclusions can be drawn from the preliminary experiments: 1) elastography is feasible without knowing initial boundary conditions, and plausible marginal posteriors on material parameters are obtained, 2) it is possible to characterize posterior distributions on deformations and 3) improved registration results are achieved with finer elastic sampling resolution. Hence, this registration framework is particularly useful whenever it is difficult to establish precise values for the elastic parameters.

Acknowledgments. We are grateful to Prof. Hal Stern, University of California, Irvine, for his helpful advice. This research was supported by NIH grants R01CA138419 and U41RR019703.

References

1. Ferrant, M., Nabavi, A., Macq, B., Black, P., Jolesz, F., Kikinis, R., Warfield, S.: Serial registration of intraoperative mr images of the brain. *Medical Image Analysis* 6(4), 337–359 (2002)
2. Hastreiter, P., Rezk-Salama, C., Soza, G., Greiner, G., Fahlbusch, R., Ganslandt, O., Nimsky, C.: Strategies for Brain Shift Evaluation. *Medical Image Analysis* 8(4), 447–464 (2004)
3. Hartkens, T., Hill, D.L.G., Castellano-Smith, A.D., Hawkes, D.J., Maurer, C.R., Martin, A.J., Hall, W.A., Liu, H., Truwit, C.L.: Measurement and analysis of brain deformation during neurosurgery. *IEEE Trans. Med. Imaging* 22(1), 82–92 (2003)
4. Hagemann, A., Rohr, K., Stiehl, H., Spetzger, U., Gilsbach, J.: Biomechanical modeling of the human head for physically based. *Nonrigid Image Registration* 18(10), 875–884 (1999)
5. Rueckert, D., Sonoda, L., Hayes, I., Hill, D., Leach, M., Hawkes, D.: Nonrigid registration using free-form deformations: Application to breast mr images. *IEEE Trans. Med. Imaging* 18(8), 712–721 (1999)
6. Thirion, J.P.: Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis* 2(3), 243–260 (1998)
7. Clatz, O., Delingette, H., Talos, I.F., Golby, A.J., Kikinis, R., Jolesz, F.A., Ayache, N., Warfield, S.K.: Robust nonrigid registration to capture brain shift from intraoperative MRI. *IEEE Trans. Med. Imaging* 24(11), 1417–1427 (2005)
8. Miga, M.I., Roberts, D.W., Kennedy, F.E., Platenik, L.A., Hartov, A., Lunn, K.E., Paulsen, K.D.: Modeling of retraction and resection for intraoperative updating of images. *Neurosurgery* 49(1), 75–85 (2001)
9. Ou, J.J., Ong, R.E., Yankeelov, T.E., Miga, M.I.: Evaluation of 3d modality-independent elastography for breast imaging: a simulation study. *Physics in Medicine and Biology* 53(1), 147–163 (2008)
10. Modersitzki, J.: *Numerical Methods for Image Registration*. Oxford University Press, Oxford (2004)
11. Zienkiewicz, O.C., Taylor, R.L.: *The Finite Element Method*, 5th edn. Butterworth-Heinemann, Butterworths (2000)
12. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.: *Bayesian Data Analysis*, 2nd edn. Texts in Statistical Science. Chapman & Hall/CRC (July 2003)
13. Collins, D.L., Zijdenbos, A.P., Kollokian, V., Sled, J.G., Kabani, N.J., Holmes, C.J., Evans, A.C.: Design and construction of a realistic digital brain phantom. *IEEE Trans. Med. Imaging* 17(3), 463–468 (1998)