Published in final edited form as:

Med Image Anal. 2009 October ; 13(5): 744–756. doi:10.1016/j.media.2009.07.002.

Data Assimilation Using a Gradient Descent Method for Estimation of Intraoperative Brain Deformation

Songbai Ji¹, Alex Hartov^{1,2}, David Roberts^{2,3}, and Keith Paulsen^{1,2}

¹Thayer School of Engineering, Dartmouth College, Hanover, NH 03755 USA

²Norris Cotton Cancer Center, Lebanon, NH 03756 USA

³Dartmouth-Hitchcock Medical Center, Lebanon, NH 03756 USA

Abstract

Biomechanical models that simulate brain deformation are gaining attention as alternatives for brain shift compensation. One approach, known as the “forced-displacement method”, constrains the model to exactly match the measured data through boundary condition (BC) assignment. Although it improves model estimates and is computationally attractive, the method generates fictitious forces and may be ill-advised due to measurement uncertainty. Previously, we have shown that by assimilating intraoperatively acquired brain displacements in an inversion scheme, the Representer algorithm (REP) is able to maintain stress-free BCs and improve model estimates by 33% over those without data guidance in a controlled environment. However, REP is computationally efficient only when a few data points are used for model guidance because its costs scale linearly in the number of data points assimilated, thereby limiting its utility (and accuracy) in clinical settings. In this paper, we present a steepest gradient descent algorithm (SGD) whose computational complexity scales nearly invariantly with the number of measurements assimilated by iteratively adjusting the forcing conditions to minimize the difference between measured and model-estimated displacements (model-data misfit). Solutions of full linear systems of equations are achieved with a parallelized direct solver on a shared-memory, eight-processor Linux cluster. We summarize the error contributions from the entire process of model-updated image registration compensation and we show that SGD is able to attain model estimates comparable to or better than those obtained with REP, capturing about 74% to 82% of tumor displacement, but with a computational effort that is significantly less (a factor of 4-fold or more reduction relative to REP) and nearly invariant to the amount of sparse data involved when the number of points assimilated is large. Based on five patient cases, an average computational cost of approximately 2 minutes for estimating whole-brain deformation has been achieved with SGD using 100 sparse data points, suggesting the new algorithm is sufficiently fast with adequate accuracy for routine use in the operating room (OR).

Keywords

Biomechanical model; inverse method; image-guided neurosurgery; data assimilation; brain shift

1. Introduction

Modern neuronavigational systems rely on preoperative magnetic resonance images (pMR) and fiducial-based registration schemes to establish a rigid-body transformation between anatomical structures in the operating room (OR) and their counterparts in pMR. During the course of open cranial neurosurgery, however, brain shift induced by gravitational effects and surgical interventions often exceeds 20 mm at the cortical surface after dural opening and can reach more than 3 mm at the tumor boundary (Nimsky et al., 2000). These magnitudes in the face of the complex, dynamic and seemingly unpredictable nature of the brain motion (Hartkens et al., 2003; Nabavi et al., 2001; Sun et al., 2005a) pose a significant challenge to the accuracy of neuronavigational systems.

To compensate for the brain deformation, various intraoperative imaging techniques have been devised. Intraoperative stereovision (Skrinjar et al., 2002; Sun et al., 2005ab) and laser range scanning (Miga et al., 2003) acquire images of the parenchymal surface at the craniotomy site, but not internal anatomical structures. The dominant schemes are intraoperative magnetic resonance (iMR; e.g., Black et al., 1997; Ferrant et al., 2002; Hall et al., 2000; Nimsky et al., 2001; Wirtz et al., 1997) and ultrasound (iUS; e.g., Bonsanto et al., 2001; Bucholz et al., 2001; Comeau et al., 2000; Unsgaard et al., 2005), because they supply full or partial volume sampling of the anatomical features of interest deeper in the brain. Intraoperative MR provides unparalleled delineation of parenchymal tissue and is able to scan the full volume of the head (e.g., Maurer et al., 1998; Nabavi et al., 2001; Nimsky et al., 2000). However, its capital cost and impact on neurosurgical workflow in the OR as well as its relatively long image acquisition are potential barriers to gaining wide acceptance. By comparison, iUS is easy to implement, low in cost and achieves real time image acquisition compared with iMR. However, iUS image noise, artifacts and limited feature localization accuracy make anatomical structures more difficult to delineate. In addition, tumor motion may only be partially measured when two-dimensional (2D) iUS is applied because of discontinuous and limited imaging plane sampling. These respective characteristics of MR and US make their combination particularly attractive in image-guided neurosurgery, where anatomical features (e.g., tumor) deep in the brain captured in iUS (i.e., in deformed state) can be compared with their counterparts in pMR (i.e., in undeformed, original state) to generate brain displacement maps locally to compensate for intraoperative motion.

When incomplete brain motion is measured, biomechanical models are effective alternatives to estimate whole brain deformation for brain shift compensation (e.g., Clatz et al., 2005; Skrinjar et al., 2002; Sun et al., 2005b; Lunn et al., 2005, 2006; Ferrant et al., 2001; Wittek et al., 2005). These models assimilate intraoperative displacement data measured from the cortical surface using stereo vision (Skrinjar et al., 2002; Sun et al., 2005b) or laser range scanning (Dumpry et al., 2007), or from deep brain anatomical structures (e.g., tumor, ventricle) using iMR (e.g., Clatz et al., 2005; Ferrant et al., 2001; Hu et al., 2007; Wittek et al., 2005) or iUS (e.g., Lunn et al., 2005, 2006), and estimate brain deformation subject to loading and boundary conditions in the OR. One approach, known as the “forced-displacement method”, constrains the model to exactly match the measured data through boundary condition assignment (Clatz et al., 2005; Ferrant et al., 2001; Skrinjar et al., 2002; Sun et al., 2005b; Wittek et al., 2005). Although this method improves model displacement estimates and is computationally attractive, fictitious forces at otherwise stress-free locations are generated that may degrade its performance (Liu et al., 2006). Further, matching the measured data exactly may be ill-advised because of measurement errors that are always present (Carter et al., 2005; Lunn et al., 2005).

Alternatively, a least squares inversion scheme using an adjoint equations method (AEM) is able to maintain stress-free boundary conditions while fitting the model to the measured data

approximately by adjusting the forcing conditions (Lynch, 2004; Lunn et al., 2005, 2006). Using the Representer (REP) algorithm, model response subject to a unit misfit is computed, and the overall model deformation is then the weighted superposition of these individual responses determined through a least squares fitting scheme. As shown in Lunn et al. (2005), this algorithm solves the set of adjoint equations directly (at a cost of $6N_{sprs}+1$ full linear system matrix solutions in total for a 3D model, where N_{sprs} is the number of sparse data points used for model guidance) to minimize the difference between the model-predicted displacements and the measured data (model-data misfit). We have recently shown that using the REP algorithm to solve the adjoint equations, displacement estimates can be improved on average 33% over that without data guidance in a porcine brain retraction study (Lunn et al., 2006). Using clinical patient cases, we have also demonstrated that tumor displacement estimates degraded (i.e., model-data misfit at the tumor boundary became *worse* than that initially) when nodal displacements at the exposed cortical surface were constrained in a forced-displacement method, but improved by 20–43% when these displacements were otherwise assimilated through an inverse method (Liu et al., 2006). The amount of improvement can be even larger when displacement maps at *both* the cortical surface and the tumor boundaries are used for model guidance. These findings suggest that inversion approaches which approximate the whole-brain deformation based on locally measured displacement maps are indeed capable of outperforming forced- displacement methods.

However, the REP algorithm applied in these studies is computationally efficient only when the number of sparse data points is small (Lynch, 2004; e.g., $N_{sprs}=10/4$ in the simulated/ clinical 2D model in Lunn et al. (2005), and $N_{sprs}=7$ in the 3D porcine brain study in Lunn et al. (2006)). Although a small amount of data may be sufficient for accurate model guidance in a controlled environment (Lunn et al., 2005, 2006), it is questionable whether a few data points are adequate in clinical settings, where brain motion is significant, dynamic and complex (Hartkens et al., 2003; Nimsy et al., 2001; Nabavi et al., 2001), and has less predictable or unknown measurement errors. While increasing N_{sprs} to guide the model is expected to result in more accurate estimates, the computational effort associated with REP (Lunn et al., 2006) is likely to limit its ultimate application in the OR, especially as the amount of measurement data grows.

Therefore, developing and implementing a computationally efficient, yet sufficiently accurate algorithm for estimating whole-brain deformation is *critical* for application of the model-based scheme for brain shift compensation in the OR. The major steps involved in such a scheme typically include patient registration, intraoperative imaging to sample the exposed cortical surface and/or features deep in the brain (e.g., tumor), measurement of local brain displacements in the sampled regions, data assimilation using a biomechanical model, and visualization feedback based on the computed whole-brain deformation. The most computationally intensive steps appear to be the measurement of local brain displacement and the computational data assimilation to estimate whole-brain deformation (model-updated image visualizations are fast and not a time-limiting step). Given that we have recently developed a completely automatic and efficient (1–2 min) image-based re-registration scheme that can be used to measure tumor displacement by comparing tumor boundaries captured in iUS and pMR (Ji et al., 2008a), the efficiency of the data assimilation to compute whole-brain deformation is, therefore, of critical importance to the overall procedure for application in the OR (e.g., the computational cost of the overall procedure to be within 5 min).

For this purpose, we have incorporated an algorithm using the steepest gradient descent (SGD) to dramatically reduce the computational burden relative to the REP algorithm without sacrificing accuracy in the model estimate. The SGD algorithm solves the identical set of adjoint equations by iteratively adjusting the forcing conditions according to the gradient of the objective function with respect to the forcing vector. In contrast to REP, the computational

complexity of SGD scales nearly invariantly with the number of sparse data points, N_{sprs} , used for model guidance, making it computationally much more efficient than REP when N_{sprs} is large, which is important because model accuracy improves with increases in the amount of data assimilated. This scaling behavior of the algorithm is ideal for computations in the OR, especially when volumetric 3D iUS is utilized that fully samples the target tissue, in which case large amounts of uniformly distributed displacement data are generated that need to be assimilated into the estimation of whole-brain deformation.

To demonstrate that the SGD algorithm is able to produce model estimates of comparable or improved accuracy relative to REP but with significantly less computational effort that is nearly invariant to the amount of data assimilated, we applied the technique first to a clinical brain model with simulated tumor displacement data and then to five (N=5) clinical datasets with measured intraoperative tumor displacement after dural opening, and compared the accuracy and computational efficiency with the results obtained from REP.

2. Material and Methods

2.1 Biomechanical Brain Model Formulation

Biot's consolidation theory (Biot, 1941) was employed to model the brain as a porous deformable medium to incorporate a pressure gradient from incompressible interstitial fluid as a hydrodynamic driving force. The detailed mathematical formulation has been developed elsewhere (Paulsen et al., 1999), and is summarized by the following coupled equations:

$$\nabla \cdot G \nabla \mathbf{u} + \nabla \frac{G}{1 - 2\nu} (\nabla \cdot \mathbf{u}) - \alpha \nabla p = (\rho_t - \rho_f) \mathbf{g} \quad (1)$$

$$\alpha \frac{\partial}{\partial t} (\nabla \cdot \mathbf{u}) + \frac{1}{S} \frac{\partial p}{\partial t} - \nabla \cdot k \nabla p = \Psi \quad (2)$$

where $\mathbf{u}$ and p are the displacement and pore fluid pressure to be computed, respectively. Other material properties are summarized in Table 1. The values of the parameters used in this study were obtained from the literature (Miga, 1998), and were identical to those in Lunn et al. (2005). The utility of these model parameters has been demonstrated in a previous porcine brain retraction study (Platenik et al., 2002). Using the finite element discretization method, the above equations are reduced to a linear algebraic system expressed in matrix form (Paulsen et al., 1999)

$$\mathbf{Kx} = \mathbf{b} \quad (3)$$

where $\mathbf{K}$ is the stiffness matrix that embeds the discretized model equations, $\mathbf{x}$ is the displacement and pore fluid pressure vector (4 degrees-of-freedom (DOF) per node), and $\mathbf{b}$ includes the surface forces and displacement boundary conditions due to surgical intervention, interior forces due to the pressure gradient of the interstitial fluid and tissue buoyancy forces caused by gravitational effects. In addition, when brain-skull contact boundary conditions are incorporated, forces at the points of contact between the brain and skull surfaces are also included (Ji et al., 2007).

2.2 The Adjoint Equations and the Steepest Gradient Descent Method

When incomplete displacement measurements are available, a generalized least squares inversion scheme is suitable for minimizing the difference () between the model estimate ($\mathbf{x}$) and measured data ($\mathbf{d}$) through a minimization process constrained by the model equations (Eqn. 3; Lynch, 2004). This is achieved by using the well-known Lagrange multiplier method that embeds the model equations into the minimizing objective function. Specifically, a data error covariance model, a forcing condition correlation covariance model and a set of Lagrange multipliers () are introduced to construct an augmented quadratic objective function, (Lynch, 2004):

$$\Omega = (\mathbf{d} - \mathbf{Ax})^T \mathbf{W}_\varepsilon (\mathbf{d} - \mathbf{Ax}) + \mathbf{b}^T \mathbf{W}_b \mathbf{b} + \lambda^T (\mathbf{Kx} - \mathbf{b}) \quad (4)$$

where $\mathbf{A}$ is the sampling matrix producing values at given locations of the data. The matrices $\mathbf{W}$ and $\mathbf{W}_b$ (the inverse of the covariances of and $\mathbf{b}$, respectively) describe the covariances of the misfit between any two data points and the forcing conditions between any two locations, and are controlled by the standard deviation (std) in the measurement error () and the correlation length (L), respectively (Lunn et al., 2005; Lynch, 2004). Since the optimal covariance structures may be case-dependent, the influence of these values was explored in a parametric study using synthetic data to investigate their effect on model estimation accuracy (see Section 2.6). By embedding the model equations (Eqn. 3) directly into the objective function, we have converted the constrained optimization problem into an unconstrained one. Eqn. 4 includes the unknown control variables $\mathbf{b}$ and $\mathbf{u}$ that must be constrained through the model equations. The Lagrange multipliers, λ , are additional scalar variables that serve as the constraint weighting factors, adding to the list of unknowns. The value of each Lagrange multiplier denotes the resistance of the objective function minimum to the violation of the constraint as determined by the corresponding model equation. Using the first-order conditions to determine the minimum of the objective function, the derivatives of with respect to $\mathbf{x}$, $\mathbf{b}$ and λ are set to zero, resulting in the following set of adjoint equations after some simplification (Eqn. 8 rewrites the definition of the model-data misfit; Lunn et al., 2005, 2006; Lynch, 2004):

$$\frac{\partial \Omega}{\partial \mathbf{x}} = \mathbf{K}^T \lambda - 2\mathbf{A}^T \mathbf{W}_\varepsilon \varepsilon = 0 \quad (5)$$

$$\frac{\partial \Omega}{\partial \mathbf{b}} = 2\mathbf{W}_b \mathbf{b} - \lambda = 0 \quad (6)$$

$$\frac{\partial \Omega}{\partial \lambda} = \mathbf{Kx} - \mathbf{b} = 0 \quad (7)$$

$$\varepsilon - (\mathbf{d} - \mathbf{Ax}) = 0 \quad (8)$$

Unlike REP that solves these equations directly for each component of displacement misfit, the iterative method adjusts the forcing conditions to search for the minimum of the scalar function () in the multidimensional space, $\mathbf{b}$. Suppose that the forcing vector at the k^{th} iteration is $\mathbf{b}_k$, and is adjusted at the next iteration according to:

$$\mathbf{b}_{k+1} = \mathbf{b}_k + \Delta \mathbf{b} \quad (9)$$

where the change in the forcing conditions, $\Delta \mathbf{b}$, is divided into its direction ($\partial \mathbf{b}$) and magnitude (α). Using the steepest gradient, the directional component is determined through:

$$\partial \mathbf{b} = -\partial \Omega / \partial \mathbf{b} \quad (10)$$

which specifies a line direction in the n -dimensional space (n is the total number of model constraint equations in Eqn. 3 or $4 \times m$, where m is the number of nodes in the model) along which the minimum of the scalar function, Ω , is sought. The magnitude of the forcing vector adjustment, α , is computed next. First, the effect of a change in the forcing vector, $\mathbf{b}$, is defined with the following equations:

$$\partial \mathbf{x} = [\mathbf{K}^{-1}] \partial \mathbf{b}, \text{ and } \partial \varepsilon = -[\mathbf{A}] \partial \mathbf{x} \quad (11)$$

which are linearly related to the magnitude, α :

$$\Delta \mathbf{b} = \alpha \times \partial \mathbf{b} \quad (12)$$

$$\Delta \mathbf{x} = [\mathbf{K}^{-1}] \Delta \mathbf{b} = [\mathbf{K}^{-1}] \alpha \times \partial \mathbf{b} \quad (13)$$

$$\Delta \varepsilon = \alpha \times \partial \varepsilon \quad (14)$$

Consequently, the value of the objective function, Ω , at the $(k+1)^{\text{st}}$ iteration is given by:

$$\begin{aligned} \Omega_{k+1} &= (\varepsilon_k + \Delta \varepsilon)^T \mathbf{W}_\varepsilon (\varepsilon_k + \Delta \varepsilon) + (\mathbf{b}_k + \Delta \mathbf{b})^T \mathbf{W}_b (\mathbf{b}_k + \Delta \mathbf{b}) \\ &= \Omega_k + \Delta \varepsilon^T \mathbf{W}_\varepsilon \varepsilon_k + \varepsilon_k^T \mathbf{W}_\varepsilon \Delta \varepsilon + \Delta \varepsilon^T \mathbf{W}_\varepsilon \Delta \varepsilon \\ &\quad + \Delta \mathbf{b}^T \mathbf{W}_b \mathbf{b}_k + \mathbf{b}_k^T \mathbf{W}_b \Delta \mathbf{b} + \Delta \mathbf{b}^T \mathbf{W}_b \Delta \mathbf{b} \end{aligned} \quad (15)$$

Assuming that $[\mathbf{W}_b]$ and $[\mathbf{W}_\varepsilon]$ are symmetric, the above expression is further simplified:

$$\begin{aligned} \Delta \Omega &= 2\alpha \cdot (\partial \varepsilon^T \mathbf{W}_\varepsilon \varepsilon_k) + \alpha^2 \cdot (\partial \varepsilon^T \mathbf{W}_\varepsilon \partial \varepsilon) \\ &\quad + 2\alpha \cdot (\partial \mathbf{b}^T \mathbf{W}_b \mathbf{b}_k) + \alpha^2 \cdot (\partial \mathbf{b}^T \mathbf{W}_b \partial \mathbf{b}) \end{aligned} \quad (16)$$

The optimal scalar magnitude is then determined by setting $\partial \{\Delta \Omega\} / \partial \alpha = 0$, leading to:

$$\alpha = - \frac{\{\partial \mathbf{b}\}^T [\mathbf{W}_b] \{\mathbf{b}\} + \{\partial \varepsilon\}^T [\mathbf{W}_\varepsilon] \{\varepsilon\}}{\{\partial \mathbf{b}\}^T [\mathbf{W}_b] \{\partial \mathbf{b}\} + \{\partial \varepsilon\}^T [\mathbf{W}_\varepsilon] \{\partial \varepsilon\}} \quad (17)$$

which dictates the amount of forcing vector advancement along the line parallel to $\partial \mathbf{b}$. The overall solution algorithm becomes:

1. Initialize stiffness, covariance and sampling matrices ($[\mathbf{K}]$, $[\mathbf{W}_\varepsilon]^{-1}$, $[\mathbf{W}_b]^{-1}$, $[\mathbf{A}]$).

2. Compute the best prior estimate of $\mathbf{b}$; set to $\mathbf{0}$ if unavailable.
3. Solve the forward model for $\mathbf{x}$: $\mathbf{K}\mathbf{x} = \mathbf{b}$
4. Sample the model-data misfit: $\varepsilon = \mathbf{d} - \mathbf{A}\mathbf{x}$.
5. Solve the adjoint model for λ : $[\mathbf{K}]^T \lambda = 2[\mathbf{S}^T \mathbf{W}_\varepsilon] \varepsilon$.
6. Evaluate $\nabla \Omega$ with respect to $\mathbf{b}$; $\mathbf{g} = 2[\mathbf{W}_b] \mathbf{b} - \lambda$; STOP if $\|\mathbf{g}\| \approx 0$.
7. Set the direction of descent (steepest gradient descent): $\partial \mathbf{b} = -\mathbf{g}$.
8. Project in the resulting direction: (a.) $\partial \mathbf{x} = [\mathbf{K}]^{-1} \partial \mathbf{b}$; (b.) $\partial \varepsilon = -[\mathbf{A}] \partial \mathbf{x}$.
9. Set the step size: $\alpha = - \frac{\{\partial \mathbf{b}\}^T [\mathbf{W}_b] \{\mathbf{b}\} + \{\partial \varepsilon\}^T [\mathbf{W}_\varepsilon] \{\varepsilon\}}{\{\partial \mathbf{b}\}^T [\mathbf{W}_b] \{\partial \mathbf{b}\} + \{\partial \varepsilon\}^T [\mathbf{W}_\varepsilon] \{\partial \varepsilon\}}$.
10. Assemble increments: $\mathbf{b} = \mathbf{b} + \alpha \times \partial \mathbf{b}$.

The most computationally intensive steps (3, 5 and 8a) involve the solution of full linear systems of equations ($3 \times Iter$ equations in total, with $Iter$ iterations). The number of sparse data points used to construct the estimate of $\mathbf{b}$ only affects the computation of the model-data misfit (step 4) and the calculation of the effect of $\mathbf{x}$ on the misfit (step 8b). Thus, the computational costs only increase marginally with the increase in N_{sprs} because no change in the rank of the full linear systems of equations is involved. Therefore, the overall solution computation with SGD does not scale directly with N_{sprs} . Since the linear least squares formulation has a single minimum, SGD will converge to the correct minimum, although the convergence rate may be slow, necessitating the use of a stopping criterion to terminate the iteration (Lynch, 2004). Consequently, instead of searching for the global minimum that requires infinite iterations, we used the minimum iterations ($Iter_{min}$) required for SGD to reach a model-data misfit at the sparse data locations (spr_s) statistically equal to or less than that obtained with REP for comparison purposes. In the following sections, we show that SGD is able to achieve such a minimization but with considerably fewer computations than REP, especially when N_{sprs} is large. In practice, since the REP solution would not generally be available, we recommend a stopping criterion of a change of less than 0.1% (of the initial model-data misfit) in successive iterations and show that this is effectively equivalent to the REP solution in the cases considered here.

To further reduce the computational costs of solutions of linear systems of equations that are the most computationally extensive in both SGD and REP algorithms, parallelization was exploited. A direct solver, PARDISO package (Schenk and Gartner, 2004), was employed for this purpose that explicitly factorizes the system matrix and uses forward elimination and backward substitution to complete the solution process (Gould et al., 2007), thereby eliminating the need for any iterations required in otherwise iterative solvers (e.g., SPLIB; Bramley R., Wang X., 1997). All model computations in this work were executed with the PARDISO direct solver on a shared-memory eight-processor Linux cluster (8G RAM in total and 2.33GHz for each processor), a computer platform to which we have ready access in the OR.

2.3 Patient Population and Model Development

To demonstrate the applicability of the SGD algorithm for minimizing the misfit between the model estimated displacement data and those measured, five ($N=5$) patients undergoing open cranial tumor resection were recruited at the Mary Hitchcock Memorial Hospital, under approval by the Institutional Review Board at the Dartmouth-Hitchcock Medical Center. The only criterion for patient inclusion was the presence of well-defined tumor (or cyst) boundaries in both iUS and pMR, which allows for accurate feature segmentation to determine intraoperative displacement. Patient age, gender, and type and location of tumor (cyst) are

summarized in Table 2. For each patient, preoperative T1-weighted axial MRIs (pMR; $N=124$; field of view: 240 mm; voxel size: $1.016 \text{ mm} \times 1.016 \text{ mm} \times 1.5 \text{ mm}$) were obtained from a 1.5 Tesla GE MR scanner with the use of a head coil. The brain was segmented through a two-step level set scheme (Wu et al., 2005) and meshed into tetrahedral elements (Sullivan et al., 1997), from which the triangular boundary elements were extracted. The average number of mesh nodes (in thousand) was 17.18 ± 3.05 , while it was 94.98 ± 17.48 and 8.58 ± 0.59 for the number of volume and boundary elements, respectively. Depending on the image contrast, either the tumor (patient 1, 3 and 5) or cyst (patient 2 and 4) in each pMR image was segmented semi-automatically using isointensity contours at a level determined by an expert (typical intensity values ranged 4000–6000), and a triangulated surface was generated to represent the undeformed tumor (or cyst; same as below) configuration.

2.4 Simulated and Clinical Sparse Data for Model Guidance

Patient 3 was randomly chosen to benchmark the performance of the SGD algorithm using a noise-free synthetic displacement field generated from an imposed forced-displacement solution. This was achieved by assigning the boundary nodes at the craniotomy site ($N=68$) to be free to move while those at the brainstem were stress free. All remaining boundary nodes were fixed in the direction normal to the skull, but allowed to freely move tangentially. Fluid pressure at the craniotomy nodes was fixed (zero pressure), while free at all remaining positions. A fluid plane passing through the “lowest” craniotomy node but perpendicular to gravity was defined. Elements below/above the fluid plane were saturated/unsaturated, and the surrounding fluid density was assigned accordingly (1000 kg/m^3 and 1 kg/m^3 for saturated and unsaturated regions, respectively; Ji et al., 2007; Lunn et al., 2005), which allowed the brain to deform under a gravitational load.

To simulate the measurements, displacements at the triangulated tumor surface (generated from pMR) boundary nodes were interpolated from the noise-free displacement field (average displacement of $0.75 \pm 0.17 \text{ mm}$ for all tumor boundary nodes). These nodes and their displacements were used as a pooled sample ($N_{tot}=2890$), from which a subset was randomly selected for model guidance. In order to parametrically study the influence of the covariance in measurement error and the correlation length on model performance, and L were logarithmically varied from 0.01 mm to 3.0 mm and from 0.1 mm to 10 mm, respectively (similar to the ranges used in Lunn et al., 2006), and the resulting model-data misfits at all mesh nodes obtained from SGD (results at the 50th iteration were used) and REP were compared. The optimal values of the covariance structures were obtained, and were then used to test the influence of measurement noise in the sparse data on model estimation performance. This was achieved by superimposing the sparse data displacements with noise from a normally distributed magnitude (zero mean and a parametrically varied std from 0.0 mm to 1.0 mm) and a uniformly distributed direction. These parametric investigations ($N_{sprs}=200$ in these simulations) allow for evaluation of the algorithm’s tolerance to the covariance structures and measurement noise, which is important in clinical settings, where the statistical behavior of the measurement noise is generally less predictable or unknown.

Finally, we investigated the influence of N_{sprs} on the accuracy of the model estimate. This was accomplished by using the optimal covariance structures and a fixed level of measurement noise (normally distributed zero mean and a standard deviation of 0.5 mm in the magnitude of displacement error and uniform distribution in direction) but different numbers of sparse data points ($N_{sprs}=10, 100$ and 200) to generate whole-brain deformations using SGD and REP that were compared.

2.5 Intraoperative Tumor Displacement

Prior to surgery, the patient's head was rigidly registered with its pMR through a group of MR-compatible fiducial markers (typically $N=10-15$) glued externally to the scalp. The markers were manually identified in the pMR image space through a dedicated interactive software tool that allows the user to move a cursor until the identified marker is visually symmetrical in all orthogonal cross-sectional images. The centers of the fiducials were subsequently digitized in the OR using a stylus tracked by a Polaris system (Northern Digital Inc., London ON, Canada). Although a rigid-body transformation can be determined from two sets of paired homologous points through closed-form solution (e.g., singular value decomposition (Arun et al., 1987)), it is not always practical or very convenient to pair the points in the surgical setting. The closed-form approach is prone to errors and inefficiencies even for experienced operators because the precise order of fiducial identification matters and often some fiducials are not even visible to the tracking system because of the relative orientation of the patient's head, making the number successfully identified in the OR (typically, $N=7-10$) unpredictable and generally smaller than those found in pMR. Thus, keeping track of the order of fiducial identification in the OR requires extra surgeon/operator interaction and communication at the time of the procedure. The additional task of documenting order, which was prone to error because of the lack of a natural or well-defined order for localizing fiducials on the patient's head coupled to the uncertainty of which fiducials would be visible in any given surgery, motivated us to develop an alternative patient registration method that does *not* require the same degree of manual record-keeping and verbal interaction with the surgeon. Specifically, instead of manually pairing the points to be used for registration, we applied a genetic search (Hartov et al., 2007) to find the 3D locations in the OR that best matched with those in pMR by computing the closed-form solution for rigid-body registration (through SVD) for each of the candidate pairings presented during the search and keeping the one that produced the minimum registration error. Because of the potential for a mis-paired set of points to occasionally yield the minimum registration error (e.g., in cases where the fiducials possess a particularly high degree of medial-lateral and/or anterior-posterior spatial symmetry), the resulting pre-operative patient registration was always visually verified by comparing the actual head orientation in the OR with the relative orientation of the gravitational vector (obtained by placing the stylus at two vertical locations in the OR) transformed into the pMR image space with respect to the patient's scalp surface (rendered from the pMR images).

For each patient, a series of 2D B-mode iUS images before and after dural opening (typically $N=20-200$; Siemens Sonoline Sienna ultrasound system with C8-5 transducer model; image size: 640×480 ; 8-bit grayscale) were digitized through a frame grabber (DT 3155, Data Translations Inc., Marlboro, MA; image acquisition time: 50–60 ms). The depth scale of the iUS images in general could vary from 30 mm to 160 mm, and was chosen depending on the optimal imaging window. All images used in this study had an iUS scan depth of 60 mm, resulting in a pixel size of $0.15 \text{ mm} \times 0.16 \text{ mm}$. The US transducer was rigidly attached to a tracker that was visible to the Polaris system through a group of infrared light emitting diodes, and the position and orientation of the transducer were continuously monitored. The US image coordinate system has previously been calibrated against the “world” coordinates determined by the Polaris (Hartov et al., 1999). This calibration, together with the rigid-body registration between the patient head in the OR and pMR allowed anatomical features in the iUS images to be spatially merged into the pMR image space.

However, tumor boundaries in iUS images before dural opening do not necessarily coincide with their pMR counterparts when using the rigid fiducial-based registration alone, due to errors from the fiducial registration itself, ultrasound scan-head calibration and even brain shift at this early stage of surgery, thereby degrading the accuracy of local tumor alignment in the two image sets. The mean fiducial registration error (FRE; West et al., 1997) was 3.7 mm for the

pooled patients (range of 2.6–5.1 mm). Similar magnitude of tumor boundary misalignment (average of 2.8 mm, range of 1.8–4.9 mm) was present in terms of closest point projection distance (CPPD; Ji et al., 2008a; similar to that used in Maurer et al., 1998). To account for these errors, a mutual-information-based re-registration scheme (Ji et al., 2008a) was employed that adjusted the transformation between iUS and pMR to minimize the misalignment of tumor before dural opening, leading to an improved or a last-known-correct local tumor registration (Hartov et al., 2007). With the pooled patient cases, tumor boundary misalignment was reduced to an average of 1.0 mm (0.7–1.3 mm) in terms of CPPD. Following the re-registration scheme, manually segmented tumor boundaries in iUS images after dural opening were spatially merged with the undeformed pMR tumor surface. An iterative closest point procedure (ICP; Besl and McKay, 1992) was employed to determine a rigid-body transformation such that the overall distance from the transformed tumor boundary points in iUS relative to the pMR tumor surface was minimized. Tumor displacement after dural opening was essentially the difference between the tumor boundary points at this stage of surgery and their counterparts after applying the transformation. The single transformation generated locally rigid-body motion of the tumor (both translation and rotation), which was used to approximate the nonrigid motion in this region (Lunn et al., 2005). To further minimize the adverse effect of segmentation errors especially from iUS, points only in regions of well-defined tumor boundaries were included. In addition, cyst boundaries were used for patients 2 and 4, due to their high contrast. The average tumor misalignment as a result of ICP in terms of CPPD was 1.4 mm after dural opening for the pooled sample of 5 patient cases. The total number of data points with measured displacement for each patient (N_{tot} ; i.e., the pooled sample) and the resulting initial model-data misfit before model correction ($init$) are given in Table 4 in the Results section. Similarly to the simulated data, N_{sprs} ($N_{sprs} = 10, 100$ and 200) sparse data points were randomly selected from the pooled sample for model guidance.

2.6 System Accuracy

The ultimate goal of estimating brain deformation intraoperatively is to compensate for brain shift in the registered images displayed in the OR. Because multiple steps are involved in model-based image compensation, the accuracies of the individual components of the approach contribute to the overall fidelity of the system and need to be evaluated. The steps that influence the overall accuracy of the method include patient registration, calibration of intraoperative imaging techniques (e.g., US scan-head and stereovision), intraoperative data acquisition of features either deep in the brain (e.g., tumor) or on the exposed cortical surface, measurement of local brain displacements, and computational data assimilation and motion estimation. The accuracies of these components of our compensation scheme have been evaluated in detail in previously published studies, and are summarized here as a frame of reference for judging the computational results reported in Section 3.

The typical accuracy of skin-affixed fiducial-based patient registration has been well documented to be 2–5 mm (Helm and Eckel, 1998; Ammirati et al., 2002; and Wolfsberger et al., 2002), and we have found similar results in our experience (Roberts et al., 1998; Roberts et al., 1999; Hartov et al., 2008; Ji et al., 2008a). US scan-head calibration error can be as high as 2–3 mm (Hartov et al., 1999). When US images are spatially transformed into the pMR image space, these errors combine and potentially compromise the accuracy of the inter-modality image registration. Controlled studies in the porcine brain (Lunn et al., 2003) showed that features in iUS could be transformed into coregistered preoperative scans with an accuracy of 1–2 mm. However, in clinical cases, we have found this degree of fidelity to degrade to 2–3 mm or more (Hartov et al., 2008). To improve the coregistration accuracy, we have recently developed an image-based re-registration technique that collectively minimizes errors due to fiducial-based registration as well as US scan-head calibration at the start of surgery. In six patient cases, we showed that this technique was able to minimize tumor misalignment in US

and pMR to achieve an accuracy of approximately 1 mm with 1–2 minutes of additional computation (Ji et al., 2008a). The resulting “last-known-correct” registration at the start of surgery is critical for maintaining the accuracy of subsequent measurements of tumor displacements, the details of which have been described in Section 2.5 as well as in (Ji et al., 2009a).

Another important intraoperative imaging technique that is frequently employed in our group is stereovision that captures the profile of the exposed cortical surface during surgery. The accuracy of the stereopsis system was found to be within 1 mm in a brain-shaped phantom (Sun et al., 2005b) and between 1–2 mm in patient cases where the stereopsis was compared to gold standard validation measurements obtained with fiducial markers implanted in the bone comprising the rim of the craniotomy (Sun et al., 2005a). This intraoperative imaging scheme has enriched the amount of displacement data available, which can be used as information either to drive the model deformation or to verify the accuracy of the model estimates (Ji et al., 2009b).

The accuracy of the computational data assimilation has also been previously evaluated in animals where a group of stainless-steel beads were implanted in the brain (Lunn et al., 2006). By simulating a 10-mm retraction, the mean model-data misfit between the measured bead locations ($N=21$) in actual CT scans and those calculated from the model deformation was only 0.9 mm. The mean model-data misfit was 1.6 mm when optimal values of the covariance length, L , and the covariance in measurement error, σ , were used to simulate balloon inflation (maximum displacement at the balloon surface was 14 mm). In both experiments, the inverse data assimilation scheme was able to recover approximately 90% of the induced brain deformations, demonstrating the effectiveness of the process.

Although a separate phantom study undergoing the equivalent clinical procedures investigated here has not been repeated, an overall system accuracy of 1–2 mm can be assumed based on the individual accuracies of the major steps involved (e.g., error in tumor displacements after dural opening of 1.4 mm in terms of CPPD; Section 2.5). Consequently, we expect that the model-based brain shift compensation scheme that estimates whole-brain deformation by assimilating intraoperative displacement data to be well worth the effort given that the accuracy and efficiency of the system components are 1–2 mm; yet, the uncompensated brain shift is typically greater than 5 mm and often reaches 10 mm and more in many surgeries. A detailed analysis of the composite system performance where the errors in its individual components are allowed to fully interact in order to determine the overall accuracy and computational efficiency of the approach remains to be completed and is important work to be reported in the future.

2.7 Data Analysis

For both the synthetic and clinical displacement data, the remainder of the pooled sample that was not selected for model guidance was defined as the cross-validation locations ($N_{cross} = N_{tot} - N_{sprs}$), whose simulated or “ground-truth” (clinical) displacements were used to evaluate the accuracy of model estimates generated by SGD and REP. The boundary conditions used in both inverse algorithms were identical to those in the forced-displacement method described in Section 2.4, in which a gravitational load induced the model deformation.

The synthetic data were used to parametrically investigate the influence of the covariance in measurement error and the correlation length on model estimation accuracy, from which the optimal covariance structures were obtained. Using these optimal covariance structures, the influence of measurement noise on model performance was also investigated, which was achieved by varying the level of noise (i.e., standard deviation of the magnitude of displacement error) and comparing the resulting model-data misfits at sparse and cross-validation locations

(defined as ϵ_{cross} for the latter). Finally, a fixed level of measurement noise was used to investigate the influence of the number of sparse data points on model performance by randomly selecting 10, 100, and 200 sparse data points from the pooled sample to guide the model. Statistical comparisons were performed using a one-tailed paired t -test, with a null hypothesis that the mean values of the two samples being compared were equal.

To determine $Iter_{min}$ (the minimum iterations required for SGD to achieve a statistically smaller ϵ_{sprs} than that produced by REP), we also performed paired t -tests with a null hypothesis that the mean ϵ_{sprs} given by SGD was equal to that produced by REP. Because in general $N_{cross} \gg N_{sprs}$, ϵ_{cross} was used to evaluate model performance. Similarly, a one-tailed paired t -test was performed to determine whether ϵ_{sprs} and ϵ_{cross} (using the intersected set of two cross-validation locations) decreased with the increase in N_{sprs} . Statistical significance was defined at $p=0.05$ for all tests.

REP and SGD solutions were computed on an Intel Xeon shared-memory (8G RAM), eight-processor (2.33GHz for each processor) Linux cluster running Ubuntu 6.10, and all data analysis was performed in Matlab (Matlab 7.3; The Mathworks, Natick, MA).

3. Results

The mean model-data misfit obtained from SGD and REP by parametrically varying the covariance in measurement error (σ) and the correlation length (L ; Fig. 1) demonstrate that model performance depends on across the tested values (similarly to that reported in Lunn et al. 2006), regardless of which algorithm was applied. However, model performance depended to a lesser degree on L except when L was large (>2 mm). The optimal values of σ and L based on the previous investigation (0.2 mm and 5 mm, respectively) appear to be equally applicable in this study as the mean model-data misfit was approximately the minimum for both algorithms at these values, which is reassuring since SGD is only expected to influence the overall computational efficiency.

Using these optimal values of σ and L but varying the level of measurement noise (i.e., the standard deviation of the magnitude of the imposed displacement error that followed a normal distribution), ϵ_{sprs} , the model-data misfits at the sparse data locations, increased with the increase in the noise level, as expected (Fig. 2a). In addition, SGD and REP generated similar ϵ_{sprs} across the range of tested noise levels, although ϵ_{cross} , the misfit at the cross-validation locations, generated from SGD was significantly smaller than that from REP, regardless of the level of noise, indicating an improved tolerance to measurement noise with the SGD algorithm. This was also evident when ϵ_{sprs} and ϵ_{cross} from SGD were plotted against the number of SGD iterations, and compared with those obtained from REP (Fig. 2b; noise level of 0.5 mm). Despite the relatively high level of noise (considering the average displacement around the tumor was only 0.75 mm) that led to large values of ϵ_{sprs} (0.41 mm), much smaller values of ϵ_{cross} were generated from SGD (0.05 mm) and REP (0.13 mm; Fig. 2b). Both ϵ_{sprs} and ϵ_{cross} from SGD decreased negligibly after $Iter_{min}$ (38 in this case).

The difference in model-estimation accuracy was also apparent when the whole-brain deformation generated from SGD and REP were compared using the identical set of sparse data (with standard deviation of 0.5 mm in the magnitude of measurement noise) and the optimal covariance structures for model guidance. The average difference between the ground-truth and the model estimation by SGD and REP was 0.10 ± 0.06 mm and 0.32 ± 0.12 mm, respectively. The difference was smallest near the tumor boundary for both algorithms (0.04 ± 0.03 mm and 0.13 ± 0.08 mm for SGD and REP, respectively), while larger in regions far away (0.15 ± 0.07 mm and 0.53 ± 0.23 mm for SGD and REP, respectively), which was not surprising because only sparse data from the tumor boundary were generated and used for model-

guidance. Nevertheless, SGD achieved similar model-data misfit at the sparse data points relative to REP, but improved the overall model estimation accuracy, especially at points more distant from the guiding data, when measurement noise was present, indicating it is well suited to clinical applications.

Using different numbers of sparse data points to guide the brain model, spr_{rs} showed a statistically significant decrease with increase in $N_{spr_{rs}}$ for both algorithms, as expected ($p < 0.01$), although modest in size due to the high level of noise in the measurement data compared to the initial model-data misfit (decreased by 31% and 33% for SGD and REP, leading to a residual spr_{rs} of 0.41 mm and 0.42 mm, respectively, slightly less than the level of simulated noise in the data). Similarly, $cross$ decreased with increase in $N_{spr_{rs}}$ ($p < 0.01$; by 82% and 29% for SGD and REP, respectively), recovering 94% and 82% of brain deformation using SGD and REP, respectively, when $N_{spr_{rs}}$ was sufficiently large (e.g., $N_{spr_{rs}} = 200$; residual $cross$ was reduced to about 0.046 mm and 0.132 mm, respectively, much smaller than the noise level in the data). These results demonstrate the robustness of both inverse algorithms in minimizing model-data misfit, especially considering the high level of noise in the data relative to the initial model-data misfit of 0.75 mm. Interestingly, the mean $cross$ given by SGD was consistently smaller than its REP counterpart at any given $N_{spr_{rs}}$, and the difference was significant (up to 0.10 mm or 29% at $N_{spr_{rs}} = 10$), indicating that SGD may be more tolerant of measurement error than REP. While this finding may be specific to the data, covariance structures and the measurement noise used in the study, it does suggest that SGD may offer benefits in clinical settings where the amount of measurement error is generally unknown.

$Iter_{min}$ as well as the computational costs of SGD and REP increased with $N_{spr_{rs}}$ (Table 3). However, when normalized with respect to the total computational cost of REP, the relative computation time of SGD decreased with $N_{spr_{rs}}$ and was less than 25% of that of REP when a large number of sparse data points was used for model guidance (e.g., $N_{spr_{rs}} > 100$; Table 3). These results indicate that SGD is able to achieve improved accuracy in model estimates relative to REP with significantly fewer computations.

With clinical data assimilated for model guidance, the residual spr_{rs} and $cross$ for each $N_{spr_{rs}}$ were normalized with respect to $init$ (initial model-data misfit before model correction), and were pooled for all datasets. Both the normalized spr_{rs} and $cross$ decreased with increase in $N_{spr_{rs}}$ regardless of the algorithm used, as evidenced in Fig. 3 ($p < 0.01$; average normalized spr_{rs} reduced by 18% and 12%, and $cross$ reduced by 19% and 15% for REP and SGD, respectively). Similarly with the synthetic data, both normalized spr_{rs} and $cross$ produced by SGD were consistently smaller than those given by REP regardless of the number of sparse data points used (Fig. 3), indicating that SGD may still outperform REP when a non-optimal covariance model is used in clinical settings. The model performances using SGD and REP for each patient with $N_{spr_{rs}} = 200$ are summarized in Table 4. Both SGD and REP minimized model data misfit and the recovered deformation ranged from 70% to 87% (mean of 80% and 75% across all cases, respectively).

$Iter_{min}$ for each dataset with each $N_{spr_{rs}}$ was obtained to determine the resulting execution time of SGD. The computational costs of SGD and REP for each individual patient are reported in Table 5 along with the numbers of mesh nodes. With larger $N_{spr_{rs}}$ (e.g., $N_{spr_{rs}} \bullet 100$), the run times of SGD were approximately 2.5 minutes or less (except for patient 4 which had the largest mesh), whereas those of REP almost all exceeded 5 minutes (up to 13 minutes for patient 4). The run times of SGD were also normalized by those of REP (tabulated in Table 6 and plotted in Fig. 4). Although $Iter_{min}$ increased with the increase in $N_{spr_{rs}}$, the corresponding normalized execution time of SGD decreased. With a sufficiently large number of sparse data points ($N_{spr_{rs}} > 100$), SGD execution time was about 25% of that required by REP, representing a 4-fold reduction in computational effort. This finding is consistent regardless of whether

simulated or clinical data were used for model guidance, suggesting that SGD is well-suited to use in the OR when more sparse data are available, such as when applying 3D ultrasound. When only a few sparse data points are available, however, REP may be the better candidate.

The purpose of data assimilation is to generate a deformation field where the model-updated location and shape of the tumor match the iUS. Patient 3 was used to illustrate the procedures involved in the model-based brain shift compensation employed in this study. In order to begin with the last-known-correct registration between iUS and pMR, tumor in iUS before dural opening was rigidly displaced to match the pMR through maximization of the normalized mutual-information. IUS images after dural opening were subsequently transformed into the pMR image space, and the misalignment between tumor boundaries in the two image volumes demonstrated the gross tumor displacement as the result of dural opening (Fig. 5a). The resulting displacement data were randomly selected for model guidance using SGD and REP ($N_{sprs}=200$). Based on the computed whole-brain deformation, the pMR image was modified (with a computational cost of approximately 6sec) and overlaid with iUS images after dural opening (Fig. 5cd). Clearly, both SGD and REP generated comparable deformations near the tumor, as evidenced by the alignment of the model-updated tumor boundaries and their counterparts in iUS. However, model solutions from SGD and REP differed in regions farther away from the tumor, as indicated by the varying misalignment of the tentorium (Fig. 5cd). Prominent non-rigid tumor deformation was clearly demonstrated with the large discrepancy between the model-updated and pMR tumor cross-sections (solid and dashed lines in Fig. 5cd), the latter of which was *rigidly* transformed to the coordinate system of iUS after dural opening.

For all other patients, a typical iUS image after dural opening was overlaid against its corresponding pMR oblique slice (Fig. 6). The undeformed tumor (dashed thin lines in Fig. 6) is represented by the cross-section of the pMR tumor surface intersected by the iUS imaging plane. When the tumor was displaced according to the model estimated deformation field generated by SGD or REP with 200 sparse data points ($N_{sprs}=200$), the resulting tumor surface cross-section (solid/dashed thick lines as generated by SGD/REP) represented the model updated tumor location and shape. Both algorithms produced similar magnitudes of residual $_{cross}$. However, visual inspection also indicated that in certain regions (arrows in Fig. 6), SGD appears to have outperformed REP in matching the tumor surface in iUS images.

4. Discussion

Computational efficiency is important for effective deployment of biomechanical brain models for intraoperative brain shift compensation, especial for inverse methods where the deformation is obtained through an estimation scheme to match model response with measurements at sparse data points rather than computed directly. The computational inefficiency of the REP algorithm with large numbers of sparse data points (more than an hour with $N_{sprs}=7$; Lunn et al., 2006) has prompted us to investigate and evaluate an alternative algorithm – SGD to solve the identical set of adjoint equations and to increase the amount of sparse data used to improve model performance. Using a simulated dataset and five actual clinical cases, we have demonstrated that increasing the amount of sparse data used for model guidance significantly improves the accuracy of the model deformation estimates, and that SGD is able to achieve an estimate comparable to or better than that generated by REP but with far fewer computations, especially when the amount of sparse data is large. With sufficient sparse data ($N_{sprs}=200$), the SGD and REP algorithms captured 74% to 82% and 69% to 82% of the subsurface tumor displacement in the clinical datasets, respectively. The SGD and REP algorithms achieved a mean model-data misfit of 1.1 mm and 1.4 mm in the five patient cases, respectively, which is in line with the accuracy of the overall system estimated to be 1–2 mm, indicating that both inversion algorithms are acceptable for clinical application in the OR from the accuracy point-of-view. However, the computational cost of SGD was only 25% of that of

REP, resulting in a 4-fold reduction in run-time with a comparable or improved model estimate when a sufficient amount of sparse data was used. The average run time for SGD and REP was 2 min and 5.5 min with N_{sprs} of 100, respectively. When the number of sparse data points used for model guidance was doubled (i.e., N_{sprs} of 200), the average run time for SGD was increased by only 0.4 min whereas it was almost doubled for REP to approximately 10 min as expected. These results indicate that although SGD and REP are both able to achieve comparable model estimation accuracies, only SGD is practical for clinical application in the OR from the efficiency point-of-view, especially as the amount of intraoperatively-acquired displacement data is increased.

Because the most time consuming part of both inversion algorithms is the solution of systems of equations that contribute over 90% of the total computational cost when N_{sprs} is large, it is advantageous to exploit parallelization. In this work, a parallelized direct solver, PARDISO package (Schenk and Gartner, 2004), was employed for this purpose. When a sequential, uni-processor iterative sparse solver (e.g., SPLIB; Bramley and Wang, 1997) was used that does not take advantage of multiprocessor parallelization, the run times were prohibitively long (e.g., up to 3.6 hours for SGD and 18 hours for REP with N_{sprs} of 200). The SGD computational efficiency with the PARDISO package was similar to that achieved in Dumpuri et al. (2007) with the Portable, Extensible Toolkit for Scientific Computation (PETsc; Balay et al., 2004) package (140 sec with 8 CPUs, and 117 sec with 16 CPUs). In this paper, we have used REP to benchmark the computational efficiency of SGD, whereas in practice a stopping criterion is needed when SGD is used alone in the OR. In this case, the minimum number of iterations ($Iter'_{min}$) needed for SGD to reach an approximate convergence can be empirically defined as the one where the net change in the model-data misfit relative to the previous iteration is less than 0.1% of the initial model-data misfit. Based on the five patient cases in this study, we have found that $Iter'_{min}$ is comparable to $Iter_{min}$ (maximum difference of 6 iterations and maximum difference in residual error of 0.1 mm or less), and therefore, similar run times (difference within 12 seconds) are achieved with this practical stopping criterion. Finally, the overall computational cost also depends on the number of nodes or the degrees-of-freedom in the brain mesh (e.g., patient 4 had the largest number of nodes, which led to the longest run time; Table 5). Therefore, one simple strategy to achieve additional reduction in computational cost is to refine the brain mesh *locally* around the surgical area of interest rather than *globally* as in this study. These findings suggest that routine computational model updates can be achieved in the OR with SGD. The efficiency of the computational scheme can be improved further through code optimization (e.g., by parallelizing other parts of the code in addition to the solutions of systems of equations and by eliminating disk I/O used to output nodal displacements at each SGD iteration as in this study), applying a practical stopping criterion, and by using locally refined meshes for model computation.

The relative merits of using methods that enforce displacements as a type of boundary condition or invoke inversion processes that approximate the measured displacements remain an open question. Forced-displacement methods are usually simple to code and are computationally efficient, while inversion schemes are more difficult to implement and involve more extensive computations. However, forced-displacement methods are sensitive to measurement noise because measurement errors that are always present are directly encoded into the computed whole-brain deformation. When ICP or other rigid-body registration methods (e.g., mutual-information-based rigid registration; Ji et al., 2008a, Ji et al., 2009a) are used to measure tumor displacement, non-rigid tumor deformation cannot be recovered through forced-displacement methods (because the rigidly-acquired displacements are directly imposed in the model solution), which potentially would degrade the accuracy of the model-updated pMR images, especially when non-rigid deformation is prominent. In these cases, a non-rigid registration must be applied instead to generate more faithful tumor displacements in order to correctly assign the data constraints (e.g., see Fig. 5) Because the computational cost of non-rigid

registrations (Holden, 2008) can easily be more expensive than the SGD inversion algorithm used in this study (the computational cost of the latter is similar to that of mutual-information-based *rigid* registration between pMR and iUS (1–2 min); Ji et al., 2008a), the computational effectiveness of forced-displacement methods may be compromised. By contrast, inversion schemes are less prone to measurement errors, and tumor displacements generated from rigid (as opposed to non-rigid) registrations can be directly applied to *guide* the model in order to achieve non-rigid tumor deformation (see Fig. 5).

Our inversion methods actively construct the right-hand-side vector and potentially introduce non-zero forcing terms for nodes that we have assigned as having stress-free boundary conditions (e.g., surface nodes exposed at the craniotomy or in the brainstem region). However, the magnitude and directionality of these terms are determined by the approximation scheme in order to minimize the model-data misfit at the sparse data points, and are, therefore, *post priori*. For example, the magnitudes of the force vectors estimated to exist on the surface nodes exposed at the craniotomy typically range from 10^{-5} to 10^{-3} Newtons after SGD convergence, whereas the range is 10^{-14} to 10^{-13} for nodes in the brainstem region where no or minimal displacements are present. Apparently, the values of the reconstructed right-hand-side forcing terms depend on the exact nature of the brain motion (e.g., sagging, bulging, or both). Forced-displacement methods essentially construct the right-hand-side terms corresponding to the nodes where displacement is measured through boundary condition assignment, and the nature of these terms is, therefore, *a priori*. Finally, it is important to note that the right-hand-side forcing vectors produced by the inversion process are expected to be realistic under the presumption that the underlying model equations and mechanical properties are realistic representations of brain tissue because the computed displacement fields arise from the model system driven by the same inferred forcing conditions which matched the model response to the measurement data. We do not currently constrain nodes associated with the exposed parenchyma resulting from craniotomy to remain stress free, but it would be relatively easy to introduce such a constraint during the data assimilation and displacement estimation process which is an idea well worth exploring in the future.

The distinction between forced-displacement and inversion methods can be minimized when both are applied for brain shift compensation. Specifically, using displacements generated from a forced-displacement method as the initial starting point for SGD solution may reduce the number of SGD iterations at a cost of one additional forward model solution (both methods share essentially identical pre-processing, e.g., global stiffness matrix assembly), leading to further reduction in the overall computational cost in this case.

Our inversion approach does not rely on heavy pre-computations as in the case of the very promising technique reported in Dumpuri et al. 2007 (10 hours required to compute an atlas with 320 basis solutions with 64 different gravitational vectors and 5 levels of fluid saturation). Although pre-computations are often desirable (if they significantly reduce the computational burden in the OR as in Dumpuri et al., 2007), the cost of these computations grows exponentially when more variables (e.g., administration of mannitol, size of craniotomy) and denser variations are considered, which may ultimately become prohibitive. In addition, the Dumpuri et al. technique essentially interpolates motion of the brain from pre-computed basis solutions. It is important, therefore, that the intraoperative values of the variables (e.g., gravitational direction) lie within the ranges of the quantities considered in the construction of the basis solutions. When the pre-computed basis solutions incompletely represent the true motion of the brain, *extrapolation* (as opposed to interpolation) is otherwise needed to generate brain deformation, the accuracy of which remains to be explored.

The accuracy and efficiency of the overall procedure of brain shift compensation depends on the composite accuracy and efficiency of all of the steps involved (e.g., brain segmentation

and mesh generation, measurement of local brain displacements at the tumor boundary and/or cortical surface, fidelity of the computational data assimilation as well as the brain-skull boundary conditions). We are currently integrating the overall model-based scheme for brain shift compensation in the OR, and we have planned a controlled study to validate the accuracy of the overall approach. Results from these investigations will be reported in the future.

5. Conclusion

We have introduced an iterative SGD algorithm to solve the set of adjoint equations for recovering whole-brain deformation and have compared the model performance and computational efficiency of SGD relative to the REP algorithm. With five clinical datasets, we show that the average initial model-data misfit of 5.2 mm was reduced to 1.1 mm and 1.4 mm by applying SGD and REP, respectively, suggesting that SGD is able to recover whole-brain deformation with accuracy comparable to or better than that determined with REP. Importantly, the computational cost of SGD is significantly less than that of REP (a 4-fold reduction) when the amount of sparse data is large. The average run time for the pooled patient cases was approximately 2 min with SGD when 100 sparse data points were assimilated. In addition, the accuracy of the model estimates increases as the amount of measurement data used to guide the model increases and SGD has enabled the incorporation of hundreds of data points that was not previously practical or possible with REP. Together with further code parallelization, local mesh refinement and a practical stopping criterion for SGD as well as with the availability of volumetric 3D ultrasound that will generate more complete volumetric data on intraoperative tumor displacement, we expect that SGD will play an important role in intraoperative brain shift compensation in the OR in the near future.

Acknowledgement

Funding from the National Institutes of Health grant R01 EB002082-11 is acknowledged. The authors are grateful to Dr. Andrea Borsic at the Dartmouth College for helpful discussion on utilizing the PARDISO software package.

References

- Ammirati M, Gross JD, Ammirati G, Dugan S. Comparison of Registration Accuracy of Skin-and Bone-Implanted Fiducials for Frameless Stereotaxis of the Brain: A Prospective Study. *Skull Base* 2002;129(3):125–130. [PubMed: 17167664]
- Arun KS, Huang TS, Blostein SD. Least squares fitting of two 3D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1987;9:698–700.
- Balay, S.; Buschelman, K.; Eijkhout, V.; Groppe, WD.; Kaushik, D.; Knepley, MG.; McInnes, LC.; Smith, BF.; Zhang, H. 2001. Petsc webpage: <http://www.mcs.anl.gov/petsc>
- Besl PJ, McKay ND. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 1992;14(2):239–256.
- Biot M. General theory of three-dimensional consolidation. *J. Appl. Phys* 1941;12:155–164.
- Black PM, Moriarty T, Alexander E III, et al. Development and implementation of intraoperative magnetic resonance imaging and its neurosurgical applications. *Neurosurgery* 1997;41:831–845. [PubMed: 9316044]
- Bonsanto MM, Staubert A, Wirtz CR, Tronnier V, Kunze S. Initial experience with an ultrasound-integrated single-rack neuronavigation system. *Acta Neurochirurgica* 2001;143(11):1127–1132. [PubMed: 11731863]
- Bramley, R.; Wang, X. SPLIB: a library of iterative methods for sparse linear systems. Bloomington, IN: Department of Computer Science, Indiana University; 1997. (<http://www.cs.indiana.edu/ftp/bramley/splib.tar.gz>)
- Bucholz RD, Smith KR, Laycock KA, McDurmont LL. Three-dimensional localization: From image-guided surgery to information-guided therapy. *Methods (Duluth)* 2001;25(2):186–200.

- Carter TJ, Sermesant M, Cash DM, Barratt DC, Tanner C, Hawkes DJ. Application of soft tissue modeling to image-guided surgery. *Med. Eng. & Physics* 2005;27:893–909.
- Clatz O, Delingette H, Talos IF, Golby AJ, Kikinis R, Jolesz FA, Ayache N, Warfield SK. Robust nonrigid registration to capture brain shift from intraoperative MRI. *IEEE Trans Med Imaging* 2005;24(11):1417–1427. [PubMed: 16279079]
- Comeau RM, Sadikot AF, Fenster A, Peters TM. Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. *Medical Physics* 2000;27:787–800. [PubMed: 10798702]
- Dumpuri P, Thompson RC, Dawant BM, Cao A, Miga MI. An atlas-based method to compensate for brain shift: Preliminary results. *Medical Image Analysis* 2007;11:128–145. [PubMed: 17336133]
- Ferrant M, Nabavi A, Macq B, Jolesz F, et al. Registration of 3D intraoperative MR images of the brain using a finite element biomechanical model. *IEEE Trans Medical Imaging* 2001;20:1384–1397.
- Ferrant M, Nabavi A, Macq B, Black PM, Joles FA, Kinkinis R, Warfield SK. Serial registration of intraoperative MR images of the brain. *Medical Image Analysis* 2002;6:337–359. [PubMed: 12426109]
- Gould NIM, Scott JA, Hu Y. A numerical evaluation of sparse direct solvers for the solution of large sparse symmetric linear systems of equations. *ACM Trans. Math. Softw* 2007;33(2)Article 10
- Hall WA, Liu H, Martin AJ, Pozza CH, Maxwell RE, Truwit CL. Safety, efficacy, and functionality of high-field strength interventional magnetic resonance imaging for neurosurgery. *Neurosurgery* 2000;46:632–642. [PubMed: 10719860]
- Hajnal, JV.; Hill, DLG.; Hawkes, DJ., editors. *Medical Image Registration*. Biomedical Engineering Series. M. N. Neuman Ed. CRC Press; 2001.
- Hartkens T, Hill DLG, Castellano-Smith AD, Hawkes DJ Jr, Maurer CR, Martin MJ, Hall WA, Liu H, Truwit CL. Measurement and analysis of brain deformation during neurosurgery. *IEEE Transactions on Medical Imaging* 2003;22(1):82–92. [PubMed: 12703762]
- Hartov A, Roberts DW, Paulsen KD. A comparative analysis of co-registered ultrasound and MRI in neurosurgery. *Neurosurgery* 2008;62(3 Suppl 1):91–101. [PubMed: 18424971]
- Hartov A, Eisner SD, Roberts DW, Paulsen KD, Platenik LA, Miga MI. Error analysis for a free-hand three-dimensional ultrasound system for neuronavigation. *Neurosurg Focus* 1999;6(3)Article 5
- Helm PA, Eckel TS. Accuracy of Registration Methods in Frameless Stereotaxis. *Computer Aided Surgery* 1998;3:51–56. [PubMed: 9784952]
- Holden M. A Review of Geometric Transformations for Nonrigid Body Registration. *IEEE Trans. Med. Imaging* 2008;27(1):111–128. [PubMed: 18270067]
- Hu J, Jin X, Lee JB, Zhang L, Chaudhary V, Guthikonda M, Yang KH, King AI. Intraoperative brain shift prediction using a 3D inhomogeneous patient-specific finite element model. *J. Neurosurg* 2007;106:164–169. [PubMed: 17236503]
- Liu, F.; Paulsen, K.; Lunn, K.; Sun, H.; Hartov, A.; Wu, Z.; Roberts, D. Comparative Study of Brain Deformation Estimation Methods. In: Cleary, Kevin R.; Galloway, Robert L., Jr, editors. *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*; Proceedings of SPIE Vol. 6141 (SPIE, Bellingham, WA 2006) 61411D: 2006.
- Lunn KE, Paulsen KD, Roberts DW, Kennedy FE, Hartov A, West JD. Displacement Estimation with Co-Registered Ultrasound for Image Guided Neurosurgery: A Quantitative In Vivo Porcine. *IEEE Trans. Med. Imaging* 2003;22(11):1358–1368. [PubMed: 14606670]
- Lunn KE, Paulsen KD, Lynch DR, Roberts DW, Kennedy FE, Hartov A. Assimilating intraoperative data with brain shift modeling using the adjoint equations. *Medical Image Analysis* 2005;9:281–293. [PubMed: 15854847]
- Lunn KE, Paulsen KD, Liu F, Kennedy FE, Hartov A, Roberts DW. Data-guided brain deformation modeling: Evaluation of a 3-D adjoint inversion method in porcine studies. *IEEE Transactions on Biomedical Engineering* 2006;53(10):1893–1900. [PubMed: 17019852]
- Lynch, D. *Numerical Partial Differential Equations for Environmental Scientists and Engineers*. Springer: Berlin; 2004.
- Ji, S.; Liu, F.; Hartov, A.; Roberts, DW.; Paulsen, KD. Brain-Skull Boundary Conditions in a Computational Deformation Model; *Medical Imaging 2007: Visualization, Display and Image-Guided Procedures*; Proceedings of SPIE, Vol. 6509, 65092J. San Diego, USA: 2007.

- Ji S, Wu Z, Hartov A, Roberts DW, Paulsen KD. Mutual-Information-Based Image to Patient Registration Using Intraoperative Ultrasound in Image-Guided Neurosurgery. *Medical Physics* 2008a;35(10):4612–4624. [PubMed: 18975707]
- Ji, S.; Hartov, A.; Fontaine, K.; Borsic, A.; Roberts, D.; Paulsen, K. Coregistered Volumetric True 3D Ultrasonography in Image-Guided Neurosurgery; *Medical Imaging 2008: Visualization, Display and Image-Guided Procedures*; Proceedings of SPIE, Vol. 6918, 69180F. San Diego, USA: 2008b.
- Ji S, Roberts DW, Hartov A, Paulsen KD. Intraoperative ultrasound for brain shift compensation in image-guided neurosurgery. 2009a(under review)
- Ji S, Roberts DW, Hartov A, Paulsen KD. Brain-skull contact boundary conditions in an inverse computational deformation model. *Medical Image Analysis*. 2009b(in press)
- Maurer CR Jr, Hill DLG, Maciunas RJ, Barwise JA, Fitzpatrick JM, Wang MY. Measurement of intraoperative brain surface deformation under a craniotomy. *MICCAI 1998*:51–62. LNCS 1496
- Miga, MI. Development and Quantification of a 3D Brain Deformation Model for Model-Updated Image-Guided Stereotactic Neurosurgery. Hanover, NH: Dartmouth College, Thayer School of Engineering; 1998.
- Miga MI, Paulsen KD, Kennedy FE, Hartov A, Roberts DW. Model-updated image-guided neurosurgery using the finite element method: incorporation of the falx cerebri. *MICCAI 1999*:900–909.
- Miga MI, Sinha TK, Cash DM, Galloway RL, Weil RJ. Cortical surface registration for image-guided neurosurgery using laser-range scanning. *IEEE Trans. Med. Imag* 2003;22(8):973–985.
- Nabavi A, Black PM, Gering DT, Westin CF, Mehta V, Pergolizzi RS Jr, Ferrant M, Warfield SK, Hata N, Schwartz RB, Wells WM III, Kikinis R, Jolesz FA. Serial intraoperative magnetic resonance imaging of brain shift. *Neurosurgery* 2001;48(4):787–798. [PubMed: 11322439]
- Nimsky C, Ganslandt O, Cerny S, Hastreiter P, Greiner G, Falbusch R. Quantification of, visualization of and compensation for brain shift using intraoperative magnetic resonance imaging. *Neurosurgery* 2000;47(5):1070–1080. [PubMed: 11063099]
- Nimsky C, Ganslandt O, Hastreiter P, Fahlbush R. Intraoperative compensation for brain shift. *Surgical Neurology* 2001;56:357–365. [PubMed: 11755962]
- Paulsen KD, Miga MI, Kennedy FE, Hoopes PJ, Hartov A, Roberts DW. A computational model for tracking subsurface tissue deformation during stereotactic neurosurgery. *IEEE Trans. Biomed. Eng* 1999;46:213–225. [PubMed: 9932343]
- Platenik LA, Miga MI, Roberts DW, Lunn KE, Kennedy FE, Hartov A, Paulsen KD. In vivo quantification of retraction deformation modeling for updated image-guidance during neurosurgery. *IEEE Transactions on Biomedical Engineering* 2002;49(8):823–835. [PubMed: 12148821]
- Roberts DW, Hartov A, Kennedy FE, Miga MI, Paulsen KD. Intraoperative brain shift and deformation: A quantitative analysis of cortical displacement in 28 cases. *Neurosurgery* 1998;43(4):749–760. [PubMed: 9766300]
- Roberts DW, Miga MI, Kennedy FE, Hartov A, Paulsen KD. Intraoperatively updated neuroimaging using brain modeling and sparse data. *Neurosurgery* 1999;45:1199–1207. [PubMed: 10549938]
- Schenk O, Gartner K. Solving unsymmetric sparse systems of linear equations with PARDISO. *J. Future Gen. Compu. Syst* 2004;20(3):475–487.
- Skrinjar O, Nabavi A, Duncan J. Model-driven brain shift compensation. *Medical Image Analysis* 2002;6:361–373. [PubMed: 12494947]
- Sun H, Roberts DW, Farid H, Wu Z, Hartov A, Paulsen KD. Cortical surface tracking using an operating microscope. *Neurosurgery* 2005a;56(1):86–97. [PubMed: 15799796]
- Sun H, Lunn KE, Faird H, Wu Z, Roberts DW, Hartov A, Paulsen KD. Stereopsis-Guided Brain Shift Compensation. *IEEE Transactions on Medical Imaging* 2005b;24(8):1039–1052. [PubMed: 16092335]
- Sullivan JM Jr, Charron G, Paulsen KD. A three dimensional mesh generator for arbitrary multiple material domains. *Finite Element Anal. and Design* 1997;25(3):219–241.
- Unsgaard G, Selbekk T, Muller TB, Ommedal S, Torp SH, Myhr G, Bang J, Hernes TAN. Ability of navigated 3D ultrasound to delineate gliomas and metastases – comparison of image interpretation with histopathology. *Acta Neurochir (Wien)* 2005;147:1259–1269. [PubMed: 16172831]

- West J, Fitzpatrick JM, Wang MY, Dawant BM, Maurer CR, et al. Comparison and evaluation of retrospective intermodality brain image registration techniques. *J. of Computer Assisted Tomography* 1997;21(4):554–566.
- Wirtz CR, Bonsanto MM, Knauth M, Tronnier VM, Albert FK, Staubert A, Kunze S. Intraoperative magnetic resonance imaging to update interactive navigation in neurosurgery: Method and preliminary experience. *Computer Aided Surgery* 1997;2:172–179. [PubMed: 9377718]
- Wittek A, Kikinis R, Warfield SK, Miller K. Brain shift computation using a fully nonlinear biomechanical model. *Med Image Comput Comput Assist Interv Int Conf Med Image Comput Comput Assist Interv* 2005;8(Pt 2):583–590. [PubMed: 16686007]
- Wolfsberger S, Rössler K, Regatschnig R, Ungersböck K. Anatomical landmarks for image registration in frameless stereotactic neuronavigation. *Neurosurg Rev* 2002;25:68–72. [PubMed: 11954768]
- Wu Z, Paulsen KD, Sullivan JM Jr. Model initialization and deformation for automatic segmentation of T1-weighted brain MRI data. *IEEE Transaction on Biomedical Engr* 2005;52:1128–1131.

Fig. 1. Mean model-data misfit at the sparse data points as a function of the standard deviation in measurement error (σ_ϵ) and correlation length (L). The sparse data used for model guidance were noise-free.

(a)

(b)

Fig. 2.

Using the optimal values of σ and L (0.2 mm and 5 mm, respectively), the model-data misfit as a function of the level of measurement noise (i.e., the standard deviation of the magnitude of displacement error that followed a normal distribution; a). Model-data misfit as a function of the number of iterations with a fixed level of measurement noise of 0.5 mm (b). The initial misfit around the tumor boundary and the minimum number of iterations for SGD to reach a statistically equivalent ϵ_{sprs} relative to REP are also shown. Because of the synthetic noise added to the sparse data (whereas the cross-validation points were noise free), ϵ_{cross} was smaller than ϵ_{sprs} .

(a)

(b)

Fig. 3.

The residual model-data misfit at the sparse data (a; spr_{s}) and cross-validation locations (b; $cross$) normalized with respect to $init$ as a function N_{sprs} for the pooled clinical datasets. Both normalized spr_{s} and $cross$ decreased with the increase in N_{sprs} , regardless of which inverse algorithm was used, once again demonstrating that model performance improves with increased number of data points available for model guidance.

Fig. 4.

With all five clinical datasets pooled, the SGD execution time normalized by that required for REP (solid line; in percentage) is shown as a function of N_{sprs} . Although the actual computational cost increased with the increase in N_{sprs} (Table 6), the normalized run-time of SGD relative to REP decreased when N_{sprs} increased. Also shown are the minimum iterations required for SGD to reach an spr_{rs} statistically equal to or smaller than that produced by REP (dashed line) for each N_{sprs} .

Fig. 5.

A typical iUS image after dural opening transformed into the pMR space (patient 3) using the last-known-correct registration obtained between iUS before dural opening and pMR (a). Tumor displacement was achieved by first rigidly transforming pMR to match tumor boundaries in iUS after dural opening, and then by calculating the spatial change in the pMR tumor boundary nodal positions (b). The whole-brain deformation was subsequently computed using SGD and REP with a subset of the resulting displacement data ($N_{sprs}=200$). The pMR was modified accordingly and overlaid with iUS ((c) and (d), respectively), along with projected brain displacements on the iUS plane. Dashed lines in (c) and (d) are identical to the solid line in (b).

Fig. 6.

Typical iUS (green) overlay with the corresponding oblique pMR image (red) captured immediately after dural opening in four patients. The undeformed tumor (or cyst for patient 2 and 4) surface cross-section (dashed thin lines) intersected by the iUS imaging plane and the corresponding deformed counterparts given by SGD (solid thick lines; at $Iter_{min}$) and REP (dashed thick lines) with $N_{sprs}=200$ are shown. Both algorithms had similar magnitudes of residual $_{cross}$, and the respective deformed MR tumor surfaces were comparable in matching that presented in the iUS images, although it was also evident from visual inspection that SGD outperformed REP in regions indicated by the arrows.

Table 1

Model parameters and their values used in this study.

Parameter	Description	Value
G	Brain shear modulus	2100 Pa
ν	Brain Poisson's ratio	0.46
ρ_t	Brain tissue density	1000 kg/m^3
$\rho_{f,s}$	Surrounding fluid density for saturated nodes	1000 kg/m^3
$\rho_{f,u}$	Surrounding fluid density for unsaturated nodes	1.0 kg/m^3
k	Hydraulic conductivity	$1e^{-7} m^3 s/kg$
α	Ratio of fluid volume extracted to change in solid volume	1
$1/S$	Fluid that can be forced into the tissue under constant volume	0
ψ	Pressure source strength	0 Pa/s

Table 2

Summary of patient population and type, location and size of tumor.

	Age (gender)	Type of tumor	Location of tumor	Tumor volume (cc)
1	61 (F)	meningioma	right parietal	0.62
2	38 (M)	low grade glioma	left temporal	4.7 (25.73) ¹
3	62 (M)	metastasis	right posterior temporal	10.99
4	49 (F)	ganglioglioma	right parietal	0.2 (10.51) ²
5	62 (F)	meningioma	right parietal	13.77

^{1,2} Cyst volume in parenthesis when applicable.

Table 3

$Iter_{min}$ for SGD and the resulting SGD computational cost (in seconds) compared with that of REP with different N_{sprs} when the synthetic data were used for model guidance. Also shown are the SGD3 computational costs normalized by those produced by REP.

N_{sprs}	$Iter_{min}$	t_{SGD} (sec)	t_{REP} (sec)	Normalized t_{SGD}
10	6	43.1	54.2	79.5%
100	26	91.3	235.3	38.8%
200	38	104.1	436.4	23.9%

Table 4

Total number of data points with measured displacement (N_{tot}), initial model-data misfit ($_{init}$) before model correction, $_{sprs}$ and $_{cross}$ with $N_{sprs}=200$, and the percentage of model deformation recovery for each individual patient using either the SGD or REP algorithm.

	Patient 1	Patient 2	Patient 3	Patient 4	Patient 5
N_{tot}	524	1180	3366	4014	1186
$_{init}$ (mm)	3.6±0.56	5.8±0.31	5.9±2.4	5.1±1.3	5.8±0.48
$_{sprs}$ (mm; $N_{sprs}=200$)	REP	0.90±0.47	1.36±0.89	1.75±0.63	1.07±0.58
	SGD	0.87±0.35	1.16±0.36	1.59±0.78	1.05±0.66
$_{cross}$ (mm; $N_{sprs}=200$)	REP	0.97±0.46	1.46±0.91	1.63±0.61	1.06±0.95
	SGD	0.79±0.31	0.78±0.43	1.52±0.81	1.06±0.63
Percentage of recovery	REP	73.1%	74.9%	77.0%	81.7%
	SGD	78.1%	86.6%	74.3%	81.8%

Summary of computational cost for REP (t_{REP}) and SGD (t_{SGD}) in seconds, and $Iter_{min}$ for SGD with different N_{sprs} for each individual patient. Also shown is the number of nodes (in thousands) of the corresponding brain mesh (the DOF is then four times the number of mesh nodes).

	Patient 1	Patient 2	Patient 3	Patient 4	Patient 5
Number of nodes	14.1	15.2	13.2	18.8	14.4
t_{REP}	75.3	99.8	64.6	102.8	64.0
t_{SGD}	73.8	99.0	41.4	119.0	51.6
$Iter_{min}$	9	11	5	24	6
t_{REP}	312.5	397.7	236.2	422.4	266.7
t_{SGD}	125.7	117.0	102.9	180.3	80.4
$Iter_{min}$	25	19	37	35	15
t_{REP}	577.0	729.5	436.4	776.5	492.5
t_{SGD}	131.2	139.1	122.0	205.2	117.4
$Iter_{min}$	37	27	47	47	27

Table 6

Summary of the computational cost for REP (t_{REP}) and SGD (t_{SGD}) in seconds, and $Iter_{min}$ for SGD with different N_{sprs} for the pooled patient cases using clinical data for model guidance. The SGD computational costs are also normalized by those produced by REP for comparison.

N_{sprs}	$Iter_{min}$	t_{SGD} (sec)	t_{REP} (sec)	Normalized t_{SGD}
10	11±7.6	77±32	81±19	91.6%±19.8%
100	26±9.7	121±37	327±81	37.2%±6.9%
200	37±10	143±36	602±147	24.0%±3.4%