Learning Nonrigid Deformations for Constrained Multi-modal Image Registration

John A. Onofrey[1], Lawrence H. Staib[1,2,3], and Xenophon Papademetris[1,3]

[1] Departments of Biomedical Engineering
[2] Electrical Engineering
[3] Diagnostic Radiology
Yale University, New Haven CT 06520, USA
{john.onofrey,lawrence.staib,xenophon.papademetris}@yale.edu

Abstract. We present a new strategy to constrain nonrigid registrations of multi-modal images using a low-dimensional statistical deformation model and test this in registering pre-operative and post-operative images from epilepsy patients. For those patients who may undergo surgical resection for treatment, the current gold-standard to identify regions of seizure involves craniotomy and implantation of intracranial electrodes. To guide surgical resection, surgeons utilize pre-op anatomical and functional MR images in conjunction with post-electrode implantation MR and CT images. The electrode positions from the CT image need to be registered to pre-op functional and structural MR images. The post-op MRI serves as an intermediate registration step between the pre-op MR and CT images. In this work, we propose to bypass the post-op MR image registration step and directly register the pre-op MR and post-op CT images using a low-dimensional nonrigid registration that captures the gross deformation after electrode implantation. We learn the nonrigid deformation characteristics from a principal component analysis of a set of training deformations and demonstrate results using clinical data. We show that our technique significantly outperforms both standard rigid and nonrigid intensity-based registration methods in terms of mean and maximum registration error.

Keywords: nonrigid registration, multi-modal, statistical deformation model, principal component analysis, image-guided surgery.

1 Introduction

Nonrigid, multi-modal image registration is a challenging task. The ability of nonrigid registration algorithms to successfully find a globally optimal deformation is made difficult by the high dimensionality of the deformations being modeled. In multi-modal registration tasks, nonlinear intensity relationships exacerbate the problem by causing similarity metrics to have many local minima. With an unconstrained search space, algorithms fail to escape the local minima and inaccurately register the images. In this paper, we present a solution to this problem that uses a training set to learn a low-dimensionality parameterization

K. Mori et al. (Eds.): MICCAI 2013, Part III, LNCS 8151, pp. 171–178, 2013.

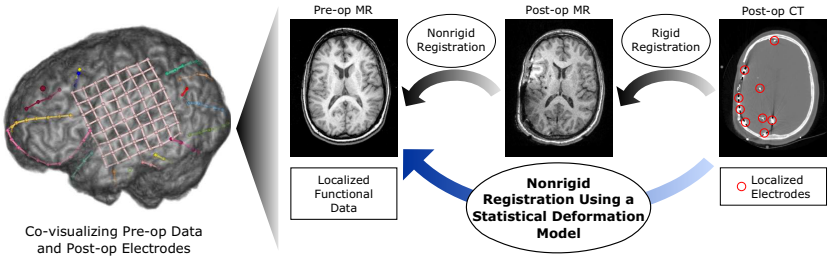

Fig. 1. To co-visualize icEEG electrodes with pre-op imaging data, the currently practiced registration framework rigidly registers the post-op CT image and the localized electrodes to a post-op MRI, and then nonrigidly registers the post-op MRI to the pre-op MRI. Our proposed method directly registers the post-op CT image to pre-op MRI using a learned statistical deformation model. Without having a constrained model of deformation, missing anatomical correspondences, *e.g.* removal of the skull during surgery, as well as imaging artifacts caused by the presence of the electrodes make direct nonrigid registration of the CT to the pre-op MRI inaccurate, even less accurate than rigid registration (see Figure 2).

of the deformation space for a specific task. We present results showing our initial application to nonrigidly map post-surgical electrode implantation CT images to pre-operative MR images acquired as part of epilepsy treatment.

For epilepsy patients whose seizures do not adequately respond to medication, surgical treatment is often an effective method to reduce or to eliminate seizure activity. Intracranial electroencephalography (icEEG), in which surgeons perform a craniotomy and implant (typically over 200) electrodes in suspected regions the brain, is the current "gold-standard" for localizing the focus of seizure activity [8]. Following electrode implantation, clinicians constantly monitor the icEEG for several days to identify, if possible, the electrodes nearest to the seizure focus. The surgeons localize the physical locations of these electrodes using post-op CT imaging. Successful electrode localization in combination with pre-operatively acquired functional brain images can be used to determine the feasibility of surgical tissue resection. Therefore, accurate spatial registration of electrodes with respect to functionally eloquent areas of the brain is critical.

The best, current approach to co-register the icEEG electrodes with the pre-op MR data involves the acquisition of a post-op MR image. The post-op MRI, in this case, serves as an intermediate registration step between the post-op CT and pre-op MR images. This registration framework projects the electrodes to the pre-op MRI space by first rigidly registering the CT image to the post-op MRI, and then nonrigidly registering the post-op MRI to the pre-op MRI to account for post-surgical deformations. Once the electrodes have been transformed to pre-op imaging space, they may be co-visualized with any other functional imaging studies. However, not all institutions are capable of acquiring the post-op MR image, and instead rely upon rigid registration of the post-op CT and pre-op MR images, which leads to inaccurate electrode localization. Given that post-op

MRIs are available at our institution, we propose to use this information to learn the nonrigid deformations from the post-op CT space to pre-op MR space, as Figure 1 illustrates.

Our proposed method leverages the currently used mono-modal post-op MRI to pre-op MRI nonrigid registrations to learn a statistical deformation model (SDM). The SDM is surgical-site dependent because larger brain deformations generally occur ipsilateral to the craniotomy site [1], thus we assume that patients with craniotomies in similar locations will experience similar deformation characteristics. We perform a principal component analysis (PCA) of the nonrigid deformations to construct our SDM [6]. While PCA has difficulty with high-dimensional data, we hypothesize that the post-surgical brain deformation is of low enough dimension that our SDM can capture the gross, intra-subject deformations observed after surgery. Other authors have made use of PCA SDMs to register mono-modal medical images for inter-subject registration [10,4] and for intra-subject motion compensation [2]. In contrast, we use our SDM trained on MR images to directly register each patient's post-op CT image to their pre-op MRI. Our SDM models intra-subject deformations that result from surgical intervention, and *not* inter-subject anatomical variability. To the best of our knowledge, training a SDM on a mono-modality registration task and using that SDM to perform a multi-modality registration is a novel application. By doing so, the SDM can model subcortical deformations in the CT image that would otherwise not be possible with standard, intensity-only registration to the pre-op MR. Our results show our approach significantly reduces both mean and maximum registration error compared to standard rigid and nonrigid intensity-only MR-CT registration methods. We emphasize that standard intensity-based nonrigid methods perform worse than rigid ones, and our approach is the first, to our knowledge, to successfully directly register pre-op MRIs and post-electrode implantation CT images.

2 Methods

2.1 Training the Deformation Model

Given a database of surgical epilepsy patients, we select N patients with craniotomies at similar locations to train our SDM. Each patient's dataset consists of a pre-op MR image $I^{(1)}_{\mathrm{MR}}$, a post-op MR image $I^{(2)}_{\mathrm{MR}}$, and a post-op CT image $I^{(2)}_{\mathrm{CT}}$, where $I^{(t)}$ denotes pre-op images acquired at time $t = 1$ and post-op images at time $t = 2$. To create our SDM, we need to transform all N images into a common reference space. As per current practice, we first rigidly register the post-op MR and CT images by maximizing their normalized mutual information (NMI) [9], and obtain the transformation $T_{\mathrm{CT}\rightarrow\mathrm{MR}}$ (here, $i \rightarrow j$ denotes rigid transformation from space i to j). Next, we rigidly register the pre- and post-op MR images using NMI to produce a transformation $T_{\mathrm{MR}\rightarrow\mathrm{MR}}$. Finally, we nonrigidly register all pre-op MR images to the MNI Colin 27 brain, I_{MNI}, using a free-form deformation (FFD) [7] with NMI and write this transformation

$T_{\mathrm{MR}\rightsquigarrow\mathrm{MNI}}$ (we denote nonrigid transformation from space i to j as $i \rightsquigarrow j$). Once we have computed all transformations, we reslice all images from each of the N patients into MNI space ($181 \times 217 \times 181$ volume with $1\mathrm{mm}^3$ resolution) by concatenating the transformations. Spatial normalization to MNI space is necessary for defining a common reference space for intra-subject nonrigid deformations.

In MNI space, the post-op MR has yet to be nonrigidly registered to the pre-op MR. In current practice, this registration task uses skull-stripped brains to mitigate the effects of missing anatomical correspondences after surgery. However, since we are planning to register the CT images directly to the pre-op MR, the skull is actually one of the most informative structures to register, even if it is lacking some correspondence. Therefore, to accurately register the two non-skull stripped MRIs, we first create brain surface masks with some manual refinement, and then utilize an integrated intensity and point-feature registration algorithm [5]. This algorithm uses a FFD transformation model with 15mm control point spacing and minimizes the NMI similarity metric. With points weighting parameter set to 0.1, the brain surface points constrain the algorithm to align the cortical surface that otherwise has difficulty being accurately registered using intensity registration by itself. We denote the resulting transformations $T_{\mathrm{MR}\rightsquigarrow\mathrm{MR}}$. It is these transformations that we use to train our SDM.

For each patient $i = 1, \ldots, N$, we rewrite the transformation $T_{\mathrm{MR}\rightsquigarrow\mathrm{MR}}$ as a column vector of P concatenated FFD control point displacements in 3D, $\mathbf{d}_i \in \mathbb{R}^{3P}$. Using PCA, we linearly approximate the deformation distribution [6]

$$\mathbf{d} = \bar{\mathbf{d}} + \mathbf{\Phi}\mathbf{w} \tag{1}$$

where $\bar{\mathbf{d}} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{d}_i$ is the mean deformation of the N training registrations, $\mathbf{\Phi}$ is the matrix of orthogonal principal components, and $\mathbf{w}$ is the vector of model variation coefficients. We compute the principal components from the eigensystem decomposition of the covariance matrix $\mathbf{C} = \frac{1}{N-1}\sum_{i=1}^{N}(\mathbf{d}_i - \bar{\mathbf{d}})(\mathbf{d}_i - \bar{\mathbf{d}})^T$. Using this formulation, $\mathbf{\Phi} = (\phi_1|\phi_2|\ldots|\phi_K) \in \mathbb{R}^{3P\times K}$ is a matrix of $K \leq \min\{N, 3P\}$ eigenvectors $\phi_k \in \mathbb{R}^{3P}$ with corresponding eigenvalues in decreasing order $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_K$, and $\mathbf{w} \in \mathbb{R}^K$. The eigenvectors with the k largest eigenvalues define a SDM using k principal modes of variation, with $1 \leq k \leq K$.

2.2 Nonrigid SDM MR-CT Registration

Given a previously unseen pre-op MR and post-op CT image pair for a new patient, we use the learned SDM in Equation 1 to drive the nonrigid registration of post-op CT images to pre-op MR without the use of a post-op MR image. Before we nonrigidly register the two images, we must first transform the images into our model reference space. We do so by rigidly registering $I^{(2)}_{\mathrm{CT}}$ to $I^{(1)}_{\mathrm{MR}}$ and then nonrigidly registering $I^{(1)}_{\mathrm{MR}}$ to I_{MNI}. In both cases we maximize NMI, and use a FFD transformation for nonrigid registration. We use the resulting transformations to reslice both images into MNI space such that $I^{(1)'}_{\mathrm{MR}} = T_{\mathrm{MR}\rightarrow\mathrm{MNI}} \circ I^{(1)}_{\mathrm{MR}}$

and $I_{\mathrm{CT}}^{(2)'} = T_{\mathrm{CT}\to\mathrm{MR}} \circ T_{\mathrm{MR}\to\mathrm{MNI}} \circ I_{\mathrm{CT}}^{(2)}$, where $\circ$ is the transformation operator and I' indicates a resliced image.

With the images now transformed to our SDM reference space, we nonrigidly register $I_{\mathrm{CT}}^{(2)'}$ to $I_{\mathrm{MR}}^{(1)'}$. The SDM model coefficients $\mathbf{w}$ in Equation 1 are a low-dimensional parameterization of a high-dimensional FFD $\mathbf{d}$, and we denote this transformation $T_{\mathrm{SDM}}(\mathbf{x};\mathbf{d})$ for points $\mathbf{x}$ in the reference image domain $\Omega_{\mathrm{MNI}} \subset \mathbb{R}^3$. We register $I_{\mathrm{CT}}^{(2)'}$ and $I_{\mathrm{MR}}^{(1)'}$ by optimizing the cost function

$$\hat{T}_{\mathrm{SDM}} = \arg\max_{\mathbf{w}} J(I_{\mathrm{MR}}^{(1)'}(\mathbf{x}), T_{\mathrm{SDM}}(\mathbf{x};\bar{\mathbf{d}} + \mathbf{\Phi}\mathbf{w}) \circ I_{\mathrm{CT}}^{(2)'}(\mathbf{x})), \quad \forall \mathbf{x} \in \Omega_{\mathrm{MNI}} \qquad (2)$$

where J is the NMI similarity metric. We use a multi-resolution image pyramid and conjugate gradient optimization to solve Equation 2.

3 Results and Discussion

From the database of surgical epilepsy patients available at our institution, we manually identified 18 patients with lateral craniotomies (10 on the ride side, 8 on the left). In order to increase our dataset size, we flipped the left-side craniotomy images to be right-side craniotomies under the assumption that the direction of gross brain deformation correlates to craniotomy location. For each patient in the database, we have images $I_{\mathrm{MR}}^{(1)}$ ($256 \times 256 \times 106$ at $0.977 \times 0.977 \times 1.5$mm resolution), $I_{\mathrm{MR}}^{(2)}$ ($256 \times 256 \times 110$ at $0.977 \times 0.977 \times 1.5$mm resolution), and $I_{\mathrm{CT}}^{(2)}$ ($512 \times 512 \times 137$ at $0.488 \times 0.488 \times 1.25$mm resolution).

We performed *leave-one-out* testing to demonstrate our approach. For each patient $i = 1, \ldots, 18$, we trained the SDM as described in Section 2.1 by omitting the i-th patient from the training set, which consists of the $N = 17$ remaining samples. We then register the i-th patient's post-op CT to their pre-op MR using the SDM in MNI space as described in Section 2.2. We repeated our registration method using different numbers of modes of variation in $\mathbf{w}$, *i.e.* different values for k. We compared our method to rigid registration, which in our case $\hat{T}_{\mathrm{R}} = I$, and intensity-only FFD $\hat{T}_{\mathrm{FFD}}$ with 15mm control point spacing. We also compared to FFD using a lower-dimensional 30mm control point spacing, but the results were similar to 15mm. Thus, for brevity, we reported only the results using 15mm FFD. We implemented and ran our algorithm as part of BioImage Suite [3].

To evaluate registration performance, we treated the nonrigid transformations $T_{\mathrm{MR}\rightsquigarrow\mathrm{MR}}$ found during training with the post-op MR image, and as used in current practice, as a ground-truth. For an estimated transformation $\hat{T}$, we calculated the magnitude of transformation error

$$\varepsilon(\mathbf{x}) = \| \hat{T}(\mathbf{x}) - T_{\mathrm{MR}\rightsquigarrow\mathrm{MR}}(\mathbf{x}) \|, \quad \mathbf{x} \in \Omega_{\mathrm{MNI}}.$$

While this method evaluates performance against the current standard, we assessed our MR-CT registration results with the rigid registration $\hat{T}_{\mathrm{R}}$ in mind.

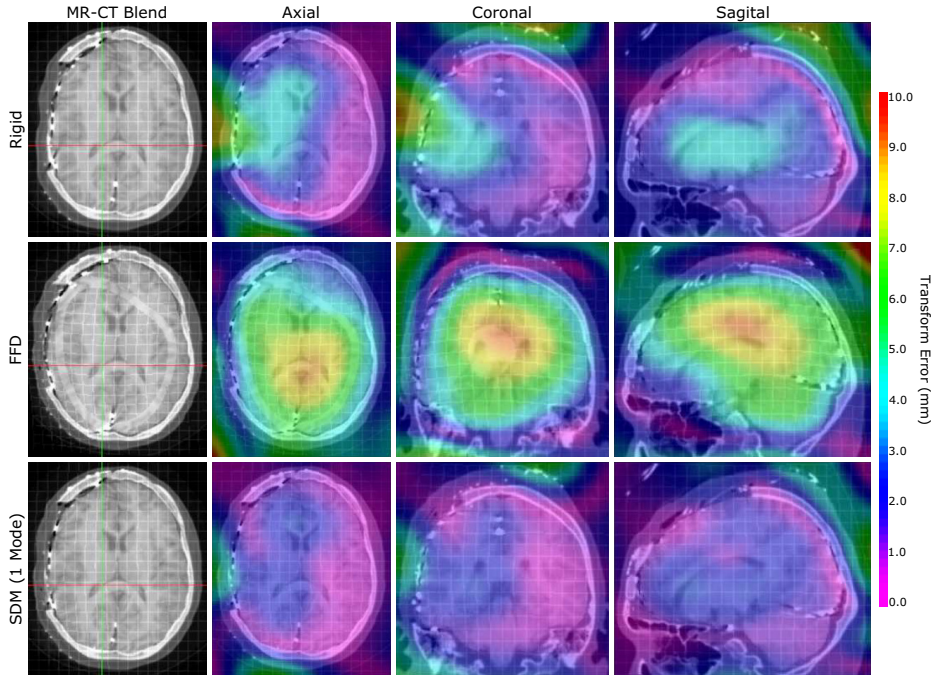

Fig. 2. Our nonrigid SDM registration method using a single mode of variation has lower transformation error throughout much of the brain, particularly in the areas around the craniotomy and around the ventricles, than both standard rigid and FFD intensity-only registration methods. We spatially visualize voxel-wise transformation error using a colormap overlay for an exemplar patient with a craniotomy on right side of the skull (left side of the axial images in radiological convention). The left column shows intensity-blended MR-CT images that highlight electrodes and skull location with respect to the skull-stripped MRI brain, along with a deformation grid.

It is our experience that surgeons generally only trust rigid registration for volumetric MR-CT registration. Furthermore, we evaluated maximum error because surgeons are most interested in quantifying worst-case performance.

We spatially visualized $\varepsilon(\mathbf{x}), \forall \mathbf{x} \in \Omega_{\text{MNI}}$ using a colormap overlay, as shown in Figure 2, to compare our estimates for $\hat{T}_{\text{R}}$, $\hat{T}_{\text{FFD}}$, and $\hat{T}_{\text{SDM}}$. In comparison to rigid registration, our method reduced error throughout the brain, and particularly so around the craniotomy and ventricles. We highlight the poor performance of the intensity-only FFD MR-CT registration. Due to the poor soft-tissue contrast in the CT, the FFD failed to accurately register the interior of the brain. Even though both FFD and SDM used the same NMI similarity metric, the SDM constrained the transformation to accurately mimic the interior deformations.

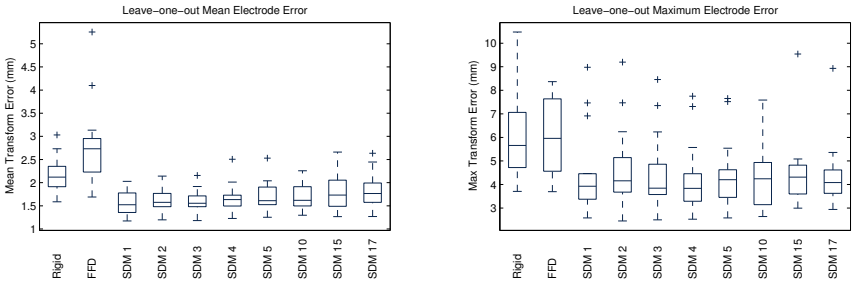

Fig. 3. Our proposed nonrigid SDM registration method significantly reduced transformation error compared to standard rigid and FFD intensity registration methods. SDM k denotes registration using the first $k = 1, 2, 3, \ldots$ principal modes of variation. We plot the distributions of both mean (left plot) and maximum (right plot) transformation error at electrode locations for the 18 leave-one-out MR-CT registrations. The boxplots show median, inner quartile, extremes, and outlier values.

We next quantified $\varepsilon(\mathbf{x}_e)$ at each patient's electrode locations $\mathbf{x}_e \in \Omega_{\mathrm{E}} \subset \Omega_{\mathrm{MNI}}$, where the mean number of electrodes for a patient was 197. Figure 3 summarizes patient mean error $\bar{\varepsilon}$ and maximum error $\varepsilon_{\max}$ over all $\mathbf{x}_e \in \Omega_{\mathrm{E}}$. Our proposed SDM registration significantly ($p \leq 0.05$, paired t-test) reduced both $\bar{\varepsilon}$ and $\varepsilon_{\max}$ for all modes of variation with respect to rigid registration. The mean and maximum ε across all 18 patients using SDM with 1 mode was 1.58 ± 0.24mm and 4.39 ± 1.70mm, respectively, which compared to 2.12 ± 0.37mm and 6.08 ± 1.84mm for rigid registration and 2.75 ± 0.85mm and 5.95 ± 1.64mm for FFD (all reported values are mean $\pm$ std). FFD significantly increased $\bar{\varepsilon}$ with respect to rigid registration.

Additionally, we computed mean error throughout 4 different volumes of interest (VOIs): the right brain hemisphere Ω_{RB}, the right skull Ω_{RS}, the left brain hemisphere Ω_{LB}, and left skull Ω_{LS}, such that $\Omega_i \subset \Omega_{\mathrm{MNI}}$ and $\Omega_i \bigcap \Omega_j = \emptyset, i \neq j$. The Ω_{RB} and Ω_{RS} VOIs were of particular interest since they were ipsilateral to the craniotomy. Compared to rigid registration, our method significantly ($p \leq 0.05$) reduced $\bar{\varepsilon}$ in Ω_{RB} using 1-5 modes of variation and in Ω_{RS} using 1 and 3 modes. SDM registration significantly increased $\bar{\varepsilon}$ in Ω_{LB} when using 15 and 17 modes, but otherwise there were no significant increases in $\bar{\varepsilon}$ in Ω_{LB} or Ω_{LS} compared to rigid registration. We noted that $\bar{\varepsilon}$ generally increased as we used more modes of variation in our SDM, which is most likely explained by the PCA over-fitting to the training set.

4 Conclusion

Our proposed method models post-surgical nonrigid deformations to significantly reduce both mean and maximum transformation errors in multi-modality nonrigid MR-CT registration compared to standard rigid and unconstrained

nonrigid registrations. As shown in Figures 2 and 3, standard nonrigid registration methods perform worse than rigid methods, at least in this dataset where the CT images had both deformation and significant artifacts. Although we present results for training a SDM at only a single craniotomy location, in the future, we aim to create craniotomy site-specific SDMs to model the corresponding deformations at different locations. Furthermore, we could improve our training set registrations by including additional labeled anatomical structures, *e.g.* the ventricles, to improve PCA model construction. The ideas presented in this paper constitute a general framework to effectively register multi-modal volumetric images nonrigidly. In particular we demonstrate how we can use a small subset of high quality training data (in this case the rare availability of a post-electrode implantation MRI) to learn about the properties of the deformation model in a given case of nonrigid deformation, and to subsequently use this knowledge to solve the nonrigid registration problem in the more general case with lesser quality data (*i.e.* the direct nonrigid multimodal registration of pre-op MRI to post-implantation CT). Similar principles could be applied to nonrigidly register, for example, interventional ultrasound images to pre-operative MRI.

References

1. Hartkens, T., Hill, D., Castellano-Smith, A., Hawkes, D., Maurer Jr., C.R., Martin, A., Hall, W., Liu, H., Truwit, C.: Measurement and analysis of brain deformation during neurosurgery. IEEE Trans. on Medical Imaging 22(1), 82–92 (2003)
2. He, T., Xue, Z., Xie, W., Wong, S.T.C.: Online 4-D CT estimation for patient-specific respiratory motion based on real-time breathing signals. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010, Part III. LNCS, vol. 6363, pp. 392–399. Springer, Heidelberg (2010)
3. Joshi, A., Scheinost, D., Okuda, H., Belhachemi, D., Murphy, I., Staib, L., Papademetris, X.: Unified framework for development, deployment and robust testing of neuroimaging algorithms. Neuroinformatics 9, 69–84 (2011)
4. Kim, M.J., Kim, M.H., Shen, D.: Learning-based deformation estimation for fast non-rigid registration. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPRW 2008, pp. 1–6 (June 2008)
5. Papademetris, X., Jackowski, A.P., Schultz, R.T., Staib, L.H., Duncan, J.S.: Integrated intensity and point-feature nonrigid registration. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) MICCAI 2004. LNCS, vol. 3216, pp. 763–770. Springer, Heidelberg (2004)
6. Rueckert, D., Frangi, A., Schnabel, J.: Automatic construction of 3-d statistical deformation models of the brain using nonrigid registration. IEEE Trans. on Medical Imaging 22(8), 1014–1025 (2003)
7. Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D.: Nonrigid registration using free-form deformations: application to breast mr images. IEEE Trans. on Medical Imaging 18(8), 712–721 (1999)
8. Spencer, S.S., Sperling, M., Shewmon, A.: Intracranial electrodes. In: Epilepsy, A Comprehensive Textbook, pp. 1719–1748 (1998)
9. Studholme, C., Hill, D., Hawkes, D.: An overlap invariant entropy measure of 3d medical image alignment. Pattern Recognition 32(1), 71–86 (1999)
10. Xue, Z., Shen, D., Davatzikos, C.: Statistical representation of high-dimensional deformation fields with application to statistically constrained 3d warping. Medical Image Analysis 10(5), 740–751 (2006)