

Simulating deformations of MR brain images for validation of atlas-based segmentation and registration algorithms

Zhong Xue,^{a,*} Dinggang Shen,^a Bilge Karacali,^a Joshua Stern,^b
David Rottenberg,^b and Christos Davatzikos^a

^aSection of Biomedical Image Analysis (SBIA), Department of Radiology, University of Pennsylvania, Philadelphia, Pennsylvania, USA

^bUniversity of Minnesota, Veterans Affairs Medical Center, Minneapolis, Minnesota, USA

Received 26 September 2005; revised 17 March 2006; accepted 7 August 2006
Available online 25 September 2006

Simulated deformations and images can act as the gold standard for evaluating various template-based image segmentation and registration algorithms. Traditional deformable simulation methods, such as the use of analytic deformation fields or the displacement of landmarks followed by some form of interpolation, are often unable to construct rich (complex) and/or realistic deformations of anatomical organs. This paper presents new methods aiming to automatically simulate realistic inter- and intra-individual deformations. The paper first describes a statistical approach to capturing inter-individual variability of high-deformation fields from a number of examples (training samples). In this approach, Wavelet–Packet Transform (WPT) of the training deformations and their Jacobians, in conjunction with a Markov random field (MRF) spatial regularization, are used to capture both coarse and fine characteristics of the training deformations in a statistical fashion. Simulated deformations can then be constructed by randomly sampling the resultant statistical distribution in an unconstrained or a landmark-constrained fashion. The paper also describes a model for generating tissue atrophy or growth in order to simulate intra-individual brain deformations. Several sets of simulated deformation fields and respective images are generated, which can be used in the future for systematic and extensive validation studies of automated atlas-based segmentation and deformable registration methods. The code and simulated data are available through our Web site.

© 2006 Elsevier Inc. All rights reserved.

Introduction

The plethora of automated methods for deformable registration of brain images, and their widespread use for template-based segmentation and labeling, spatial normalization, shape analysis and a variety of other tasks, has necessitated the construction of gold or “silver” standards for validation of different methods and evaluation of their relative merits. Although various methods for

synthesizing anatomical deformations have occasionally been adopted in the literature, they typically rely on analytic forms of deformations, such as linear, polynomial and spline models, or on the displacement of a number of landmarks followed by the thin plate spline (TPS) (Bookstein, 1989) or other interpolations. Such approaches generate relatively simple and/or unrealistic deformations, and they are unable to capture/simulate the true and full range of anatomical variability. This paper aims to develop approaches that simulate anatomically valid deformations, using brain MRI as a focus, and to generate realistic data sets that could be used for validation of atlas-based segmentation and registration algorithms.

Two classes of deformations and images are simulated. The first class includes deformations between different individuals (inter-individual deformations) and intends to capture inter-individual morphological variability; it therefore aims to develop a foundation for generating gold standard for evaluating atlas registration methods. The second class simulates morphological changes that can be encountered within the same individual due to growth or atrophy of brain tissue; it therefore aims to develop a foundation for generating gold standard for brain development as well as brain pathology, as these are manifested via change of tissue volume.

With respect to the first class, the basic premise is that a number of deformation fields are available to be used for estimating the probability density function (pdf) of high-dimensional deformation fields, which can then be randomly sampled to generate any number of synthesized deformations. For the purposes of demonstrating and testing our methodology, we used a number of deformations that were obtained using the high-dimensional deformable registration methods described in (Shen and Davatzikos, 2002; Davatzikos et al., 2003a). This inevitably biases the generated deformations toward the family of deformations that can be generated by this particular warping algorithm, and it is somewhat inevitable because there is no truly unbiased way to define deformation fields that are realistic and representative of true anatomical variability. Our rationale was that, in order to eventually evaluate a number of warping methods, each of them could be used for training our simulator, and then the rest can be

* Corresponding author.

E-mail address: zhong.xue@uphs.upenn.edu (Z. Xue).

Available online on ScienceDirect (www.sciencedirect.com).

used for testing via comparison against the simulated ground truth. A warping method that consistently performs better in capturing the synthesized (known) deformations is likely to be more accurate. More accurate training sample deformations could also be generated by first extensively labeling and landmarking a number of images (Kristi Boesen et al., 2005), and then applying a high-dimensional warping algorithm constrained by these manual labels and landmarks. Such adequately constrained warping algorithms are likely to generate deformations that are close to a gold standard, and therefore appropriate for training.

As described above, central to the simulator of inter-individual morphological variability is the estimation of the pdf of deformation fields from a limited number of training samples. This is a very challenging task and is the main focus of this paper. One of the most popular methods has been the application of the principal component analysis (PCA) in order to estimate a number of principal components that are frequently called principal modes of variation. However, this approach yields poor results when applied to 3D dense deformation fields, due to under-training in practical settings. For example, accurately estimating a dense 3D deformation field of the entire brain could require tens of thousands of training deformations, if not more. The dramatic failure of standard methods for estimating covariance matrices and the associated pdfs has been well known in the signal estimation literature (Mallat and Wavelet, 1998). Accordingly, in the 1990s, methods based on scale-space decompositions were investigated. Our approach, referred to as statistical simulation of deformations (SSD), builds upon the methods described in Davatzikos et al. (2003a,b) and Xue et al. (2005), which use wavelet-based decompositions in order to more accurately estimate and sample pdfs of high-dimensional deformation fields, when only a relatively small number of training samples are available (e.g., tens, or in the order of 100). Random sampling of the pdfs estimated by SSD generates synthesized deformations. Moreover, a method for constrained sampling is also presented, and is applicable when anatomical landmarks from an individual are also available, thereby generating deformations that conform to both group-based statistics of morphological variation and (limited) anatomical information pertaining to the individual. Compared to the traditional landmark-based interpolation algorithms, the landmark-constrained SSD can generate both richer and anatomically valid deformations.

The second class of deformations simulates intra-individual deformations. Unlike inter-individual deformations, intra-individual brain deformations reflect the structural changes of an individual brain at different time points, e.g., tissue atrophy/growth of a selected structure or within a selected region of a brain. In this paper, we use the method proposed in Karacali and Davatzikos (2006) to simulate the atrophy and growth. That algorithm generates a deformation field by minimizing the difference between its Jacobian determinants and the desired ones, subject to some smoothness constraints on the deformation field. The desired Jacobian determinants describe the volumetric changes of different tissues or different regions.

As one of the primary goals of this work is to develop a testbed for validation of deformable registration algorithms, various inter-individual and intra-individual brain deformations are simulated using the above approaches and a variety of data sets are generated for future use as ground truth in testing template-based segmentation and warping algorithms, and to further perform parameter optimization of these algorithms.

The remainder of the paper is organized as follows. The Methods section describes the detailed methods used to simulated deformations and the corresponding images, and the Stimulation results section introduces the simulated data sets and performs evaluations on the simulated data. The Summary and conclusion section summarizes and concludes with additional discussions and future directions.

Methods

In this section, we first describe our method for estimating the multivariate statistics of high-dimensional deformation fields that reflect inter-individual variability of brain structures, and for generating new deformations by sampling the statistical models. Then, we describe how additional anatomical landmarks can be used to further constrain this sampling process. Finally, we describe a different type of anatomical deformation that represents intra-individual anatomical changes, which might occur with normal development and aging, or with development of disease, i.e., simulation of tissue growth and tissue atrophy.

Statistical simulation of deformations (SSD)

Description of SSD

Let $\mathbf{f}(x)$ be a scalar or vector field defined over the template image domain Ω , $x \in \Omega$. Estimating the pdf of \mathbf{f} from a relatively small number of training samples is necessary in order to randomly simulate new deformations via random sampling, or in order to evaluate the likelihood of a given deformation field in the context of Bayesian reasoning and/or the context of statistically constrained deformable registration. However, the commonly used PCA method (e.g., Cootes et al., 1994; Miller et al., 1997) performs poorly when \mathbf{f} is of very high dimensionality and only a limited training set is available to estimate the PCA model. This is because a global PCA model is able to capture mainly global size and shape characteristics that are of limited interest and value, especially for the purposes of simulating complex deformations to be used for validation purposes. In order to capture finer and more localized variations of \mathbf{f} we follow and extend the framework proposed in Davatzikos et al. (2003a), which is referred to as the wavelet-based PCA (W-PCA) model. The W-PCA model decomposes \mathbf{f} using the Wavelet-Packet Transform (WPT) and subsequently captures within-scale statistics via hierarchically organized PCA models. These PCA models are estimated from statistical distributions that are both of lower dimensionality, and more compact due to correlations among variables. For example, the PCA model derived from a coarse-scale representation of \mathbf{f} represents a very compact distribution due to the smoothing and down-sampling applied at each level of the wavelet-packet decomposition; the distribution of high frequency detail within a local window is also easier to estimate due to its low dimensionality owing to the small window size. Performing PCA within each band at a given scale is important due to correlations among wavelet coefficients corresponding to adjacent locations, something which is particularly prominent in smooth elastic-type of deformations, in contrast to, for example, acoustic signals in which wavelet coefficients are typically assumed to be statistically independent. The fundamental assumption in W-PCA is that the wavelet-based rotation renders the covariance matrix of \mathbf{f} close to block-diagonal, thereby enabling a more accurate estimation from a limited set of examples, compared to the usual sample covariance

estimation. Generating a random sample from the pdf of \mathbf{f} is then achieved by randomly sampling each PCA model in this hierarchy, and then using the inverse WPT.

In theory, if the W-PCA model described above captures the statistics of deformation \mathbf{f} accurately, we can just generate sample deformations as described above. In practice, however, the assumption that the covariance matrix of \mathbf{f} is block-diagonal in the wavelet-packet basis does not hold exactly. Although it is well known that for broad classes of signals, correlations across scales diminish rapidly, they are nonetheless non-negligible for adjacent scales. Therefore, the resulting deformation fields might have unrealistic discontinuities. In order to alleviate this problem, we observe that additional constraints imposed on the estimated deformation fields can be used to define subspaces in which the deformation must belong to. Therefore, we require that a valid deformation field simultaneously satisfies several constraints, i.e., it belongs to the intersection of a number of subspaces, each of which satisfies some constraints on the deformation. Fig. 1 shows the three subspaces used on our approach. The W-PCA model applied to the deformation field specifies one such subspace, the W-PCA model of the Jacobian determinants of the deformation fields specifies the second subspace, and the Markov random field (MRF) regularization represents the third subspace.

The second subspace is represented by the W-PCA model of the Jacobian determinants of the deformation fields. The Jacobian determinants are included separately from the deformations themselves because they reflect local volumes of anatomical structures, which are important from the perspective of spatial distribution of the amount of brain tissue. Fig. 2 shows two precentral gyri from different subjects, which differ quite dramatically in shape, but not in volume. More generally, it would be reasonable to assume that, although the cortical folding patterns can vary significantly across individuals, the need of different cortical structures to occupy certain tissue volume renders the Jacobian determinant, which is directly related to tissue volume, relatively less variable across individuals and therefore easier to estimate from a limited number of samples. Hence, the Jacobian is considered here as an additional quantity to be constrained in Fig. 1. It is worth noting that the complementary property of the statistical models of deformation fields and those of the Jacobian determinants allows us to combine them together by requiring that a valid deformation field be within the intersection of the subspaces defined by them. Therefore, a valid deformation field has to be first sampled randomly from the statistical model of deformations, and then iteratively constrained by both of the statistical models. We note, here, that converting a constraint on the




Fig. 1. The space of valid deformations is represented as the intersection of different subspaces reflecting different aspects of a warping field.




Fig. 2. Jacobian determinants, reflecting volumetric changes, could be expected to be less variable than the displacement fields. For example, the displacement fields of the precentral gyri are very different for these two brains, whereas the volumes of these gyri are very similar, leading to much less variable Jacobian determinants.

Jacobian determinant to a constraint on the deformation field, i.e., finding the displacement field that satisfies certain conditions on the Jacobian determinant, is a challenging task with no unique solution. Herein we use (Karacali and Davatzikos, 2004), which utilizes an iterative projection scheme that minimally, according to some distance criteria, modifies a given displacement field so that it satisfies certain conditions on the Jacobian determinant. This algorithm is used to realize the projection of a given displacement field to the subspace of “valid Jacobians”.

The third subspace is represented by a nested Markov random field (MRF) regularization, which imposes spatial smoothness at different scales in conjunction with the inverse WPT. The purpose of MRF regularization is to eliminate potential discontinuities emanating from the assumption of independence across wavelet bands.

In the following subsections, we first describe the W-PCA model, which can be used to effectively capture the statistics of both deformation fields and their Jacobian determinants. Then, we introduce the MRF regularization in detail. Finally, we summarize the SSD algorithm for deformable simulation.

The wavelet-PCA (W-PCA) models for estimating the pdfs of deformation fields and their Jacobian determinants

The W-PCA model is used to estimate the pdf of \mathbf{f} , which can be a deformation field or a Jacobian determinant field, using N samples. It first applies an L -level WPT to \mathbf{f} , and then constructs a PCA model of the wavelet coefficients of each wavelet band at level L . Finally, it combines these pdfs of different wavelet bands together. Fig. 3 illustrates the structure of 1D WPT. For 3D WPT, the wavelet coefficients at level l are represented by $\mathbf{w}^{(l,b)}$, $b=0, 1, \dots, B_l-1$, where $B_l=8^l$ and $l=1, 2, \dots, L$. At each level, $\mathbf{w}^{(l,0)}$ always




Fig. 3. Illustration of Wavelet-Packet Transform (WPT).

represents the low-pass wavelet coefficients. For simplicity, \mathbf{f} is also referred to as $\mathbf{w}^{(0,0)}$.

After L -level WPT, \mathbf{f} can be represented by all the wavelet coefficients at level L , i.e., $\mathbf{w}^{(L,b)}$. Assuming that different bands in the wavelet subspaces are independent, then the pdf of deformation \mathbf{f} is,

$$p(\mathbf{f}) = \prod_{b=0}^{B_L-1} p(\mathbf{w}^{(L,b)}). \quad (1)$$

The pdf of each band (L,b) , $p(\mathbf{w}^{(L,b)})$, can be estimated by applying PCA to the wavelet coefficients of N sample deformations \mathbf{f}_s at that band, denoted as $\mathbf{w}_s^{(L,b)}$, $s=1, 2, \dots, N$. After performing PCA, we obtain the mean of the wavelet coefficients $\bar{\mathbf{w}}^{(L,b)}$ and the matrix $\Phi^{(L,b)}$ formed by the eigenvectors of the covariance matrix of these coefficients, which correspond to the largest $K_{(L,b)}$ eigenvalues $\lambda_j^{(L,b)}$, $j=1, \dots, K_{(L,b)}$, of that matrix. Therefore, $\mathbf{w}^{(L,b)}$ can be represented by its projected vector or feature vector $\mathbf{v}^{(L,b)}$ in the space spanned by the $K_{(L,b)}$ eigenvectors,

$$\mathbf{v}^{(L,b)} = \Phi^{(L,b)T} (\mathbf{w}^{(L,b)} - \bar{\mathbf{w}}^{(L,b)}). \quad (2)$$

Then, the pdf of \mathbf{f} in Eq. (1) is represented by,

$$p(\mathbf{f}) = \prod_{b=0}^{B_L-1} c_{(L,b)} \exp \left\{ - \sum_{j=1}^{K_{(L,b)}} \frac{v_j^{(L,b)^2}}{2\lambda_j^{(L,b)}} \right\}, \quad (3)$$

where $c_{(L,b)}$ is the normalization coefficient and $v_j^{(L,b)}$ is the j th element of $\mathbf{v}^{(L,b)}$. After obtaining $p(\mathbf{f})$, we can now randomly generate new vectors $\hat{\mathbf{v}}^{(L,b)}$ conforming to the distribution in Eq. (3) and synthesize the wavelet coefficients of different wavelet bands using,

$$\hat{\mathbf{w}}^{(L,b)} = \Phi^{(L,b)} \hat{\mathbf{v}}^{(L,b)} + \bar{\mathbf{w}}^{(L,b)}. \quad (4)$$

A simulated deformation can therefore be generated by performing L -level inverse WPT.

This W-PCA model cannot only be used to model statistics of deformation fields, but also to model other fields like the Jacobian determinants of deformations. We stress the importance of using PCA within each wavelet band, which is in contrast to the commonly used independence assumption for wavelet coefficients. In particular, PCA is known to be the optimal linear expansion, provided that a good estimate of the covariance matrix is available. By decomposing deformation fields into different wavelet bands, we can estimate more efficiently the sample covariance matrices at various bands, for reasons that were detailed in the Wavelet-PCA (W-PCA) models for estimating the pdfs of deformation fields and their Jacobian determinants section. As a result, the W-PCA model can capture correlations between adjacent spatial locations at a given scale.

Hierarchical MRF regularization

As mentioned in the Wavelet-PCA (W-PCA) models for estimating the pdfs of deformation fields and their Jacobian determinants section, if deformations are synthesized directly using the W-PCA model, some unrealistic discontinuities emanating from the assumption of independence across wavelet bands may occur. In order to eliminate such potential discontinuities, a nested MRF regularization scheme that imposes spatial smoothness at different

scales is applied in conjunction with the inverse WPT. That is, \mathbf{f} is regularized at different scales: at level l , $l=L-1, \dots, 1, 0$, its wavelet coefficients $\hat{\mathbf{w}}^{(l,0)}$ are regularized. (Although for uniformity in the notation we call these wavelet coefficients, they actually result from the scaling functions and form a smooth approximation of \mathbf{f} at various scales.)

Denoting the input low-pass coefficients as $\hat{\mathbf{w}}$ (which can be any of $\hat{\mathbf{w}}^{(l,0)}$, $l=L-1, \dots, 0$), the MRF regularization estimates a “true” \mathbf{w}_r , by assuming that low-pass wavelet coefficients form an MRF and $\hat{\mathbf{w}}$ is a degraded observation of \mathbf{w}_r ($\hat{\mathbf{w}} = \mathbf{w}_r + \mathbf{n}$, where \mathbf{n} is the disturbance assumed to be zero-mean Gaussian noise with standard deviation (SD) σ_N), and by using the maximum a posteriori (MAP) framework (Gelge et al., 2000; Geman and Geman, 1984),

$$\mathbf{w}_r = \underset{\mathbf{w}}{\operatorname{argmax}} \{p(\mathbf{w} | \hat{\mathbf{w}})\} = \underset{\mathbf{w}}{\operatorname{argmax}} \{p(\hat{\mathbf{w}} | \mathbf{w})p(\mathbf{w})/p(\hat{\mathbf{w}})\}. \quad (5)$$

Assuming the priors $p(\hat{\mathbf{w}} | \mathbf{w})$ and $p(\mathbf{w})$ are Gaussian distributions, we have $p(\hat{\mathbf{w}} | \mathbf{w}) \propto \exp \left\{ -\frac{1}{2\sigma_N^2} \|\mathbf{w} - \hat{\mathbf{w}}\|^2 \right\}$ and $p(\mathbf{w}) \propto \exp \{-\Psi(\mathbf{w})\}$, where $\Psi(\mathbf{w}) = \frac{1}{2} (\mathbf{w} - \bar{\mathbf{w}})^T \chi^{-1} (\mathbf{w} - \bar{\mathbf{w}})$. $\bar{\mathbf{w}}$ and χ refer to the mean and the covariance matrix of \mathbf{w} , respectively, and the structure of χ meets the MRF property. Thus \mathbf{w}_r is solved by minimizing an energy function $E_r(\mathbf{w})$,

$$E_r(\mathbf{w}) = \frac{1}{2\sigma_N^2} \|\mathbf{w} - \hat{\mathbf{w}}\|^2 + \Psi(\mathbf{w}). \quad (6)$$

We use a simplified approach similar to Gelge et al. (2000) and Shen and Ip (1998) to minimize $E_r(\mathbf{w})$. First, we estimate $p(\mathbf{w})$ as a product of all the local (marginal) pdfs across the location x , i.e., $p(\mathbf{w}) = \prod_x G(\mathbf{w}_x, \mu_x, \sigma_x)$, where $G(\cdot)$ represents a single Gaussian distribution with mean μ_x and standard deviation σ_x . Then $\Psi(\mathbf{w})$ in Eq. (6) is estimated by $\hat{\Psi}(\mathbf{w}) = \sum_x \left\{ \frac{\|\mathbf{w}_x - \mu_x\|^2}{2\sigma_x^2} \right\}$, where $\mu_x = \frac{1}{|\delta(x)|} \sum_{y \in \delta(x)} \mathbf{w}_y$ and $\sigma_x^2 = \frac{1}{|\delta(x)|-1} \sum_{y \in \delta(x)} \|\mathbf{w}_y - \mu_x\|^2$. $\delta(x)$ refers to a neighborhood centered on x but not including x , and $|\delta(x)|$ is the cardinality of $\delta(x)$. Therefore, the regularized wavelet coefficients \mathbf{w}_r can be obtained by minimizing Eq. (6) using the Newton’s method.

Summary of SSD

In summary, the SSD algorithm samples the estimated pdf of deformations and then iteratively projects the deformation onto each of the three subspaces:

- Step 1. Initialization: randomly sample the W-PCA model of deformation fields, thereby generating a tentative deformation.
- Step 2. Project the deformation field onto the W-PCA model of valid deformation fields (this step is superfluous at the first iteration).
- Step 3. Project the Jacobian of the deformation field onto the W-PCA model of valid Jacobian determinants.
- Step 4. Modify the deformation field so that its Jacobian determinants match the target Jacobians generated in Step 3, and at the same time, subject to certain smoothness constraints on the deformation field (refer to Karacali and Davatzikos, 2004, for details).
- Step 5. Apply the nested MRF regularization to impose spatial smoothness on the deformation at all scales.

Step 6. Go to step 2 and iterate until convergence, i.e., until the smoothed deformation field belongs to the subspaces of valid Jacobians and deformations.

The landmark-constrained SSD

The SSD algorithm presented in the Statistical simulation of deformations (SSD) section can be used to randomly synthesize an unlimited number of deformations and respective images. In this section, we describe how to extend this approach in order to generate a new deformation using not only the statistical priors on the deformation fields, but also simultaneously satisfying constraints on a number of landmarks for which the deformation field is known. We call this method landmark-constrained SSD. The landmark-constrained SSD algorithm does not try to exactly match the input landmark displacements, and in fact it generates a deformation field by minimizing the target function, which, on one hand, must have higher likelihood from the statistical models, and on the other hand, must match the input landmarks as well as possible. Compared to other interpolation algorithms, such as the TPS-based interpolation and the Gaussian kernel-based interpolation, the advantage of the landmark-constrained SSD is that it can generate both richer and anatomically valid deformations.

Particularly, denote the desired displacement vector $\hat{\mathbf{f}}_M^{(d)}$ as the vector formed by all the known/prescribed displacements on M landmarks, the purpose of the landmark-constrained SSD is to estimate the whole deformation $\hat{\mathbf{f}}$ so that (1) it is subject to the statistics of deformations, or the likelihood of $\hat{\mathbf{f}}$ is maximized; and (2) the distance between the desired displacement vector $\hat{\mathbf{f}}_M^{(d)}$ and the actual displacement vector $\hat{\mathbf{f}}_M$ on the M landmarks is minimized. This can be achieved by finding the feature vectors $\hat{\mathbf{v}}^{(L,b)}$, $b=0, 1, \dots, B_L-1$, by minimizing the following objective function,

$$E\left(\hat{\mathbf{v}}(L, 0), \hat{\mathbf{v}}^{L,1}, \dots, \hat{\mathbf{v}}^{L,B_L-1}\right) = \|\hat{\mathbf{f}}_M^{(d)} - \hat{\mathbf{f}}_M\|^2 + \lambda \sum_{b=0}^{B_L-1} \sum_{j=1}^{K(L,b)} \frac{\hat{\mathbf{v}}_j^{(L,b)2}}{2\lambda_j^{(L,b)}}, \quad (7)$$

where the first term ensures that the actual displacements on the M landmarks are close to the desired ones, and the second term reflects the statistical constraint of the feature vectors $\hat{\mathbf{v}}^{(L,b)}$, and it ensures that the resultant deformation field $\hat{\mathbf{f}}$ has higher likelihood according to the distribution.

Because WPT, inverse WPT and PCA are all linear transformations, and because they can be represented by matrix forms, the actual displacement vector $\hat{\mathbf{f}}_M$ on the M landmarks can be directly

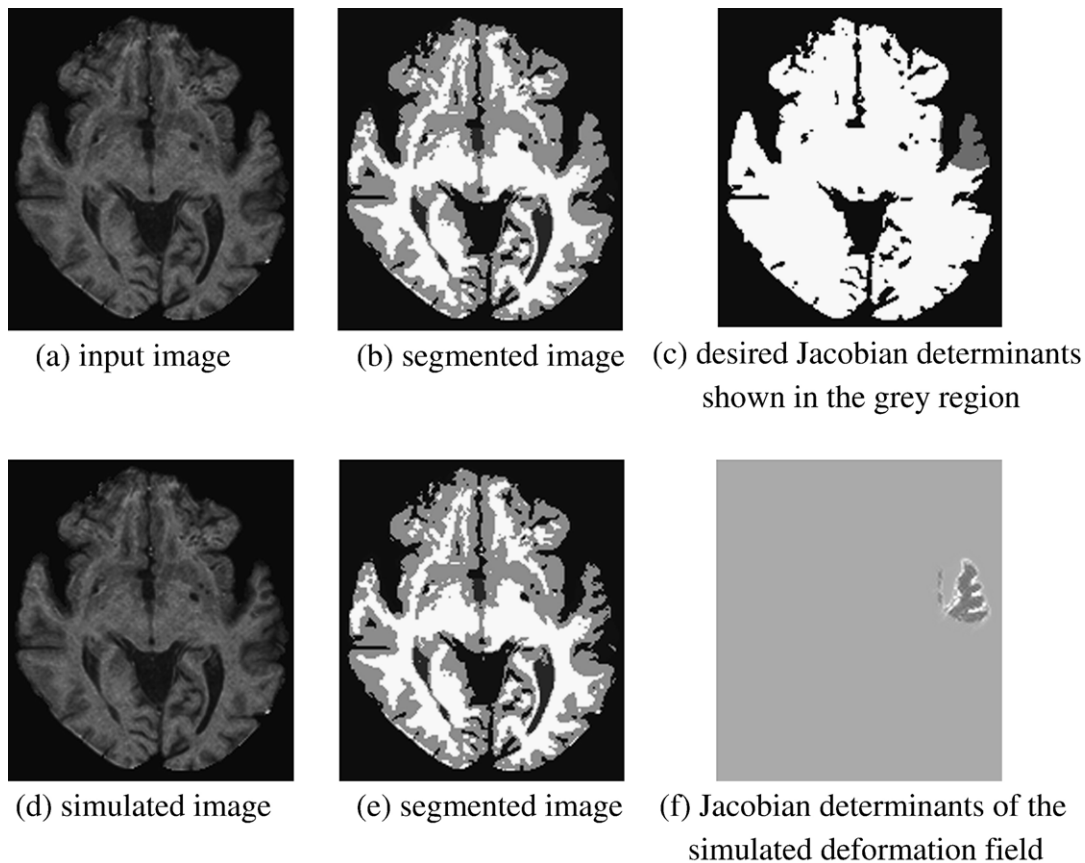


Fig. 4. Illustration of atrophy simulation. Panels a and b show the input image and its segmented image; panel c shows the prescribed volume change or Jacobian determinants. The grey regions give the desired rate of volume change within a spherical region; white indicates brain regions that do not require volume changes and black means CSF and background whose volume can be changed; panels d and e show the simulated image with atrophy and the segmented image; panel f gives the Jacobian determinants of the simulated deformation field, where Jacobian determinants are one within the large grey area, and they are smaller than one within darker grey regions and greater than one within lighter grey or white regions. It can be seen that shrinkage of the brain tissue is achieved in the desired atrophy region, and the surrounding CSF and background region has been enlarged to fill the shrinkage. Also, there are no deformations in other regions.

expressed as a function of the feature vectors $\hat{\mathbf{v}}^{(L,b)}$. First, according to inverse WPT, a deformation field $\hat{\mathbf{f}}$ can be reconstructed by B_L wavelet coefficients $\hat{\mathbf{w}}^{(L,b)}$, $b=0, 1, \dots, B_L-1$,

$$\hat{\mathbf{f}} = \sum_{b=0}^{B_L-1} W^{(L,b)} \hat{\mathbf{w}}^{(L,b)}, \quad (8)$$

where $W^{(L,b)}$ represents the inverse WPT matrix for band (L,b) . Because the displacement vector $\hat{\mathbf{f}}_M$ on the M landmarks can be calculated from the corresponding rows of matrix $W^{(L,b)}$, we can form a new matrix $W_M^{(L,b)}$ by selecting only the corresponding rows from $W^{(L,b)}$ that make contribution to $\hat{\mathbf{f}}_M$. Thus, based on Eqs. (8) and (4), $\hat{\mathbf{f}}_M$ can be expressed as,

$$\begin{aligned} \hat{\mathbf{f}}_M &= \sum_{b=0}^{B_L-1} W_M^{(L,b)} \hat{\mathbf{w}}^{(L,b)} = \sum_{b=0}^{B_L-1} W_M^{(L,b)} (\Phi^{(L,b)} \hat{\mathbf{v}}^{(L,b)} + \bar{\mathbf{w}}^{(L,b)}) \quad (9) \\ &= \sum_{b=0}^{B_L-1} (W_{M,\Phi}^{(L,b)} \hat{\mathbf{v}}^{(L,b)} + \mathbf{w}_{M,\bar{\mathbf{w}}}^{(L,b)}), \end{aligned}$$

where $W_{M,\Phi}^{(L,b)} = W_M^{(L,b)} \Phi^{(L,b)}$ and $\mathbf{w}_{M,\bar{\mathbf{w}}}^{(L,b)} = W_M^{(L,b)} \bar{\mathbf{w}}^{(L,b)}$. By substituting Eq. (9) into Eq. (7), the objective function can be explicitly expressed as a function of the feature vectors $\hat{\mathbf{v}}^{(L,b)}$, $b=0, 1, \dots, B_L-1$. We can then solve the minimization problem using the gradient descent method. After obtaining the feature vectors that minimize Eq. (7), the inverse WPT is used to reconstruct the deformation field $\hat{\mathbf{f}}$, and then the iterative mechanism of SSD is used to regularize $\hat{\mathbf{f}}$ to obtain the final estimated deformation.

Atrophy and growth simulation

In the previous subsections, we described how to generate deformation fields that reflect inter-individual variability in brain structure. In this subsection, we briefly discuss how to generate intra-individual deformations, which are important when algorithms for tracking growth or atrophy of brain tissue are to be evaluated. We used an extension of the approach described in (Karacali and Davatzikos, 2004). Specifically, the desired level of tissue growth or shrinkage within a pre-specified brain region is provided to the algorithm. An iterative procedure then tries to achieve this level of volumetric change, by seeking a deformation

field, whose Jacobian determinants have the appropriate values everywhere in the brain. For example, if tissue atrophy of 10% within a specific ROI is desired, then the determinant of the Jacobian of the underlying deformation field must be equal to 0.9 in that ROI, and equal to 1 everywhere else. An iterative projection is used to find a deformation field that not only has a valid gradient of a deformation field (integrability conditions in Karacali and Davatzikos, 2004), but also meets the aforementioned target values. Fig. 4 illustrates an example of the atrophy simulation in detail.

Simulation results

In our first experiment, we partially tested our assumption that it is important to constrain both the deformation fields and their Jacobian determinants. We used 68 deformation fields from a template of the corpus callosum to the corpus callosum of each of 68 individuals, as well as their Jacobian determinants. These deformations can be considered to be quite accurate representations of the underlying callosal morphology because they were based on outlining the callosal boundary and parameterizing piece-wise via two landmarks placed on its anterior and posterior ends; an elastic transformation interpolated the deformation in the interior of the callosum. Moreover, because the corpus callosum midsagittal section is a single, relatively small, and 2D structure, it is reasonable to expect that its pdf can be estimated reasonably well from 68 samples, at least compared to 3D whole-brain deformations. Accordingly, we constructed PCA models for the deformation fields and their Jacobians. Then, we simulated a large number (300 shown in Fig. 5) of deformation fields by randomly sampling the statistical distribution of deformation fields, thereby generating “valid” deformation fields, and then calculated the likelihood of the resultant Jacobian determinants based on their own PCA-estimated pdfs. Fig. 5 shows the resulting likelihood values of the Jacobian determinant of these “valid” deformation fields (the solid horizontal line defines the threshold on the likelihood based on which “valid” deformation fields were randomly generated). The results of Fig. 5 indicate that the Jacobians of about 54% of the randomly generated deformation fields are outside the subspace defined by the statistical model of Jacobians. In other words, many deformation fields that are generated by sampling their own estimated pdf and are “valid” according to this pdf have very unlikely Jacobian determinants. This shows that further restricting

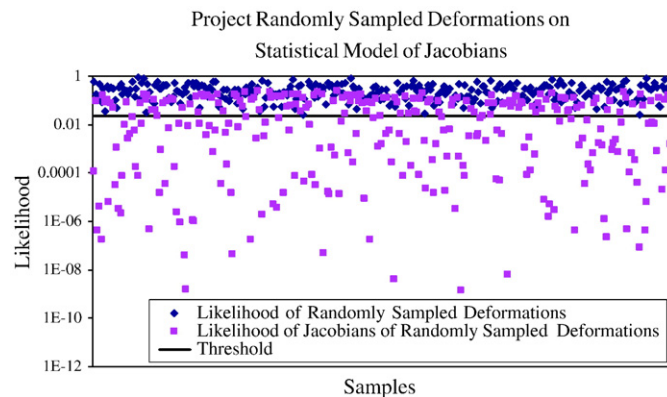


Fig. 5. Experiments showing the complementary property by combining the statistical models of deformation fields with those of the Jacobian determinants. The vertical axis shows the likelihood of the Jacobian of a synthesized deformation. Pink deformations were generating by sampling only the pdf of the deformation fields, whereas blue deformation fields were generated by constraining both the deformation and the Jacobian. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the Jacobian determinant is important for generating realistic deformations. After constraining both the deformation field and the Jacobian, the simulated deformation fields are all above the threshold line in Fig. 5, by construction, and thus have both valid deformations and Jacobians.

Original data

BLSA data

A set of 158 T1-weighted MR brain images of older healthy adults participating in the Baltimore Longitudinal Study of Aging (BLSA), as previously reported in (Resnick et al., 2000, 2003), were registered to an image selected as the template using the HAMMER registration algorithm (Shen and Davatzikos, 2002). These deformation fields were used as a starting point in order to obtain statistics of brain deformations, after the smoothness constraints and correction of Jacobians were applied to them (Karacali and Davatzikos, 2004).

MNI data

Forty T1-weighted MR brain images of young healthy adults from the McConnell Brain Imaging Centre, Montréal Neurological Institute, McGill University, were used as a second set of images (Kristi Boesen et al., 2005). Each image has 36 manually generated landmarks. One image was randomly selected as the template image and the HAMMER registration algorithm was used to obtain the deformation fields between two images.

Simulation results

Data set 1: randomly simulated data

Generally, the feature vectors $\mathbf{v}^{(L,b)}$, $b=0, 1, \dots, B_L-1$, are randomly sampled assuming that they follow a zero-mean Gaussian distribution described in Eq. (3). Inevitably, many of the resulting deformation fields are relatively close to the origin, i.e., to the mean deformation field, hence they are not very interesting. In order to better visualize and appreciate the properties of the deformations synthesized by our model, we biased our random sampling to select values further away from the origin, so that considerable deformations are generated.

Data set 1 was generated using this sampling method and using the SSD. It consists of 50 randomly synthesized deformations and respective images using the statistical models trained from the

BLSA data. The template and several representative synthesized images are shown in Fig. 6. The feature vectors of the simulated deformations are selected to be around 2 times the SD of their statistics.

Evaluation of W-PCA. Because W-PCA model is the key component of SSD to estimate statistics of deformations in SSD, we used the representation error to evaluate the performance of the W-PCA by comparing it with the global PCA method using different number of training samples. In the experiments, 3, 10, 20, 50, 100 and 150 deformations are selected as the training samples, respectively, and 16 deformations are used as the testing data (8 testing deformation fields are used for 150 training samples because we have a total number of 158 deformation fields). After training the W-PCA and PCA using the same training sample set, we project each of the testing samples onto the space of W-PCA and PCA, respectively, and then reconstruct the deformation fields using a number of principal components. The representation error between the testing deformation field and a reconstructed deformation field is defined as the mean of the voxel-wise deformation differences. Smaller representation errors imply that the statistical model accurately represents not only the training deformation samples, but other brain deformations, as well. Fig. 7 shows the plots of average representation error with respect to different numbers of samples. We can see that W-PCA yields smaller representation errors as compared with PCA. The performance of PCA and W-PCA is similar when the number of training samples is less than 10, simply because too few samples are used for capturing the statistics of high-dimensional data. It can be seen from the figure that W-PCA does present the testing deformation fields better than PCA.

Data set 2: simulation using landmark constraints

Given a number of anatomical landmarks for the 40 MNI images, which were defined in Kristi Boesen et al. (2005), we also used the landmark-constrained SSD presented in the Hierarchical MRF regularization section to generate simulated data sets that conformed to the prior statistics as well as to the user-defined landmarks. In order to generate simulated deformations with these data sets, we first randomly selected one of the 40 scans as the template (Fig. 8 shows the template image, as well as some manually marked landmarks), and then selected one




Fig. 6. Examples of the randomly simulated images. (a) A cross-section of the template image; (b–d) representative examples of cross-sections of the volumetric images synthesized by randomly sampling the statistical model of the deformations.




Fig. 7. Comparison of the representation errors of PCA and W-PCA by using different number of training samples.

of the rest 39 images as the test image. The statistical models of deformation fields and Jacobian determinants are then constructed by using the deformation fields that register the template image with the rest 38 subject images. Then, we used these statistical models, together with the 36 manual landmarks defined on the template image and the test image, respectively, to generate a new deformation using the landmark-constrained SSD. Fig. 9 shows the representative data sets by setting different subject images as the test image. It can be seen from the examples that the synthesized images using the landmark-constrained SSD are more realistic than those generated using other interpolation methods. This is because that the landmark-constrained SSD utilizes the statistical models of the deformation fields, and thus it picks up the one from the statistics that match the landmarks best. As a result, it comes up with more realistic deformation fields and images. On the other hand, the two interpolation methods used in conjunction with these landmarks (TPS and Gaussian kernel) generated deformation fields that were not only limited in flexibility, due to the limited number of landmarks but were often quite unrealistic (Figs. 9c, d). Notable is the similarity between the second row in Fig 9, which shows the warped template images using the landmark-constrained SSD, and the first row, which shows the actual images of the subjects from which the landmarks were derived, even though only 36 landmarks were defined for the entire brain.

Data set 3 and Data set 4: atrophy/growth simulation

We used the model of tissue atrophy/growth simulation that was briefly described in the Atrophy and growth simulation section to generate two sets of data, as described next.

Data set 3. For each image of the 40 MNI data, we simulated 7% atrophy and 7% growth within spherical regions centered on three manually selected locations. Fig. 10 gives three examples of the simulated images. The left column shows the image with three spherical regions at different locations of the brain, within which the atrophy and growth will be simulated. The middle column and the right column show the simulated atrophy and growth images, respectively. It can be seen that atrophy and growth are achieved only for the brain tissues within the selected spherical region, as expected.

Because the center points of the spherical regions are manually selected on each subject image, and the anatomical structures of the selected regions are different across different subjects, the atrophy/growth simulator does not necessarily achieve exactly the pre-specified 7% level. In fact, the simulation iteratively solves an optimization problem that can converge to a local minimum, and the brain tissue regions to be shrunk or grown are different for different subjects and different locations. Fig. 11 illustrates the average and SD of actual volume changes of the simulated atrophy and growth across the 40 images. We can see from the results that the actual atrophy and growth rates get quite close for different subjects and different locations.

Data set 4. In Data set 3, we simulated atrophy and growth on real subject images. In this data set, we simulate atrophy and growth on the simulated images of Data set 1 (50 images). Similar to Data set 3, for each image, we simulated 7% atrophy and growth of the brain tissues within spherical regions centered on three selected locations. Because simulated images are used and their deformations onto the template image are known, the three locations of atrophy and growth are selected only from the template image, and their corresponding locations in all other simulated images are automatically determined. The reason why we performed these additional simulations is that we have a relatively consistent placement of the sphere of atrophy/growth across different simulated images because they are simulated brain images emanating from the same template; whereas in Data set 3 the regions of simulated atrophy and growth were selected manually in each image, which could potentially involve some human error. Moreover, as many as desired brain images can be synthesized in this way, whereas Data set 3 can only generate a limited number of brains.




Fig. 8. The template image, which was randomly chosen from the 40 MNI images. The left image shows slice no. 79 of the template image, and the remaining images show a representative subset of the 36 landmarks.




Fig. 9. Three representative images from the 40 MNI data sets, along with three different sets of simulated images (warped template). It can be seen from the examples that the synthesized images using the landmark-constrained SSD are not only more realistic than those generated using other interpolation methods, but also quite close to the actual images from which the 36 constraining landmarks were derived.

Fig. 12 gives two examples of the simulated images, and Fig. 13 illustrates the average and SD of volume changes of the simulated atrophy and growth across the 50 simulated

images. Comparing Fig. 13 with Fig. 11, we can see that the actual volume changes of the simulated atrophy and growth for Data set 4 are more consistent than those in Data set 3

across different images, i.e., Fig. 13 has relatively small SD values.

Summary and conclusion

We have presented methods for intra- and inter-individual brain deformation simulation, which can be used in the future for systematic and extensive evaluation of registration algorithms. The statistical simulation of deformation (SSD) technique is applied in order to capture the statistics of deformations and to simulate new deformations and respective images. In SSD, a Wavelet–Packet Transform (WPT) decomposition of the training deformations, in conjunction with a Markov random field (MRF) spatial regularization, is used to capture both coarse and fine characteristics of the training deformations in a statistical fashion. Simulated deformations are constructed by randomly sampling the resultant statistical distribution. In addition, when landmark points are available, the landmark-constrained SSD is applied in order to utilize both landmark and statistical

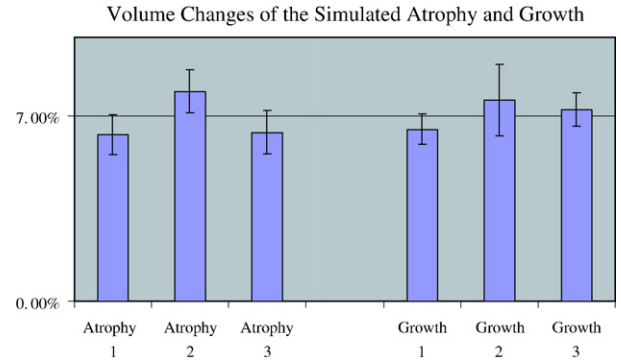


Fig. 11. Average and SD values of the actual volume changes measured from the simulated atrophy and growth of the 40 MNI images.

information. Finally, the proposed deformation simulation approach is coupled with a model for generating tissue atrophy or growth.




Fig. 10. Examples of simulated atrophy and growth within spherical regions centered on three selected locations. Tissue atrophy and growth are simulated within the spherical region illustrated by the red circles. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

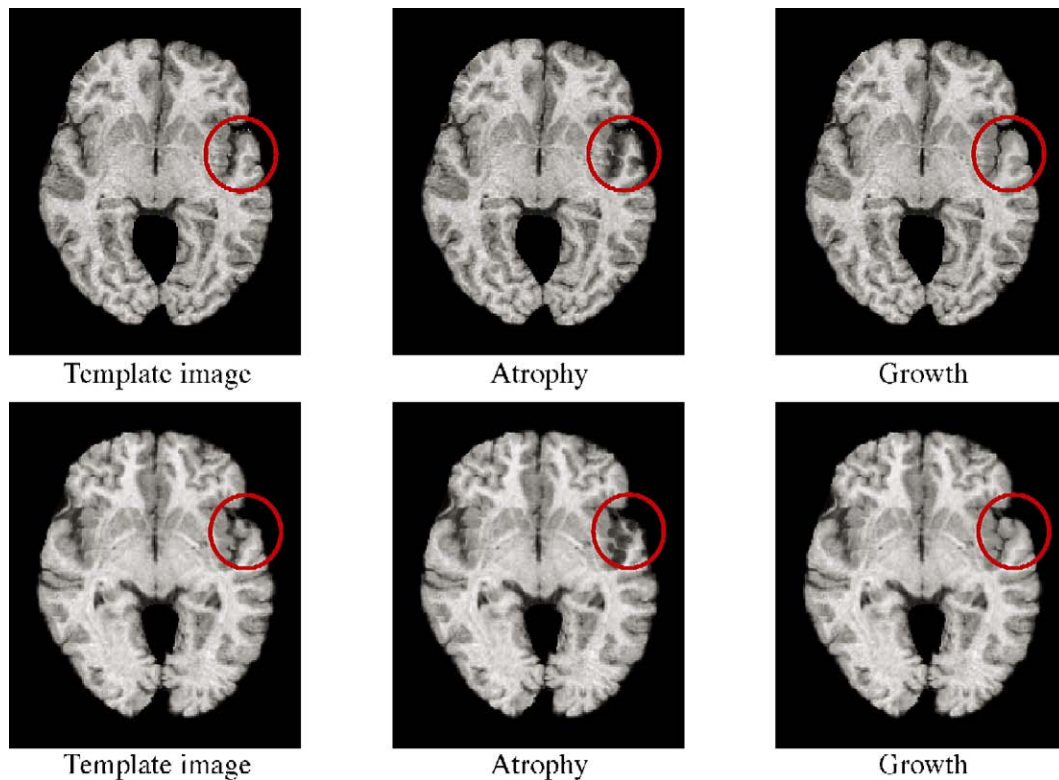


Fig. 12. Examples of the simulated atrophy and growth of the template image and a simulated image.

A main assumption in the W-PCA method is that a number of training samples, i.e., of example deformation fields, are available for pdf estimation purposes. For the purposes of describing our simulation method herein we used a specific deformable registration method, which biases the estimated pdf towards the deformations that this method produces. Our rationale has been that, in future studies aiming at validation of different warping methods, a set of well-labeled images can be first constructed and used to determine the training deformations in a rather unbiased way. This is because if one constrains the warping adequately via landmarks, curves or regions, highly accurate deformations can be obtained with minimal bias towards the warping method. Secondly, our approach to simulating deformations can be used in conjunction with different warping methods, each time using one of them as ground truth and testing the others. A warping method

that consistently outperforms others, regardless of the training process, is likely to be a more accurate method.

The assumption of independence between different bands, after the WPT, is only an approximation. Although our method significantly expands upon the common assumption in many signal-processing paradigms of independence of different wavelet coefficients, by grouping together coefficients in each band and performing PCA, it ignores possible correlations across different bands. Future work in this direction will aim at grouping different wavelet coefficients into groups that have minimal statistical dependence on each other, and which might span across different frequency bands.

The framework of iterative projections onto different subspaces (Fig. 1) can be extended to include additional subspaces to the ones we currently consider. More generally seen, this framework is based on the fact that in situations in which very accurate estimation of a pdf is not directly feasible, the pdf can be estimated by evaluating various (hopefully complementary) aspects of it and finding the solution that satisfies all. Examples of potential extensions include prior expert knowledge defining regions that are expected, from an anatomical perspective, to be correlated to each other, where the joint marginal distribution of these regions can be estimated. The respective subspace would then include all the pdfs of anatomical structures, whose marginal distributions relative to the expert-defined anatomical structures agree with what was estimated from the training samples.

In this paper, we generated several data sets of simulated deformations for future use in validation studies of automated segmentation and deformable registration methods. The software and data sets are available for distribution online at <http://www.rad.upenn.edu/sbia/>.




Fig. 13. Average and SD values of the actual volume changes of the simulated atrophy and growth of the 50 simulated images.

References

- Bookstein, F.L., 1989. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* 11, 567–585.
- Cootes, T.F., Hill, A., Taylor, C.J., Haslam, J., 1994. Use of active shape models for locating structures in medical images. *Image Vis. Comput.* 12, 355–365.
- Davatzikos, C., Tao, X., Shen, D.G., 2003a. Hierarchical active shape models using the wavelet transform. *IEEE Trans. Med. Imag.* 22, 414–423.
- Davatzikos, C., Tao, X., Shen, D., 2003b. Applications of Wavelets in Morphometric Analysis of Medical Images, presented at SPIE Wavelets X, San Diego, CA.
- Gelge, M., Kilmer, M., Miller, E., 2000. Wavelet domain image restoration with adaptive edge-preserving regularization. *IEEE TMI* 6, 721–741.
- Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-6, 721–741.
- Karacali, B., Davatzikos, C., 2004. Estimating topology preserving and smooth displacement fields. *IEEE Trans. Med. Imag.* 23, 868–880.
- Karacali, B., Davatzikos, C., 2006. Simulation of tissue atrophy using a topology preserving transformation model. *IEEE Trans. Med. Imag.* 25, 649–652.
- Kristi Boesen, S.F., Huang, J., Germann, J., Stern, J., Louis Collins, D., Evans, A.C., Rottenberg, D.A., 2005. Inter-rater reproducibility of 3D cortical and subcortical landmark points. Presented at 11th Annual Meeting of the Organization for Human Brain Mapping, Toronto, Canada, 2005.
- Mallat, S., Wavelet, A., 1998. *Tour of Signal Processing*. Academic Press.
- Miller, M., Banerjee, A., Christensen, G., Joshi, S., Khaneja, N., Grenander, U., Matejic, L., 1997. Statistical methods in computational anatomy. *Stat. Methods Med. Res.* 6, 267–299.
- Resnick, S.M., Goldszal, A., Davatzikos, C., Golski, S., Kraut, M.A., Metter, E.J., Bryan, R.N., Zonderman, A.B., 2000. One-year age changes in MRI brain volumes in older adults. *Cereb. Cortex* 10, 464–472.
- Resnick, S.M., Pham, D.L., Kraut, M.A., Zonderman, A.B., Davatzikos, C., 2003. Longitudinal magnetic resonance imaging studies of older adults: a shrinking brain. *J. Neurosci.* 23, 3295–3301.
- Shen, D., Davatzikos, C., 2002. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imag.* 21, 1421–1439.
- Shen, D., Ip, H., 1998. Markov random field regularization models for adaptive binarization of nonuniform images. *IEE Proc., Vis. Image Signal Process.* 145, 322–332.
- Xue, Z., Shen, D., Karacali, B., Davatzikos, C., 2005. Statistical representation and simulation of high-dimensional deformations: Application to synthesizing brain deformations. Presented at Medical Image Computing and Computer Assisted Intervention (MICCAI 2005), Palm Springs, California, USA.