

Statistical Representation and Simulation of High-Dimensional Deformations: Application to Synthesizing Brain Deformations

Zhong Xue, Dinggang Shen, Bilge Karacali, Christos Davatzikos

Section of Biomedical Image Analysis, Department of Radiology
University of Pennsylvania, Philadelphia, PA 19104
{zhong.xue,dinggang.shen,christos.davatzikos}@uphs.upenn.edu

Abstract. This paper proposes an approach to effectively representing the statistics of high-dimensional deformations, when relatively few training samples are available, and commonly used methods, like PCA, fail due to insufficient training. Based on previous work on scale-space decomposition of deformation fields, herein we represent the space of “valid deformations” as the intersection of three subspaces: one that satisfies constraints on deformations themselves, one that satisfies constraints on Jacobian determinants of deformations, and one that represents smooth deformations via a Markov Random Field (MRF). The first two are extensions of PCA-based statistical shape models. They are based on a wavelet packet basis decomposition that allows for more accurate estimation of the covariance structure of deformation or Jacobian fields, and they are used jointly due to their complementary strengths and limitations. The third is a nested MRF regularization aiming at eliminating potential discontinuities introduced by assumptions in the Wavelet-PCA model. A randomly sampled deformation field is projected onto the space of valid deformations via iterative projections on each of these subspaces until convergence, upon which all three constraints are met. A deformation field simulator uses this process to generate random samples of deformation fields that are not only realistic but also representative of the full range of anatomical variability. These simulated deformations can be used for validation of deformable registration methods. Other potential uses of this approach include representation of shape priors in statistical shape models as well as various estimation and hypothesis testing paradigms in the general fields of computational anatomy and pattern recognition.

1 Introduction

Representing prior statistical knowledge of high-dimensional scalar or vector fields is of fundamental importance in a variety of scientific areas including computational anatomy, shape analysis, pattern recognition, and hypothesis testing applied to images or their deformations [1–4]. For instance, statistical study of deformations can be used to provide voxel-based morphological characterization

of different groups; to incorporate prior knowledge of deformations from training samples into image segmentation and registration algorithms; to provide an efficient way of synthesizing new deformation fields for validation of registration and segmentation methods; to regularize deformations according to prior knowledge of sample deformations; and to estimate the missing parts of a deformation from parts that are observed.

The goal of this paper is to construct a statistical model of deformations from a limited number of training samples and to simulate deformations of MR brain images by sampling this model. By simulating deformations, we can synthesize respective images, which can then be used for validation of various segmentation and registration algorithms since the deformations are known [5]. Although many statistical shape modeling methods have been proposed in literature, they are often designed for 2-D or 3-D shapes that can be represented by a relatively small number of landmarks or outline points [2]. However, the high dimensionality of a variety of image warping methods of 3-D structural images renders simple PCA-based methods unable to properly estimate the statistics of deformation fields, from the typically limited number of training samples [6]. This is especially true for the finer local detail of a deformation.

This paper builds upon previous methods for estimating the covariance structures of high-dimensional distributions via scale-space decompositions [7–9]. In particular, the approach described herein utilizes an expansion to a wavelet packet basis [6] to rotate the high-dimensional coordinate system in which the discretized deformation field is defined, so that its covariance structure is close to a block-diagonal one; each block of this covariance matrix is approximated by its principal components via a scale-wise PCA, which captures spatial correlations among wavelet coefficients within a specific scale and band. The probability density function (pdf) of deformations can therefore be approximated by a product of the pdfs derived from the different wavelet-based PCA subspaces. This approach, referred to as the Wavelet-PCA (W-PCA) model, effectively implies the assumption of independence across different scales/bands. Although this assumption is not strictly true, it is found to lead to far more accurate estimation of the pdf of the deformation field, compared to global PCA. It is also motivated by the structural and functional correlations found among adjacent anatomical regions, at various scales. For example, in the brain adjacent neurons typically display dense connections among each other, and the same holds for adjacent anatomical regions at the substructure as well as at the between-structures level.

In this paper, the W-PCA model is applied not only to deformation fields, but also to Jacobian determinants. The rationale is that the Jacobian determinant field is much smoother than the deformation field itself. Cortical gyri of the brain can have quite variable curvature patterns, but their volumes do not differ greatly, thereby leading to tighter pdfs of volumetric measurements and making them easier to estimate. In addition, since the W-PCA model of deformations can introduce some unrealistic discontinuities emanating from the assumption of independence across wavelet bands, a valid deformation should also be a smooth field represented by a MRF. Therefore, a deformation field sim-

ulator is designed to synthesize a valid deformation via iteratively constraining a randomly sampled deformation so that it lies in the subspaces defined by the statistics of deformations, the statistics of Jacobian determinants, as well as the smoothness constraints.

The experiments synthesize deformations and respective images by randomly sampling the pdfs determined as above. Quantitative measures of generalization of the approach are evaluated and compared with those of the global PCA method. The results show that the proposed statistical model of deformations captures the prior knowledge of sample deformations well and generates realistic deformations and simulated warped images.

2 Methods

2.1 General Description

Wavelet-PCA (W-PCA) statistical priors Let $\mathbf{f}(\mathbf{x})$ be a scalar or vector field defined over the template image domain Ω_t , $\mathbf{x} \in \Omega_t$. Trying to estimate the pdf of \mathbf{f} from a relatively small number of training samples is a very difficult task. The commonly used PCA method (e.g. [1, 2, 4]) fails miserably when \mathbf{f} is of very high dimensionality [6]. If \mathbf{f} represents a 3-D warping transformation and it is to be estimated from 100 training samples, a global PCA model will 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 [9], which 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 (e.g. the PCA model derived from a high-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 emanating from the small window size). 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. The fundamental assumption here is that the wavelet-based rotation renders the covariance matrix of \mathbf{f} close to block-diagonal, thereby leading to a more accurate estimation 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.

“Valid deformations” as the intersection of different subspaces 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 all available constraints, *i.e.* it belongs to the intersection of a number of subspaces, each of which satisfies some constraints on the deformation. The W-PCA model applied to the deformation field specifies one such subspace. In order to describe the second subspace, we first observe that if \mathbf{f} represents a linear or nonlinear function of the deformation field that can be estimated more accurately, we can use it to further constrain the form of the deformation field. Herein we also use the determinant of the Jacobian of the deformation field in place of \mathbf{f} in the W-PCA formulation above. The Jacobian determinant is not only of lower dimensionality, since it is a scalar and not a vector field, but it is also much less variable across individuals, for the reason that is pictorially shown in Fig. 2. Accordingly, it can be estimated much more accurately from the typically limited training samples, using the W-PCA framework. In order to find the deformation field that satisfies a given Jacobian determinant, we use [10].




Fig. 1. The space of valid deformations are represented as the intersection of different subspaces reflecting different aspects of deformations.




Fig. 2. Jacobians are less variable than displacement fields. *E.g.* 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.

Hierarchical regularization The W-PCA model of deformations can introduce some unrealistic discontinuities emanating from the assumption of independence across different scales and bands. Such potential discontinuities can be eliminated by using a nested regularization scheme, which is applied sequentially to each level of the wavelet decomposition via a respective MRF that imposes spatial smoothness at different scales.

Summary of the algorithm An algorithm that samples the resulting pdfs is described next and detailed in the following sections.

- Step 1. Randomly sample the W-PCA model of deformation fields, thereby generating a tentative deformation.

- Step 2. Project the Jacobian of the deformation field onto the W-PCA model of valid Jacobian determinants.
- Step 3. Find the deformation field whose Jacobian matches the one generated in Step 2, using [10].
- Step 4. Apply the nested MRF regularization to impose spatial smoothness on the deformation at all scales.
- Step 5. Iterate the above steps until convergence, *i.e.* until the smoothed deformation field belongs to the subspaces of valid Jacobians and valid deformations.

2.2 The Wavelet-PCA (W-PCA) Model

The W-PCA model described in Section 2.1 is detailed in this section. It 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} , it then constructs a PCA model of the wavelet coefficients of each wavelet band at level L , and finally it combines these pdfs together. Fig.3 illustrates the structure of 1-D WPT. For 3-D 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, $\mathbf{w}^{(l,0)}$ represents the low-pass wavelet coefficients. For simplicity, \mathbf{f} is also referred to




Fig. 3. Illustration of Wavelet Packet Transformation (WPT).

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, the pdf of deformation \mathbf{f} can be expressed as,

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

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 $\mathbf{v}^{(L,b)}$ 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 calculated by,

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

where $c_{(L,b)}$ is the normalization coefficient. After obtaining $p(\mathbf{f})$, we can generate randomly new vectors $\hat{\mathbf{v}}^{(L,b)}$ according to the pdf 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 can not only be used to model statistics of deformations, but also be used to model other fields like the Jacobian determinants of deformations.

We stress the importance of using PCA within each 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. Although sample covariance is a very inaccurate estimate of the covariance of \mathbf{f} , the sample covariance at various scales provides a much better estimate of the covariance at that scale, for reasons that were detailed in Section 2.1. As a result, the W-PCA model can capture correlations between adjacent spatial locations at a given scale.

2.3 Hierarchical MRF Regularization

As mentioned in Section 2.1, 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 coefficient $\hat{\mathbf{w}}^{(l,0)}$ is regularized.

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 obey MRF conditions and $\hat{\mathbf{w}}$ is a degraded observation of \mathbf{w}_r ($\hat{\mathbf{w}} = \mathbf{w}_r + \mathbf{n}$, \mathbf{n} is the disturbance assumed to be zero-mean Gaussian noise with standard deviation (std) σ_N), and by using the Maximum *a posteriori* (MAP) framework [11, 12], *i.e.*

$$\begin{aligned} \mathbf{w}_r &= \operatorname{argmax}_{\mathbf{w}} \{p(\mathbf{w}|\hat{\mathbf{w}})\} \\ &= \operatorname{argmax}_{\mathbf{w}} \{p(\hat{\mathbf{w}}|\mathbf{w})p(\mathbf{w})/p(\hat{\mathbf{w}})\}. \end{aligned} \quad (5)$$

Assuming the priors $p(\hat{\mathbf{w}}|\mathbf{w})$ and $p(\mathbf{w})$ are Gaussian distributions, we have $p(\hat{\mathbf{w}}|\mathbf{w}) \propto \exp\{-\frac{1}{2\sigma_N^2}\|\mathbf{w} - \hat{\mathbf{w}}\|^2\}$ 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 [11, 13] to minimize $E_r(\mathbf{w})$. First, we estimate $p(\mathbf{w})$ as a product of all the local (marginal) pdfs across the locations \mathbf{x} ,

i.e. $p(\mathbf{w}) \simeq \prod_{\mathbf{x}} G(\mathbf{w}_{\mathbf{x}}, \boldsymbol{\mu}_{\mathbf{x}}, \sigma_{\mathbf{x}})$, where $G(\cdot, \cdot)$ represents a single Gaussian distribution with mean $\boldsymbol{\mu}_{\mathbf{x}}$ and std $\sigma_{\mathbf{x}}$. Then $\Psi(\mathbf{w})$ in Eq.(6) is estimated by $\hat{\Psi}(\mathbf{w}) = \sum_{\mathbf{x}} \left\{ \frac{\|\mathbf{w}_{\mathbf{x}} - \boldsymbol{\mu}_{\mathbf{x}}\|^2}{2\sigma_{\mathbf{x}}^2} \right\}$, where $\boldsymbol{\mu}_{\mathbf{x}} = \frac{1}{|\delta(\mathbf{x})|} \sum_{\mathbf{y} \in \delta(\mathbf{x})} \mathbf{w}_{\mathbf{y}}$ and $\sigma_{\mathbf{x}}^2 = \frac{1}{|\delta(\mathbf{x})|} \sum_{\mathbf{y} \in \delta(\mathbf{x})} \|\mathbf{w}_{\mathbf{y}} - \boldsymbol{\mu}_{\mathbf{x}}\|^2$. $\delta(\mathbf{x})$ refers to a neighborhood centered on \mathbf{x} but not including \mathbf{x} , and $|\delta(\mathbf{x})|$ is the cardinality of $\delta(\mathbf{x})$. Therefore, the regularized wavelet coefficients \mathbf{w}_r can be obtained by minimizing Eq.(6) iteratively using Newton's method.

3 Results

In the experiments, the W-PCA models of deformations and the W-PCA models of their Jacobian determinants are first constructed, and then new deformations and their respective images are synthesized using the proposed deformation simulator. The dataset used in this experiment includes $N = 158$ MR brain images of different subjects. All the subject images were first rigidly transferred onto the space of the template image, and then the registration [14] was used to obtain the deformations from the template image to all the 158 rigidly-aligned subject images. After constructing the W-PCA models, we simulated randomly a large number (more than 50) of deformations and respective images using the deformation simulator described in Section 2.1. These deformations and images were visually evaluated and the simulated deformations were found realistic, smooth and also with quite significant variations. An example of the synthesized deformation fields and the respective image is given in Fig.4.



(a) the template image (b) a simulated deformation (c) the simulated image

Fig. 4. An example of the simulated deformation.

We also used the generalization measure in [9, 15] to evaluate the performance of our approach. The average and std of the generalization errors are calculated using the leave-one out method. The generalization error is defined as the mean of the absolute voxel-wise differences between a testing deformation field and its counterpart after projection to the intersection of the three subspaces (see Fig.1). We also compared the results using our approach with those using the global PCA. It turned out that the average and std of the generalization errors of the proposed approach are 2.3mm and 0.6mm respectively, which is a significant (P-value close to zero) improvement comparing to the results of the global PCA (mean 7.6mm and std 1.5mm).

4 Conclusion

We presented a method for estimating the statistics of deformations, which is then utilized by a deformation simulator to generate realistic deformations and images. The W-PCA models capture the statistics of deformations and their Jacobian determinants, and the nested MRF regularization eliminates possible discontinuities of deformations. A valid deformation is simulated by iteratively constraining a randomly generated deformation, so that it belongs to the intersection of these three subspaces. Better performance is observed by comparing to the global PCA approach using subjective and quantitative evaluations.

References

1. Cootes, T., Hill, A., Taylor, C., Haslam, J.: Use of active shape models for locating structures in medical images. *Image and Vision Computing* **6** (1994) 355–365
2. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models - their training and application. *Computer Vision and Image Understanding* **61** (1995) 38–59
3. Staib, L., Duncan, J.: Boundary finding with parametrically deformable models. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **14** (1992) 1061–1075
4. Miller, M., Banerjee, A., Christensen, G., Joshi, S., Khaneja, N., Grenander, U., Matejic, L.: Statistical methods in computational anatomy. *Statistical Methods in Medical Research* **6** (1997) 267–299
5. Schnabel, J., Tanner, C., Castellano-Smith, A., Degenhard, A., Leach, M., Hose, D.: Validation of non-rigid registration using finite element methods: application to breast MR images. *IEEE Trans. on Medical Imaging* **22** (2003) 238–247
6. Mallat, S.: A wavelet tour of signal processing. Academic Press (1998)
7. Coifman, R., Wickerhauser, M.: Entropy-based algorithms for best basis selection. *IEEE Trans. on Information Theory* **38** (1992) 713–718
8. Davatzikos, C., Tao, X., Shen, D.: Hierarchical active shape models using the wavelet transformation. *IEEE Trans. on Medical Imaging* **22** (2003) 414–423
9. Mohamed, A., Davatzikos, C.: Shape representation via best orthogonal basis selection. In: MICCAI 2004, St. Malo, France (2004) 225–233
10. Karacali, B., Davatzikos, C.: Simulation of tissue atrophy using a topology preserving transformation model. Submit to *IEEE Trans. on Medical Imaging* (2004)
11. Belge, M., Kilmer, M., Miller, E.: Wavelet domain image restoration with adaptive edge-preserving regularization. *IEEE Trans. on Image Processing* **9** (2000) 597–608
12. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **6** (1984) 721–741
13. Shen, D., Ip, H.: Markov random field regularization models for adaptive binarization of nonuniform images. *IEE Proc. Vis. Image Signal Process.* **145** (1998) 322–332
14. Shen, D., Davatzikos, C.: HAMMER: Hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. on Medical Imaging* **21** (2002) 1421–1439
15. Styner, M.A., Rajamani, K.T., Nolte, L.P., Zsemlye, G., Székely, G., Taylor, C.J., Davies, R.H.: Evaluation of 3D correspondence methods for model building. In: *Information Processing in Medical Imaging*. (2003)