

# A Hypothesis Testing Framework for High-Dimensional Shape Models

Joshua Cates, Thomas Fletcher, Ross Whitaker

► **To cite this version:**

Joshua Cates, Thomas Fletcher, Ross Whitaker. A Hypothesis Testing Framework for High-Dimensional Shape Models. Xavier Pennec. 2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy, Oct 2008, New-York, United States. pp.170-181, 2008. <inria-00632885>

**HAL Id: inria-00632885**

**<https://hal.inria.fr/inria-00632885>**

Submitted on 16 Oct 2011

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# A Hypothesis Testing Framework for High-Dimensional Shape Models

Joshua Cates, P. Thomas Fletcher, Ross Whitaker

Scientific Computing and Imaging Institute  
University of Utah  
Salt Lake City, Utah

**Abstract.** Statistical shape models are powerful tools for describing anatomical structures and are increasingly being used in a wide variety of clinical and biological contexts. One of the promising applications of this technology is the testing of hypotheses that entail shape differences, and visualization of those differences between cohorts. Statistical testing of shapes, however, is difficult due the large numbers of degrees of freedom and challenge of obtaining sufficient numbers of subjects to ensure statistical power. To date, research in statistical shape modeling has focused mainly on the construction of representative models, and the field has not reached a consensus on the best approach to statistical hypothesis testing. This paper illustrates some problems inherent in the statistical analysis of high-dimensional shape models, and suggests a systematic approach to hypothesis testing that avoids those problems. The proposed framework is based on research in the factor analysis statistics literature, with permutation testing in the PCA space of the model, and dimensionality reduction via a simulation-based analysis. We also describe two methods for visualizing group mean differences, first by direct visualization of the linear discriminant implicit in the hypothesis test metric, and second, by visualizing strain tensors from a deformation computed between the group means. We illustrate the proposed analysis and visualization framework on several clinical and biological datasets.

## 1 Introduction

Many important fields of basic research in medicine and biology now routinely employ statistical models of shape to quantify the anatomical variation in populations. Often, researchers are also interested in hypothesis testing to evaluate or demonstrate shape differences between populations. One such application area, for example, is the study of gene function as it pertains to human development and disease. Modern gene targeting technology allows researchers to create specific alterations in a mouse genome that result in different patterns of anatomical growth and form, or *phenotypes*, which can be modeled as shape and compared with normal populations to gain insight into the functionality of the targeted genes [1, 2]. Many areas of clinical psychiatric and neurological research also employ statistical shape analysis. The study of autism and its impact on brain

regions is one notable example [3, 2]. Statistical models that capture the variation in the shape of brain structure, including the *covariation* among multiple substructures, are increasingly necessary for researchers to gain understand into the development and progression of neurological disorders [4, 5].

Anatomical shape from images can be represented and computed using a variety of tools. Traditional representations of shape for phenotyping, for example, have relied on explicitly chosen *landmark* positions to define relatively *low-dimensional* parameterized models (e.g. [6]). Researchers in the 3D medical imaging community, however, have more recently pioneered the use of *high-dimensional* models of shape, which consist of very large collections of surface parameters that are derived automatically from the 3D images. High-dimensional shape models are appealing because they provide a much more detailed description of anatomy than landmark-based models, and do not require an a-priori choice of surface homologies. One common strategy for high-dimensional shape modeling is to consider shapes as embedded in images, and then to deformably register images and perform statistics on those deformations (e.g. [7]). Another common approach is to construct parameterized or point-based shape descriptors and compute statistics on those descriptions (e.g. [4, 3]). In the latter case, technologies exist to construct compact models that are optimized with respect to the information content of the population [8, 2], an important consideration for subsequent statistical analysis.

While high-dimensional models offer significant advantages for shape representation, their statistical analysis is not straightforward. The large number of degrees of freedom in the shape space, often coupled with a relatively low sample size (HDLSS), means that traditional low-dimensional statistical metrics cannot be directly applied [9]. While the shape modeling literature has proposed methods for analysis, it has not reached a consensus regarding a systematic approach that addresses the HDLSS problem. The statistics literature addresses the HDLSS problem with a variety of useful techniques, but these tools must be applied systematically in order to avoid either under-powered studies or over-optimistic conclusions. Through examples on clinical and biological datasets, this paper illustrates some of the potential difficulties that are encountered in high-dimensional shape analysis. We focus on the problems of hypothesis testing for group differences, and the visualization of those group differences. For hypothesis testing, we suggest permutation testing in a lower-dimensional PCA subspace of the model. For the dimensionality reduction, we propose using a simulation-based method to choose dimensions whose variance is distinguishable from noise. For visualization of group differences, we describe two approaches. The first is a direct visualization of the linear discriminant vector implicit in the hypothesis test, and the second is a visualization of strain tensors derived from a thin-plate spline deformation between the group mean shapes.

## 2 Related Work


Hypothesis testing for group differences using high-dimensional shape models has been most extensively investigated for comparative studies of brain anatomy. To date, however, researchers have typically each chosen different strategies for analysis, and there has not been a systematic treatment of the effects of HDLSS and dimensionality reduction choices on the statistical tests. Davies, et al. investigate hippocampus shape in schizophrenia populations with spherical harmonic and MDL-based shape models, and perform classification trials based on linear discriminant analysis in the high-dimensional shape space [8]. Terriberry [10] proposes a framework for multivariate, permutation-based hypothesis testing using nonlinear m-rep shape models, with an application to the lateral ventricles of twins. Styner, et al. propose point-wise hypothesis testing on shape correspondence points derived from spherical harmonic parameterizations [11, 12]. Golland [13] proposes a Support Vector Machine algorithm for training group classifiers of distance-transform shape representations. More recently, Gorczowski, et al. use medial axis shape representations and distance-weighted discriminants to compare complexes of brain structures in autistic and normal populations [3].

Several researchers have explored hypothesis testing using high-dimensional models for mouse phenotyping. Chen, et al. perform hypothesis testing using image-based metrics on deformation-based shape models to illustrate differences between strains of mutant mice [7]. In a phenotyping study of *Hoxd11* knock-out mice, the authors employ univariate hypothesis testing with regression analysis on PCA bases of point-based models of the mouse forepaw [14].

## 3 Challenges for HDLSS Shape Statistics

This section describes some problems and solutions in HDLSS shape statistics as they relate to shape analysis. The concepts presented here are applicable to linear statistics using any high-dimensional model, though we will use surface-point correspondence models as examples, computed by the particle-based optimization described in [2]. Point-based models represent shape by sampling each shape surface in a consistently ordered fashion so as to define homologous object surface points called *correspondences*. The set of 3D positions of all  $m$  correspondences on a shape is  $3m$  shape vector, and the positions of the individual shapes in this  $3m$ -dimensional shape space give rise to the statistical analysis. Hypothesis testing is done on a single shape model constructed from all data without knowledge of the group classification, which we refer to as a *combined* model.

In the context of point-based models, or surface samplings obtained from parameterized models, one approach to shape statistics is point-wise analysis of correspondences, which are elements of  $R^3$ . These data are drawn from the *marginal distributions* of the full shape space, and the mean shape is obtained by computing the Euclidean averages of correspondence positions, with point-wise differences defining local shape variation [11]. Hypothesis tests in this case reveal *regions* of significant differences between groups, which can be directly visualized




**Fig. 1.** Point-wise hypothesis test results for the putamen

as  $p$ -value maps on the mean shapes. Styner [11], for example, proposes statistical analysis of the correspondences that uses a nonparametric permutation test with the Hotelling  $T^2$  metric with an FDR correction for the multiple-comparison problem inherent in the analysis.

Figure 1 is an illustration of point-wise hypothesis testing as proposed in [11] on a combined model of the right putamen from normal control subjects and autism patients. The data is taken from an ongoing longitudinal pediatric autism study[15]. For the test, we had 10 autism structures available with 15 matched normals, and used 1024 correspondence points and 20,000 test permutations. The uncorrected  $p$ -values that indicate significance at the 5% level are colored in red on the mean normal putamen surface in Fig 1a, and suggest several distinct areas of shape differences. Fig. 1b shows that in this case, however, which is not uncommon in neurological shape analysis, no significant  $p$ -values remain after FDR correction (5% bound). This example illustrates a major difficulty encountered in point-wise analysis: the large number of comparisons results in a very conservative correction of the hypothesis test results, significantly reducing the statistical power of the test.

To avoid the multiple-comparisons problem, we can analyze high-dimensional shape model data in the full shape space, i.e. the joint-space of the correspondences. The analysis in this case, however, is also difficult because traditional statistical metrics no longer apply [9]. At issue is the fact that the convergence of any estimator in very high dimensional space is prohibitively slow with the respect to the number of samples. A common solution is to employ dimensionality reduction by choosing a subspace of the  $3m$ -dimensional shape in which to project the data for traditional multivariate analysis. Principal component analysis (PCA) is often an attractive choice because the basis vectors are orthogonal and determined solely from the data. With PCA, we can find no more than  $n - 1$  modes that have non-zero variance, meaning that the problem is reduced to  $d < n$  without loss of information. Other basis functions such as wavelets [12] have also been used for dimensionality reduction, with the difference being that they impose an a-priori choice of how the space of the model should be decomposed.



**Fig. 2.** Hotelling  $T^2$  test results with increasing numbers of PCA modes for 3 brain structures from a pediatric autism study.

In a suitably low-dimensional shape space, such as basis vectors from PCA, we can apply traditional statistical methods such as nonparametric Hotelling  $T^2$  testing. There are two challenges for dimensionality reduction in shape analysis, however. First, is the the factor analysis problem of *how many* basis vectors to choose, which can be hard to resolve when the choice of different numbers of factors leads to different statistical results. The second challenge is how to visualize group differences, which is important for researchers in order to relate the outcomes of statistical tests to scientific hypotheses regarding the growth and form of anatomy. The remainder of this section addresses the problem of choosing the number of bases, and then proposes two visualization strategies for understanding the group differences.

To illustrate the number-of-bases problem, Figure 2 shows the  $p$ -value results of Hotelling  $T^2$  permutation tests using increasing numbers of PCA modes on three brain structures from the pediatric autism study referenced above. Several trends can be observed that pose a challenge for the analysis. First, is the trend at higher numbers of modes towards increasing  $p$ -values, which is due to the cumulative effects of noise in the these lower-variance modes of the PCA. The second trend is that the  $p$ -value curves do not smoothly decrease to a minimum value, but rather, tend to bounce around with lower numbers of higher-variance modes. The challenge is to choose as many modes as possible that contain meaningful variation, i.e. variation that is distinguishable from noise, with the caveat that too few modes may result in the loss of information that is useful for group discrimination.

Many methodologies have been proposed to address the number-of-bases problem, and good reviews such as [16], are available. Methodologies range from simple conventions, such as choosing only PCA modes that account for at least 5% of the total variance, to more systematic approaches that attempt to model the problem based on its dimensionality and sample sizes. Notable among these latter approaches, *parallel analysis* is commonly recommended as the best method for determining modes with variance that is distinguishable from noise[17], and is described in more detail in Section 4.


In contrast to the point-wise statistical method illustrated in Figure 1, a significant drawback of hypothesis testing in a PCA subspace is that the group differences in this space are not necessarily easy to visualize and interpret from an anatomical perspective. The hypothesis test poses the question of whether there are significant group differences. The next logical question of interest to researchers is *what* are the group differences. One possible approach to gain insight into this question is to transform the group differences measured in the PCA space back into the full shape space, where they can be visualized on the mean shape surfaces. Implicit in the Hotelling  $T^2$  metric, for example, is a linear discriminant which indicates the direction in the PCA space along which the maximum group difference is observed. This discriminant vector can be rotated back into the full shape space for visualization. Another standard approach for visualizing group differences is a comparison of the differences in the mean shapes. Thin-plate spline analysis is commonly used in morphometric studies, for example, to visualize the deformations between shapes parameterized with landmarks (sparse correspondences) [18], and a similar approach can be applied in the context of high-dimensional point-based shape models. We discuss these visualization strategies further in the next section, along with the development of the dimensionality reduction and hypothesis testing.

## 4 Methodology

For a correspondence point shape model in 3D, we have a  $3m \times n$  shape matrix  $\mathbf{P}$ , where columns of  $\mathbf{P}$  are the shape vectors of correspondence points for the set of all samples. For dimensionality reduction, we first project  $\mathbf{P}$  onto the basis vectors determined by PCA analysis, i.e.  $\tilde{\mathbf{P}} = \mathbf{E}\mathbf{P}$ , where columns of  $\mathbf{E}$  are the eigenvectors of the covariance matrix of correspondences, in decreasing order of the magnitude of their eigenvalues.

Following projection into the PCA space, we perform parallel analysis to choose the number of PCA bases for hypothesis testing. In the context of principal components analysis (PCA) of  $n$ , vector-valued data samples of dimensionality  $3m$ , the goal of parallel analysis is to identify the subset of the components that contain variation distinguishable from the expected variation resulting from noise, where noise is modeled by an isotropic, multivariate unit Gaussian, i.e. a random  $3m$ -vector  $X \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . To make such a distinction, we need an estimator  $\mathbf{E}$  for the expected values of the variances in the *ordered* PCA modes of random samplings from  $X$ , given the fixed sample size  $n$ . Due to the ordering problem, there is no obvious closed-form expression for  $E$ , so it is estimated using Monte Carlo simulation. Many random sample sets of size  $n$  are independently drawn from  $X$ , followed by PCA on each sample set and ordering of the associated eigenvalues. The ordered eigenvalues are then averaged to produce an estimate of the Gaussian noise variance profile across modes. Note that the eigenvalues in this case quantify variance, and the percentage of total variance for a PCA mode is equivalent to the ratio of its eigenvalue to the sum of all eigenvalues.

In order to determine the number of modes to use from parallel analysis, the percent-total-variance profiles from the Monte Carlo simulation and the PCA of the true data are compared, and only modes where the percent-total-variance in the true data is greater than the simulation data are retained. Figure 3, for example, is a scree plot of the percent-variances associated with shape data of a putamen brain structure ( $n = 25, 3m = 3000$ ) [19] (solid line) and the variances from the Monte Carlo noise variance simulation (dashed line). The two lines intersect just before mode 6, and so we would consider only modes 1-5 in the analysis.



**Fig. 3.** Parallel analysis for the putamen data

Once we have chosen the set of  $k$  PCA basis vectors by parallel analysis, we project the correspondences into that subspace. Hypothesis testing for group differences can now be done using a nonparametric, permutation test with the Hotelling  $T^2$  metric, with the null hypothesis that the two groups are drawn from the same distribution.

The Hotelling  $T^2$ , two-sample metric is given by

$$T^2 = \frac{(n_a n_b)(n_a + n_b - 2)}{n_a + n_b} (\mu_{\mathbf{a}} - \mu_{\mathbf{b}})^T \mathbf{w}, \quad (1)$$

$$\mathbf{w} = (\Sigma_a + \Sigma_b)^{-1} (\mu_{\mathbf{a}} - \mu_{\mathbf{b}})$$

where  $\mu_{\mathbf{a}}$  and  $\mu_{\mathbf{b}}$  are the means,  $\Sigma_a$  and  $\Sigma_b$  are the covariances, and  $n_a$  and  $n_b$  are the sample sizes of the two groups, respectively. Note that the vector  $\mathbf{w}$  is also Fisher's linear discriminant, which is well known to be the line along which the between-group variance is maximized with respect to within-group variance[20]. The Hotelling  $T^2$  metric is therefore a scaled projection of the group difference onto the discriminant line. We therefore propose to visualize the morphological differences that are driving the statistical test results by transforming  $\mathbf{w}$  back from PCA space into the full-dimensional shape space, i.e.  $\hat{\mathbf{w}} = \mathbf{E}^{-1} \mathbf{w}$ , where  $\hat{\mathbf{w}}$






Fig. 4. LDA visualization for the *HoxD11* phenotype study

is  $\mathbf{w}$  padded to  $n$ -dimensions with  $n - k$  zeros. The components of the  $3m$  vector  $\hat{\mathbf{w}}$  can then be mapped onto a visualization of the  $n$  mean correspondence point positions, such as a surface mesh. This resulting linear discriminant analysis (LDA) visualization indicates group differences in shape that the test metric identified as the most effective for discrimination.


To visualize deformations between the group mean shapes, we can compute metrics on the displacement field describing the mapping from points  $\mathbf{x}$  on one group mean to corresponding points  $\mathbf{x}'$  on the another. Using the set of correspondence points, a smooth transformation  $T(\mathbf{x}) = x'$ , can be computed using a thin-plate spline interpolation. Details for computing  $T(\mathbf{x})$  are omitted here for brevity, and good descriptions can be found elsewhere (e.g. [18, 21]). We propose to visualize strain, a measure on the Jacobian  $J$  of the deformation field  $x - T(x)$  that describes the local stretching and compression caused by the deformation. The Lagrangian strain tensor is a symmetric, second order tensor given by

$$E = \frac{1}{2}(J + J^T + J^T J). \quad (2)$$

The eigenvectors of  $E$  indicate the principal directions of strain, and the eigenvalues of  $E$  indicate the unit elongations in those directions. An effective visualization for the strain tensor is an ellipsoid with principal axes given by the eigenvalues and oriented along the eigenvector directions.

## 5 Results and Discussion

This section presents two shape analysis experiments that illustrate the effectiveness of the proposed hypothesis testing and visualization methodology for phenotyping studies of gene function, and for the analysis of clinical neurological datasets. As an example of the application of our method to phenotyping, we analyzed differences between wild-type mice and a population of mice deficient in the gene *Hoxd11*. The data for this study are segmentations of the second phalange of the first digit of the right forepaw, derived from micro-CT images of normal and mutant strains acquired by the authors [14]. We computed a combined shape model of the normal ( $n=20$ ) and mutant population ( $n=20$ ) using the particle system method from [2], and applied the hypothesis testing framework from Section 4.




**Fig. 5.** LDA visualization of the right putamen from an autism study

For the mouse data, we have a clear biological hypothesis as to the group differences between mutant and normal mouse strains. Through a series of gene-targeting experiments, Boulet, Davis and Capecchi have shown that *Hoxd11* is important for the normal development and patterning of the appendicular skeleton, and suggest group differences between wild-type and mutant mice in the length of several of the digits of the forepaw [1, 22]. The proposed hypothesis test method applied to the phalange model indicates highly significant group differences ( $p \ll .01$ ), with the parallel analysis choosing the first two PCA modes for the hypothesis test. Figure 4 depicts the length in the surface normal direction of each of the point-wise vector components of the discriminant  $\hat{\mathbf{w}}$  on a surface reconstruction of the mean wild-type population correspondence points. In the figure, the direction of the arrows are from the wild-type to the mutant population. The linear discriminant visualization reveals two clear morphological differences: a reduction in bone length and an increase in bone thickness in the *Hoxd11*-deficient population. This analysis has, therefore, quantified and statistically validated one of the major conclusions from empirical studies of the *Hoxd11* gene, as well as revealing a new significant phenotypic effect.

As a second example, we present the analysis of the three brain structure model described from the pediatric autism study described in Sect.3. The  $p$ -value results for the amygdala, putamen, and hippocampus models, respectively, are 0.003, 0.046, and 0.100, with the number of PCA modes chosen as 5, 6, and 5. Of particular interest is the fact that the result for the putamen indicates group differences at the 5% significant level, which is in contrast to the point-wise hypothesis testing shown in Fig. 1 that indicates no significance. This difference illustrates the increased statistical power of the proposed testing method, which avoids the multiple-comparisons problem. The discriminant vector is visualized for the putamen in Fig.5 for the mean normal population correspondence points, with arrows indicating the direction from patient to the normal control populations. The visualization indicates a shortening of the anterior and posterior regions of the putamen, with a thickening in the medial region.

Figure 6 is a visualization of the strain tensors computed from the deformation from the mean patient shape to the mean normal control shape for



**Fig. 6.** Strain tensors for the right putamen from an autism study. Tensor scale is exaggerated for visualization.

the putamen data. The strain tensors were computed from a thin-plate spline deformation, as described in the previous section. In the figure, the three principal axes of each ellipsoid are scaled by the three principal eigenvalues of the strain tensor at each correspondence position, and oriented according to their corresponding eigenvectors. Ellipsoids and the surrounding surface are colored according to the value of the first principal eigenvector (the longest axis), with yellow indicating negative (compression) and blue indicating positive (stretching). While a clinical interpretation of this result is beyond the scope of this paper, this visualization clearly offers a more detailed insight into *how* groups differ than, for example, the LDA visualization. Note, however, that in this case we have no indication of the statistical significance of these differences.

In summary, the proposed hypothesis testing and visualization methodologies offer an intuitive approach to analysis of anatomical shape from images for biological and clinical research, and avoid the problems inherent in HDLSS statistics. We have illustrated the effectiveness of the framework by statistically validating a biological hypothesis regarding *Hoxd11* gene function that was previously only based on empirical evidence, and we have shown how the method can be useful in exploring new patterns in shape from clinical data which have not previously been observed, or that are not observable with lower-powered statistical methods.

### Acknowledgments

This work was funded by the Center for Integrative Biomedical Computing, National Institutes of Health (NIH) NCRR Project 2-P41-RR12553-07. This work is also part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149.

## References

1. Davis, A., Capecchi, M.: Axial homeosis and appendicular skeleton defects in mice with targeted disruption of *hoxd-11*. *Development* **120** (1995) 2187–2198
2. Cates, J., Fletcher, P.T., Styner, M., Shenton, M., Whitaker, R.: Shape modeling and analysis with entropy-based particle systems. In: *Information Processing in Medical Imaging (IPMI 2007)*, LNCS 4584. (2007) 333–345
3. Gorczowski, K., Styner, M., Jeong, J., Marron, J., Piven, J., Hazlett, H., Pizer, S., Gerig, G.: Statistical shape analysis of multi-object complexes. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE (2007)* 1–8
4. Styner, M., Lieberman, J.A., Pantazis, D., Gerig, G.: Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis* (2004)
5. Pizer, S.M., Jeong, J.Y., Lu, C., Muller, K.E., Joshi, S.C.: Estimating the statistics of multi-object anatomic geometry using inter-object relationships. In: *Deep Structure, Singularities, and Computer Vision*. Volume 3753 of LNCS. (2005) 60–71
6. Klingenberg, C.P.: Morphometrics and the role of the phenotype in studies of the evolution of developmental mechanisms. *Gene* **287** (2002) 3–10
7. Chen, X.J., Kovacevic, N., Lobaugh, N.J., Sled, J.G., Henkelman, R.M., Henderson, J.T.: Neuroanatomical differences between mouse strains as shown by high-resolution 3d mri. *NeuroImage* **29** (2005) 99–105
8. Davies, R.H., Twining, C.J., Allen, P.D., Cootes, T.F., Taylor, C.J.: Shape discrimination in the hippocampus using an mdl model. In: *IPMI*. (2003) 38–50
9. Ahn, J., Marron, J.S., Muller, K.M., Chi, Y.: The high-dimension, low-sample-size geometric representation holds under mild conditions. *Biometrika* **94**(3) (2007) 760–766
10. Terriberry, T., Joshi, S., Gerig, G.: Hypothesis testing with nonlinear shape models. In: *IPMI'05*. (2005) 15–26
11. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J., Shenton, M., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. *The Insight Journal* (2006)
12. Nain, D., Niethammer, M., Levitt, J., Shenton, M., Gerig, G., Bobick, A., Tannenbaum, A.: Statistical shape analysis of brain structures using spherical wavelets. In: *IEEE Symposium on Biomedical Imaging ISBI*. (2007) in print
13. Golland, P., Grimson, W., Shenton, M., Kikinis, R.: Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis* **9** (2005) 69–86
14. Cates, J., Fletcher, P.T., Warnock, Z., Whitaker, R.: A shape analysis framework for small animal phenotyping with application to mice with a targeted disruption of *hoxd11*. In: *Proc. 5th IEEE International Symposium on Biomedical Imaging (ISBI '08)*. (2008) 512–516
15. Hazlett, H., Poe, M., Gerig, G., Smith, R., Provenzale, J., Ross, A., Gilmore, J., Piven, J.: Magnetic resonance imaging and head circumference study of brain size in autism: Birth through age 2 years. *Arch Gen Psych* **62** (2005) 1366–1376
16. Fabrigar, L.R., Wegener, D.T., MacCallum, R.C., Strahan, E.J.: Evaluating the use of exploratory factor analysis in psychological research. *Psychological Methods* **4** (1999) 272–299
17. Glorfeld, L.W.: An improvement on horn's parallel analysis methodology for selecting the correct number of factors to retain. *Educational and Psychological Measurement* **55** (1995) 377–393
18. Bookstein, F.: Principal warps: Thin plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(6) (1989)

19. Cates, J., Fletcher, P., Styner, M., Hazlett, H., Whitaker, R.: Particle-based shape analysis of multi-object complexes. In: Proceedings of the 11th International Conference on Medical Image Computing and Computer Assisted Intervention, MIC-CAI (2008) to appear
20. Timm, N.H.: Applied Multivariate Analysis. Springer-Verlag (2002)
21. Whitbeck, M., Guo, H.: Multiple landmark warping using thin-plate splines. In: IPCV. (2006) 256–263
22. Boulet, A.M., Capecchi, M.R.: Duplication of the *hoxd11* gene causes alterations in the axial and appendicular skeleton of the mouse. *Developmental Biology* **249** (2002) 96–107