

Particle-Based Shape Analysis of Multi-object Complexes

Joshua Cates¹, P. Thomas Fletcher¹, Martin Styner², Heather Cody Hazlett²,
and Ross Whitaker¹

¹ Scientific Computing and Imaging Institute, Univ. Utah, Salt Lake City, UT, USA

² Depts. Psychiatry & Computer Science, Univ. North Carolina,
Chapel Hill, NC, USA

Abstract. This paper presents a new method for optimizing surface point correspondences for shape modeling of multiobject anatomy, or shape *complexes*. The proposed method is novel in that it optimizes correspondence positions in the full, joint shape space of the object complex. Researchers have previously only considered the correspondence problem separately for each structure, thus ignoring the interstructural shape correlations that are increasingly of interest in many clinical contexts, such as the study of the effects of disease on groups of neuroanatomical structures. The proposed method uses a nonparametric, dynamic particle system to simultaneously sample object surfaces and optimize correspondence point positions. This paper also suggests a principled approach to hypothesis testing using the Hotelling T^2 test in the PCA space of the correspondence model, with a simulation-based choice of the number of PCA modes. We also consider statistical analysis of object poses. The modeling and analysis methods are illustrated on brain structure complexes from an ongoing clinical study of pediatric autism.

1 Introduction

Statistical shape modeling is an increasingly important tool for the analysis of anatomical objects derived from 3D medical images. In many areas of clinical psychiatric and neurological research, the joint analysis of *complexes* of multiple anatomical structures is of increasing interest because certain spectrum disorders, such as autism, are thought to represent a confluence of several underlying abnormalities, impacting the *relationships* between brain regions [1]. Shape models of anatomical complexes are also important tools for geneticists and developmental biologists, who rely on quantifications of phenotype in gene targeting studies (e.g., [2]).

We define a *multiobject complex* as a set of solid shapes, each representing a single, connected biological structure, assembled into a scene within a common coordinate frame. A multiobject complex contains shape, pose, scale, and positional information for each structure. Some examples include the segmentations of multiple brain structures from a single MRI of a patient and sets of bones segmented from a CT scan. Point-based models, which we consider in this work,

represent shape by sampling each shape surface in a consistently ordered fashion in order to define homologous object surface points called *correspondences*. The set of correspondences for a population is then used for statistical analysis, including hypothesis testing for group differences[3,4].

The choice of correspondence positions is a critical step for point-based modeling. State-of-the-art methods typically rely on parameterized surface models that assume a spherical or toroidal topology, and are thus not suitable for multiobject complexes, which are nonmanifold, and consist of disconnected sets of discrete surfaces. Some parametric methods have been applied to shape complexes by finding correspondences for each structure independently, and then treating those correspondences as the *marginal* distributions of the multiobject complex [5]. This particular approach, however, is not consistent with standard methods in statistics, which generally seek to use the simplest model that explains the observed data. For point-based modeling, this means that ideally one should seek a compact distribution for the correspondences in the full, joint shape space. Several methods for optimizing correspondence positions for collections of single objects have been proposed (e.g.,[6,4]), but a *joint optimization* for a multiobject model has yet to be demonstrated and analyzed.

Optimization in the full, joint shape space of complexes is important for several reasons. If variabilities between individual shapes in a complex are correlated, for example, the marginal variabilities can appear small, and might not otherwise be preserved. By modeling these correlations among variabilities, optimization in the joint space may also produce more compact distributions for correspondences. The specific choice of optimization methods is also an important factor. Parametric approaches, for example, are potentially limited for multiobject modeling because they typically rely on individual *anchor shapes* to regularize the optimization process, and thus would restrict the degree to which parameterizations of different objects in the ensemble could interact.

In this paper we propose a novel, nonparametric approach to multiobject shape modeling that is an extension of the entropy-based particle system method given for single objects in [4]. The proposed method optimizes correspondence positions in the full, joint shape space of the object complex. Because statistical analysis of shape models is in itself a difficult problem due to the very high dimensionality of the shape space and the relatively low numbers of samples, we also present a systematic approach to shape analysis using Hotelling T^2 tests in the PCA space of the correspondences, with a simulation-based approach to the dimensionality reduction. Additionally, we show how point-based models can be used to analyze group differences in object position, pose, and scale. The proposed modeling framework is applied to a proof-of-concept analysis of brain structure complexes from a study of pediatric autism. Our analysis shows group differences in shape between normal and patient populations that have not been seen previously in this data, and show how an optimized joint model yields results with a higher statistical power than a model constructed by simply optimizing the marginal distributions.

This paper builds on previous work on statistical shape parameterization and analysis that has mostly focused on single objects. For instance, Styner et al. [3] use geometric considerations on individual shapes, without ensemble statistics. Finding correspondence positions that minimize information content (log-determinant of the covariance matrix) across an ensemble of simple, closed, 2D objects was first proposed by Kotcheff and Taylor[7], and extended by Davies et al. [6] using a minimum description length (MDL) minimization to compute 3D surface parameterizations. In previous work, Cates et al. [4] propose a non-parametric, particle-based method, which minimizes the entropy of the resulting description. This strategy includes an explicit penalty for geometric regularity, and thus does not require a regularization based on a-priori anchor shapes, as is typically done with MDL. Previous results on statistical analysis of multiobject complexes have been shown using sampled medial mesh (m-rep) representations [8,5], but do not include an ensemble-based parameterization. The contributions of this paper are, therefore, the generalization of an ensemble-based, statistical correspondence algorithm to shape complexes, and a systematic statistical analyses of multiobject shape, size, and pose, as well as a demonstration of hypothesis testing with this framework in a compelling clinical application.

2 Methodology

This section gives a brief overview of the particle-system correspondence optimization method for single object surfaces given in [4], and then describes its extension to multiobject complexes of surfaces. We also present our approach to the statistical analysis of the shape, scale, and pose in the resulting models.

Correspondence Optimization for Single Objects. We define a surface as a smooth, closed manifold of codimension one, which is a subset of $\mathbb{R}^d$ (e.g., $d = 3$ for volumes). We sample a surface $\mathcal{S} \subset \mathbb{R}^d$ using a discrete set of N points that are considered random variables $Z = (X_1, X_2, \dots, X_N)$ drawn from a probability density function (PDF), $p(X)$. We denote a realization of this PDF with lower case, and thus we have $z = (x_1, x_2, \dots, x_N)$, where $z \in \mathcal{S}^N$. The probability of a realization x is $p(X = x)$, which we denote simply as $p(x)$.

The amount of information contained in such a random sampling is, in the limit, the differential entropy of the PDF, which is

$$H[X] = - \int_{\mathcal{S}} p(x) \log p(x) dx = -E\{\log p(X)\}, \quad (1)$$

where $E\{\cdot\}$ is the expectation. Approximating the expectation by the sample mean, we have $H[X] \approx -\frac{1}{N-1} \sum_i \log p(x_i)$. To estimate $p(x_i)$, we use a non-parametric Parzen windowing estimation from the particle positions, modified to adaptively oversample in regions of higher curvature by the inclusion of a scaling term that is proportional to local curvature magnitude. This results in a set of points on the surface that repel each other with Gaussian-weighted forces.

Now consider an ensemble $\mathcal{E}$, which is a collection of M surfaces, each with their own set of particles, i.e., $\mathcal{E} = z^1, \dots, z^M$. The ordering of the particles on

each shape implies a correspondence among shapes, and thus we have a matrix of particle positions $P = x_j^k$, with particle positions along the rows and shapes across the columns. We model $z^k \in \mathbb{R}^{Nd}$ as an instance of a random variable Z , and minimize a combined ensemble and shape cost function

$$Q = H(Z) - \sum_k H(P^k), \quad (2)$$

which favors a compact ensemble representation balanced against a uniform distribution of particles on each surface. Given the low number of samples relative to the dimensionality of the space, we use a parametric approach for density estimation in the space of shapes, modeling $p(Z)$ parametrically as a Gaussian with covariance Σ . The entropy is then given by

$$H(Z) \approx \frac{1}{2} \log |\Sigma| = \frac{1}{2} \sum_{j=1}^{Nd} \log \lambda_j, \quad (3)$$

where $\lambda_1, \dots, \lambda_{Nd}$ are the eigenvalues of Σ .

The cost function Q is minimized using a gradient descent strategy to manipulate particle positions. The negative gradient $-\partial H(Z)/\partial P$ gives a vector of updates for the entire system, which is recomputed once per iteration of the entire particle system. This gradient term is added to the individual shape-based updates $\partial H(P^k)/\partial P^k$ to give the update for each particle. The surface constraint is specified by the zero set of a scalar function $F(x)$, and maintained by projecting the gradient of the cost function onto the tangent plane of the surface, followed by iterative reprojection of the particle onto the nearest root of F by the method of Newton-Raphson. The optimization function balances entropy of individual surface samplings with the entropy of the shape model, maximizing the former for geometric accuracy (a good sampling) and minimizing the latter to produce a compact model.

Correspondences Across MultiObject Complexes. The particle-based correspond method outlined above *can* be directly applied to multiobject complexes by treating all of the objects in the complex as one. However, if the objects themselves have distinct identities (i.e., object-level correspondence is known a priori), we can assign each particle to a specific object, decouple the spatial interactions between particles on different shapes, and constrain each particle to its associated object, thereby ensuring that each correspondence stays on a particular anatomical structure. The shape-space statistics remain coupled, however, and the covariance Σ (Eqn. 3) includes all particle positions across the entire complex, so that optimization takes place on the joint, multiobject model.

Any set of implicitly defined surfaces is appropriate as input to this framework. In the case of binary segmentations, the input is a set of M segmentations of N -object complexes, which contains $N \times M$ distinct, volumetric label masks. A binary mask contains an implicit shape surface at the interface of the labeled pixels and the background, but contains aliasing artifacts that must first be removed. We have found that the r -tightening algorithm given by Williams

et al. [9] is effective in removing these artifacts without compromising the precision of the segmentation. Typically we follow the antialiasing step with a very slight Gaussian blurring to remove the high-frequency artifacts that can occur as a result of numerical approximations. We initialize the optimization using the splitting strategy described in [4], starting with a single particle on each object. For the hypothesis testing that will follow, it is important that correspondences be computed without knowledge of the group classification of the shapes. We therefore compute shape models for the shape complexes that include both the control and study data, which we will refer to as *combined* models.

Statistical Analysis for Correspondence-Based Shape Models. The goal of statistical analysis in the context of this paper is to quantify shape differences of the targeted anatomy between control and study populations, and to perform hypothesis testing for statistical significance of those differences. An important consideration is to quantify differences in a way that accounts for desirable invariances, which is typically done by explicitly normalizing for size and pose variation. In a multiobject setting, we must decide the level of granularity at which to align shapes in order to analyze pose and scale. Previous work [5,10] employs a hierarchical strategy, with a *global* coordinate frame for the entire complex, followed by a set of *local* coordinate frames for each object. The global frame is established by alignment of the entire complex, resulting in M sets of global pose parameters. Remaining pose discrepancies among the individual objects constitute the local coordinate frames, and are determined by alignment of each ensemble of individual shapes, to give a set of N local pose parameters for each of the M complexes.

For the analysis of object pose and scale, we first align shapes with respect to their centers of mass and the orientation of their first principal eigenvectors. We then align shapes with respect to rotation, translation, and scaling using a Procrustes algorithm, which is run at regular intervals between the correspondence optimization updates (see [4]). Hypothesis testing on object scale can now be done using standard, two-tailed parametric t -test, and group differences in relative position analyzed with a parametric Hotelling T^2 test. For relative pose, we use a general nonparametric hypothesis test for metric spaces [11], which relies only on pairwise distances between the data, and we use geodesic distances in the rotation group.

We refer to the differences in correspondences that remain in the population after pose and scale alignment as *shape*. The high dimensionality of the shape space, coupled with the relatively low sample size of our data, precludes the use of traditional low-dimensional statistical metrics directly in the full shape space. Instead, we use a standard, data-driven approach to dimensionality reduction and project the correspondences into a lower dimensional space determined by choosing a number of basis vectors from principal component analysis (PCA). Ideally, we would like to choose only PCA modes that account for variance that cannot be explained by random noise. *Parallel analysis* is commonly recommended for this purpose [12]. Parallel analysis works by comparing the percent variances of each of the PCA modes with the average percent variances obtained




Fig. 1. Mean brain structure complexes with average pose. Colormap indicates the magnitude and direction of the linear discriminant.

via PCA of Monte Carlo simulations of samplings from isotropic, multivariate, unit Gaussian distributions. We choose only modes with greater variance than the simulated modes for the dimensionality reduction, and use a standard, parametric Hotelling T^2 test to test for group differences, with the null hypothesis that the two groups are drawn from the same distribution.

To visualize group differences that are driving the statistical result, we compute the linear discriminant vector implicit in the Hotelling T^2 statistic, which is the line along which the between-class variance is maximized with respect to the within-class variance. This line is also known as Fisher’s linear discriminant, and is given by

$$\mathbf{w} = (\Sigma_{\mathbf{a}} + \Sigma_{\mathbf{b}})^{-1}(\mu_{\mathbf{a}} - \mu_{\mathbf{b}}), \quad (4)$$

where μ are the group means for groups a and b , and Σ are their covariance matrices. Vector $\mathbf{w}$ can be rotated back from PCA space into the full dimensional shape space, and then mapped onto the mean group shape visualizations to give an indication of the significant morphological differences between groups.

3 Results and Discussion

For the experimental analysis, we used multiobject segmentation data taken from an ongoing longitudinal pediatric autism study[13], which includes MRI brain scans of autistic subjects and typically-developing controls at time points of 2 and 4 years of age. The data consists of binary segmentations of 10 subcortical brain structures (see Fig 1), which were done by trained experts using semi-automated procedures[5]. For this analysis, we had 10, 2-year old male controls available, and chose 15 matched autism subjects for comparison. Multiobject correspondences were computed from the segmentations as described in Section 2 to produce a combined model of the groups. We sampled each complex of segmentations with 10,240 correspondence points, using 1024 particles per structure. For comparison, we also computed point-correspondence models for each of the 10 structures

separately and concatenated their correspondences together to form a *marginally optimized* joint model.

Scale and pose were calculated for structures in the complex as described in Sect. 2. Hypothesis testing indicates significant group differences in scale only for the right and left amygdala, with p -values of 0.0017 and 0.018, respectively. Hotelling T^2 tests on mean structure positions do not suggest any differences between the groups, with $p > 0.05$ for all structures. Similarly, group difference in pose are not indicated by the statistical pose analysis, and we obtained $p > 0.05$ for all structures. This result for pose is also consistent with results given in [5] on this data.

The hypothesis test method outlined in Sect. 2 gives a highly significant p -value of 0.0087, with 8 PCA modes chosen by parallel analysis. This result is the first evidence shown for this data for group differences in shape alone. Gorczowski [5] reports group differences when scale is included with shape, but reports insignificant shape discrimination between groups when the shapes are normalized to the same size. Parallel analysis of the marginally-optimized model indicates that the first 6 modes should be used, which gives a p -value of 0.0480. While the test still suggests group differences at the 5% significance level, we note that the result is an order of magnitude lower in statistical power.

Previous work in shape analysis has suggested hypothesis testing on individual correspondence point positions instead of working in the full dimensional shape space [14]. For comparison, we ran statistical tests at every correspondence point location, using an open-source implementation of the nonparametric Hotelling T^2 method described in [14], with 20,000 permutations among groups and an FDR bound set to 5%. Uncorrected p -values show widespread differences, but no significance remains after FDR correction. This result is in contrast to the global shape result, and illustrates one of the difficulties with point-based shape analysis at a local feature scale: the unavoidable reduction in statistical power due to the necessary correction for multiple comparisons.

Figure 1 shows the mean shape surfaces for the normal and autistic groups, as reconstructed from the Euclidean averages of the correspondence points. Each structure is displayed in its mean orientation, position, and scale in the global coordinate frame. We computed the average orientation for each structure using methods for averaging in curved spaces [15]. We used the arithmetic mean of position and the geometric mean of scale. Mean pose differences between the two groups appear small, as might be expected from their statistical analyses.

To illustrate the morphological differences that are driving the global shape result, we visualize the linear discriminant vector $\mathbf{w}$ (Equation 4) in Fig. 1, as described in the previous section. The length in the surface normal direction of each of the point-wise discriminant vector components for the autism data is given by the colormap. Yellow indicates a negative (inward) direction, and blue




Fig. 2. Detail from Fig. 1

indicates a positive (outward) direction. The right amygdala for the normals is shown in Fig. 2 as a more detailed example, with the vectors depicted as arrows. Note the clear trend towards a shortening of the anterior end of the amygdala in the autistic versus the normal population.

In summary, our results suggest that the proposed modeling and analysis framework can effectively model group differences in the autism data that have not been seen with other methods, and more powerful statistical results are obtained by optimization in the joint space than by optimization in the marginal space. This analysis, however, is only a proof-of-concept example of how the particle method may be applied to multiobject data. A more rigorous study, which remains for future work, is required in order draw clinical conclusions.

Acknowledgments

This work is supported by the Center for Integrative Biomedical Computing, (NIH 2-P41-RR12553-07), the National Alliance for Medical Image Computing (NIH U54-EB005149), NIH grant HD03110, and the Autism Speaks Foundation.

References

1. Chultz, R., Robins, D.: Functional neuroimaging studies of autism spectrum disorders. In: Handbook of autism and pervasive developmental disorders, 3rd edn., pp. 515–533. J. Wiley and Sons, Chichester (2005)
2. Davis, A., Capecchi, M.: Axial homeosis and appendicular skeleton defects in mice with targeted disruption of *hoxd-11*. *Development* 120, 2187–2198 (1995)
3. Styner, M., Lieberman, J.A., Pantazis, D., Gerig, G.: Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis* (2004)
4. Cates, J., Fletcher, P.T., Styner, M., Shenton, M., Whitaker, R.: Shape modeling and analysis with entropy-based particle systems. In: Karssemeijer, N., Lelieveldt, B. (eds.) *IPMI 2007*. LNCS, vol. 4584, pp. 333–345. Springer, Heidelberg (2007)
5. 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*, pp. 1–8. IEEE, Los Alamitos (2007)
6. Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J.: 3D statistical shape models using direct optimisation of description length. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002*. LNCS, vol. 2352, pp. 3–20. Springer, Heidelberg (2002)
7. Kotcheff, A., Taylor, C.: Automatic Construction of Eigenshape Models by Direct Optimization. *Medical Image Analysis* 2, 303–314 (1998)
8. 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: Fogh Olsen, O., Florack, L.M.J., Kuijper, A. (eds.) *DSSCV 2005*. LNCS, vol. 3753, pp. 60–71. Springer, Heidelberg (2005)
9. Williams, J., Rossignac, J.: Tightening: curvature-limiting morphological simplification. In: *Proc. 9th ACM Symp. on Solid and Physical Modeling*, pp. 107–112 (2005)

10. Gerig, G., Joshi, S., Fletcher, T., Gorczowski, K., Xu, S., Pizer, S., Styner, M.: Statistics of populations of images and its embedded objects: Driving applications in neuroimaging. In: IEEE Symp. on Biomed. Imaging ISBI, pp. 1120–1123 (2006)
11. Hall, P., Tajvidi, N.: Permutation tests for equality of distributions in high-dimensional settings. *Biometrika* 89(2), 359–374 (2002)
12. 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, 377–393 (1995)
13. 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, 1366–1376 (2005)
14. 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)
15. Fletcher, P., Lu, C., Pizer, S., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Med. Imaging* 23(8), 995–1005 (2004)