Particle Based Shape Regression of Open Surfaces with Applications to Developmental Neuroimaging

Manasi Datar, Joshua Cates, P. Thomas Fletcher, Sylvain Gouttard, Guido Gerig, and Ross Whitaker

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

Abstract. Shape regression promises to be an important tool to study the relationship between anatomy and underlying clinical or biological parameters, such as age. In this paper we propose a new method to building shape models that incorporates regression analysis in the process of optimizing correspondences on a set of open surfaces. The statistical significance of the dependence is evaluated using permutation tests designed to estimate the likelihood of achieving the observed statistics under numerous rearrangements of the shape parameters with respect to the explanatory variable. We demonstrate the method on synthetic data and provide a new results on clinical MRI data related to early development of the human head.

1 Introduction

Technologies for shape representation and statistical shape analysis are important for several problems in medical imaging including image segmentation, quantitative analysis of anatomy, and group comparisons. A widely used approach is to evaluating shapes is assign correspondences or landmarks to shapes (curves, or surfaces) and to compare the positions or configurations of these landmarks. This approach has benefitted in recent years from methods for the automatic placement of landmarks in a way that captures the statistical properties of an ensemble of images [1,2,3]. Finding correspondences that minimize description length [2] or entropy [1] has been shown to generate shape models that systematically capture the underlying variability of the population and conform, qualitatively, to the underlying anatomy. This paper extends the method of Cates et al. [1], which uses an variational formulation of ensemble entropy to position dense collections of landmarks, or *particles*.

On the clinical front, quantitative magnetic resonance imaging has significantly advanced our understanding of brain development during childhood and adolescence. Courchesne et al. [4] describe differences in growth patterns in autism compared to controls. However, these studies do not include children below the age of 4 years. Data measured in infants from birth to 4 years are mostly *volumetric measurements*, such as intracranial volume and volumes of

brain lobes and subcortical structures [5]. Whereas this selection of previous work demonstrates very active research towards determining brain growth at early stage of development, there is little data on modelling head and brain growth across a continuum of time and almost no work on the study of how development influences shape.

In developmental analyses, such as paediatric neurodevelopment, *shape regression* gives aggregate models of growth, with variability. Thus shape analysis promises to give not only basic insights into the process of development, but also allow comparisons of individuals against normative models. Of course, precise characterizations of these relationships will require shape models that can tease apart those aspects of shape variability that are explained by the underlying variables and those that are not. Likewise, in order to understand the statistical significance of such relationships we will need a systematic, unbiased method for testing these correlations. These are the topics addressed in this paper.

2   Methodology

This section gives a brief overview of the particle-system correspondence optimization method, which is first described in [1]. The general strategy of this method is to represent correspondences as point sets that are distributed across an ensemble of similar shapes by a gradient descent optimization of an objective function that quantifies the entropy of the system. Our proposed extension to this method incorporates a linear regression model into the correspondence optimization. We also present a new methodology for correspondence optimization on open surfaces where surface boundaries are defined by arbitrary geometric constraints—which is important for studying paediatric head shape.

Correspondence Optimization
We define a surface as a smooth, closed manifold of codimension one, which is a subset of $\Re^d$ (e.g., $d = 3$ for volumes). We sample a surface $\mathcal{S} \subset \Re^d$ using a discrete set of N points that are considered random variables $\mathbf{Z} = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_N)^T, \mathbf{X} \in \Re^d$ drawn from a probability density function (PDF), $p(\mathbf{X})$. We denote a realization of this PDF with lower case, and thus we have $\mathbf{z} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N)^T$, where $\mathbf{z} \in \mathcal{S}^N$. The probability of a realization $\mathbf{x}$ is $p(\mathbf{X} = \mathbf{x})$, which we denote simply as $p(\mathbf{x})$.

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

$$H[\mathbf{X}] = -\int_S p(\mathbf{x}) \log p(\mathbf{x}) dx = -E\{\log p(\mathbf{X})\}, \tag{1}$$

where $E\{\cdot\}$ is the expectation. Approximating the expectation by the sample mean, we have $H[\mathbf{X}] \approx -\frac{1}{Nd} \sum_i \log p(\mathbf{x}_i)$. To estimate $p(\mathbf{x}_i)$, we use a non-parametric Parzen windowing estimation, modified to scale density estimation in proportion to local curvature magnitude. The kernel width σ is chosen adaptively at each $\mathbf{x}_i$ to maximize the likelihood of that position. We refer to the positions $\mathbf{x}$ as *particles*, and a set of particles as a *particle system*.

Now consider an ensemble $\mathcal{E}$, which is a collection of M surfaces, each with their own set of particles, i.e., $\mathcal{E} = \mathbf{z}^1, \ldots, \mathbf{z}^M$. The ordering of the particles on each shape implies a correspondence among shapes, and thus we have a matrix of particle positions $P = \mathbf{x}_j^k$, with particle positions along the rows and shapes across the columns. We model $\mathbf{z}^k \in \Re^{Nd}$ as an instance of a random variable $\mathbf{Z}$, and minimize a combined ensemble and shape cost function

$$Q = H(\mathbf{Z}) - \sum_k H(P^k), \tag{2}$$

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 described in [1] for density estimation in the space of shapes. The entropy cost function Q is minimized using a gradient descent strategy to manipulate particle positions (and, thus, also correspondence positions). The surface constraint is specified by the zero set of a scalar function $F(x)$. This optimization strategy 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.

Any set of implicitly defined surfaces is appropriate as input to this framework. For this paper, we use binary segmentations, which contain an implicit shape surface at the interface of the labeled pixels and the background. To remove aliasing artifacts in these segmentations, we use the r-tightening algorithm given by Williams et al. [6]. Correspondence optimizations are initialized with the splitting strategy described in [1], starting with a single particle on each object. We use a Procrustes algorithm, applied at regular intervals during the optimization, to align shapes with respect to rotation and translation, and to normalize with respect to scale.

Correspondence with Regression Against Explanatory Variables. With the assumption of a Gaussian distribution in the space of shapes, we can introduce a generative statistical model

$$\mathbf{z} = \mu + \epsilon, \epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma) \tag{3}$$

for particle correspondence positions, where μ is the vector of mean correspondences, and ϵ is normally-distributed error. Replacing μ in this model with a function of an explanatory variable t gives the more general, *regression* model

$$\mathbf{z} = f(t) + \hat{\epsilon}, \hat{\epsilon} \sim \mathcal{N}(\mathbf{0}, \hat{\Sigma}). \tag{4}$$

The optimization described in the previous section minimizes the entropy associated with ϵ, which is the difference from the mean. In this paper, we propose to optimize correspondences under the regression model in Eqn. 4 by instead minimizing entropy associated with $\hat{\epsilon}$, the residual from the model. Considering particle correspondence to be a linear function of t, given as $f(t) = \mathbf{a} + \mathbf{b}t$, we

need an estimate of parameters $\mathbf{a}$ and $\mathbf{b}$ to compute $\hat{\epsilon}$. We estimate these with a least-squares fit to the correspondence data,
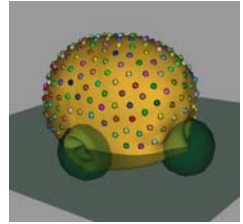
$$\arg\min_{\mathbf{a},\mathbf{b}} E(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \sum_k \left[(\mathbf{a} + \mathbf{b}t_k) - \mathbf{z}_k \right]^T \Sigma^{-1} \left[(\mathbf{a} + \mathbf{b}t_k) - \mathbf{z}_k \right]. \qquad (5)$$

Setting $\frac{\delta E}{\delta \mathbf{a}} = \frac{\delta E}{\delta \mathbf{b}} = 0$ and solving for $\mathbf{a}$ and $\mathbf{b}$, we have $\mathbf{a} = \frac{1}{n}(\sum_k \mathbf{z}_k - \sum_k \mathbf{b}t_k)$, and $\mathbf{b} = \left(\sum_k t_k \mathbf{z}_k - \sum_k \mathbf{z}_k \sum_k t_k \right) / \left(\sum_k t_k^2 - (\sum t_k)^2 \right).$

The proposed regression model optimization algorithm then proceeds as follows. Correspondences are first optimized under the nonregression model (Eqn 3) to minimize the entropy associated with the total error ϵ, and to establish an initial estimate for $\mathbf{a}$ and $\mathbf{b}$. We then follow the same optimization procedure as described in Section. 2, but replace the covariance of the model with the covariance of the underlying residual relative to the generative model. We interleave the two estimation problems, and thus the parameters $\mathbf{a}$ and $\mathbf{b}$ are re-estimated after each iteration of the gradient descent on the particle positions.

Correspondences on Open Surfaces. To compute correspondence positions on a set of *open* surfaces, we propose an extension to the sampling method reviewed in Section. 2. The proposed method is to define the boundary as the intersection of the surface $\mathcal{S}$ with a set of geometric primitives, such as cutting planes and spheres. Our goal is to formulate the interactions with these boundaries so that the positions of these constraints has as little influence as possible on the statistical shape model.

For each geometric primitive, we construct a *virtual* particle distribution that consists of all of the closest points on its surface to the particles $\mathbf{x}_i$ on $\mathcal{S}$. During the gradient descent optimization, particles $\mathbf{x}_i$ interact with the virtual particles, and are therefore effectively repelled from the geometric primitives, and thus from the open surface boundary. The virtual distributions are updated after each iteration as the particles on $\mathcal{S}$ redistribute under the optimization. Because the virtual particles are allowed to factor into the Parzen windowing kernel size estimation, particles $\mathbf{x}_i$ maintain a distance from the

Fig. 1. Particle system with geometric primitives defining the boundary

boundary proportional to their density on the surface $\mathcal{S}$. In this way, features near the boundary may be sampled, but particles are never allowed to lie on the actual boundary itself. One such configuration is shown in Figure. 1

Permutation Test of Significance. Analysis of variance (ANOVA) is the standard parametric test for testing if the explanatory variables have a significant effect in a linear regression. The test statistic used is

$$T = \frac{R^2/(m-1)}{(1-R^2)/(n-m)}, \qquad (6)$$

where R^2 is Pearson's coefficient of regression, generally defined as $R^2 = 1 - \frac{SS_{err}}{SS_{tot}}$, where SS_{err} is the sum-squared residual error, and SS_{tot} represents total variance in the data. In general, R^2 can be related to the unexplained variance of the generated model, and is used to measure the *goodness-of-fit* for the regression model. When the residuals of the linear model are iid Gaussian, the statistic T follows an F distribution with $m - 1$ and $n - m$ degrees of freedom under the null hypothesis.

In this case where the outcome variables are correspondence-optimized shape parameters, the underlying assumptions of the parametric F-test may not hold. Furthermore, optimization with knowledge of the underlying parameter could lead to optimistic estimates of significance, because we are explicitly minimizing the residual. To overcome this, we propose a nonparametric permutation test for significance. Permutation tests for regression work by permuting the values of the explanatory variables. This allows us to compute a distribution of our test statistic under the null hypothesis that the explanatory variable has no relationship to the dependent variable. Given data (z_i, t_i), we generate the kth permuted data set as $(z_i, t_{\pi_k(i)})$, where π_k is a permutation of $1, \ldots, n$. For each permutation we compute a test statistic T_k using (6). Then comparing our unpermuted test statistic T to the distribution of T_k, we can compute the p-value as the percentage of T_k that are greater than T. Notice, that for the case of regression-optimized correspondences, described in Section 2, we perform a the correspondence optimization on *each* permutation separately, and thus *the results of our permutation test are not biased by the correspondence method.*

3   Results and Discussion

This section details experiments designed to illustrate and validate the proposed method. First, we present an experiment with synthetically generated tori to illustrate the applicability of the method and validation based on permutation tests. Next, we present an application to the study of early growth of head shapes extracted from structural MRI data.

To illustrate and validate the proposed methods, we performed two experiments on sets of 40 synthetically generated tori, parameterized by the small radius r and the large radius R. The values for the shape parameters were chosen as independent functions of a uniformly distributed explanatory variable t. The definition of R^2, used to compute the test statistic as explained in Section. 2, is extended to include the two independent variables for this experiment:

$$R^2 = 1 - \frac{(SS_{err})_r + (SS_{err})_R}{(SS_{tot})_r + (SS_{tot})_R}. \tag{7}$$

We examine sets of time-dependent shapes with p-values $\{0.01, 0.1\}$ in order to examine the performance of the system with and without significance. To construct these example data sets, we use the value for the statistic T (look up from the F-distribution) to generate a target R^2. The values of r and R are chosen such that the R^2 of the generated set is approximately equal to the

target R^2 for that experiment. Along with explicit correspondences generated from the standard torus parametrization, we use the correspondence methods from Section. 2, optimization with and without an underlying regression model, to optimize correspondences using 256 particles on each shape. An analysis of the resulting models showed that all three sets of correspondences exhibited two pure modes of variation.

Synthetic Data (Tori). Here we present the results of the statistical analysis of the tori test data using permutation tests consisting of 1000 permutations of the explanatory variable t. For the correspondences we compute the test statistics using the two dominant modes from a PCA on the set of correspondences. The procedure described in Section. 2 is then applied to get the corresponding p-values. Table. 1 shows the results of the two permutation tests for the explicit correspondences, and correspondences generated using the proposed methods. A comparison of the parametric p-value with the p-values obtained by the permutation tests confirms that the proposed methods preserve the relationship between the explanatory variable and the dependent variables. The correspondence-based approaches, particularly with the regression model, show greater significance than the parametric case. This might be an inherent property of the statistic or it could be an artifact due to the limited number of example datasets and the limited number of permutations. Future work will include more datasets, more permutations, and a bootstrapping procedure to analyze variability of the p-values computed by the various methods.

Table 1. Results of permutation tests (1000 permutations)

p-value (theory)	p-value (parametric)	Correspondence Type		
		Explicit	Min. Entropy	Regression-based
0.01	0.011	0.011	0.007	0.004
0.1	0.095	0.095	0.067	0.066

Head Shape Regression. The proposed regression-based correspondence method is also used to study the growth of head shape from structural MRI data obtained from clinical studies spanning the age range from neonate to 5 year old. The 40 cases include 1.5T, T1-weighted MRI scans with resolutions of $1mm \times 1mm \times 1mm$ and $0.4mm \times 0.4mm \times 3.6mm$. The scans are preprocessed and segmented to obtain the head surfaces, which are input to the optimization process. Manually placed landmarks on the bridge of the nose and the openings of the left and right ear canals define a cutting plane and a pair of spheres that we use as constraints, as in Section. 2, to repel the correspondences from the neck, face, and ears, in order to restrict the analysis to the cranium, which is most interesting from a neurological point of view. Figure. 1 shows the particle system distributed across one of the head shapes after optimizing 500 particles.

Head size, measured in volume or circumference is well known to correlate with age. This is confirmed by the linear regression plot (size versus *log* of age) with $p < 2 \times 10^{-16}$, shown in Figure. 2. Next, the shapes were preprocessed

using methods mentioned in Section. 2 to remove the effects of size. Changes in head shape along the linear regression line (shape versus *log* of age) are shown in Figure. 3. Note the relative lengthening of the head, and the narrowing at the temples with increasing age. These shape changes are consistent with clinical observations that neonatal brain growth proceeds more rapidly in the forebrain. These results tie head shape to age in the paediatric setting.
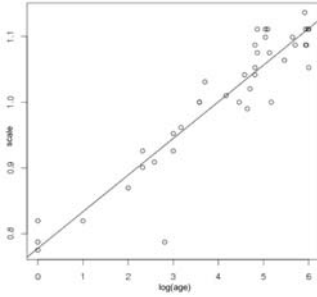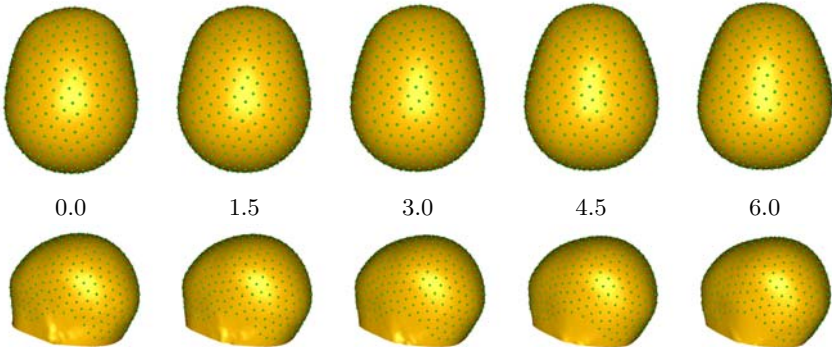

Fig. 2. Changes in head size with age

The permutation tests for both the proposed methods for this example showed that none of 1000 permutations gave a better correlation than the input data. While this $p = 0$ result is not conclusive, it does give strong evidence for significance. Future work will include more permutations to more accurately evaluate the significance.

The experiments were run on a 2GHz processor with run times of approximately 15 minutes for the tori (256 particles) and 40 minutes for the head shapes (500 particles). The permutation tests (1000 permutations) were run as parallel processes on a 16-processor machine.

Fig. 3. Overview of head shape regression: Changes in head shape with age

4 Conclusion

This paper describes a method for *shape regression* that accounts for explanatory variables in the placement of correspondences and allows for open surfaces with arbitrary geometric constraints, and presents a mechanism for hypothesis testing of the role of underlying variables in shape. Results from a study of head shape growth indicate that the proposed method can be applied to quantitative characterization of the relationship between age and head shape in young

children. Such analysis will generate data beyond the currently established standard of head circumference measurements as an index of growth. Moreover, it will generate normative data as a continuous growth model of shape, which can be useful in building optimal MRI head coils for young infants. The continuous shape model could also find use in population studies where two groups are compared with respect to growth trajectory rather than differences at individual time points.

References

1. Cates, J.E., Fletcher, P.T., Styner, M.A., Shenton, M.E., Whitaker, R.T.: 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)
2. Davies, R., Twining, C., Allen, P., Cootes, T., Taylor, C.: Shape discrimination in the hippocampus using an MDL model. In: Information Processing in Medical Imaging, pp. 38–50 (2003)
3. Gerig, G., Styner, M., Jones, D., Weinberger, D., Leiberman, J.: Shape analysis of brain ventricles using spharm. In: MMBIA, pp. 171–178. IEEE Press, Los Alamitos (2001)
4. Courchesne, E., Karns, C.M., Davis, H.R., Ziccardi, R., Carper, R.A., Tigue, Z.D., Chisum, H.J., Moses, P., Pierce, K.: Unusual brain growth patterns in early life in patients with autistic disorder: An MRI study. Neurology 57
5. Knickmeyer, R.C., Gouttard, S., Kang, C., Evans, D., Wilber, K., Smith, K.J., Hamer, R.M., Lin, W., Gerig, G., Gilmore, J.H.: A structural MRI study of human brain development from birth to two years. J. Neurosci. 28(47), 12176–12182 (2008)
6. Williams, J., Rossignac, J.: Tightening: curvature-limiting morphological simplification. In: Proc. Ninth ACM Symposium on Solid and Physical Modeling, pp. 107–112 (2005)