

CHAPTER 10

ShapeWorks

Particle-Based Shape Correspondence and Visualization Software

Joshua Cates^{*,†,‡}, Shireen Elhabian^{*,†,¶}, Ross Whitaker^{*,§}

^{*}Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, USA

[†]Biomedical Image and Data Analysis Core, University of Utah, Salt Lake City, USA

[‡]Comprehensive Arrhythmia Research and Management Center, University of Utah, Salt Lake City, USA

[§]School of Computing, University of Utah, Salt Lake City, USA

[¶]Faculty of Computers and Information, Cairo University, Cairo, Egypt

Contents

10.1	Introducing ShapeWorks	258
10.1.1	The Potential and the Challenge	258
10.1.2	Rising to the Challenge	259
10.1.2.1	Particle Systems for a Flexible Shape Representation	259
10.1.2.2	Optimized Correspondences Address the Model Selection Problem	260
10.1.3	ShapeWorks: An Open-Source Implementation of PBM	261
10.2	Particle-Based Modeling	262
10.2.1	Overview	262
10.2.2	Surface Representation	263
10.2.3	Adaptive Distributions on Surface Features	265
10.2.4	Surface Constraint	266
10.2.5	The Kernel Width σ for PDF Estimation	267
10.2.6	Numerical Considerations	267
10.2.7	Correspondence: Entropy Minimization in Shape Space	268
10.2.8	Setting Parameters	270
10.2.9	Illustration of the Properties and Interpretability of the PBM Optimization	271
10.3	PBM Extensions	273
10.3.1	Modeling Shape with Open Surfaces	274
10.3.2	Modeling Shape Complexes	276
10.3.3	Correspondence Based on Functions of Position	277
10.3.4	Correspondence with Regression Against Explanatory Variables	279
10.3.5	Dense PBM Correspondence Models	281
10.4	ShapeWorks Software Implementation and Workflow	284
10.4.1	The ShapeWorks Shape Modeling Workflow	284
10.4.1.1	Segmentation Preprocessing and Alignment	285
10.4.1.2	Initialization and Optimization	286
10.4.1.3	Analysis	287
10.4.2	The PBM Code Library	288
10.4.3	The ShapeWorks Software Tool Suite	289
10.5	ShapeWorks in Biomedical Applications	290

10.1 INTRODUCING SHAPEWORKS

10.1.1 The Potential and the Challenge

A revolution in shape analysis is underway. While morphometrics have been important for the study of biology and medicine for 100 years, recent advances in multivariate shape representation and statistics [1–8], coupled with the increased availability of computed tomography (CT) and magnetic resonance imaging (MRI), have enabled a whole new generation of morphometric tools. Characterized by the use of modern computational techniques to automatically construct detailed three-dimensional shape representations, these new morphometrics are called Statistical Shape Modeling, and they have the potential to measure anatomy and its variability with an unprecedented level of precision and statistical power.

Statistical Shape Modeling (SSM) is beginning to impact a wide spectrum of basic scientific and clinical applications, including the study of mechanisms of disease in neurobiology [9–11], the design of optimal patient-specific implants and bone substitutes [12–16], reconstruction of anatomical structures from both two- and three-dimensional medical images [9,17–27], computer-aided surgeries through pre- and postoperative surgical planning [28–31,20,19], and reconstructive surgery [32–39]. In addition, research involving large cohorts of image data, an emerging “big data” problem, is also seeing benefit from population-level SSM to better understand disease etiology, monitor pathology progression, and study treatment [16,40–43].

While SSM is poised to revolutionize morphometry, its widespread adoption by the biological and medical research communities has been hindered by a lack of both effective software implementations and generic approaches that can be applied across a wide variety of anatomies. The relative complexity of SSM algorithms makes it challenging to engineer software implementations that are user friendly. The increased computational requirements of these approaches also make it difficult to deploy SSM algorithms on a standard desktop computer, which is the only hardware widely available in research labs. Thus, to maximize the scientific and clinical impact of SSM, we need approaches that are designed to be both simple to use, from an algorithmic perspective, and that are scalable to commodity computer hardware. Furthermore, approaches that are robust to a variety of anatomies will avoid costly development and validation of multiple custom solutions. General approaches also support a larger user base, which leads to more standardization and acceptance of SSM methodologies within research communities.

10.1.2 Rising to the Challenge

It is in response to the preceding challenges that we introduce ShapeWorks. ShapeWorks is an open source software implementation of a surface correspondence approach to SSM called Particle Based Modeling (PBM). PBM is designed to be simple to use, robust, and applicable to general anatomy. Simplicity and generalizability are achieved through the use of a particle system representation of shape, instead of relying on specific shape parameterizations. Correspondence points are modeling as interacting sets of particles that redistribute themselves under an energy optimization. The optimization finds correspondence configurations that minimize the entropy of the model, which is a metric of information content. Thus, the optimization learns the shape parameters that are the most efficient descriptors of the geometry of the anatomy, which increases the robustness and statistical power of the model. This energy minimization is then balanced by a parameter-free regularization strategy that maximizes entropy of correspondence positions, in order to ensure good shape representations through efficient surface sampling.

10.1.2.1 Particle Systems for a Flexible Shape Representation

The choice of a specific shape representation defines, and may also limit, the class of shapes that can be modeled. To be applicable to the full range of shape analysis problems in biomedicine, a modeling methodology must be capable of representing different topological classes of shape. A correspondence-based model that relies on spherical parameterizations of shape, for example, can only represent manifold surfaces with spherical topologies [44,10]. By contrast, many important structures in the body, such as the heart, for example, consist of multiple interconnected chambers, with shared boundaries and open surfaces. The boundaries of other structures may be somewhat arbitrarily defined by specific landmarks or regions of interest. Many problems in orthopedics, for example, are concerned only with variability in specific areas of bone or the interaction of bone and cartilage surfaces at joints. In short, human anatomy can be very complex, and the simplifying assumptions of parametric models may even lead to erroneous or misleading results. Another important consideration is that medical or biological shapes are typically derived from the interfaces between organs or tissue types and usually defined implicitly in the form of segmented volumes, rather than explicit parameterizations, triangulations, or surface point samples. Such representations therefore require additional preprocessing steps that may limit the fidelity of the model and introduce error, especially as hypotheses regarding shape become more complex.

Instead of a parametric shape representation, PBM uses the idea of the particle system surface representation first proposed by Witkin and Heckbert, who introduced the idea of modeling a point set as a system of interacting particles that are constrained to lie on an implicit surface. Particles interact with one another with mutually repelling forces, such as electrostatic charge, so that they find distributions that optimally cover,




Figure 10.1 Femur shape representation using increasing numbers of particles.

and therefore describe, the surface geometry [45]. Meyer et al. proposed numerically robust extensions to this approach, including a new class of radial-basis energy functions and methods for curvature-adaptive surface sampling [46]. PBM adapts the numerical approaches of Meyer et al. for correspondence models by using a set of interacting particle systems, one for each shape in the sample, to produce optimal sets of surface correspondences. Adopting a point-based surface representation avoids many of the limitations and complexities inherent in parametric representations, such as the limitation to specific topologies and processing steps necessary to construct parameterizations. Another advantage is that, unlike representations that rely on surface meshes, particles do not have fixed neighbors and are free to move past one another to form different neighborhood configurations during the optimization process. This property means that the result is less constrained by the initialization and can potentially produce a less biased, more fully optimized model.

Fig. 10.1 illustrates the concept of a particle system representation of an implicit surface on a femur bone shape. The panels from left to right show an increasing number of particles placed on the surface and the resulting surface reconstruction from the particles. The surface reconstruction is done using the method for unorganized sets of points given by Hoppe et al. [47]. The number of particles doubles in each panel (256, 512, and 1024). As the particle count increases, so does the detail of the corresponding surface reconstruction.

10.1.2.2 Optimized Correspondences Address the Model Selection Problem

Correspondences offer a flexible shape representation and are intuitively the computational extension of traditional landmark models. However, the problem of how to automatically choose correspondence positions is difficult and ill-posed because of the fact that there are potentially an infinite number of possible configurations that can be

chosen for the correspondences on the shapes in the sample set. Thus, an important advance in SSM technology is the idea of choosing correspondence positions based on the minimization of an energy function of their positions (or associated features). Such optimized correspondence models are motivated by the Occam's razor principle of parsimony: given a number of possible models for the data, choose the simplest model. This idea has historically been widely applied to model selection problems in statistics, in order to find models with minimal numbers of parameters and greater predictive power [48]. The idea of an optimized correspondence model of shape was first proposed by Kotcheff and Taylor, who developed an algorithm that minimizes the magnitude of the covariance of the correspondences. Davies, Cootes and Taylor later expanded on this idea, using an information-theoretic cost function of correspondence positions based on minimum description length (MDL) [49,50]. The PBM approach to optimized correspondence uses a minimization based on entropy, which is a related measure to minimum description length, in that it seeks to minimize information content of the model.

A major consideration in the optimization process is to avoid overfitting to the data. If correspondence placement is completely unconstrained, for example, the optimal solution with respect to the information content is to place all correspondences as close to one another as possible. For this reason, optimized SSM approaches usually incorporate a constraint that ensures that correspondences faithfully represent the geometries of the samples. Several regularization strategies in the basic MDL formulation have been proposed that entail additional free parameters and assumptions about the quality of the initial parameterizations. One strategy, for example, constrains the solution so that it remains close to an anchor shape [50]. Such approaches, however, artificially limit the minimization process and bias the solution toward the anchor shape. To avoid assumptions about the initial quality of sample surface representations and ad-hoc regularizations, the PBM algorithm instead explicitly constructs good shape representations during the optimization procedure by maximizing an entropy measure on their distributions.

10.1.3 ShapeWorks: An Open-Source Implementation of PBM

Developed at the University of Utah, the underlying methods and theory of PBM have been described in a series of papers over the last 10 years [51–55]. The scientific and clinical effectiveness of ShapeWorks has been demonstrated in a range of applications including neuroscience [53,56–59], biological phenotyping [60,61], orthopedics [61–63] and cardiology [64,65]. The ShapeWorks software itself consists both a flexible C++ code library for the PBM optimization and an evolving suite of software applications that implement workflows for applying PBM to image data. The latter includes command line executables suitable for batch processing large cohorts of image data, as well

as desktop applications that support the complete shape analysis workflow, from preprocessing of image segmentations to analysis and visualization of PBM shape models.

The remainder of this chapter describes the theory and implementation of ShapeWorks, starting with the mathematical formulation and numerical implementation of the PBM algorithm and its extensions (Sections 10.2–10.3). Following the PBM development, we describe the ShapeWorks code distribution, which includes a C++ library of PBM code, and the workflow that is implemented in the ShapeWorks software applications (Section 10.4). Finally, we describe some scientific and clinical applications of ShapeWorks (Section 10.5) and conclude with a brief discussion of future work and ShapeWorks developments (Sections 10.6).

10.2 PARTICLE-BASED MODELING

10.2.1 Overview

The Particle Based Modeling (PBM) approach to SSM constructs a correspondence-point model of shape, which describes shape variation by choosing a discrete set of corresponding points on shape surfaces whose relative positions can be statistically analyzed. The correspondence model is analogous to a dense landmark model and is defined as follows. Consider a statistical sample of N surface representations drawn from a population of surfaces. The surface representations are embedded in a d -dimensional Cartesian space (typically, $d = 2$ or $d = 3$). A model for shape variation is constructed by choosing a set of M , d -dimensional points on each of the N surfaces. Each of the points is called a *correspondence* point. Collectively, the set of M points is known as the *configuration*, after Dryden and Mardia [5], and the space of all possible configurations is the *configuration space*. The *configuration matrix*, \mathbf{C} , is the $M \times d$ matrix of Cartesian coordinates in a configuration. The ordering of the points in the N configurations and, equivalently, the rows in the configuration matrices explicitly define the correspondences among the surfaces. Row $k \leq M$ in configuration matrix i , for example, corresponds to row k in configuration matrix j . The variation of the positional information encoded in the rows of the configuration matrices describe geometric variation in shape. Each configuration can be mapped to a single point X in a $d \times M$ -dimensional *shape space* by concatenating the correspondence coordinate positions into a single vector. The mapping to the dual shape space is invertible. The sample set forms a distribution in shape space, whose statistical properties can be estimated.

The PBM algorithm models the correspondence positions as sets of dynamic particles that are constrained to lie on the surface of the sample set, as in the surface sampling methods described in Section 10.1.2.1. The optimization is based on the idea of treating correspondence position in configuration space as a random variable, while simultaneously treating correspondence configuration as a random variable. Correspondence positions are optimized by gradient descent on an energy function that balances

the negative entropy of the distribution of particles in configuration space with the positive entropy of the distribution of the configurations in shape space. The method is to consider $\mathbf{z}_k \in \mathfrak{R}^{dM}$, $k = \{1, 2, \dots, N\}$ both as observations on a $dM \times 1$ vector random variable \mathbf{Z} and as N samples of M observations on N , $d \times 1$ vector random variables \mathbf{X}_k . The optimization to establish correspondence minimizes the energy function

$$Q = H(\mathbf{Z}) - \sum_{k=1}^N H(\mathbf{X}_k), \quad (10.1)$$

where H is an estimation of differential entropy. Minimization of the first term in Q produces a compact distribution of samples in shape space, while the second term seeks uniformly-distributed correspondence positions on the shape surfaces for accurate shape representation. Each term is given in commensurate units of entropy, avoiding the need for a separate regularization strategy. Because correspondence points in this formulation are not tied to a specific parameterization, the method operates directly on volumetric data and extends easily to arbitrary shapes, even nonmanifold surfaces.

10.2.2 Surface Representation

Consider a single configuration for a shape surface $\mathcal{S} \subset \mathfrak{R}^d$. The configuration consists of a discrete set of M points, which are the correspondence positions. The PBM formulation represents these positions with a set of particles, whose positions are considered a sample on a vector random variable $\mathbf{X} \in \mathfrak{R}^d$, with an associated probability density function describing their distribution. This probability density function $p(\mathbf{X} = \mathbf{x})$ gives the probability of an observation \mathbf{x} on \mathbf{X} , denoted as $p(\mathbf{x})$. In the limit, the amount of information contained in the sample on \mathbf{X} is the differential entropy of $p(\mathbf{X})$,

$$H(\mathbf{X}) = - \int_{\mathcal{S}} p(\mathbf{X}) \log p(\mathbf{X}) dx = -E\{\log p(\mathbf{X})\}, \quad (10.2)$$

where $E\{\cdot\}$ is the expectation. When there are a sufficient number of points sampled from p , the expectation can be approximated by the sample mean [66], which gives

$$H(\mathbf{X}) \approx -\frac{1}{M} \sum_{i=1}^M \log p(\mathbf{x}_i). \quad (10.3)$$

The PBM algorithm manipulates particle positions using a gradient-descent optimization on a cost function C , that is an approximation of negative entropy,

$$C(\mathbf{x}_1, \dots, \mathbf{x}_M) \approx -H(\mathbf{X}). \quad (10.4)$$

The optimization problem is given by

$$\mathbf{z} = \arg \min_{\mathbf{z}} C(\mathbf{x}_1, \dots, \mathbf{x}_M) \text{ s.t. } \mathbf{x}_1, \dots, \mathbf{x}_M \in \mathcal{S}, \quad (10.5)$$

and uses a Gauss–Seidel update with forward differences. Each particle therefore moves with a time parameter and positional update,

$$\mathbf{x}_i \leftarrow \mathbf{x}_i - \gamma \frac{\partial C}{\partial \mathbf{x}_i}, \quad (10.6)$$

where γ is a time step. The partial gradient of C for particle i is

$$\frac{\partial C}{\partial \mathbf{x}_i} = \frac{\partial}{\partial \mathbf{x}_i} \frac{1}{M} \sum_{j=1}^M \log p(\mathbf{x}_j) = \frac{1}{M} \sum_{j=1}^M \frac{\frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_j)}{p(\mathbf{x}_j)}. \quad (10.7)$$

The gradient requires estimates of the probability $p(\mathbf{X} = \mathbf{x}_j)$. For distributions of particles on surfaces, a probability density function may be quite complex, which suggests a nonparametric, kernel-based approach. The PBM algorithm uses a Parzen windowing density estimation [67] that is based on the particle configurations. The probability of the position of a particle in this formulation is given by the mixture of multivariate Gaussian kernels,

$$p(\mathbf{x}, \sigma) \approx \frac{1}{M} \sum_{j=1}^M G(\mathbf{x} - \mathbf{x}_j, \sigma), \quad (10.8)$$

where $G(\mathbf{x} - \mathbf{x}_j, \sigma)$ is a d -dimensional, isotropic Gaussian with standard deviation σ . When $j = i$ in (10.7), the partial derivative of p with respect to particle position is

$$\frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_i, \sigma_i) = \frac{1}{\sigma_i^2 M} \sum_{j=1}^M G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i) (\mathbf{x}_i - \mathbf{x}_j). \quad (10.9)$$

When $i \neq j$, the derivative is

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_j, \sigma_j) &= \frac{1}{M} \left[\frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_1, \sigma_j) + \frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_2, \sigma_j) + \dots \right. \\ &\quad \left. + \frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_i, \sigma_j) + \dots + \frac{\partial}{\partial \mathbf{x}_i} G(\mathbf{x}_j - \mathbf{x}_M, \sigma_j) \right] \\ &= \frac{1}{M} \left[0 + 0 + \dots - \sigma_j^{-2} G(\mathbf{x}_i - \mathbf{x}_j, \sigma_j) (\mathbf{x}_i - \mathbf{x}_j) + \dots + 0 \right] \\ &= \frac{1}{\sigma_j^2 M} G(\mathbf{x}_j - \mathbf{x}_i, \sigma_j) (\mathbf{x}_i - \mathbf{x}_j). \end{aligned} \quad (10.10)$$

Substituting (10.9) and (10.10) into (10.7) gives

$$\frac{\partial C}{\partial \mathbf{x}_i} = \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i)(\mathbf{x}_i - \mathbf{x}_j)}{\sigma_i^2 p(\mathbf{x}_i, \sigma_i)} + \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_j - \mathbf{x}_i, \sigma_j)(\mathbf{x}_i - \mathbf{x}_j)}{\sigma_j^2 p(\mathbf{x}_j, \sigma_j)}. \quad (10.11)$$

The computational complexity for Eq. (10.11) is $\mathcal{O}(M^2)$, since the entire density function p must be recomputed for each particle update. To simplify the computation, the PBM formulation instead considers p to be fixed for a given particle update: for $j \neq i$ in (10.11), the estimation of the density function at j is allowed to lag behind the update of particle position i . Under this assumption, $\frac{\partial}{\partial \mathbf{x}_i} p(\mathbf{x}_j, \sigma_j) = 0$, and the second term in (10.11) drops out, simplifying the gradient computation to only $\mathcal{O}(M)$.

After dropping the second term, the final approximation to the gradient of particle positional entropy is given by

$$\begin{aligned} \frac{\partial C}{\partial \mathbf{x}_i} &\approx \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i)(\mathbf{x}_i - \mathbf{x}_j)}{\sigma_i^2 p(\mathbf{x}_i, \sigma_i)} \\ &= \frac{1}{M} \sum_{j=1}^M \frac{G(\mathbf{x}_i - \mathbf{x}_j, \sigma_i)}{\sigma_i^2 \frac{1}{M} \sum_{k=1}^M G(\mathbf{x}_i - \mathbf{x}_k, \sigma_i)} (\mathbf{x}_i - \mathbf{x}_j) \\ &= \frac{1}{M} \sum_{j=1}^M w_{ij} (\mathbf{x}_i - \mathbf{x}_j), \end{aligned} \quad (10.12)$$

where w_{ij} are Gaussian weights based on interparticle distance and $\sum_j w_{ij} = 1$. To minimize C , the particles must move away from each other. Thus, we have a set of particles moving under a repulsive force and constrained to lie on the surface, with $\gamma < \sigma^2$ in (10.6) for stability. The motion of each particle is away from all of the other particles, but interactions are effectively local for sufficiently small σ , where w_{ij} vanishes with increasing interparticle distance.

10.2.3 Adaptive Distributions on Surface Features

The preceding minimization produces a uniform sampling of a surface. For some applications, a strategy that samples adaptively in response to higher order shape information is more effective for several reasons. From a numerical point of view, the minimization strategy relies on a degree of regularity in the tangent planes between adjacent particles, which argues for sampling more densely in high curvature regions. An adaptive sampling strategy also produces a more efficient representation of geometric detail by reducing redundant samples in flatter regions. Adaptive sampling with PBM is implemented by

modifying the Parzen windowing in Eq. (10.8) as follows:

$$\tilde{p}(x_i) \approx \frac{1}{M} \sum_{j=1, j \neq i}^M G\left(\frac{\mathbf{x}_i - \mathbf{x}_j}{k_j}, \sigma_i\right) \quad (10.13)$$

where k_j is a scaling term proportional to the curvature magnitude computed at each neighbor particle j . The effect of this scaling is to expand space in response to local curvature. A uniform sampling based on maximum entropy in the warped space translates into an adaptive sampling in unwarped space, where points pack more densely in higher curvature regions. The extension of Eq. (10.12) to incorporate the curvature-adaptive Parzen windowing is straightforward to compute. Since k_j is not a function of x_i , the modified gradient is

$$\frac{\partial C}{\partial \mathbf{x}_i} \approx \frac{1}{M} \sum_{j=1}^M \frac{G((\mathbf{x}_i - \mathbf{x}_j)/k_j, \sigma_i)(\mathbf{x}_i - \mathbf{x}_j)}{\sigma_i^2 k_j p(x_i, \sigma_i)}. \quad (10.14)$$

There are many possible choices for the scaling term k . Meyer et al. [68] describe an adaptive surface sampling that uses the scaling

$$k_i = \frac{1 + \rho \kappa_i (\frac{s}{2\pi})}{\frac{1}{2}s \cos(\pi/6)}, \quad (10.15)$$

where κ_i is the root sum-of-squares of the principal curvatures at surface location x_i . The user-defined variables s and ρ specify the ideal distance between particles on a planar surface and the ideal density of particles per unit angle on a curved surface, respectively. Note that the scaling term in this formulation could easily be modified to include surface properties other than curvature.

10.2.4 Surface Constraint

The surface constraint in both the uniform and adaptive optimizations is specified by the zero set of a scalar function $F(x)$. This constraint is maintained, as described in several papers [46], by projecting the gradient of the cost function onto the tangent plane of the surface, as prescribed by the method of Lagrange multipliers. The projection operator is given by

$$\mathbf{I} - \mathbf{n} \otimes \mathbf{n}, \quad (10.16)$$

where \mathbf{I} is the identity matrix, \mathbf{n} is the normal to the surface, and \otimes denotes the outer, or tensor, product. The tangent-plane projection is followed by iterative reprojection of the particle onto the nearest root of F by the Newton–Raphson method. Principal curvatures are computed analytically from the implicit function, as described in [69].

10.2.5 The Kernel Width σ for PDF Estimation

Finally, the kernel width σ of the Parzen windowing estimation of particle density must be chosen at each particle. This is done automatically, before the positional update, using a maximum likelihood optimality criterion. The contribution to C of the i th particle is simply the probability of that particle position. Optimizing that quantity with respect to σ therefore gives a maximum likelihood estimate of σ for the current particle configuration. Using the Newton–Raphson method, the strategy is to find σ such that

$$\partial p(\mathbf{x}, \sigma)/\partial \sigma = 0, \quad (10.17)$$

which typically converges to machine precision in several iterations. For the adaptive sampling case, we find σ such that

$$\partial \tilde{p}(\mathbf{x}, \sigma)/\partial \sigma = 0, \quad (10.18)$$

so that the optimal σ is scaled locally based on the curvature. The iteration is given by

$$\sigma^{t+1} \leftarrow \sigma^t + \frac{\frac{\partial p}{\partial \sigma}}{\frac{\partial^2 p}{\partial \sigma^2}}, \quad (10.19)$$

and the first derivative of p with respect to σ , from (10.8), is

$$\begin{aligned} \frac{\partial}{\partial \sigma} \sum_{j=1}^M \frac{1}{M} G(\mathbf{x} - \mathbf{x}_j, \sigma) &= \frac{\partial}{\partial \sigma} \sum_{j=1}^M \frac{1}{M(2\pi)^{d/2}\sigma^d} e^{-\frac{r_j}{2\sigma^2}} \\ &= \frac{1}{M(2\pi)^{d/2}\sigma^{d+3}} \sum_{j=1}^M e^{-\frac{r_j}{2\sigma^2}} (r_j - d\sigma^2), \end{aligned} \quad (10.20)$$

where $r_j = (\mathbf{x} - \mathbf{x}_j)^T(\mathbf{x} - \mathbf{x}_j)$ is the distance from x to x_j . The second derivative follows from (10.20), and is given by

$$\begin{aligned} \frac{\partial^2}{\partial \sigma^2} \sum_{j=1}^M \frac{1}{M} G(\mathbf{x} - \mathbf{x}_j, \sigma) &= \frac{1}{M(2\pi)^{d/2}\sigma^{d+6}} \\ &\times \left[\sum_{j=1}^M e^{-\frac{r_j}{2\sigma^2}} (r_j^2 - (3+2d)\sigma^2 r_j + d(1+d)\sigma^4) \right]. \end{aligned} \quad (10.21)$$

10.2.6 Numerical Considerations

There are a few important numerical considerations in computing the particle-based surface representation. First, the Gaussian kernels must be truncated. Typically, we truncate kernels so that $G(x, \sigma) = 0$ for $|x| > 3\sigma$. This means that each particle has a finite radius of influence, and a spatial binning structure to identify neighboring particles can

be used to reduce the computational burden associated with particle interactions. A second consideration is the case where σ for a particle is too small to allow the particle to interact with its neighbors, and updates of σ or position cannot be computed. When σ is small, kernel size is updated using $\sigma \leftarrow 2 \times \sigma$, until σ is large enough for the particle to interact with its neighbors. A final numerical consideration is that the system must include bounds σ_{\min} and σ_{\max} to account for anomalies such as bad initial conditions or too few particles. These are not critical parameters, and as long as they are set to include the minimum and maximum resolutions, the system operates reliably.

One final aspect of the particle formulation to consider is that it computes the Euclidean distance between particles, rather than the geodesic distance on the surface. The PBM algorithm therefore assumes sufficiently dense samples so that nearby particles lie in the tangent planes of the zero sets of F . This is an important consideration; in cases where this assumption is not valid, such as highly convoluted surfaces, the distribution of particles may be affected by neighbors that are outside of the true manifold neighborhood. Limiting the influence of neighbors whose normals differ by some threshold value (e.g., 90 degrees) does limit these effects.

10.2.7 Correspondence: Entropy Minimization in Shape Space

A sample set, \mathcal{E} , is a collection of N surfaces, each with their own set of M particles mapped to a single, dM -dimensional vector in shape space, i.e. $\mathcal{E} = \mathbf{z}^1, \dots, \mathbf{z}^N$. The sample set in vector form may be collected into a single matrix $\mathbf{P} = \mathbf{z}_j^k$, with particle positions along the rows and shape samples across the columns. Modeling $\mathbf{z}^k \in \mathfrak{R}^{dM}$ as an instance of random variable \mathbf{Z} , the PBM method for correspondence minimizes the combined sample and shape cost function

$$Q = H(\mathbf{Z}) - \sum_{k=1}^N H(\mathbf{X}^k), \quad (10.22)$$

which favors a compact representation of the sample, and is balanced against a uniform distribution of particles on each surface.

For this discussion we assume that the complexity of each shape is greater than the number of samples, and so normally $dM > N$. Given the low number of examples relative to the dimensionality of the space, the density estimation requires some assumptions. The PBM algorithm therefore assumes a normal distribution and models $p(\mathbf{Z} = \mathbf{z})$ parametrically using an anisotropic Gaussian with covariance Σ . The entropy is then given by

$$H(\mathbf{Z}) \approx \frac{1}{2} \log |\Sigma| = \frac{1}{2} \sum_{j=1}^{dM} \log \lambda_j, \quad (10.23)$$

where $\mathbf{e}_k, \lambda_k, j = 1, \dots, dM$ are the eigenvalues of Σ .

In practice, Σ will not have full rank, in which case the entropy is not finite. The problem must therefore be regularized with the addition of a diagonal matrix $\alpha \mathbf{I}$ to introduce a lower bound on the eigenvalues. The covariance is estimated from the data, and is given by

$$\Sigma = (dMN - 1)^{-1} \mathbf{Y} \mathbf{Y}^T, \quad (10.24)$$

where

$$\mathbf{y}^k = \mathbf{z}^k - \boldsymbol{\mu}, \text{ and } \boldsymbol{\mu} = \frac{1}{N} \sum_{k=1}^N \mathbf{z}^k. \quad (10.25)$$

Thus, \mathbf{Y} denotes the matrix of sample vectors \mathbf{P} minus the sample mean $\boldsymbol{\mu}$, i.e. $\mathbf{Y} = \mathbf{P} - \boldsymbol{\mu} \mathbf{1}^T$, where $\mathbf{1}$ is a $dM \times 1$ vector of ones. Because $N < dM$, the eigenanalysis in (10.23) is done on the dual space of the $N \times N$ covariance matrix $\Sigma^T = (dMN - 1)^{-1} \mathbf{Y}^T \mathbf{Y}$. The nonzero eigenvalues of Σ can be obtained from Σ^T by noting the following relationships (see also [50]). For eigenvalues and eigenvectors $\{\mathbf{e}_k, \lambda_k\}$ of Σ , $\Sigma \mathbf{e}_k = \lambda_k \mathbf{e}_k$. Similarly, for eigenvalues and eigenvectors $\{\mathbf{e}'_k, \lambda'_k\}$ of Σ^T , $\Sigma^T \mathbf{e}'_k = \lambda'_k \mathbf{e}'_k$. Substituting for Σ^T , we have

$$(dMN - 1)^{-1} \mathbf{Y}^T \mathbf{Y} \mathbf{e}'_k = \lambda'_k \mathbf{e}'_k, \quad (10.26)$$

and premultiplying each side by \mathbf{P} gives

$$(dMN - 1)^{-1} \mathbf{Y} \mathbf{Y}^T \mathbf{Y} \mathbf{e}'_k = \lambda'_k \mathbf{Y} \mathbf{e}'_k, \quad (10.27)$$

which is equivalent to

$$\Sigma (\mathbf{Y} \mathbf{e}'_k) = \lambda'_k (\mathbf{Y} \mathbf{e}'_k). \quad (10.28)$$

Thus, $\mathbf{e}_k = \mathbf{Y} \mathbf{e}'_k$, and $\lambda_k = \lambda'_k$, for nonzero eigenvectors of Σ . The covariances $|\Sigma|$ and $|\Sigma^T|$ are therefore equivalent (up to a constant factor of α), and the final cost function G associated with the sample entropy is given by

$$G(\mathbf{P}) = \frac{1}{2} \log |\Sigma| = \frac{1}{2} \log \left| \frac{1}{dMN - 1} \mathbf{Y}^T \mathbf{Y} + \alpha \mathbf{I} \right|. \quad (10.29)$$

To compute the gradient of G , we follow a logic similar to that used in the derivation of (10.11), and allow the estimation of the mean $\boldsymbol{\mu}$ of the distribution \mathbf{Z} to lag behind the updates $\frac{\partial G}{\partial \mathbf{P}}$. This allows for the simplifying assumption $\frac{\partial G}{\partial \mathbf{P}} \approx \frac{\partial G}{\partial \mathbf{Y}}$. This approximation becomes more accurate as the number of shape samples is increased, and changes in individual particle positions have increasingly less of an effect on the sample mean. The

matrix of partial derivatives of G with respect to \mathbf{Y} is derived as follows.

$$\begin{aligned}
\frac{\partial G}{\partial \mathbf{Y}} &= \frac{\partial}{\partial \mathbf{Y}} \left(\frac{1}{2} \log \frac{1}{(dMN - 1)^n} + \frac{1}{2} \log |\mathbf{Y}^T \mathbf{Y}| \right) \\
&= 0 + \frac{1}{2} |\mathbf{Y}^T \mathbf{Y}|^{-1} \frac{\partial}{\partial \mathbf{Y}} |\mathbf{Y}^T \mathbf{Y}| \\
&= |\mathbf{Y}^T \mathbf{Y}|^{-1} |\mathbf{Y}^T \mathbf{Y}| \mathbf{Y} (\mathbf{Y}^T \mathbf{Y})^{-1} \\
&= \mathbf{Y} (\mathbf{Y}^T \mathbf{Y})^{-1}.
\end{aligned} \tag{10.30}$$

Adding the regularization to the covariance, we have the following equation for the updates

$$\frac{\partial G}{\partial \mathbf{P}} \approx \mathbf{Y} (\mathbf{Y}^T \mathbf{Y} + \alpha \mathbf{I})^{-1}. \tag{10.31}$$

The regularization α on the inverse of $\mathbf{Y}^T \mathbf{Y}$ can now be seen to account for the possibility of a diminishing determinant. The negative gradient $-\partial G / \partial \mathbf{P}$ gives a vector of updates for the entire system, which is recomputed once per system update. This term is added to the shape-based updates described in the previous section to give the update of each particle:

$$\mathbf{z}_j^k \leftarrow \gamma \left[-\partial G / \partial \mathbf{z}_j^k + \partial E^k / \partial \mathbf{z}_j^k \right]. \tag{10.32}$$

The stability of this update places an additional restriction on the time steps, requiring γ to be less than the reciprocal of the maximum eigenvalue of $(\mathbf{Y}^T \mathbf{Y} + \alpha \mathbf{I})^{-1}$, which is bounded by α . Thus, we have $\gamma < \alpha$ and note that α has the practical effect of preventing the system from slowing too much as it tries to reduce the thinnest dimensions of the sample distribution. This also suggests an annealing approach for computational efficiency in which α starts off somewhat large (e.g., the size of the shapes) and is incrementally reduced as the system iterates.

10.2.8 Setting Parameters

The SSM approach described thus far is a self-tuning system of particles that distribute themselves across the shape surface using repulsive forces to achieve optimal distributions. Particles may also optionally adjust their sampling frequency locally in response to surface curvature. Free parameters of the system are limited to the choice of the number of particles (M), and the parameters s and ρ from (10.15), if adaptive sampling is used. In practice, adaptivity parameters are typically determined empirically based on the data under analysis. The number of particles is also typically chosen empirically by adding particles until the representation is deemed to capture enough details for the given application.



Figure 10.2 PBM correspondences on example tori from a random distribution on r and R . Colors indicate correspondence. (For interpretation of the references to color in this figure, the reader is referred to the web version of this chapter.)

In order to explicitly manage the tradeoff between model compactness and the geometric regularization, an additional free parameter β may be introduced into Eq. (10.1) as follows:

$$Q = H(\mathbf{Z}) - \beta \sum_{k=1}^N H(\mathbf{X}_k). \quad (10.33)$$

Empirical results, however, suggest that the two terms in this function are already well balanced and $\beta = 1$ represents a good default setting.

10.2.9 Illustration of the Properties and Interpretability of the PBM Optimization

Finally, we present two examples that illustrate the properties of the optimization. The first is an experiment on a class of nonspherical shapes, for which the PBM optimization recovers the optimal ground truth shape parameters. The second example illustrates how PBM is able to discover the underlying mode of variation in a box ensemble with a moving bump in comparison to diffeomorphism-based shape modeling.

To illustrate PBM optimization results for ground truth synthetic data, we applied the algorithm to sample set of 40 randomly generated tori, which are nonspherical shapes that can be described by exactly two shape parameters, a small radius r and the large radius R (see also the example given in Fig. 10.3). Tori were randomly chosen from a distribution with mean $r = 1$, $R = 2$ and $\sigma_r = 0.30$, $\sigma_R = 0.15$. A rejection policy was used to exclude invalid tori (e.g., $r > R$). Correspondences were optimized using 1024 particles per shape, and a uniform sampling (no adaptivity). Fig. 10.2 shows the particle system distribution across several of the torus shapes in the sample set with 1024 correspondences. Correspondence positions are indicated by spherical glyphs and correspondence across shapes is indicated by the color of the glyph. Surface reconstructions for each sample were done using the correspondence positions and the algorithm given by Hoppe et al. for collections of unorganized points [47]. A principal component analysis (PCA) of the resulting correspondence positions indicates that the particle system method discovered two pure modes of variation. PCA mode 1 contains 69.7870% of




Figure 10.3 PCA modes 1 and 2 illustrated for the PBM torus shape model at -3 to $+3$ standard deviations from the mean shape. PCA 1 (top row) corresponds to r and PCA 2 (bottom row) correspondence to R .

total variation and PCA mode 2 contains 30.2076% of total variation. Less than 0.006% of total variation remains in the smaller, “error” modes.

Empirical observation of PCA modes 1 and 2 suggests that they correspond well to variation in r and R , respectively, from the parametric model that was used to generate the sample data. Fig. 10.3 shows the mean correspondence positions from the model moved along each of the top two PCA modes. Torus shapes along each mode are reconstructed from the learned PBM model parameters at -3 to $+3$ standard deviations from the mean. The top row illustrates variation in PCA 1, which corresponds to r , and the bottom row indicates variation in PCA 2, which corresponds to R . In this experiment, the PBM method appears to have estimated the true orthogonal modes of variation of the torus shape sample.

In the second experiment, we constructed an ensemble of 15 three-dimensional “box-bump” shapes with a bump at a varying location (see top row of Fig. 10.4). Each shape was constructed as a fast-marching distance transform of a union of a rounded-corner cuboid and an ellipsoid representing a bump added at a random location along the top side of the cuboid. This example is interesting because we would, in principle, expect a correspondence algorithm that is minimizing information content to discover this single mode of variability in the sample set. We used the PBM method to optimize 1024 particles per shape under uniform sampling on the “box-bump” shapes (see middle row of Fig. 10.4).

As opposed to optimized correspondences on shape surfaces, the shape geometry can be embedded in the image intensity values at pixels or voxels and then nonlinear registration can be used to map all sample images to a reference image, or atlas. The variation in shape is then considered to be captured by the nonlinear registration parameters. Of the image registration methods, the diffeomorphic methods are in most widespread use. Hence, it is important to show whether modeling shape variations using diffeomorphic warps are able to recover the known mode of variation (i.e., a moving bump) as compared to the analysis of the optimized correspondence model. In this regard, we generated an unbiased atlas of the “box-bump” ensemble as proposed in [70] and implemented in AtlasWerks [71]. We then constructed a parametric statistical model of shape variation using PCA on the deformation fields that map each box-bump sample to the estimated unbiased atlas (see bottom row of Fig. 10.4).




Figure 10.4 The box-bump experiment: PBM-based and diffeomorphic-based shape models.

Fig. 10.4 illustrates the mean and three standard deviations of the first mode of the PBM- and diffeomorphic-based models. Shapes from the particle method remain more faithful to those described by the original training set, even out to three standard deviations where the diffeomorphic description breaks down. In particular, one can observe that PCA identified a single dominant mode of variation for the PBM method (see middle row of Fig. 10.4). However, diffeomorphic warps recovered incorrect shape model in which the mean shape showed a box with two bumps rather than a single bump and the first mode of variation represented the relative height of the two “artificial” bumps. Further, diffeomorphic-based shape modeling showed five dominant shape modes in an ensemble of a single mode.

10.3 PBM EXTENSIONS

This section describes several mathematical extensions of the particle-based modeling (PBM) algorithm that are designed to make it more robust to realistically complex shape analysis problems. This work is motivated by the needs of the biomedical research community for tools to model more complicated anatomical shapes and statistical designs, such as joint variability of multiple structures and optimization with respect to explanatory variables. Cardiac anatomy, for example, consists of multiple interconnected chambers with shared openings, valve annuli, and septa. In orthopedics, researchers are often concerned with the mechanical interactions of multiple bone surfaces, in order to understand dysfunction in joints. In other cases, the geometric features of an anatomical object are not sufficient to properly establish correspondence. Some anatomy is highly variable across subjects and additional information, such as functional data, is helpful in determining how surface regions correspond. The cortical surface of the brain represents one such example. Cortical folding patterns are highly variable among individuals,

and neuroanatomists typically rely on information such as sulcal depth and vascular connectivity for correspondence, rather than geometric information alone.

Extensions to the PBM algorithm that are developed in this section include the following. Section 10.3.1 describes an approach for modeling surface with open boundaries. The open surface method allows an arbitrary boundary to be defined as the intersection of a closed surface with a set of shape primitives. Section 10.3.2 describes how the PBM algorithm can be used to optimize the joint correspondence among shapes that consist of *multiple* anatomical objects. This capability is important in the study of shape covariance among anatomy that is functionally or structurally correlated. Section 10.3.3 describes a generalization of the PBM optimization criteria to correspondence in arbitrary, multivariate functions of position, rather than only considering positional information. This approach is useful for problems where there is data other than geometric information that indicates correspondence, such as multimodal imaging studies, and studies with functional imaging data. Finally, Section 10.3.4 describes a methodology for including a regression model on independent variables into the PBM correspondence optimization. Shape regression modeling can improve statistical power when controlling for correlations between shape and factors such as age or other clinical variables.

10.3.1 Modeling Shape with Open Surfaces

Conceptually, there are two ways to handle a surface boundary when optimizing correspondences. The first approach is to explicitly represent and model the boundary, which requires that correspondences must be allowed to lie on the boundary, and the optimization must track particle movement on and off of the boundary. This approach is appropriate for applications where the boundary shape is of specific interest to the problem, such as specifically modeling the shapes of valve annuli or the ostia of vascular openings in conjunction with chamber shape in the heart. In many cases, however, it is not important, or even desirable, to model the variation in the shape of the boundary. A segmentation, for example, may contain noise in the boundary shape due to ambiguities in its specification during the segmentation process. Such examples often arise in orthopedics, for example, where only the proximal or distal end of a bone may be of interest. In this situation, where the boundary is considered noisy, it can simply be treated as a constraint on the particle optimization, which is the approach currently implemented for ShapeWorks. Explicit modeling of open shape boundaries is left for future work.

The PBM algorithm for correspondence on open surfaces represents the surface boundary as the intersection of a closed surface (e.g., \mathcal{S} in Section 10.2.2) with a set of geometric primitives, such as cutting planes and spheres. The boundary representation is then used to influence the entropy maximization of the PBM algorithm particle position (Section 10.2.2), so that it indirectly constrains the positions of particles to lie within




Figure 10.5 Increasing numbers of PBM particles on an open surface, where the boundary is defined by the intersection of an implicit bone surface with a sphere and a cutting plane.

the surface boundary. The goal of the open surface modeling algorithm is to formulate particle interactions with the boundaries so that the positions of the constraints have as little influence as possible on the statistical shape model. This approach is consistent with the idea that the boundary shape may contain noise, and we wish to minimize the influence of this noise on the model.

The algorithm proceeds as follows. For each geometric primitive in the surface boundary representation, the algorithm constructs a *virtual* particle distribution that consists of all of the closest points on its surface to the particles with positions \mathbf{x}_j on \mathcal{S} . During the gradient descent optimization, particles \mathbf{x}_j 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 (Eq. (10.8)), particles \mathbf{x}_j maintain a distance from the 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 boundary itself, limiting the effect of errors in the boundary specification on the configuration. Note that the virtual particle distributions are also not used in the correspondence optimization term (the sample entropy from Eq. (10.23)) and therefore do not directly affect the distribution of samples in shape space.

Fig. 10.5 illustrates a particle configuration using the method outlined above for open surfaces, and shows the effect of increasing the number of particles. In the figure, the open surface boundary is defined by the intersection of an implicit bone surface, a cutting plane, and a sphere. As the number of particles is increased, the distribution samples regions of the bone become closer and closer to the surface boundary. Note, however, that the particle distribution never touches or crosses the boundaries.

Further examples are given in Fig. 10.6. Fig. 10.6A illustrates placement of particles on a neonatal head surface (from MRI), with a cutting plane defined by anatomical landmarks at the tip of the nose and the center of the ears, and spheres placed at those same landmarks to exclude both nose and ears from the analysis. This approach was used in work characterizing shape change during neonatal head development [60]. Another example is that of the left atrium, which is shown in Fig. 10.6B. The left atrium is a complex structure with openings to the mitral valve and the pulmonary veins. For




Figure 10.6 Examples of PBM optimization on open surfaces, including pediatric head shape (A) and the left atrium (B).

studies involving left atrial shape in atrial fibrillation [64], we masked out the mitral valve region using a collection of exclusion spheres.

10.3.2 Modeling Shape Complexes

In general, biological function and phenotype is not explained by any single anatomical structure, but are instead the result of complex systems of functionally- or structurally-related anatomy. This section presents an extension of the PBM algorithm for shape modeling of multiple, disconnected anatomical surfaces, or what we will call shape complexes. A *multiobject complex* is defined as a set of solid shapes, each representing a single, connected biological structure. The complex of structures are assembled into a scene within a common coordinate frame. Each structure in a multiobject complex contains shape, pose, scale, and positional information. 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. The proposed correspondence method for establishing correspondence on multiobject complexes 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 in the optimization process. By explicitly modeling the correlations among variabilities, however, optimization in the joint space may produce more compact distributions for correspondences, resulting in fewer model parameters and greater statistical power.

The particle-based correspondence method described in Section 10.2 can be directly applied to multiobject complexes by treating all of the objects in the complex as defining a single surface. 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. In this way, each correspondence is guaranteed to stay on a particular anatomical structure, and the surface sampling is not influenced by regions where structures in the complex happen to be near to one

another. The shape-space statistics using this method, however, remain coupled, and the covariance Σ (Eq. (10.23)) includes all particle positions across the entire complex, so that optimization takes place on the joint, multiobject model.

As with the single-object framework, any set of implicitly defined surfaces is appropriate as input to the multiobject framework, with similar preprocessing considerations as those discussed in Section 10.4.1. In the case of binary segmentations, the input is now a set of N segmentations of K -object complexes, which contains $N \times K$ distinct, volumetric label masks.

10.3.3 Correspondence Based on Functions of Position

Geometric features of an anatomical object are often not sufficient to properly establish correspondence. The basic PBM algorithm described in Section 10.2 only considers particle position information in the optimization, which only represents the geometric, or structural, information of the shape surface. Here we describe an extension to the PBM algorithm to establish correspondence by minimization of the entropy of arbitrary, vector-valued *functions* of position. This more general method is useful in cases where the notion of correspondence is not well defined by the surface geometry, but can be described by other metrics.

The extension to the PBM algorithm to incorporate functional data, which we refer to as the *generalized* PBM algorithm, is straightforward. It consists of substituting the entropy estimation of the matrix of particle positions with an entropy estimation on an arbitrary, vector-valued function of the particle position. From Section 10.2, the energy term for the basic PBM optimization is given by

$$Q = H(\mathbf{Z}) - \beta \sum_{k=1}^N H(\mathbf{X}_k), \quad (10.34)$$

where H is an estimation of entropy, X_k is a vector random variable with the distribution of particle configuration k , and Z is the vector random variable with the distribution of the shape samples in the dM -dimensional shape space. The extension to the generalized PBM algorithm only modifies the correspondence term $H(\mathbf{Z})$. The entropy associated with individual correspondence configurations, $H(\mathbf{X}_k)$, is not modified, and still operates on positional information. In other words, particles are still constrained to lie on the surface of the shape and distribute themselves across shape surfaces using the maximization of positional entropy, but their correspondence is established using a function of positional information. Note that a function of position could be designed to also include particle position, so that both structural and functional data influence the correspondences.

Recall from Section 10.2.7, that the entropy estimation of the sample distribution in shape space is given by

$$H(\mathbf{Z}) \approx \frac{1}{2} \log |\Sigma|, \text{ and } \Sigma = (dMN - 1)^{-1} \mathbf{Y}\mathbf{Y}^T, \quad (10.35)$$

where Σ is the covariance matrix, and \mathbf{Y} is the $dM \times N$ data matrix P of sample vectors $\mathbf{z}^k, k = \{1, \dots, N\}$, minus the sample mean $\boldsymbol{\mu}$, and each vector \mathbf{z}^k consists of the positional information from M particles on the shape surface k . In the case of computing entropy of vector-valued functions of the correspondence positions, the extension to functional data considers the more general case where columns of the data matrix are instead given by

$$\tilde{\mathbf{p}}^k = \begin{bmatrix} f(\mathbf{x}_0^k) \\ f(\mathbf{x}_1^k) \\ \vdots \\ f(\mathbf{x}_j^k) \\ \vdots \\ f(\mathbf{x}_{M-1}^k) \end{bmatrix}, \quad (10.36)$$

where \mathbf{x}_j^k is the positional information of particle j for shape k , and $f: \mathfrak{R}^d \rightarrow \mathfrak{R}^q$.

The matrix \mathbf{Y} now becomes a matrix $\tilde{\mathbf{Y}}$ of the function values at the particle points, minus the means of those functions at the points. Columns of $\tilde{\mathbf{Y}}$ are given by

$$\tilde{\mathbf{y}}^k = \begin{bmatrix} f(\mathbf{x}_0^k) - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_0^i) \\ f(\mathbf{x}_1^k) - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_1^i) \\ \vdots \\ f(\mathbf{x}_{M-1}^k) - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_{M-1}^i) \end{bmatrix}. \quad (10.37)$$

The new cost function \tilde{G} is the estimation of entropy of the samples $\tilde{\mathbf{y}}^k$. With the same assumption of a Gaussian distribution in shape space, by the same logic as for the derivation of the cost function G in (10.29), we have

$$\tilde{G}(\tilde{\mathbf{z}}) = \log \left| c \tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}} \right|, \quad (10.38)$$

with c a constant.

Let $\mathbf{Q} = (\tilde{\mathbf{Y}}^T \tilde{\mathbf{Y}} + \alpha \mathbf{I})^{-1}$. By the chain rule, the partial derivative of \tilde{G} with respect to the data \mathbf{y}^k becomes

$$-\frac{\partial \tilde{G}}{\partial \tilde{\mathbf{p}}^k} = \mathbf{J}_k^T \mathbf{Q}^k, \quad (10.39)$$

where \mathbf{J}_k is the Jacobian of the functional data for shape k . The matrix \mathbf{J}_k has the structure of a block diagonal matrix with $M \times M$ blocks, with diagonal blocks the $q \times d$ submatrices of the function gradients at particle j . Specifically, for each shape k , we have function data

$$\mathbf{y}^k = \left[f_1^0, f_2^0, \dots, f_q^0, f_1^1, f_2^1, \dots, f_q^1, \dots, f_1^{M-1}, f_2^{M-1}, \dots, f_q^{M-1} \right]^T, \quad (10.40)$$

and a diagonal submatrix block of the Jacobian $\mathbf{J}_k = \nabla_{\mathbf{z}^k} \mathbf{y}^k$ has the structure

$$\begin{bmatrix} \frac{\partial f_1^j}{\partial x_{dj+1}} & \frac{\partial f_1^j}{\partial x_{dj+2}} & \dots & \frac{\partial f_1^j}{\partial x_{dj+d}} \\ \frac{\partial f_2^j}{\partial x_{dj+1}} & \frac{\partial f_2^j}{\partial x_{dj+2}} & \dots & \frac{\partial f_2^j}{\partial x_{dj+d}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_q^j}{\partial x_{dj+1}} & \frac{\partial f_q^j}{\partial x_{dj+2}} & \dots & \frac{\partial f_q^j}{\partial x_{dj+d}} \end{bmatrix}, \quad (10.41)$$

where $j = \{0, 1, 2, \dots, M-1\}$ is the block number, which corresponds to a single particle, and $\{x_1, x_2, \dots, x_{dM}\}$ are the directional components of the full set of M particles. The correspondence optimization proceeds by gradient descent, as described in Section 10.2, with the substitution of the gradient of the new cost function \tilde{G} for the original cost function G in Eq. (10.32).

In summary, the generalized PBM algorithm replaces the entropy of positional information with entropy of an arbitrary function of positional information. This modification offers a much more generalized framework for optimizing the statistical properties of an ensemble of shapes. Note that the standard PBM algorithm from Section 10.2 is now just a special case of the generalized PBM algorithm, where $f(\mathbf{z}) = \mathbf{z}$.

10.3.4 Correspondence with Regression Against Explanatory Variables

In general, the design of a scientific study in biology or medicine cannot control for all confounding variables. The variability in shape due to such factors as age, differential growth rates, or clinical variables, for example, must be accounted for during the analysis phase. In other cases, this variability is the specific focus of the study, and researchers want to examine the correlation of an explanatory variable with shape. A typical experiment, for example, might examine the correlation of disease progression with the shape of anatomical structures or the change in the shape of anatomy with age. If such correlations can be established, they may lead to new diagnostic protocols or interventional planning.

This section extends the PBM algorithm to the problem of establishing correspondence in the presence of confounding variables and examining the correlation of shape with explanatory variables. Like in the previous section, this method allows for a more general notion of correspondence that takes into account additional information about

the data under study. The algorithm works by expanding the point-based correspondence model from Section 10.2 to include a regression against the independent variables. The optimization of correspondence position is then done on the *residual* to the regression model. Of course, the alternative is simply to use statistical methods in the analysis of post-optimization PBM shape parameters. However, the motivation to instead optimize the parameters of the residual model itself is the same principle of parsimony behind the basic PBM formulation: to further minimize model parameters and further maximize statistical power.

Under the assumption of a Gaussian distribution for the random variable \mathbf{Z} from Eq. (10.33), which is the distribution of shape samples in shape space, we can write the generative statistical model

$$\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}) \quad (10.42)$$

for particle correspondence positions, where $\boldsymbol{\epsilon}$ is normally-distributed error. Replacing $\boldsymbol{\mu}$ in this model with a function of an explanatory variable t gives the more general, regression model

$$\mathbf{z} = f(t) + \hat{\boldsymbol{\epsilon}}, \quad \hat{\boldsymbol{\epsilon}} \sim \mathcal{N}(\mathbf{0}, \hat{\boldsymbol{\Sigma}}). \quad (10.43)$$

The optimization described for the basic PBM algorithm minimizes the entropy associated with $\boldsymbol{\epsilon}$, which is the difference from the mean. In this section, the goal is to optimize correspondences under the regression model in Eq. (10.43) by instead minimizing entropy associated with $\hat{\boldsymbol{\epsilon}}$, the residual from the regression model. For the simple case where particle correspondence is a linear function of t , given as $f(t) = \mathbf{a} + \mathbf{b}t$, parameters \mathbf{a} and \mathbf{b} must be estimated to compute $\hat{\boldsymbol{\epsilon}}$. These parameters are estimated with a least-squares fit to the correspondence data,

$$\arg \min_{\mathbf{a}, \mathbf{b}} E(\mathbf{a}, \mathbf{b}) = \frac{1}{2} \sum_k [(\mathbf{a} + \mathbf{b}t_k) - \mathbf{z}_k]^T \boldsymbol{\Sigma}^{-1} [(\mathbf{a} + \mathbf{b}t_k) - \mathbf{z}_k]. \quad (10.44)$$

Setting $\frac{\delta E}{\delta \mathbf{a}} = \frac{\delta E}{\delta \mathbf{b}} = 0$ and solving for \mathbf{a} and \mathbf{b} , we have

$$\mathbf{a} = \frac{1}{n} \left(\sum_k \mathbf{z}_k - \sum_k \mathbf{b}t_k \right), \quad (10.45)$$

and

$$\mathbf{b} = \left(\sum_k t_k \mathbf{z}_k - \sum_k \mathbf{z}_k \sum_k t_k \right) / \left(\sum_k t_k^2 - (\sum_k t_k)^2 \right). \quad (10.46)$$

The proposed regression model optimization algorithm proceeds as follows. Correspondences are first optimized under the nonregression model (Eq. (10.42)) to minimize

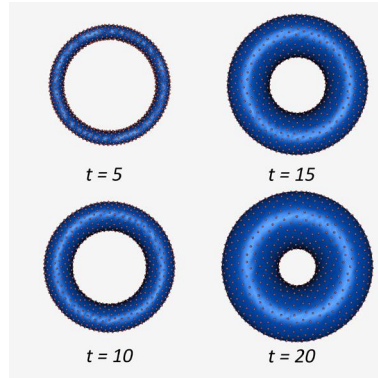


Figure 10.7 Shape regression on an explanatory variable t correlated with the small radius of a set of tori with $r \sim \mathcal{U}(5, 20)$.

the entropy associated with the total error ϵ . This process also establishes an initial estimate for \mathbf{a} and \mathbf{b} . The next step is to optimize under the regression model, which proceeds by gradient descent on $H(\mathbf{Z}) \approx \frac{1}{2} \log |\hat{\Sigma}| + H(P^k)$. In other words, the method follows the same optimization procedure as the basic PBM framework (Section 10.2), but replaces the covariance of the model with the covariance of the underlying residual, relative to the generative model. The two estimation problems are interwoven: the parameters \mathbf{a} and \mathbf{b} are re-estimated after each iteration of the gradient descent on the particle positions.

As an example of a correspondence optimization, consider the regression method applied to a set of $N = 40$ tori. To generate each torus, the large radius R was randomly drawn from a Gaussian distribution $R \sim \mathcal{N}(35, 3)$, and small radius r randomly drawn from a uniform distribution $r \sim \mathcal{U}(5, 20)$. An explanatory variable $t_i = r_i + \epsilon$, with $\epsilon \sim \mathcal{N}(0, .3)$, was assigned to each shape sample $i \in \{1, \dots, 40\}$ to establish a good correlation with variation in the small torus radius. Correspondences were optimized using 1024 particles per shape and the PBM regression algorithm outlined above. In the resulting correspondences, variation in the residuals to the regression line exhibits one major mode that empirically corresponds to r . Empirical observation of the regression line, which is shown in Fig. 10.7 suggests good correlation with R .

10.3.5 Dense PBM Correspondence Models

PBM yields relatively sparse correspondence models that may be inadequate to reconstruct thin structures and high curvature regions of the underlying anatomical surfaces. However, for many applications, we require a denser correspondence model, for example, to construct better surface meshes, make more detailed measurements, or conduct biomechanical or other simulations on mesh surfaces. One option for denser modeling




Figure 10.8 ShapeWorksView: pelvis surface reconstruction as a function of the number of particles.

is to increase the complexity of the PBM model via increasing its number of particles per shape sample. However, this approach necessarily increases the computational overhead, especially when modeling large clinical cohorts. Fig. 10.8 shows a pelvis surface reconstruction with an increased number of particles in which poor reconstructions are observed with smaller number of particles, especially along the pubic arch and iliac crest. Hence, an extension is needed to recover anatomically plausible and accurate 3D shapes at both the population and sample levels from a sparse set of particles.

In this extension, we adopt a template-deformation approach to establish an inter-sample dense surface correspondence, given a sparse set of optimized particles. To avoid introducing bias due to the template choice, we propose an unbiased framework for template mesh construction that includes three steps. First, generalized Procrustes alignment [72] is used to define the mean particle system from the PBM model, while estimating the rigid transformation that maps the sample-level particle system to the population-level counterpart. Second, the distance transform (DT) of each shape is deformed based on a nonlinear warping function that is built using the sample's particle system and the mean particle system as control points. Third, the warped DTs are then averaged to compute an average DT whose zero level set represents the geometry and topology of the mean shape in the population space. The dense *template mesh* is then constructed by triangulating the isosurface of this mean DT. This unbiased strategy will preserve the topology of the desired anatomy by taking into account the shape population of interest. In order to recover a sample-specific surface mesh, a warping function is constructed using the sample-level particle system and the mean/template particle system as control points. This warping function is then used to deform the template dense mesh to the sample space. A core ingredient of this PBM-based surface reconstruction is the form of the warping function being guided by the sparse particle system. One option is to use a thin-plate spline (TPS) [73] that defines a spatial mapping with global support, i.e., any perturbation in a single correspondence affects the whole warping function. Fig. 10.9 shows a sample PBM-based pelvis surface reconstruction using TPS with 512 particles. Compared to the current ShapeWorksView, our proposed reconstruction is able to recover the pubic arc and iliac crest with few particles. Nonetheless,




Figure 10.9 Surface reconstruction of a pelvis mean shape with 512 particles using (A) ShapeWorksView and (B) the proposed warping-based reconstruction (with thin plate splines).




Figure 10.10 TPS vs. RBF: histogram of vertex-wise average distance (in mm) between original patient-specific femur meshes and the warped template femur mesh.

TPS poses timing and memory challenges to process PBM models with more particles. To speed up the reconstruction process while minimizing the memory footprint, we propose the use of compactly supported radial basis functions (RBF) [74], which results in a sparse matrix that can be solved using sparse solvers. For preliminary results, we used a cohort of 70 femur shapes with 0.7 mm resolution. The computation time for template mesh construction was reduced from 4 hours using TPS to 9 minutes using RBF. Sequentially deforming the template mesh to the space of each subject was reduced from 4.5 hours using TPS to 5 minutes using RBF. To ensure that we are not sacrificing accuracy for this speed-up, Fig. 10.10 shows histograms of the warping error (in mm) between the groundtruth and reconstructed patients meshes in which both TPS and RBF attain similar trends and the majority of error is less than a seventh of voxel size. With a femur cohort at a resolution of 0.24 mm, Fig. 10.11 shows the RBF-based warping error being quantified at different numbers of particles where *less than half voxel size* warping error is achieved with 1K particles or more. This PBM-based warping can be further used to deform clinical measurements from patient to population space and vice versa [63].




Figure 10.11 Mean and standard deviation of RBF-based warping error in mm as a function of number of particles. Vertex-wise mean error along all the 70 femur shapes was computed and visualized as colormaps on the template dense mesh.

10.4 SHAPEWORKS SOFTWARE IMPLEMENTATION AND WORKFLOW

ShapeWorks is an open-source software distribution of the PBM approach to SSM. The software includes a careful numerical implementation of the approaches developed in Sections 10.2–10.3 in a modular C++ library. ShapeWorks also includes a suite of software applications for applying PBM to image segmentations. The software distribution consists of a set of command line tools for preprocessing binary segmentations (ShapeWorksGroom) and computing landmark-based shape models (ShapeWorksRun). It also includes a simple user interface to analyze and visualize the optimized shape models (ShapeWorksView). The current shape modeling pipeline for establishing shape correspondences from a set of binary segmentation image volumes is outlined in Fig. 10.12, with reference to the software tools that implement each step. In addition to the command line tools and visualization software, ShapeWorks also includes a full graphical user interface called ShapeWorksStudio, which can be used to interactively run all of the steps in the shape modeling pipeline. All ShapeWorks source code, binaries, and user documentation are freely available from the ShapeWorks website (www.sci.utah.edu/software/shapeworks.html) under the MIT license that is General Public License (GPL) compatible. According to Google Scholar, ShapeWork technology has received over 260 citations and over 2500 downloads since its initial release in 2009.

The remainder of this section describes each of the elements of the ShapeWorks software implementation in more detail, including the PBM code framework, an overview of the steps in the ShapeWorks SSM workflow, and an introduction to the ShapeWorks software applications.

10.4.1 The ShapeWorks Shape Modeling Workflow

A typical workflow for establishing shape correspondence from binary image volume inputs is outlined in Fig. 10.12. The preprocessing steps in the pipeline establish an initial alignment of the segmentations and generate suitable distance transforms. The

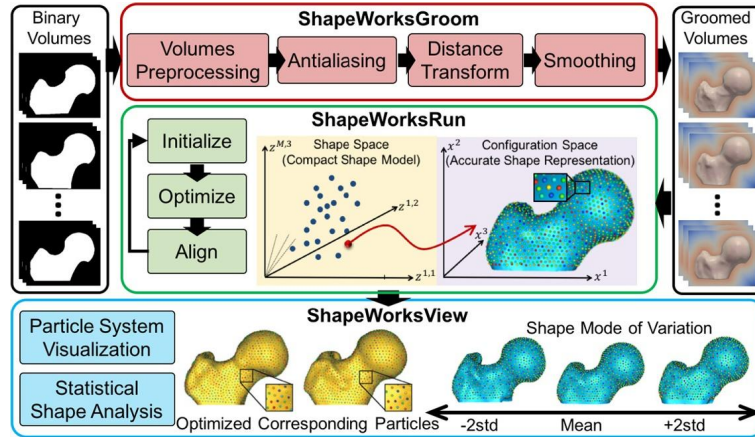


Figure 10.12 ShapeWorks shape modeling pipeline.

optimization phase consists of initializing the particle system, running the PBM optimization, and optionally refining the alignment using the Procrustes algorithm to iteratively remove residual nonshape data. Iterations of the correspondence optimization are interleaved with alignment steps until convergence. The remainder of this section discusses each of these steps in more detail.

10.4.1.1 Segmentation Preprocessing and Alignment

Any set of implicitly defined surfaces, such as a set of binary segmentations, is appropriate as input to the PBM algorithm. The algorithm, however, can be applied directly to binary segmentation volumes, which are often the output of a manual or automated segmentation process. Binary volumes contain an implicit shape surface at the interface of the labeled pixels and the background. Any suitably accurate distance transform from that interface may be used to form the implicit surface necessary for the particle optimization.

Segmentation data typically requires some processing to remove aliasing artifacts in the binary mask. Aliasing artifacts can adversely affect numerical approximations of surface features and the computations required to maintain the surface constraint in the PBM algorithm. One effective method for antialiasing binary volumes is given by Whitaker in [75], who describes a method for fitting an antialiased, level-set surface to a binary volume through an iterative relaxation process. The process uses curvature flow of the surface, with constraints on the flow dictated by the binary voxel locations of the segmentation. Another effective antialiasing method is the r -tightening algorithm given by Williams et al. [76]. The surface tightening method follows a similar approach to that of Whitaker, but constrains the level-set relaxation process using bi-

nary volumes that result from morphological opening and closing of the targeted binary surface. This method has proven to be particularly effective at removing aliasing artifacts without compromising the precision of the segmentation. As a final preprocessing step, the distance transform is typically followed by a slight Gaussian blurring to remove the high-frequency artifacts that can occur as a result of numerical approximations.

A collection of shape segmentations must often be aligned in a common coordinate frame for modeling and analysis. Where no information exists to specify a correct alignment, one approach is to first align segmentations with respect to their centers of mass and the orientation of their first principal eigenvectors. Then, during the optimization, the PBM method may optionally further align shapes with respect to rotation, translation, and scale using generalized Procrustes analysis (GPA) [77]. The GPA alignment is applied at regular intervals after particle updates in order to remove any residual, non-shape information from the model. GPA alignment during the optimization process is only enabled once the full set of M particles have been initialized on all surfaces. Where the true shape alignments are known, however, the GPA iterations may be omitted. A subset of the GPA alignment parameters may also be applied, such as only the rotational and translational components, leaving the scale unaffected.

10.4.1.2 Initialization and Optimization

There are number of possibilities for initializing the particle systems on the sample shapes, including manual specification of points and regular surface sampling. For spherical topologies, Paniagua et al. have proposed initialization of PBM with parametric SPHARM-PDM models [10]. One effective approach for general categories of shape, however, is to use an iterative, *particle splitting* strategy, in which the full set of particles is initialized in a multiscale fashion as follows. First, the PBM system is initialized with a single particle on each shape that finds the nearest zero of the implicit surface. This single particle is then split to produce a new, nearby particle. The two-particle (per shape) system is then optimized for correspondence until a steady state is reached. The splitting process, followed by optimization, is then repeated until a specific number of particles have been produced. Thus, the initialization proceeds simultaneously with the optimization in a multiscale fashion, generating progressively more detailed correspondence models with each split.

Typically, we set the numerical parameters for the PBM optimization automatically as follows. The numerical parameter σ_{\min} is set to machine precision and σ_{\max} is set to the size of the domain. The annealing parameter α starts with a value roughly equal to the diameter of an average shape and is reduced to machine precision over several hundred iterations. Particles are initialized on each shape using the splitting procedure described above. These default settings have been found to produce reliably good results that are very robust to the initialization, although some degree of parameter tuning is typically warranted in practice.

Processing time for the PBM algorithm on a modern desktop computer averages around 1/16,000 second/particle per iteration. This translates to full optimization times that scale linearly with the number of particles in the system and are on the order of minutes for small systems of a few thousand particles to several hours for larger systems of tens-of-thousands of particles. Optimizations of very large systems of hundreds-of-thousands to millions of particles may take processing times of several dozen hours.

10.4.1.3 Analysis

For analysis, sets of configurations are usually aligned within a common d -dimensional coordinate frame by a rotation, translation, and scaling to remove the geometric information unrelated to shape variation. Goodall's model of shape [77,78] describes each of these nonshape components and the residual variation around the mean correspondence configuration. For configuration matrix C_i , the model is given by

$$\mathbf{C}_i = a_i \boldsymbol{\mu} + \mathbf{E}_i \mathbf{R}_i + \mathbf{1} \mathbf{t}_i, \quad (10.47)$$

where a_i is a scalar representing the relative size of specimen i relative to the mean size, \mathbf{E}_i are the residuals from the mean configuration $\boldsymbol{\mu}$, \mathbf{R}_i is a rotation matrix describing the orientation of sample i , $\mathbf{1}$ is a d -dimensional vector of 1s, and \mathbf{t}_i is a translation vector describing the locational information for sample i .

The most common method for estimating $\boldsymbol{\mu}$, and the nonshape components \mathbf{R}_i , \mathbf{t}_i , and a_i is generalized Procrustes analysis (GPA) [7,72,79,5]. When transformed using GPA, correspondences are said to be in *Procrustes space*. Statistical analysis is commonly done in Procrustes space because, for reasonably similar sets of shapes, distance measures between Procrustes coordinates have been shown to be good linear approximations to the geodesic distances in Kendall's shape space [5,6].

The remaining geometric variability of a correspondence model after GPA is often summarized as a set of shape parameters that are the orthogonal directions of a principal components analysis (PCA) of the correspondence point positions. A complete mathematical description of this process can be found in, for example, [77,80], with application to PBM in [81]. PCA-based shape parameters allow us to compress the very large amount of geometric information into a much smaller representation of shape that is suitable for traditional statistics, while still retaining most of the geometric information of the shapes. Typically, we choose a finite number of shape parameters m for analysis either empirically, or by picking a set that accounts for most of the variability in the model. A more objective approach for PCA model selection is to use a method called parallel analysis to automatically determine a finite set of PCA modes that are distinguishable from Gaussian noise in the model [82].

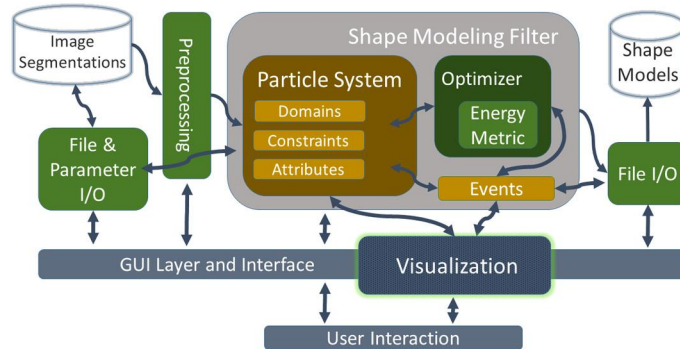


Figure 10.13 Abstract implementation of PBM shape modeling.

10.4.2 The PBM Code Library

A numerical implementation of the PBM algorithms from Sections 10.2–10.3 is available as an open source library of C++ code (github.com/joshcates/ITKParticleShapeModeling). This library is part of the ShapeWorks distribution and is used in all ShapeWorks applications. The PBM code library is built using the Insight Toolkit (ITK, www.itk.org) and can be easily compiled with ITK as an external module. The PBM codebase also conforms to ITK coding standards and includes Doxygen-based documentation (www.doxygen.org) and unit regression tests for all C++ classes. The PBM code library is templated on dimensionality, so that the methods operate equally for two- and three-dimensional image segmentation data.

Fig. 10.13 depicts an abstract view of the main code modules in the PBM library and how they are combined to implement a software application for shape analysis, including the data flow among these components. With reference to Fig. 10.13, the PBM library includes ITK filters for preprocessing image segmentations and converting them to suitable distance transform inputs (see Section 10.4.1.1). Distance transforms are input into a PBM Shape Modeling Filter object, which manages the construction of the multiple Particle Systems and executes the optimization process using a suitable Optimizer and Energy Metric. The main output of the Shape Modeling Filter is the set of correspondence point positions for all N input shapes. The Particle System data container is central to the processing. It is implemented as a facade class [83] that stores and manipulates all of the point-based representations of the input implicit shape surfaces, their local coordinate domains, and the mappings between those domains. The Particle System class, along with most other objects that maintain state in the PBM framework, can make use of ITK’s command/observer framework to allow state changes to trigger attached processes at the application level. For example, visualization code can be written to listen for particle position changes that are broadcast from the Particle System object and then trigger graphical updates in response.

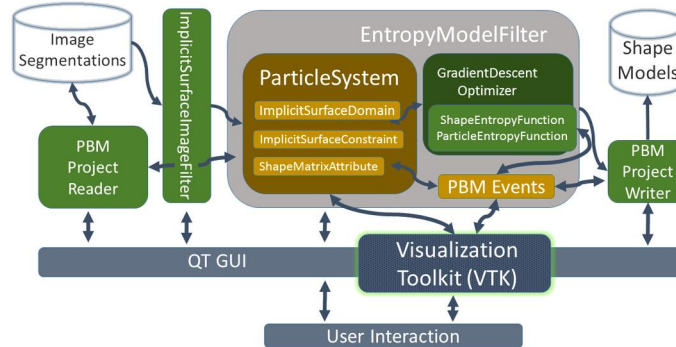


Figure 10.14 Concrete implementation of PBM shape modeling.

Fig. 10.14 shows the specific instantiation of the abstract framework from Fig. 10.13 that implements a basic PBM correspondence optimization. The generic class types have been replaced with specific instantiations in each box. Alternative PBM modeling approaches can be implemented by substituting alternative implementations of the various modular components in this figure. For example, the PBM regression optimization described in Section 10.3.4 can be implemented by simply substituting the “RegressionShapeMatrixAttribute” for the more general “ShapeMatrixAttribute” class in Fig. 10.14. For more details on the PBM code library and its use, the reader is referred to the documentation included with the ShapeWorks distribution.

10.4.3 The ShapeWorks Software Tool Suite

With reference to Fig. 10.12, the ShapeWorks software suite includes the following applications, which collectively implement the complete workflow described in Section 10.4.1. **ShapeWorksGroom** is a command line tool for batch processing binary segmentations, as described in Section 10.4.1.1. It can be used to perform simple alignment of segmentations, basic quality control, and to generate appropriate distance transform inputs for the PBM optimization. **ShapeWorksRun** implements the PBM correspondence optimization algorithm described in Section 10.2, including multiscale initialization via particle splitting and iterative GPA (Section 10.4.1.2). ShapeWorksRun also includes the extensions to PBM described in Sections 10.3.1–10.3.4. **ShapeWorksView** is an application for visualizing the results of the PBM optimization. It is built using the open-source Visualization Toolkit (VTK, www.vtk.org) and includes visualization of correspondence positions on shape surfaces, reconstructions of the mean shape of the correspondence model, and reconstructions of shapes along PCA modes (Section 10.4.1.3) and regression lines (Section 10.3.4).

ShapeWorks Studio is a new desktop application that encapsulates the entire processing pipeline from Fig. 10.12 under a single user interface. The different steps of the

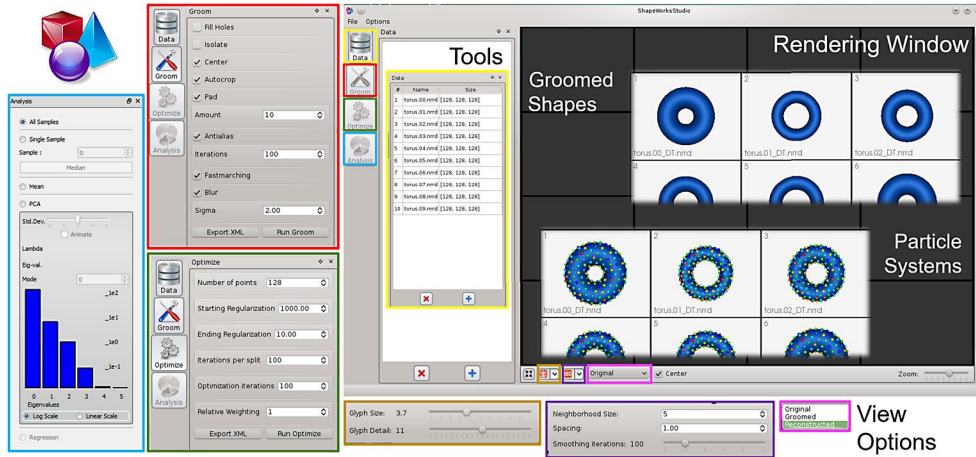


Figure 10.15 The ShapeWorks Studio desktop application.

pipeline are organized under separate tabs, as shown in Fig. 10.15. ShapeWorks Studio also includes postprocessing to produce dense correspondence models, as described in Section 10.3.5. ShapeWorks Studio is built using the PBM C++ library, ITK, Qt (www.qt.io) for graphical user interface elements, and VTK for visualization components. The design of the software and its dataflow are shown in Fig. 10.14.

10.5 SHAPEWORKS IN BIOMEDICAL APPLICATIONS

ShapeWorks has been applied to a wide variety of problems in medicine and biology, including neurobiology, biological phenotyping, orthopedics, and cardiology. In this section, we review some of the applications in these areas.

The PBM algorithm and many of its extensions were originally developed in the context of neurobiology. The method of modeling ensembles of shape complexes (Section 10.3.2), for example, was developed with application to a study of the shape of subcortical brain structures in pediatric autism. In that study, we identified statistically significant differences between the joint mean shape of 10 subcortical structures in autism patients and those of normal controls [60]. The PBM regression modeling approach outlined in Section 10.3.4 was used to describe a longitudinal model for neonatal head shape development [58]. Oguz et al. have applied the generalized method of PBM correspondence (Section 10.3.3) for human cortical surface correspondence [56,57], and have shown improved correspondence with respect to cortical thickness and sulcal depth over more commonly used approaches such as FreeSurfer (<http://www.freesurfer.net>).

Gene targeting is one of the most important tools for genetic study, and is widely used to examine the role that specific genes play in human development and disease.

Because gene targeting studies often rely on metrics of shape to quantify phenotypic expression, more comprehensive and detailed representations of shape from SSM allow for observations of genetic expression that have not been possible with traditional morphometrics. We applied ShapeWorks to quantify phenotype of the forepaw in mice deficient in the *Hoxd11* gene and compare to normal control mice. *Hoxd11* is known to play a role in the normal patterning of the appendicular skeleton. Our results showed both significant gross shape changes associated with bone length and thickness, but also more subtle local shape abnormalities at the distal end of the bones [60].

In orthopedics, ShapeWorks has been applied to the quantification of the spectrum of hip deformities associated with cam-type femoroacetabular impingement (FAI). We have analyzed variation in the thickness of the femur cortex between asymptomatic controls and cam-FAI patients, for whom we had quantitative evidence that repetitive impingement would induce bone hypertrophy [62,84,63]. We have further used ShapeWorks to study the resection of the lesion in cam-FAI patients to reduce the likelihood of damage to chondrolabral tissue and have developed guidelines to assist surgeons with resection of cam lesions as an effective intraoperative guide [85]. We established the limitations of radiographic measurements used in the clinical diagnosis of cam-FAI for which shape modeling greatly promise in deriving measurements to best describe cam-FAI deformities [86]. ShapeWorks has further helped in developing cost-effective patient-specific finite element (FE) models (which require hundreds of man-hours otherwise) of the cartilage and labrum to advance our understanding of contact mechanics and the pathogenesis of osteoarthritis [87].

In cardiology, ShapeWorks has been applied to study the maladaptive remodeling of the left atrium (LA) and clinical outcomes in atrial fibrillation (AF). PBM is uniquely suited for modeling the LA (and other heart structures) because its particle-system formulation can accommodate holes in a surface, such as the mitral valve openings and pulmonary vein openings. LA shape and size changes have previously been associated with AF progression and decreased response to catheter ablation, a first-line therapy for symptomatic drug refractory AF [88–91]. Using ShapeWorks models of the LA, we described the statistically significant differences in LA shape in populations with increasing severity of the disease [64]. In that same study, we also identified characteristic shape changes of the left-atrial appendage that are associated with an increased likelihood of thrombus, a major risk factor for cardioembolic stroke. These results suggest the possibility of a shape-based clinical measure for stroke risk, as well as other AF outcomes. In another study, we applied the PBM correspondence models in the LA to develop a population atlas of the distribution of fibrosis in AF sufferers [65]. Fibrosis is a well-known degradation of the tissue structure in AF that can be measured using late-gadolinium enhancement MRI [92].

10.6 CONCLUSIONS AND FUTURE WORK

Statistical shape modeling (SSM) represents a revolutionary new approach to morphometry, but the widespread adoption of these new tools will require careful engineering to make them accessible to the average user and applicable to a wide variety of biological shapes. ShapeWorks is an implementation of Particle Based Modeling (PBM), which is designed to be a general tool for SSM. By modeling surfaces nonparametrically as collections of dynamic particle systems, ShapeWorks is not limited to any particular topology and can even represent shapes with multiple disconnected surfaces and surfaces with arbitrary open boundaries. The particle system approach also avoids the algorithmic complexity and parameter tuning associated with constructing parameterizations.

The ShapeWorks software suite implements a general SSM workflow that supports interactive processing on smaller datasets and offline processing of very large cohorts. ShapeWorks also includes preprocessing and analysis tools and, for many applications, can be a complete end-to-end SSM solution. The PBM C++ code library represents a careful numerical implementation of the particle system framework, and is extensible and reusable in any application supporting ITK. The ShapeWorks approach has been used successfully for investigation in many areas of biomedical investigation, including neurobiology, genetic phenotyping, orthopedics, and cardiology.

Despite proven success for many biomedical applications, the rapid advances and growing use of medical imaging technologies, and the associated need to model more complex morphological variations, require significant functionality and usability improvements of all modern SSM approaches, including ShapeWorks. In order to meet today's real-world shape modeling problems from a diverse biomedical community, we are developing new algorithmic extensions and new software versions. For example, motivated by studies of the highly variable anatomy of the atria of the heart, we are expanding our PBM approach to allow mixtures of Gaussians, in order to find natural clusterings of shapes in the populations and better discrimination when seeking shape-based indicators of treatment outcome. To support ever increasing cohort sizes and reduce ShapeWorks' run-times and memory requirements, we are developing parallel versions of the PBM optimization and modifying our algorithms to operate directly on mesh-based representations of surfaces. Finally, more comprehensive user interfaces for our software and domain-specific customizations of the PBM workflows are in progress, in order to make the PBM approach more accessible to the larger morphometrics community.

REFERENCES

- [1] D.G. Kendall, The diffusion of shape, *Adv. Appl. Probab.* 9 (1977) 428–430.
- [2] D.G. Kendall, Shape-manifolds, procrustean metrics and complex projective spaces, *Bull. Lond. Math. Soc.* 16 (1984) 81–121.

- [3] J.T. Kent, The complex Bingham distribution and shape analysis, *J. R. Stat. Soc. B* 56 (1994) 285–299.
- [4] F.L. Bookstein, Biometrics, biomathematics, and the morphometric synthesis, *Bull. Math. Biol.* 58 (1996) 313–365.
- [5] Ian Dryden, Kanti Mardia, *Statistical Shape Analysis*, John Wiley and Sons, 1998.
- [6] D.G. Kendall, *Shape and Shape Theory*, Wiley, 1999.
- [7] C.P. Klingenberg, Morphometrics and the role of the phenotype in studies of the evolution of developmental mechanisms, *Gene* 287 (2002) 3–10.
- [8] D. Adams, F.J. Rolf, D. Slice, Geometric morphometrics: ten years of progress following the “revolution”, *Ital. J. Zool.* 71 (2004) 5–16.
- [9] Tobias Heimann, Hans-Peter Meinzer, Statistical shape models for 3D medical image segmentation: a review, *Med. Image Anal.* 13 (4) (2009) 543–563.
- [10] Martin Styner, Ipek Oguz, Shun Xu, Christian Brechbühler, Dimitrios Pantazis, James J. Levitt, Martha E. Shenton, Guido Gerig, Framework for the statistical shape analysis of brain structures using SPHARM-PDM, *Insight J.* 1071 (2006) 242.
- [11] Martin Styner, Jeffrey A. Lieberman, Dimitrios Pantazis, Guido Gerig, Boundary and medial shape analysis of the hippocampus in schizophrenia, *Med. Image Anal.* 8 (3) (2004) 197–203.
- [12] Amir A. Zadpoor, Harrie Weinans, Patient-specific bone modeling and analysis: the role of integration and automation in clinical adoption, *J. Biomech.* 48 (5) (2015) 750–760.
- [13] Daniel P. Nicoletta, Todd L. Bredbenner, Development of a parametric finite element model of the proximal femur using statistical shape and density modelling, *Comput. Methods Biomech. Biomed. Eng.* 15 (2) (2012) 101–110.
- [14] Rebecca Bryan, P. Surya Mohan, Andrew Hopkins, Francis Galloway, Mark Taylor, Prasanth B. Nair, Statistical modelling of the whole human femur incorporating geometric and material properties, *Med. Eng. Phys.* 32 (1) (2010) 57–65.
- [15] Laura Belenguer Querol, Philippe Büchler, Daniel Rueckert, Lutz P. Nolte, Miguel Á. González Ballester, Statistical finite element model for bone shape and biomechanical properties, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, Springer, 2006, pp. 405–411.
- [16] Nina Kozic, Stefan Weber, Philippe Büchler, Christian Lutz, Nils Reimers, Miguel Á. González Ballester, Mauricio Reyes, Optimisation of orthopaedic implant design using statistical shape space analysis based on level sets, *Med. Image Anal.* 14 (3) (2010) 265–275.
- [17] Heiko Seim, Dagmar Kainmueller, Hans Lamecker, Matthias Bindernagel, Jana Malinowski, Stefan Zachow, Model-based auto-segmentation of knee bones and cartilage in MRI data, in: *Medical Image Analysis for the Clinic: A Grand Challenge*, Beijing, 2010.
- [18] Hans Lamecker, Thomas H. Wenckeback, Hans-Christian Hege, Atlas-based 3D-shape reconstruction from X-ray images, in: *18th International Conference on Pattern Recognition, ICPR 2006*, vol. 1, IEEE, 2006, pp. 371–374.
- [19] Primoz Markelj, Dejan Tomažević, Bostjan Likar, Franjo Pernuš, A review of 3D/2D registration methods for image-guided interventions, *Med. Image Anal.* 16 (3) (2012) 642–661.
- [20] Guoyan Zheng, Sebastian Gollmer, Steffen Schumann, Xiao Dong, Thomas Feilkas, Miguel A. González Ballester, A 2D/3D correspondence building method for reconstruction of a patient-specific 3D bone surface model using point distribution models and calibrated X-ray images, *Med. Image Anal.* 13 (6) (2009) 883–899.
- [21] Nora Baka, B.L. Kaptein, Marleen de Bruijne, Theo van Walsum, J.E. Giphart, Wiro J. Niessen, Boudewijn P.F. Lelieveldt, 2D–3D shape reconstruction of the distal femur from stereo X-ray imaging using statistical shape models, *Med. Image Anal.* 15 (6) (2011) 840–850.
- [22] Guoyan Zheng, Xiao Dong, Kumar T. Rajamani, Xuan Zhang, Martin Styner, Ramesh U. Thoranaghatte, Lutz-Peter Nolte, Miguel A. González Ballester, Accurate and robust reconstruction of a surface model of the proximal femur from sparse-point data and a dense-point distribution model for surgical navigation, *IEEE Trans. Biomed. Eng.* 54 (12) (2007) 2109–2122.

- [23] Jalda Dworzak, Hans Lamecker, Jens von Berg, Tobias Klinder, Cristian Lorenz, Dagmar Kainmüller, Heiko Seim, Hans-Christian Hege, Stefan Zachow, 3D reconstruction of the human rib cage from 2D projection images using a statistical shape model, *Int. J. Comput. Assisted Radiol. Surg.* 5 (2) (2010) 111–124.
- [24] Guoyan Zheng, Steffen Schumann, Steven Balestra, Benedikt Thelen, Lutz-P. Nolte, 2D–3D reconstruction-based planning of total hip arthroplasty, in: *Computational Radiology for Orthopaedic Interventions*, Springer, 2016, pp. 197–215.
- [25] Moritz Ehlke, Heiko Ramm, Hans Lamecker, Hans-Christian Hege, Stefan Zachow, Fast generation of virtual X-ray images for reconstruction of 3D anatomy, *IEEE Trans. Vis. Comput. Graph.* 19 (12) (2013) 2673–2682.
- [26] Aurélie Carlier, Liesbet Geris, Johan Lammens, Hans Van Oosterwyck, Bringing computational models of bone regeneration to the clinic, *Wiley Interdiscip. Rev., Syst. Biol. Med.* 7 (4) (2015) 183–194.
- [27] P. Sami Väänänen, Lorenzo Grassi, Gunnar Flivik, Jukka S. Jurvelin, Hanna Isaksson, Generation of 3D shape, density, cortical thickness and finite element mesh of proximal femur from a DXA image, *Med. Image Anal.* 24 (1) (2015) 125–134.
- [28] B. Reggiani, L. Cristofolini, E. Varini, M. Viceconti, Predicting the subject-specific primary stability of cementless implants during pre-operative planning: preliminary validation of subject-specific finite-element models, *J. Biomech.* 40 (11) (2007) 2552–2558.
- [29] C. Dean Barratt, Carolyn S.K. Chan, Philip J. Edwards, Graeme P. Penney, Mike Slomczykowski, Timothy J. Carter, David J. Hawkes, Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging, *Med. Image Anal.* 12 (3) (2008) 358–374.
- [30] Lawrence M. Specht, Kenneth J. Koval, Robotics and computer-assisted orthopaedic surgery, *Bull. Hosp. Joint Dis. Orthop. Inst.* 60 (3–4) (2001) 168–172.
- [31] T. Kumar Rajamani, Martin A. Styner, Haydar Talib, Guoyan Zheng, Lutz P. Nolte, Miguel A. González Ballester, Statistical deformable bone models for robust 3D surface extrapolation from sparse data, *Med. Image Anal.* 11 (2) (2007) 99–109.
- [32] Stefan Zachow, Hans Lamecker, Barbara Elsholtz, Michael Stiller, Reconstruction of mandibular dysplasia using a statistical 3D shape model, in: *International Congress Series*, vol. 1281, Elsevier, 2005, pp. 1238–1243.
- [33] Stefan Zachow, Hans-Christian Hege, Peter Deuffhard, Computer assisted planning in cranio-maxillofacial surgery, *CIT, J. Comput. Inf. Technol.* 14 (1) (2006) 53–64.
- [34] Li Wang, Yi Ren, Yaozong Gao, Zhen Tang, Ken-Chung Chen, Jianfu Li, Steve G.F. Shen, Jin Yan, Philip K.M. Lee, Ben Chow, et al., Estimating patient-specific and anatomically correct reference model for craniomaxillofacial deformity via sparse representation, *Med. Phys.* 42 (10) (2015) 5809–5816.
- [35] Kun Zhang, Wee Kheng Leow, Yuan Cheng, Performance analysis of active shape reconstruction of fractured, incomplete skulls, in: *Computer Analysis of Images and Patterns*, Springer, 2015, pp. 312–324.
- [36] H. Lamecker, S. Zachow, H. Hege, H. Haberl, Surgical treatment of craniosynostosis based on a statistical 3D-shape model: first clinical application, *Int. J. Comput. Assisted Radiol. Surg.* 1 (10) (2006) 253.
- [37] Carlos S. Mendoza, Nabile Safdar, Kazunori Okada, Emmarie Myers, Gary F. Rogers, Marius George Linguraru, Personalized assessment of craniosynostosis via statistical shape modeling, *Med. Image Anal.* 18 (4) (2014) 635–646.
- [38] Mascha Hochfeld, Hans Lamecker, Ulrich-W. Thomale, Matthias Schulz, Stefan Zachow, Hannes Haberl, Frame-based cranial reconstruction: technical note, *J. Neurosurg. Pediatrics* 13 (3) (2014) 319–323.
- [39] Stefan Zachow, Computational planning in facial surgery, *Facial Plast. Surg.* 31 (5) (2015) 446.

- [40] Francis Galloway, Max Kahnt, Heiko Ramm, Peter Worsley, Stefan Zachow, Prasanth Nair, Mark Taylor, A large scale finite element study of a cementless osseointegrated tibial tray, *J. Biomech.* 46 (11) (2013) 1900–1906.
- [41] Jeffrey E. Bischoff, Yifei Dai, Casey Goodlett, Brad Davis, Marc Bandi, Incorporating population-level variability in orthopedic biomechanical analysis: a review, *J. Biomech. Eng.* 136 (2) (2014) 021004.
- [42] Rebecca Bryan, Prasanth B. Nair, Mark Taylor, Use of a statistical model of the whole femur in a large scale, multi-model study of femoral neck fracture risk, *J. Biomech.* 42 (13) (2009) 2171–2176.
- [43] C. Merle, W. Waldstein, J.S. Gregory, S.R. Goodyear, R.M. Aspden, P.R. Aldinger, D.W. Murray, H.S. Gill, How many different types of femora are there in primary hip osteoarthritis? An active shape modeling study, *J. Orthop. Res.* 32 (3) (2014) 413–422.
- [44] Ch. Brechbühler, Guido Gerig, Olaf Kübler, Parametrization of closed surfaces for 3-D shape description, *Comput. Vis. Image Underst.* 61 (2) (1995) 154–170.
- [45] Andrew P. Witkin, Paul S. Heckbert, Using particles to sample and control implicit surfaces, in: *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press, 1994, pp. 269–277.
- [46] Miriah D. Meyer, Pierre Georgel, Ross T. Whitaker, Robust particle systems for curvature dependent sampling of implicit surfaces, in: *Proceedings of the International Conference on Shape Modeling and Applications*, June 2005, pp. 124–133.
- [47] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, W. Stuetzle, Surface reconstruction from unorganized points, in: *ACM SIGGRAPH 1992 Conference Proceedings*, 1992, pp. 71–78.
- [48] Mark Hansen, Bin Yu, Model selection and the principle of minimum description length, *J. Am. Stat. Assoc.* 96 (454) (June 2001) 746–774.
- [49] Rhodri H. Davies, Carole J. Twining, Timothy F. Cootes, John C. Waterton, Christopher J. Taylor, A minimum description length approach to statistical shape modeling, *IEEE Trans. Med. Imaging* 21 (5) (2002) 525–537.
- [50] Rhodri H. Davies, Carole J. Twining, Timothy F. Cootes, John C. Waterton, Christopher J. Taylor, 3D statistical shape models using direct optimisation of description length, in: *ECCV* (3), 2002, pp. 3–20.
- [51] Joshua Cates, Miriah Meyer, Thomas Fletcher, Ross Whitaker, et al., Entropy-based particle systems for shape correspondence, in: *1st MICCAI Workshop on Mathematical Foundations of Computational Anatomy: Geometrical, Statistical and Registration Methods for Modeling Biological Shape Variability*, 2006, pp. 90–99.
- [52] Joshua Cates, P. Thomas Fletcher, Martin Styner, Martha Shenton, Ross Whitaker, Shape modeling and analysis with entropy-based particle systems, in: *Information Processing in Medical Imaging*, Springer, Berlin, Heidelberg, 2007, pp. 333–345.
- [53] Joshua Cates, P. Thomas Fletcher, Martin Styner, Heather Cody Hazlett, Ross Whitaker, Particle-based shape analysis of multi-object complexes, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2008*, Springer, Berlin, Heidelberg, 2008, pp. 477–485.
- [54] Manasi Datar, Yaniv Gur, Beatriz Paniagua, Martin Styner, Ross Whitaker, Geometric correspondence for ensembles of nonregular shapes, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011*, Springer, Berlin, Heidelberg, 2011, pp. 368–375.
- [55] M. Datar, P. Muralidharan, A. Kumar, S. Gouttard, J. Piven, G. Gerig, R.T. Whitaker, P.T. Fletcher, Mixed-effects shape models for estimating longitudinal changes in anatomy, in: *Stanley Durrleman, P. Thomas Fletcher, Guido Gerig, Marc Niethammer (Eds.), Spatio-Temporal Image Analysis for Longitudinal and Time-Series Image Data*, in: *Lect. Notes Comput. Sci.*, vol. 7570, Springer, Berlin/Heidelberg, 2012, pp. 76–87.
- [56] Ipek Oguz, Joshua Cates, Thomas Fletcher, Ross Whitaker, Derek Cool, Stephen Aylward, Martin Styner, Cortical correspondence using entropy-based particle systems and local features, in: *5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, ISBI 2008*, IEEE, 2008, pp. 1637–1640.

- [57] Ipek Oguz, Marc Niethammer, Josh Cates, Ross Whitaker, Thomas Fletcher, Clement Vachet, Martin Styner, Cortical correspondence with probabilistic fiber connectivity, in: *Information Processing in Medical Imaging*, Springer, Berlin, Heidelberg, 2009, pp. 651–663.
- [58] Manasi Datar, Joshua Cates, P. Thomas Fletcher, Sylvain Gouttard, Guido Gerig, Ross Whitaker, Particle based shape regression of open surfaces with applications to developmental neuroimaging, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009*, Springer, Berlin, Heidelberg, 2009, pp. 167–174.
- [59] Manasi Datar, Ilwoo Lyu, SunHyung Kim, Joshua Cates, Martin A. Styner, Ross Whitaker, Geodesic distances to landmarks for dense correspondence on ensembles of complex shapes, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013*, Springer, 2013, pp. 19–26.
- [60] Joshua Cates, P. Thomas Fletcher, Zachary Warnock, Ross Whitaker, A shape analysis framework for small animal phenotyping with application to mice with a targeted disruption of *hoxd11*, in: *5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, ISBI 2008*, IEEE, 2008, pp. 512–515.
- [61] Kevin B. Jones, Manasi Datar, Sandhya Ravichandran, Huifeng Jin, Elizabeth Jurrus, Ross Whitaker, Mario R. Capecchi, Toward an understanding of the short bone phenotype associated with multiple osteochondromas, *J. Orthop. Res.* 31 (4) (2013) 651–657.
- [62] Michael D. Harris, Manasi Datar, Ross T. Whitaker, Elizabeth R. Jurrus, Christopher L. Peters, Andrew E. Anderson, Statistical shape modeling of cam femoroacetabular impingement, *J. Orthop. Res.* (2013).
- [63] Penny B. Atkins, Shireen Elhabian, Praful Agrawal, Mike Harris, Jeffery Weiss, Chris Peters, Ross Whitaker, Andrew Anderson, Quantitative comparison of cortical bone thickness using correspondence-based shape modeling in patients with cam femoroacetabular impingement, *J. Orthop. Res.* (2016).
- [64] Joshua Cates, Erik Bieging, Alan Morris, Gregory Gardner, Nazem Akoum, Eugene Kholmovski, Nassir Marrouche, Christopher McGann, Rob S. MacLeod, Computational shape models characterize shape change of the left atrium in atrial fibrillation, *Clin. Med. Insights, Cardiol.* 8 (Suppl. 1) (2014) 99.
- [65] Gregory Gardner, Alan Morris, Koji Higuchi, Robert MacLeod, Joshua Cates, A point-correspondence approach to describing the distribution of image features on anatomical surfaces, with application to atrial fibrillation, in: *2013 IEEE 10th International Symposium on Biomedical Imaging (ISBI)*, IEEE, 2013, pp. 226–229.
- [66] T. Cover, J. Thomas, *Elements of Information Theory*, Wiley and Sons, 1991.
- [67] Emanuel Parzen, On estimation of a probability density function and mode, *Ann. Math. Stat.* 33 (3) (1962) 1065–1076.
- [68] M. Meyer, B. Nelson, R.M. Kirby, R. Whitaker, Particle systems for efficient and accurate high-order finite element visualization, *IEEE Trans. Vis. Comput. Graph.* 13 (5) (2007) 1015–1026.
- [69] G. Kindlmann, R. Whitaker, T. Tasdizen, T. Moller, Curvature-based transfer functions for direct volume rendering, in: *Proceedings of IEEE Visualization*, 2003, pp. 512–520.
- [70] Sarang Joshi, Brad Davis, Matthieu Jomier, Guido Gerig, Unbiased diffeomorphic atlas construction for computational anatomy, *NeuroImage* 23 (2004) S151–S160.
- [71] SCI Institute, AtlasWerks: an open-source (BSD license) software package for medical image atlas generation, Scientific Computing and Imaging Institute (SCI), 2016, available from <http://www.sci.utah.edu/software/atlaswerks.html>.
- [72] John C. Gower, Generalized Procrustes analysis, *Psychometrika* 40 (1) (1975) 33–51.
- [73] Fred L. Bookstein, Principal warps: thin-plate splines and the decomposition of deformations, *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (6) (1989) 567–585.
- [74] Shengxin Zhu, Compactly supported radial basis functions: how and why?, *SIAM Rev.* (2012).

- [75] Ross Whitaker, Reducing aliasing artifacts in iso-surfaces of binary volumes, in: IEEE Volume Visualization and Graphics Symposium, 2000, pp. 22–23.
- [76] J. Williams, J. Rossignac, Tightening: curvature-limiting morphological simplification, in: Proc. Ninth ACM Symposium on Solid and Physical Modeling, 2005, pp. 107–112.
- [77] C. Goodall, Procrustes methods in the statistical analysis of shape, *J. R. Stat. Soc. B* 53 (1991) 285–339.
- [78] F.J. Rohlf, Bias and error in estimates of mean shape in geometric morphometrics, *J. Hum. Evol.* 44 (2003) 665–683.
- [79] J.M.F. Ten Berge, Orthogonal Procrustes rotation for two or more matrices, *Psychometrika* 42 (1977) 267–276.
- [80] T.F. Cootes, C.J. Taylor, D.H. Cooper, J. Graham, Active shape models – their training and application, *Comput. Vis. Image Underst.* 61 (1) (January 1995) 38–59.
- [81] Joshua Cates, Thomas Fletcher, Ross Whitaker, et al., A hypothesis testing framework for high-dimensional shape models, in: 2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy, 2008, pp. 170–181.
- [82] L.W. Glorfeld, An improvement on Horn's parallel analysis methodology for selecting the correct number of factors to retain, *Educ. Psychol. Meas.* 55 (1995) 377–393.
- [83] Erich Gamma, *Design Patterns: Elements of Reusable Object-Oriented Software*, Pearson Education India, 1995.
- [84] Penny Atkins, Prateep Mukherjee, Shireen Elhabian, Sumedha Singla, Michael Harris, Jeffery Weiss, Ross Whitaker, Andrew Anderson, Proximal femoral cortical bone thickness in patients with femoroacetabular impingement and normal hips analyzed using statistical shape modeling, in: Summer Biomechanics, Bioengineering and Biotransport Conference, 2015, oral presentation.
- [85] Penny Atkins, Shireen Elhabian, Praful Agrawal, Ross Whitaker, Jeffery Weiss, Stephen Aoki, Chris Peters, Andrew Anderson, Can the sclerotic subchondral bone of the proximal femur cam lesion be used as a surgical resection guide? An objective analysis using 3D computed tomography and statistical shape modeling, in: International Society of Hip Arthroscopy Annual Scientific Meeting, 2016, submitted abstract.
- [86] Penny Atkins, Shireen Elhabian, Praful Agrawal, Ross Whitaker, Jeffery Weiss, Chris Peters, Stephen Aoki, Andrew Anderson, Which radiographic measurements best identify anatomical variation in femoral head anatomy? Analysis using 3D computed tomography and statistical shape modeling, in: International Society of Hip Arthroscopy Annual Scientific Meeting, 2016, submitted abstract.
- [87] Penny Atkins, Prateep Mukherjee, Shireen Elhabian, Sumedha Singla, Ross Whitaker, Jeffery Weiss, Andrew Anderson, Warping of template meshes for efficient subject-specific FE mesh generation, in: International Symposium of Computer Methods in Biomechanics and Biomedical Engineering, 2015, oral presentation.
- [88] Toshiya Kurotobi, Katsuomi Iwakura, Koichi Inoue, Ryusuke Kimura, Yuko Toyoshima, Norihisa Ito, Hiroya Mizuno, Yoshihisa Shimada, Kenshi Fujii, Shinsuke Nanto, et al., The significance of the shape of the left atrial roof as a novel index for determining the electrophysiological and structural characteristics in patients with atrial fibrillation, *Europace* 13 (6) (2011) 803–808.
- [89] Luigi Di Biase, Pasquale Santangeli, Matteo Anselmino, Prasant Mohanty, Ilaria Salvetti, Sebastiano Gili, Rodney Horton, Javier E. Sanchez, Rong Bai, Sanghamitra Mohanty, et al., Does the left atrial appendage morphology correlate with the risk of stroke in patients with atrial fibrillation? Results from a multicenter study, *J. Am. Coll. Cardiol.* 60 (6) (2012) 531–538.
- [90] Felipe Bisbal, Esther Guiu, Naiara Calvo, David Marin, Antonio Berruezo, Elena Arbelo, José Ortiz-Pérez, Teresa María Caralt, José María Tolosana, Roger Borràs, et al., Left atrial sphericity: a new method to assess atrial remodeling. Impact on the outcome of atrial fibrillation ablation, *J. Cardiovasc. Electrophysiol.* 24 (7) (2013) 752–759.
- [91] Teresa S.M. Tsang, Marion E. Barnes, Kent R. Bailey, Cynthia L. Leibson, Samantha C. Montgomery, Yasuhiko Takemoto, Pauline M. Diamond, Marisa A. Marra, Bernard J. Gersh, David O. Wiebers, et

- al., Left atrial volume: important risk marker of incident atrial fibrillation in 1655 older men and women, in: *Mayo Clin. Proc.*, vol. 76, Elsevier, 2001, pp. 467–475.
- [92] Christopher J. McGann, Eugene G. Kholmovski, Robert S. Oakes, Joshua J.E. Blauer, Marcos Daccarett, Nathan Segerson, Kelly J. Airey, Nazem Akoum, Eric Fish, Troy J. Badger, et al., New magnetic resonance imaging-based method for defining the extent of left atrial wall injury after the ablation of atrial fibrillation, *J. Am. Coll. Cardiol.* 52 (15) (2008) 1263–1271.