

Statistical Multi-object Shape Models

Conglin Lu, Stephen M. Pizer, Sarang Joshi, Ja-Yeon Jeong

Medical Image Display & Analysis Group
University of North Carolina, Chapel Hill

Abstract The shape of a population of geometric entities is characterized by both the common geometry of the population and the variability among instances. In the deformable model approach, it is described by a probabilistic model on the deformations that transform a common template into various instances. To capture shape features at various scale levels, we have been developing an object-based multi-scale framework, in which a geometric entity is represented by a series of deformations with different locations and degrees of locality. Each deformation describes a residue from the information provided by previous steps. In this paper, we present how to build statistical shape models of multi-object complexes with such properties based on medial representations and explain how this may lead to more effective shape descriptions as well as more efficient statistical training procedures. We illustrate these ideas with a statistical shape model for a pair of pubic bones and show some preliminary results on using it as a prior in medical image segmentation.

Keywords: Statistical shape models, multi-object shape models, deformable model, multi-scale shape analysis, medial representation

1 Introduction

In many shape analysis areas, including deformable models [1, 2], it is desirable to consider shape as a property of a *population* of geometric entities, such as object sections (e.g., the left lobe of livers), objects (e.g., livers), or groups of objects (e.g., liver-kidney complexes). As such, it describes both the common geometric features and the geometric variability among instances of the population. With the deformable model approach, a representative instance S_0 , usually referred to as the *model* or *template*, is chosen for a given shape, say \mathbf{S} . The geometry of S_0 is essentially specified by transformations on certain canonical geometric primitives, for instance, the positions of boundary/surface points or positions and boundary normal directions of medial structures. The overall transformation is a “summary” of the shape geometry and captures the typical geometric conformation of \mathbf{S} at various levels of detail, including the geometry of objects, parts, sections, and their relative configurations. Any instance of \mathbf{S} , say S , is then described via a further geometric deformation F applied to S_0 , i.e. $S = F(S_0)$. There is a natural correspondence between

the space of all such geometric transformations, \mathcal{F} , and the set of all geometric entities, \mathcal{G} , through the following map:

$$\begin{aligned}\psi : \mathcal{F} &\rightarrow \mathcal{G}, \\ F &\mapsto F(S_0).\end{aligned}$$

The shape \mathbf{S} can be regarded as a random quantity and be effectively described by a probability measure ν on \mathcal{G} . An instance of \mathbf{S} is then considered to be a sample from (\mathcal{G}, ν) . Intuitively, ν puts soft constraints on “what an instance of the shape \mathbf{S} should look like”. One way to define ν is to introduce a probability measure μ on \mathcal{F} and to let ν be the probability measure induced by the mapping ψ . This is the approach we will adopt in this paper. The measure μ can be effectively used as shape priors to make statistical inferences on what one observes [3, 4, 5, 6].

Existing geometric representations include representation by dense sample points, landmarks, diffeomorphisms on displacement velocity fields[7], distance functions or their level sets [8], skeletons [9], shocks [10], multi-grid and scale-space based methods [11], etc. Examples of statistical shape models are deformable templates [3], point distribution models [12, 13], spherical harmonic descriptors [14] for 3D objects, to name a few. Most of these methods do not respect object-based scale levels, so they lack the ability to explicitly describe important geometric information such as locality, inter-object relations, etc. The statistical models typically face the high-dimension-low-sample-size (HDLSS) problem: on the one hand, the number of parameters needed to accurately describe the geometry is usually very large; on the other hand, the available training samples are often limited, as is the case in many medical imaging applications.

Because geometric features have various degrees of locality, i.e., locally relevant size and distance, we argue that it is advantageous not to think of F as a single *global* deformation, but rather as a collection of deformations with various *localities* [15]. In our study of geometric entities such as objects and boundaries, locality must be taken with respect to the components of which an entity is formed: the relevant sizes and distances for an object complex are determined by various objects in it, while those for an object must be determined by the natural sections making up the object, and so on. Each feature summarizes the geometric information contained in a certain spatial domain of size relevant to a particular scale. The differences in description between scale levels reflect different levels of detail. Moreover, the relevant distance within a scale level induces a notion of neighbors, i.e., nearby geometric entities at that scale: nearby objects, nearby object sections, etc. This neighbor relation, together with the spatial extent of features and levels of detail in description, realize the notion of locality and form the basis of our representation. From a statistical point of view, locality reflects the spatial correlation between different geometric components of an entity. For many shapes, a particular feature at a certain scale level tends to be more correlated with the features that are close by, compared to those that are further away. This is an important assumption that we make in developing statistical shape models, an assumption based on measurements we have made in certain sample populations.

To effectively capture locality, we use object-based scale levels and describe *residues* within and between scales similar to the wavelet approaches [16, 17]. Rather than representing a shape instance \mathcal{S} by a single

transformation F , i.e., $S = F(S_0)$, we represent it in a coarse to fine manner by decomposing F into a series of, say K , deformations, i.e., $F = F_K \circ F_{K-1} \circ \dots \circ F_1$. This allows \mathcal{S} to be described by geometric entities at various scale levels and their relative conformations. For each $k = 1, \dots, K$, $\mathcal{S}_k = F_k \circ F_{k-1} \circ \dots \circ F_1(S_0)$ is the representation of S at step k , and the deformation F_k describes a *residual* transformation relative to \mathcal{S}_{k-1} , measured in terms of both across-scale and within-scale primitive relationships. In other words, the smaller scale features are described as *residues* of larger scale features. Consider, for example, the shape of a pair of pubic bones (see Fig. 1(b)). In this case F_1 may be a single global level deformation that is applied to the both bones simultaneously, F_2 may represent a *further* deformation of the left bone, F_3 may represent a collection of even further local deformations of the left bone, each applied to a very small part of it, and so on.

In addition to providing more intuitive and accurate geometric descriptions, the above multi-step approach also has the potential of leading to more efficient probabilistic models. We believe that by taking advantage of locality one can effectively characterize each deformation by a much smaller number of parameters. Because the process of estimating these parameters suffer a lot less from the HDLSS problem, it might be possible to obtain more reliable and stable estimations of statistical shape models from limited training samples.

We have been developing a methodology for describing 3D entities using medial representations combined with boundary displacements, which together form a representation called *m-reps* [5, 18]. In this framework, a geometric entity is represented at discrete scales and locations. At each scale it is described by a set of geometric primitives and their relative transformations. For a population of geometric entities, we establish correspondence among instances by fixing the topology of our representation. As a result, what differs among members of the population is the quantitative, geometric parameters and not qualitative properties of structure or topology. One method for determining the fixed topology from a population is described in [19].

We describe shapes at multiple scale levels to provide locality and efficiency in statistical analysis. The relationships between adjacent scale levels and among intra-scale neighbors make Markov models the natural choice. In this framework, we learn a series of probability models, each of which can be effectively characterized by a reasonably small number of parameters and thus can be estimated with limited number of training cases.

In what follows, we briefly describe m-reps in section 2. Details of the multi-object m-rep models are presented in sections 3 - 5. Section 6 shows an example of using statistical shape models to segment pubic bones. We finish with some discussion in section 7.

2 Multi-scale Representation with M-reps

Medial-based representations [20, 9, 21] provide a method of explicitly describing geometric deformations such as elongation, bending, and widening of object interiors. To obtain stable medial and boundary structures, it is important to build them in a multi-scale fashion, including a boundary displacement component. In our framework, called m-reps, at all but the boundary scale level, an object is described

by a set of continuous medial manifolds, which are sampled to yield discrete representations. Each sample point is called a *medial atom*, which describes a through section of an object (see Fig. 1(a)), and it is at this locality that the geometric deformation can be applied. An internal atom (*i.e.*, one that is not on the mesh boundary) is represented as a four-tuple $\mathbf{A} = (\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1})$ consisting of

- a translation $\mathbf{p} \in \mathbb{R}^3$, specifying the position of the medial point; we can consider this translation in units of the medial width r (defined below);
- a magnification scalar $r \in \mathbb{R}^+$, defined as the distance from the medial point to the implied boundary points $\mathbf{b}^{+1}, \mathbf{b}^{-1}$ (the local width);
- $\mathbf{U}^{+1}, \mathbf{U}^{-1} \in S^2$ are two unit vectors pointing from the medial position to the two implied boundary points $\mathbf{b}^{+1}, \mathbf{b}^{-1}$. Equivalently, they are the surface normal vectors at \mathbf{b}^{+1} and \mathbf{b}^{-1} ;

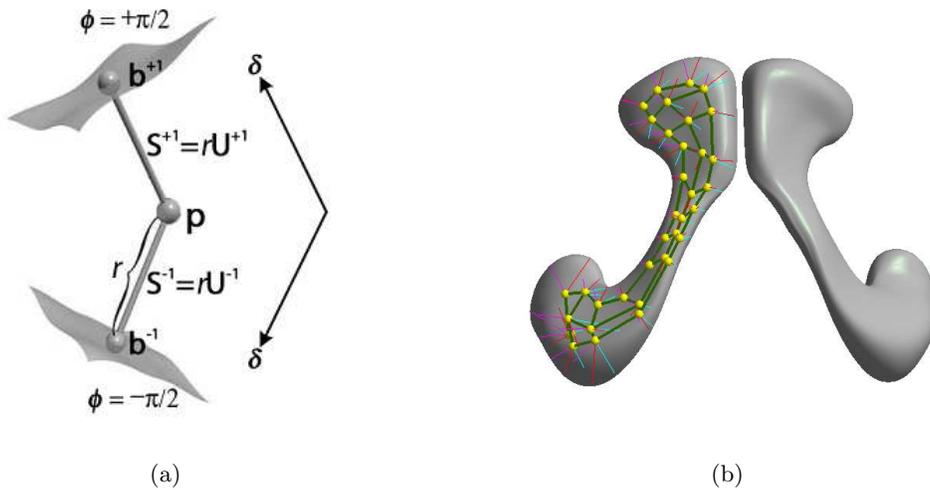


Figure 1: M-reps. (a) A medial atom with a cross-section of the boundary surface it implies; (b) An object complex containing two pubic bones. Each bone is represented by one m-rep figure.

As such, in m-reps each internal medial atom can be identified as an element of the eight-dimensional space $G = \mathbb{R}^3 \times \mathbb{R}^+ \times S^2 \times S^2$. For atoms on the boundary of the mesh, there is an extra parameter (with value in \mathbb{R}^+ . See [21]). For simplicity, in this paper we assume all atoms are elements of G .

In our current scheme, an m-rep *figure* is a sheet of medial atoms represented by a quadrilateral mesh of atoms with spacing determined through the analysis of the training population [19]. It describes a slab-like object or object part. The four-adjacency in the mesh determines the atom neighbor relationship (see 1(b)). A smooth boundary surface of a figure is generated by an algorithm (we presently use a subdivision surface method [22]) that approximates the boundary positions and normals implied by each atom. An m-rep object corresponds to a geometric object and is generally represented by a linked figural model:

a main figure describes the main section of an object; various sub-figures, each of which described by a single medial sheet, represent different branches, protrusions or indentations. Finally, an object complex is made up of individual objects. With m-reps, the inter-object, inter-figure, and inter-atom relations can be effectively described by appropriate sets of atom transformations, which are in turn describable by basic transformations such as translation, scaling, and rotation. Fig. 1(b) illustrates an example of this hierarchy of representations for a pair of pubic bones.

The m-rep framework allows geometric features at different positions and scale levels to be explicitly described. Furthermore, the medial structure, which is determined by a training population, provides a multi-scale intrinsic coordinate system [23]. This establishes correspondence among a population and makes m-reps extremely well suited for statistical analysis of shapes. The issue of how to choose the spacing of the medial atoms so that all the members of a population are in optimal correspondence is beyond the scope of this paper. The ideas in [24, 25] are certainly relevant.

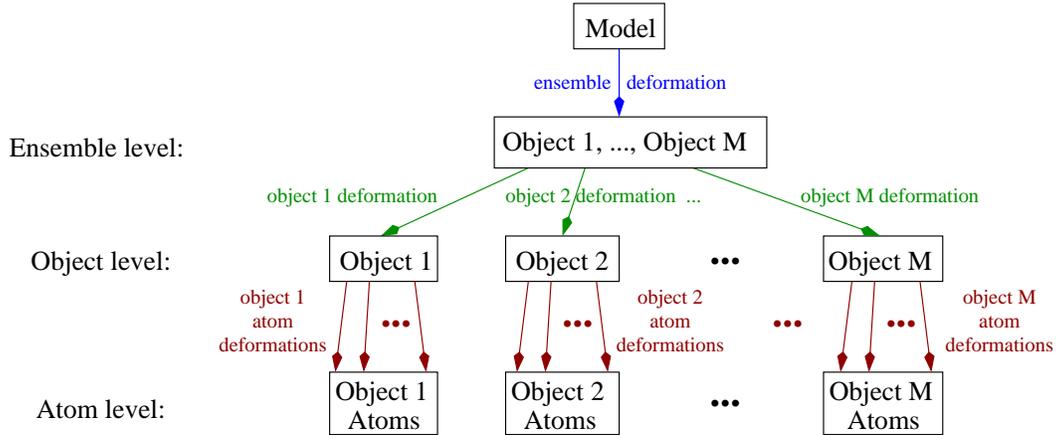


Figure 2: Hierarchical representation of an object ensemble consisting of single-figure objects.

In this paper, we consider object ensembles that contain only single-figure objects. The hierarchy of representation of such a geometric entity instance is obtained by applying a series of deformations to the shape model (template), illustrated in the diagram in Fig. 2. After the atom level, there can be a separate boundary displacement level where each medially implied boundary point moves along its associated surface normal direction to “fine tune” the representation. We will not discuss this step in detail here. A Markov random field model for boundary displacement fields can be found in [26].

Notice that in Fig. 2, the number of medial atoms used at each level of representation does not change. For instance, object 1 is described by the same number of atoms at the ensemble, object, and atom level. In other words, this is not a type of multi-grid representation. Instead, the number of deformations that are applied increase as one goes from coarse level to fine level. For example, suppose for $i = 1, 2, \dots, M$, object i has n_i atoms. Then at the ensemble level, there is one transformation that acts on all $\sum_{i=1}^M n_i$

atoms; at the object level, there are M different deformations; and at the atom level, there are $\sum_{i=1}^M n_i$ atom deformations. This set of deformations is what we need to model in order to characterize a shape.

3 Probabilistic Models for Object Ensembles

Let \mathcal{S}_0 be the shape model for an object ensemble. The underlying geometric entity, considered as a random quantity, is denoted by \mathbf{S} . We assume it can be represented by an m-rep model consisting of single-figure objects. Let \mathbf{A}_i denote the i -th medial atom of \mathbf{S} , assuming all the atoms are indexed by an index set \mathcal{I} . For any index set \mathcal{J} , $\mathbf{A}_{\mathcal{J}} := \{\mathbf{A}_i | i \in \mathcal{J}\}$. Therefore, the shape \mathbf{S} is just $\mathbf{A}_{\mathcal{I}}$. For simplicity, we identify \mathbf{A}_i with $\mathbf{A}_{\{i\}}$ when the index set has only one element, namely i .

Let $\mathbf{A}_{\mathcal{I}}^0$ be the shape model. A random sample of the geometric entity $\mathbf{A}_{\mathcal{I}}$ can be obtained by applying a random deformation \mathbf{F} to $\mathbf{A}_{\mathcal{I}}^0$. We describe this deformation by a series of, say K , random deformations, $\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_K$. The representation of $\mathbf{A}_{\mathcal{I}}$ at the k -th step is

$$\mathbf{A}_{\mathcal{I}}^k = \mathbf{F}_k \circ \mathbf{F}_{k-1} \circ \dots \circ \mathbf{F}_1(\mathbf{A}_{\mathcal{I}}^0) = \mathbf{F}_k(\mathbf{A}_{\mathcal{I}}^{k-1}), \quad \text{for } k = 1, 2, \dots, K.$$

The sequence $\mathbf{A}_{\mathcal{I}}^0, \mathbf{A}_{\mathcal{I}}^1, \dots, \mathbf{A}_{\mathcal{I}}^K$ are representations of $\mathbf{A}_{\mathcal{I}}$ at different scale levels and provide a series of approximations to $\mathbf{A}_{\mathcal{I}}$. Each representation $\mathbf{A}_{\mathcal{I}}^k$ corresponds to a node of the tree in Fig. 2. There are actually a collection of deformations in each atom deformation step. In these cases we use one single transformation to describe the net effect of all individual atom deformations.

Our goal is to describe the shape by defining a probability measure on $\mathbf{A}_{\mathcal{I}}$, which is usually considered as the shape prior. By virtue of our multi-step representation, we instead seek to define a joint probability distribution $\Pr(\mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K)$. The marginal distribution of the final step, $\Pr(\mathbf{A}_{\mathcal{I}}^K)$, is an approximation to the “true” distribution $\Pr(\mathbf{A}_{\mathcal{I}})$ (Strictly speaking, these probability distributions are defined on appropriate spaces of deformations. The mapping $\mathbf{A}_{\mathcal{I}}^{k-1} \rightarrow \mathbf{A}_{\mathcal{I}}^k$ determines \mathbf{F}_k .) Notice that we may equivalently define the joint distribution as follows:

$$\Pr(\mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K) = \Pr(\mathbf{A}_{\mathcal{I}}^1) \prod_{k=2}^K \Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1}, \dots, \mathbf{A}_{\mathcal{I}}^1). \quad (1)$$

At a first glance, this joint distribution seems to be more complicated than $\Pr(\mathbf{A}_{\mathcal{I}})$ and thus harder to estimate. However, there are several reasons why we consider that this alternative may be a better approach.

Firstly, because geometric features have intrinsic scales associated with them, it is more advantageous to describe a shape at multiple scale levels, so that features with various degrees of locality can be revealed. This also enables one to investigate and answer intuitive questions such as “which part of the object tends to vary more” or “how likely it is for one object to be bigger than a neighboring object”. In many applications, especially in medical image analysis, these kinds of questions are very crucial for users. The series of probability distributions on the right hand side of Eq. (1) allows one to explicitly describe features at different scale levels.

Secondly, the multi-step description may be more effective in certain applications. One such example is the image segmentation of prostates. Because of the variability in prostate geometry and the lack of contrast in image intensity from 3D CT images, experience of ours and others suggests that the accuracy of a direct automatic segmentation has so far been inadequate. To achieve better results, we can exploit the geometric relationship between the prostate and the surrounding organs. For instance, we may build a statistical model for an object ensemble that contains two pubic bones and the prostate. The bones have better image contrast and are easier to segment. Once they are in place, the statistical inter-object relations between the prostate and the bones can be used to predict where the prostate should be, giving us a better chance of locating it. Probabilistically, this multi-step process is best characterized by a marginal probability measure of the bones and a conditional probability measure of the prostate given the bones (ref. Eq. (1)).

Thirdly, to estimate $\Pr(\mathbf{A}_{\mathcal{T}})$ directly, we are likely to run into the HDLSS problem. On the other hand, we can estimate $\Pr(\mathbf{A}_{\mathcal{T}}^1, \dots, \mathbf{A}_{\mathcal{T}}^K)$ by estimating the set of marginal and conditional distributions given by the right hand side of Eq. (1). Under certain Markov assumptions (explained in detail later), each conditional distribution may be specified by a much smaller number of parameters and thus may be estimated more accurately and efficiently.

In order to use Eq. (1), an explicit order in which different objects and scales are described has to be chosen. In other words, we need to decide how to index the objects in the ensemble and how different nodes of the tree in Fig. 2 are traversed. Obviously, our estimation of the probability distribution would depend on these decisions. In many situations, such as in the earlier prostate example, the geometric variability and/or image quality for some objects (e.g., the bones) are more stable than those of others (e.g., the prostate). This may give us a hypothesis on how to order the objects, for example, to describe the bones first, then the prostate. As for traversing the tree in Fig. 2, some natural choices include the depth-first order and the width-first order. In the prostate example, since we have more confidence in describing and identifying bones, we may choose to describe them all the way through the atom level, then go to the object level for the prostate – which ends up in a depth-first order. It is certainly possible that there is more than one sensible choice. We argue that these decisions depend on the specific problems one faces. As a rule of thumb, objects should be described in an order of decreasing stability, i.e. objects with less combined variation in geometry and image contrast should be described before those with more. In order to achieve this, a measure of such variation needs to be established, and comparisons on the effects of choosing different orders need to be carried out. We will not discuss the details of these aspects further here. In the rest of this paper, we assume an appropriate order for the objects has been chosen, and we describe them in a depth-first manner, as shown in Fig. 3.

As an example, the two-bone shape model (see Fig. 1(b)) can be described in the following sequence of steps: the two-bone ensemble, the right bone (object 1) object level, the right bone atom level, the left bone (object 2) object level, and the left bone atom level. Therefore in this case, $M = 2$ and $K = 5$ (refer to Fig. 3.)

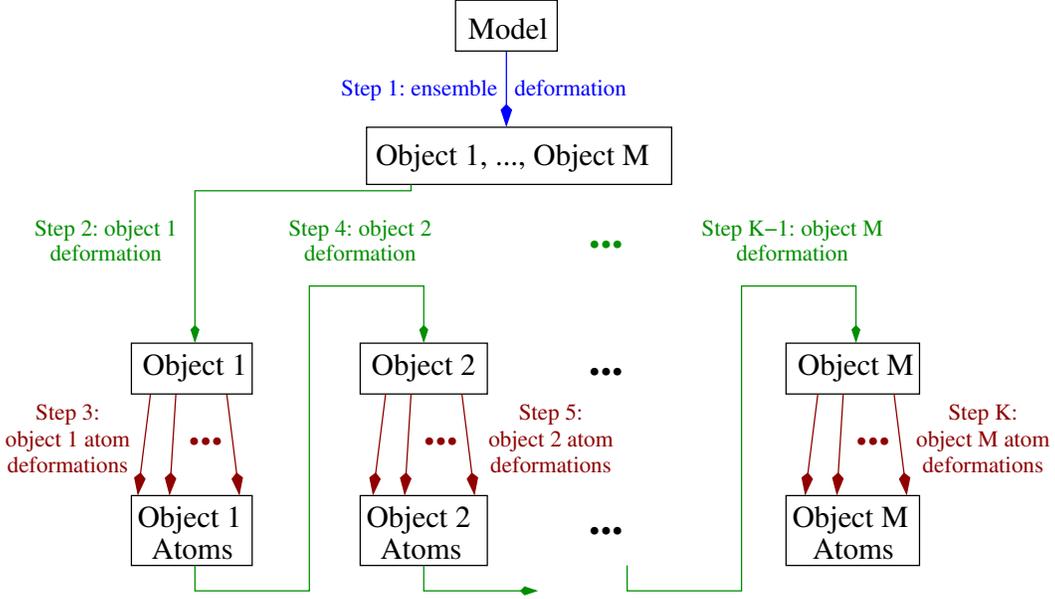


Figure 3: Depth-first representation of an object ensemble consisting of single-figure objects.

4 Markov Models for Object Ensembles

As we have seen, instances of a geometric entity $\mathbf{S} = \mathbf{A}_{\mathcal{I}}$ can be regarded as random samples drawn from an underlying population, according to certain probability distribution defined on it. This probability can be characterized by $\Pr(\mathbf{A}_{\mathcal{I}}^0, \dots, \mathbf{A}_{\mathcal{I}}^K)$, where $\mathbf{A}_{\mathcal{I}}^0, \dots, \mathbf{A}_{\mathcal{I}}^K$ are descriptions of $\mathbf{A}_{\mathcal{I}}$ at different levels. The shape model $\mathbf{A}_{\mathcal{I}}^0$ is supposed to be fixed, although it needs to be optimized. Thus in essence we need to define $\Pr(\mathbf{A}_{\mathcal{I}}^1, \dots, \mathbf{A}_{\mathcal{I}}^K | \mathbf{A}_{\mathcal{I}}^0)$, which for simplicity we will denote by $\Pr(\mathbf{A}_{\mathcal{I}}^1, \dots, \mathbf{A}_{\mathcal{I}}^K)$.

4.1 Description of Inter-object Relations

Because of inter-object relationships, the deformation of one object typically has an effect on other objects as well. We represent inter-object relations by an augmentation-prediction mechanism that is described in more detail in [27]. The basic idea is that for each object, say object k , we select from other objects a small subset of the atoms whose residual deformations are most correlated with those in object k . This group of selected atoms are called the *augmenting atoms* for object k . When describing the deformation on an object, we actually think of it as a deformation applied on both the object and its augmenting atoms. The rest of the atoms on other objects are deformed by a so-called *predicted deformation* which is intended to lower the overall shape change of these objects.

Suppose the shape template $\mathbf{A}_{\mathcal{I}}^0$ is fixed. For $k \geq 1$, let $\mathbf{F}_k : \mathbf{A}_{\mathcal{I}}^{k-1} \rightarrow \mathbf{A}_{\mathcal{I}}^k$ be the k -th step deformation.

At the first step, a global deformation \mathbf{F}_1 is applied to all the atoms in $\mathbf{A}_{\mathcal{I}}^0$ to yield $\mathbf{A}_{\mathcal{I}}^1$. At each of the following steps ($k \geq 2$), there is a deformation \mathbf{f}_k that is either an object-level deformation or a collection of atom-level deformations on the atoms of a particular object, say object m_k (refer to Fig. 3), *plus* its augmenting atoms. Since we describe objects in sequential order, we can assume the atoms of object 1 through object $m_k - 1$ *do not change* at this step. Therefore, the augmenting atoms are in objects $m_k + 1$ through object M . \mathbf{F}_k thus has three different types of effects on $\mathbf{A}_{\mathcal{I}}^{k-1}$:

- object 1 through object $m_k - 1$ do not deform;
- object m_k and its augmenting atoms are deformed by \mathbf{f}_k ;
- the rest of the atoms in object $m_k + 1$ through object M deform according to the predicted deformation Pred_k , as a result of \mathbf{f}_k .

If we define the index sets for these three sets of atoms by

$$\mathcal{L}^k := \bigcup_{m=1}^{m_k-1} \mathcal{I}_m, \quad \mathcal{M}^k := \mathcal{I}_{m_k} \cup \{\text{indices of augmenting atoms}\}, \quad \mathcal{N}^k := \mathcal{I} \setminus (\mathcal{L}^k \cup \mathcal{M}^k),$$

where “ \setminus ” denotes set difference, then in step k , $\mathbf{A}_{\mathcal{L}^k}$ do not change, $\mathbf{A}_{\mathcal{M}^k}$ deform according to \mathbf{f}_k , and $\mathbf{A}_{\mathcal{N}^k}$ deform according to Pred_k , which is based on (a) the previous-step relationship between $\mathbf{A}_{\mathcal{M}^k}^{k-1}$ and $\mathbf{A}_{\mathcal{N}^k}^{k-1}$ and (b) the current-step deformation on the augmenting atoms.

For a concrete example, consider the two-bone example. In this case \mathbf{F}_1 is a deformation on atoms of both bones. When $k = 2$, $m_k = m_2 = 1$. This is the the right bone (object 1) object step. We choose a subset of atoms on the left bone whose residual deformations are most highly correlated with those of the right bone atoms. These are the *augmenting atoms* at step $k = 2$. By modeling the correlation between the right bone and the left bone augmented atoms, we can predict where the latter should be when the right bone is deformed. The remaining atoms of the left bone are then deformed by an appropriate transformation so that the overall shape change of the left bone is minimized. In other words, the prediction function (a) reflects the effect of right bone deformation through the deformations of the augmented atoms, and (b) propagates this effect to all other atoms of the left bone in a way such that the overall geometry (modulo rigid or similarity transforms) of the left bone has minimal change.

In summary, the k -th step deformation $\mathbf{F}_k : \mathbf{A}_{\mathcal{I}}^{k-1} \rightarrow \mathbf{A}_{\mathcal{I}}^k$ is characterized as follows:

$$\mathbf{F}_1 : \mathbf{A}_{\mathcal{I}}^0 \mapsto \mathbf{A}_{\mathcal{I}}^1, \quad \text{for } k = 1,$$

$$\mathbf{F}_k : \begin{cases} \mathbf{A}_{\mathcal{L}^k}^{k-1} & \mapsto \mathbf{A}_{\mathcal{L}^k}^{k-1}, \\ \mathbf{A}_{\mathcal{M}^k}^{k-1} & \mapsto \mathbf{f}_k(\mathbf{A}_{\mathcal{M}^k}^{k-1}), \\ \mathbf{A}_{\mathcal{N}^k}^{k-1} & \mapsto \text{Pred}_k(\mathbf{A}_{\mathcal{N}^k}^{k-1}, \mathbf{A}_{\mathcal{M}^k}^{k-1}, \mathbf{f}_k(\mathbf{A}_{\mathcal{M}^k}^{k-1})), \end{cases} \quad \text{for } k \geq 2. \quad (2)$$

Because of the way we define each \mathbf{F}_k , the sequence $\mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K$ has the following Markov property: given the geometry of an object ensemble at all previous steps, the geometric description at any particular

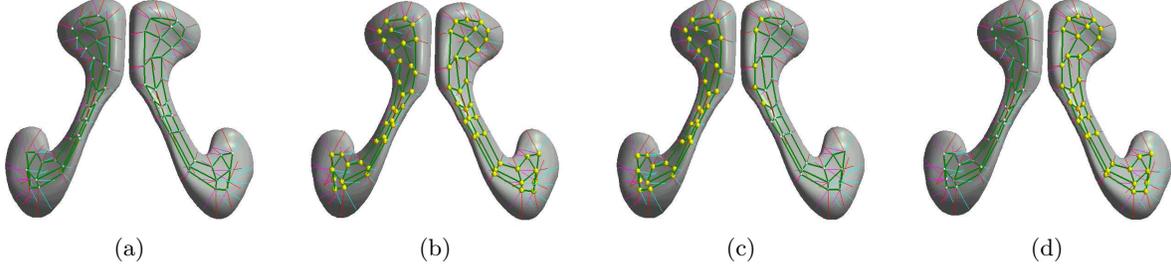


Figure 4: Multi-step representation of the two-bone shape. The deformation at each stage is applied to highlighted atoms. (a) The template. (b) The ensemble stage deformation is applied to all atoms. (c) The right bone deformation is applied to atoms on the right bone and the augmenting atoms on the left bone (highlighted). The rest of the atoms on the left bone are deformed by the prediction function. (d) The left bone deformation is applied to all atoms on the left bone. The right bone does not change.

step only depends on that of the immediate previous step. In other words,

$$\Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1}, \dots, \mathbf{A}_{\mathcal{I}}^1) = \Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1}), \quad \text{for } k \geq 2,$$

and therefore

$$\begin{aligned} \Pr(\mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K) &= \Pr(\mathbf{A}_{\mathcal{I}}^1) \prod_{k=2}^K \Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1}, \dots, \mathbf{A}_{\mathcal{I}}^1) \\ &= \Pr(\mathbf{A}_{\mathcal{I}}^1) \prod_{k=2}^K \Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1}). \end{aligned} \tag{3}$$

In light of Eq. (3), we will focus on defining $\Pr(\mathbf{A}_{\mathcal{I}}^1)$ and $\Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1})$ for $k \geq 2$. This set of probabilities describe (a) the geometric variability of each object at different scale levels; and (b) the inter-object relations, using augmentation and prediction discussed earlier.

4.2 Residues and Residue Deformations

Our approach is to define $\Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1})$ in terms of the difference between $\mathbf{A}_{\mathcal{I}}^k$ and $\mathbf{A}_{\mathcal{I}}^{k-1}$. Recall that a medial atom is an element of the space $G = \mathbb{R}^3 \times \mathbb{R}^+ \times S^2 \times S^2$. In fact, the space G is a *symmetric space*. A more thorough discussion about symmetric spaces and their properties can be found in [28] and the references therein. Given any two medial atoms $\mathbf{A}_1, \mathbf{A}_2 \in G$, their difference can be described by using the following operator:

$$\begin{aligned} \ominus : G \times G &\rightarrow G, \\ (\mathbf{A}_1, \mathbf{A}_2) &\mapsto \left(\mathbf{x}_1 - \mathbf{x}_2, \frac{r_1}{r_2}, \mathbf{R}_{\mathbf{u}_2}(\mathbf{u}_1), \mathbf{R}_{\mathbf{v}_2}(\mathbf{v}_1) \right), \end{aligned} \tag{4}$$

where $\mathbf{A}_i = (\mathbf{x}_i, r_i, \mathbf{u}_i, \mathbf{v}_i), i = 1, 2$, and for any $\mathbf{w} = (w_1, w_2, w_3) \in S^2$, $\mathbf{R}_{\mathbf{w}} \in SO(3)$ denotes the rotation around the axis that passes through $(0, 0, 0)$ and $(w_2, -w_1, 0)$, with the rotation angle θ being the spherical distance between $\mathbf{p} = (0, 0, 1)$ and \mathbf{w} , i.e., $\theta = \arccos(w_3)$. Also, Let $\mathbf{R}_{\mathbf{w}}^{-1}$ be the inverse of $\mathbf{R}_{\mathbf{w}}$. It is easy to see that $\mathbf{R}_{\mathbf{w}}(\mathbf{w}) = \mathbf{p}$ and that $\mathbf{R}_{\mathbf{w}}$ maps a neighborhood of \mathbf{w} to a neighborhood of \mathbf{p} . $\mathbf{R}_{\mathbf{w}}^{-1}$ does just the opposite.

Let $\Delta\mathbf{A} = \mathbf{A}_1 \ominus \mathbf{A}_2$. It is also an element of G and is called the *residue* of \mathbf{A}_1 relative to \mathbf{A}_2 . It is a measure of the difference between \mathbf{A}_1 and \mathbf{A}_2 , *relative to \mathbf{A}_2 's coordinates*. When $\mathbf{A}_1 = \mathbf{A}_2$, we get $\Delta\mathbf{A} = (\mathbf{0}, 1, \mathbf{p}, \mathbf{p})$, which is the identity of G (if we choose \mathbf{p} to be the identity of S^2). On the other hand, given an atom $\mathbf{A} = (\mathbf{x}, r, \mathbf{u}, \mathbf{v})$ and a residue $\Delta\mathbf{A} = (\Delta\mathbf{x}, \Delta r, \Delta\mathbf{u}, \Delta\mathbf{v})$, we can obtain the atom that has residue $\Delta\mathbf{A}$ relative to \mathbf{A} using the operator \oplus defined as

$$\begin{aligned} \oplus : G \times G &\rightarrow G, \\ (\mathbf{A}, \Delta\mathbf{A}) &\mapsto \left(\mathbf{x} + \Delta\mathbf{x}, r \cdot \Delta r, \mathbf{R}_{\mathbf{u}}^{-1}(\Delta\mathbf{u}), \mathbf{R}_{\mathbf{v}}^{-1}(\Delta\mathbf{v}) \right). \end{aligned} \quad (5)$$

Notice that this operator is neither commutable nor associative. The second operand is always understood as a residue measured relative to the coordinates of the first operand.

For two sets of corresponding atoms with a common index set \mathcal{J} , say $\mathbf{A}_{\mathcal{J}}^1$ and $\mathbf{A}_{\mathcal{J}}^2$, the difference between them is defined to be

$$\mathbf{A}_{\mathcal{J}}^1 \ominus \mathbf{A}_{\mathcal{J}}^2 := \{\mathbf{A}_j^1 \ominus \mathbf{A}_j^2 : j \in \mathcal{J}\}.$$

In our multi-scale representation, the k -th step representation $\mathbf{A}_{\mathcal{I}}^k$ is an approximation to the ‘‘true’’ representation $\mathbf{A}_{\mathcal{I}}$. The *residue* at the k -th step is defined to be

$$\mathbf{R}_{\mathcal{I}}^k := \mathbf{A}_{\mathcal{I}} \ominus \mathbf{A}_{\mathcal{I}}^k, \quad k \geq 1. \quad (6)$$

It represents the residual geometric information that needs to be described at later steps. On the other hand, the effect of the k -th step deformation \mathbf{F}_k is given by the difference

$$\Delta\mathbf{A}_{\mathcal{I}}^k := \mathbf{A}_{\mathcal{I}}^k \ominus \mathbf{A}_{\mathcal{I}}^{k-1}, \quad k \geq 1. \quad (7)$$

Therefore, $\Delta\mathbf{A}_{\mathcal{I}}^k$ can be regarded as an approximation to $\mathbf{R}_{\mathcal{I}}^{k-1}$. We will call \mathbf{F}_k (or equivalently, $\Delta\mathbf{A}_{\mathcal{I}}^k$) the *residue deformation* at step k .

In the two-bone shape example, let $\mathbf{A}_{\mathcal{I}}$ be a random two-bone shape, and $\mathbf{A}_{\mathcal{I}}^0$ be the shape template. At each step $k(1 \leq k \leq 5)$, the residue is $\mathbf{R}_{\mathcal{I}}^k = \mathbf{A}_{\mathcal{I}} \ominus \mathbf{A}_{\mathcal{I}}^k$, and the residue deformation is given by $\Delta\mathbf{A}_{\mathcal{I}}^k = \mathbf{A}_{\mathcal{I}}^k \ominus \mathbf{A}_{\mathcal{I}}^{k-1}$. We emphasize again that each residue deformation is described relative to the geometry of the previous step, not relative to the shape template. For instance, at step 4, the residue deformation is a further object level deformation of the left bone after step 3.

In general, we construct the sequence $\mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K$ by choosing statistics at any step such that interactions at the corresponding scale level is accounted for and thus the conditional distribution $\Pr(\mathbf{A}_{\mathcal{I}}^k | \mathbf{A}_{\mathcal{I}}^{k-1})$ can be induced by a probability distribution on $\Delta\mathbf{A}_{\mathcal{I}}^k$, and that this distribution does not explicitly depend

Step	Representation	Residue	Residue deformation
shape template	$\mathbf{A}_{\mathcal{I}}^0$	$\mathbf{R}_{\mathcal{I}}^0$	—
two-bone ensemble	$\mathbf{A}_{\mathcal{I}}^1$	$\mathbf{R}_{\mathcal{I}}^1$	$\Delta\mathbf{A}_{\mathcal{I}}^1$: global deformation of both bones
right bone object	$\mathbf{A}_{\mathcal{I}}^2$	$\mathbf{R}_{\mathcal{I}}^2$	$\Delta\mathbf{A}_{\mathcal{I}}^2$: right bone object level deformation, plus prediction of left bone
right bone atom	$\mathbf{A}_{\mathcal{I}}^3$	$\mathbf{R}_{\mathcal{I}}^3$	$\Delta\mathbf{A}_{\mathcal{I}}^3$: right bone atom level deformation, plus prediction of left bone
left bone object	$\mathbf{A}_{\mathcal{I}}^4$	$\mathbf{R}_{\mathcal{I}}^4$	$\Delta\mathbf{A}_{\mathcal{I}}^4$: left bone object level deformation
left bone atom	$\mathbf{A}_{\mathcal{I}}^5$	$\mathbf{R}_{\mathcal{I}}^5$	$\Delta\mathbf{A}_{\mathcal{I}}^5$: left bone atom level deformation

Table 1: Multi-step representation of the two-bone shape.

on $\mathbf{A}_{\mathcal{I}}^{k-1}$. As a result, the original joint distribution $\Pr(\mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K)$ can be induced by $\mathbf{A}_{\mathcal{I}}^0$ and the product

$$\prod_{k=1}^K \Pr(\Delta\mathbf{A}_{\mathcal{I}}^k). \quad (8)$$

In stochastic process terminology, we have designed our approach so that the sequence of representations $\mathbf{A}_{\mathcal{I}}^0, \mathbf{A}_{\mathcal{I}}^1, \mathbf{A}_{\mathcal{I}}^2, \dots, \mathbf{A}_{\mathcal{I}}^K$ has “independent increments”. The multi-step representation of the two-bone shape $\mathbf{A}_{\mathcal{I}}$ is summarized in Table 1. In this case, we need to estimate $\mathbf{A}_{\mathcal{I}}^0$ as well as the (independent) probabilities $\Pr(\Delta\mathbf{A}_{\mathcal{I}}^1), \dots, \Pr(\Delta\mathbf{A}_{\mathcal{I}}^5)$.

4.3 Atom Level Residue Deformations

As mentioned earlier, if $\Delta\mathbf{A}_{\mathcal{I}}^k$ corresponds to an atom residue transformation step (e.g., if $k = 3$ or $k = 5$ in the previous example), it actually represents a collection of individual atom residue deformations, each of which characterizes the deformation of an atom given those of all other atoms and is described by the probability $\Pr(\Delta\mathbf{A}_i^k | \Delta\mathbf{A}_{\mathcal{M}^k \setminus \{i\}}^k)$, where $\mathbf{A}_{\mathcal{M}^k}$ contains (a) the atoms that are subject to atom residue deformations at step k , and (b) the augmenting atoms (on neighboring objects). At this finer scale level, we assume that the direct long range dependency among $\mathbf{A}_{\mathcal{M}^k}$ have already been described at the (larger) object level, so that the atom residue deformations are localized and can be described by the following Markov random field (MRF) model:

$$\Pr(\Delta\mathbf{A}_i^k | \Delta\mathbf{A}_{\mathcal{M}^k \setminus \{i\}}^k) = \Pr(\Delta\mathbf{A}_i^k | \Delta\mathbf{A}_{N(i)}^k), \quad i \in \mathcal{M}^k, \quad (9)$$

where $N(i)$ denotes the index set of the atoms that are neighbors of atom i . Because the medial sheet is sampled by a quadrilateral array of atoms, we use the neighborhood structure induced by the 4-adjacency graph of the quad-mesh (see Fig. 5.)

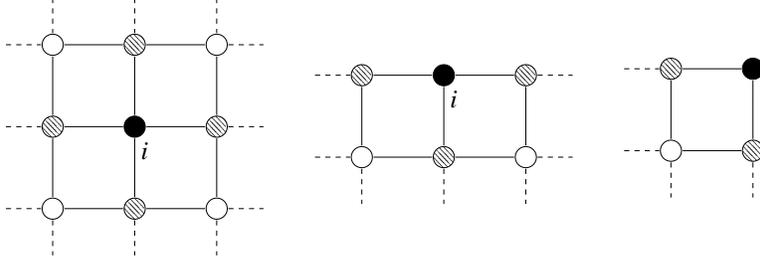


Figure 5: The 4-neighbor structure for quad-mesh. A typical atom i has 4 neighbors (shaded), whereas an atom on an edge or at a corner has fewer neighbors.

We can define these MRF models by specifying potentials on the cliques in the 4-adjacency graph and estimating parameters using standard MRF techniques such as Monte Carlo Markov Chain (MCMC) methods. An example of this approach can be found in [26]. In many situations, this type of models still yield a large number of parameters that need to be estimated simultaneously. An alternative approach is to select an order in which we go through the atoms one by one and decompose the joint distribution $\Pr(\Delta \mathbf{A}_{\mathcal{M}^k})$ into a product of marginal and conditional distributions, similar to Eq. (1). However, the quadrilateral structure of the atom mesh does not provide a natural traversing order, making this method somewhat ad hoc and hard to analyze. A third method, which we adopt here, is to estimate the conditional distributions $\Pr(\Delta \mathbf{A}_i^k | \Delta \mathbf{A}_{N(i)}^k)$ directly. To deal with the low sample size problem, we assume that the conditional probability distribution $\Pr(\Delta \mathbf{A}_i^k | \Delta \mathbf{A}_{N(i)}^k)$ does not explicitly depend on $\Delta \mathbf{A}_{N(i)}^k$, but rather is induced by a probability distribution on the difference $\Delta \mathbf{A}_i^k \ominus \Delta \mathbf{A}_{i,p}^k$, where for each i , the term $\Delta \mathbf{A}_{i,p}^k$ represents a residue deformations of atom i that is *predicted* by the deformations of its neighbors and is calculated as a weighted average of $\Delta \mathbf{A}_{N(i)}^k$. The problem then is to estimate the probabilities

$$\Pr(\Delta \mathbf{A}_i^k \ominus \Delta \mathbf{A}_{i,p}^k), i \in \mathcal{M}^k. \quad (10)$$

From (8) and (10), we see that the probability distribution of the object ensemble is induced from a product of inter-scale and intra-scale residue distributions, each of which can be estimated separately. In the next section we discuss how to learn these distributions from a training set.

5 Estimation of the Probabilistic Model

Suppose the training set $\{A_{1,\mathcal{I}}, A_{2,\mathcal{I}}, \dots, A_{N,\mathcal{I}}\}$ contains N instances of the random object ensemble $\mathbf{A}_{\mathcal{I}}$. Let $\mathcal{T} = \{1, 2, \dots, N\}$. We use $A_{n,\mathcal{I}}$ to denote an instance. The first subscript n denotes the index in the training set, i.e., $n \in \mathcal{T}$, and the second one denotes the index for medial atoms as before. Therefore, the whole training set can be denoted by $A_{\mathcal{T},\mathcal{I}}$. The m-rep representation provides correspondence among the training set. In particular, the corresponding atoms across the population have the same index in \mathcal{I} . As

before, we assume that the object ensemble has K single-figure objects, and that we use the depth-first order to traverse the objects as shown in Fig. 3.

5.1 M-rep Alignment and Principal Geodesic Analysis (PGA)

The first step in training is to align the shape instances by an m-rep alignment procedure that is similar to Procrustes analysis. In the standard Procrustes analysis, a rigid or similarity transformation is applied to each instance so that the total sum-of-squared Euclidean distances between corresponding points is minimized. In contrast, in the m-rep alignment, described in [29], the sum-of-squared geodesic distances between corresponding medial atoms is minimized. In what follows we assume the shape instances are aligned.

The probabilities $\Pr(\Delta\mathbf{A}_{\mathcal{I}}^1), \dots, \Pr(\Delta\mathbf{A}_{\mathcal{I}}^K)$ are defined on m-rep parameter spaces, which are nonlinear Riemannian symmetric spaces. In these spaces, one can define the so-called exponential map, Exp_x , and log map, Log_x , which establish correspondence between tangent vectors and geodesics at any given point x . Based on the properties of these maps, Fletcher *et. al.* have developed a method called principal geodesic analysis (PGA), which is a generalization of principal component analysis to nonlinear Lie groups and symmetric spaces. The details are laid out in [29, 30]. Here we briefly summarize the basic ideas.

Suppose M is a symmetric space or Lie group. Given a set of samples $x_1, \dots, x_N \in M$, we want to estimate the underlying probability distribution from which they are drawn. The PGA of $\{x_i\}_{i=1}^N$ essentially involves finding the mean μ of $\{x_i\}_{i=1}^N$, projecting $\{x_i\}_{i=1}^N$ to $T_\mu M$ (the tangent space at μ) by the log map, and doing a standard principal component analysis (PCA) in that tangent space, which yields a set of principal directions $\{v_k\}$ and corresponding eigenvalues $\{\lambda_k\}$. The ‘‘principal geodesics’’ are determined by applying the exponential map to the principal components obtained from PCA. Analogous to principal components, the principal geodesics describe the major modes of variation among $\{x_i\}_{i=1}^N$ in the curved space M (see Fig. 6).

The space M can be approximated by a submanifold $H \subset M$ that is generated by the first, say h , principal geodesics. Elements in H can be represented as

$$\text{Exp}_\mu\left(\sum_{k=1}^h \alpha_k v_k\right), \tag{11}$$

where $\alpha_k \in \mathbb{R}$. $\forall x \in M$, the projection of x onto H is

$$\pi_H(x) = \arg \min_{y \in H} \text{dist}^2(x, y) \approx \text{Exp}_\mu\left(\sum_{k=1}^h \langle \text{Log}_\mu(x), v_k \rangle \cdot v_k\right),$$

where $\langle \cdot, \cdot \rangle$ denotes inner product in $T_\mu M$.

The density

$$\frac{1}{Z} \exp\left(-\frac{1}{2} \sum_{k=1}^h \frac{\langle v_k, y \rangle^2}{\lambda_k}\right), \quad y \in N(\mu) \subset T_\mu M, \tag{12}$$

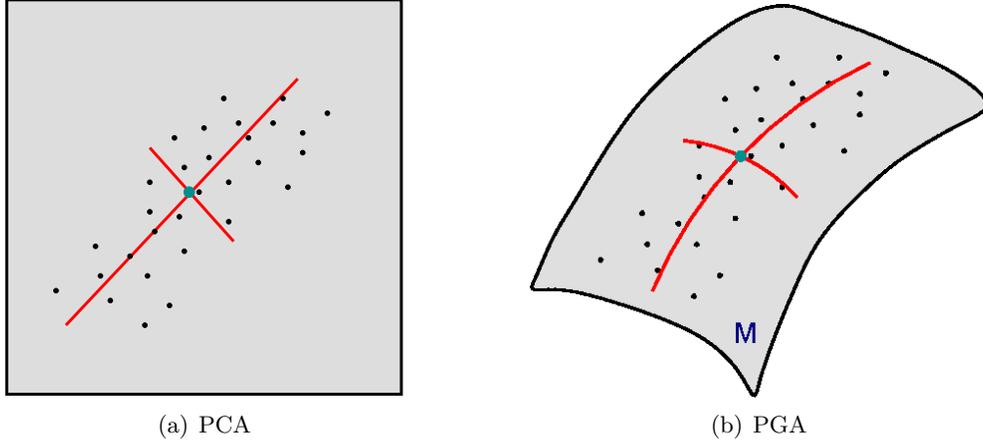


Figure 6: Comparison of PCA in linear spaces and PGA in curved spaces. Each small dot represents a sample point. The big dot indicates the Fréchet mean. The thick lines/curves represent principal components/geodesics. Both surfaces represent n -dimensional spaces.

where Z is a normalizing constant and $N(\mu)$ is a neighborhood of μ in the tangent plane, defines a probability distribution on $N(\mu)$. It induces a probability distribution on M by the exponential map Exp_μ . This induced distribution is an approximation to the probability distribution on M from which $\{x_i\}_{i=1}^N$ are drawn.

In what follows we describe the process for estimating $\Pr(\Delta \mathbf{A}_{\mathcal{I}}^1), \dots, \Pr(\Delta \mathbf{A}_{\mathcal{I}}^K)$. The two-bone shape will be used as an example. The ensemble has a total of 72 medial atoms, with 36 atoms on each bone.

5.2 Estimation of the Shape Model

The shape model $\mathbf{A}_{\mathcal{I}}^0$ is set to be the centroid of the training population $A_{\mathcal{T}, \mathcal{I}}$. This is obtained by computing the Fréchet mean of each atom. In other words,

$$\mathbf{A}_{\mathcal{I}}^0 = \text{mean}_{\mathcal{T}}(A_{\mathcal{T}, \mathcal{I}}) = \left\{ \text{mean}_{\mathcal{T}}(A_{\mathcal{T}, i}) : i \in \mathcal{I} \right\}. \quad (13)$$

The Fréchet mean of the set of atoms $A_{\mathcal{T}, i}$ is defined to be an atom \mathbf{A}_i^0 that satisfies

$$\mathbf{A}_i^0 = \arg \min_{A \in G} \sum_{n \in \mathcal{T}} \text{dist}^2(A, A_{n, i}),$$

where $\text{dist}(\cdot, \cdot)$ is the geodesic distance on G , the atom parameter space. The Fréchet mean can be computed using a gradient descent method. For details refer to [30, 29].

The residues at this step are the difference between the instances and the model, i.e.,

$$R_{n, \mathcal{I}}^0 := A_{n, \mathcal{I}} \ominus \mathbf{A}_{\mathcal{I}}^0, \quad n \in \mathcal{T}.$$

5.3 Object Ensemble Residue Statistics

The ensemble deformation $\Delta\mathbf{A}_{\mathcal{I}}^1$ is an approximation to $R_{n,\mathcal{I}}^0$. We estimate $\Pr(\Delta\mathbf{A}_{\mathcal{I}}^1)$ by doing a principal geodesic analysis (PGA) on $R_{\mathcal{I},\mathcal{I}}^0$. In the two-bone example, this amounts to a PGA of a population of tuples consisting of 72 atoms, or equivalently, a 576-dimensional PGA since each atom has 8 parameters.

Let H be the m-rep parameter space (in fact, if there are a total of d medial atoms in the object ensemble, then $H = G^d$. Recall that $G = \mathbb{R}^3 \times \mathbb{R}^+ \times S^2 \times S^2$.) $\Pr(\Delta\mathbf{A}_{\mathcal{I}}^1)$ is approximated by the first h_1 principal geodesics, which generate a sub-manifold $H^1 \subset H$. The ensemble deformation for each instance is described by the projection of $R_{\mathcal{I},\mathcal{I}}^0$ onto H^1 , i.e.,

$$\Delta A_{n,\mathcal{I}}^1 := \text{Proj}_{H^1}(R_{n,\mathcal{I}}^0), \quad n \in \mathcal{T}.$$

The representations and residues at this step are

$$\begin{aligned} A_{n,\mathcal{I}}^1 &:= \mathbf{A}_{\mathcal{I}}^0 \oplus \Delta A_{n,\mathcal{I}}^1, \\ R_{n,\mathcal{I}}^1 &:= A_{n,\mathcal{I}} \ominus A_{n,\mathcal{I}}^1, \end{aligned} \quad n \in \mathcal{T},$$

where \oplus is defined by Eq. (5). The above procedure is described by Algorithm 1.

Input: $A_{\mathcal{I},\mathcal{I}}, \mathbf{A}_{\mathcal{I}}^0$
Output: P^1 = estimates of $\Pr(\Delta\mathbf{A}_{\mathcal{I}}^1)$, as well as $A_{\mathcal{I},\mathcal{I}}^1$
 $R_{\mathcal{I},\mathcal{I}}^0 \leftarrow A_{\mathcal{I},\mathcal{I}} \ominus \mathbf{A}_{\mathcal{I}}^0$
 $P^1 \leftarrow$ PGA results on $R_{\mathcal{I},\mathcal{I}}^0$ (principal geodesics and eigenvalues)
 $H^1 \leftarrow$ sub-manifold generated by a chosen number of principal geodesics of $R_{\mathcal{I},\mathcal{I}}^0$
 $\Delta A_{\mathcal{I},\mathcal{I}}^1 \leftarrow \text{Proj}_{H^1}(R_{\mathcal{I},\mathcal{I}}^0), A_{\mathcal{I},\mathcal{I}}^1 \leftarrow \mathbf{A}_{\mathcal{I}}^0 \oplus \Delta A_{\mathcal{I},\mathcal{I}}^1$
Output $P^1, A_{\mathcal{I},\mathcal{I}}^1$

Algorithm 1: Object ensemble statistics.

As with PCA, we need to decide the value for h_1 , which is the number of principal geodesics to use for approximation. While too few geodesics might not be adequate, too many may be unstable due to the small sample size. In principle, we want to choose the smallest value such that the residues possess certain Markov properties at the (next) object level. In practice, this is not easy to verify, and we use some heuristics to determine h_1 . For example, we can choose h_1 to be the smallest number such that (a) the percentage of variance explained by the first h_1 principal geodesics is above a certain threshold, and/or (b) the inter-object correlations among residues are below a certain threshold.

5.4 Object Residue Statistics for Object One

First, suppose object 1's residue deformation does not have an effect on neighboring objects. Let the index set for the atoms in object 1 be \mathcal{I}_1 . Then at this step, the probability $\Pr(\Delta\mathbf{A}_{\mathcal{I}}^2)$ actually describes

the object residue deformation $\Delta \mathbf{A}_{\mathcal{I}_1}^2$. We estimate this probability by performing a PGA on $R_{\mathcal{I}, \mathcal{I}_1}^1$ and using the first h_2 principal geodesics to approximate it. The procedure is similar to Algorithm 1 and is summarized in Algorithm 2. For the two-bone case, this corresponds to a $36 * 8 = 288$ dimensional PGA on the right bone.

Input: $A_{\mathcal{T}, \mathcal{J}}, A_{\mathcal{T}, \mathcal{J}}^{k-1}$, where \mathcal{J} is the index set for the atoms in the object

Output: $P^k =$ estimates of $\Pr(\Delta \mathbf{A}_{\mathcal{J}}^k)$, as well as $A_{\mathcal{T}, \mathcal{J}}^k$

$$R_{\mathcal{T}, \mathcal{J}}^{k-1} \leftarrow A_{\mathcal{T}, \mathcal{J}} \ominus A_{\mathcal{T}, \mathcal{J}}^{k-1}$$

$P^k \leftarrow$ PGA results on $R_{\mathcal{T}, \mathcal{J}}^{k-1}$ (Fréchet mean $\mu_{\mathcal{J}}^k$, principal geodesics and eigenvalues)

$H^k \leftarrow$ sub-manifold generated by the principal geodesics of $R_{\mathcal{T}, \mathcal{J}}^{k-1} \ominus \mu_{\mathcal{J}}^k$

$$\Delta A_{\mathcal{T}, \mathcal{J}}^k \leftarrow \text{Proj}_{H^k}(R_{\mathcal{T}, \mathcal{J}}^{k-1}), A_{\mathcal{T}, \mathcal{J}}^k \leftarrow A_{\mathcal{T}, \mathcal{J}}^{k-1} \oplus \mu_{\mathcal{J}}^k \oplus \Delta A_{\mathcal{T}, \mathcal{J}}^k$$

Output $P^k, A_{\mathcal{T}, \mathcal{J}}^k$

Algorithm 2: Object residue statistics for a single-figure object at step k .

To capture the effect of object 1’s deformation on other objects, we need to update all other objects so that the inter-object relationship is reflected. This is achieved with an augmentation-prediction mechanism. We use the two-bone shape to illustrate how it works. In this case object 1 is the right bone, and object 2 is the left bone. The index set for the atoms of these objects are $\mathcal{I}_1, \mathcal{I}_2$, respectively. We pick a subset of object 2 (left bone)’s atoms that are most highly correlated with the atoms in object 1 (right bone); these typically correspond to the section of object 2 that is closest to object 1 (see Fig. 4). Let \mathcal{U}^2 be the index set for the augmenting atoms and \mathcal{M}^2 denote the union of \mathcal{I}_1 and \mathcal{U}^2 . We can then perform an *augmented* PGA analysis on $R_{\mathcal{T}, \mathcal{M}^2}^1$ instead of $R_{\mathcal{T}, \mathcal{I}_1}^1$, using Algorithm 2 with $\mathcal{J} = \mathcal{M}^2$. In the two-bone example, if there are 5 augmented atoms on the left bone, then this leads to a $(36 + 5) \cdot 8 = 328$ dimensional PGA.

The deformation of the augmenting atoms $\Delta \mathbf{A}_{\mathcal{U}^2}^2 : \mathbf{A}_{\mathcal{U}^2}^1 \rightarrow \mathbf{A}_{\mathcal{U}^2}^2$ also predict how the rest of the atoms on object 2 will deform. Our approach is as follows. We find that member of the principal geodesic shape space of the remaining objects that agrees best on the augmenting atoms, and we remove that member from the remaining objects’ atoms. In detail, first, we perform a PGA on $R_{\mathcal{T}, \mathcal{I}_2}^1$, the residue of object 2, and denote the mean and PGA submanifold by ρ^2 and \mathcal{H}^2 , respectively. The *predicted* residual deformation of object 2 is then defined to be

$$\text{Pred}_2(\Delta \mathbf{A}_{\mathcal{I}_2}^2; \Delta \mathbf{A}_{\mathcal{U}^2}^2) = \text{Proj}_{\mathcal{H}^2}(\Delta \mathbf{A}_{\mathcal{U}^2}^2) = \exp_{\rho^2} \left(\sum_{j=1}^h \langle \text{Log}_{\rho^2}(\Delta \mathbf{A}_{\mathcal{U}^2}^2), v_j \rangle \cdot v_j \right), \quad (14)$$

where $\{v_j\}_{j=1}^h$ are principal directions in the tangent space of ρ^2 corresponding to the principal geodesics in \mathcal{H}^2 and the dimension of $\text{Log}_{\rho^2}(\Delta \mathbf{A}_{\mathcal{U}^2}^2)$ is adjusted to match with that of v_j by adding zeros to $\text{Log}_{\rho^2}(\Delta \mathbf{A}_{\mathcal{U}^2}^2)$ for parameters corresponding to $\mathbf{A}_{\mathcal{I}_2 \setminus \mathcal{U}^2}$.

It is straightforward to adapt these procedures to situations where the object ensemble has more than

two objects. The net effect of these predictions is summarized by a prediction function $\text{Pred}_2()^1$, which is determined using statistics from a training set. The general algorithm is described in Alorithm 3.

Input: $A_{\mathcal{T},\mathcal{J}}, A_{\mathcal{T},\mathcal{J}}^{k-1}, A_{\mathcal{T},\mathcal{U}'}^k$, where \mathcal{J} is the index set for atoms in objects that need to be predicted and $\mathcal{U} \subset \mathcal{J}$ is the index set for augmenting atoms at step k

Output: $A_{\mathcal{T},\mathcal{J}}^k =$ predicted objects

$R_{\mathcal{T},\mathcal{J}}^{k-1} \leftarrow A_{\mathcal{T},\mathcal{J}} \ominus A_{\mathcal{T},\mathcal{J}}^{k-1}$

$\mathcal{P}^k \leftarrow$ PGA results on $R_{\mathcal{T},\mathcal{J}}^{k-1}$ (Fréchet mean $\rho_{\mathcal{J}}^k$, principal geodesics and eigenvalues)

$\mathcal{H}^k \leftarrow$ sub-manifold generated by the principal geodesics of $R_{\mathcal{T},\mathcal{J}}^{k-1} \ominus \rho_{\mathcal{J}}^k$

$A_{\mathcal{T},\mathcal{J}}^k \leftarrow A_{\mathcal{T},\mathcal{J}}^{k-1} \oplus \rho_{\mathcal{J}}^k \oplus \text{Proj}_{\mathcal{H}^k}(\Delta A_{\mathcal{T},\mathcal{U}'}^k)$

Output $\mathcal{P}^k, A_{\mathcal{T},\mathcal{J}}^k$

Algorithm 3: Statistical prediction at step k .

In the two-bone example, after the object 1 (right bone) deformation step the we have

$$\begin{aligned}
 A_{\mathcal{T},\mathcal{I}_1}^2 &= A_{\mathcal{T},\mathcal{I}_1}^1 \oplus \mu_{\mathcal{I}_1}^2 \oplus \text{Proj}_{H^2}(R_{\mathcal{T},\mathcal{I}_1}^1), \\
 A_{\mathcal{T},\mathcal{I}_2}^2 &= A_{\mathcal{T},\mathcal{I}_2}^1 \oplus \rho_{\mathcal{I}_2}^2 \oplus \text{Proj}_{H^2}(\Delta A_{\mathcal{T},\mathcal{U}'}^2), \\
 R_{\mathcal{T},\mathcal{I}}^2 &= A_{\mathcal{T},\mathcal{I}} \ominus A_{\mathcal{T},\mathcal{I}}^2.
 \end{aligned} \tag{15}$$

5.5 Atom Residue Statistics for Object One

The probabilities to be estimated at this step are $\Pr(\Delta \mathbf{A}_i^3 \ominus \Delta \mathbf{A}_{i,p}^3), i \in \mathcal{J}^{3,1}$. Recall that $\Delta \mathbf{A}_{i,p}^3$ is the residue deformation of atom i that is predicted by those of its neighbors. For each i , we estimate $\Pr(\Delta \mathbf{A}_i^3 \ominus \Delta \mathbf{A}_{i,p}^3)$ by doing a PGA on $R_{\mathcal{T},i}^2 \ominus R_{\mathcal{T},i,p}^2$, where $R_{\mathcal{T},i,p}^2$ denotes the residue of atom i that is predicted by those of its neighbors. Let H_i^3 be the submanifold generated by a chosen number of stable principal geodesics. Then

$$A_{\mathcal{T},i}^3 = A_{\mathcal{T},i}^2 \oplus R_{\mathcal{T},i,p}^2 \oplus \mu_i^3 \oplus \text{Proj}_{H_i^3}(R_{\mathcal{T},i}^2 \ominus R_{\mathcal{T},i,p}^2). \tag{16}$$

This process is summarized in Algorithm 4 (with \mathcal{J} being the augmented set). For the two-bone model, this step corresponds to estimating atom residue deformations for the 36 atoms on the right bone and 5 augmenting atoms on the left bone. Each estimation involves an 8-dimensional PGA.

The atom residue deformations for object 1 will have an effect on other objects. This can be described by the augmentation-prediction scheme similar to the one described in the previous step. For the two-bone

¹Notice that the subscript 2 here corresponds to the step number, not the object index.

Input: $A_{\mathcal{T},\mathcal{J}}, A_{\mathcal{T},\mathcal{J}}^{k-1}$, where \mathcal{J} is the index set for the atoms in the current object

Output: $P^k =$ estimates of $\left\{ \Pr(\Delta \mathbf{A}_i^k \ominus \Delta \mathbf{A}_{i,p}^k) \right\}_{i \in \mathcal{I}}$, as well as $A_{\mathcal{T},\mathcal{J}}^k$

$$R_{\mathcal{T},\mathcal{J}}^{k-1} \leftarrow A_{\mathcal{T},\mathcal{J}} \ominus A_{\mathcal{T},\mathcal{J}}^{k-1}, \quad R_{\mathcal{T},\mathcal{J},p}^{k-1} \leftarrow A_{\mathcal{T},\mathcal{J},p} \ominus A_{\mathcal{T},\mathcal{J},p}^{k-1}$$

for all $i \in \mathcal{J}$ **do**

$P_i^k \leftarrow$ PGA results on $R_{\mathcal{T},i}^{k-1} \ominus R_{\mathcal{T},i,p}^{k-1}$ (mean μ_i^k , principal geodesics and eigenvalues)

$H_i^k \leftarrow$ sub-manifold generated by the first few principal geodesics

$$A_{\mathcal{T},i}^k \leftarrow A_{\mathcal{T},i}^{k-1} \oplus R_{\mathcal{T},i,p}^{k-1} \oplus \mu_i^k \oplus \text{Proj}_{H_i^k}(R_{\mathcal{T},i}^{k-1} \ominus R_{\mathcal{T},i,p}^{k-1})$$

end for

$$\text{Output } P^k := \{P_i^k\}_{i \in \mathcal{J}}, A_{\mathcal{T},\mathcal{J}}^k$$

Algorithm 4: Atom residue statistics for a single-figure object at step k .

example we have

$$\begin{aligned} A_{\mathcal{T},i}^3 &= A_{\mathcal{T},i}^2 \oplus R_{\mathcal{T},i,p}^2 \oplus \mu_i^3 \oplus \text{Proj}_{H_i^3}(R_{\mathcal{T},i}^2 \ominus R_{\mathcal{T},i,p}^2), \quad i \in \mathcal{I}_1 \\ A_{\mathcal{T},\mathcal{I}_2}^3 &= A_{\mathcal{T},\mathcal{I}_2}^2 \oplus \rho_{\mathcal{I}_2}^3 \oplus \text{Proj}_{\mathcal{H}^3}(\Delta A_{\mathcal{T},\mathcal{U}^3}^2) = A_{\mathcal{T},\mathcal{I}_2}^2 \oplus \rho_{\mathcal{I}_2}^3 \oplus \text{Pred}_3(\Delta \mathbf{A}_{\mathcal{I}_2}^3; \Delta \mathbf{A}_{\mathcal{U}^3}^3) \\ R_{\mathcal{T},\mathcal{I}}^3 &= A_{\mathcal{T},\mathcal{I}} \ominus A_{\mathcal{T},\mathcal{I}}^3 \end{aligned} \quad (17)$$

5.6 Object and Atom Residue Statistics for Other Objects

Suppose we describe object 2 next. Let \mathcal{I}_2 be the index set for its atoms. The object residue statistics can be obtained using Algorithm 2, with $k = 4$ and $\mathcal{J} = \mathcal{I}_2 \cup \mathcal{U}^4$, where \mathcal{U}^4 is the index set for augmenting atoms. The difference is that in this step, there is no deformation on object 1. Therefore

$$\begin{aligned} A_{\mathcal{T},\mathcal{I}_1}^4 &= A_{\mathcal{T},\mathcal{I}_1}^3 \\ A_{\mathcal{T},\mathcal{I}_2}^4 &= \text{results from Algorithm 2 (with } k = 4 \text{ and } \mathcal{J} = \mathcal{I}_2 \cup \mathcal{U}^4) \\ A_{\mathcal{T},\mathcal{I}'}^4 &= \text{results from Algorithm 3 (with } k = 4 \text{ and } \mathcal{J} = \mathcal{I}' = \mathcal{I} \setminus (\mathcal{I}_1 \cup \mathcal{I}_2)) \\ R_{\mathcal{T},\mathcal{I}}^4 &= A_{\mathcal{T},\mathcal{I}} \ominus A_{\mathcal{T},\mathcal{I}}^4. \end{aligned} \quad (18)$$

For the two-bone example, the left bone object residue deformation is estimated by a $36 * 8 = 288$ dimensional PGA. Since there are only two objects, there is no augmentation-prediction at this step (\mathcal{U}^4 and \mathcal{I}' are empty).

The atom residue statistics for object 2 is computed by Algorithm 4, with $k = 5$ and $\mathcal{J} = \mathcal{I}_2 \cup \mathcal{U}^5$. The prediction works the same way as in the previous step. Therefore,

$$\begin{aligned} A_{\mathcal{T},\mathcal{I}_1}^5 &= A_{\mathcal{T},\mathcal{I}_1}^4 \\ A_{\mathcal{T},i}^5 &= A_{\mathcal{T},i}^4 \oplus R_{\mathcal{T},i,p}^4 \oplus \mu_i^5 \oplus \text{Proj}_{H_i^5}(R_{\mathcal{T},i}^4 \ominus R_{\mathcal{T},i,p}^4), \quad i \in \mathcal{I}_2 \\ A_{\mathcal{T},\mathcal{I}'}^5 &= A_{\mathcal{T},\mathcal{I}'}^4 \oplus \rho_{\mathcal{I}'}^5 \oplus \text{Proj}_{\mathcal{H}^5}(\Delta A_{\mathcal{T},\mathcal{U}^5}^4), \quad \text{where } \mathcal{I}' = \mathcal{I} \setminus (\mathcal{I}_1 \cup \mathcal{I}_2) \\ R_{\mathcal{T},\mathcal{I}}^5 &= A_{\mathcal{T},\mathcal{I}} \ominus A_{\mathcal{T},\mathcal{I}}^5. \end{aligned} \quad (19)$$

Again, in the two-bone case, the left bone atom residue deformations are described by 36 individual 8-dimensional PGA's. There is no augmentation-prediction at this step (\mathcal{U}^5 and \mathcal{I}' are empty).

If there are more than two objects in the object complex, we can perform the same procedures and go through each of them one by one to generate the statistics we need.

5.7 Summary

The entire estimation process is summarized in the Algorithm 5.

Input: A training set $A_{\mathcal{T},\mathcal{I}}$, where \mathcal{I} is the index set for all atoms in the object ensemble, and \mathcal{T} is the index set for the training samples

Output: estimates of mean $\mathbf{A}_{\mathcal{T}}^0$ and $\Pr(\Delta\mathbf{A}_{\mathcal{T}}^k), k = 1, \dots, K$

Determine an order to traverse the objects (see Fig. 3). $K \leftarrow$ number of steps

$\mathbf{A}_{\mathcal{T}}^0 \leftarrow \text{mean}(A_{\mathcal{T},\mathcal{I}})$

$\{P^1, A_{\mathcal{T},\mathcal{I}}^1\} \leftarrow$ Algorithm 1

for $k = 2$ to K **do**

$\mathcal{M}^k \leftarrow$ index set for atoms in the current object

$\mathcal{N}^k \leftarrow$ index set for atoms in the rest of the objects that have not been visited

if step k is an object residue step **then**

$\{P^k, A_{\mathcal{T},\mathcal{M}^k}^k\} \leftarrow$ Algorithm 2

else {step k is an atom residue step}

$\{P^k, A_{\mathcal{T},\mathcal{M}^k}^k\} \leftarrow$ Algorithm 4

end if

$\{P^k, A_{\mathcal{T},\mathcal{N}^k}^k\} \leftarrow$ Algorithm 3

end for

Output $\mathbf{A}_{\mathcal{T}}^0$ and $\{P^k : k = 1, \dots, K\}$

Algorithm 5: Statistical training for multi-object ensemble.

5.8 Examples of Multi-object Statistics

We now present some statistics of a pair of pubic bones using the training procedure just described. The training population has 15 shape instances from different patients. For this illustration, we omit the atom residue stages and describe each instance at the ensemble, right bone object, and left bone object stages. Each bone is represented by 36 medial atoms on a 3×12 grid. The results are shown in Fig. 7.

At each stage, the base shape is the shape obtained by applying the mean residue at this stage to the shape template. In other words, the base shape at step k is $\mathbf{A}_{\mathcal{T}}^0 \oplus \mu^k$. The principal modes are shown relative to the corresponding base shape. These modes typically have intuitive interpretations. For instance, Fig. 7(b) shows the first three PGA modes for the ensemble residues. The first mode is mainly a size variation; the second mode reflects a simultaneous bending of the two bones; and the third mode corresponds to a

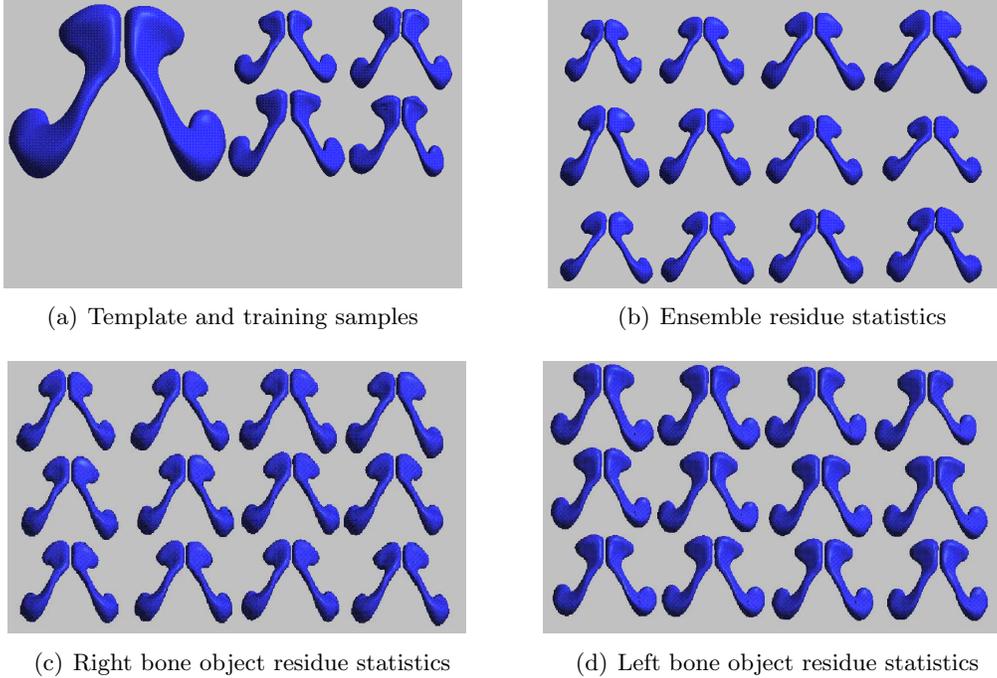


Figure 7: Statistics of the two-bone ensemble. (a) Some training samples (right) and the estimated shape template (left). (b)-(d) The first three PGA modes of variation for the different stages. In each picture, from top to bottom: the first, second, and the third mode. From left to right, PGA deformations corresponding to -2 , -1 , $+1$, and $+2$ standard deviations.

twisting of the lower part of the two bones. In Fig. 7(c), the residual deformation of the right bone takes place mostly at both ends, especially at the one that is closer to the left bone. In contrary, the left bone deformation mostly happens at the end that is further from the right bone. This is to be expected because the deformation at the other end has been mostly described in the previous stage through the augmenting atoms.

In the second example, we show statistics for a bladder-prostate-rectum complex across different days within a particular patient. The medial atom grids used to fit the organs are: 4×6 for the bladder, 3×4 for the prostate, and 3×7 for the rectum, and we described the objects in the above order. The training set consists of 12 cases of one patient at different days. Fig. 8 shows the ensemble statistics of the complex (a) and the residue statistics for each object (b)-(d).

The statistics can be used to generate random samples of the shape. We assume each principal mode follows the standard Gaussian distribution after we scale each principal direction by the square root of the corresponding eigenvalue in the tangent space. Although this is not exactly true, it is a reasonable approximation which gives us a way of quickly generating samples (refer to section 5.1). Thus, random

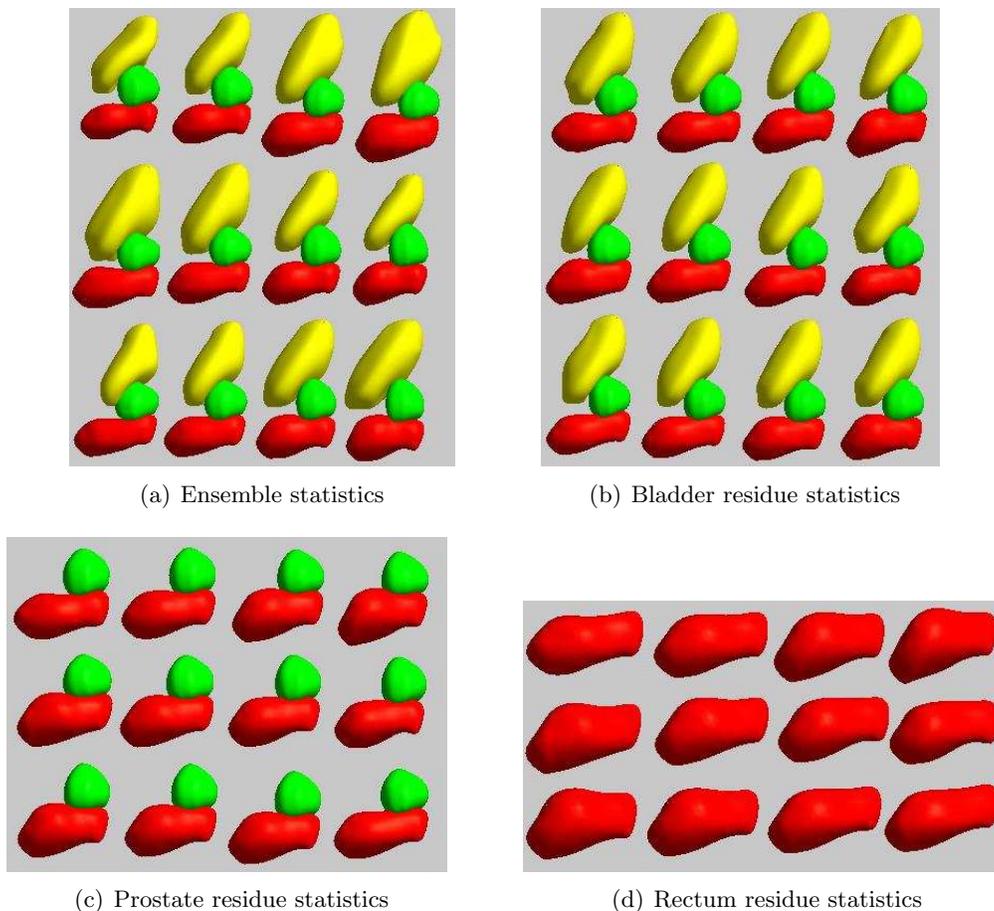


Figure 8: Statistics of the bladder-rectum-prostate complex. (a)-(d) The first three PGA modes of variation for the different stages. In each picture, from top to bottom: the first, second, and the third mode. From left to right, PGA deformations corresponding to -2 , -1 , $+1$, and $+2$ standard deviations.

shape samples can be obtained by starting from the template and successively going through each stage as follows: at step k , generate multivariate normal samples on the tangent space about the corresponding mean μ^k , take the exponential map to produce m -reps residues, and concatenate them to the previous-step m -reps using the \oplus operation. Fig. 9 shows a few samples for the two-bone shape. We have found that when we stay within $[-2,+2]$ standard deviations most of the samples preserve topology, do not have geometric singularities and have very little inter-penetration among objects. In other words, with high likelihood the samples are *geometrically proper* [27].

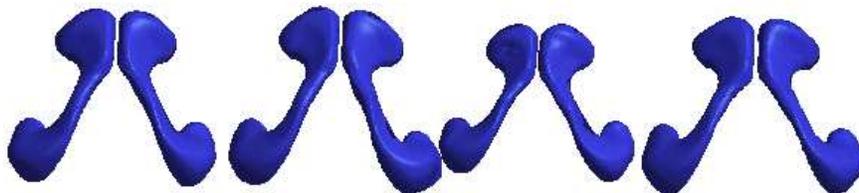


Figure 9: Random samples of the two-bone shape.

6 Segmentation using the Multi-object Statistics

The above framework can be used in image segmentation, where we seek an m-rep model that best fits a given image I . In the Bayesian framework, this can be done by maximizing the log-posterior

$$\log p(\mathbf{A}|I) \propto \log p(I|\mathbf{A}) + \log p(\mathbf{A}).$$

The prior distribution $p(\mathbf{A})$ is given by Eq. (8), and we use PGA statistics $\{P^k : k = 1, \dots, K\}$ generated by the training process to approximate this prior. Here we assume a likelihood distribution $p(I|\mathbf{A})$ is available. For details on estimating this probability refer to [31, 32].

To segment an image, we first apply a similarity transformation on \mathbf{A}_I^0 so that it roughly fits the image. Then we go through the object complex in the same order as when we obtained the statistics. At the object ensemble step, the ensemble geometry is deformed simultaneously by an overall similarity transformation and a transformation in the ensemble principal geodesic space H^1 . At each subsequent step, we deform a particular object either at the object level or at the atom level, using the appropriate PGA statistics as prior, until all objects have been deformed.

Since each atom residue step is described by a set of conditional distributions

$$\left\{ \Pr(\Delta \mathbf{A}_i^k | \Delta \mathbf{A}_{N(i)}^k) \right\} \sim \left\{ \Pr(\Delta \mathbf{A}_i^k \ominus \Delta \mathbf{A}_{i,p}^k) \right\},$$

the optimization process at this step uses an Iterative Conditional Modes (ICM) algorithm.

Fig. 10 shows an example of segmenting a two-bone shape from a CT image. Although this is a relatively simple application, it does illustrate how the method works and provides with a basic test case. Our methodology has also been applied to some other more challenging problems, in particular segmentation of male pelvis, and have produced promising results [27, 33].

7 Discussion

We illustrated a new method of describing shapes that consist of multiple geometric objects. Based on our m-reps representation, the geometric features of various scales and inter-object relationships are characterized by a multi-scale probabilistic model, which is in turn characterized by a series of probability distributions on residues. This method has the advantage of being able to provide accurate, intuitive

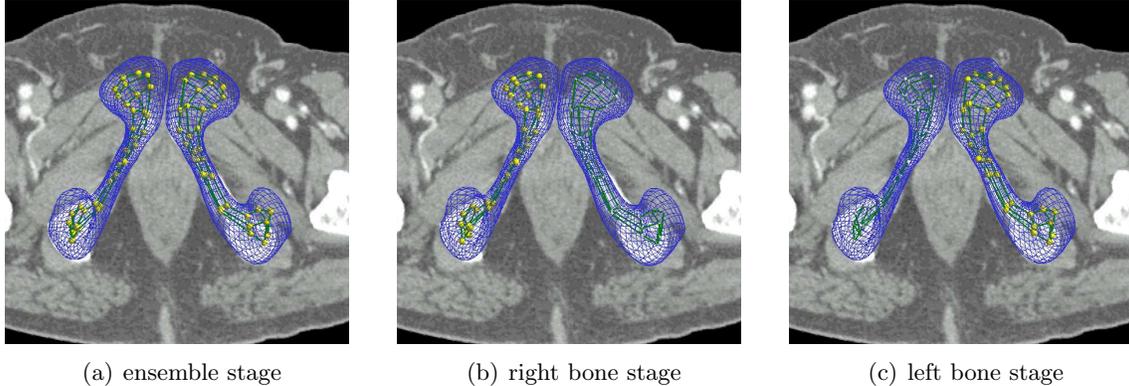


Figure 10: Segmentation of the two-bone shape.

geometric information and efficiency in statistical training. The shape model can be effectively used in a variety of image analysis applications, such as segmentation, shape discrimination, and atlas building.

As we have mentioned, the method described in the previous sections depends on the order we choose to traverse the multi-object complex. Now we briefly describe an alternative that does not specifically requires a pre-determined order.

To explain, let us consider how we might use the probability model for segmentation. The idea is to maximize the joint posterior distribution of all objects, $\Pr(O_1, O_2, \dots, O_M | I)$. In the method described earlier, we break this distribution into a product of distributions (see Eq. (3)). An alternative approach is to use the iterative conditional mode (ICM) method [34], where one iteratively maximizes the conditional distributions $\Pr(O_i | \{O_{j \neq i}\})$. It has been shown that under suitable conditions, the solution to this approach provides a good approximation to the maximizer of the original joint distribution.

Therefore, we can focus on learning the conditional distributions $\Pr(O_i | \{O_{j \neq i}\})$, which can be described by the difference between the actual deformation of object i and the predicted deformation of it based on those of its neighbors. Clearly, this approach does not depend on the order in which we choose to traverse the objects.

The basic assumptions in our methodology are that geometric entities can be described by a series of successive residual deformations from a common template and that these residual deformations can be treated as conditionally independent. Ultimately, the questions of in which order we describe objects and how many PGA modes to use at each step come down to whether our choices make those assumptions valid. Depending on the situation, it may make sense to have schemes that have only global deformations or only local deformations, or have schemes going back and forth between scale levels.

The ideas of residual deformation and prediction can also be applied to objects made up of multiple m-reps figures, in which case the hinge atoms on the subfigures are the natural choices for augmenting atoms. For details refer to [35].

Our tests have suggested that the statistical shape model learned from training sets produce visually

valid samples. However, we need a more rigorous way of evaluating these probability models and comparing them with other shape models. This is still work in progress. Other questions that require further investigation include the following: method for determining augmenting atoms, prediction algorithms, and object-based image match models.

Acknowledgements

We thank Gregg Tracton, Derek Merck, Julian Rosenman for providing image data and preparing some of the m-reps segmentations of them. We also thank Qiong Han, P. Thomas Fletcher, and Keith Muller for their help. This work was done with support from NIH grant P01 EB02779.

References

- [1] U. Grenander. *Elements of Pattern Theory*. Johns Hopkins University Press, 1995.
- [2] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1987.
- [3] U. Grenander, Y. Chow, and D.M. Keenan. *HANDS: A Pattern Theoretic Study of Biological Shapes*. Springer-Verlag, New York, 1991.
- [4] S. C. Zhu. Embedding Gestalt laws in the Markov random fields – a theory for shape modeling and perceptual organization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(11), 1999.
- [5] S. M. Pizer, J. Z. Chen, P. T. Fletcher, Y. Fridman, D. S. Fritch, A. G. Gash, J. M. Glotzer, M. R. Jiroutek, S. Joshi, C. Lu, K. E. Muller, A. Thall, G. Tracton, P. Yushkevich, and E. L. Chaney. Deformable m-reps for 3D medical image segmentation. *International Journal of Computer Vision*, 55(2):85–106, November-December 2003.
- [6] G. Gerig, M. Styner, M.E. Shenton, and J.A. Lieberman. Shape versus size: improved understanding of the morphology of brain structures. In *Proc. MICCAI 2001*, volume 2208 of *Springer LNCS*, pages 24–32. Springer-Verlag, 2001.
- [7] S. Joshi. *Large deformation diffeomorphisms and Gaussian random fields for statistical characterization of brain submanifolds*. PhD thesis, Washington University, 1997.
- [8] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, E. Grimson, and A. Willsky. A shape-based approach to curve evolution for segmentation of medical imagery. *IEEE T-MI*, 22(2):137–154, February 2003.
- [9] K. Siddiqi, S. Bouix, A. R. Tannenbaum, and S. W. Zucker. Hamilton-Jacobi skeletons. *Int. J. Computer Vision*, 48(3):215–231, 2002.

- [10] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks, and deformations I: the components of two-dimensional shape and the reaction-diffusion space. *Int. J. Comput. Vision*, 15:189–224, 1995.
- [11] C. Lu, Y. Cao, and D. Mumford. Surface evolution under curvature flows. *J. Visual Communication and Image Representation*, 13:65–81, 2002.
- [12] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *Fifth European Conference on Computer Vision*, pages 484–498, 1998.
- [13] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
- [14] A. Kelemen, G. Szekely, and G. Gerig. Three-dimensional model-based segmentation. *IEEE-TMI*, 18(10):828–839, 1999.
- [15] P. Yushkevich, S.M. Pizer, S. Joshi, and J.S. Marron. Intuitive, localized analysis of shape variability. In M.F. Insana and R.M. Leahy, editors, *Information Processing in Medical Imaging (IPMI)*, pages 402–408, 2001.
- [16] S. G. Mallat. Multifrequency channel decompositions of images and wavelet models. *IEEE Trans. Acoust. Speech, Signal Processing*, 37(12):2091–2110, 1989.
- [17] M. Unser. A review of wavelets in biomedical applications. *Proceedings of the IEEE*, 84(4):626–638, April 1996.
- [18] S. Joshi, S. Pizer, P. T. Fletcher, P. Yushkevich, A. Thall, and J. S. Marron. Multi-scale deformable model segmentation and statistical shape analysis using medial descriptions. *IEEE-TMI*, 21(5), 2002.
- [19] M. Styner and G. Gerig. *IPMI '01*, volume 2082 of *LNCS*, chapter Medial models incorporating object variability for 3D shape analysis, pages 502–516. Springer, 2001.
- [20] H. Blum and R. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10(3):167–180, 1978.
- [21] S. M. Pizer, K. Siddiqi, G. Szekely, J.N. Damon, and S.W. Zucker. Multiscale medial loci and their properties. *Int. J. Computer Vision*, 55(2):155–179, November-December 2003.
- [22] A. Thall. Fast C^2 interpolating subdivision surfaces using iterative inversion of stationary subdivision rules. http://midag.cs.unc.edu/pub/papers/Thall_TR02-001.pdf.
- [23] S. M. Pizer, P. T. Fletcher, A. Thall, M. Styner, G. Gerig, and S. Joshi. Object models in multi-scale intrinsic coordinates via m-reps. *Image and Vision Computing, Special Issue on Generative Model-based Vision*, 21(1):5–15, 2003.

- [24] D. Merck, S. Joshi, G. Tracton, and S.M. Pizer. On single figure statistical m-rep model construction. In preparation, 2005.
- [25] R.H. Davies, C.J. Twining, T.F. Cootes, J.C. Waterton, and C.J. Taylor. A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging*, 21(5):525–537, 2002.
- [26] C. Lu, S.M. Pizer, and S. Joshi. A markov random field approach to multi-scale shape analysis. In L.D.Griffin and M. Lillholm, editors, *Scale Space Methods in Computer Vision*, volume 2695 of *LNCS*, pages 416–431. Springer-Verlag, 2003.
- [27] S.M. Pizer, J. Jeong, C. Lu, K. Muller, and S. Joshi. Estimating the statistics of multi-object anatomic geometry using inter-object relationships. To appear, International Workshop on Deep Structure, Singularities and Computer Vision (DSSCV) 2005.
- [28] P.T. Fletcher and S. Joshi. Principal geodesic analysis on symmetric spaces: statistics of diffusion tensors. In *ECCV 2004 Workshop on Computer Vision Approaches to Medical Image Analysis (CVAMIA)*, 2004.
- [29] P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Medical Imaging*, 23(8):995–1005, August 2004.
- [30] P.T. Fletcher, C. Lu, and S. Joshi. Statistics of shape via principal component analysis on Lie groups. http://midag.cs.unc.edu/pub/papers/CVPR03_Fletcher_StatsShape.pdf.
- [31] S. Ho and G. Gerig. Profile scale-space for multi-scale image match. In C. Barillot, D. Haynor, and P. Hellier, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 176–183, Germany, 2004.
- [32] J. Stough, S.M. Pizer, E. Chaney, and M. Rao. Clustering on image boundary regions for deformable model segmentation. In *International Symposium on Biomedical Imaging (ISBI)*, pages 436–439, Piscataway, NJ, April 2004.
- [33] E. Chaney, S.M. Pizer, R. Joshi, S. Broadhurst, P.T. Fletcher, G. Gash, Q. Han, J. Jeong, C. Lu, D. Merck, J. Stough, G. Tracton, J. Rosenman, Y. Chi, and K. Muller. Automatic male pelvis segmentation from ct images via statistically trained multi-object deformable m-rep models. In *American Society for Therapeutic Radiology and Oncology (ASTRO)*, 2004.
- [34] J. Besag. On the statistical analysis of dirty pictures. *J. Royal Stat. Soc. B*, 48(3):259–302, 1986.
- [35] Q. Han, C. Lu, G. Liu, S. Pizer, S. Joshi, and A. Thall. Representing multi-figure anatomical objects. In *International Symposium on Biomedical Imaging*, pages 1251–1254, 2004.