Proper Shape Representation of Single Figure and Multi-Figure Anatomical Objects

Qiong Han

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
2009

Approved by:

Stephen M. Pizer, Advisor

Edward L. Chaney, Reader

Sarang Joshi, Reader

James N. Damon, Committee Member

Keith Muller, Committee Member

Jack Snoeyink, Committee Member

Abstract

**Qiong Han: Proper Shape Representation of Single Figure and
Multi-Figure Anatomical Objects
(Under the direction of Stephen M. Pizer)**

Extracting anatomic objects from medical images is an important process in various
medical applications. This extraction, called image segmentation, is often realized by
deformable models. Among deformable model methods, medial deformable models have
the unique advantage of representing not only the object boundary surfaces but also the
object interior volume.

Based on one medial deformable model called the m-rep, the main goal of this dis-
sertation is to provide proper shape representations of simple anatomical objects of one
part and complex anatomical objects of multiple parts in a population. This dissertation
focuses on several challenges in the existing medially based deformable model method:
1. how to derive a proper continuous form by interpolating a discrete medial shape rep-
resentation; 2. how to represent complex objects with several parts and do statistical
analysis on them; 3. how to avoid local shape defects, such as folding or creasing, in
shapes represented by the deformable model. The proposed methods in this dissertation
address these challenges in more detail:

1. An interpolation method for a discrete medial shape model is proposed to guaran-
    tee the legality of the interpolated shape. This method is based on the integration
    of medial shape operators.

2. A medially based representation with hierarchy is proposed to represent complex
    objects with multiple parts by explicitly modeling interrelations between object
    parts and modeling smooth transitions between each pair of connected parts. A
    hierarchical statistical analysis is also proposed for these complex objects.

3. A method to fit a medial model to binary images is proposed to use an explicit legality penalty derived from the medial shape operators. Probability distributions learned from the fitted shape models by the proposed fitting method have proven to yield better image segmentation results.

Acknowledgments

First and foremost, I want to thank Steve Pizer, for his support and guidance. His almost "contagious" passion in our research of the Medical Image Display and Analysis Group (MIDAG) partially contributes to the existence of this dissertation. My thanks also go to Ed Chaney, who provided the important driving problems for this dissertation, Jim Damon, who taught me the medial mathematics in great detail, Sarang Joshi, who guided me on the work of the multi-figure m-rep, Keith Muller, who suggested a systematic way to conduct an evaluation, starting from a well designed test sample, and Jack Snoeyink, who helped me on writing this document.

I feel fortunate to work in the MIDAG. The support provided by Graham Gash and Gregg Tracton has been impeccable. I also want to thank Guido Gerig, Mark Styner, Julian Rosenman, and all my former colleagues in the MIDAG, especially Andrew Thall, Thomas Fletcher, Ja-Yeon Jeong, Xiaoxiao Liu, Derek Merck, and Rohit Saboo. The discussions with you, the comments from you, the opportunity to shadow you around have all become valuable pieces in my life as a graduate student in Chapel Hill.

I also want to thank Janet Jones and Timm Quigg for their help when I needed it the most.

Finally, I thank my family: my sister Lan in my hometown, my parents in Beijing, my dear wife Jinze, and little Erik, for their unconditional love and support. Thank them for being there for me for all these years.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

Medical image processing is becoming an important component in modern medicine because it assists physicians in diagnosing and treating diseases. The following list shows some medical image processing applications:

- Computer-assisted nodule detection in chest CT images helps physicians to screen for early-stage lung cancer, which significantly increases the survival rate of patients with early-stage lung cancer;

- Statistical shape analysis applied to neurological images sheds light on early diagnosis of children with autism;

- Image-guided radiation therapy (IGRT) helps radiation oncologists to locate tumors and organs at risk more precisely and plan better radiation treatment plans.

These medical image processing applications are based on two fundamental tasks: image segmentation and statistical shape analysis. The first task, image segmentation, tries to separate image regions corresponding to certain anatomical structures from the background image. Various image segmentation methods have been proposed. Deformable model methods tackle the problem by representing anatomical objects as

geometric shape models and deforming shape models into target images via object shape and object-relative image intensity (*appearance*) information inferred from a population of anatomical objects. The use of the inferred information from a population enables deformable model methods to deal with challenging images, in which noise is apparent or contrast across object boundaries is lacking. The second task, statistical shape analysis, applies learned shape statistics on two populations of shapes to infer differences between the populations, e.g., by hypothesis testing.

Both tasks require the understanding of object shapes or appearance in one or more populations. This is achieved by learning object shape or appearance variations of a training set of anatomical objects from a respective object population. There are two common approaches to use a set of training objects: 1. estimating shape or appearance probability distributions from training objects or images from a population and then computing the estimated probability on a new object for statistical shape analysis or on a new image for image segmentation; 2. using the training set of anatomical objects to solve problems without learned shape or appearance probability distributions. The first approach is often used in posterior optimization for image segmentation or Bayesian classification, and the second approach is used in shape discrimination, via support vector machine (SVM) or distance weighted discriminant methods [Vapnik (1993)]. This dissertation focuses on the first approach, which summarizes shape and appearance statistics from a training set of objects and applies the estimated probability distributions to a new object.

In the framework of UNC MIDAG (Medical Image Display and Analysis Group), various projects use learned statistics: shape and appearance statistics are used to guide image segmentation [Pizer et al. (2003)]; shape statistics are also used in the discrimination between populations of object complexes by statistical shape analysis [Gorczowski et al. (2007)]; in [Terriberry et al. (2005)], shape statistics of two populations are used in a two-sample hypothesis testing to find significant differences between the

two populations. Our experiences indicate that quality of the shape and appearance statistics is crucial to the success of both image segmentation and shape analysis tasks.

Shape statistics are learned from shape parameters of a set of pre-extracted shape models. Appearance statistics are learned from shape-relative volumetric regions around shape boundaries. How the shapes are represented directly affects how shape and appearance information is obtained for each training object. Therefore the quality of both shape and appearance statistics depends on the quality of the shape models.

Various shape model representations [Cootes et al. (1995); Pizer et al. (2003); Styner et al. (2006); Yang et al. (2003)] have been proposed for anatomical objects, with modeling an object population as the goal in mind. In particular, the UNC MIDAG group has developed a discrete medial shape representation called the *m-rep* [Pizer et al. (2003)] that is especially suited to this population modeling task: discrete m-reps have localized shape control allowing accurate and precise representation of anatomical structures; discrete m-reps can represent objects with a fixed topology from a population and can capture non-linear shape variations in these objects via learned shape statistics [Fletcher (2005)]. However, the following challenges still remain:

1. Undesired shape defects, such as creasing, can not be efficiently prevented;

2. Objects with more than one component, such as a liver with two lobes, are only represented as a whole, not in terms of its components;

3. Interpolation within discrete medial shape models is not well understood and implemented.

The following subsections detail these challenges, followed my thesis objectives and claims to face these challenges.

1.1.1 Proper Shape for Population Modeling

Any undesired shape defects in training shape models are unnatural shape features, such as local creasing or folding. These shape defects will be reflected in inferred statistics based on those models. The "tainted" statistics will yield undesirable results.

Shape representations in deformable model methods provide not only the geometric basis to infer shape statistics but also the object-relative coordinates via which the image intensities are extracted to infer appearance statistics. A shape representation that guarantees shapes as *legal* or *proper*, i.e. free of any local creasing and folding, are necessary to obtain both good shape and good appearance statistics. These shape and appearance statistics are used as priors and likelihoods, respectively, in image segmentations, via posterior optimization on deformable models. Shape statistics are also used in statistical shape analysis applications such as hypothesis testing. Shape legality is thus important to the success of image segmentation using posterior optimization and statistical shape analysis.

However, it is not trivial to maintain shape legality at feasible computational cost for most proposed shape representations. Among shape representations, medially based ones are richer in descriptive power than boundary surface representations by representing not only the object boundary but also the volumetric object interior. Medial representations can be understood as multiscale descriptions that represent not only global shape deformations but also local translation, bending, scaling, and twisting of the object interior. Focusing on the object interior has led to mathematics that provide explicit legality constraints on medially based shape models [Damon (2003, 2004, 2005)]. Constraints are key to acquiring proper shapes and generating correct statistics of an object population. Providing and applying such shape legality constraints is one main goal of this dissertation [Han et al. (2007)].
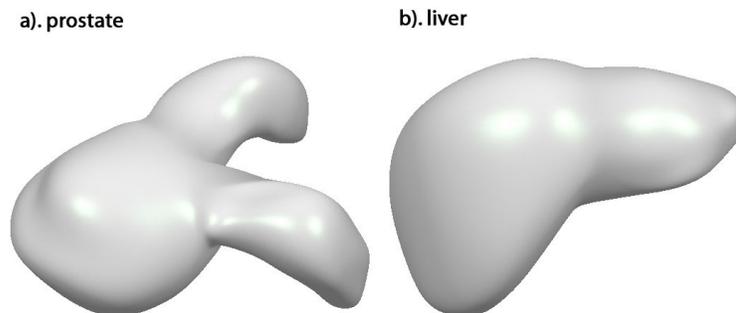
Figure 1.1: **a**) A prostate with two seminal vesicle protrusions. **b**) A liver represented by the union of the left and right lobes. Object a has three single-figure parts while object b has two such parts.

1.1.2 Multi-part Objects

Shape representation and statistics of simple one-part objects have been widely studied. Methods using various representations have been proposed and shown to be effective [Cootes et al. (1995); Fletcher et al. (2004); Yang et al. (2003)]. However, many anatomical objects have multiple named parts. For example, a prostate (Figure 1.1-**a**) has two seminal vesicles attached to it, and a liver (Figure 1.1-**b**) has both left and right lobes.

The difficulty of representing multi-part objects lies in the fact that only modeling each individual part is not sufficient. The inter-relations among parts have been elusive for boundary based shape representations. The inter-relations include part-to-part connection, global deformations to all the connected parts, and implied deformations from one part to an adjacent part.

Due to the inherent complexity of objects with multiple parts, previous shape and statistical descriptions of such objects concentrated on their global structures [Cootes et al. (1995); Gerig et al. (2001)] or on the extremely local behavior of geometric primitives, such as points, without references to the parts' inter-relations [Csernansky et al. (1998); Styner et al. (2003)].

One focus of this dissertation is to represent complex objects with multiple parts so that statistical analysis can be effectively conducted on such objects from a population.

5

Medially based shape representations parameterize both the object interior and the adjacent exterior volumetric regions. This allows us to model the relations between pairs of adjacent object parts because each part can be known in the local coordinates of its connected neighboring part. The local coordinates also allow us to describe the inter-part relations by a small set of parameters. This makes it feasible to explicitly model multi-part objects as a connected set of parts [Han et al. (2004)] and to conduct hierarchical statistical analysis on the parts of the objects from a population [Han et al. (2005)].

1.1.3 Interpolation of Discrete Medial Representations

A parametric shape representation captures an object shape via continuous functions of arguments on topologically equivalent shapes. For shapes with spherical topology, an example is a shape representation based on piece-wise polynomial splines or a shape representation based on Fourier basis functions [Duncan (1992); Staib and Duncan (1996)] that globally represents an object shape with a set of parameters as the coefficients of the basis functions. Such representations, however, do not provide a means to model localized shape variations.

A discrete medial shape representation, such as a discrete *m-rep*, is composed of a set of discrete medial samples, called *medial atoms* [Pizer et al. (2003)]. Each medial atom has a hub position point and two to three spoke vectors from the hub position to the boundary points, as shown in figure 2.1. Although the discrete m-rep does not give a continuously parameterized representation of the object, it conveniently provides localized control over the object shape. By interpolating or approximating the set of discrete primitives, we can have a continuous and parameterized representation derived from the discrete object, and we can still benefit from capturing the localized shape variations by the discrete primitives.

An m-rep medial atom has the form of $(\mathbf{p}, r, \mathbf{U}^{+1,-1})$, where $\mathbf{p}$ is the hub position

as a 3D point, r is the radius or the length of spokes, and $\mathbf{U}^{+1,-1}$ are the directions of the spokes (more detail about medially based representations can be found in section 2.1). An interpolation or approximation of $(\mathbf{p}, r, \mathbf{U}^{+1,-1})$ can provide a parameterized medial representation $(\mathbf{p}(u,v), r(u,v), \mathbf{U}(u,v)^{+1,-1})$. In order for the interpolated atoms to stay medial [Blum and Nagel (1978)], the interpolated or approximated atoms must satisfy a set of constraints as follows:

1. The bisector spoke of each atom is tangential to *the medial sheet surface* formed by $\mathbf{p}(u,v)$ at the point of which the atom hub $\mathbf{p}$ lies;

2. The normal of the medial sheet surface is in the direction of the difference vector between the pair of spokes in a medial atom;

3. The derivatives of the spoke along the medial sheet and the medial sheet surface fulfill certain conditions to keep the spokes from crossing one another.

Although interpolation of points in $\mathbb{R}^3$ has been well studied, a new interpolation method is needed for medial atoms because the interpolated medial atoms must fulfill the set of constraints above. In particular, each discrete medial primitive, i.e., each medial atom lies in a curved Riemannian space, and the dimensions of each atom are not independent because its spoke directions must reflect the derivatives of the medial sheet surface and of the radius field on the medial sheet. There has not been an interpolation method on discrete m-reps that fulfill such constraints. A method to interpolate m-rep atom hubs into a smooth medial sheet surface and to interpolate m-rep spokes into a smooth spoke field is proposed in this dissertation. The key is to maintain the relation among all the atom dimensions so the legality of the interpolated medial sheet and spoke fields on the sheet is upheld.

A direct application of the proposed interpolation method is to improve correspondences between sample objects represented by m-reps in a population. When modeling a population of objects, correspondence across models becomes important to efficiently

represent the shape variations in the population by shape statistics. Shape representations using dense boundary points have already approached this correspondence problem via an interpolation of boundary points [Davies et al. (2002); Gerig et al. (2001)]. The improvement to the correspondence across objects represented by m-reps relies on this proposed interpolation method in this dissertation.

Complex objects, such as livers with more than one component, are more challenging to represent. This dissertation will propose a means to represent such objects by establishing a hierarchy of connected parts. The connection between each pair of object parts also requires interpolations, and this dissertation will propose a method for such interpolation by a similar interpolation scheme as that within a medially represented object part.

This dissertation thus concentrates on proper shape representations of anatomical objects, including simple objects and complex objects with multiple parts, based on an interpolation method of medial spokes.

1.2   Thesis Objectives and Claims

Thesis claim: *An interpolation method, based on sound mathematics, and the hinge geometry, defining the interrelations between connected object parts, provide a powerful tool to derive a medial representation of both single part and multi-part objects from discrete m-reps. When incorporated with explicit geometric constraints, based on the same mathematics, the shape models based on discrete m-reps guarantee proper shape representations of both single and multi-part objects. These shapes allow statistical modeling of both shape and appearances in a population.*

The contributions of this dissertation are as follows:

- A method to interpolate a discrete medial model into a continuously parameterized one for both simple objects with one part and complex objects with multiple parts

- A description of geometric interrelations among adjacent parts for an object with multiple parts

- A means to calculate hierarchical statistics of an object of multiple parts, from the global level to the local level of each parts and their interrelations

- An explicit use of an illegality penalty in binary fitting to improve the smoothness of fitted shapes, thereby leading to properly trained shape and appearance statistics and eventually better segmentation results

- A way to generate standard test cases of deformable m-reps via synthesizing warped ellipsoids or via sampling learned shape space

1.3   Dissertation Outline

The rest of my dissertation is organized as follows.

Chapter 2 reviews the related background literature including deformable model methods in general for modeling a population of objects, medially based shape representations, and the mathematics for medial structures.

Chapter 3 details the interpolation method on a medial shape representation. For single part objects, the interpolation is applied to internal and end primitives. For objects with multiple parts, the interpolation is based on the interpolation of each part, with smooth transitions between each pair of adjacent parts generated via so-called *spoke interpolation*.

Chapter 4 describes the representation of objects with multiple parts in detail, including the definition of the interrelations among object parts, the coordinate system to describe that interrelations, part deformation implied by adjacent part, and self-deformation of parts without affecting adjacent parts.

Chapter 5 shows an application of the proposed method: using explicit geometric

legality constraints in the binary fitting process of fitting a template model into target binary images.

Chapter 6 describes the validation of the atom interpolation methods in the binary training process. The methods are evaluated on both synthetic and real world objects.

Chapter 7 begins with a discussion, provides conclusions, and indicates future work.

Chapter 2

Background

As the driving task of this dissertation, medical image segmentation is the process of delineating anatomical structures from the background in an image. Segmenting anatomical objects from a medical image containing information of complex anatomical structures, already a challenge, is made even more difficult by image noise, sampling artifacts, lack of contrast at object boundaries, and neighboring anatomical structures with confusing intensities.

Deformable model methods are designed to overcome many of these difficulties in image segmentation. A deformable model method segments an image by deforming a geometric shape model into the image via the optimization on an objective function. Various deformable model methods will be reviewed in section 2.1. Among the proposed methods, medially based methods are the focus of this dissertation. Section 2.2 reviews three medially based deformable model methods. Sections 2.3 and 2.4 review the mathematics related to the contributions of this dissertation including medial *atom interpolation*, representation of complex objects by *multi-figure m-reps*, and preparation of m-rep models for statistical training, described in chapters 3 to 5 respectively.

2.1 Deformable Model Methods

Early segmentation methods, such as the one proposed in [Marr and Nishihara (1978)], used local image features such as edges and corners to construct a shape model in a bottom-up fashion. Local edge or corner detectors applied to an image can, however, produce false edges or gaps. These false edges can become parts of a segmented object boundary that do not exist in a real anatomical object; these false gaps can distort or break a segmented boundary. Such a problem can be overcome by incorporating larger scale information in the image or a global model of the target object. Although such information has been used for boundary delineation via grouping, labeling, and scale-space analysis in [Kass et al. (1988)], this method did not guarantee finding a complete object boundary. Furthermore, these early segmentation methods were mostly applied to 2D image slices and are conceptually and computationally expensive to extend to 3D.

Where the local feature based methods fail is exactly where deformable model methods [McInerney and Terzopoulos (1996); Caselles et al. (1997)] shine. Deformable model methods use information gathered from different scale levels. Recent deformable model methods use learned shape probability distributions to guide their segmentation process. In general, deformable model methods segment images in a top-down fashion, which is more likely to find the global optimal object boundary in an efficient way.

Deformable model methods can be categorized by their model representations and the means by which their image segmentation, i.e., the process of fitting a template model into a target image, is conducted.

Among deformable model representations, the point distribution model (PDM) [Dryden and Mardia (1998); Kass et al. (1988)] using a set of landmarks is among the earliest. Numerous boundary points formed landmarks in the shape representations of later methods, such as the active shape model (ASM) [Cootes et al. (1995)] and active appearance model (AAM) [Cootes et al. (2001)] methods.

Some other deformable model methods use parametric models based on projection

onto orthogonal functions. They represent object boundaries by the coefficients of orthogonal basis functions. Choices of the orthogonal basis functions include sinusoidal functions [Duncan (1992)] and spherical harmonic functions [Gerig et al. (2001)].

There are also deformable model methods using medially based shape representations, including the ones following Blum's notation for medial axis [Blum and Nagel (1978)], such as in [Yushkevich et al. (2003); Terriberry et al. (2007)], and the discrete *m-rep* method [Pizer et al. (2003)], in which each shape representation primitive is the combination of a sample from a Blum medial axis and extra components.

These deformable model methods use quite different shape representations. They also approach the image segmentation problem differently. The earlier deformable model methods [Dryden and Mardia (1998); Kass et al. (1988)] segment medical images by deforming a template model $\mathbf{M}$ into an image $\mathbf{I}$ via optimizing an objective function summing two terms: image match F_{img} and geometric typicality F_{geom}. The image match term F_{img} measures how a deformed model $\mathbf{M}'$ fits into a given target image $\mathbf{I}$, and the geometric typicality term penalizes peculiar shapes with creased or folded surfaces to prevent or reduce shape defects in a fit shape model. I call this type of deformable model methods the geometric type.

To guide their segmentation process, recent deformable model methods [Cootes et al. (1995); Mitchell et al. (2002); Pizer et al. (2003)] use learned probability distributions of both shape and appearance, i.e., image intensity relative to the object model. Based on the Bayes' theorem [Bayes (1958)], the posterior probability $p(\mathbf{M}|\mathbf{I})$ can be factored into two terms: the shape prior $p(\mathbf{M})$ and the image likelihood $p(\mathbf{I}|\mathbf{M})$, shown in equation 2.1, where $p(\mathbf{I})$ is the marginal probability of a given image and is constant for a given $\mathbf{M}$:

$$\log p(\mathbf{M}|\mathbf{I}) = \log \frac{p(\mathbf{M})p(\mathbf{I}|\mathbf{M})}{p(\mathbf{I})} = \log p(\mathbf{M}) + \log p(\mathbf{I}|\mathbf{M}) - \log(\mathbf{I}) \qquad (2.1)$$

Equation 2.1 suggests a probabilistic framework for image segmentation. Deformable

13

Method	Shape representation	Segmentation
Method using boundary point	Boundary points representation (B-rep)	Geometric
Active contour	Point distribution model (PDM B-rep)	Geometric
Geodesic active contour	PDM B-rep	Geometric or probabilistic
ASM/AAM	PDM B-rep	Probabilistic
Fourier basis function based method	Coefficients of Fourier basis functions	Probabilistic
SPHARM	Coefficients of spherical harmonics	Probabilistic
Cm-rep	Splines on $\{x, y, z, rad\}^*$, only for 2D objects	Probabilistic
Cm-rep based on subdivision surfaces	Subdivided surface of of $\{x, y, z, rad\}^*$	Probabilistic Probabilistic
Discrete m-rep	Sampled medial atoms	Probabilistic

Table 2.1: Categorization of deformable model methods. * rad indicates a radius scalar field on the 3D surface $\{x, y, z\}$

model methods within such a framework segment an image by optimizing the posterior probability, calculated as the sum of the shape prior and image likelihood terms from the learned shape and appearance statistics. I call such group of deformable model methods the probabilistic type.

The geometric type of deformable model methods is still widely used to extract shape models for learning shape and appearance statistics, because no shape or appearance statistics are available at this step of learning. However, the probabilistic type of deformable model methods are used to crack the most difficult image segmentation problem using the learned probability distributions.

Some previously proposed deformable methods are listed in table 2.1. Details of these methods will be reviewed in the coming subsections 2.1.1 to 2.1.4.

2.1.1 Active Contour and Geodesic Active Contour Methods

The discrete active contour model, also known as the *snake*, was first proposed by Kass, et al. in [Kass et al. (1988)]. The method is named "snake" because the way an active contour deforms resembles the movement of a snake.

An active contour $\mathbf{V}$ is a set of vertices $\mathbf{p}_i$ on an object boundary. An energy function is defined for $\mathbf{V}$ with two parts (shown in 2.5) that the authors of [Kass et al. (1988)] called the internal and external part, which respectively correspond to the geometric typicality term and image match term in an objective function for the geometric type of deformable model methods.

$$E(\mathbf{V}) = \alpha E_{int}(\mathbf{V}) + \beta E_{ext}(\mathbf{V}); \qquad (2.2)$$

$$E_{int}(\mathbf{V}) = \gamma E_{continuity}(\mathbf{V}) + \delta E_{balloon}(\mathbf{V}); \qquad (2.3)$$

$$E_{ext}(\mathbf{V}) = \eta E_{intensity}(\mathbf{V}) + \phi E_{gradient}(\mathbf{V}); \qquad (2.4)$$

The internal energy E_{int}, i.e., the geometric typicality term, enforces $\mathbf{V}$ to stay smooth and propagate in a desired direction. There are two sub-terms in E_{int}: the continuity (or smoothness) term and the balloon force term. The weight γ for the continuity term controls the smoothness of a deformed contour $\mathbf{V}'$. If γ is too big, the resulting contour $\mathbf{V}'$ will be over-smoothed and fail to fit into the image well enough; if γ is too small, $\mathbf{V}'$ will be too jaggy with sharp features such as corners. The balloon force term ensures that a snake curve keeps propagating in a desired direction, either expanding or shrinking depending on the specific target object. These two internal energy terms force a snake curve to propagate smoothly regardless of the image intensity information.

The external energy E_{ext}, i.e., the image match term, depends on the neighboring intensity pattern around each vertex $p_i \in \mathbf{V}$. $E_{intensity}$ determines whether $\mathbf{V}$ is attracted

to a region of low or high image intensities; $E_{gradient}$ attracts $\mathbf{V}$ to edges in an image. The key is that the direction of the image gradient at an object boundary point should be similar to the unit normal direction of the active contour at that same point. In practice, for noisy images, *geometry-limited diffusion* [Whitaker (1993)] helps $\mathbf{V}$ to converge to the correct boundary at appropriate image scales.

In a continuous form of the energy function for active contours, the energy terms E_{int} and E_{ext} are both integrals of local energy measures on the contour parameterized by q:

$$E(\mathbf{V}) = \alpha \int_0^1 (e_{int}(\mathbf{V}(q)))dq + \int_0^1 (e_{ext}(\mathbf{V}(q)))dq \qquad (2.5)$$

where the integrand in the first term e_{int} corresponds to the local internal energy at the contour point $\mathbf{V}(q)$, i.e., the geometric typicality term, and the integrand in the second term e_{ext} corresponds to the local external energy there, i.e., the image match term.

Geometric models of active contours using the continuous form of the energy functions were proposed in [Malladi et al. (1995)]. Their method is based on the theory of curve evolution and geometric flow. $E(\mathbf{V})$ sums two terms: one related to the regularity of the contour curve and the other related to the shrinking or expanding of the contour towards object boundaries. Their method is driven by mean curvature flow, implemented by solving partial differential equations (PDEs). Different from the energy minimization process in the original snake method, the geometric active contours allow automatic changes in the shape topology if implemented by level-sets [Kimmel et al. (1995)]. As a result, multiple objects can be segmented simultaneously without any pre-knowledge of a given image. This is arguably both the strength and weakness of a geometric active contour method: it can segment multiple objects without guidance but can also undesirably break a single object into multiple parts in the presence of image artifacts.

Based on the relation between active contours and the computation of geodesics,

16

i.e., minimal distance curves, the geodesic active contour method proposed in [Caselles et al. (1997)] further pushed forward the methods in the active contour category. This geodesic approach allows the connection of classical snake methods based on energy minimization to geometric active contour methods based on curve evolution theory. The geodesic contour method improves the geometric active contour methods by the use of a new velocity term based on image appearance, allowing stable boundary detection when image gradients suffer from large variations, e.g., gaps in image edges.

The deformable model methods in the active contour category share the same vertex point based shape representation. These methods solely depend on local geometry and image features, making them both efficient for image segmentations and sensitive to local image defects, which might lead to undesired leaking of the contours or being stuck at local sub-optimal solutions.

Attempts have been made to integrate global shape priors into these active contour based methods [Yang et al. (2003)]. The proposed method applies principal component analysis (PCA), a linear statistical analysis method, on distance maps implied by contours. Problems can arise from the fact that a linear combination of two distance maps does not guarantee to be a valid distance map. Methods, such as those in [Pohl et al. (2007)], that are more theoretically sound have therefore been developed to solve the problem of [Yang et al. (2003)].

2.1.2 ASM/AAM

Active shape model (ASM) [Cootes et al. (1995)] and active appearance model (AAM) [Cootes et al. (2001)] methods use the object boundary point distribution model (PDM) as their shape representation. As probabilistic type of deformable model methods, the ASM and AAM both use learned shape prior statistics in image segmentation tasks. PCA is applied to the acquired PDM models to form a Gaussian probability distribution as the shape prior $p(\mathbf{M})$. In the statistical training stage to learn such shape prior

statistics, significant landmarks from an object shape model are carefully picked by humans, and intermediate landmarks are then uniformly placed between the hand-picked landmarks. Human intervention was a built-in step in both the ASM and AAM methods until automatic correspondence building methods [Davies et al. (2002); Styner et al. (2003)] were proposed and used as a refinement step before the statistics are calculated. These methods update where to place the intermediate landmarks for optimal landmark correspondences across a training set of objects from a population.

Both ASM and AAM methods incorporate image intensities in image segmentation. In the ASM method image intensity profile vectors as appearance features are centered at the landmarks and sampled along the surface normals at those landmarks. A segmentation starts with a coarse alignment of a template PDM model into a target image. The derivative of the image profile vector at each landmark determines the movement of that landmark. After all landmarks are moved once as in one iteration, the entire PDM as the set of landmarks is projected into the pre-learned shape prior space. By this means the segmented shape is always constrained in the learned shape space. The above process of moving and projecting the landmarks are repeated until an equilibrium is reached.

The AAM method improves the ASM framework by learning image appearance statistics as well as shape prior statistics. The shape prior statistics are built from applying the PCA to a set of training shape models of landmarks. A thin-plate spline deformation is calculated to warp landmarks in each training shape model to match the landmarks in the mean shape model of the shape statistics. The calculated deformation is then applied to each training image, and each warped training image is rasterized and normalized into an appearance vector. A second PCA is applied to the pooled and normalized appearance vectors to form an appearance probability distribution. In order to reflect the correlation between the shape and appearance probability distributions, a third PCA is applied to the combined feature of the shape and appearance PCA param-

eters. The final combined PCA model is used in image segmentations. A multi-scale search has also been proposed in the AAM framework [Mitchell et al. (2002)] to improve the robustness and efficiency of the method.

Both ASM and AAM have proven to be effective in image segmentation. They are also widely applied to computer vision tasks such as object tracking, detection, and recognition.

2.1.3   Methods Based on the Projection onto Orthogonal Functions

Deformable model methods in this category represent shapes by projecting an original shape model to a space spanned by a set of orthogonal basis functions. The coefficients of the basis functions are used to represent the original shape model.

The choices of the basis functions include sinusoidal basis functions [Duncan (1992)] and spherical harmonic (SPHARM) functions [Brechbühler et al. (1995); Kelemen et al. (1999)]. Compared with the methods using the PDM, these methods efficiently represent a shape model by a small number of parameters. However, they tend to have trouble in capturing localized shape variations.

The method proposed in [Staib and Duncan (1996)] projects a surface in the form of contour meshes to a space spanned by sinusoidal basis functions and uses the coefficients of the sinusoidal basis functions as the shape parameters. Four types of shapes can be represented: tori, open surfaces, closed surfaces, and tubes. In the segmentation framework of [Staib and Duncan (1996)], learned shape statistics of the basis function coefficients are used as the shape prior. The image match term is measured via either Gaussian derivatives applied to image profiles along model surface normals or a Gaussian model for the image noise after pooling training image intensities from a population. The segmentation is done by finding the optimal shape parameters for a target image via optimizing an objective function summing geometric typicality and image match terms.

Being a global shape representation, these shapes in [Staib and Duncan (1996)] have fewer parameters to optimize for image segmentation, which makes this method run faster but yield less accurate results because of the lack of local shape control. However, better segmentation results might be achieved by increasing the dimension of their shape space at higher computational cost.

Another method uses spherical harmonic basis functions (SPHARM) to form a projection space, which guarantees shape legality at expensive computational cost but does not handle surface locality [Csernansky et al. (1998); Styner et al. (2006)], a common problem of this type of deformable model methods.

2.1.4 Medially Based Methods

Some probabilistic deformable model methods are based on medial shape representations. Among these methods are the continuous m-rep method [Yushkevich et al. (2003, 2005)], the parameterized medial representation using Catmull-Clark subdivision scheme [Terriberry et al. (2005)], and the discrete m-rep method [Pizer et al. (2003)]. The first two medially based methods use (x, y, z, r) as their shape primitive, where (x, y, z) forms a surface called the *medial sheet* and r is a radius field on both sides of the medial sheet surface. The discrete m-rep method uses samples from the medial sheet surface and the radius field on the medial sheet, plus two or three spokes connecting both sides of the medial sheet to their corresponding object surface. All these three medially based methods follow the Bayesian framework to segment images. In section 2.3 I will review medial shape representations and these three medially based deformable model methods in detail.

Next, I will show the training and segmentation framework shared by the probabilistic type of deformable model methods.

2.1.5 Summary of Probabilistic Deformable Model Methods

The shape and appearance statistics used by probabilistic deformable model methods are learned from a set of training shape models and their corresponding images. In image segmentation by these deformable model methods, fitting a model $\mathbf{m}$ into a target image $\mathbf{I}_t$ is driven by maximizing a posterior probability in the learned shape prior space.

The two major components in probabilistic deformable model methods are thus the training and the segmentation. The general procedures of these two components are shown as follows.

1. Learning shape and appearance probability distributions

   (a) Acquire a set of segmented training images, often by experts' manual contouring (the accuracy and reliability of human manual segmentation, though an important issue, is not the focus of this dissertation);

   (b) Initialize a template deformable model into each training image by estimating a global transformation for the model;

   (c) Deform the respective initialized deformable model into each training image by optimizing an objective function summing the geometric typicality and image match terms;

   (d) Calculate the shape probability distribution from the fitted deformable models from the training images, yielding a mean shape model and modes of shape variations;

   (e) Use the extracted shape models and their corresponding original images to collect the image intensity information relative to the models, and calculate the appearance probability distribution.

2. Segmentation of target image(s) using the learned shape and appearance probability distributions

(a) Initialize a template deformable model into the target image, often by selecting the position and orientation of the mean shape model in the learned shape prior statistics;

(b) Repeat the following steps until convergence:

   i. Calculate a sample image $\mathbf{I}_s$ in the relative coordinates to the current model;

   ii. Compare the sample image $\mathbf{I}_s$ against the target image $\mathbf{I}_t$, and the difference between the two, together with the shape prior, generates deformation forces to deform the template mean model.

   In other words, the template model is deformed into a target image by optimizing the posterior probability as a combination of the shape prior and image likelihood probabilities;

The probabilistic deformable model methods must handle the following two issues in both the training and segmentation:

1. Initialization: the initial placement of a deformable model into a target image. Often this is carried out by applying either a rigid or similarity transformation to the initial deformable model. Various attempts have been made, including manual or semi-automatic initializations based on landmarks or contours picked or drawn by an human expert. Fully automatic initialization seems more attractive to release the burden on humans. Methods based on evolutionary algorithms have been proposed in [Heimann et al. (2007)] to provide such luxury at the cost of longer computational time.

2. Locally optimal solutions: the objective function or the posterior probability function optimized for fitting training models or segmenting images often have multiple local extreme points. The optimization of these functions will frequently trapped

at such local extreme points to yield sub-optimal results. Given the high dimension of the shape parameter space, the objective function for fitting training models or the posterior probability function for image segmentation is almost inevitably "bumpy", which makes the optimization quite likely to yield suboptimal results. Methods have been proposed to approach this issue, including brute force global search, simulated annealing, and evolutionary algorithms. A common shortcoming of these proposed methods is their speed.

In contrast to earlier segmentation methods using local image features, more recent deformable methods make efficient use of the Bayes's statistical framework shown in equation 2.1. These deformable model methods appear to be less sensitive to image noise and defects and more likely to find the global optimal object boundary without being trapped by local image features. Some of these deformable model methods are based on medial shape representations. Next section 2.2 starts with an introduction to medially based shape representations and then reviews three deformable model methods based on the medial shape representations.

2.2 Medially Based Deformable Models

In contrast to a shape representation using boundary landmarks or the point distribution model (PDM), a medially based shape representation represents not only the boundary but also the interior volume of the shape. It generates a local object-relative coordinate system. This coordinate system improves the robustness of volumetric image access in the object interior and in the adjacent object exterior regions. These medially based shape representations arise from the Blum medial axis [Blum and Nagel (1978)].

2.2.1 The Blum Medial Axis and Blum Condition

A Blum medial axis describes a 3D object by a medial locus $\mathbf{p} = x, y, z$ as the center of a maximal sphere with a radius r, which is entirely contained in the interior of the object and tangential to two or more points on the object boundary. The radius r of the maximal sphere is a part of the Blum medial axis representation. The set of points $\mathbf{p}$ forms a smooth surface called a *medial sheet*, which is a smooth surface for a simple object but a complicated branching structure for a complex object. Given an object, its medial sheet surface plus a smooth radius field r attached to the medial sheet form a *Blum medial axis*. Samples from a Blum medial axis are used in medially based shape representations.

A sample $(\mathbf{p}, r)$ from a medial axis is a primary building block of a medially based shape representation. A medial primitive $(\mathbf{p}, r)$ is called a 0th-order medial atom. A set of such atoms can be used as the control points to describe the medial axis of an object via an appropriate interpolation or approximation scheme. The object boundary can be reconstructed as the envelope of the overlapping tangential spheres implied by the interpolated 0th-order medial atoms $(\mathbf{p}(u_1, u_2), r(u_1, u_2))$.

Beside $(\mathbf{p}, r)$, more information can be included in a medial primitive, such as a pair of vectors $\mathbf{S}^i$ from an internal medial sheet point $\mathbf{p}$ to the two object boundary points at which the corresponding maximal sphere $\mathbf{Sphere}(\mathbf{p}, r)$ is tangential to the object boundary. These two vectors $\mathbf{S}^i$ $(i = -1, +1)$ are called *medial spokes*, which have the same length r and the directions $\mathbf{U}^i$, where $|\mathbf{U}^i| = 1$. Each medial primitive $(\mathbf{p}, r, \{\mathbf{U}^i | i = -1, +1\})$ is called a 1st-order medial atom.

The medial spokes $\mathbf{S}^i$ can be derived from the first-order derivatives of $\mathbf{p}$ and r, that is how a 1st-order medial atom is named. For simplicity, I will call each 0th-order medial atom $(\mathbf{p}, r)$ a *medial control point* and call each 1st-order medial atom a *medial atom* from now on. So a medial atom is a hub position $\mathbf{p}$ plus two spokes $\mathbf{S}^i = r\mathbf{U}^i$ for $i = -1, +1$, as shown in figure 2.1.

24
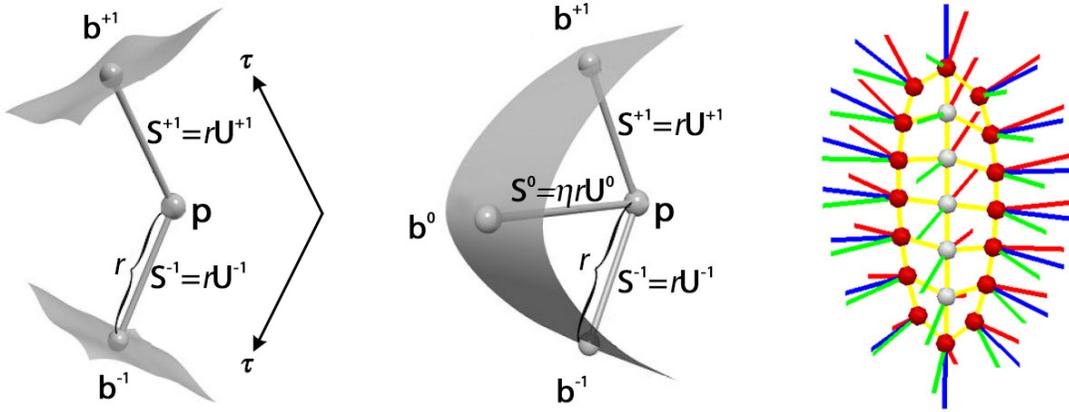
Figure 2.1: From left to right: an internal atom of two spokes $\mathbf{S}^{+1}$ and $\mathbf{S}^{-1}$, with τ to parameterize the object interior along the spokes; an end atom with an extra bisector spoke $\mathbf{S}^0$ and parameter η that controls the shape of the boundary crest region; a discrete m-rep as a mesh of internal atoms (with white hubs) and end atoms (with red hubs).

For an object in 2D, its medial sheet degenerates to a 2D curve, but the medial control points or medial atoms to represent the 2D object have the same configuration as that of a 3D object: $(\mathbf{p}_{2D}, r)$ or $(\mathbf{p}, r, \{\mathbf{U}_{2D}^i\})$ for a medial control point or medial atom in 2D, respectively. This dissertation focuses on representing objects in 3D.

There is a mapping between a medial sheet and its object boundary in each direction, and the mappings in both directions have been studied:

1. From an object boundary to its medial axis - calculating the medial axis given an object boundary

2. From a medial axis to its object boundary - generating an object boundary given the medial axis of the object

The first mapping from an object boundary to its medial axis is sensitive to boundary noise, and various methods have been proposed. Reviews of these methods can be found in chapters 4-8 of [Siddiqi and Pizer (2007)]. The second mapping from a medial axis to its implied boundary is stable and is a focus of this dissertation. Next, I will review the mathematics behind a medial axis, leading to a set of medial constraints called the Blum condition.

25

Assume that a medial axis is represented by a pair of functions ($\mathbf{p}(u_1, u_2)$ and $r(u_1, u_2)$), which corresponds to the medial sheet and the radius field on the medial sheet, respectively. The medial spokes can be calculated at each point on the medial sheet by the first-order partial derivatives $\mathbf{p}_{u_1}, \mathbf{p}_{u_2}, r_{u_1}, r_{u_2}$ of the functions $\mathbf{p}(u_1, u_2)$ and $r(u_1, u_2)$).

The gradient of the radius is given by $\nabla r = \begin{pmatrix} \mathbf{p}_u & \mathbf{p}_v \end{pmatrix} \mathbf{I_p}^{-1} \begin{pmatrix} r_u \\ r_v \end{pmatrix}$, where $\mathbf{I_p}$ is the metric tensor of the medial sheet surface $\mathbf{p}$:

$$\mathbf{I_p} = \begin{pmatrix} < \mathbf{p}_u, \mathbf{p}_u > & < \mathbf{p}_u, \mathbf{p}_v > \\ < \mathbf{p}_v, \mathbf{p}_u > & < \mathbf{p}_v, \mathbf{p}_v > \end{pmatrix}$$

In order to have a proper (legal), i.e., non-folding, boundary and spoke field, a set of constraints called the Blum condition must hold as follows.

1. $\|\nabla r\| < 1$ for interior points on the medial sheet

2. $\|\nabla r\| = 1$ for edge medial points on the medial sheet

3. The Jacobian determinant of the medial-to-boundary mapping remains positive

4. $\|\nabla r\| = 1$ for all branching medial points shared by a set of one-piece partial medial axes in a branching medial axis of a complex object

The second constraint is called the **Blum boundary condition** and is crucial for a simple object with a non-branching medial axis, and the fourth constraint is important for a complex object with a branching medial axis.

2.2.2 Three Medially Based Shape Representations

This subsection reviews three series of medially based shape representations in sections 2.2.3-2.2.5:

1. The cm-rep based on B-splines in different forms, which uses B-spline patches of medial control points [Yushkevich et al. (2003, 2005)];

2. The cm-rep based on the Catmull-Clark subdivision [Catmull and Clark (1978)], which uses subdivision patches of medial control points and uses spline control curves on the edges of the patches [Terriberry et al. (2007)];

3. The discrete m-rep, which uses 1st order medial atoms as shape representation primitives [Pizer et al. (2003)].

The first two medial representations both use medial control points, i.e., 0th-order medial atoms, as their shape primitives but apply different approximation schemes to the control points. The derivatives of the approximated medial controls points are used to calculate medial spokes, and the end points of the calculated medial spokes are used to generate an object boundary. Both the medial representations handle types of branching medial axes.

The third medial representation, the m-rep, uses a set of 1st-order medial atoms as its shape primitives, with the medial spokes explicitly included in the shape representation. Approximation is applied to the control boundary points [Thall (2004)] implied by medial atoms and to medial atoms directly [Han et al. (2006)]. An m-rep can represent certain types of branching structure in 3D.

In order for a medially based shape representation to imply a proper (legal) boundary surface without local shape defects, such as creasing or folding, the first-order derivatives of all the medial points must meet the *Blum condition*, described in section 2.2.1. The Blum condition directly affects approximation schemes used in medial shape representations, including the three medial shape representations to be reviewed in the following subsection.

2.2.3 Cm-reps by B-splines

Yushkevich et al. (2003) proposed a medial representation implemented by continuous B-spline patches of medial control points. This medial representation generates a continuous boundary and volumetric interior of an object. The pair of continuous functions $(\mathbf{p}(u_1, u_2), r(u_1, u_2))$ are in the the form of cubic B-spline patches, controlled by medial control points $(\mathbf{p}_{ij}, r_{ij})$, $i \in [1, 4], j \in [1, 4]$.

As stated in subsection 2.2.1, the two functions $\mathbf{p}$ and r representing a medial axis must satisfy the Blum condition to ensure a non-folding surface and spoke field. Yushkevich's cm-rep achieved first three constraints in the Blum condition by setting large negative r to the edge medial control points and positive r to the interior medial control points. This guarantees the level curve of $\|\nabla r\| = 1$ lie within the b-spline patches. The b-spline patches are then trimmed along this level curve and the edge resulting from the trimming forms the real edge of the medial sheet surface.

A side effect of this treatment to the real medial sheet edge is that the domain of the trimmed level curve varies across different objects, so different objects do not share a common parameter space. This varying parameter space for different objects makes it difficult to apply statistical analysis on shapes represented by Yushkevich's cm-rep. Branching medial axes in 3D also remain a challenge to the proposed cm-rep in [Yushkevich et al. (2003)] because of the lack of enough free parameters to meet the fourth constraint in the Blum condition.

Yushkevich et al. (2005) proposed a second form of the cm-rep based on solutions to Poisson equations that fulfill the Blum condition. This representation enforces a fixed parameter space on different objects, which allows statistical analysis on different objects. However, this form of cm-rep is limited to represent 2D structures because it does not provide enough degrees of freedom to handle branching medial structures in 3D.

Sun et al. (2008) proposed a new form of the cm-rep by applying the Loop subdivision

scheme Loop and DeRose (1990)] to B-spline patches of medial control points plus soft constraints to maintain the Blum condition for edge and branch medial points. This form of the cm-rep has been shown to work with 3D branching medial axes. The soft constraints used in this method were inspired by the control "curve" idea proposed in the next cm-rep method based on Catmull-Clark subdivision patches.

2.2.4   Cm-reps by Subdivision Patches

A Catmull-Clark subdivision based medial shape representation was proposed by Terriberry et al. in [Terriberry et al. (2007)] to form the functions ($\mathbf{p}(u_1, u_2)$ and $r(u_1, u_2)$) of a medial axis. Medial control points ($\mathbf{p}, r$) are the primitives in this representation, and the Catmull-Clark subdivision is used to approximate the medial control points. The Blum condition is met by modifying the edge subdivision patches with the help of interpolating spline curves. This representation applies a degree reduction to the polynomial equations implied by the Catmull-Clark subdivision scheme and shows a direct solution to the Blum boundary condition, i.e., $\|\nabla r\| = 1$ at the edge of a medial sheet. The solutions to the Blum boundary condition are then used to construct a b-spline "control curve" [Terriberry et al. (2007)] to replace the outer layer of medial control points for the Catmull-Clark subdivision patches. Detailed descriptions on this control curve approach to the Blum boundary condition are in [Terriberry et al. (2007)]. Furthermore, the same framework using control curves can be extended to handle the fourth constraint in the Blum condition for branching medial axes. Terriberry's cm-rep can represent 2D and 3D objects of both non-branching and branching medial axes.

2.2.5   Discrete M-reps

The cm-reps of Yushkevich and Terriberry both use 0th-order medial control points ($\mathbf{p}, r$) in their shape representations. Medial spokes have to be implied by the derivatives of smooth functions ($\mathbf{p}(u_1, u_2), r(u_1, u_2)$). Pizer et al. developed the discrete m-rep, using

1st-order medial atoms $(\mathbf{p}, r, \{\mathbf{U}^i\})$. In a discrete m-rep, medial spokes are explicitly included in its shape representation. By doing so, a discrete m-rep can capture localized shape variations, including bending, twisting, and tapering, and provide local shape control via the discrete medial atoms.

The explicit inclusion of medial spokes in a discrete m-rep atom and the allowed atom spokes that can be non-Blum by relaxing the Blum condition make an m-rep a variant of medial models that is approximately but not precisely Blum. Such relaxations are shown as follows.

- The included spokes in an m-rep can be non-perpendicular to the m-rep implied boundary, which will be elaborated in section 3.4;

- The m-rep spokes in the crest region, i.e., at the edge of the medial sheet are handled specially as a part of m-rep end atoms for reasons of computational stability, to be described in next paragraph;

- The m-rep spokes in a branching structure are handled differently from a Blum medial axis, to be described after the description of m-rep end atoms.

The m-rep handles spokes at the edge of the medial sheet differently from a Blum medial axis. In a Blum medial axis, an internal medial point on the medial sheet has a pair of spokes with equal length attached to the medial sheet at that point. When approaching the edge of the medial sheet from an internal point, the pair of spokes swing close to each other and finally collapse infinitely fast into each other to form one single spoke at the very end, i.e., the edge, of the medial sheet. Such a collapsing process is generic but extremely sensitive to small boundary noise, and it is reflected in the second Blum condition, i.e., the Blum boundary condition, that $\|\nabla r\| = 1$. In an m-rep, the basic shape representation primitive is a medial atom. A medial atom containing the spoke at the edge of the m-rep medial sheet is called an end atom, in which there is a pair of regular spokes $\mathbf{S}^{+1,-1}$ and a third spoke $\mathbf{S}^0 = \eta r \mathbf{U}^0$ that bisects the pair of regular

spokes. This third spoke represents the spoke at the edge of the m-rep medial sheet, and the extra parameter $\eta > 1$ controls the shape of the object crest. The hub position $\mathbf{p}$ of each end atom is not exactly an edge point on the m-rep medial sheet. Instead, the edge point of an m-rep medial sheet is implied by an end atom to be $\mathbf{p} + r(\eta - 1)\mathbf{U}^0$, which is on the bisector spoke $\mathbf{S}^0 = \eta r \mathbf{U}^0$. As a result, each end atom in an m-rep is a combination of an internal m-rep atom $\{\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1}\}$ with the pair of regular spokes and a bisector spoke $\mathbf{S}^0$ corresponding to the crest of the object boundary. By adding $\eta > 1$ as a free parameter for each end atom to control the shape of the crest, the burden to fulfill the Blum boundary condition at the edge of a medial sheet is relieved by allowing the edge of the medial sheet to slide on the bisector spoke of an end atom.

An m-rep handles spokes at a branching structure differently from a Blum medial axis as well. In a Blum medial axis with a typical 3D "fin" branching, the main medial sheet, to which the fin part is attached, contains a curve shared by the main medial sheet and the fin part. In order for the branching medial axis to meet the Blum condition, all the medial points on this curve must meet the fourth constraint in the Blum condition such that $\|\nabla r\| = 1$. This requirement greatly complicates any approximation or interpolation method on a branching medial axis [Yushkevich et al. (2005); Terriberry et al. (2007)]. Furthermore, the branching part of a 3D Blum medial axis represents little object volume, which reduces the stability in a shape representation following the exact Blum branching constraint. In an m-rep, a set of single-piece medial axes is used to represent a complex object. The set of medial axes are treated hierarchically and connected geometrically. Such an m-rep with connected medial figures is called a multi-figure m-rep. The hierarchy in a multi-figure m-rep is organized by connecting each pair of adjacent figures via *hinge geometry*, to be detailed in section 4.2.1. Within each pair of connected medial figures there are a host figure and a subfigure, as shown in figure 2.2. This multi-figure m-rep framework follows in the footsteps of how the natural anatomy of a complex object is explained in medicine, by hierarchy and connectivity.
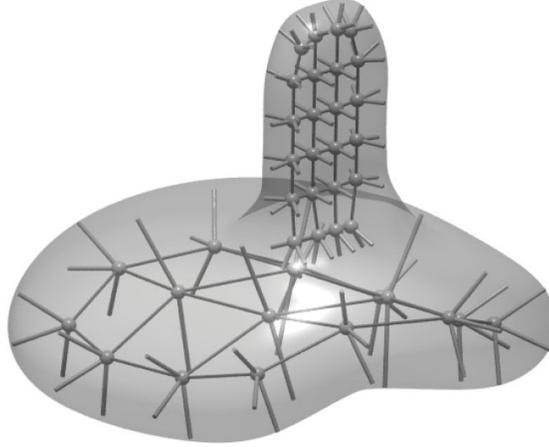
Figure 2.2: An object with two parts represented by a two-figure m-rep.

A multi-figure m-rep relieves the burden of the exact Blum condition for a branching axis by its geometric definition of connections between adjacent object parts and its replacement of the branching part of a Blum medial axis representing little volume by a *skeletal structure* [Damon (2005)], which is the underlying medial structure for a smooth transition between each pair of connected parts, to be detailed in section 4.2.3.

Overall, the relaxed branching structure in a multi-figure m-rep has the following advantages:

1. The interrelations and transformations between a pair of adjacent figures are explicitly defined by the hinge geometry, which allows each one of the two medial figures to deform with or without affecting the other;

2. A subtractive part as an indentation subfigure connects in a completely equivalent way as an additive part as a protrusion subfigure to its host figure, shown in figure 2.2. This significantly simplifies the statistical analysis on objects with indentations or protrusions;

3. The representation of an object as a set of m-rep figures with each figure consisting of a sheet of medial atoms yields a multiscale shape representation that enables a multiscale statistical analysis on such complex objects, to be detailed in chapter

4. The fact that a host figure stands for the majority of an object volume and is visited before its subfigure in a multi-scale fitting or segmentation process provides more stability.

Given a medial shape representation, we often need to calculate its implied boundary. In the cm-reps of Yushkevich and Terriberry, different approximation schemes are applied directly to their shape primitives, i.e., medial control points, and the approximation schemes are strictly constrained by the Blum condition. In the cm-reps, no implied boundary points can be calculated, not even for the medial control points, until smooth approximations are derived from the control points for both the medial sheet and the radial field because a cm-rep implied boundary solely depends on the underlying spokes, and the underlying spokes can only be calculated from the derivatives of the medial sheet and radius field. It is different for a discrete m-rep because the included spokes in an m-rep imply a coarse boundary mesh connecting all the spoke ends. As a result, there are two types of methods to calculate an m-rep implied boundary by an interpolation either on the mesh of medial atom spoke ends or directly on medial atoms to form an interpolated medial sheet and spoke field.

The method based on the Catmull-Clark subdivision [Thall (2004)] was the first one proposed to generate the implied boundary of an m-rep. This method uses an interpolating variation of the Catmull-Clark subdivision scheme and belongs to the first type of interpolation methods on m-reps. This interpolation method is applied to the surface points and approximate normals implied by medial atoms in an m-rep figure: medial spoke end points as the surface points and spoke directions as the surface normals at those surface points.

A new method will be proposed in this dissertation to construct the boundary from an m-rep, by a constrained medial *atom interpolation*, which directly interpolates m-rep atoms and belongs to the second type of interpolation methods on m-reps.

At this moment, the Catmull-Clark subdivision based method is more efficient in

computation. However, the new method based on atom interpolation has the advantage of generating a spoke field on a medial sheet, which allows us to apply constraints to enforce a shape model to stay proper (legal) or to implement medially based integrations on the interpolated medial spoke field to assist deformable model fitting or image segmentation. Furthermore, the new method provides a full parameterization of the entire object volume by the interpolated spokes. The atom interpolation will be described in detail by chapter 3.

I have reviewed three series of medial shape representations. The main advantages of them are as follows.

1. They are volumetric representations: they model not only the boundary but also the interior of an object and the immediately adjacent exterior of the object. This provides robust accesses to volumetric image intensities.

2. They are multiscale representations: they can represent the geometry of objects or object complexes in a coarse-to-fine fashion. Multiscale approaches for deformable models have proven to be more robust to image defects such as noise, aliasing, and missing data. Also, multiscale methods are more robust to the local optimum problem during the maximum a posterior (MAP) optimization for image segmentation;

3. They provide a figural coordinate system: such a coordinate system can be used to describe the local interrelationships among adjacent object parts or adjacent objects in an object complex.

The discrete m-rep framework has proven to yield one of the best segmentation results on bladder and prostate CT images in the literature. In the rest of this chapter, section 2.3 reviews the mathematical background of a medial axis that enables us to design explicit legality constraints on m-reps, and section 2.4 reviews how to conduct non-linear statistical analysis on m-reps.

2.3 Mathematics of a Medial Structure

This section reviews the radial shape operator and its $r\mathbf{S}_{rad}$ matrix, and the edge shape operator and its $r\mathbf{S}_E$ matrix. These algebraic operators characterize the geometry behind a medial structure. The rest of this section lays out the definition and calculation of both the shape operators and their extensions on a Blum medial structure, the application of these shape operators to any smooth spoke field attached to a smooth medial sheet, called a skeletal structure [Damon (2005)], and two legality conditions based on the geometric property of these shape operators.

Recall that a shape operator defined on a 3D surface describes how the surface normal swings along the surface, i.e., how a surface curves locally. Analogously, a *radial shape operator* $\mathbf{S}_{rad}$ tells how a unit length medial spoke $\mathbf{U}$ changes while walking on the interior of a medial sheet, and an *edge shape operator* $\mathbf{S}_E$ tells how $\mathbf{U}$ changes while walking along the edge curve of a medial sheet. Both the radial shape operator and the edge shape operator are defined by Damon in [Damon (2003, 2004, 2005)]. These medial shape operators record derivative information of a medial sheet surface and medial spokes attached to it. A radial shape operator and an edge shape operator can also be extended to an $r\mathbf{S}_{rad}$ matrix and an $r\mathbf{S}_E$ matrix, respectively, to represent the rate of change of full spokes $\mathbf{S}$ instead of unit length spokes. Next I will show how to calculate the shape operators and their extended matrices, starting with the radial shape operator and the $r\mathbf{S}_{rad}$ matrix followed by the edge shape operator and the $r\mathbf{S}_E$ matrix.

Each internal medial point on the medial sheet has a spoke at each side of a medial sheet; thus the spoke fields on a medial sheet are double-valued. So there are two radial shape operators defined for each medial atom: one for each spoke at one side of the medial sheet. The following description focuses on one radial shape operator because the same description applies to the other radial shape operator of an m-rep atom.

Consider one side of a medial sheet, and assume that there is a continuous spoke

field $\mathbf{S}(\mathbf{u})$ with unit length spoke direction $\mathbf{U}(\mathbf{u})$ and spoke length $r(\mathbf{u})$ on the side of the continuous medial sheet $\mathbf{p}(\mathbf{u})$. $\mathbf{S}(\mathbf{u}) = r(\mathbf{u}) \cdot \mathbf{U}(\mathbf{u})$, and $\mathbf{u} = (u_1, u_2)$ parameterizes the two dimensional medial sheet and spoke field. The derivatives of the unit length spoke direction $\mathbf{U}(\mathbf{u})$ by (u_1, u_2) are calculated as follows, with $\mathbf{U}$ and $\mathbf{p}_{u_1,u_2}$ being 1×3 row vectors.

$$\frac{\partial \mathbf{U}}{\partial v_i} = a_{0,i}\mathbf{U} - a_{i,1}\mathbf{p}_{u_1} - a_{i,2}\mathbf{p}_{u_2}, \text{ where } i = 1, 2, \tag{2.6}$$

or rewritten in matrix form:

$$\frac{\partial \mathbf{U}}{\partial \mathbf{u}} = \begin{pmatrix} a_{0,1} \\ a_{0,2} \end{pmatrix} \mathbf{U} - \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{p}_{u_2} \end{pmatrix} \tag{2.7}$$

where $\frac{\partial \mathbf{U}}{\partial \mathbf{u}}$ is a 2×3 matrix with row i as the vector $\frac{\partial \mathbf{U}}{\partial v_i}$, and where $\mathbf{p}_{u_1}$ and $\mathbf{p}_{u_2}$ are the derivatives of the medial sheet $\mathbf{p}$ by parameters u_1 and u_2. In these equations, the derivative of $\mathbf{U}$ is decomposed by a generally non-orthogonal projection along the spoke direction $\mathbf{U}$ to the tangent plane of the medial sheet spanned by $\mathbf{p}_{u_1}$ and $\mathbf{p}_{u_2}$.

Let $\mathbf{A_u} = \begin{pmatrix} a_{0,1} \\ a_{0,2} \end{pmatrix}$, and let $\mathbf{S}_{rad} = \begin{pmatrix} a_{1,1} & a_{2,1} \\ a_{1,2} & a_{2,2} \end{pmatrix}$. $\mathbf{S}_{rad}$ is called the radial shape operator. In general, the radial shape operator has the following properties:

- 2×2 matrix

- not self-adjoint

Then (2.17) implies

$$\frac{\partial \mathbf{U}}{\partial \mathbf{u}} = \mathbf{A_u}\mathbf{U} - \mathbf{S}_{rad}^T \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{p}_{u_2} \end{pmatrix}, \tag{2.8}$$

$\mathbf{U}(\mathbf{u})$ is of unit length, implying $\mathbf{U} \cdot \mathbf{U}^T = 1$ and $\frac{\partial \mathbf{U}}{\partial \mathbf{u}} \cdot \mathbf{U}^T = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$. So

$$\mathbf{A_u} = \mathbf{S}_{rad}^T \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{p}_{u_2} \end{pmatrix} \mathbf{U}^T \tag{2.9}$$

Substituting (2.9) into (2.17) yields the means of computing $\mathbf{S}_{rad}$ given $\frac{\partial \mathbf{U}}{\partial \mathbf{u}}$, $\mathbf{U}$ and $\begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{p}_{u_2} \end{pmatrix}$:

$$\mathbf{S}_{rad} = \left(\frac{\partial \mathbf{U}}{\partial \mathbf{u}} \cdot \mathbf{Q}^{-1} \right)^T \tag{2.10}$$

where $Q = \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{p}_{u_2} \end{pmatrix} (\mathbf{U}^T\mathbf{U} - \mathbf{I})$ is a 2×3 matrix.

That is to say, $\mathbf{S}_{rad}$ depends on the spoke direction $\mathbf{U}$, and the derivatives of $\mathbf{U}$ and $\mathbf{p}$.

Besides the radial shape operator $\mathbf{S}_{rad}$ describing the rate of change of unit length spokes, we also need to describe the rate of change of full spokes $\mathbf{S} = r\mathbf{U}$. The derivative of a full spoke $\mathbf{S} = r\mathbf{U}$ of length r is given as follows, based on the derivative of a unit length spoke $\mathbf{U}$.

$$\frac{\partial \mathbf{S}}{\partial \mathbf{u}} = \frac{\partial(r\mathbf{U})}{\partial \mathbf{u}} = r\frac{\partial \mathbf{U}}{\partial \mathbf{u}} + \begin{pmatrix} r_{u_1} \\ r_{u_2} \end{pmatrix} \mathbf{U} \tag{2.11}$$

Substituting (2.17) and (2.9) into (2.11) yields

$$\frac{\partial \mathbf{S}}{\partial \mathbf{u}} = r\mathbf{S}_{rad}^T \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{p}_{u_2} \end{pmatrix} (\mathbf{U}^T\mathbf{U} - \mathbf{I}) + \begin{pmatrix} r_{u_1} \\ r_{u_2} \end{pmatrix} \mathbf{U} \tag{2.12}$$

According to Damon, a *compatibility condition* [Damon (2004)] requires that $r_{u_1 or u_2} = -\mathbf{p}_{u_1 or u_2}\mathbf{U}^T$, which implies that a spoke $\mathbf{S}$ is perpendicular to the implied boundary at its spoke end. Then

$$rS_{rad} = \left(\left(\frac{\partial S}{\partial u} + \begin{pmatrix} p_{u_1} U^T \\ p_{u_2} U^T \end{pmatrix} U \right) Q^T (QQ^T)^{-1} \right)^T \qquad (2.13)$$

Based on the compatibility condition, equation (2.13) shows how to compute an rS_{rad} matrix given the derivatives of S, p, and U in a Blum medial axis. However in a general skeletal structure [Damon (2004, 2005)], in which spokes are not necessarily perpendicular to the object boundary, the derivatives of r are not directly related to the derivatives of the medial sheet p and have to be calculated directly instead of being implied by the compatibility condition. In the case of a skeletal structure, the calculation of rS_{rad} is shown in equation (2.14):

$$rS_{rad} = \left(\left(\frac{\partial S}{\partial u} - \begin{pmatrix} r_{u_1} \\ r_{u_2} \end{pmatrix} U \right) Q^T (QQ^T)^{-1} \right)^T \qquad (2.14)$$

Besides the regular spokes on both sides of a medial sheet, there are also end spokes lying on the edge curve of the medial sheet. As a result, each end m-rep atom has an extra bisector spoke corresponding to the boundary crest. This bisector spoke swings along the edge curve of the medial sheet, so a different shape operator is needed to record its rate of change while walking along the edge curve. An edge shape operator S_E serves this purpose. The following describes the calculation of an edge shape operator and its extension rS_E matrix.

Let $\delta(t)$ be the edge curve of a medial sheet, parameterized by t, let p_0 be an edge point on δt as an edge medial point of interest, and let $p(u_1, u_2)$ be a *local edge parameterization* of an open set W with $p(0,0) = p_0 \in \delta(t)$ (the exact definition of this local edge parameterization is out of the scope of this dissertation, and detail of this special parameterization can be found in [Damon (2005)]). Assume that $p(u_1, u_2)$ is chosen such that the derivative of p along u_1 is aligned with the tangent direction of $\delta(t)$ at p_0, i.e., $p'_{u_1} \in T_{p_0}\delta(t)$, and such that the derivative of p along u_2 maps under the

edge parameterization to $c\mathbf{U}_{tan}$, for $c\mathbf{U}_{tan}$ the tangential component of $\mathbf{U}$ and $c \geq 0$. $\mathbf{S}_E$ is then calculated by projecting the derivative of $\mathbf{U}$ via a non-orthogonal projection along $\mathbf{U}$ to a plane, spanned by $\mathbf{n}$, the normal to the medial sheet at $\mathbf{p}_0$, and $\mathbf{p}_{u_1}$:

$$\frac{\partial \mathbf{U}}{\partial v_i} = a_{0,i}\mathbf{U} - a_{i,1}\mathbf{p}_{u_1} - a_{i,2}\mathbf{n}, \text{ where } i = 1, 2, \tag{2.15}$$

or rewritten in matrix form:

$$\frac{\partial \mathbf{U}}{\partial \mathbf{u}} = \begin{pmatrix} a_{0,1} \\ a_{0,2} \end{pmatrix} \mathbf{U} - \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{n} \end{pmatrix} \tag{2.16}$$

where $\frac{\partial \mathbf{U}}{\partial \mathbf{u}}$, $\mathbf{p}_{u_1}$, and $\mathbf{p}_{u_2}$ are defined as in .

$\mathbf{S}_E = \begin{pmatrix} a_{1,1} & a_{2,1} \\ a_{1,2} & a_{2,2} \end{pmatrix}$ is defined as the edge shape operator. In general, the edge shape operator also has the following properties:

- 2×2 matrix

- not self-adjoint

Let $\mathbf{A}_{\mathbf{u},E} = \begin{pmatrix} a_{0,1} \\ a_{0,2} \end{pmatrix}$. Then (2.16) implies

$$\frac{\partial \mathbf{U}}{\partial \mathbf{u}} = \mathbf{A}_{\mathbf{u},E}\mathbf{U} - \mathbf{S}_E^T \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{n} \end{pmatrix}, \tag{2.17}$$

Following the derivation of the $\mathbf{S}_{rad}$ and its extension $r\mathbf{S}_{rad}$, the means of computing $\mathbf{S}_E$ and its extension $r\mathbf{S}_E$ given $\frac{\partial \mathbf{U}}{\partial \mathbf{u}}$, $\mathbf{U}$ and $\begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{n} \end{pmatrix}$ are as follows:

$$\mathbf{S}_E = \left(\frac{\partial \mathbf{U}}{\partial \mathbf{u}} \cdot \mathbf{Q}_E^{-1} \right)^T \tag{2.18}$$

where $\mathbf{Q}_E = \begin{pmatrix} \mathbf{p}_{u_1} \\ \mathbf{n} \end{pmatrix} (\mathbf{U}^T\mathbf{U} - \mathbf{I})$ is a 2×3 matrix.

$$rS_E = \left(\left(\frac{\partial \mathbf{S}}{\partial \mathbf{u}} + \begin{pmatrix} \mathbf{p}_{u_1} \mathbf{U}^T \\ \mathbf{p}_{u_2} \mathbf{U}^T \end{pmatrix} \mathbf{U} \right) \mathbf{Q}_E^T (\mathbf{Q}_E \mathbf{Q}_E^T)^{-1} \right)^T \qquad (2.19)$$

which is based on the compatibility condition in a Blum medial axis.

$$rS_E = \left(\left(\frac{\partial \mathbf{S}}{\partial \mathbf{u}} - \begin{pmatrix} r_{u_1} \\ r_n \end{pmatrix} \mathbf{U} \right) \mathbf{Q}_E^T (\mathbf{Q}_E \mathbf{Q}_E^T)^{-1} \right)^T \qquad (2.20)$$

which is not enforcing the compatibility condition for a general skeletal structure.

Equations (2.14) and (2.20) show the full power of an rS_{rad} matrix and an rS_E matrix lying in the fact that they are applicable to a skeletal structure, i.e., any smooth spoke field on a smooth medial sheet surface, either one-sided or two-sided. A Blum medial axis is a special case of a two-sided skeletal structure such that each medial axis point implies a maximal sphere inside the object and tangential to the object boundary at two or more points, i.e., the medial spokes implied by a Blum medial axis are perpendicular to the implied boundary at the spoke end points. The rS_{rad} and rS_E matrices allow us to work with a skeletal structure, which can be non-Blum, such as the interpolated spoke field of a blend region in a multi-figure m-rep, to be described in section 4.2.3.

So far this section has focused on the algebraic property of the two medial shape operators and their extensions, recording first order derivatives of a medial sheet and medial spokes. Both the shape operators and their extensions also have a geometric property that is crucial to two medial legality conditions, coming next.

Geometrically, the property of a radial shape operator and an edge shape operator are also analogous to the geometric property of a 3D surface shape operator. The geometric property of a 3D surface shape operator is reflected by its eigenvalues, called the principal curvatures [O'Neill (1997)]. The local curviness of a 3D surface at a surface point is described by the principal curvatures at that point. Analogously, the eigenvalues of a radial shape operator $\mathbf{S}_{rad}$ reflect how the interior of a medial axis, including the

interior of the medial sheet $\mathbf{p}$ and the corresponding spoke directions $\mathbf{U}$, curves and swings locally, and the generalized eigenvalue of $(\mathbf{S}_E, \mathbf{I}_{1,1})$, computed as $a_{2,2}^{-1} \cdot det(\mathbf{S}_E)$, reflects how the edge curve and its attached end spokes of a medial axis change.

As an extension to a radial shape operator, an $r\mathbf{S}_{rad}$ matrix automatically incorporates the curving of the interior of the medial sheet $\mathbf{p}$ when simultaneously measuring the rate of change of the full spokes $\mathbf{S}$ on the interior of the medial sheet; as an extension to an edge shape operator, an $r\mathbf{S}_E$ matrix automatically incorporates the curving of the edge curve of the medial sheet $\mathbf{p}$ when simultaneously measuring the rate of change of the full spokes $\mathbf{S}$ along the medial sheet edge curve. The eigenvalues of an $r\mathbf{S}_{rad}$ matrix, i.e., the *radial principal curvatures*, and the generalized eigenvalue of an $r\mathbf{S}_E$ matrix, i.e., the *principal edge curvature* will be used in the two legality conditions to be described next [Damon (2005)].

Consider a local radial flow from the interior of a medial sheet $\mathbf{p}$ along one of the two spokes $\mathbf{S}^{+1,-1}$ to the implied boundary as $\varphi(\mathbf{p}, t) = \mathbf{p} + t\mathbf{S}, t \in [0, 1]$. φ can be generalized to a global radial flow from the medial sheet to the entire object boundary via the doubled-sided spoke field on the medial sheet. The spoke field is *legal* if and only if the Jacobian matrix of this global radial flow φ is never singular. This implies that for a legal spoke field in the interior of a medial sheet, i.e., one free of any intersections among the spokes, it has to fulfill a legality condition called the radial curvature condition [Damon (2005)] shown below. The radial curvature condition is mathematically equivalent to the constraints in the Blum condition (listed at the end of section 2.2.1).

$$\lambda_{ri} < 1, \text{ where } \lambda_{ri} = r\kappa_{ri}, \text{ for all positive real eigenvalues } \lambda_{ri}, i = 1 \text{ or } 2 \text{ of } r\mathbf{S}_{rad}.$$

(2.21)

where r is the interior spoke length.

Based on a similar derivation of the legality of a local edge radial flow, for the spoke

field on a medial sheet edge curve to be legal, it has to fulfill a legality condition called the edge condition [Damon (2005)] shown below.

$$\lambda_E < 1, \text{ where } \lambda_E = r\kappa_E, \text{ for the generalized eigenvalue } \kappa_E \text{ of } (\mathbf{S}_E, \mathbf{I}_{1,1}) \qquad (2.22)$$

where r is the end spoke length.

If the radial curvature condition and the edge condition are met at all medial points, the legality of the implied boundary by m-rep atoms is guaranteed. Furthermore, the edge condition can be greatly simplified for a crest region formed by sweeping planar curves in an interpolated single figure m-rep (3.2.3), and the radial curvature condition can be applied to a one-sided spoke field on a medial sheet, which allows the application of this condition to the underlying one-sided spoke field of a blend region in a multi-figure m-rep (4.2.3). The radial curvature condition and the edge condition can be also converted into explicit constraints in a binary fitting process to extract shape models from segmented medical images (5), which enforces the legality of the model boundary and the model interior. The mathematics reviewed here, especially equations (2.13) and (2.19), will be used to calculate the $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices for m-rep atoms in a statistical training process using explicit legality constraints based on these extension matrices.

This section showed how to calculate the $r\mathbf{S}_{rad}$ matrix given the derivatives of $\mathbf{S}$, $\mathbf{p}$, and $\mathbf{U}$ in equations (2.13) and (2.14). Conversely, if we know the $r\mathbf{S}_{rad}$ at each interior point on a medial sheet, we can calculate the first-order derivative of $\mathbf{S}$ by equation (2.13) or (2.14) and recover a full spoke field $\mathbf{S}(\mathbf{u})$ on the interior of the medial sheet by a first order integration. This integration will be used in the interpolation of internal m-rep atoms, to be described in chapter 3.

2.4 Statistics on Single Figure M-reps

This section reviews the mathematics required to build a shape prior on single figure m-reps representing objects of one part. These mathematics are the base for building hierarchical shape priors of multi-figure m-reps in section 4.3. First, I will describe the definition and property of the manifold of single figure m-reps. Then I will review how to calculate statistics of single figure m-reps via this manifold and how to represent differences between m-reps via m-rep residues.

2.4.1 Manifold of Single Figure M-reps

As a primitive in an m-rep, each internal medial atom can be understood as a point on a smooth manifold. Each such internal atom consists of $(\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1})$, where $\mathbf{p}$ is the atom hub position, r is the spoke length, and $\mathbf{U}^{+1,-1}$ are the spoke directions of the two regular spokes. So the manifold of an internal atom is $\mathcal{M}_{int}(1) = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{S}^2 \times \mathcal{S}^2$. Each end medial atom $(\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1}, \eta)$ in an m-rep has one more free parameter $\eta > 1$ controlling the shape of the crest region of the boundary, so each end medial atom is a point on a manifold $\mathcal{M}_{end}(1) = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{S}^2 \times \mathcal{S}^2 \times \mathbb{R}^+$. Let $\mathcal{M}(1)$ be an alias for a manifold of both an internal and an end medial atom without specifying whether it is internal or end. An m-rep of n medial atoms, n_{int} of which are internal medial atoms, and n_{end} of which are end medial atoms, can be treated as a point on the manifold $\mathcal{M}(n) = \mathcal{M}(1) \times \mathcal{M}(1) \times ... \times \mathcal{M}(1) = [\mathcal{M}(1)]^n$, i.e., $\mathcal{M}(n) = [\mathcal{M}(1)]^n = [\mathcal{M}_{int}(1)]^{n_{int}} \times [\mathcal{M}_{end}(1)]^{n_{end}}$, where $n = n_{int} + n_{end}$. For simplicity, $\mathcal{M}(n)$ will be used (instead of $[\mathcal{M}_{int}(1)]^{n_{int}} \times [\mathcal{M}_{end}]^{n_{end}}$) to represent the manifold an m-rep lies on.

The space $\mathcal{M}(n)$ is a particular type of manifold known as the Riemannian symmetric space [Fletcher et al. (2004)]. Let $dis(\mathbf{y}, \mathbf{z}) : \mathcal{M}(n) \times \mathcal{M}(n) \to \mathbb{R}^+ \cup \{0\}$ denote the geodesic distance, i.e., the locally shortest distance on the manifold $\mathcal{M}(n)$, between two points $\mathbf{y}, \mathbf{z} \in \mathcal{M}(n)$. There are a pair of maps $\mathrm{Exp}_{\mathbf{y}}$ and $\mathrm{Log}_{\mathbf{y}}$ that map between $\mathcal{M}(n)$

and its tangent space $T_{\mathbf{y}}\mathcal{M}(n)$ at $\mathbf{y}$ and are inverse to each other. $T_{\mathbf{y}}\mathcal{M}(n)$ can be identified with an Euclidean space $\mathbb{R}^{8n+n_{end}}$ where n_{end} is the number of end atoms in $\mathbf{y}$ and $\mathbf{z}$.

- $\mathrm{Log}_{\mathbf{y}}(\mathbf{z})$ maps point $\mathbf{z}$ to a point in the tangent space $T_{\mathbf{y}}\mathcal{M}(n)$ at $\mathbf{y}$. $\mathbf{y}$ is the origin of $T_{\mathbf{y}}\mathcal{M}(n)$. The geodesic distance between $\mathbf{y}$ and $\mathbf{z}$ is calculated as the length of the vector in the tangent plane from the Log map.

$$dis(\mathbf{y}, \mathbf{z}) = \|\mathrm{Log}_{\mathbf{y}}(\mathbf{z})\| \tag{2.23}$$

- $\mathrm{Exp}_{\mathbf{y}}(\mathbf{v})$ maps a tangent vector $\mathbf{v} \in T_{\mathbf{y}}\mathcal{M}(n)$ to a point on $\mathcal{M}(n)$ along the geodesic curve $\gamma_{\mathbf{v}}(t), t \in [0, 1]$. The distance between $\mathbf{y}$ and the point at which the tangent vector $\mathbf{v}$ is mapped is calculated as follows:

$$dis(\mathbf{y}, \mathrm{Exp}_{\mathbf{y}}(\mathbf{v})) = \|\mathbf{v}\| \tag{2.24}$$

2.4.2 Probability Distributions on Single Figure M-reps

Given the geodesic distance function dis, we can calculate the Fréchet mean $\mathbf{M}$ of N points (m-reps) $\{\mathbf{M}_i \mid \mathbf{M}_i \in \mathcal{M}(n), i = 1, 2, ..., N\}$ by minimizing an average squared geodesic distance of the $\mathbf{M}_i$ to $\mathbf{M}$:

$$\mathbf{M} = \mathrm{FMean}(\mathbf{M}_i) = \arg\min_{\mathbf{M}\in\mathcal{M}(n)} \frac{1}{N}\sum_{i=1}^{N} dis^2(\mathbf{M}, \mathbf{M}_i) = \arg\min_{\mathbf{M}\in\mathcal{M}(n)} \frac{1}{N}\sum_{i=1}^{N} \|\mathrm{Log}_{\mathbf{M}}(\mathbf{M}_i)\|^2 \tag{2.25}$$

The Fréchet mean of N given m-reps is calculated by an iterative optimization [Fletcher et al. (2003)]. Given the Fréchet mean, a novel statistical analysis as an extension to linear principal component analysis (PCA) has been proposed by Fletcher, called principal geodesic analysis (PGA) [Fletcher et al. (2003)]. Analogously to PCA,

PGA tries to find principal geodesic directions, in terms of geodesic curves, on a curved manifold in order to maximize the variance of the original manifold data points. There are two means to do so. The first means is to define the maximization of the data variance on the manifold and to find geodesic curves on the manifold directly maximizing the data variance. Finding these exact variance-maximizing geodesic curves on an m-rep manifold is both theoretically and computationally challenging. Therefore, a practical solution is to approximate the variance-maximizing geodesic curves.

In this approximation, PGA is implemented by projecting all original data points onto the tangent plane $T_{\mathbf{M}}\mathcal{M}(n)$ of the m-rep manifold at the Fréchet mean, via the *Log* map. A PCA is then applied to those projected points in the tangent plane to form principal components in the tangent plane. The approximate principal geodesic curves are given by mapping the linear principal components in the tangent plane $T_{\mathbf{M}}\mathcal{M}(n)$ back to the m-rep manifold $\mathcal{M}(n)$, via the *Exp* map. These geodesic curves approximate the exact variance-maximizing geodesic curves on the m-rep manifold.

The fact that $\mathcal{M}(n)$ is a Riemannian symmetric space simplifies the calculation of geodesic distances a great deal [Fletcher et al. (2004)]. This allows an efficient application of PGA on m-rep models. Another important issue in calculating shape prior statistics is the shape alignment, which is described in detail in [Fletcher et al. (2004)].

In a multi-scale m-rep, PGA can be applied to the entire m-rep to form global statistics. The next natural scale level will be the level of individual figures in a multi-figure m-rep or the level of medial atoms in each m-rep figure while reflecting the hierarchical connectivity among adjacent figures or atoms. To build such hierarchical statistics, we need to represent shape variations left after the global shape variations are removed from sample objects. Left-over shape variations can be described by the differences between m-reps or atoms, called the *residue.* The definition and calculation of the residue will be explained in next subsection 2.4.3.

2.4.3 Residues on M-reps

The residue, i.e., the difference between m-reps or m-rep figures, is calculated via the difference between corresponding atoms. Let $\mathbf{m}_1, \mathbf{m}_2 \in \mathcal{M}(1)$ be two corresponding atoms, and $G(1)$ denote the group of composed transformations of hub translation, spoke magnification(s), and spoke rotations that acts smoothly on each medial atom in $\mathcal{M}(1)$. Then the difference between two atoms can also be represented as a medial atom as follows.

$$\mathbf{m}_1 \ominus \mathbf{m}_2 = g_{\mathbf{m}_2}^{-1} \circ \mathbf{m}_1 \in \mathcal{M}(1) \tag{2.26}$$

where $g_{\mathbf{m}_2}^{-1} \in G(1)$ is a composed atom transformation that transforms an atom $\mathbf{m}_2$ to a pre-determined fixed "origin" atom on $\mathcal{M}(1)$; "$\circ$" means applying the transformation $g_{\mathbf{m}_2}^{-1}$ to atom $\mathbf{m}_1$, carried out component by component: the hub translation is applied to the hub position $\mathbf{p}$; the spoke magnification is applied to the radius r; the spoke rotations are applied to the corresponding spoke directions $\mathbf{U}$.

Assume an m-rep template $\in \mathcal{M}(n)$ has n medial atoms $\{\mathbf{a}_i\}$. $G(n) = [G(1)]^n$ acts smoothly on $\mathcal{M}(n)$ as the transformation between two such m-reps $\in \mathcal{M}(n)$. The difference between two m-reps $\mathbf{M}_1, \mathbf{M}_2 \in \mathcal{M}(n)$ from the same template is defined as

$$\mathbf{M}_1 \ominus \mathbf{M}_2 = \prod_{j=1}^{n} (\mathbf{m}_{1j} \ominus \mathbf{m}_{2j}) \in \mathcal{M}(n) \tag{2.27}$$

So the residue, i.e., the difference, between a pair of m-reps can be represented as a new m-rep. The beauty of this is that we can apply PGA to m-rep or atom residues to get an estimated probability distribution of the m-rep or atom residues, which allows us to calculate multi-scale and multi-level statistics of especially multi-figure m-reps, detailed in chapter 4.

Based on the background knowledge in section 2.3, chapter 3 first reviews previous work on how to interpolate an m-rep, and then it describes how to interpolate a single

figure m-rep by *atom interpolation* on both internal and end medial atoms. The atom interpolation can be extended to multi-figure m-reps as to be described in chapter 4, which covers all the aspects of multi-figure m-reps. Recall that medial mathematics reviewed in section 2.3 can be converted to an explicit geometric constraint in a binary training process to ensure the properness (legality) of fitted shape models for statistical training. This binary training process using the explicit geometric constraint will be detailed in chapter 5.

Chapter 3

Atom Interpolation

An m-rep is composed of a set of discrete medial atoms. In order to use either a single-figure or a multi-figure m-rep in a task such as image segmentation, it is often required to calculate an implied boundary of the m-rep. The calculation of an m-rep implied boundary is realized by interpolations. This chapter will describe a new method to interpolate a single figure m-rep to generate an underlying medial structure and thus to produce an m-rep implied boundary. Next chapter 4 will describe the interpolation of a multi-figure m-rep, which is based on this new interpolation method of a single figure m-rep. There are two types of method to generate an m-rep implied boundary, which are listed as follows and described next.

1. The first type of method: interpolate spoke end points of m-rep medial atoms as boundary points, and use the interpolated boundary points to form an m-rep boundary;

2. The second type of method: interpolate m-rep medial atoms, and use end points of interpolated atom spokes to form an m-rep boundary.

In the first type of method, spoke end points are interpolated as boundary points. There can be locally improper shapes in the calculated boundary, such as surface creases or folds. Even though some methods of this type do interpolate normals at control

points, improper shapes can still result, from the lack of control on surface normals of the entire object boundary. These improper shapes are computationally expensive to detect or avoid. Because this type of method does not provide an interpolated m-rep, it is also difficult to have a valid parameterization for the object interior volume without the underlying medial structure.

In the second type of method, to be developed in this chapter, medial atoms are directly interpolated. The properness (legality) of the implied m-rep boundary by interpolated medial atoms can be guaranteed if the legality of interpolated atom spokes is maintained. This type of method produces an interpolated m-rep and thus an m-rep implied boundary. Unlike the first type of method, it also interpolates the object interior.

As summarized in section 2.2, the difficulty of interpolating medial atoms arises from the fact that the components in a medial atom need to simultaneously fulfill certain legality constraints (in section 2.2.1). The new method, belonging to this second type of method, overcomes this difficulty by basing its interpolation on differential properties, i.e., derivatives of the medial sheet and spokes. These derivatives are summarized by the 2×2 matrices called the *radial shape operator* and the *edge shape operator*. These operators were described in section 2.3. They both will be used in the new m-rep interpolation method to be detailed in this chapter.

Section 3.1 reviews previous attempts to generate m-rep implied boundaries by interpolations. Section 3.2 details the new interpolation method of single figure m-reps. Section 3.3 shows the results from the application of the new interpolation method to synthetic and real-world m-rep models. Section 3.4 lists related open problems and concludes this chapter.

3.1 Previous Work on Generating M-Rep Implied Boundaries by Interpolations

This section reviews various proposed methods of both types of method to generate an m-rep implied boundary.

In the first type of method, spoke ends of m-rep atoms are interpolated as boundary points. In computer graphics, methods have been proposed on boundary point interpolations in [Biermann et al. (2000); Solem et al. (2007); Stiller (2007)]. These methods are used to generate a smooth boundary from a sparse set of control points with or without the consideration of surface normals. These methods can be applied to a mesh formed by m-rep spoke end points.

Thall (2004) proposed to incorporate surface normals into an adapted Catmull-Clark subdivision scheme and to interpolate rather than approximate the surface control points. To interpolate the original surface control points, he shifted them an appropriate amount along their normals. To incorporate normals, he post-processed a subdivision control mesh after two levels of subdivision, by rotating local surface patches centered at m-rep spoke ends to match the surface normals at those points to the normals implied by corresponding control atoms. This processed control mesh is then plugged back in the Catmull-Clark subdivision to calculate a resulting mesh. In the resulting mesh, the surface normals at atom spoke ends approximate the normals implied by corresponding control atoms. Because of the lack of legality constraints on the resulting mesh, Thall's method can generate unexpected surface creases or folds, which is a common problem of interpolations that attempt to exact-match the control points.

The first type of method does not provide a valid parameterization of an object interior volume. An interpolated object boundary and its normals do not suffice to provide such a parameterization, which is important for gathering image intensities relative to an object geometry. When marching along surface normals into the interior

of an m-rep to gather intensity profiles, one does not know the distance to go before the normal vectors cross one another.

In the second type of method, medial atoms are directly interpolated. The interpolated atoms provide a valid parameterization of the object interior if the legality of the interpolated atoms is maintained. I will review two previous methods to interpolate m-rep atoms.

Crouch (2003) used cubic b-spline patches to interpolate both a medial sheet and an m-rep implied boundary. For each point on the B-spline patch of the medial sheet, there is a corresponding point on the B-spline patch of the m-rep implied boundary. The correspondence is determined by the same parameters of both the B-spline patches. Each pair of these corresponding points are connected to form an interpolated spoke as a part of an interpolated atom. A pair of the interpolated spokes, one at each side of the medial sheet, and their shared medial hub position are combined into an interpolated atom. Such an interpolation works efficiently for the volumetric subdivision proposed in [Crouch (2003)] - chapter 3. However, because this method is based on separate interpolations on the medial sheet and on the m-rep implied boundary, it does not guarantee the legality of the interpolated spokes. Interpolated spokes could cross one another without being detected.

Another method to interpolate m-rep atoms is based on a *weighted geodesic average.* Recall that an m-rep atom can be considered as a point on a Riemannian manifold $\mathcal{M}$. The weighted geodesic average of points on such a manifold is analogous to the weighted average of points on a unit sphere $\mathcal{S}^2$, which has been studied by Buss and Fillmore [Buss and Fillmore (2001)]. The weights used in the weighted average are provided by existing weight functions in $\mathbb{R}^3$, such as bilinear weights, cubic Bézier weights, and cubic B-spline weights. The approach in [Buss and Fillmore (2001)] was extended by Fletcher et al. to interpolate diffusion tensors on a Riemannian manifold. The weighted geodesic average of medial atoms follows the idea in [Fletcher et al. (2004)] by treating each medial atom

as a point on a Riemannian manifold $\mathcal{M}$.

Given N atoms $\{\mathbf{A}_i = (\mathbf{p}_i, r_i, \mathbf{U}_i^+, \mathbf{U}_i^-) \in$ the Riemannian manifold $\mathcal{M} = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{S}^2 \times \mathcal{S}^2, i = 1, 2, ..., N\}$, and given the weight $\omega = (\omega_1, \omega_2, ..., \omega_N)$, the interpolated atom $\mathbf{A}(\omega)$ is defined in equation 3.1:

$$\mathbf{A}(\omega) = \arg\min_{\mathbf{A}} \sum_{i=1}^{N} \omega_i \cdot dis^2(\mathbf{A}, \mathbf{A}_i) \tag{3.1}$$

The weighted average of the atoms $\{\mathbf{A}_i\}$ $(i = 1...N)$ is based on a geodesic distance $dis(\mathbf{A}, \mathbf{A}_i)$ between two atoms $\mathbf{A}$ and $\mathbf{A}_i$ on $\mathcal{M}$. When the weight components in ω are all equal to $\frac{1}{N}$, the resulting atom is exactly the Fréchet (intrinsic) mean of N atoms $\mathbf{A}_i$; when $N = 2$ and when $\omega = (1 - \omega_2, \omega_2)$, $(\omega_2 \in [0, 1])$, $\mathbf{A}(\omega)$ forms the exact geodesic path on $\mathcal{M}$ between $\mathbf{A}_0$ and $\mathbf{A}_1$. For $N \geq 2$ and ω determined by weight functions, such as cubic B-spline weights, equation (3.1) defines $\mathbf{A}(\omega)$ as the weighted geodesic average of the N atoms.

Equation (3.1) is solved by a minimization. The existence and uniqueness of the solution to the minimization are not generally guaranteed. However, it has been discussed in [Fletcher et al. (2004)] that the solution to this minimization exists and is unique if the data are well-localized on a Riemannian manifold. Because the interpolation of m-rep atoms is always applied to a set of adjacent m-rep atoms, which are indeed well localized, the interpolation of m-rep atoms by the weighted geodesic average (3.1) is thus expected to converge well and to yield unique solutions. Empirically, interpolated atoms $\mathbf{A}(\omega)$ have been successfully calculated for different control atoms $\{\mathbf{A}_i\}$ and weights ω. The solution to equation (3.1) ends up optimizing each one of the atom components separately. An interpolated m-rep by this means might still yield an improper shape because it does not adequately reflect the relations within all components of a medial atom.

Section 3.2 describes a new direct interpolation on m-rep atoms. The new interpolation method tries to incorporate all the components of a medial atom by explicitly using
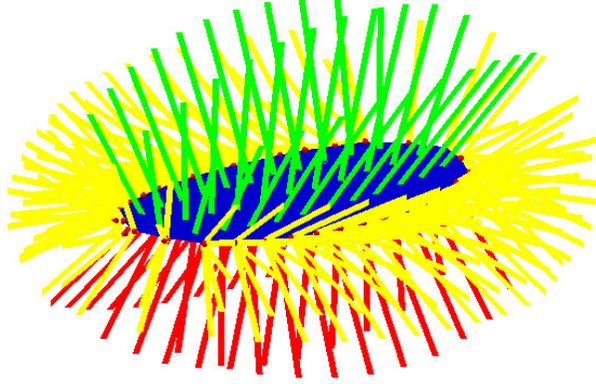
Figure 3.1: Demonstration of the interpolated spoke fields from a discrete m-rep: the medial sheet (in dark blue) with an end curve as its boundary, the interpolated spokes (red and green on each side of the medial sheet) for the internal region, and the interpolated spokes (yellow) for the crest region.

the radial and edge shape operators. An interpolated medial sheet and its attached two-sided spoke field, generated by this method, is demonstrated in figure 3.1.

3.2 Atom Interpolation in Single Figure M-reps

Each discrete single figure m-rep consists of a quad-mesh of two types of medial atom: internal atoms and end atoms. This section starts with a brief review of the two types of medial atom, and then it describes the interpolations on both of them.

3.2.1 Internal and End Medial Atoms

Recall that as shown in figure 2.1, each internal medial atom $\mathbf{A}$ has components $\{\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1}\}$, and each end medial atom is a combination of a *special* internal atom $\{\mathbf{p}, r, \mathbf{U}^{+1}, \mathbf{U}^{-1}\}$ and a bisector spoke $\mathbf{S}^0 = r\eta\mathbf{U}^0$, where $\mathbf{U}^0 = \frac{\mathbf{U}^{+1}+\mathbf{U}^{-1}}{\|\mathbf{U}^{+1}+\mathbf{U}^{-1}\|}$. In practice, $\mathbf{U}^{+1}$ and $\mathbf{U}^{-1}$ are never back to back in an end atom. Therefore, $\mathbf{U}^0$ and $\mathbf{S}^0$ are well defined.

In a single figure m-rep, its special internal atoms in end atoms and regular internal atoms form a quad-mesh of internal atoms. This quad-mesh of internal atoms consists

of quad-patches. Each quad-patch has one internal atom at one of its four corners. The entire quad-mesh of internal atoms will be interpolated quad-patch by quad-patch, taking care for some level of continuity at the patch boundaries. The interpolation on internal atoms is detailed in section 3.2.2. The generations of end atoms and the crest are detailed in section 3.2.3, based on interpolated internal atoms.

3.2.2 Interpolation of Internal Medial Atoms

The interpolation of internal medial atoms includes the interpolation of the medial sheet $\mathbf{p}$, the interpolation of $r\mathbf{S}_{rad}$ matrices, and the interpolation of spokes $\mathbf{S}$. These interpolations are applied to one quad-patch, of four control atoms, at a time. It suffices to describe these interpolations applied to one quad-patch.

Assume the four control atoms in a quad-patch are $\mathbf{A}_i = \{\mathbf{p}_i, r_i, \mathbf{U}_i^{+1}, \mathbf{U}_i^{-1}\}, i = 1, 2, 3, 4$. The hub positions $\mathbf{p}_i$ are first interpolated by a *point-and-normal interpolation* to form a continuous medial sheet $\mathbf{p}(u_1, u_2)$ patch. Then $r\mathbf{S}_{rad}$ matrices are estimated for control atom spokes, and the estimated matrices are interpolated to ensure the legality of each interpolated $r\mathbf{S}_{rad}(u_1, u_2)$ matrix. The interpolated medial sheet $\mathbf{p}(u_1, u_2)$ and interpolated $r\mathbf{S}_{rad}$ matrices are then used to calculate the first-order derivatives of spokes $\mathbf{S}$, and these derivatives are integrated to calculate interpolated full spokes. The interpolated medial sheet and interpolated spokes are combined to form interpolated internal medial atoms.

Medial Sheet $\mathbf{p}(u_1, u_2)$ Interpolation

For each internal atom in a quad-patch, its hub position gives a medial sheet point, and the difference vector between its two spokes gives the medial sheet normal at that point. The points and normals of the medial sheet, given by the four control atoms, are $\mathbf{p}_i$ and $\mathbf{n}_i = \frac{\mathbf{U}_i^{+1} - \mathbf{U}_i^{-1}}{|\mathbf{U}_i^{+1} - \mathbf{U}_i^{+1}|}$ for $i = 1, 2, 3, 4$, respectively. In practice, the two spoke directions in an internal atom are never identical, so the medial sheet normals are well defined. With

54

the points and normals known for the medial sheet at control atoms, the medial sheet $\mathbf{p}$ is interpolated by a point-and-normal interpolation via a cubic Hermite patch. A cubic Hermite patch requires 16 control elements:

$$\mathbf{H}_{control} = \begin{pmatrix} \mathbf{h}_{11} & \mathbf{h}_{12} & \mathbf{h}_{13} & \mathbf{h}_{14} \\ \mathbf{h}_{21} & \mathbf{h}_{22} & \mathbf{h}_{23} & \mathbf{h}_{24} \\ \mathbf{h}_{31} & \mathbf{h}_{32} & \mathbf{h}_{33} & \mathbf{h}_{34} \\ \mathbf{h}_{41} & \mathbf{h}_{42} & \mathbf{h}_{43} & \mathbf{h}_{44} \end{pmatrix},$$

which is written in a matrix form, including four control points, eight first-order derivatives, and four second-order cross derivatives.

The four control points are set to be points $\mathbf{p}_i$, so the cubic Hermite patch interpolates these points as its intrinsic property. In the current implementation, the four second-order cross derivatives are set to be $(0,0,0)$, so the patch also interpolates its eight first-order derivatives. The following describes how to set the remaining eight first-order derivatives.

First-order derivatives $\mathbf{p}'_{u_1,i}$ and $\mathbf{p}'_{u_2,i}$ of the medial sheet $\mathbf{p}$ are estimated for each control point, where $i = 1, 2, 3, 4$. These estimations are carried out by finite differences between neighboring hub positions $\mathbf{p}_i$, along each one of the parametric directions u_1 and u_2. $\mathbf{p}'_{u_1,i}$ and $\mathbf{p}'_{u_2,i}$ are then projected onto the medial sheet tangent planes determined by the given normals $\mathbf{n}_i$: $\mathbf{p}^T_{u_j,i} = \mathbf{p}'_{u_j,i} - (\mathbf{p}'_{u_j,i} \cdot \mathbf{n}_i)\mathbf{n}_i$, for $j = 1, 2$. Let $\mathbf{p}^T_{u_j,i}$ for $j = 1, 2; i = 1, 2, 3, 4$ be the eight first-order derivatives in the cubic Hermite patch. This patch interpolates $\mathbf{p}^T_{u_j,i}$, so it interpolates each tangent plane, spanned by $\mathbf{p}^T_{u_1,i}$ and $\mathbf{p}^T_{u_2,i}$, and determined by $\mathbf{n}_i$, at each one of the four control points. Therefore, at each control point, the normal of the cubic Hermite patch is exactly $\mathbf{n}_i$. A complete description of such a cubic Hermite patch is shown as follows:

$$\mathbf{p}(u_1, u_2) = \left(\begin{array}{cccc} H_1(u_1) & H_2(u_1) & H_3(u_1) & H_4(u_1) \end{array} \right) \cdot \mathbf{H}_{control} \cdot \left(\begin{array}{c} H_1(u_2) \\ H_2(u_2) \\ H_3(u_2) \\ H_4(u_2) \end{array} \right),$$

where $H_{1,2,3,4}(s)$ are cubic Hermite weight functions: $H_1(s) = 2s^3 - 3s^2 + 1, H_2(s) = -2s^3 + 3s^2, H_3(s) = s^3 - 2s^2 + s, H_4(s) = s^3 - s^2$.

The elements in $\mathbf{H}_{control}$ are set as follows. $\mathbf{h}_{11}, \mathbf{h}_{12}, \mathbf{h}_{21}, \mathbf{h}_{22}$ are the four control points: $\mathbf{h}_{11} = \mathbf{p}_1, \mathbf{h}_{12} = \mathbf{p}_2, \mathbf{h}_{21} = \mathbf{p}_4, \mathbf{h}_{22} = \mathbf{p}_3$; $\mathbf{h}_{31}, \mathbf{h}_{32}, \mathbf{h}_{41}, \mathbf{h}_{42}$ are the four tangential first-order derivatives of $\mathbf{p}$ along the u_1 direction: $\mathbf{h}_{31} = \mathbf{p}_{u_1,1}^T, \mathbf{h}_{32} = \mathbf{p}_{u_1,2}^T, \mathbf{h}_{41} = \mathbf{p}_{u_1,4}^T, \mathbf{h}_{42} = \mathbf{p}_{u_1,3}^T$; $\mathbf{h}_{13}, \mathbf{h}_{14}, \mathbf{h}_{23}, \mathbf{h}_{24}$ are the four tangential first-order derivatives of $\mathbf{p}$ along the u_2 direction: $\mathbf{h}_{13} = \mathbf{p}_{u_2,1}^T, \mathbf{h}_{14} = \mathbf{p}_{u_2,2}^T, \mathbf{h}_{23} = \mathbf{p}_{u_2,4}^T, \mathbf{h}_{24} = \mathbf{p}_{u_2,3}^T$; $\mathbf{h}_{33}, \mathbf{h}_{34}, \mathbf{h}_{43}, \mathbf{h}_{44}$ are the four tangential second-order cross derivatives and are set as $(0, 0, 0)$ in the current implementation.

This cubic Hermite interpolation is then applied to the entire m-rep quad-mesh, quad-patch by quad-patch. Each pair of adjacent quad-patches are set to share the same tangential first-order derivatives across their shared boundary, so C^1 continuity is achieved across boundaries between adjacent quad-patches in the interpolated medial sheet.

Because a cubic Hermite patch has an explicit form, the derivatives of the interpolated medial sheet $\mathbf{p}(u_1, u_2)'_{u_j}$ for $j = 1, 2$ can be analytically evaluated, as shown below. These derivatives will be used in the spoke interpolation.

$$\mathbf{p}(u_1, u_2)'_{u_1} = \begin{pmatrix} H'_1(u_1) & H'_2(u_1) & H'_3(u_1) & H'_4(u_1) \end{pmatrix} \cdot \mathbf{H}_{control} \cdot \begin{pmatrix} H_1(u_2) \\ H_2(u_2) \\ H_3(u_2) \\ H_4(u_2) \end{pmatrix}$$

and

$$\mathbf{p}(u_1, u_2)'_{u_2} = \begin{pmatrix} H_1(u_1) & H_2(u_1) & H_3(u_1) & H_4(u_1) \end{pmatrix} \cdot \mathbf{H}_{control} \cdot \begin{pmatrix} H'_1(u_2) \\ H'_2(u_2) \\ H'_3(u_2) \\ H'_4(u_2) \end{pmatrix},$$

where $H'_{1,2,3,4}(s)$ are the first-order derivatives of the cubic Hermite weight functions:
$H'_1(s) = 6s^2 - 6s, H'_2(s) = -6s^2 + 6s, H'_3(s) = 3s^2 - 4s + 1, H'_4(s) = 3s^2 - 2s.$

$r\mathbf{S}_{rad}$ Matrix Interpolation

Section 2.3 showed that the first-order derivatives of spokes $\mathbf{S}$ can be calculated from known $r\mathbf{S}_{rad}$ matrices and known derivatives of $\mathbf{p}$, by equations (2.9) and (2.12). These derivatives can be integrated to recover an interpolated full spoke field $\mathbf{S}(u_1, u_2)$. However, this integration requires an interpolated medial sheet, just described, and interpolated $r\mathbf{S}_{rad}$ matrices, to be described next.

$r\mathbf{S}_{rad}$ matrices of four control atoms in a quad-patch are numerically estimated by finite differences between components of adjacent atoms. The estimated $r\mathbf{S}_{rad}$ matrices are checked against the radial curvature condition, detailed in section 2.3, by comparing the greater real eigenvalue of each matrix with 1.0. Given a legal m-rep, each estimated $r\mathbf{S}_{rad}$ matrix should pass this check with its greater real eigenvalue smaller than 1.0.

57

Each internal control atom has two $r\mathbf{S}_{rad}$ matrices, one for each spoke. The estimated $r\mathbf{S}_{rad}$ matrices for spokes at the same side of the medial sheet are interpolated by the following three steps:

1. The estimated $r\mathbf{S}_{rad}$ matrices are decomposed into eigenvalues and eigenvectors;

2. The eigenvalues and eigenvectors are interpolated, with the interpolated eigenvalues enforced to stay legal, i.e., to be smaller than 1.0;

3. The interpolated eigenvalues and interpolated eigenvectors are composed back into interpolated $r\mathbf{S}_{rad}$ matrices.

where the first step is straightforward. The second step and the third step are detailed below.

Each $r\mathbf{S}_{rad}$ matrix has two eigenvalues $\lambda_{rk}, k = 1, 2$ and two eigenvectors $\mathbf{e}_{rk}, k = 1, 2$. Assume the two eigenvalues from each $r\mathbf{S}_{rad}$ matrix are ordered $\lambda_{r1} \geq \lambda_{r2}$, and assume $\mathbf{e}_{rk}$ is the corresponding eigenvector to λ_{rk}. In order to interpolate the $r\mathbf{S}_{rad}$ matrices of the four control atoms bounding a quad-patch, I separately interpolate the λ_{r1}'s, λ_{r2}'s, $\mathbf{e}_{r1}$'s, and $\mathbf{e}_{r2}$'s.

To interpolate the eigenvalues $\lambda_{r1,i} < 1 \in \mathbb{R}, i = 1, 2, 3, 4$, each $\lambda'_{r1,i} = (1 - \lambda_{r1,i}) > 0$ is treated as an element in a multiplicative Lie group $\mathbb{R}^+$. The interpolation in this group is calculated by a weighted arithmetic average of the logarithms of the elements. The $\lambda'_{r1,i}$ are interpolated in $\mathbb{R}+$ and then converted back to $\mathbb{R}$ via an exponential map. Assume the weights for the arithmetic average are $w_i(u_1, u_2)$, so the interpolated eigenvalue is $\lambda_{r1}(u_1, u_2) = 1 - e^{\sum_{i=1}^{4} w_i(u_1, u_2) \cdot ln(\lambda'_{r1,i})}$. The same method is applied to the other four eigenvalues $\lambda_{r2,i}$ to get the interpolated eigenvalue $\lambda_{r2}(u_1, u_2)$. By this means, an interpolated eigenvalue is always smaller than 1.0, fulfilling the legality condition.

To interpolate the eigenvectors $\mathbf{e}_{r1,i}, i = 1, 2, 3, 4$, each eigenvector $\mathbf{e}_{r1,i}$ is treated as an element in a Lie group $\mathcal{SO}(2)$. Geometrically, the manifold of $\mathcal{SO}(2)$ forms a 2D unit circle, and each point on this circle represents a 2D rotational angle. Each $\mathbf{e}_{r1,i}$ is then

converted to a 2D rotational angle $\alpha_i \in [0, 2\pi)$. The interpolation of four eigenvectors $\mathbf{e}_{r1,i}$ is realized by an arithmetic average of their four converted angles, with the special consideration that a 2D unit circle is a wrapped-around manifold. Assume the weighted average of the angles is $\alpha(u_1, u_2)$, given the weights $w_i(u_1, u_2)$, then the interpolated eigenvector is $\mathbf{e}_{r1}(u_1, u_2) = (cos(\alpha(u_1, u_2)), sin(\alpha(u_1, u_2)))$. The same method is applied to the other eigenvectors $\mathbf{e}_{r2,i}$ to get the interpolated eigenvector $\mathbf{e}_{r2}(u_1, u_2)$.

The interpolated eigenvalues and eigenvectors are then combined into an $r\mathbf{S}_{rad}(u_1, u_2)$ matrix by the inverse of an eigen-decomposition, as shown below. The interpolated $r\mathbf{S}_{rad}(u_1, u_2)$ matrix is legal because both of its eigenvalues are smaller than 1.0, which meets the radial curvature condition.

$$r\mathbf{S}_{rad}(u_1, u_2) = P \cdot \begin{pmatrix} \lambda_{r1}(u_1, u_2) & 0 \\ 0 & \lambda_{r2}(u_1, u_2) \end{pmatrix} \cdot P^{-1},$$

where $\mathbf{P} = \begin{pmatrix} \mathbf{e}_{r1}(u_1, u_2)^T & \mathbf{e}_{r2}(u_1, u_2)^T \end{pmatrix}$.

This $r\mathbf{S}_{rad}$ interpolation is applied to both sets of $r\mathbf{S}_{rad}$ matrices in a quad-patch, each for one side of the medial sheet. It is then applied to the m-rep quad-mesh, patch by patch. C^0 continuity is achieved for the interpolated $r\mathbf{S}_{rad}$ matrices across a shared boundary between two adjacent patches. Because $r\mathbf{S}_{rad}$ matrices record the first-order derivatives of $\mathbf{p}$ and $\mathbf{S}$, the spoke interpolation, based on an integration over these first-order derivatives, is designed to have C^1 continuity across a shared quad-patch boundary.

Problems with this $r\mathbf{S}_{rad}$ interpolation are its speed and the existence of medial umbilic points, with repeated eigenvalues. Furthermore, a numerical error in the $r\mathbf{S}_{rad}$ estimation might yield a matrix with complex eigenvalues and complex eigenvectors. These problems are addressed in section 3.4.

The following details the spoke interpolation based on the interpolations of the medial sheet and of the $r\mathbf{S}_{rad}$ matrices.

Spoke $\mathbf{S}(u_1, u_2)$ Interpolation

The first-order derivatives $\frac{\partial \mathbf{S}}{\partial \mathbf{u}}$ of the spokes are computed by equations (2.9) and (2.12) in section 2.3, and they are integrated to generate interpolated spokes. The spoke interpolation consists of the following steps, applied to one quad-patch at a time:

1. First-order derivatives $\frac{\partial \mathbf{S}^{+1}}{\partial \mathbf{u}}$ are calculated for spokes at one side of the medial sheet by equations (2.9) and (2.12), based on the interpolated medial sheet $\mathbf{p}$ and the interpolated $r\mathbf{S}_{rad}$ matrices in a quad-patch;

2. Spokes $\mathbf{S}^{+1}(u_1, u_2)$ are interpolated, for one side of the medial sheet, by an integration of $\frac{\partial \mathbf{S}^{+1}}{\partial \mathbf{u}}$ along a straight path between $(0,0)$ and (u_1, u_2):

$$
\mathbf{S}^{+1}_{(u_1,u_2)} = \mathbf{S}^{+1}_{(0,0)} + \int_{0:(0,0)}^{1:(u_1,u_2)} \left(\frac{\partial \mathbf{S}^{+1}}{\partial u_1} du_1 + \frac{\partial \mathbf{S}^{+1}}{\partial u_2} du_2 \right) \tag{3.2}
$$

   where $\mathbf{S}_{(0,0)}$ is $\mathbf{S}^{+1}_1$, the spoke of one of the four control atoms.

   This integral is numerically calculated along the line between $(0,0)$ and (u_1, u_2), divided into equally spaced intervals. On each of the intervals, the partial derivatives $\frac{\partial \mathbf{S}^{+1}}{\partial u_1}$ and $\frac{\partial \mathbf{S}^{+1}}{\partial u_2}$ are considered constant. Therefore, this integration requires the integration path to be finely divided;

3. The same integration is used to generate the interpolated spokes $\mathbf{S}^{-1}(u_1, u_2)$ on the other side of the medial sheet;

4. Spoke directions and spoke length are both implied by the interpolated full spokes: $\mathbf{U}^{+1,-1}(u_1, u_2) = \frac{\mathbf{S}^{+1,-1}(u_1,u_2)}{\|\mathbf{S}^{+1,-1}(u_1,u_2)\|}$, and $r(u_1, u_2) = \|\mathbf{S}^{+1}(u_1, u_2)\|$;

Each interpolated internal atom $\mathbf{A}(u_1, u_2)$ has components $\{\mathbf{p}(u_1, u_2),\ \mathbf{U}^{+1}(u_1, u_2),\ \mathbf{U}^{-1}(u_1, u_2),\ r(u_1, u_2)\}$. The internal atom interpolation is applied to all the quad-patches in an m-rep quad-mesh, which yields a new quad-mesh of interpolated internal atoms.

60

3.2.3 Generation of Interpolated End Atoms and a Swept Crest

Generation of interpolated end atoms and a swept crest for an m-rep depends on the interpolated internal atoms created by the internal atom interpolation just described. Each interpolated internal atom on the boundary of the quad-mesh of interpolated internal atoms is extended by a bisector spoke to form an interpolated end atom. Then a crest region of sweeping planar curves, each of which is fitted to one interpolated end atom, is generated, and the legality of this crest is checked.

Generation of Interpolated End Atoms

Each interpolated end atom $\mathbf{A}_k^{end}$ is generated by adding a bisector spoke $\mathbf{S}_k^0$ to each interpolated internal atom $\mathbf{A}_k^{in} = (\mathbf{p}_k, \mathbf{S}_k^{+1,-1} = r_k \mathbf{U}_k^{+1,-1})$, $k \in [1, N]$, on the boundary of the quad-mesh of interpolated internal atoms, where N is the total number of these interpolated internal atoms. Assume $\mathbf{S}_k^0 = \eta_k r_k \mathbf{U}_k^0$, where r_k is the spoke length of $\mathbf{A}_k^{in}$. The bisector crest parameter η_k and the bisector spoke direction $\mathbf{U}_k^0$ are yet to be specified for each bisector spoke $\mathbf{S}_k^0$ to be added.

Among all the interpolated internal atoms $\{\mathbf{A}_k^{in}\}$, some are the special internal atoms of original end atoms in the m-rep, so the η's of bisector spokes to be added to these special internal atoms are set to be the η's of their corresponding end atoms in the m-rep. For each interpolated internal atom $\mathbf{A}_k^{in}$ that is not a special internal atom, it lies between two original end atoms, and the η_k of $\mathbf{S}_k^0$ is interpolated between the η's of the two original end atoms. In the current implementation, a linear interpolation is used.

Each bisector spoke direction $\mathbf{U}_k^0$ is determined by the pair of regular spokes of $\mathbf{A}_k^{in}$ as $\mathbf{U}_k^0 = \frac{\mathbf{U}_k^{+1} + \mathbf{U}_k^{-1}}{\|\mathbf{U}_k^{+1} + \mathbf{U}_k^{-1}\|}$. In practice, the pair of regular spokes of $\mathbf{A}_k^{in}$ are never back to back, so $\mathbf{U}_k^0$ is well defined.

As a combination of an interpolated internal atom $\mathbf{A}_k^{in}$ and an added bisector spoke $\mathbf{S}_k^0$, each $\mathbf{A}_k^{end}$ is one of a series of interpolated end atoms at the boundary of the quad-

mesh of interpolated atoms.

Generation of a Swept Crest

A crest is yet to be generated for the m-rep, from the series of interpolated end atoms. For each interpolated end atom with a pair of regular spokes $\mathbf{S}^{+1,-1}$ and a bisector spoke $\mathbf{S}^0$, it is necessary to generate a series of spokes smoothly filling the gaps between $\mathbf{S}^{+1}$ and $\mathbf{S}^0$, and between $\mathbf{S}^0$ and $\mathbf{S}^{-1}$, to form a part of a boundary crest.

Based on the swept surface generation proposed in [Damon (2008)], a planar curve is fitted to the three spokes of each interpolated end atom. This planar curve then sweeps along the edge curve of the interpolated medial sheet, generating a swept boundary crest. According to [Damon (2008)], the calculation of the principal edge curvature for such a swept crest is greatly simplified. This simplified calculation is used to check the legality of the swept crest. How a planar curve is fitted to each interpolated end atom and how the crest region is checked against a legality condition are detailed next.

An ellipse curve is fitted to each interpolated end atom to interpolate its three spoke ends. The long axis of each interpolating ellipse is set to align with the bisector spoke, so the medial sheet of the partial ellipse curve is a segment on the bisector spoke $\mathbf{S}^0$. Without loss of generality, the three spokes of an interpolated end atom and the interpolating partial ellipse are described in $\mathbb{R}^2$, as shown in figure 3.2-left. The resulting partial ellipse is transformed to 3D by a change of coordinate.

Assume the ellipse is given by an implicit form $\frac{(x-x_0)^2}{a^2} + \frac{y^2}{b^2} = 1$, where $a, b, x_0 \in \mathbb{R}$ are the unknowns, the constraints for the ellipse to fit an interpolated end atom are described as follows:

- The ellipse interpolates the spoke end points $\mathbf{e}^{\pm 1} = (x_L, \pm y_L)$ and $\mathbf{e}^0 = (x_R, 0)$, corresponding to spokes $\mathbf{S}^{\pm 1}$ and $\mathbf{S}^0$, respectively, where $x_L, y_L, x_R \in \mathcal{R}$ are determined by an interpolated end atom;
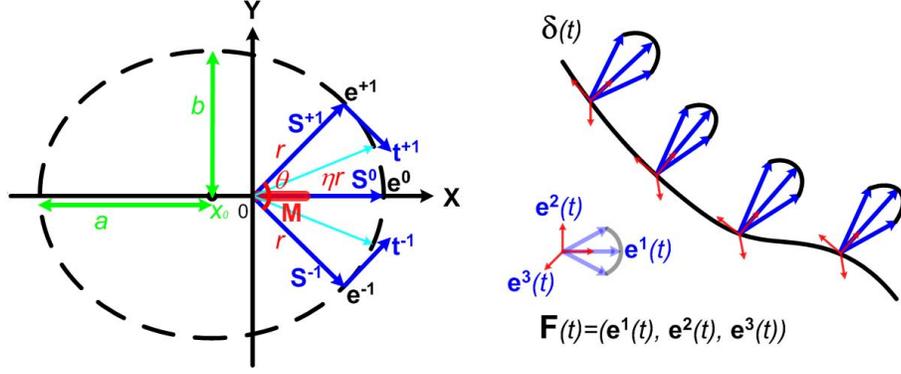
Figure 3.2: Left: three spokes $\mathbf{S}^{+1,-1,or0}$ (thick arrows in dark blue) with length $r, r, or \eta n$ of an end atom forming a plane; interpolated spokes (thin arrows in cyan) filling the gaps between the three spokes of each end atom; the medial sheet of the partial ellipse fit to the three spoke ends is $\mathbf{M}_{ellipse}$, the thick red segment on $\mathbf{S}^0$ bisector spoke. Right: planar curves sweeping along the end curve $\delta(t)$ of the medial sheet; a local frame implied by end atom spokes also sweeps along $\delta(t)$, forming a frame field $\mathbf{F}(t)$.

- The tangents of the ellipse at the two spoke ends $\mathbf{e}^{\pm 1}$ are set as $\mathbf{t}^{\pm 1} = (y_L, \mp x_L)$, so the ellipse is perpendicular to spokes $\mathbf{S}^{\pm 1}$.

Given the pose of the ellipse, it is perpendicular to the bisector spoke by default. Therefore, the constraints above do not include the perpendicularity between the ellipse and the bisector spoke. The unknowns (a, b, x_0) are then solved as follows:

$$x_0 = \frac{x_L(x_L^2 - x_R^2 + y_L^2)}{2x_L^2 - 2x_L x_R + y_L^2} \tag{3.3}$$

$$a = \frac{-x_L^3 + x_R(2x_L^2 + y_L^2) - x_L(x_R^2 + y_L^2)}{2x_L^2 - 2x_L x_R + y_L^2} \tag{3.4}$$

$$b = \sqrt{\frac{(x_L^2 - x_L x_R + y_L^2)^2}{2x_L^2 - 2x_L x_R - y_L^2}} \tag{3.5}$$

if two conditions hold: $2x_L^2 - 2x_L x_R - y_L^2 > 0$ and $-x_L^2 + x_R(2x_L + y_L^2) - x_L(x_R^2 + y_L^2) > 0$.

Let θ be the angle between two regular spokes $\mathbf{S}^{+1,-1}$, and let η be the bisector length parameter in an interpolated end atom. The two conditions might be violated when θ and η have extreme values. In practice, I maintain $\theta \in (\frac{\pi}{3}, \frac{2\pi}{3})$ and $\eta \in (1.05, 1.2)$. With

these empirically determined ranges for θ and η, a, b, and x_0 have been successfully solved for different m-reps.

Only the part of the ellipse connecting the three spoke ends of an interpolated end atom is used as a part of the boundary crest. The medial sheet of the partial ellipse is a segment lying on the bisector spoke, as the thick red segment $\mathbf{M}_{ellipse}$ shown in figure 3.2-left. $\mathbf{M}_{ellipse}$ and its attached filling spokes $\mathbf{S}$ are analytically calculated from the solved equation of the ellipse: $\mathbf{M}_{ellipse} = \{(\frac{a^2-b^2}{a}\cos\phi + x_0, 0)|\phi \in [0, \frac{\theta}{2}]\}$, $\mathbf{S}^{\pm1}(\phi) = ((a - \frac{a^2-b^2}{a})\cos\phi, \pm b\sin\phi), \phi \in (0, \frac{\theta}{2}]$, and $\mathbf{S}^0(0) = (a - \frac{a^2-b^2}{a}, 0)$. The filling spokes for an interpolated end atom are shown as the cyan spokes in figure 3.2-left. Each fitted partial ellipse forms a planar cross section of an m-rep crest. All the fitted partial ellipses form a swept m-rep crest. The following describes how to check the legality of this swept crest.

Recall that in section 2.3 a principal edge curvature was defined for the legality condition of edge spokes on a medial sheet edge curve. For a crest region formed by sweeping planar curves, the calculation of this principal edge curvature is greatly simplified [Damon (2008)].

Assume the edge curve of the interpolated medial sheet is $\mathbf{p}_\delta(v)$, parameterized by v. Each interpolated end atom $\mathbf{p}_i, \mathbf{U}^{+1,0,-1}, r_i, \eta_i$ lying on this edge curve also determines a local frame $\mathbf{F}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$, where $\mathbf{e}_1$ is the bisector spoke direction $\mathbf{U}^0$, where $\mathbf{e}_2 = \frac{\mathbf{U}^{+1}-\mathbf{U}^{-1}}{\|\mathbf{U}^{+1}-\mathbf{U}^{-1}\|}$, and where $\mathbf{e}_3 = \mathbf{e}_1 \times \mathbf{e}_2$, as shown in figure 3.2-right. The local frames of interpolated end atoms form a frame field $\mathbf{F}(v) = (\mathbf{e}_1(v), \mathbf{e}_2(v), \mathbf{e}_3(v))$ on $\mathbf{p}_\delta(v)$.

Given this local frame field $\mathbf{F}(v)$ on the edge curve $\mathbf{p}_\delta(v)$, the derivative of the edge curve $\mathbf{p}_\delta(v)$ can be written in $\mathbf{F}(v)$ as:

$$\mathbf{p}_\delta(v)'_v = r_1\mathbf{e}_1(v) + r_2\mathbf{e}_2(v) + r_3\mathbf{e}_3(v) \tag{3.6}$$

In $\mathbf{F}(v)$, $\mathbf{e}_1(v) = \mathbf{U}^0$. The derivative of $\mathbf{U}^0$ can be written in terms of $\mathbf{e}_2$ and $\mathbf{e}_3$:

64

$$\mathbf{U}^{0\prime} = \mathbf{e}'_1 = \omega_{12}\mathbf{e}_2 + \omega_{13}\mathbf{e}_3 \tag{3.7}$$

The derivative of $\mathbf{U}^0$, represented in the local frame field $\mathbf{F}(v)$, is used to simplify the calculation of the principal edge curvature of a crest spoke [Damon (2008)]:

$$r\kappa_E = -\frac{\omega_{13}}{r_3} \tag{3.8}$$

where r is the bisector spoke length.

The resulting κ_E from this simplified calculation is used to check the edge condition: $r\kappa_E < 1 \; iff$ the crest region is locally legal at a crest spoke. This condition is checked against sample spokes on each partial ellipse curve. If this legality condition fails to hold at any sample, it can be resolved by changing the η, i.e., by changing the bisector spoke length of an interpolated end atom to maintain the local legality of the implied crest.

The overall atom interpolation method includes the interpolation on internal atoms, the generation of interpolated end atoms based on the interpolated internal atoms, and the generation of a swept crest. An implied boundary, as well as an interpolated spoke field, is generated by this atom interpolation method. In the following section, the atom interpolation method is applied to synthetic and real-world m-rep models, and the results are shown. A careful evaluation of the atom interpolation method will be carried out in section 6.1.1.

3.3   Results

The application of the atom interpolation to a single figure m-rep yields a double-sided spoke field on a smoothly interpolated medial sheet. Both synthetic warped ellipsoid m-reps and real-world kidney m-reps are used to test the atom interpolation method.

The synthetic ellipsoids are generated by applying three statistically independent

global deformations to medial atoms in an m-rep, namely bending, twisting, and tapering. Each global deformation is controlled by a variable sampled from a normal distribution. A detailed description of the synthetic ellipsoids is included in section 6.1.1. There are also 39 kidney m-reps segmented from real-patient images as the real-world test objects.

The interpolation results for both the synthetic warped ellipsoids and kidneys are shown in figure 3.3: the underlying interpolated spokes and their implied boundary meshes. In each interpolated m-rep, the red and green boundary meshes are generated from the interpolated internal atoms, by connecting their interpolated internal spoke ends at the same side of the medial sheet, quad-patch by quad-patch; the yellow boundary mesh for the crest is generated from the sweeping planar curves fitted to interpolated end atoms, by connecting corresponding points on adjacent planar curves. Recall that each interpolated end atom has a pair of regular internal spokes as a part of a special internal atom. The yellow mesh for the crest is connected to the red and green meshes at the ends of both the regular spokes in interpolated end atoms. The red, green, and yellow meshes form an overall implied boundary mesh of each interpolated m-rep.

The legality of the interpolated spokes are maintained. As a result, the overall implied boundary mesh of each interpolated m-rep is free of any local shape defects, such as creasing and folding, as shown in figure 3.3. The smoothness of an implied boundary mesh depends on the continuity of the interpolated spokes, which is C^0 in this case. Further improvements may be desirable to achieve better smoothness of the implied boundary meshes.

3.4   Discussion

There are open problems in the proposed internal and end atom interpolations. This section lists these problems as follows, and the potential solutions will be discussed in
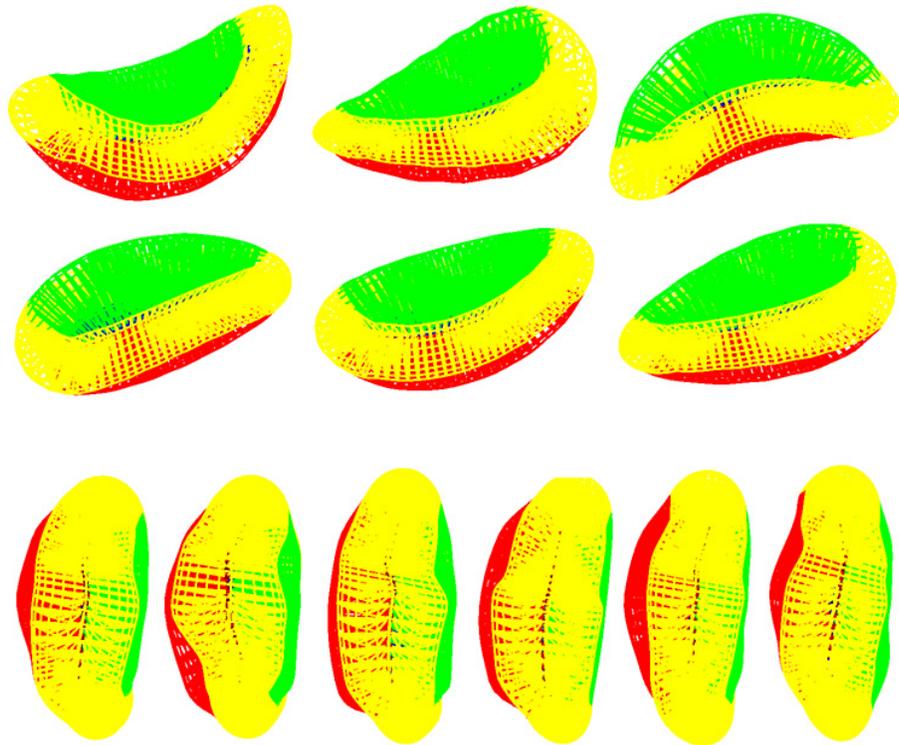
Figure 3.3: Top: interpolation results on synthetic ellipsoid m-reps. Bottom: interpolation results on real-patient kidney m-reps.

section 7.2.1.

- Speed of the internal atom interpolation: each interpolation is calculated by an integration, the speed of which can be improved.

- Continuity: currently C^0 continuity is enforced across interpolated patches of internal atoms. Higher order continuity might be desirable.

- Medial umbilic points: a medial umbilic point having repeated principal radial curvatures could be problematic for the $r\mathbf{S}_{rad}$ interpolation and the spoke interpolation.

- A numerical error in computing the $r\mathbf{S}_{rad}$ matrix can lead to a matrix with complex-valued eigenvalues and eigenvectors. The method to interpolate $r\mathbf{S}_{rad}$ matrices by interpolating their real eigenvalues and real eigenvectors needs to be changed to accommodate an $r\mathbf{S}_{rad}$ matrix of complex eigenvalues and complex eigenvectors.

Potential solutions to these open problems will be discussed in chapter 7, section 7.2.1.

This chapter has proposed a new interpolation method on a single figure m-rep. The legality of interpolated internal atoms is implied by the interpolation, and the legality of the boundary crest implied by interpolated end atoms is checked. Results have shown that the atom interpolation method applies well to both synthetic and real-world single figure objects.

The legality of an interpolated m-rep, including its implied boundary, is important for an m-rep to represent anatomical objects. Local shape defects, such as creasing and folding, rarely exist in real-world anatomical objects. These shape defects must be avoided in a shape representation to model a population of such anatomical object. Therefore, the proposed atom interpolation method is an important component of this dissertation, with modeling a population of anatomical objects as a goal.

This chapter has focused on the interpolation of single figure m-reps. The next chapter describes multi-figure m-reps, representing complex objects of more than one part. In order to handle the transition, i.e., the blend, between a pair of connected object parts in a multi-figure m-rep, an interpolation on a one-sided spoke field is used. This one-sided spoke interpolation is directly derived from the interpolation on internal atoms described in this chapter.

Chapter 4

Representing Multi-figure Anatomical Objects

This chapter focuses on representing anatomical objects of multiple named parts. For example, a prostate has two seminal vesicles attached to it, a liver has left and right lobes, and a kidney has a renal pelvis region. Figure 4.1-a and -b show a prostate with two attached seminal vesicles and a liver of two lobes, respectively, both as objects having one or two additive parts (protrusions); figure 4.1-c shows a kidney with its renal pelvis as an object having a subtractive part (indentation).

Due to the inherent complexity of these complex objects, previous shape representations of them concentrated on their global structure [Cootes et al. (1995); Gerig et al. (2001)] or on the extremely local behavior of the geometric primitives, such as points, without reference to the parts' inter-relations [Csernansky et al. (1998); Styner et al. (2003)]. The challenge of representing objects of multiple parts lies in the fact that modeling each individual part alone is not sufficient. The inter-relations among object parts have to be modeled, as well, to capture both global and local shape variations in these complex objects. Such inter-relations include a host part and subpart connection, global deformations of all connected parts, and the implied deformation from one part to the part(s) connected to it. The word "deformation" has two meanings: a change of shape (form) or a deffeomorphism, such as a warping. In this chapter, the deformation
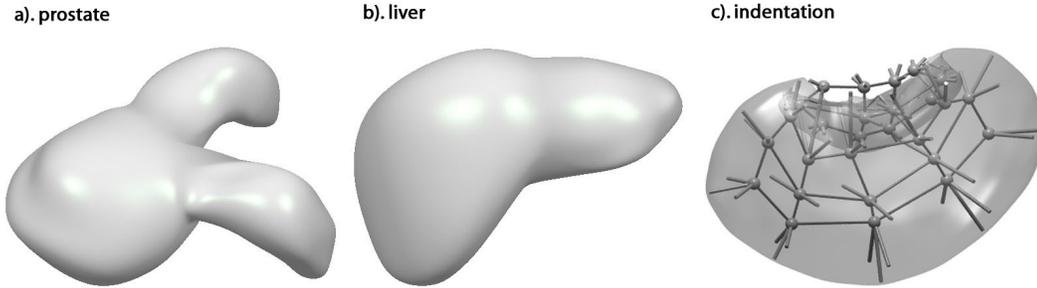
Figure 4.1: **a**) A prostate with two seminal vesicle protrusions. **b**) A liver represented by the union of the left and right lobes. **c**) A kidney with the renal pelvis as an indentation subfigure. Object a has three single-figure parts while objects b and c have two such parts.

of an m-rep means the former, i.e., the change of the m-rep shape.

The multi-figure m-rep to be described in this chapter overcomes this challenge of representing a complex object by treating the object as a set of connected parts and by representing each object part with a separate single figure m-rep. The set of m-rep figures is connected hierarchically, which allows an complex objects to be represented at multiple levels of scale. A multi-figure m-rep allows us to study the object as a whole, to study the inter-relations among object parts, and to study each individual part. The main components of a multi-figure m-rep are listed as follows:

- The hinge geometry to connect adjacent figures;

- The multi-scale deformations to deform a multi-figure m-rep of connected figures;

- The smooth transition, i.e., the blend region, between each pair of connected figures;

- The multi-scale statistics on multi-figure m-reps.

The rest of this chapter covers all these components. It starts with a review of a single figure m-rep and a proposal of a spherical parameterization, on which a multi-figure m-rep is based. It then describes the hinge geometry defined to connect a pair

of adjacent m-rep figures in section 4.2. Based on the hinge geometry a set of multi-scale deformations are also defined in section 4.2 to deform a multi-figure m-rep in consecutively smaller scale levels. A blend is described to smoothly transit between a pair of connected figures in section 4.2.3. A multi-scale way to conduct statistical analysis on multi-figure m-reps is described in section 4.3. Section 4.4 shows statistics resulting from the application of the hierarchical statistical analysis, and section 4.5 concludes this chapter with discussions.

4.1 Single Figure M-reps and a Spherical Parameterization

Each single figure m-rep is a quad-mesh of both internal and end medial atoms over a spatially regular lattice, as shown in figure 2.1-right. An internal atom and an end atom are both shown in figure 2.1.

Given a single figure m-rep, a smooth boundary can be generated both by a method based on subdivision, proposed in Thall's dissertation Thall (2004), and by a method based on atom interpolation, proposed in chapter 3. The method, based on atom interpolation, also provides a spoke field for the volumetric object interior. Both Thall's method and the proposed method provide a parameterized object boundary for a single figure m-rep. A previously proposed parameterization of an m-rep implied boundary is described next, followed by a newly proposed spherical parametrization.

A parameterization of an m-rep implied boundary is described in Han et al. (2004). In this parameterization, an m-rep implied boundary is divided into three patches: two internal patches implied by internal atom spokes, one on each of the two sides of a medial sheet, and a crest patch implied by end atom spokes. An m-rep implied boundary is then parameterized by three parameters (u_1, u_2, ϕ). Each one of the two internal patches is parameterized by (u_1, u_2), with ϕ indicating on which side of the medial sheet the
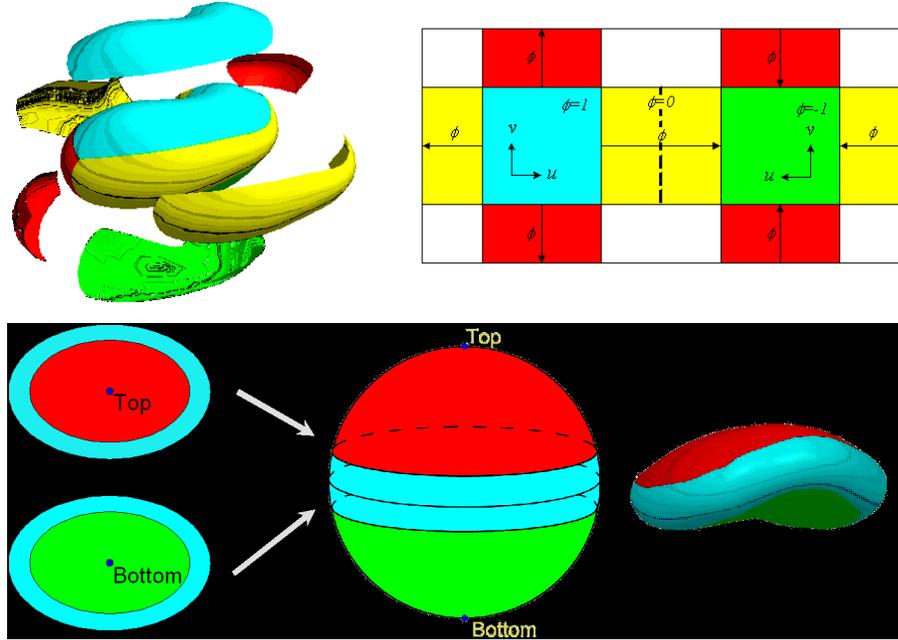
Figure 4.2: From top to bottom: **a**). the (u, v, φ) parameterization of a single figure m-rep with discontinuities in its parametric space; **b**). the two ellipses, one for each side of the medial sheet, mapped into a spherical coordinate system, with the two collar regions corresponding to the crest region of the object.

patch stays: $\phi = +1$ or $\phi = -1$ indicates the top or bottom side of the medial sheet, respectively. The crest patch is parameterized by $(u_1$ or $u_2)$ and ϕ, where either u_1 or u_2 is used depending on the position of a point on the crest patch, relative to the two internal patches, and where ϕ changes smoothly from the top side $(\phi = +1)$ to the bottom side $(\phi = -1)$ of the medial sheet. ϕ is zero at the exact boundary crest line, formed by bisector spoke ends of end atoms.

The space of (u_1, u_2, ϕ) parameterizing an m-rep can be unfolded to form a wrapped-around parametric space, as shown in figure 4.2-a. Such a parametric space has the disadvantage of being discontinuous at its boundary. It requires special care to traverse across the boundary of this parametric space.

A new spherical parametrization is proposed for an m-rep boundary. In this parameterization, a single figure m-rep implied boundary is divided into two parts by its boundary crest line: a top part and a bottom part, corresponding to the top side and

the bottom side of an m-rep medial sheet, respectively. The top part of the boundary is further divided into two regions: a center region implied by internal atom spokes and a half crest region implied by end atom spokes. The half crest region is collar-shaped, surrounding the center region. The top part of the boundary, including the center and the half crest regions, is mapped to an ellipse with a collar corresponding to the half crest region. The same mapping is applied to the bottom part of the boundary. Both of the ellipses with collars, as shown in figure 4.2-b, are then symmetrically mapped to two hemispheres of a unit sphere, with a spherical coordinate (θ, ϕ), as shown in figure 4.2-b. This spherical parameterization is a natural choice, given the spherical topology of the implied boundary of a single figure m-rep.

This spherical parameterization of each m-rep figure provides a basis for connecting a subfigure to its host figure in a multi-figure m-rep. A subfigure is fully known in the figural spherical coordinates of its host figure boundary at which the subfigure is attached to its host. This spherical parameterization can be applied to an m-rep boundary generated both by the subdivision based method and by the atom interpolation.

Based on a single figure m-rep and its parameterization, section 4.2 describes hinge geometry and multi-scale deformations of a multi-figure m-rep.

4.2   Multi-figure M-reps

An multi-figure m-rep consists of a set of m-rep figures organized as a hierarchy of connected figure pairs. Each connected pair contains a host figure and a subfigure. The entire set of m-rep figures are then recursively connected to form a desired object hierarchy. Such a hierarchy allows a representation of a subset of complex objects that encompasses many anatomical objects. Most of these objects can be adequately represented by two to three m-rep figures in two to three hierarchical levels, assuming that each connection is defined via a host figure and a subfigure pair. With such a
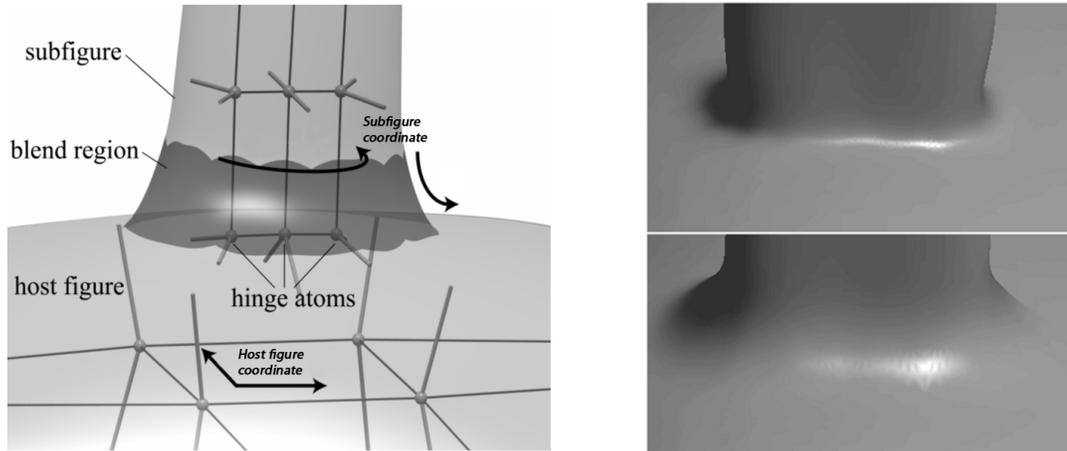
Figure 4.3: **Left** The host figure and subfigure arrangement, with the subfigure (six medial atoms appearing) on top, the host figure (four medial atoms showing) on bottom, and the blend region shown darker. **Right** Different shapes of the blend region.

hierarchy, neither a 6-junction connection [Damon (2003)] nor subfigures connected to host figures along multiple curves are supported by the multi-figure m-rep, which will be further discussed in section 7.2.1. The relation between a host figure and subfigure are determined by the anatomical structure of a target object and by the tightness of posterior probabilities of object parts.

The rest of this section describes how each pair of adjacent figures are connected by hinge geometry, how a multi-figure m-rep is deformed, and how a smooth boundary is generated for a multi-figure m-rep by a method called *blending*. A two-figure m-rep is used as an example in the following definitions and descriptions. A two-figure m-rep has one host figure and one subfigure, e.g., a liver with its right lobe as a host figure and its left lobe as a subfigure. All the definitions and descriptions of a two-figure m-rep can be recursively applied to a multi-figure m-rep. Note that here an m-rep "deformation" means the change in an m-rep shape.

4.2.1 Hinge Geometry

Hinge geometry defines how a subfigure is connected to its host figure. In a two-figure m-rep $\mathbf{M}$, assume its subfigure $\mathbf{F}_{sub}$ has a quad-mesh grid of $n_{row} \times n_{column}$ atoms, and assume these atoms are $\mathbf{A}_{i,j}, i = 1...n_{row}, j = 1...n_{column}$. In $\mathbf{M}$ its subfigure $\mathbf{F}_{sub}$ is attached to its host figure by a 1D curve of *hinge atoms*, which, when sampled, forms a row of atoms $\mathbf{A}_{1,j}, j = 1...n_{column}$ in the subfigure atom mesh. Generally, hinge atoms form a curve of atoms. However, in the multi-figure m-rep I describe hinge atoms are defined as a row of atoms. Each hinge atom $\mathbf{A}_{1,j}$ rides on the medially implied boundary of the host figure, with a known spherical figural coordinate (θ_j, ϕ_j) of the host figure boundary. At each (θ_j, ϕ_j), there is a boundary point $\mathbf{p}_{host,j}$ and a boundary surface normal $\mathbf{n}_{host,j}$ of the host figure. The points $\{\mathbf{p}_{host,j}\}$ and the corresponding normals $\{\mathbf{n}_{host,j}\}$ are key to connecting a subfigure to its host figure because they allow the method to reflect the relative position and orientation of a subfigure to its host figure.

The hinge geometry to represent a branching medial structure is a relaxation from the exact Blum condition to avoid instability against boundary noise of the low-volume portion of a branching medial axis, summarized in section 2.2.5. A host figure and subfigure pair connected by hinge geometry is demonstrated in figure 4.3-left. An m-rep allows both additive and subtractive subfigures, as shown in figure 4.1-a-c. Hinge geometry, defined as a connection via hinge atoms between a host figure and a subfigure, can be recursively applied to a multi-figure m-rep of more than two figures.

4.2.2 Multi-scale Deformations of a Multi-figure M-rep

Capturing the natural hierarchy in a complex anatomical object, a multi-figure m-rep deforms in a hierarchical way, consisting of multiple scale levels. At the subfigure scale level, hinge geometry is essential to how a multi-figure m-rep deforms.

Consider a two-figure m-rep of one host figure and one subfigure. At a global level, the m-rep deforms as a whole, i.e., as a collection of medial atoms from both of its figures,

and after a global deformation is applied to the m-rep, the subfigure hinge atoms might be re-projected back to the boundary of the host figure to maintain the hinge geometry connection; at a host figure level, the host figure deforms, and its deformation must be propagated to the subfigure via the hinge atoms connecting the two figures; at a subfigure level, the subfigure deforms on its own without affecting its host figure, while staying connected to its host figure, and the subfigure deformation originating from the hinge atoms is propagated to the rest of the subfigure. The global deformation of a two-figure m-rep and the figural deformation of a host figure are carried out in the same way as the deformation of a single figure m-rep, reviewed in section 2.4.1. A subfigure, however, goes through two types of deformation: an implied deformation by its host figure and a spontaneous deformation without affecting its host. Both of the subfigure deformations are realized via hinge atoms.

Before detailing the two subfigure deformations based on hinge geometry, I will describe *deformation propagation* in a subfigure, from its hinge atoms to the rest of the subfigure.

Deformation Propagation in a Subfigure

Deformation propagation in a subfigure includes the deformation to hinge atoms, either implied by a host figure deformation or implied by a spontaneous subfigure deformation, and the propagation of the hinge atom deformations to the rest of a subfigure.

When a host figure deforms, its boundary changes; when a subfigure deforms spontaneously, its hinge atoms transform on the implied boundary of the host figure. In both of these cases, the points $\mathbf{p}_{host,j}$ and the normals $\mathbf{n}_{host,j}$ at these points, connecting the host figure and the subfigure, change. These changes must affect the hinge atoms and then the rest of the subfigure.

Assume $\mathbf{p}_{host,j}$ and $\mathbf{n}_{host,j}$ change to $\mathbf{p}'_{host,j}$ and $\mathbf{n}'_{host,j}$, respectively. There is a composed transformation $g_j \in \mathbb{R}^3 \times \mathcal{SO}(3)$ of a translation and a rotation implied by

77

the changes in $\mathbf{p}'_{host,j}$ and $\mathbf{n}'_{host,j}$ for each hinge atom such that $g_j \circ (\mathbf{p}_{host,j}, \mathbf{n}_{host,j}) = (\mathbf{p}'_{host,j}, \mathbf{n}'_{host,j})$, where $j \in [1, n_{column}]$. These composed transformations $g_j, j = 1...n_{column}$ are applied to the corresponding hinge atoms one by one, such that $\mathbf{A}'_{1,j} = g_j \circ \mathbf{A}_{1,j}$.

The changes applied to the hinge atoms are then propagated to the rest of the subfigure. This propagation is carried out row by row: the changes to the hinge atoms in row $i = 1$ are propagated to the atoms in row $i = 2$, and the consequential changes of the atoms in row $i, n_{row} - 1 \geq i \geq 2$ are propagated to the atoms in row $i+1$. It suffices to show the propagation from the hinge atoms to the second row of atoms because the rest of the propagation follows. The process of this propagation includes representing the relation between the two rows of atoms by atom residues, and applying the atom residues to the deformed hinge atoms to acquire the deformed second row of atoms. By this means, we propagate the changes in the hinge atoms to the second row of atoms, while maintaining the relation between these two rows of atoms.

Recall that atom residues were defined in 2.4. Equation 2.26 suggests the atom residues between the second row of atoms $\mathbf{A}_{2,j}$ and the hinge atoms $\mathbf{A}_{1,j}$ can be calculated as $\Delta\mathbf{A}_{1,j} = \mathbf{A}_{2,j} \ominus \mathbf{A}_{1,j} = g^{-1}_{\mathbf{A}_{1,j}}(\in G(1)) \circ \mathbf{A}_{2,j}(\in \mathcal{M}(1))$, where $G(1)$ denotes the group of composed transformations of hub translation, spoke magnification(s), and spoke rotations that acts smoothly on each medial atom as a point on the Riemannian manifold $\mathcal{M}(1)$, where $g_{\mathbf{A}}$ denotes one such composed transformation that transforms a pre-fixed point as the origin on $\mathcal{M}(1)$ to $\mathbf{A} \in \mathcal{M}(1)$, and where $g^{-1}_{\mathbf{A}}$ is the inverse transformation of g that transforms $\mathbf{A}$ to the origin on $\mathcal{M}(1)$.

Given the atom residues $\Delta\mathbf{A}_{1,j}$ representing the relation between the second row of atoms and the hinge atoms, and given the deformed hinge atoms $\mathbf{A}'_{1,j}$, the deformed second row of atoms are calculated as $\mathbf{A}'_{2,j} = \mathbf{A}'_{1,j} \oplus \Delta\mathbf{A}_{1,j}$, where $j = 1...n_{column}$, and where "$\oplus$" is the inverse operation of "$\ominus$".

The same propagation is sequentially applied to consecutive rows of subfigure atoms: from row $i, n_{column} - 1 \geq i \geq 2$ to row $i + 1$ and one row at a time. This concludes

the deformation propagation. The two types of subfigure deformation, based on this propagation, are described next.

Host Figure Implied Subfigure Deformation

As the host figure deforms, its surface points and normals change, at the host figural coordinates $\{(\theta_j, \phi_j)\}$. These changes are applied to the hinge atoms and then the rest of the subfigure, by the deformation propagation just described. The advantage of the host figure implied subfigure deformation is that the relation between the host figure and the subfigure is maintained when the host figure deforms.

Spontaneous Subfigure Deformations

The subfigure can also translate, rotate, hinge, scale, and elongate on the host figure boundary while the host stays put. These basic hinge-relative subfigure deformations start in the figural coordinate of the host figure, described in the list below, with its first item detailed and the rest briefed:

- Translation or rotation: implemented by translating or rotating the hinge atoms in the host figure spherical coordinate and by propagating the deformations of the hinge atoms to the rest of the subfigure. In detail, there are several steps:

  1. Each hinge atom $\mathbf{A}_{1,j}, j \in [1, n_{column}]$ corresponds to a point (θ_j, ϕ_j) in the host figural coordinate, the space of which is a unit sphere, as described in section 4.1. All the corresponding points of the hinge atoms form an open curve on the unit sphere;

  2. The curve is translated or rotated on the unit sphere. The translation or rotation is implemented by a series of small steps. In each step, each point (θ_j, ϕ_j) on the curve changes to a new point along a geodesic curve of the sphere. After the translation or rotation, (θ_j, ϕ_j) changes to (θ'_j, ϕ'_j);

3. Accordingly, each boundary point $\mathbf{p}_{host,j}$ and its normal $\mathbf{n}_{host,j}$, at which the subfigure is connected to is host figure, change to $\mathbf{p}'_{host,j}$ and $\mathbf{n}'_{host,j}$, respectively;

4. The changes in $\mathbf{p}'_{host,j}$ and $\mathbf{n}'_{host,j}$ are applied as transformations to the hinge atoms;

5. The deformations of the hinge atoms are then propagated to the rest of the subfigure, via deformation propagation.

As a result, where a subfigure is connected to its host and how the subfigure orients relatively to its host both change according to the changes in $\mathbf{p}'_{host,j}$ and $\mathbf{n}'_{host,j}$;

- Hinging: implemented by rotating the row of hinge atoms around the curve of the hinge atoms and by propagating the hinge atom deformations to the rest of the subfigure;

- Scaling: implemented by scaling the curve of the hinge atoms in the host figural coordinate, by scaling the radii of the hinge atoms, and by propagating the hinge atom deformations to the rest of the subfigure;

- Elongation: implemented by modifying the distance between the second row of atoms and hinge atoms, represented by atom residues $\Delta\mathbf{A}_{1,j}$, and by propagating the changed relation to the rest of the subfigure.

These spontaneous deformations of a subfigure combined with the host figure implied subfigure deformation conclude the definition of subfigure deformations in a two-figure m-rep.

As a summary, the multi-scale deformations to a two-figure m-rep include three levels: a level of a global deformation to the entire m-rep, a level of a host figure deformation with its implied deformation to the subfigure, and the level of spontaneous subfigure deformations.

Section 4.2.3 describes the generation of a blend region as a transition between a pair of connected figures. Both of the methods to generate an m-rep implied boundary, the Catmull-Clark subdivision based method and the atom interpolation based method, will be extended to a multi-figure m-rep.

4.2.3 Blending

In order to generate the boundary of a multi-figure m-rep, a smooth transition between each pair of connected figures is required. The process of generating such a smooth transition is called blending. The overall process of generating the boundary of a two-figure m-rep, including a blend region, has the following key steps:

1. Generate the figural boundaries of both the host figure and the subfigure;

2. Intersect the two boundaries to find an intersection curve;

3. Transform the intersection curve into one cut curve on each one of the two figures;

4. Cut both the host figure and subfigure via the cut curves;

5. Create the smooth transition, i.e., the blend between the cut host figure and the cut subfigure;

6. Combine the boundaries of both the cut figures and the boundary of the blend region into an overall boundary of the object.

This subsection introduces two methods of blending. The first method is an extension to Thall's adapted Catmull-Clark subdivision method. The second method is an extension to atom interpolation, introduced in chapter 3. Given the efficiency and simplicity of the Catmull-Clark subdivision scheme, the first method is computationally faster. However, it does not provide an underlying medial structure for the blend region, and it does not guarantee the calculated boundary to be free of local shape defects.

Unlike the first method, the second method generates a medial structure for the blend region: between each pair of a host and subfigure, the medial sheet of the subfigure splits and merges into the cut version of the medial sheet of the host. Importantly, the interpolated spokes from the second method allow us to apply the available mathematics, especially those reviewed in section 2.3 and in section 3.3, to the blend region. Again without loss of generality, a two-figure m-rep will be used as an example in the rest of this subsection.

Blending Based on Catmull-Clark Subdivision

The idea is to build an overall control mesh for a two-figure m-rep, including the blend region, and to apply Thall's adapted Catmull-Clark subdivision scheme to the overall control mesh.

The original subdivision control meshes of both the host figure and subfigure, implied by their m-rep atoms, are used as starting meshes. The two starting meshes meet and merge into each other: designated sections from both figures' initial control meshes are removed, and the rest of both control meshes are re-meshed, as shown in figure 4.3-left. Two parameters d_{host} and d_{sub}, delimiting the top and the bottom of the blend region, control the shape of the blend, as shown in figure 4.3-right. The detailed algorithm is shown as follows:

1. Intersect two initial control meshes of the host figure and subfigure to form an intersection curve $\mathbf{C}_{int}$, which is mapped to the figural spherical coordinates of both figures;

2. Generate a cut curve for each figure:

   - Generate a cut curve for the host figure:

     - Map the intersection curve $\mathbf{C}_{int}$ to the spherical parametric space of the host figure;

- Iteratively dilate the mapped curve by a smooth curvature flow in the spherical parametric space. The amount of dilation is controlled by a parameter d_{host};

- Map the dilated curve from the spherical parametric space of the host figure back to its initial control mesh as a cut curve $\mathbf{C}_{cut}^{host}$;

- Generate a cut curve for the subfigure:

  - Map the intersection curve $\mathbf{C}_{int}$ to the spherical parametric space of the subfigure;

  - Iteratively dilate the mapped curve by a smooth curvature flow in the spherical parametric space. The amount of dilation is controlled by a parameter d_{sub};

  - Map the dilated curve from the spherical parametric space of the subfigure back to its initial control mesh as a cut curve $\mathbf{C}_{cut}^{sub}$;

3. Re-sample and smooth the cut curves $\mathbf{C}_{cut}^{host}$ and $\mathbf{C}_{cut}^{sub}$ if necessary to maintain the number of points and length of each segment to stay in empirically determined ranges;

4. Remove a designated section inside the cut curve from each of the two initial control meshes;

5. Build point correspondences between the two cut curves $\mathbf{C}_{cut}^{host}$ and $\mathbf{C}_{cut}^{sub}$ by minimizing the sum of the squared distances between corresponding points on the two curves;

6. Mesh between the two cut curves via the built correspondences, to connect the cut control meshes of the host figure and subfigure; This results in an overall control mesh for the entire object;

7. Apply Thall's adapted Catmull-Clark subdivision to the overall control mesh of the object to generate a smooth boundary, including the blend region, for the two-figure m-rep.

A second blending method, based on the spoke interpolation directly derived from the atom interpolation, is introduced next.

Blending Based on Spoke Interpolation

The second blending method uses a one-sided spoke interpolation, which is directly derived from the atom interpolation method. The atom interpolation is first applied to both the host figure and the subfigure to obtain figural boundaries with smooth medial sheets and attached spokes (step 1 below). The two figural boundaries are intersected, and based on the intersection curve both of the figural boundaries and both of the figural medial sheets are cut (steps 2-7 below). A one-sided collar-shaped medial sheet is created to connect the two cut figural medial sheets (steps 8-10 below). Then, a blend region is calculated by a one-sided spoke interpolation (steps 11-13 below). Here again, there are two parameters that control the shape of the blend region. The overall algorithm is shown as follows:

1. Interpolate the m-reps of both the host figure and the subfigure by atom interpolation:

   - Each interpolated figure has a boundary mesh $\mathbf{B}^{host}$ or $\mathbf{B}^{sub}$, parameterized by a spherical coordinate;

   - Each interpolated figure has an interpolated medial sheet $\mathbf{p}^{host}$ or $\mathbf{p}^{sub}$ and the implied normal $\mathbf{n}^{host}$ or $\mathbf{n}^{sub}$ of the sheet;

   - Each interpolated figure has a spoke field $\mathbf{S}^{host}$ or $\mathbf{S}^{sub}$;

2. Intersect the two boundary meshes $\mathbf{B}^{host}$ and $\mathbf{B}^{sub}$ to form an intersection curve $\mathbf{C}_{int}$, which is mapped to the figural coordinates of both figures;

3. Generate a cut curve for each figure:

   - Generate a cut curve for the host figure:

     (a) Map the intersection curve $\mathbf{C}_{int}$ to the spherical parametric space of the host figure;

     (b) Iteratively dilate the mapped curve by a smooth curvature flow in the spherical parametric space. The amount of dilation is controlled by a parameter d_{host};

     (c) Map the dilated curve from the parametric space of the host figure back to its interpolated mesh as a cut curve $\mathbf{C}_{cut}^{host}$;

   - Generate a cut curve for the subfigure:

     (a) Map the intersection curve $\mathbf{C}_{int}$ to the spherical parametric space of the subfigure;

     (b) Iteratively dilate the mapped curve by a smooth curvature flow in the spherical parametric space. The amount of dilation is controlled by a parameter d_{sub};

     (c) Map the dilated curve from the parametric space of the subfigure back to its interpolated mesh as a cut curve $\mathbf{C}_{cut}^{sub}$;

4. Re-sample or smooth the cut curves $\mathbf{C}_{cut}^{host}$ and $\mathbf{C}_{cut}^{sub}$ if necessary, in order to maintain the number of points to be same for both curves, and to maintain the length of each segment on both curves to be in empirically determined ranges;

5. Remove designated sections from both boundary meshes inside the cut curves $\mathbf{C}_{cut}^{host}$ and $\mathbf{C}_{cut}^{sub}$ to form a cut version of the interpolated host figure boundary and a cut version of the interpolated subfigure boundary, respectively;

6. Convert each boundary cut curve to a cut curve $\mathbf{C}_{cut}^{host\_medial}$ or $\mathbf{C}_{cut}^{sub\_medial}$ on the interpolated medial sheet of the host figure or of the subfigure, respectively. This

85

conversion is done by the correspondences between the interpolated medial sheet and the m-rep implied boundary;

7. Remove designated sections from both interpolated medial sheets inside the cut curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$ to form a cut version of the interpolated medial sheet of the host and of the subfigure, respectively;

8. Build correspondences between points on the curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$ by minimizing the sum of the squared distances between corresponding points from both the curves;

9. Based on the built correspondences, mesh between these two curves by a quad-mesh of a series of quad-patches. This is made possible by the same number of points on both the medial sheet cut curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$;

10. Interpolate the medial sheet for the blend region between the two medial sheet cut curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$ by a point-and-normal interpolation:

    (a) Each point on one of the medial sheet cut curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$ has a point and a normal, given by the interpolated medial sheets of the host and of the subfigure, respectively;

    (b) A point-and-normal interpolation is applied to the quad-mesh between the two medial cut curves, quad-patch by quad-patch, to generate a smooth transitional medial sheet for the blend region.

    This point-and-normal interpolation by cubic Hermite patches was detailed in section 3.2.2.

11. Interpolate the spokes attached to the two medial sheet cut curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$ by a one-sided spoke interpolation:

86

(a) Each point on the medial sheet cut curve $\mathbf{C}_{cut}^{host\_medial}$ or $\mathbf{C}_{cut}^{sub\_medial}$ has an attached spoke, given by the interpolated spoke field $\mathbf{S}^{host}$ or $\mathbf{S}^{sub}$, respectively;

(b) In each quad-patch of the quad-mesh between the two medial sheet cut curves, estimate the $r\mathbf{S}_{rad}$ matrix of each spoke at one of four corners. Apply the $r\mathbf{S}_{rad}$ interpolation to the four estimated $r\mathbf{S}_{rad}$ matrices, described in section 3.2.2;

(c) Based on the interpolated medial sheet and the interpolated $r\mathbf{S}_{rad}$ matrices, the spokes attached to the two curves $\mathbf{C}_{cut}^{host\_medial}$ and $\mathbf{C}_{cut}^{sub\_medial}$ are interpolated. This interpolation is realized by a one-sided spoke interpolation in each quad-patch, via an integration of the first-order derivatives of the spokes, calculated from the derivatives of the interpolated medial sheet and the interpolated $r\mathbf{S}_{rad}$ matrices. This spoke interpolation is derived directly from the two-sided spoke interpolation, detailed in section 3.2.2;

12. Connect the spoke ends of the interpolated spokes for the blend to form a smooth boundary for the blend;

13. Combine the interpolated boundary of the blend and the cut boundaries of both the figures into an overall boundary for the two-figure m-rep. The underlying medial sheet and spokes of the blend are also given by this method.

As shown in figure 4.4, in this spoke interpolation based method, the cut medial sheet of the subfigure splits into a one-sided collar-shaped surface. This split collar-shaped medial sheet then merges with C^1 continuity, almost everywhere, into the cut medial sheet of the host figure. The only exceptions of C^0 continuity exist at the medial points that are on the cut curve of the host figure medial sheet and also on the original edge curve of the host figure medial sheet. The resulting medial representation, as shown in figure 4.5 for the blend region, is a skeletal structure, with the medial sheet appearing as
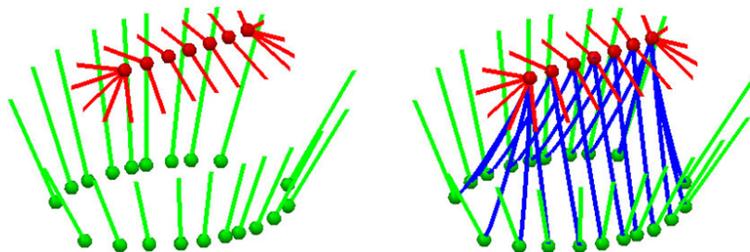
Figure 4.4: **Left**: the spokes attached to the two medial sheet cut curves. **Right**: the blue curves as a part of the one-sided collar-shaped medial sheet for the blend region.

a one-sided collar-shaped surface instead of a two-sided medial sheet in a Blum medial axis.

The mathematics, described in section 2.3 can be applied to a such a skeletal structure, of which a Blum medial axis is a special case. In spite of the one-sided skeletal structure for the blend, the legality of the blend region is implied by the spoke interpolation, based on one assumption that the estimated $r\mathbf{S}_{rad}$ matrices of the attached spokes on both medial sheet cut curves are legal. This might not always hold for an arbitrary multi-figure m-rep. However, if we start with a multi-figure m-rep that does uphold this assumption, we can maintain it by penalizing deformations on the multi-figure m-rep that might violate this assumption.

Both the subdivision based method and the atom interpolation based method for a single figure m-rep have been extended to generate a boundary of a two-figure m-rep. Furthermore, both the extended methods can be applied recursively to a multi-figure m-rep of more than two figures, assuming that no two subfigures attached to the same host figure intersect each other, even in their blend regions. Section 4.3 describes hierarchical statistics of multi-figure m-reps.
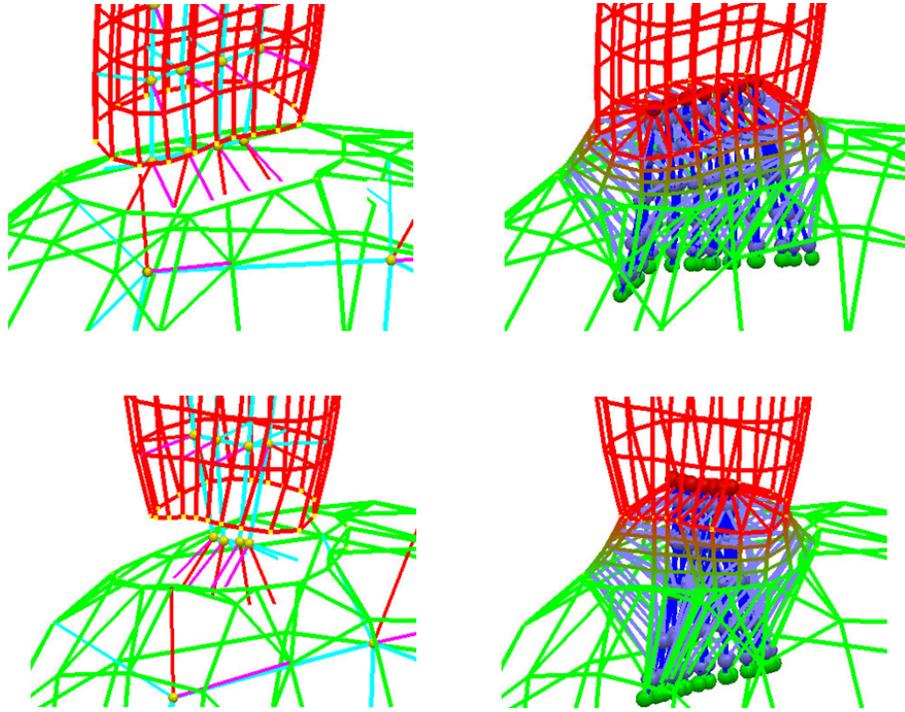
Figure 4.5: Demonstration of two interpolated multi-figure m-reps: each row of two figure components shows a two-figure m-rep with its subfigure connected to its host figure in a different orientation. In each of the two rows, the left figure component shows the cut boundary meshes of both figures, and the right figure component shows the interpolated medial sheet and medial spokes and the implied boundary mesh of the blend region. In both of the left figure components, the medial atoms (hubs in yellow and spokes in magenta and cyan) and the medial sheet meshes (in cyan) of the host figure and subfigure are also shown.

89

4.3 Hierarchical Statistics of Multi-figure Objects

The hierarchy in a multi-figure m-rep provides an extension of PGA statistics of single figure m-reps, reviewed in section 2.4, to hierarchical PGA statistics of multi-figure m-reps.

Hierarchical statistics of multi-figure m-reps are calculated in a similar way as that in the hierarchical statistical framework for a multi-object complex, detailed in [Pizer et al. (2005)]. Without loss of generality, a two-figure object is used as an example. The hierarchical statistics of two-figure m-reps include the following parts: the global statistics, the residue statistics at the host figural level, and the residue statistics at the subfigural level, to be described in next two subsections. Recall that PGA of single figure m-reps is applied to the union of all atoms in each m-rep. The global statistics of two-figure m-reps are calculated by a PGA on the union of all the atoms from both figures.

4.3.1 Global Statistics

Assume that each two-figure m-rep object $\mathbf{O}$ has two figures $\{\mathbf{F}_{host}, \mathbf{F}_{sub}\}$, of n_{host} and n_{sub} atoms, respectively. Let n_O be the total number of atoms in $\mathbf{O}$. The global statistics of such two-figure m-reps are computed as a Fréchet mean m-rep $\overline{\mathbf{O}}$ and PGA statistics on the m-reps, each as a union of all its n_O atoms.

4.3.2 Hierarchical Statistics Based on Residues

In each two-figure object $\mathbf{O}$, its host figure and subfigure are treated as the single figure objects in the case of a multi-object complex, and the hinge atoms, which relate the host figure's deformation to the sympathetic subfigure deformation, are treated as the *augmenting atoms* in the case of a multi-object complex, detailed in [Pizer et al. (2005)]. Let $n_O = n_{host} + n_{sub}$. Three definitions are briefed as follows:

- M-rep residue: difference between two m-reps, calculated by the operation $\ominus$ (equations 2.26 and 2.27 in section 2.4.3). Because the m-rep residue can also be represented as an m-rep, PGA can be directly applied to m-rep residues;

- Augmentation: $\mathbf{U}_1 = \mathbf{F}_{host} \cup \mathbf{A}_{hinge}$ denotes the augmented union of host figure atoms and subfigure hinge atoms $\mathbf{A}_{hinge}$ in $\mathbf{F}_{sub}$;

- Projection: an m-rep $\mathbf{M}$ can be projected to a PGA subspace by $\pi_H(\mathbf{M}) \approx \mathrm{Exp}_{\overline{\mathbf{M}}} \sum_{i=1}^{k} \langle \mathbf{p}_i, \mathrm{Log}_{\overline{\mathbf{M}}}(\mathcal{M}) \rangle \mathbf{p}_i$, as detailed in [Pizer et al. (2005)].

Based on these definitions, the three parts in the hierarchical statistics of two-figure m-reps are as follows:

1. PGA_{global}: global PGA statistics of all the n_O atoms making up the entire object. PGA_{global} describes the global shape variation of the original objects.

   This shape variation is then removed from both the host figure atoms and the subfigure atoms before steps 2 and 3, via the atom residues between the original atoms and the atoms projected to a global PGA subspace, spanned by the first a few principal geodesic modes in the global PGA;

2. PGA_{host}: PGA statistics of the residues of the union $\mathbf{U}_1$ of the host figure atoms augmented by the subfigure hinge atoms. PGA_{host} describes the remaining shape variation of the augmented $\mathbf{U}_1$ after the global shape variation has been removed in step 1.

   The changes in the hinge atoms, from the projection in step 1, are propagated to the rest of the subfigure, via the hinge geometry and deformation propagation, detailed in section 4.2.1.

   Then the shape variation in $\mathbf{U}_1$ is removed, via atom residues between atoms in $\mathbf{U}_1$ and the projection of $\mathbf{U}_1$ to a PGA subspace of PGA_{host};
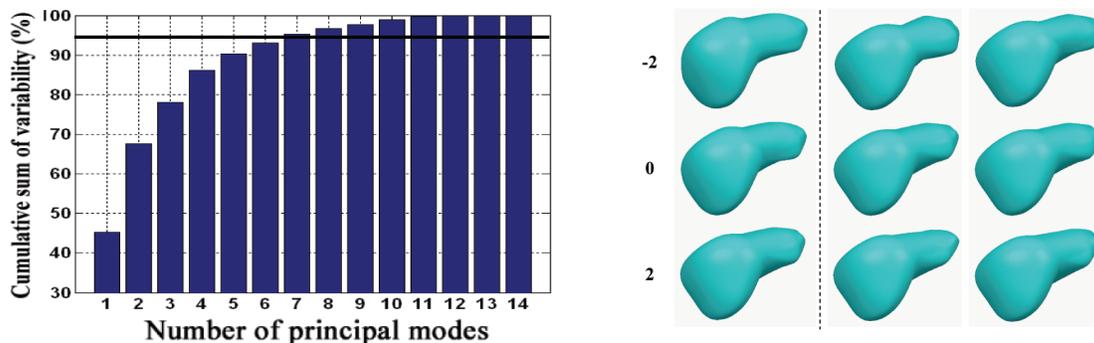
Figure 4.6: **Left**: accumulated sum of the variances from the global statistics PGA_g: the first **7** modes capture over **95%** of the total variability. **Right**: the residue shape variation after the global variation is removed: each column shows the liver -2 standard deviations from the residue mean along the respective eigenmode, the residue mean, and the liver $+2$ standard deviations from the mean. The left column shows the first principal mode of the host residue statistics PGA_h; the other two columns show the first two modes of the subfigure residue statistics: PGA_s describes the remaining shape variation of the subfigure after the global and host-implied variation have been removed.

3. PGA_{sub}: PGA statistics of the residues of subfigure $\mathbf{F}_2$ atoms after the changes in the hinge atoms, from the projection in step 2, are propagated to the rest of the subfigure.

The hierarchical statistics on two-figure m-reps can be recursively extended to multi-figure m-reps of more than two figures. The hierarchical statistics will be applied to liver m-reps next.

4.4 Results

The results from applying the hierarchical statistics to a set of two-figure liver m-reps are shown in figure 4.6. The hierarchical statistics of the two-figure livers include three levels: the global level, the host figure level, and the subfigure level. As shown in figure 4.6, the shape variation in the three levels gradually diminishes, which is consistent with the decreasing scale levels in the m-rep hierarchy.

92

4.5   Discussion

The ability of representing complex objects, while reflecting the inter-part relations in such objects, is a valuable asset of the m-rep to model anatomical objects in a population. A multi-figure m-rep captures the natural hierarchy within a complex object, so it allows learned shape statistics to follow the same anatomical hierarchy in such objects. Relative to the object geometry, learned appearance statistics of multi-figure m-reps can potentially follow this hierarchy as well.

A multi-figure m-rep also has the advantage that a subtractive part (indentation) or an additive part (protrusion) connects to its host part in a geometrically equivalent way. This uniform representation greatly simplifies the calculation of statistics of such objects, which makes it feasible to model complex shapes in a population having additive and subtractive parts.

At this moment, multi-figure m-rep statistics only handle objects of a fixed connection topology, i.e., a fixed number of figures and fixed connection relations among figures. Extending this framework to handle objects of varying connection topologies is potentially valuable, to be discussed in section 7.2.2.

Chapter 5

Geometrically Proper Models in Statistical Training

In a deformable model framework, a statistical training process includes two steps: the first step is to extract shape models from a set of training images, and the second step is to calculate shape statistics from the extracted shape models and to calculate appearance statistics relative to these shape models. This chapter focuses on the first step, also called a fitting step.

The fitting step is to deform a template model into a set of binary characteristic images for training, usually produced from a set of grey scale images by human experts. The lack of any shape prior statistics to guide the model fitting can lead to shape models with local shape illegalities. Local shape illegalities, such as surface folding or creasing as shown in figure 5.1, are defects in shape models, which do not exist in anatomical objects. These shape defects will taint learned shape and appearance statistics. Any use of these tainted statistics can yield undesirable results. In order to ensure the quality of the shape and appearance statistics, shape illegalities in extracted shape models must be avoided.

On the one hand, it is desirable to avoid shape illegalities, but on the other hand, it is not trivial to do so at feasible computational cost for existing shape representations. Existing solutions to avoiding shape illegalities in a fitting step either require manual
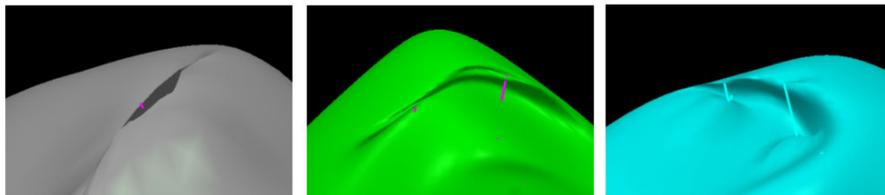
Figure 5.1: Illegal local shape defects: folding and creasing.

enforcements on the smoothness of extracted deformable models [Cootes et al. (1995)] or require expensive computation [Brechbuhler et al. (1995)]. Human interferences may bring artificial biases into extracted shape models and into learned statistics as well.

As reviewed in section 2.3, a particular strength of the m-rep as a medial shape representation is that powerful mathematics exist for medial geometry. This chapter shows how to adapt these mathematics to a geometric shape penalty, which can be explicitly applied to deformable m-reps in an automatic fitting step of reasonable computational cost. The legality of boundary surfaces and interiors of extracted m-rep models is maintained in this fitting step, also called a proper fitting step.

The rest of this chapter is organized in the following way. Section 5.1 details the proposed method for achieving medial legality and thus proper statistics in a training process using m-reps. Section 5.2 shows the generation of synthetic test models and the results from the application of the proper fitting to both synthetic and real world data. Section 5.3 concludes this chapter with discussions.

5.1 A Proper Binary Fitting Step

The goal of a fitting step is to find a best model $\mathbf{M}$ in a shape space for each training image $\mathbf{I}$. This goal can be written in the form of $\arg\min_{\mathbf{M}} F_{obj}(\mathbf{M}|\mathbf{I})$, where $F_{obj}(\mathbf{M}|\mathbf{I})$ is an objective function to be minimized to find the best shape parameters for $\mathbf{M}$. This objective function is typically the sum of two parts: a data match part measuring the quality of the fit of $\mathbf{M}$ to target data and a geometric penalty part penalizing any

inappropriate $\mathbf{M}$.

The target data in a deformable model fitting step are what a template deformable model is aimed to fit into. In our case, the data include a binary image and a small set of selected image landmarks. As a result, the data match part used in the proper fitting step has two terms: an image match term F_{img} and an image landmark match term F_{ilm}. The use of selected image landmarks as a part of the data not only improves the data fitting quality but also enforces correspondences between the model and identifiable anatomical features in the image.

Maintaining appropriate geometry of fitted models is a common goal of a deformable model fitting step. This goal is typically achieved by maintaining regular placements of shape features of a deformable model and by maintaining shape smoothness. In our case, the geometric penalty has two terms to penalize inappropriate geometry of a fitted m-rep: an irregularity penalty term F_{reg} and an illegality penalty term F_{leg}, where F_{leg} is a new geometric penalty term proposed by this dissertation. It is used to maintain the legality of a fitted m-rep and therefore improve shape smoothness of both the boundary and the interior of an m-rep.

When minimized, the data match part and the geometric penalty part, summed up as the objective function in our proper fitting, yield an m-rep that fits well to a target image and its corresponding image landmarks without any irregular atom arrangement and without illegal local shapes.

5.1.1   Data Match

Image Match

F_{img}, the image match term in our objective function, measures whether an m-rep implied boundary $\Omega(\mathbf{M})$ is in accordance with the voxel boundary $\mathbf{B(I)}$ in a binary image $\mathbf{I}$. An m-rep implied boundary, generated either by the subdivision based method or by the atom interpolation method, consists of a set of discrete surface points $\omega_i \in \Omega(\mathbf{M})$;

the voxel boundary of $\mathbf{I}$ consists of a set of discrete voxel points $\mathbf{b}_i \in \mathbf{B(I)}$. In a discrete form, F_{img} is defined as the sum of the squared distances $d^2(\omega_i, \mathbf{B(I)})$ from surface points $\omega_i \in \Omega(\mathbf{M})\, i \in [1, N]$ to their closest voxel points on $\mathbf{B(I)}$:

$$F_{img}(\mathbf{M}, \mathbf{I}) = \sum_{\omega_i \in \Omega(\mathbf{M})} d^2(\omega_i, \mathbf{B(I)}) \tag{5.1}$$

where $d^2(\omega_i, \mathbf{B(I)}) = \arg\min_{\mathbf{b} \in \mathbf{B(I)}} \|\mathbf{b} - \omega_i\|^2$.

In order to calculate F_{img} efficiently, a space filling lookup table is pre-calculated for a voxelized unit cube, into which the binary volumetric image and the model are normalized. The lookup table, in the form of a distance map, contains the distance from each voxel center in the unit cube to its closest voxel point $\mathbf{b} \in \mathbf{B(I)}$ on the voxel boundary $\mathbf{B(I)}$. The lookup table is created by Danielsson's distance map propagation from the original binary image [Danielsson (1980)], and the table is saved for later uses. A trilinear interpolation in the lookup table gives a fast calculation of $d^2(\omega_i, \mathbf{B(I)})$ and a fast calculation of F_{img}.

There are two special cases to be handled carefully when using the image match term F_{img} in our binary fitting:

- If at a boundary point ω_i, the surface normal $\mathbf{n}_i$ of Ω differs from the distance gradient of the distance map at that point by more than a certain empirically determined angle, $d^2(\omega_i, \mathbf{B(I)})$ will be replaced by the distance from ω_i along $\mathbf{n}_i$ to the nearest binary boundary voxel on $\mathbf{B}$, detailed in [Xiaoxiao Liu and Pizer (2008)];

- If the object in $\mathbf{I}$ is very thin, i.e., less than a voxel in thickness, $d^2(\omega_i, \mathbf{B(I)})$ is not accurate enough to guide the model $\mathbf{M}$ to deform into $\mathbf{I}$. A solution is to fit an initially dilated model, by scaling up the spoke length of m-rep atoms, into a dilated version of the binary boundary in the training image and then to contract the fitted model by the same amount as the dilation, via shrinking the

spoke length. The contracted m-rep is the fitted m-rep for the original image $\mathbf{I}$.

Landmark Match

F_{ilm}, the landmark match term in our objective function, measures how good explicit feature correspondences are between the m-rep implied boundary $\Omega(\mathbf{M})$ and the voxel boundary $\mathbf{B(I)}$. An expert identifies a few anatomically important and easily identifiable landmarks on each model boundary and in each binary image. Each image landmark $\mathbf{LI}_k$ identified in $\mathbf{I}$ is provided with a real tolerance value ϕ_k, which reciprocally weights the associated distance from the image landmark to its corresponding model landmark $\mathbf{LM}_k$, where $k = 1..N$, and N is the total number of image (model) landmarks.

A model landmark $\mathbf{LM}_k$ on an m-rep model is identified as a spoke end of an medial atom, and a corresponding image landmark in a binary image is identified as a point $\mathbf{LI}_k$ in the image volume. In our current implementation, a squared Euclidean distance from a model landmark to its corresponding image landmark is computed as the measure of fit. The squared distances are then summed up for all the model landmarks, weighted by tolerance factors $1/\phi_k^2$:

$$F_{ilm}(\mathbf{M}, \mathbf{I}) = \sum_{k=1}^{N} \frac{1}{\phi_k^2} d^2(\mathbf{LM}_k, \mathbf{LI}_k) \qquad (5.2)$$

5.1.2 Geometric Penalty

Besides the measurement of the dissimilarity between a model and image data, including the binary image and the corresponding image landmarks, by the data match, the geometric penalty part in our objective function penalizes any inappropriate $\mathbf{M}$ when being minimized.

The geometric penalty part penalizes any irregular arrangement of m-rep atoms in $\mathbf{M}$ via an irregularity penalty term F_{reg}, and it penalizes any illegality of the interior and boundary implied by $\mathbf{M}$ via an illegality penalty term F_{leg}.

Irregularity Penalty

The irregularity penalty term F_{reg} for an m-rep measures the deviation of every atom from the average of its neighboring atoms. This term penalizes non-uniform spacing of medial atoms and inconsistent changes in the spoke length and direction of medial atoms. It contributes to proper object geometry and to good correspondences among shape models across training cases, which are especially important for modeling anatomical objects in a population. For each medial atom $\mathbf{m}_i$, its irregularity penalty is calculated as the squared Riemannian distance, as described in section 2.4.1, between $\mathbf{m}_i$ and the Fréchet mean of its neighboring atoms $\mathbf{N}(\mathbf{m}_i)$, defined by equation (2.25), section 2.3. The penalties of m-rep atoms are then summed up to form the irregularity penalty for an m-rep:

$$F_{reg}(\mathbf{M}) = \sum_{i=1}^{M} dis^2(\mathbf{m}_i, FMean(\mathbf{N}(\mathbf{m}_i))) \tag{5.3}$$

where M is the total number of atoms in the m-rep.

Illegality Penalty

The illegality penalty term F_{leg} for an m-rep measures the likelihood of an m-rep to have a local shape defect. The calculation of this penalty term depends on the calculation of $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices, for internal atom spokes and edge atom spokes, respectively. The legality conditions based on these two matrices, detailed in section 2.3, are converted into explicit penalty functions. The calculations of these two matrices are summarized first, followed by the definition of a penalty function and the definition of the illegality penalty term.

According to equations (2.13) and (2.19) in section 2.3, the $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices are calculated from the derivatives of the medial sheet $\mathbf{p}$ and spokes $\mathbf{S}$. Briefly, the derivatives of the medial sheet $\mathbf{p}$ and the regular spokes $\mathbf{S}^{+1,-1}$ are calculated by the
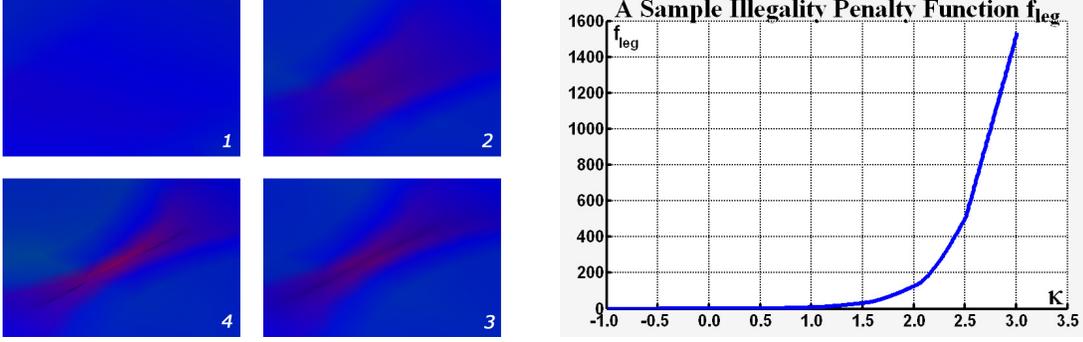
Figure 5.2: **Left**: gradual formation of a self-intersection on a surface portion, rendered by the maximal λ_{ri} of corresponding spokes to surface points. Blue or red means legal($\lambda_{ri} \leq 0$) or illegal($\lambda_{ri} \geq 1$), respectively, and any intermediate color shows the tendency for corresponding surface points to become illegal; **Right**: a sample illegality function f_{leg} as a cubic Hermite curve.

finite differences between adjacent internal atom, and an $r\mathbf{S}_{rad}$ matrix is estimated from these derivatives for each internal spoke by (2.13); then the derivatives of the medial sheet $\mathbf{p}$ and the bisector spokes $\mathbf{S}^0$ along the edge curve of the medial sheet are calculated by the finite differences between adjacent end atoms, and an $r\mathbf{S}_E$ matrix is estimated for each end bisector spoke by (2.19).

An eigen-decomposition is applied to each estimated $r\mathbf{S}_{rad}$ to calculate both its eigenvalues $\lambda_{ri}, i = 1, 2$ for each regular spoke $\mathbf{S}^{+1}$ or $\mathbf{S}^{-1}$. λ_E of each estimated $r\mathbf{S}_E$ is calculated for each end spoke $\mathbf{S}^0$ as well. Figure 5.2-Left shows a visualization of λ_{ri} on a portion of an m-rep implied boundary by internal m-rep atoms. λ_{ri} and λ_E both serve as distinct indicators of local shape illegalities. The following defines a penalty function on λ_{ri}'s and λ_E's.

Assume that the two eigenvalues $\{\lambda_{ri}, i = 1, 2\}$ of an estimated $r\mathbf{S}_{rad}$ matrix are always ordered as $\lambda_{r1} > \lambda_{r2}$. The local shape legality holds iff $\lambda_{r1} < 1$ and $\lambda_E < 1$, detailed in section 2.3. Theoretically, this legality condition is binary because the signs of $\lambda_{r1} - 1$ and $\lambda_E - 1$ determine the legality of the implied boundary, as well as the object interior. However, smaller λ_{r1} or λ_E desirably leads to a smoother surface in practice. In order to incorporate the legality condition into the illegality penalty term

F_{leg}, a smooth function $f_{leg}(\lambda_{r1}$ or $\lambda_E)$ is defined, which has the following properties:

1. Monotonically increasing;

2. Increasingly penalizing when λ_{r1} or λ_E approaches and passes 1.

f_{leg} is set to be smooth at $\lambda_{r1}, \lambda_E \in [1, \infty)$ in order to simplify the gradient-based optimization on the overall objective function $F_{obj}(\mathbf{M}|\mathbf{I})$. Cubic Hermite curves are used to form the penalty function f_{leg}, providing the luxury of freely picking the desired function slopes at certain control points of the argument of f_{leg}, which in our case are 0 and 1. A cubic Hermite curve is C^1 continuous everywhere, and it can be set to have both the properties above. A sample f_{leg} is shown in figure 5.2-Right. The penalties of λ_{r1}'s of all internal spokes and λ_E's of all edge spokes in an m-rep $\mathbf{M}$ are summed up as one explicit illegality penalty term F_{leg} in the objective function $F_{obj}(\mathbf{M}|\mathbf{I})$:

$$F_{leg}(\mathbf{M}) = \sum_{\mathbf{S}^{+1} \in \mathbf{M}} f_{leg}(\lambda_{r1}(\mathbf{S}^{+1})) + \sum_{\mathbf{S}^{-1} \in \mathbf{M}} f_{leg}(\lambda_{r1}(\mathbf{S}^{-1})) + \sum_{\mathbf{S}^0 \in \mathbf{M}} f_{leg}(\lambda_E(\mathbf{S}^0)) \qquad (5.4)$$

5.1.3 Objective Function

The overall objective function for our binary fitting step using the illegality penalty is defined as follows:

$$F_{obj}(\mathbf{M}|\mathbf{I}) = \alpha F_{img}(\mathbf{M}, \mathbf{I}) + \beta F_{ilm}(\mathbf{M}, \mathbf{I}) + \gamma F_{reg}(\mathbf{M}) + \delta F_{leg}(\mathbf{M}) \qquad (5.5)$$

where $\alpha, \beta, \gamma, \delta$ are parameters empirically determined to control the relative weights of all the four terms.

5.1.4   A Proper Statistical Training Process Based on the Proper Fitting

The proper fitting step in our training process requires an initial normalization, a gross alignment, and a fitting via the optimization (minimization) on the objective function defined in equation (5.5).

In the initial normalization, each binary image $\mathbf{I}$ is uniformly scaled into a unit cube, so the four terms in the objective function are unitless and commensurate to one another. The binary fitting then continues with the gross alignment, typically via a similarity transformation implied by the moments of a template m-rep model and the moments of the binary image. The gross alignment is followed by the optimization (minimization) on the objective function over m-rep parameters. In other words, a template m-rep model is aligned and then fitted into each normalized target binary image in the proper fitting step, which is the first step in our statistical training process.

Our statistical training process include two main steps. In the first step, a set of m-rep models are extracted from a set of binary training images by the proper fitting. In the second step, shape and appearance statistics are calculated from the fitted m-reps. The shape statistics are calculated by applying PGA [Fletcher et al. (2003)] to the fitted m-rep models. The fitted m-rep models are also mapped back to their corresponding gray scale images, which are divided into regions using model relative coordinates provided by the m-rep models. Intensity histograms of these divided image regions are collected and converted into regional intensity quantile functions (RIQFs) [Broadhurst et al. (2005); Stough et al. (2007)]. The appearance statistics are calculated from these RIQFs. The calculated PGA shape statistics and RIQF appearance statistics form the learned statistics that are used in applications, such as image segmentation and shape discrimination.

This dissertation focuses on the task of image segmentation. Section 5.2 shows training and image segmentation results based on the proposed proper fitting step applied

to both synthetic and real world images. The synthetic binary images are generated by warping a standard ellipsoid, and the real world images are male pelvic CT images containing target objects of prostates and bladders.

5.2   Results

5.2.1   Synthetic Images

A diffeomorphic deformation to the ambient space $\mathbb{R}^3$ of a standard ellipsoid was used to generate synthetic binary images of warped ellipsoids. The proposed proper fitting method was then used to fit a template m-rep model into these binary images, and the quality of the fitted models was measured.

A detailed description of the generation of warped ellipsoids can be found in section 6.1.1. Briefly, three deformations, each controlled by an independent and normally distributed variable, are applied to the ambient space of a standard ellipsoid. Figure 5.3-Left shows the three types of deformation: bending, twisting, and tapering. The standard ellipsoid is used to generate an initial binary volumetric image. The initial binary image is then warped by a sample set of the three deformations to a new sample binary image. The standard ellipsoid is therefore transformed into a warped ellipsoid in the new sample binary image. By this means, a set of 150 sample binary images containing warped ellipsoids was generated. The ground truth boundary of each warped ellipsoid was analytically calculated for each sample binary image.

In each sample binary image, image landmarks were also automatically identified, corresponding to the four extreme points in the middle section of a warped ellipsoid boundary, plus two vertices as the two warped tips of the warped ellipsoid. A template m-rep model was generated by sampling the medial axis of the standard ellipsoid. This template m-rep was then fitted into all the sample binary images using the proposed binary fitting method with the geometric illegality penalty. The quality of the fit was

103
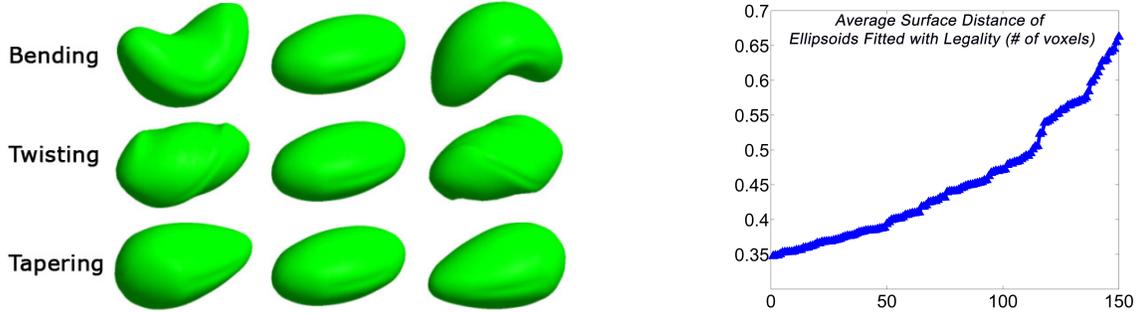
Figure 5.3: **Left**: warped ellipsoids by three deformations of bending, twisting and tapering. Each deformation is shown as -2λ, 0, +2λ away from the mean, where λ is the standard deviation; **Right**: fitting results of 150 sample images shown as the average surface distances (in number of voxels) from model boundaries to binary boundaries.

measured by the average surface distance between the fitted m-rep boundary and the ground truth binary boundary.

Since the synthetic binary images are generated independently from our medial representation, this study of the fitting quality serves as a valid consistency test of the proposed proper fitting. The results in figure 5.3-Right show that the m-rep implied boundaries are, on average, less than one image voxel from the binary boundaries as the ground truth. This measurement should be compared to the three principal radii, (a, b, c), of the original ellipsoid, approximately $(50, 30, 23)$ voxels in a binary image, indicating that the fitted m-rep boundaries are close to the the ground truth boundaries. Furthermore, in some test cases, the proper binary fitting shows more robustness by providing good fitted shape models, which we failed to generate without using the illegality penalty term.

5.2.2   Real World Shape Models

CT images $(1 \times 1 \times 3mm)$ $\mathbf{I}_i, i \in [1, 80]$ of prostates and bladders from five patients (80 images in total from different days) were used as the real world test data [Han et al. (2007)]. Each patient has from 13 to 18 images from multiple days. All 80 images were segmented by a human expert to generate a set of segmented binary images

104

$\{\mathbf{I}_i^b\}, i \in [1, 80]$. Both the proposed proper fitting and a binary fitting without using the illegality penalty term were used to generate fitted m-reps $\{\mathbf{M}_i^{with\_F_{leg}}\}$ and $\{\mathbf{M}_i^{wo\_F_{leg}}\}$ from the segmented binary images $\{\mathbf{I}_i^b\}$.

Both of the sets of fitted m-reps $\{\mathbf{M}_i^{with\_F_{leg}}\}$ and $\{\mathbf{M}_i^{wo\_F_{leg}}\}$ were used to learn leave-one-day-out PGA shape statistics and RIQF appearance statistics for each patient: each day was successively left out to learn on all remaining days for one patient, and the CT image of the left-out day of that patient was segmented using the learned shape and appearance statistics of all other days. The results based on the fittings with and without using the illegality penalty were compared with each other. The comparison results are shown in the next paragraphs.

The first comparison is conducted on the quality of fitted m-rep models $\{\mathbf{M}_i^{with\_F_{leg}}\}$ and $\{\mathbf{M}_i^{wo\_F_{leg}}\}$. The results show that the models $\{\mathbf{M}_i^{with\_F_{leg}}\}$ using the illegality penalty not only have smoother surfaces (in figure 5.4) but also fit better (figure 5.5 - curves represented by triangle signs) into the binary images than the fitted models $\{\mathbf{M}_i^{wo\_F_{leg}}\}$ without using the illegality penalty. The robustness of the proper fitting provides good fitted models even from the binary images that we failed to fit without using the illegality penalty.

The second comparison is conducted on the quality of the leave-one-day-out segmentations. Sorted statistics over all 80 image segmentations are given in figure 5.5 - curves represented by circle signs. The segmentations based on learned statistics from fitted models using the illegality penalty $\{\mathbf{M}_i^{with\_F_{leg}}\}$ are better than the segmentations based on learned statistics from fitted models without using the illegality penalty $\{\mathbf{M}_i^{wo\_F_{leg}}\}$. For the majority of cases, the segmentation results of both bladders and prostates based on the proper fitting were judged qualitatively good by experts.

The average agreement between the leave-one-day-out segmentations, based on the proper training, and their corresponding segmentations by the human expert is 93% volume overlap and $1.6mm$ average closest point surface separation. A second human
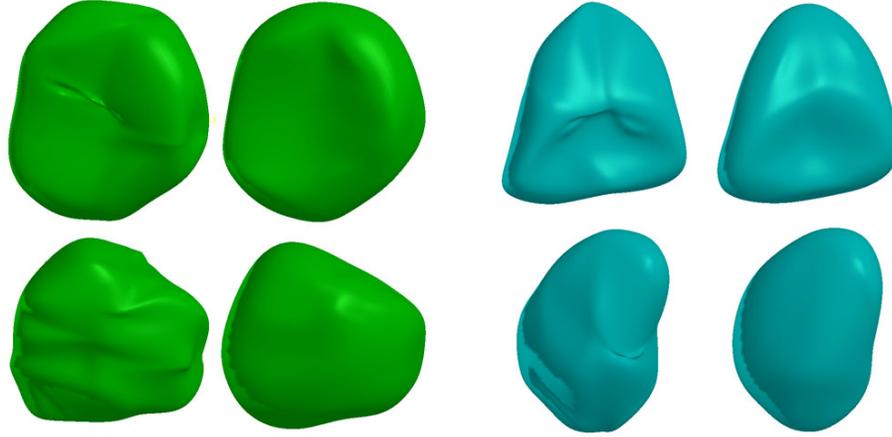
Figure 5.4: Sample fitted models from fittings without and with using the illegality penalty are shown. Green objects are for prostates, and cyan objects are for bladders. The left column for each object shows the fitted models without using the illegality penalty, and the right column shows the fitted models using the legality penalty. It shows that the fitted models using the illegality penalty are smoother and free of local shape defects.

expert also segmented 16 images of the total 80 images to get 16 prostate models, and the average agreement between two humans' segmentations of these 16 prostates is 81% volume overlap and $1.9mm$ average closest point surface separation.

The results, as shown in figure 5.4 and 5.5, indicate that adding an illegality penalty in our proper fitting step yields better fitted models and hence yields better statistics on geometry and intensity patterns, which consequently generate considerably better segmentation results. The segmentation results are indeed so good that on bladder and prostate segmentations from CT images, the segmentations by the computer based on the proper fitting were closer to the manual segmentations by the expert trainer, who created the training binary images, than another expert's manual segmentations were to the trainer's. The results show that the proposed proper fitting not only is mathematically sound but also improves segmentation results in practice. Training shape models with their illegalities penalized have proven to be crucial to acquiring proper statistics, summarizing the shape variation of anatomical objects in a population.
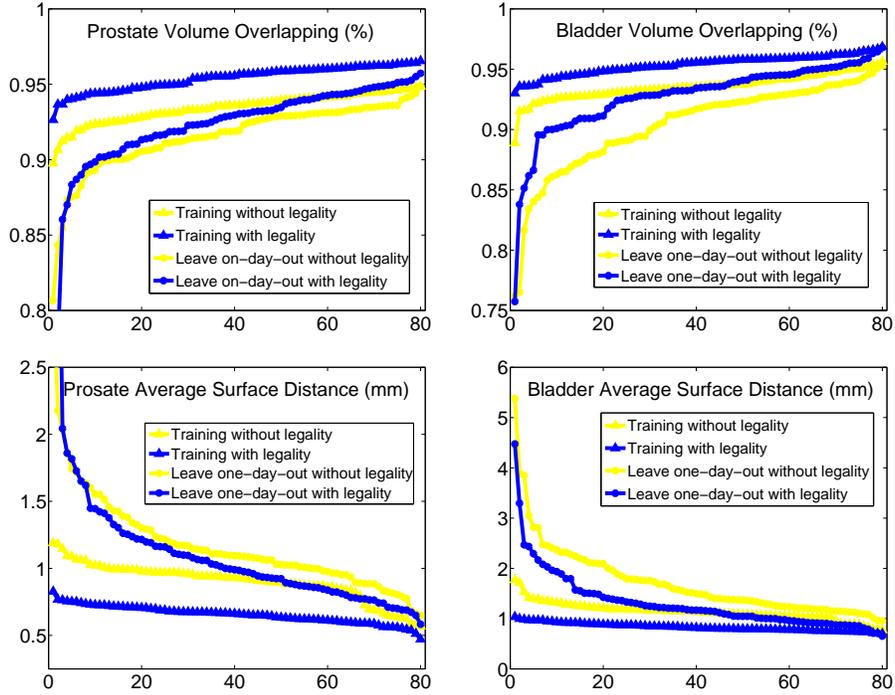
Figure 5.5: Sorted measurements (volume overlap measured by *Dice similarity coefficients* [Dice (1945)] in the top row and average surface distance in *mm* in the bottom row) for prostates (in the left column) and bladders (in the right column), comparing fitted m-reps from binary images segmented by a human expert (represented by triangle) and segmented m-reps from a leave-one-day-out experiment (represented by circle) using the fitting without the illegality penalty (curves in yellow) and the fitting with the illegality penalty (curves in blue). The measurements demonstrated in each curve are sorted independently for that curve over all 80 test cases. The average surface distance is measured for each test m-rep as the arithmetic average of the shortest distances between the m-rep implied boundary points to the target boundary mesh.

5.3   Discussion

The proposed proper fitting on single figure m-reps can be extended to fit multi-figure m-reps, representing anatomical objects with multiple parts, such as livers and hearts. This extension is further discussed in section 6.2.

In the current implementation of the proper fitting, the adapted Catmull-Clark subdivision is used to generate an m-rep implied boundary, and the illegality penalty term is applied to original atom spokes in the m-rep. The atom interpolation method, described in chapter 3, generates not only an m-rep implied boundary but also a smooth spoke field, and it can be used in the proper fitting to further improve shape legality of fitted m-rep models and learned statistics based on these models, which is discussed in section 7.2.2.

Chapter 6

Evaluation of Atom Interpolation and Multi-figure M-rep Fitting

Chapters 3-5 described the main contributions of this dissertation: atom and spoke interpolation used in the interpolation of a single figure or multi-figure m-rep, formation of multi-figure m-reps with blend regions to represent complex objects, and proper fitting using an illegality penalty in the objective function. This chapter further evaluates atom interpolation used in the interpolation of single figure m-reps and evaluates formation of multi-figure m-reps.

The evaluation of atom interpolation in this chapter focuses on the quality of the m-rep boundary implied by interpolated atoms. As described in chapter 3, atom interpolation was used to generate interpolated spokes from an m-rep and hence to generate an m-rep implied boundary. Section 3.3 qualitatively demonstrated the result of the atom interpolation method. In the evaluation here, each m-rep implied boundary generated by the atom interpolation method was compared with the boundary generated by the subdivision based method, reviewed in section 3.1. The quantitative comparison result is shown in section 6.1.2.

The evaluation of the multi-figure m-rep is carried out by using a multi-figure m-rep in a binary fitting. The binary fitting of a multi-figure m-rep was implemented via a multi-resolution and multi-step method, based on the binary fitting of a single figure

m-rep, detailed in section 5.1. The quality of fitted multi-figure m-reps was measured, and the result is shown in section 6.2.

6.1   Evaluation of Atom Interpolation

Atom interpolation was evaluated on samples of single figure m-reps from object populations. It is desirable for a test sample of m-reps to have the following properties:

1. The test sample well represents an object population, i.e., the sample captures the shape variation in the population;

2. The test sample has a desirable size, reflecting the dimensionality of the original shape models. Empirically the number of objects in a test sample should be proportional to the dimension of the shape model representing the population of objects;

3. The test sample is feasible to generate, with a small set of control parameters.

In section 6.1.1, two test samples that have the two properties are described: the first test sample consists of synthetic m-reps, and the second test sample consists of m-reps sampled from a learned shape space of real world objects. In section 6.1.2, implied boundary surfaces of the sample m-reps generated by atom interpolation are compared with implied boundary surfaces generated by the subdivision method.

6.1.1   Generation of Test Samples

Two test samples of m-reps were generated for the evaluation. The first test sample was generated by warping a standard ellipsoid with three simple transformations, each controlled by one independent and normally distributed parameter. The second test sample was generated by sampling a pre-learned shape space of anatomical objects,

kidneys in this case. The pre-learned shape space is given by learned PGA statistics from a set of training kidney m-reps.

Warped Ellipsoids with Ground Truth Boundaries

In the method I developed, warped ellipsoids are generated from a standard ellipsoid $\mathbf{E}_0$ given as $\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} \leq 1$, where $a = 0.75, b = 0.5$, and $c = 0.375$. A diffeomorphic transformation is applied to the ambient space of the standard ellipsoid volume, so the standard ellipsoid is warped to form a warped ellipsoid.

Each such diffeomorphic transformation consists of three sub-transformations: bending, twisting, and tapering, each controlled by an independent random variable α, β, and γ, respectively. Each of these three random variables belongs to a zero-mean normal distribution. The standard deviations of the three normal distributions are pre-determined to guarantee that the first three major modes of shape variations of the warped ellipsoids are distinctly represented by variances in a decreasing order: 1.5 for bending, $\frac{\pi}{3}$ for twisting, and $\frac{3}{2}\sqrt{2}$ for tapering. The application order of these three transformations is twisting first, tapering second, and bending last. The overall diffeomorphic transformation F is defined as follows:

$$x' = x \tag{6.1}$$

$$y' = (y\cos(\beta x) - z\sin(\beta x))e^{\gamma x} \tag{6.2}$$

$$z' = (y\sin(\beta x) + z\cos(\beta x))e^{\gamma x} + \alpha x^2 \tag{6.3}$$

Because the Jacobian determinant of F is $e^{2\gamma x}$, which is always positive, F is indeed a diffeomorphic transformation. The following outlines the overall method to generate a sample of warped ellipsoids:

1. α, β, γ are randomly sampled from their corresponding normal distributions;

111

2. Each set of sampled $(\alpha_i, \beta_i, \gamma_i), i \in [1, N]$ is plugged into F_i, where N is the total number of samples;

3. F_i is applied to a binary characteristic image $\mathbf{I}_0$ containing the volume of the standard ellipsoid $\mathbf{E}_0$ in order to generate a warped binary image $\mathbf{I}_i$ containing a warped ellipsoid $\mathbf{E}_i$;

4. The marching cubes method [Lorensen and Cline (1987)] is applied to each warped binary image $\mathbf{I}_i$ to reconstruct a ground truth voxel boundary $\mathbf{B}_i$ of the warped ellipsoid $\mathbf{E}_i$, with sub-voxel accuracy.

An m-rep $\mathbf{M}_i$ was extracted from each binary image $\mathbf{I}_i$, implemented by proper binary fitting, detailed in chapter 5, section 5.1. Each fitting requires the following steps:

1. The medial axis, $\frac{x^2}{\frac{a^2}{a^2-c^2}} + \frac{y^2}{\frac{b^2}{b^2-c^2}} \leq 1$, of the standard ellipsoid $\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1$ is sampled to generate an m-rep $\mathbf{M}_0$ for the standard ellipsoid;

2. F_i is applied to $\mathbf{M}_0$, atom by atom, to generate a new m-rep $\mathbf{M}_{i,0}$, which is a close approximation to the underlying m-rep of the warped ellipsoid $\mathbf{E}_i$ in $\mathbf{I}_i$;

3. $\mathbf{M}_{i,0}$ is fitted into the binary image of $\mathbf{I}_i$ of $\mathbf{E}_i$ via the binary fitting step to generate the optimal m-rep $\mathbf{M}_i$ for $\mathbf{E}_i$.

The fitted m-reps $\mathbf{M}_i, i \in [1, N]$ formed the first test sample for the evaluation. The generation of a second test sample is described next before the evaluation results are shown.

Sampled Kidneys

M-reps were sampled from a pre-learned shape space of kidney m-reps. This shape space is given by a mean m-rep $\overline{\mathbf{M}}$ and first five principal geodesic modes $\mathbf{v}_k, k \in [1, 5]$. These five principal geodesic modes represent over 98% of the total shape variation in the

112

original training sample of kidneys from a population. Monte Carlo sampling technique [Metropolis and Ulam (1949)] was used to sample the pre-learned shape space:

1. A set of weights $\mathbf{w}_i = (w_{i,k}), k \in [1, 5], i \in [1, N]$ is sampled, with each $w_{i,k}$ sampled from a normal distribution $\mathcal{N}(0, 1)$, where N is the total number of sample kidney m-reps;

2. $\overline{\mathbf{M}}$ and $\sum_{k=1}^{5} w_{i,k} \mathbf{v}_k$ are used to reconstruct a sample kidney m-rep $\mathbf{M}_i, i \in [1, N]$ via the exponential map defined on the manifold of m-reps, reviewed in section 2.4.1.

The kidney m-reps $\mathbf{M}_i$ formed the second test sample for the evaluation. Both of the test samples allow an arbitrary number of sample objects to be generated. Controlled by a small set of parameters, the generation of both of the test samples is computationally feasible. These sample m-reps capture the types of shape variation in real world anatomical objects. They uphold the desirable properties of a test sample. Evaluation results based on both the test samples are shown next.

6.1.2   Evaluation Results of Atom Interpolation

150 warped ellipsoid m-reps and 150 sample kidney m-reps were generated as two test samples for the evaluation.

For the sample of warped ellipsoids, the implied m-rep boundary surfaces implied both by the atom interpolation and by the subdivision based method were compared with the ground truth voxel boundary surfaces re-constructed from the warped binary images. The results are shown in figure 6.1.

For the sample of kidney m-reps, the implied boundary surfaces by atom interpolation were compared with the implied boundary surfaces by the subdivision based method. The results are shown in figure 6.2.
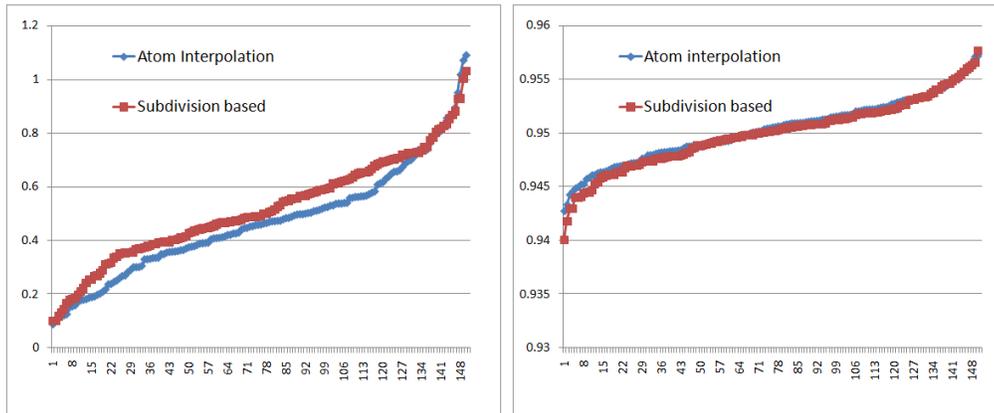
Figure 6.1: **Left**: comparison results of warped ellipsoids shown in terms of average surface distance between two surface meshes. The distance is shown in number of voxels; **Right**: comparison results of warped ellipsoids shown in terms of the Dice coefficient, i.e., the volume overlap between two surface meshes. The measurements demonstrated in each curve are sorted independently for that curve over all 150 test cases. The average surface distance is measured for each test m-rep as the arithmetic average of the shortest distances between the m-rep implied boundary points to the target boundary mesh.
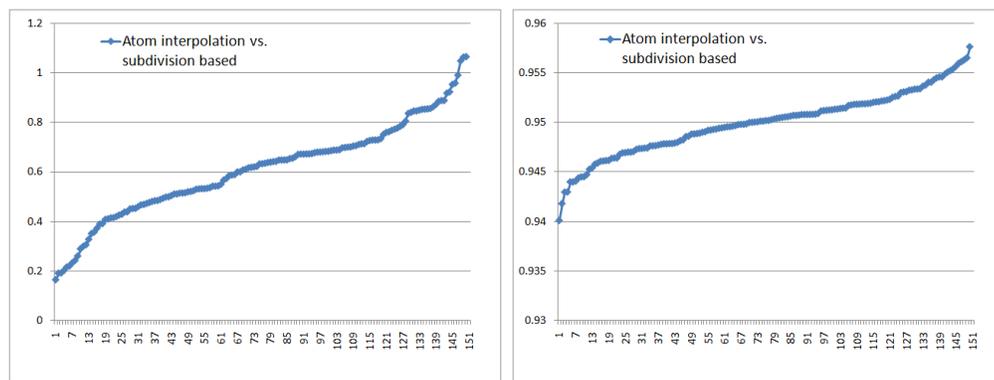
Figure 6.2: **Left**: comparison results of sampled kidney m-reps shown in terms of average surface distance between two surface meshes. The distance is shown in number of voxels; **Right**: comparison results of sampled kidney m-reps shown in terms of the Dice coefficient, i.e., the volume overlap between two surface meshes. The measurements demonstrated in each curve are sorted independently for that curve over all 150 test cases. The average surface distance is measured for each test m-rep as the arithmetic average of the shortest distances between the m-rep implied boundary points to the target boundary mesh.

The left parts of figures 6.1 and 6.2 show the comparison results in terms of average surface distances between pairs of compared surface meshes, which are within sub-voxel accuracy for both the warped ellipsoid m-reps and kidney m-reps. The right parts of the two figures show the comparison results in terms of the volume overlap measured by the Dice coefficient [Dice (1945); Zijdenbos et al. (1994)] between object volumes within surface meshes, which are over 94% for both the warped ellipsoid m-reps and the kidney m-reps.

Given the adequate sample size of both the test samples, the comparison results show both good validity and reliability of atom interpolation: atom interpolation works as well as the subdivision method in generating implied boundary surfaces of single figure m-reps. As an advantage, atom interpolation also provides the underlying spokes of each interpolated m-rep.

6.2   Evaluation of the Multi-figure M-rep

The multi-figure m-rep was used in binary fitting to extract multi-figure m-reps from a set of binary images of livers, and the quality of the fitted multi-figure m-reps was evaluated. A two-figure m-rep is used to demonstrate the binary fitting.

In order to fit a two-figure m-rep into a set of binary images of livers, an objective function was optimized by a series of large-scale-to-small-scale optimizations over transformations applied to a two-figure m-rep template. The fitting at a respective scale corresponds to a fitting stage in the multi-scale and multi-stage process. The objective function is based on the objective function 5.5 used in the binary fitting of single figure m-reps, described in chapter 5. The overall objective function of fitting a two-figure m-rep was evaluated on both its figures, shown as follows:

$$F_{obj}^{multi}(\mathbf{M}|\mathbf{I}) = \alpha F_{img}(\mathbf{M}, \mathbf{I}) + \gamma F_{reg}(\mathbf{M}) + \delta F_{ref}(\mathbf{M}, \mathbf{M}_0) \qquad (6.4)$$

where F_{img} was the image match term defined in equation 5.1, measured on the overall boundary surface of $\mathbf{M}$ with blending, where F_{reg} was the regularity penalty term defined in equation 5.3, measured on the host figure $\mathbf{M}_{host}$ and subfigure $\mathbf{M}_{sub}$ of $\mathbf{M}$, where F_{ref} is a reference penalty term to be described next, and where $\alpha, \gamma, \delta > 0$ are empirically determined parameters such that $\alpha + \gamma + \delta = 1$.

F_{ref}, the reference penalty term, penalizes any big shape change of $\mathbf{M}$ from a reference m-rep $\mathbf{M}_0$. F_{ref} is defined as the geodesic distance between the current $\mathbf{M}$ and a reference m-rep $\mathbf{M}_0$, with which the binary fitting starts:

$$F_{ref}(\mathbf{M}) = dis(\mathbf{M}, \mathbf{M}_0)$$

where dis is the geodesic distance between $\mathbf{M}$ and $\mathbf{M}_0$, defined by equation 2.23 in section 2.4.

The illegality penalty term was not included in the objective function. Adding it to the fitting of the multi-figure m-rep will improve the fitting results. In this evaluation, the Catmull-Clark subdivision based method was used to generate the implied boundary meshes including blend regions of liver m-reps. The transformation associated with each fitting stage and the overall binary fitting are described in subsection 6.2.1. The evaluation result is shown in subsection 6.2.2.

6.2.1 Multi-scale Fitting Process

The binary fitting of a two-figure m-rep is realized by optimizing the objective function defined in equation 6.4 for the m-rep. This objective function is optimized over a sequence of transformations, successively finer in scales, applied to a two-figure m-rep template $\mathbf{M}_0$. $\mathbf{M}_0$ has n atoms in total: n_1 atoms in its host figure $\mathbf{M}_{host}$ and n_2 atoms in its subfigure $\mathbf{M}_{sub}$, where $n = n_1 + n_2$. Recall that $G(n)$ stands for the composed transformations applied smoothly to n m-rep atoms, reviewed in section 2.4.3. In the

following description of transformations to be applied to a two-figure m-rep, $\mathcal{SO}(3)$ is the special orthogonal group of all $3D$ rotations as a subgroup of $G(1)$ for one m-rep atom.

- **Initial alignment**: $T_1 \in \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{SO}(3)$ applied to align the template $\mathbf{M}_0$ to a target binary image $\mathbf{I}$;

- **object stage**: $T_{obj} \in \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{SO}(3)$ applied to the entire m-rep object as a whole;

- **host figure**: the host figure $\mathbf{M}_{host}$ is the target of this stage, and the subfigure is deformed by a deformation propagation from the hinge atoms to the rest of the subfigure, described in section 4.2.1;

  - figure stage: $T_{host\_fig} \in \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{SO}(3)$ applied to the host figure;

  - atom stage: $T_{host\_atom} \in G(n_1)$ applied to host figure atoms $\mathbf{A}^1_{1,2,...,n_1}$ one by one;

- **subfigure stage**: the subfigure $\mathbf{M}_{sub}$ is the target of this stage;

  - figure stage: $T_{sub\_fig} \in \mathbb{R}^3 \times \mathbb{R}^+ \times \mathcal{SO}(3)$ applied to the subfigure;

  - atom stage: $T_{sub\_atom} \in G(n_2)$ applied to subfigure atoms $\mathbf{A}^2_{1,2,...,n_2}$ one by one.

The detailed fitting algorithm is shown as follows:

input:

a two-figure m-rep template $\mathbf{M}_0$ with a host figure $\mathbf{M}_{host}$ and a subfigure $\mathbf{M}_{sub}$;

a binary target image $\mathbf{I}$.

output:

a two-figure m-rep $\mathbf{M}$ extracted from $\mathbf{I}$.

framework:

for each binary target image I {

1. Find an optimal T_1 to initially align $\mathbf{M}_0$ to $\mathbf{I}$, $\mathbf{M}_1 = T_1 \circ \mathbf{M}_0$;

2. $T_2 = \arg \min_{T_{obj}} F_{obj}^{multi}(T_{obj} \circ \mathbf{M}_1 | \mathbf{I}), \mathbf{M}_2 = T_2 \circ \mathbf{M}_1$;

3. $T_3 = \arg \min_{T_{host\_fig}} F_{obj}^{multi}(T_{host\_fig} \circ \mathbf{M}_{2,host} \subset \mathbf{M}_2 | \mathbf{I}), \ \mathbf{M}_3 = T_3 \circ \mathbf{M}_2$;

4. $T_4 = \arg \min_{T_{host\_atom}} F_{obj}^{multi}(T_{host\_atom} \circ \mathbf{M}_{3,host} \subset \mathbf{M}_3 | \mathbf{I}), \ \mathbf{M}_4 = T_4 \circ \mathbf{M}_3$;

5. Propagate the changes in hinge atoms to the rest of the subfigure, let the new m-rep be $\mathbf{M}_4'$;

6. $T_5 = \arg \min_{T_{sub\_fig}} F_{obj}^{multi}(T_{sub\_fig} \circ \mathbf{M}_{4,sub}' \subset \mathbf{M}_4' | \mathbf{I}), \ \mathbf{M}_5 = T_5 \circ \mathbf{M}_4'$

7. $T_6 = \arg \min_{T_{sub\_atom}} F_{obj}^{multi}(T_{sub\_atom} \circ \mathbf{M}_{5,sub} \subset \mathbf{M}_5 | \mathbf{I}), \ \mathbf{M} = T_6 \circ \mathbf{M}_5$;

}

The overall binary fitting of a two-figure m-rep consists of fitting stages at increasingly finer scales. The iteration of all fitting stages is repeated to improve the quality of fitting. Starting from the second iteration, the resulting m-rep from the previous iteration is used as the template m-rep in the current iteration. This framework can be extended to a multi-figure m-rep of more than two figures. However, this evaluation and its result focused on m-reps of two-figures.

6.2.2   Test Sample Generation and Evaluation Result

The diagram in figure 6.3 lays out the evaluation process based on a test sample of liver m-reps (figure 6.4-left), which are generated using Monte Carlo sampling method, described in section 6.1.1. A learned PGA shape space of livers was sampled to generate 50 liver m-reps of two figures in the test sample. Each liver m-rep was used to generate a binary image and a ground truth boundary surface. The mean m-rep in the PGA

| $p(M)$	The Monte Carlo method is used to generate N target m-reps M_i ;
①	
$M_i, i = 1,2,..., N$	Danielsson distance mapping, from the implied boundary of each sampled M_i, is used to generate the target binary image I_i of the interior of M_i ;
②	
$I_i, i = 1,2,..., N$	
③	Deform the mean m-rep $\overline{M}$ into each target image I_i using the extraction method. The extracted m-rep is M_i' ;
$M_i', i = 1,2,..., N$	
④	The geodesic distance $\|\text{Log}_{M_i}(M_i')\|$ is then used as the measure of the extraction quality.
$\frac{1}{N}\sum_{i=1}^{N}\left\|\text{Log}_{M_i}(M_i')\right\|^2$	

Figure 6.3: Diagram flow to lay out the evaluate process given learned shape space statistics $p(\mathbf{M})$ from a population of livers
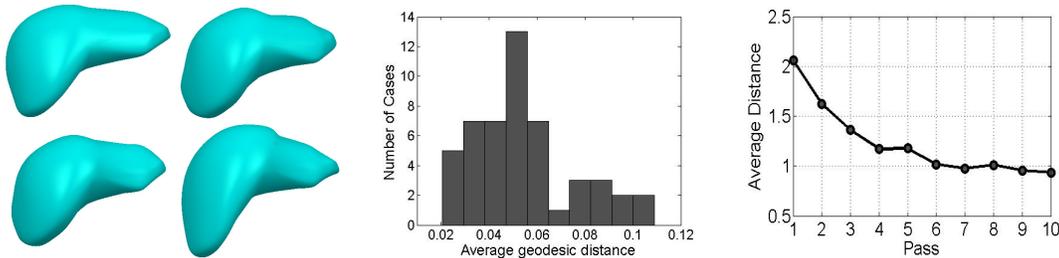
Figure 6.4: **Left**: **4** of the **50** sample m-reps used in the evaluation. **Middle**: evaluation result of the binary fitting shown as a histogram of geodesic distances between the fitted m-reps $\mathbf{M}_i$ and the ground truth m-reps. **Right**: in the first ten iterations of repeated binary fittings, the quality of fitted m-reps improves while the average distances between the m-rep implied surfaces and the ground truth surfaces decrease. The distances are shown in number of image voxels.

statistics, used as the template m-rep, was fitted into each of the 50 binary images. Each fitted m-rep was then compared with its corresponding ground truth surface.

The evaluation result is shown in figure 6.4. Figure 6.4-middle shows a histogram of the geodesic distances, i.e., mismatches, between the fitted m-reps and their corresponding ground truth liver m-reps, from which the target binary images are generated. The average geodesic distance across all the livers is 0.054, which is equivalent to the width of one image voxel in each binary image. Figure 6.4-right shows the trend of the average surface distances between m-rep implied surfaces and the ground truth surfaces in ten iterations of binary fittings. The average surface distance is 0.674 image voxels

over all 50 test cases after ten iterations of binary fitting. The panels in figure 6.4 show that repeated binary fittings improve the fitting quality and that the fitted two-figure m-reps are close approximations to the ground truth.

This binary fitting can be extended to extract objects of more than two parts from binary images. Fittings of this form can produce m-reps for most anatomic objects, even complex ones.

Chapter 7

Summary and Discussion

A general goal of shape representations in a deformable model framework is to form shape models and probability distributions on these models in a population for use in applications, such as image segmentation and statistical shape characterization. The discrete m-rep, as a medial shape representation used in the deformable model framework, is no exception. This dissertation focuses on enhancing the m-rep to better achieve this goal.

This chapter revisits and discusses the main contributions of this dissertation and their relations to the goal of forming models and model probability distributions in section 7.1. Section 7.2 discusses open problems encountered in the previous chapters 3 and 4, and it discusses potential future research based on the proposed methods in this dissertation.

7.1 Summary of Contributions

The contributions of this dissertation are reviewed as follows, followed by a summary of my thesis claims.

1. *A method to interpolate a discrete medial model into a continuously parameterized one for both simple objects with one single part and complex objects with multiple parts;*

The advantage of a medial model over a boundary-based model is that the medial model represents not only the boundary but also the interior of an object. Interpolation of a discrete medial model, if conducted properly, can provide a proper object boundary and a valid parameterization of the object interior volume and its adjacent exterior volume. Such a parameterization is crucial to forming proper appearance probability distributions relative to the shape models. The challenge of interpolating a medial shape model in a discrete form is that all the dimensions of each medial component are strongly related and must fulfill certain legality conditions.

The proposed method to interpolate a single figure m-rep overcomes this challenge by atom interpolation, described in chapter 3, section 3.2. The proposed atom interpolation explicitly uses the $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices. These matrices, as extensions to the radial shape operator $\mathbf{S}_{rad}$ and the edge shape operator $\mathbf{S}_E$ as reviewed in section 2.3, record the rate of changes of both the medial sheet and the medial spokes in an m-rep. Based on the $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices, the proposed atom interpolation reflects the relations among all dimensions in an m-rep atom. By making use of the powerful mathematics in terms of legality conditions based on the same $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices, atom interpolation enforces the local legality of an interpolated m-rep. The result in section 3.3 and the evaluation in section 6.1.2 showed that atom interpolation applied well to both synthetic and real-world objects in order to form proper continuous shape models for these objects.

Based on atom interpolation on a single figure m-rep, spoke interpolation is used to generate a blend region between each pair of connected figures in a multi-figure m-rep, detailed in section 4.2.3. Briefly, atom interpolation is first applied to a host figure and its connected subfigure, and then a one-sided spoke interpolation is used to generate the underlying medial structure for a blend region as a transition between the two connected figures. An interpolated multi-figure m-rep is a skeletal

structure: the two-sided medial sheet of the subfigure splits into a one-sided collar-shaped medial sheet, and this collar-shaped medial sheet for the blend merges into the medial sheet of the host figure. As described in section 4.2.3, this skeletal structure is a relaxation from a branching Blum medial axis for better stability. The legality of the interpolated blend region is also maintained.

The proposed interpolation method on a discrete single figure or multi-figure m-rep provides a continuous volumetric shape representation and still allows localized shape controls via discrete m-rep atoms. Local shape changes, such as bending, twisting, and tapering, are well captured. Furthermore, the enforced legality of interpolated m-reps is important to forming proper probability distributions on the appearance of anatomical objects in a population.

2. *A description of geometric interrelations among adjacent parts in an object of multiple parts*

Many anatomical objects have more than one named part. The challenge of forming shape models and shape probability distributions for these complex objects lies in the fact that modeling each individual part alone is not sufficient. A complex object is more than just the sum of all its parts. The inter-relations among parts should be modeled explicitly to capture both the global and local shape variations in these complex objects. Such relations include the host- and sub- part connection, the global deformation to all connected object parts, and the implied deformation from one part to its connected part(s).

The multi-figure m-rep described in chapter 4 overcomes this challenge by representing each object part with a single figure m-rep and by connecting the parts, pair by pair, via hinge geometry, as detailed in section 4.2.1. Hinge geometry explicitly models the interrelation in a host figure and subfigure pair. Based on hinge geometry, global and local deformations are defined for a multi-figure m-rep

in section 4.2.2 and 4.2.2, respectively. Hinge geometry and the series of deformations defined in a multi-figure m-rep allow us to look at a complex object at multiple levels of scale, from a global scale for the entire object, to an intermediate scale for the interrelation between each connected pair, and to a local scale for each individual object part. Combined with figure and atom scales in each m-rep figure, a multi-figure m-rep provides a multi-scale way to form shape models and probability distributions on shape models for complex objects.

3. *A means to calculate the hierarchical statistics of an object with multiple parts, from the global level for the entire object to the local level for each individual figure*

The hierarchy in a multi-figure m-rep follows the natural hierarchy in most anatomical objects. This hierarchy provides the basis for the hierarchical statistical analysis on complex objects of multiple parts, detailed in section 4.3. Each multi-part object is understood as a whole entity, as interrelations between each pair of connected m-rep figures, and as individual figures, at decreasing levels of scale. The formed multi-scale statistics in terms of probability distributions on shape models, as shown in section 4.4, capture the shape variations in complex objects at different scales.

The multi-figure m-rep also has the advantage that a subtractive object part, represented by an indentation subfigure, connects to its host figure in the same way as an additive part, represented by a protrusion subfigure. This advantage enables a uniform statistical analysis on multi-figure objects containing indentation or protrusion subfigures, making it feasible to form probability distributions by one work flow on the shapes and appearances of these objects from a population.

4. *An explicit use of an illegality penalty in binary fitting to improve the smoothness of fitted shapes, thereby leading to properly trained shape and appearance statistics and eventually better segmentation results*

The $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices used in atom interpolation are also used in an explicit illegality penalty term of the objective function for binary fitting, detailed in chapter 5, section 5.1. The illegality penalty term is based on the radial and edge legality conditions, reviewed in section 2.3, and hence this penalty term enforces the legalities of fitted shapes in a binary fitting step, in which there is no shape prior available to guide the fitting. This binary fitting using the illegality penalty is an important component in a proper training process, providing proper shapes for learning shape and appearance statistics.

The results, as shown in section 5.2, indicate that adding an illegality penalty yields better fitted models and hence better statistics on geometry and intensity patterns, which consequently generate considerably better segmentation results. The segmentation results are indeed so good that the segmentations by the computer based on the proper training are closer to the manual segmentations by an expert trainer than another expert's manual segmentations were to the trainer's. The results show that the proposed proper fitting is both mathematically sound and practically useful. Training shape models with their illegalities penalized have proven to be crucial to forming proper shape models and to forming proper statistics, summarizing the shape and appearance variations of anatomical objects in a population, for use in image segmentation.

5. *A way to generate standard test cases with known truth via synthesizing warped ellipsoids or via sampling learned shape space*

In order to evaluate atom interpolation and formation of multi-figure m-reps, a means to synthesize warped ellipsoids and to generate objects by sampling a learned shape space was designed. The method is detailed in chapter 6, section 6.1.1. The proposed method allows an arbitrary number of sample objects to be generated. Controlled by only a small set of parameters, the generation of a test

sample is computationally efficient. Furthermore, the generated m-reps in a test sample capture the types of shape variation in a real anatomical object population.

Here is a revisit of my thesis:

Thesis claim: *An interpolation method, based on sound mathematics, and the hinge geometry, defining the interrelations between connected object parts, provide a powerful tool to derive a medial representation of both single part and multi-part objects from discrete m-reps. When incorporated with explicit geometric constraints, based on the same mathematics, the shape models based on discrete m-reps guarantee proper shape representations of both single and multi-part objects. These shapes allow statistical modeling of both shape and appearances in a population.*

The main goal of my dissertation was to develop proper shape models and proper probability distributions on the shape models for use in applications, such as image segmentation, of simple anatomical objects of one part and complex anatomical objects of more than one part in a population. The discrete m-rep is a powerful medial shape representation, and it was chosen to be the underlying shape representation in this dissertation. This dissertation focused on providing the means of representing complex objects of multiple parts in a multi-scale way, the means of interpolating a discrete single figure or multi-figure m-rep into a continuous legal spoke field on a medial sheet, and the means of maintaining local shape legalities in fitted m-reps for training. These contributions have proven to be crucial to achieving the goal of this dissertation.

The first thesis claim is supported by the mathematics behind the spoke and atom interpolations applied to single figure and multi-figure m-reps and by hinge geometry that defines interrelations between connected pairs of figures in a multi-figure m-rep. Furthermore, the sound mathematics, in terms of the $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices and their related legality conditions, provide an efficient way to enforce the legality of an interpolated m-rep.

The second and third thesis claims are supported by the importance of maintaining the shape legality of fitted m-reps for training. As emphasized in chapter 1, shape defects result in tainted probability distributions on shape models and their corresponding appearances and eventually impair results of applications, such as image segmentation and statistical shape analysis, based on these probability distributions. The same $r\mathbf{S}_{rad}$ and $r\mathbf{S}_E$ matrices used in atom and spoke interpolation are also used in an explicit geometric illegality penalty on m-reps in binary fitting. This explicit legality penalty improves the smoothness of fitted m-reps and consequently improves the learned shape and appearance statistics. As a result, image segmentation results based on these learned statistics are substantially improved.

This dissertation provides new means for the discrete m-rep, including representing complex objects of multiple parts, conducting hierarchical statistics on anatomical objects of multiple parts, interpolating discrete m-reps into continuous forms, and enforcing legalities on interpolated m-reps and on fitted m-reps from binary fitting. These new means allow us to represent not only simple anatomical objects but also complex anatomical objects, by reflecting their parts and their inter-part relations, and to enforce the proper quality of m-rep shape models to learn proper shape statistics and proper appearance statistics relative to these shape models. These means are crucial to modeling anatomical objects in a population by forming proper shape models for these objects and by forming proper probability distributions on these shape models for use in image segmentation or shape characterization.

Compared with boundary-based representations in [Kass et al. (1988); McInerney and Terzopoulos (1996); Caselles et al. (1997); Dryden and Mardia (1998); Cootes et al. (1995, 2001)], the discrete m-rep and interpolated form have its distinct advantages by providing not only a proper boundary surface mesh for an object but also a legal underlying spoke field parameterizing the entire object interior volume and its adjacent exterior volume. A valid parameterization of the object volume provides consistent

access to object appearance relative to the object geometry in order to form proper appearance probability distributions.

Compared with the proposed cm-rep methods in [Yushkevich et al. (2003, 2005); Terriberry et al. (2007); Sun et al. (2008)], the discrete m-rep and interpolated form have medial spokes as an explicit part in its shape representation. As a result, the first-order medial atoms with the medial spokes in the discrete m-rep and interpolated form allow intuitive shape controls over local shape changes, such as bending, twisting, and tapering. Following the exact Blum condition, the cm-rep proposed in [Terriberry et al. (2007)] is able to represent 3D branching medial structures, including the 6-junction branch. The multi-figure m-rep does not support the 6-junction branch, but it provides a robust relaxation from the Blum condition for medial branching structures. The open problem of extending the multi-figure m-rep to support a 6-junction branch will be discussed in the next section.

7.2 Discussion

7.2.1 Remaining Open Problems

A list of open problems and their potential solutions are briefly described as follows:

- Spoke and atom interpolation:

  1. Speed of the internal atom interpolation: each interpolation is calculated by a numerical integration;

     An ideal speed-up lies in an analytical integration instead of the current numerical implementation. In practice such an analytical integration is not available. However, the current numerical implementation can be improved by reducing the calculation for the integrand: in each step of the numerical integration, the integrand can be incrementally evaluated based on the

integrand already evaluated in the previous step. By this means, a certain amount of calculation can be saved in each integration step to speed up the overall integration;

2. Continuity of interpolated spokes: currently C^0 continuity is enforced between interpolated spokes across interpolated patches of internal atoms;

   Recall that the original $r\mathbf{S}_{rad}$ matrices at control atoms are numerically estimated via finite differences. A potential plan to improve this continuity of interpolated spokes is to optimize these original $r\mathbf{S}_{rad}$ matrices, on which the spoke interpolation depends. An iterative process might be applied to optimize these $r\mathbf{S}_{rad}$ matrices for a higher order continuity, such as C^1;

3. Medial umbilic points: a medial umbilic point in a medial axis having repeated principal radial curvatures is problematic for the eigen-decomposition of an $r\mathbf{S}_{rad}$ matrix and thus for the $r\mathbf{S}_{rad}$ and spoke interpolation;

4. A numerical estimation of an initial $r\mathbf{S}_{rad}$ matrix may lead to a matrix with complex-valued eigenvalues and eigenvectors;

   Problems 3 and 4 can be solved by a different $r\mathbf{S}_{rad}$ interpolation method. Recall that the proposed method interpolates $r\mathbf{S}_{rad}$ matrices by interpolating their real eigenvalues and real eigenvectors separately, in section 3.2.2. An $r\mathbf{S}_{rad}$ matrix at a medial umbilic point can be legal but cannot be successfully decomposed into eigenvalues and eigenvectors; an $r\mathbf{S}_{rad}$ matrix with complex eigenvalues and eigenvectors are considered legal because in this case the Jacobian of the medial flow along medial spokes is still non-singular [Damon (2005)]. Therefore, the proposed method to interpolate real eigenvalues and real eigenvectors cannot be applied to these two cases.

   An alternative way to interpolate $r\mathbf{S}_{rad}$ matrices can be used to solve these two open problems: to directly interpolate the matrix entries one by one. As long as the original $r\mathbf{S}_{rad}$ matrices are legal, and the real weights used to in-

129

terpolate these matrix entries are convex, the interpolated $r\mathbf{S}_{rad}$ matrices will stay legal. No eigen-decomposition is necessary in this new interpolation, and the integration based on the interpolated $r\mathbf{S}_{rad}$ matrices will not be affected by any medial umbilic point. This new interpolation on $r\mathbf{S}_{rad}$ matrices is also computationally more efficient because no trigonometric functions need to be evaluated;

- Multi-figure m-reps: the 6-junction branch;

  Although a 6-junction branching medial axis is theoretically generic, it appears to be rare in anatomical structures. It is not directly supported by the current multi-figure m-rep, but a potential solution lies in an extension of hinge geometry between a host figure and subfigure pair to a more complicated connection of more than two figures. Alternatively, a small-scaled branching volume within a 6-junction branch might be well approximated by a multi-figure m-rep and its boundary surface displacements. The boundary displacement method is detailed in [Pizer et al. (2003)], and it can be easily adapted to a multi-figure m-rep given the interpolated underlying spokes.

7.2.2 Potential Extensions

Atom Interpolation of Single and Multi-figure M-reps

Atom interpolation, described in chapter 3 and evaluated in chapter 6, has the potential to generate the surface boundary with the following advantages:

1. It generates not only a boundary but also interior spokes, providing a parameterization of the object interior;

2. It has more robust geometric legality constraints based on sound mathematics;

The subdivision based interpolation is the standard method to generate m-rep implied boundaries in the daily research of our MIDAG. Atom interpolation is currently under further development as an alternative method to generate m-rep implied boundaries.

Extensions of Multi-figure M-reps

In this dissertation, a multi-figure m-rep makes the following assumptions:

1. A multi-figure m-rep hierarchy can have an arbitrary number of levels;

2. In each connection in a multi-figure m-rep hierarchy, there are only two figures connected as the host and subfigure. If a host figure has two or more subfigures, its subfigures have to be sufficiently apart to avoid inter-interferences;

3. The representation does not prohibit the case in which two connected figures share more than one connection or in which more than two figures form a cyclic series of connections;

4. All figures are regular "slab" discrete m-reps.

In order to make such a multi-figure m-rep hierarchy more general, the following extensions are necessary:

1. More complex connections should be allowed to handle the circumstance such as two figures being very close to or even intersecting each other, e.g., two seminal vesicles, as two subfigures, closely attached to a same prostate, as their host figure;

2. The quasi-tubular m-rep proposed in [Saboo et al. (2008)] should be allowed in a multi-figure hierarchy, not only in a simple connection between a pair of m-rep figures but also in a more complicated connection among more than two figures.

The combination of these two extensions allows a representation of complex objects with arbitrary branching structures, such as a lung vessel tree.

On the statistical side, the existing statistical analysis focuses on complexes of multiple single figure m-reps. Incorporating multi-figure m-reps into such a complex is a natural extension to the current statistical analysis. Such an extension allows a statistical analysis on a combination of both single figure and multi-figure objects, which typically exists in real human anatomy but has been a great challenge for boundary-based shape representations. This extension can further improve the discrete m-rep to form more comprehensive probability distributions on object complexes with multi-figure objects while not requiring significant changes to the existing statistical framework of the discrete m-rep.

Investigation into Spoke Correspondences

Correspondences among 3D shape features, such as landmarks or boundary points, across different sample shapes from a population are important for forming proper probability distributions on shape models [Davies et al. (2001); Gerig et al. (2001); Styner et al. (2003)]. By default, the discrete m-rep provides certain correspondences among atom spokes across m-reps, which have been adapted in numerous statistical training processes and have proven to be a reasonable choice. However, the natural questions to be raised are whether the optimality of such spoke correspondences can be defined and how to achieve such optimality of spoke correspondences.

I believe the answer is "yes" to the first question. With the proposed atom interpolation in this dissertation the second question can be answered. Existing methods on improving correspondences across sample shapes concentrated on PDM representations [Davies et al. (2001); Gerig et al. (2001); Styner et al. (2003)]. These proposed methods require a surface mesh to be interpolated into a continuous form in order to move the boundary points within the boundary surface without significantly changing

the shape of the original object. Atom interpolation can be used to derive a continuous spoke field on a medial sheet surface from a discrete m-rep, making an optimization on spoke correspondences possible. Although the exact definition of the optimality of spoke correspondences still requires substantial research, the investigation into spoke correspondences via atom interpolation will be a challenging but rewarding topic.

Extension of the Multi-resolution Framework

The task to model geometric objects is better accomplished under the context of modeling a population of objects instead of any individual object. Modeling a shape space for a population is less sensitive to noise in individual shapes than modeling each individual shape. The m-rep assumes that the objects to be modeled from a population are represented at a fixed resolution, i.e., the size of the atom grid in a single figure m-rep, or the number of subfigures in a multi-figure m-rep. Under such assumption, there are two questions to be answered:

1. What is the optimal resolution for a single figure m-rep to model a specific population of simple anatomical objects?

2. What is the optimal number of subfigures for a multi-figure m-rep to model a specific population of complex anatomical objects?

In the single figure case, if one wants to capture the overall global deformation of a human kidney, a 3×5 m-rep might be sufficient with the advantage of having a lower dimensionality in the feature space; if one wants to achieve a segmentation of high quality for a kidney CT image, a 5×7 version of the kidney m-rep might be required instead.

In the multi-figure case, the optimal number of subfigures depends on the significance of all existing branching structures shared by all the objects in a population.

133

Under such circumstances, the grid size or the number of subfigures in an m-rep can be considered as parameters of an objective function to be optimized. Such an optimization on the m-rep grid size or on the number of subfigures requires the ability to interpolate a discrete m-rep and the ability to represent complex objects with a multi-figure m-rep. The proposed atom interpolation and multi-figure m-rep in this dissertation make this optimization feasible for a discrete m-rep. A shape scale-space of m-reps at varying resolutions can be a useful extension to the existing multi-resolution m-rep framework.

M-reps with a Dynamic Topology

An m-rep with a fixed medial topology, i.e., a fixed hierarchy of connected figures, might be sufficient if the shape variation of an object population can be well characterized by the m-rep. However, the assumption of a fixed topology might be violated if there are necessary changes in the medial topology, and such topology changes pose a significant challenge to existing methods for shape representation and statistical analysis.

For example, lung vessel trees may have varying topologies from patient to patient in a population; significant pathological changes in shapes of anatomical objects may change the topology of underlying medial structures of the objects. The m-rep with a fixed topology may not be adequate for these cases. As a result, an m-rep with a dynamic topology will be inevitable, and novel statistical analysis will be necessary to build a shape space reflecting these medial topology changes. Still, the multi-figure m-rep framework provides a solid base as a reasonable starting point for potential research in this direction.

Bibliography

Bayes, T. (1958). Studies in the history of probability and statistics: Ix. thomas bayes' essay towards solving a problem in the doctrine of chances. *Biometrika (1763)*, 45:296–315. 13

Biermann, H., Levin, A., and Zorin, D. (2000). Piecewise smooth subdivision surfaces with normal control. In *SIGGRAPH 2000*, pages 113–120. 50

Blum, H. and Nagel, R. (1978). Shape description using weighted symmetric axis features. *Pattern Recognition*, 10:167–180. 7, 13, 23

Brechbühler, C., Gerig, G., and Kübler, O. (1995). Parametrization of closed surfaces for 3-d shape description. *Comput. Vis. Image Underst.*, 61(2):154–170. 19

Brechbuhler, C., Gerig, G., and Kubler, O. (1995). Parametrization of closed surfaces for 3-d shape description. *Computer Vision and Image Understanding*, 61(2):154–170. 95

Broadhurst, R. E., Stough, J., Pizer, S. M., and Chaney, E. L. (2005). Histogram statistics of local model-relative image regions. In *International Workshop on Deep Structure, Singularities and Computer Vision (DSSCV)*, pages 72–83. Springer-Verlag. 102

Buss, S. R. and Fillmore, J. P. (2001). Spherical averages and applications to spherical splines and interpolation. *ACM Transactions on Graphics*, 20:95–126. 51

Caselles, V., Kimmel, R., and Sapiro, G. (1997). Geodesic active contours. *IJCV*, 1(2):3–4. 12, 17, 127

Catmull, E. and Clark, J. (1978). Recursively generated bspline surfaces on arbitrary topological meshes. *Computer Aided Design*, 10:350–355. 27

Cootes, T., Cooper, D., Taylor, C., and Graham, J. (1995). Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):18–59. 3, 5, 12, 13, 17, 70, 95, 127

Cootes, T., Edwards, G. J., and Taylor, C. (2001). Active appearance models. *IEEE TPAMI*, 23(6):681–685. 12, 17, 127

Crouch, J. (2003). Medial techniques for automating finite element analysis. *PdD thesis, Dept. of Comp. Sci., UNC @ Chapel Hill*. 51

Csernansky, J., Joshi, S., Wang, L., Gado, M., PhilipMiller, J., Grenander, U., and Miller, M. (1998). Hippocampal morphometry in schizophrenia by high dimensional brain mapping. *National Academy of Science*, 95:11406–11411. 5, 20, 70

Damon, J. N. (2003). Smoothness and geometry of boundaries associated to skeletal structures i: Sufficient conditions for smoothness. *Annales de Institut Fourier*, 53. 4, 35, 75

Damon, J. N. (2004). Smoothness and geometry of boundaries associated to skeletal structures ii: Geometry in the blum case. *Composition Mathematica*, 140(6):1657–1674. 4, 35, 37, 38

Damon, J. N. (2005). Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision*, 61(1):45–64. 4, 32, 35, 38, 41, 42, 129

Damon, J. N. (2008). Swept regions and surfaces: Modeling and volumetric properties. *Theor. Comput. Sci*, 392(1). 62, 64, 65

Danielsson, P. (1980). Euclidean distance mapping. *Computer Graphics and Image Processing*, pages 227–248. 97

Davies, R., C.Twining, Cootes, T., and Taylor, C. (2002). A minimum description length approach to statistical shape modelling. *IEEE Transactions on Medical Imaging*, 21. 8, 18

Davies, R. H., Cootes, T. F., and Taylor, C. J. (2001). A minimum description length approach to statistical shape modelling. In *IPMI2001, LNCS 2082*, pages 50–63. IPMI, Springer-Verlag. 132

Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, 26:297–302. 107, 115

Dryden, I. and Mardia, K. (1998). *Statistical Shape Analysis*. New York. 12, 13, 127

Duncan, L. H. S. J. S. (1992). Boundary finding with parametrically deformable models. *PAMI*, 14(11):1061–1075. 6, 13, 19

Fletcher, P. (2005). Statistical variability in nonlinear spaces: Application to shape analysis and dt-mri. *PdD thesis, Dept. of Comp. Sci., UNC @ Chapel Hill*. 3

Fletcher, P. T., Lu, C., and Joshi, S. (2003). Statistics of shape via principal geodesic analysis on lie groups. In *CVPR*. CVPR. 44, 102

Fletcher, P. T., Lu, C., Pizer, S. M., and Joshi, S. (2004). Principal geodesic analysis for the nonlinear study of shape. *Transactions on Medical Imaging (TMI)*, 23(8):995–1005. 5, 43, 45, 51, 52

Gerig, G., Styner, M., Weinberger, D., Jones, D., and Lieberman, D. (2001). Shape analysis of brain ventricles using spharm. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, pages 171–178. 5, 8, 13, 70, 132

Gorczowski, K., Styner, M., Jeong, J., Marron, J., Piven, J., Hazlett, H., Pizer, S., and Gerig, G. (2007). Statistical shape analysis of multi-object complexes. In *CVPR*, pages 1–8. CVPR. 2

Han, Q., Lu, C., Liu, S., Pizer, S. M., Joshi, S., and Thall, A. (2004). Representing multi-figure anatomical objects. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1251–1254. ISBI. 6, 72

Han, Q., Merck, D., Levy, J., Villarruel, C., Damon, J. N., Chaney, E. L., and Pizer, S. M. (2007). Geometrically proper models in statistical training. In *IPMI2007, LNCS 4584*, pages 751–762. IPMI, Springer-Verlag. 4, 104

Han, Q., Pizer, S. M., and Damon, J. N. (2006). Interpolation in discrete single figure medial objects. In *MMBIA 2006, Conference on Computer Vision and Pattern Recognition Workshop*, page 85. 27

Han, Q., Pizer, S. M., Merck, D., Joshi, S., and Jeong, J.-Y. (2005). Multi-figure anatomical objects for shape statistics. In Christensen, G. E. and Sonka, M., editors, *IPMI2005, LNCS 3565*, pages 701–712. IPMI, Springer-Verlag. 6

Heimann, T., Münzing, S., Meinzer, H.-P., and Wolf, I. (2007). A shape-guided deformable model with evolutionary algorithm initialization for 3d soft tissue segmentation. In *IPMI*, pages 1–12. 22

Kass, M., Witkin, A., and Terzopoulos, D. (1988). Snakes: Active contour models. *IJCV*. 12, 13, 15, 127

Kelemen, A., Szekely, G., and Gerig, G. (1999). Elastic model-based segmentation of 3-d neuroradiological data sets. *Transactions on Medical Imaging (TMI)*, 18(10). 19

Kimmel, R., Kiryati, N., and Bruckstein, A. (1995). Analyzing and synthesizing images by evolving curves with the osher-sethian method. 16

Loop, C. and DeRose, T. (1990). Generalized b-spline surfaces of arbitrary topology. In *ACM SIGGRAPH*, pages 347–356. 29

Lorensen, W. E. and Cline, H. E. (1987). Marching cubes: A high resolution 3d surface construction algorithm. *Computer Graphics*, 21(4):163–169. 112

Malladi, R., Sethian, J. A., and Vemuri, B. C. (1995). Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158–175. 16

Marr, D. and Nishihara, H. (1978). Visual information processing: Artificial intelligence and the sensorium of sight. *Technology Review*, 81(1). 12

McInerney, T. and Terzopoulos, D. (1996). Deformable models in medical image analysis: A survey. *Medical Image Analysis*. 12, 127

Metropolis, N. and Ulam, S. (1949). The monte carlo method. *Journal of the American Statistical Association*, 44(247):335–341. 113

Mitchell, S. C., Bosch, J. G., Lelieveldt, B. P. F., van der Geest, R. J., Reiber, J. H. C., and Sonka, M. (2002). 3-d active appearance models: Segmentation of cardiac mr and ultrasound images. *Transactions on Medical Imaging (TMI)*, 21(9):1167–1178. 13, 19

O'Neill, B. (1997). *Elementary Differential Geometry*. Harcourt/Acedemic Press, second edition. 40

Pizer, S. M., Fletcher, T., Fridman, Y., Fritsch, D. S., Gash, A. G., Glotzer, J. M., Joshi, S., Thall, A., Tracton, G., Yushkevich, P., and Chaney, E. L. (2003). Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision - Special UNC-MIDAG issue*, 55(2):85–106. 2, 3, 6, 13, 20, 27, 130

Pizer, S. M., Jeong, J.-Y., Lu, C., Muller, K. E., and Joshi, S. C. (2005). Estimating the statistics of multi-object anatomic geometry using inter-object relationships. In *DSSCV*, pages 60–71. 90, 91

Pohl, K. M., Kikinis, R., and III, W. M. W. (2007). Active mean fields: Solving the mean field approximation in the level set framework. In *IPMI*, pages 26–37. 17

Saboo, R., Levy, J., Chaney, E. L., and Pizer, S. M. (2008). Medial models of populations of nearly tubular objects. In *Submission*. 131

Siddiqi, K. and Pizer, S. M., editors (2007). *Medial Representations: Mathematics, Algorithms and Applications*. Springer. In press. In press. 25

Solem, J. E., Aanæ, H., and Heyden, A. (2007). Variational surface interpolation from sparse point and normal data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(1):181–184. 50

Staib, L. H. and Duncan, J. S. (1996). Model-based deformable surface finding for medical images. *Transactions on Medical Imaging*, 15(5):720–731. 6, 19, 20

Stiller, J. (2007). Point-normal interpolation schemes reproducing spheres, cylinders and cones. *Comput. Aided Geom. Des.*, 24(5):286–301. 50

Stough, J., Broadhurst, R. E., Pizer, S. M., and Chaney, E. L. (2007). Regional appearance in deformable model segmentation. In Karssemeijer, N. and Lelieveldt, B., editors, *IPMI2007, LNCS 4584*, pages 532–543. IPMI, Springer-Verlag. 102

Styner, M., Gerig, G., Lieberman, J., Jones, D., and Weinberger, D. (2003). Statistical shape analysis of neuroanatomical structures based on medial models. *Medical Image Analysis (MEDIA)*, 7(3):207–220. 5, 18, 70, 132

Styner, M., Oguz, I., Xu, S., Brechbuhler, C., Pantazis, D., Levitt, J., Shenton, M., and Gerig, G. (2006). Framework for the statistical shape analysis of brain structures using spharm-pdm. In *MICCAI*. 3, 20

Sun, H., Avants, B. B., Frangi, A. F., Ordas, S., Gee, J. C., and Yushkevich, P. A. (2008). Branching medial models for cardiac shape representation. In *ISBI*, pages 1485–1488. 28, 128

Terriberry, T., Damon, J. N., Pizer, S. M., Joshi, S., and Gerig, G. (2007). Population-based fitting of medial shape models with correspondence optimization. In Karssemeijer, N. and Lelieveldt, B. P. F., editors, *IPMI*, pages 700–712. IPMI. 13, 27, 29, 31, 128

Terriberry, T., Joshi, S., and Gerig, G. (2005). Hypothesis testing with nonlinear shape models. In Rosin, P. L. and Marshall, D., editors, *IPMI*, pages 15–26. IPMI. 2, 20

Thall, A. (2004). Deformable solid modeling via medial sampling and displacement subdivision. *PdD thesis, Dept. of Comp. Sci., UNC @ Chapel Hill.* 27, 33, 50, 72

Vapnik, V. N. (1993). *Statistical Learning Theory.* 2

Whitaker, R. T. (1993). Geometry-limited diffusion in the characterization of geometric patches in image. *CVGIP: Image Understanding*, 57:111–120. 16

Xiaoxiao Liu, Ja-Yeon Jeong, J. H. L. R. R. S. E. L. C. and Pizer, S. M. (2008). A coarse-to-fine shape prior for probabilistic segmentations using a deformable m-rep. In *MMBIA 2008, Conference on Computer Vision and Pattern Recognition Workshop.* 97

Yang, J., Staib, L. H., and Duncan, J. S. (2003). Neighbor-constrained segmentation with 3d deformable models. In *IPMI2003, LNCS 2732*, pages 198–209. 3, 5, 17

Yushkevich, P., Fletcher, P. T., Joshi, S., Thall, A., and Pizer, S. M. (2003). Continuous medial representations for geometric object modeling in 2d and 3d. *Image and Vision Computing*, 21(1):17–28. 13, 20, 27, 28, 128

Yushkevich, P. A., Zhang, H., and Gee, J. C. (2005). Parametric medial shape representation in 3-d via the poisson partial differential equation with non-linear boundary conditions. In Christensen, G. E. and Sonka, M., editors, *IPMI2005, LNCS 3565*, pages 162–173. IPMI, Springer-Verlag. 20, 27, 28, 31, 128

Zijdenbos, A. P., Dawant, B. M., Margolin, R. A., and Palmer, A. C. (1994). Morphometric analysis of white matter lesions in mr images: Method and validation. *IEEE Transactions on Medical Imaging*, 13(4):716–724. 115