Committing Medical Image Analysis

Erik B. Dam

August, 2003

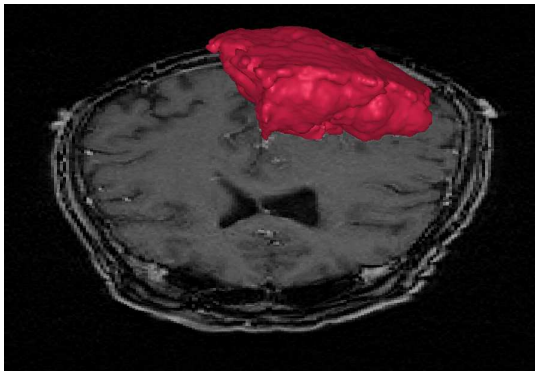

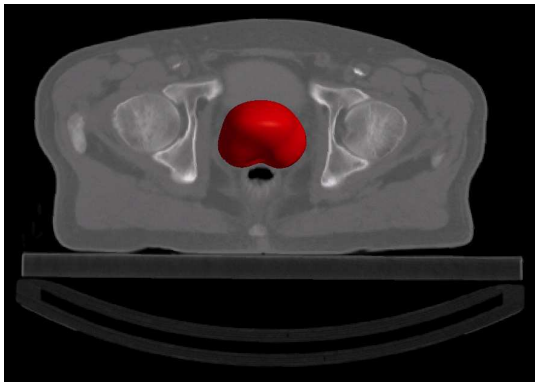

Preface

This constitutes Erik Dam's Ph.D. dissertation. The dissertation is submitted at the IT University of Copenhagen in partial fulfillment of the requirements for the degree of doctor of philosophy.

The work behind the dissertation was delivered in the period from July 2000 to June 2003 under the supervision of professor Mads Nielsen and co-supervision of assistant professor Ole Fogh Olsen from the IT University of Copenhagen. The work was carried out at the IT University of Copenhagen with the exception of a six month visit followed by a one month visit at the University of North Carolina, Chapel Hill, hosted by professor Stephen M. Pizer.

Front page, top image: A segmentation of a brain tumor. The brain scan is used in chapter 4 courtesy of the University Medical Center, Utrecht, the Netherlands. The segmentation was performed using the ∇*Vision* segmentation program [Dam et al., 2003a].

Front page, bottom image: A segmentation of a prostate. The abdominal CT scan is used in chapter 9 courtesy of the Radiation Oncology department, University of North Carolina, Chapel Hill, US. The segmentation was performed using the *Pablo* segmentation program [Fletcher et al., 2002].

The IT University of Copenhagen, Dissertation Series, D-2003-03
ISSN 1602-3536
ISBN 87-7949-044-1

Chapter 9

Prostate Shape Modeling based on Principal Geodesic Analysis Bootstrapping

The use of statistical shape models in medical image analysis is growing due to the ability to incorporate robustly prior organ shape knowledge for tasks such as segmentation, registration, and classification.

Shape models are constructed from collections of segmented organs. Though interaction can ensure correspondence, it also introduces bias and ruins reproducibility — so a high degree of automation is desirable in the training process.

We present a novel shape model construction method via a medial shape representation. The essentially automatic iterative bootstrap method is based on an iterative bootstrap method that alternates between shape representation optimization and analysis of shape mean and variations.

The method is used to create a model from 46 segmented prostates with quantitatively and intutitively good results.

9.1 Introduction

Methods based on analysis of shape variation are becoming widespread in medical imaging. These methods allow incorporation of statistical prior shape knowledge in tasks where the image information alone often is not strong enough to solve the task automatically. The obvious example is the use of deformable models in segmentation, in which the preferred deformations are determined by a statistical shape model. Another important task is shape analysis and classification, in which a statistical shape model offers information for diagnostic methods.

Most statistical shape models consists of a mean shape with deformations. The mean and the corresponding deformations are constructed through statistical analysis of shapes from a collection of training data. Each shape in the training set is represented partially by the chosen shape representation, and analysis of the parameters for the representation gives the mean and variations [Cootes et al., 1995, Cremers et al., 2002a].

The best known model from this class is the *Active Shape Model* (ASM) [Cootes et al., 1995]. Here, the shapes are represented by a *point distribution model* (PDM) with given point-wise correspondence. The mean model is achieved through Procrustes alignment of the shapes followed by mean computation of each point in the model. *Principal component analysis* (PCA) is used to provide the variations.

This work pursues the medial shape representation known as the *m-rep* [Pizer et al., 1996]. The m-rep offers an intuitive representation of the shape by means of the sheet of sampled medial atoms. Compared to PDMs this representation is less simple since the parameter space is not Euclidean but consists of a combination of position, scaling, and orientation parameters. Standard PCA is therefore not applicable. However, the analogue of PCA has been developed for a more applicable space of shape representations. This is the *Principal Geodesic Analysis* (PGA) that applies to shape representations that form Lie groups [Fletcher et al., 2003a, Fletcher et al., 2003b].

A key element in constructing shape models is the representation of the shapes in the training collection. This should be done in a manner that defines/preserves correspondence across the population. For PDMs the simplest method is manual selection of the boundary points by an expert of the specific anatomical structure. In 2D this is a time-consuming and tedious process — in 3D it is even worse. However, this process can be automated. The approach by Davies [Davies et al., 2002] starts by generating boundary points from a spherical harmonics shape representation. This set of boundary points and their correspondences are then optimized through a *Minimum Description Length* (MDL) approach.

This work presents an essentially automatic shape modeling method. The core is a fully automatic bootstrap process that iteratively optimizes the shape model on a training collection and then derives the PGA mean and modes of deformation. Through the bootstrap iterations, the PGA mean and variations are optimized to allow automatic fitting of all shapes in the training collection.

The flavor of this work resembles the MDL method in [Davies et al., 2002]. The main difference is that the MDL approach starts the optimization process from representations with good training shape fit and poor correspondence. The MDL process then keeps the shape fits while optimizing the correspondence. The PGA bootstrap starts from a generative model with explicit correspondence but with poor fit to the individual training shapes. The bootstrap process then keeps the correspondence while optimizing the fit to the training shapes.

There exists another method for generating an m-rep mean model from a set of training shapes [Styner & Gerig, 2001] that uses a spherical harmonics representation followed by generation of the mean medial sheet from pruned Voronoi skeletons. Our approach is intuitively cleaner since m-rep is the only shape representation in play. Furthermore, our approach provides modes of variation as well as the shape mean.

We evaluate the presented PGA bootstrap method for the task of constructing a shape model for a population of prostates. The training collection consists of 46 cases where the prostates were segmented in the course of prostate cancer external-beam radiation treatment. Especially in CT scans with slice thickness 2mm or larger, the boundaries of the prostate have low contrast — therefore, prior knowledge in a statistical shape model is essential to making automatic segmentation possible. This prostate shape model is a key step towards a pelvis multi-object shape model that hopefully will achieve this goal.

Furthermore, current research aims at using the shape model to analyze the prostates in order to make a shape classification used for the radiation treatment planning. This is concentrated on the problematic *saddle-back* cases where the prostate reaches around the rectum — that complicates giving radiation to the prostate without hitting the rectum.

The contributions of this work are twofold: a) The presented PGA bootstrap method that allows essentially automatic generation of a shape model with mean and corresponding main modes of variation. b) The resulting prostate model that will be central in segmentation and analysis of prostates and eventually allow better radiation treatment planning.

9.2 The UNC Pelvis Collection

In radiation treatment, accurate segmentation of the prostate and surrounding organs is vital. Low image contrast across the prostate boundary makes this a difficult task.

The segmentation programs, *MASK* [Tracton et al., 1994] and *anastruct_editor*, from the PLan-UNC suite of radiotherapy treatment tools developed at UNC-CH Radiation Oncology, have slice-based contour drawing tools and visualization of reconstructed sagittal and coronal views. Both programs have interactive 12-bit intensity-windowing, which is required to find and draw both the prostate boundaries across from the bladder and the prostate's apex (the superior tip). The contours are scan-converted to labelled images, which introduces less than one pixel of error not significant to this shape study. Prostatic fat is included in the prostate's shape, as is seen in clinical practise, both because of the difficulty of finding the border between these and the prostate and the chance that these will contain significant counts of cancer cells. Seminal vesicles are excluded from the prostate.

Clinical contours are used but adjusted when obvious errors were found, such as missing contours, overlapping contours (eg. between the rectum and prostate), or just sloppy contouring – all of these shortcuts are due to clinical time constraints and are not perceived to affect clinical care but can affect shape studies.

The ungated CT scans are acquired from non-immobilized supine patients at UNC Healthcare (Chapel Hill, NC, USA) and Western Wake Radiology (Cary, NC, USA) on Siemens Somotom 4+ scanners without administering contrast agents. Note that while the prostate is quite hard, the multi-object statistics eventually produced will be sensitive to prostate shifts based on the patient's position and the CT couch shape (flat vs. rounded) because of the surrounding tissues' malleability. These effects should not affect this shift-invariant prostate shape analysis.

Figure 9.1: *Sagittal slices of the manual segmentations of rectum, prostate, and bladder from two cases in the UNC pelvis collection.*

Retrospective patient images are selected from the patient archives based on technical criteria, such as adequate image quality and anatomical coverage (the entire bladder down through the prostate apex), as well as shape and anatomical considerations such as very large bladders, prosthetic hips, or surgical procedures proximal to the prostate, yielding "normal cancerous" prostates.

The collection has 46 sets with manual segmentations for prostate, bladder, and rectum. All cases are diagnosed with prostate cancer so the resulting shape model will not necessarily model prostates in general. For instance, an increase of the size of the prostate is common for prostate cancer patients. However, since the shape model is to be used for segmentation and analysis of patients diagnosed with prostate cancer, this bias towards cancerous prostates is desirable.

The volumes of the prostates varies from $12cm^3$ to $144cm^3$. Figure 9.1 illustrates the large variation in shape.

9.3 Medial Shape Representation: m-rep

We use a medial representation, m-rep, to model shape. Here, we briefly review the geometry of m-reps and the deformable m-rep framework for image segmentation [Pizer et al., 2003a, Joshi et al., 2002].

9.3.1 m-rep Geometry Overview

The shape representation we use is based on the medial axis of Blum [Blum & Nagel, 1978]. In this framework, a 3D geometric object is represented as a set of connected continuous medial sheets, which are formed by the centers of all spheres that are interior to the object and tangent to the object's boundary at two or more points. Here we focus on 3D objects that can be represented by a single medial figure.

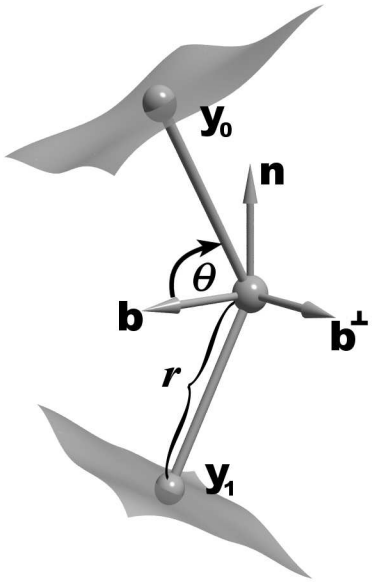

Figure 9.2: *Medial atom with a cross-section of the boundary surface it implies.*

We sample the medial sheet $\mathcal{M}$ over a spatially regular lattice. Each sample point also includes first derivative information of the medial position and radius. The elements of this lattice are called *medial atoms*. A medial atom (figure 9.2) is defined as a 4-tuple $\mathbf{m} = \{\mathbf{x}, r, \mathbf{F}, \theta\}$, consisting of: $\mathbf{x} \in \mathbb{R}^3$, the center of the inscribed sphere, $r \in \mathbb{R}^+$, the local width defined as the radius of the sphere, $\mathbf{F} \in \mathbf{SO}(3)$ an orthonormal local frame parameterized by $(\mathbf{b}, \mathbf{b}^\perp, \mathbf{n})$, where $\mathbf{n}$ is the normal to the medial sheet, $\mathbf{b}$ is the direction in the tangent plane of the fastest narrowing of the implied boundary sections, and $\theta \in [0, \pi)$ the object angle determining the angulation of the implied sections of boundary relative to $\mathbf{b}$. The medial atom implies two opposing boundary points, $\mathbf{y}_0, \mathbf{y}_1$, with respective boundary normals, $\mathbf{n}_0, \mathbf{n}_1$, which are given by

$$\mathbf{n}_0 = \cos(\theta)\mathbf{b} - \sin(\theta)\mathbf{n}, \qquad \mathbf{n}_1 = \cos(\theta)\mathbf{b} + \sin(\theta)\mathbf{n},$$
$$\mathbf{y}_0 = \mathbf{x} + r\mathbf{n}_0, \qquad \mathbf{y}_1 = \mathbf{x} + r\mathbf{n}_1.$$

Given an m-rep figure, we fit a smooth boundary surface to the model. We use a subdivision surface method [Thall, 2002] that interpolates the boundary positions and normals implied by each atom.

9.3.2 Segmentation using m-reps

Following the deformable models paradigm, an m-rep model $\mathbf{M}$ is deformed into an image I by optimizing an objective function, which we define as

$$F(\mathbf{M}, I) = L(\mathbf{M}, I) + \alpha\, G(\mathbf{M}).$$

The function L, the *image match*, measures how well the model matches the image information, while G, the *geometric typicality*, gives a prior on the possible variation of the geometry of the model. The relative importance of the two terms is weighted by $\alpha \geq 0$.

This objective function is optimized in a multiscale fashion. That is, it is optimized over a sequence of transformations that are successively finer in scale. Here we will only be concerned with two levels of scale: the figural level, and the medial atom level. At the figural level the transformation we use is a similarity transformation plus an elongation of the entire figure. At the atom level each medial atom is independently transformed by a similarity plus a rotation of the object angle.

m-rep models are fit to binary segmentation images of the prostates. These binary images are blurred slightly to smooth the objective function, which is optimized with a conjugate gradient method. The image match term of the objective function is computed as a correlation with a Gaussian derivative kernel in the normal direction to the object boundary:

$$L(\mathbf{M}, I) = \int_{\mathcal{B}(\mathbf{M})} \int_{-\epsilon}^{\epsilon} \partial_t G(t)\, I\left(\mathbf{s} + (t/r)\mathbf{n}\right)\, dt\, d\mathbf{s},$$

where $\mathbf{s}$ is a parameterization of the boundary $\mathcal{B}(\mathbf{M})$, $\partial_t G$ is the Gaussian derivative kernel, r is the radius function, and $\mathbf{n}$ is the boundary normal.

The geometric typicality term is defined as

$$G(\mathbf{M}) = (1-\beta)\, P(\mathbf{M}) + \beta\, N(\mathbf{M}), \tag{9.1}$$

where $\beta \in [0,1]$ is a weighting term. The function P measures the change in the boundary from the previous level of scale:

$$P(\mathbf{M}) = -\int_{\mathcal{B}(\mathbf{M})} \frac{||\mathbf{s} - \mathbf{s}_0||^2}{r^2} d\mathbf{s},$$

where $\mathbf{s}_0$ is the initial position of the boundary at this scale level. The function N seeks to keep medial atoms in the same relationship with their neighboring atoms. It is defined as

$$N(\mathbf{M}) = -\int_{\mathcal{B}(\mathbf{M})} \frac{||\mathbf{s} - \mathbf{s}'||^2}{r^2} d\mathbf{s},$$

where now $\mathbf{s}'$ is the boundary surface of the model in which the current medial atom is in the position predicted by its neighbors. The neighbor term is only used at the atom scale level, i.e., $\beta = 0$ during the figural level. The role of the neigbor penalty term is to keep the shape nice locally — comparable to the curvature term in active contour models.

9.4 Shape Modeling

The goal is to arrive at a shape model that describes the shape variation within the population of prostates. A shape model is a parametric shape representation with predefined rules for deformation that allow the model to represent the class of shapes encountered. In medical image analysis the shape classes are typically defined by shapes of organs or other anatomical structures. This introduces some basic modeling trade-offs:

Compactness vs Accuracy
: The model should be a more compact representation than the basic representation of the organ (typically a binary image). This introduces a trade-off between compactness and precision in the representation. The desire for compactness suggests coarse sampling of the basic elements of both the underlying shape representation (boundary points, medial points, control points in splines, spherical harmonic coefficients etc.), and the possible deformations (similarity transformation, PCA modes etc.). The desire for accuracy suggests finer sampling.

Generality vs Specificity
: The deformations should be flexible enough to allow the shape model to fit the organs up to the precision possible with a specific compactness/sampling. This ensures the generality of the model. The opposing property is specificity, that states that the model should not deform into shapes not encountered in the organs. This trade-off inspires a statistical approach that allows a soft transition between likely and un-likely shapes.

Correspondence vs Accuracy
: Another key property of shape models is the ability to give an explicit or implicit coordinate system on the shapes that offers correspondence of locations among the shapes. This introduces yet another requirement on the allowed deformations of the shape model. Being able to fit the individual organs shape is not sufficient — the model must also ensure that anatomically corresponding locations on the organs are equipped with corresponding locations given by the coordinate system of the shape model. This suggest that the shape model should restrict large deformations that violate correspondence — and thereby the attainable accuracy is limited.

9.4.1 First Attempt: The Potato Model

The segmentation program *Pablo* provides a user interface that allows construction of m-rep models and optimization of the parameters such that the constructed model is fitted to a specific training case [Pizer et al., 2003a].

The Potato in figure 9.3 is such a handcrafted m-rep model based on inspection of a subset of the prostate collection and some experimentation. If we can automatically deform this model into all the prostates in the training collection, we actually have the desired shape model.

Batch Optimization of m-reps

For this shape modeling we developed a batch non-interactive back-end to Pablo. It is of little theoretical interest, but without this automatic fitting program, our results would not be realistically possible.

The fitting program reads m-rep optimization constraints from an options file and performs the corresponding m-rep optimization steps. The process is completely free from interaction. The initial hand placement of the model is performed automatically by translating to the center of gravity and scaling to the volume of the relevant training case (represented as a binary image).

To fit the Potato in all the training prostates, the batch optimizer is run with $\alpha = 0.1$ and $\beta = 0.75$ — high confidence in the segmentations, and low importance of geometric typicality.

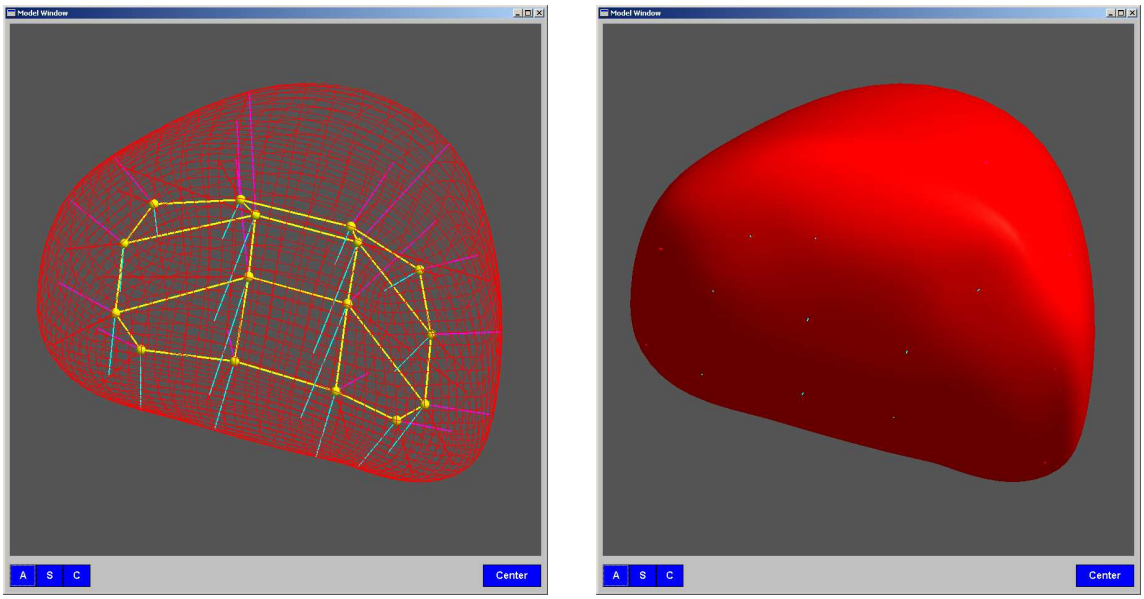

Figure 9.3: *The Potato m-rep model.* ***Left****: The medial grid with the implied boundary as wireframe.* ***Right****: The model as a solid surface. The viewpoint is from above, behind the prostate.*

Rating the Potato

In order to determine whether the Potato can be deformed satisfactorily into the training prostates we look at the image match values for the 46 cases. Recall that the match value is a normalized correlation measure in the range -1 to 1 with approximately 0.95 as the practical maximal value. Furthermore, heuristic experience from other organs show that values above 0.80 are quite good.

The Potato match values for the 46 cases are in the range 0.60 – 0.84 with a mean of 0.75. This is not acceptable for our task. Figure 9.4 illustrates the worst and best case.

The geometric penalty prevents the model from deforming enough to fit the worst cases satisfactorily. Therefore the model certainly is not general enough, and the Potato cannot be considered a good prostate prototype.

However, even though the batch optimization of the Potato does not give satisfying shape representations of all the prostates in the training collection, it does ensure a rough fit in most cases. In the following we present the theory that allows us to use these rough fits to improve the Potato.

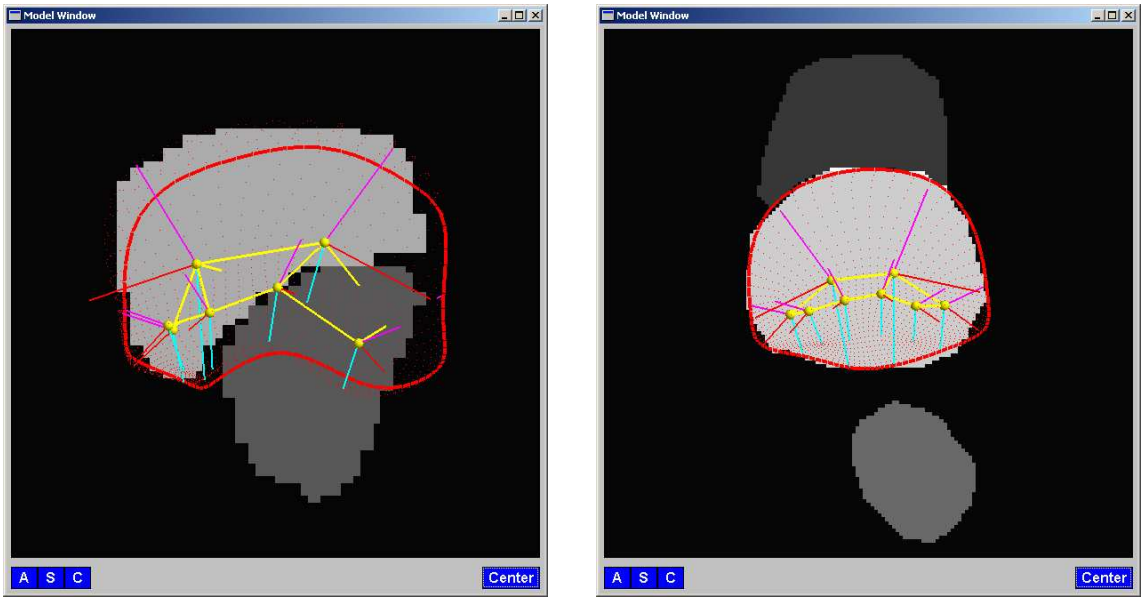

Figure 9.4: *Axial slices from the worst (left, image match 0.60, volume 27 cm^3) and best (right, image match 0.84, volume 127 cm^3) fitting cases for the deformed Potato model.*

9.5 Principal Geodesic Analysis

Principal geodesic analysis (PGA), introduced in [Fletcher et al., 2003b], is a generalization of principal component analysis (PCA) to curved manifolds. It is shown in [Fletcher et al., 2003a, Fletcher et al., 2003b] that m-rep models form a Lie group, and the necessary algorithms for computing the mean and PGA of a collection of m-rep models are given. We review these results briefly here.

9.5.1 Lie Groups

First we present a brief overview of Lie groups (see [Duistermaat & Kolk, 2000] for a detailed treatment). A Lie group G is a differentiable manifold that also forms an algebraic group, where the two group operations,

$$\begin{aligned} \mu : (x, y) \mapsto xy \quad &: \quad G \times G \to G \quad && \textit{Multiplication,} \\ \iota : x \mapsto x^{-1} \quad &: \quad G \to G \quad && \textit{Inverse,} \end{aligned}$$

are differentiable mappings.

Let e denote the identity element of a Lie group G. The tangent space at e, T_eG, forms a Lie algebra, which we will denote by $\mathfrak{g}$. The exponential map, $\exp : \mathfrak{g} \to G$, provides a method for mapping vectors in the tangent space T_eG into G. Given a vector $\mathbf{v} \in \mathfrak{g}$, the point $\exp(\mathbf{v}) \in G$ is obtained by flowing to time 1 along the unique one-parameter subgroup emanating from e with initial velocity vector $\mathbf{v}$. When the Lie group is given a compatible Riemannian metric, this one-parameter subgroup is the unique geodesic at e with velocity $\mathbf{v}$.

The exponential map is a diffeomorphism on a neighborhood of 0 in $\mathfrak{g}$ to a neighborhood of e in G. The inverse of the exponential map is called the log map. The geodesic distance between two points $g, h \in G$ is given by $||\log(g^{-1}h)||$.

As shown in [Fletcher et al., 2003b], the set of all medial atoms forms a Lie group $M = \mathbb{R}^3 \times \mathbb{R}^+ \times \mathbf{SO}(3) \times \mathbf{SO}(2)$, which we call the *medial group*. Likewise, the set of all m-rep models containing n medial atoms forms a Lie group M^n, i.e., the direct product of n copies of M.

9.5.2 m-rep Means

The Riemannian distance between m-rep models $\mathbf{M}_1, \mathbf{M}_2 \in M^n$ is given by

$$d(\mathbf{M}_1, \mathbf{M}_2) = ||\log(\mathbf{M}_1^{-1}\mathbf{M}_2)|| \tag{9.2}$$

Thus the intrinsic mean of a collection of m-rep models $\mathbf{M}_1, \ldots, \mathbf{M}_N$ is the minimizer of the sum-of-squared geodesic distances:

$$\mu = \arg\min_{\mathbf{M} \in M^n} \sum_{i=1}^{n} ||\log(\mathbf{M}_i^{-1}\mathbf{M})||^2$$

As shown in [Fletcher et al., 2003b], the mean model may be computed by the following iterative gradient descent algorithm (algorithm 4).

Algorithm 4 *m-rep Mean*

Input: $\mathbf{M}_1, \ldots, \mathbf{M}_n \in M^n$, m-rep models
Output: $\mu \in M^n$, the intrinsic mean
  $\mu = \mathbf{M}_1$
  Do
    $\Delta\mathbf{M}_i = \mu^{-1}\mathbf{M}_i$
    $\Delta\mu = \exp(\frac{1}{n}\sum_{i=1}^{n}\log(\Delta\mathbf{M}_i))$
    $\mu = \mu\Delta\mu$
  While $||\log(\Delta\mu)|| > \epsilon$.

9.5.3 PGA of m-reps

Principal components of Gaussian data in $\mathbb{R}^n$ are defined as the projection onto the linear subspace through the mean spanned by the eigenvectors of the covariance matrix. If we consider a general manifold, the counterpart of a line is a geodesic curve, that is, a curve with minimal length between two points. In the Lie group M^n geodesics can be computed via the exponential map. Given a tangent vector $\mathbf{v}$ in the Lie algebra $\mathfrak{m}^n$, the geodesic starting at the identity with initial velocity $\mathbf{v}$ is given by $\gamma : \mathbb{R} \to M^n$, where $\gamma(t) = \exp(t\mathbf{v})$. Similarly, the curve $x \cdot \gamma(t) = x \cdot \exp(t\mathbf{v})$ is a geodesic starting at the point $x \in M^n$.

Algorithm 5 *Algorithm: m-rep PGA*

Input: m-rep models, $\mathbf{M}_1, \ldots, \mathbf{M}_N \in M^n$
Output: Principal directions, $\mathbf{u}^{(k)} \in \mathfrak{m}^n$
Variances, $\lambda_k \in \mathbb{R}$
$\mu =$ mean of $\{\mathbf{M}_i\}$
$\mathbf{x}_i = \log(\mu^{-1}\mathbf{M}_i)$
$\mathbf{S} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{x}_i\mathbf{x}_i^T$
$\{\mathbf{u}^{(k)}, \lambda_k\} =$ eigenvectors/eigenvalues of $\mathbf{S}$.

As shown in [Fletcher et al., 2003a], the covariance structure of a Gaussian distribution on M^n may be approximated by a covariance matrix $\mathbf{\Sigma}$ in the Lie algebra $\mathfrak{m}^n$. The eigenvectors of this covariance matrix correspond via the exponential map to geodesics on M^n, called *principal geodesics*. The principal geodesic analysis (PGA) on a population of m-rep figures, $\mathbf{M}_i, \ldots, \mathbf{M}_N \in M^n$, is thus computed by algorithm 5.

Analogous to linear PCA models, we may choose a subset of the principal directions $\mathbf{u}^{(k)} \in \mathfrak{m}^n$ that is sufficient to describe the variability of the m-rep shape space. New m-rep models may be generated within this subspace of typical objects. Given a set of coefficients $\{\alpha_1, \ldots, \alpha_l\}$, we generate a new m-rep model by

$$\mathbf{M} = \mu \exp\Big(\sum_{k=1}^{l} \alpha_k \mathbf{u}^{(k)}\Big),$$

where α_k is chosen to be within $[-3\sqrt{\lambda_k}, 3\sqrt{\lambda_k}]$.

9.6 Shape Model Bootstrapping

The shape model bootstrapping method is now quite simple. The batch fitting process is used to give rough representations of each shape. The PGA is used to generate the mean model from the 46 fitted models. This mean model is then used as the starting model in the next iteration to fit the shapes using the batch fitting process, and the bootstrapping method is iterated.

The underlying assumption is that the mean of the roughly fitting models will be a better prototype than the initial Potato. As the bootstrap iterations progress the generated mean models will hopefully converge to a good prototype.

9.6.1 Bootstrapping the Potato using PGA Mean

The graph in figure 9.5 shows the development of the image matches during bootstrapping. The image matches become excellent as the bootstrap progresses — even the worst match is good. The resulting mean model can be deformed automatically into the shapes in the training collection, so we now have a prostate shape model as desired. Since the allowed deformations are exactly the same as in the initial attempt in section 9.4.1, the improved performance is exclusively due to the resulting mean shape being a better prostate prototype — as expected.

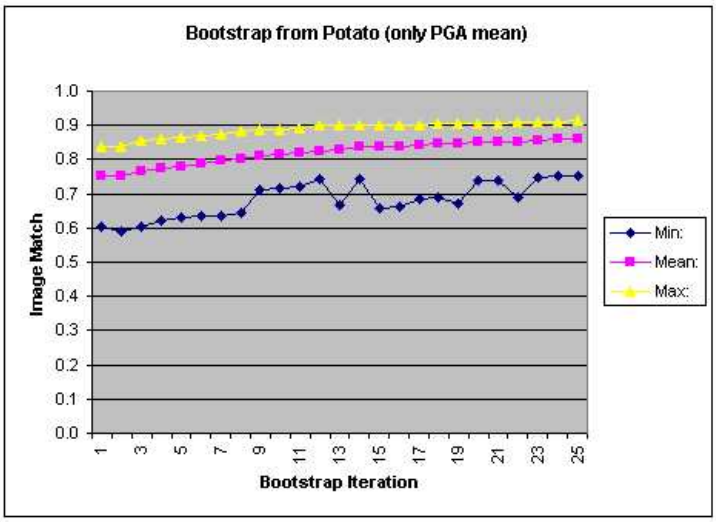

Figure 9.5: *The evolution of the worst, best, and mean image match for the 46 cases during bootstrap starting from the Potato.*

However, the method started from a reasonably good initial model being the manually constructed Potato. For the prostate, this is not a problem. However, it is not desirable if we need to divine a new suitable vegetable from which to start the bootstrapping for each new organ to be modeled.

9.6.2 Bootstrapping the Generic using PGA Mean

In order to see how dependent the bootstrap method is on the initial model, alternative m-rep models with the same 4x4 atom grid were used. The *Generic* is the default 4x4 slab model that Pablo generates as a starting model for building handcrafted models. Figure 9.6 shows how the bootstrap is virtually independent of the choice of starting model — except for the number of atoms that reflects the compactness vs accuracy issue.

9.6.3 Bootstrapping using PGA Mean and Modes

The principal geodesic analysis provides the shape mean and well as the principal modes of variation. Above, we only use the mean in the bootstrap. However, the batch fitting process can also use the PGA modes in the figural level transformation in place of the elongation.

The expectation is that these trained, global deformations steer the model to the desired image match in fewer bootstrap iterations. Figure 9.7 shows exactly that expected behavior where the 10 top PGA modes are used.

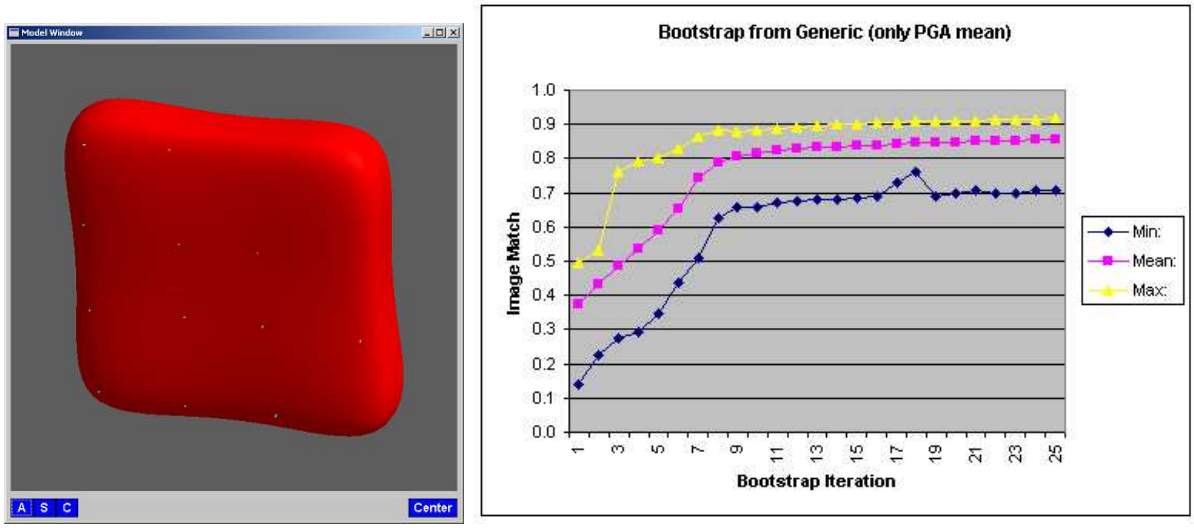

Figure 9.6: *The Generic initial model and the resulting bootstrap image match evolution.*

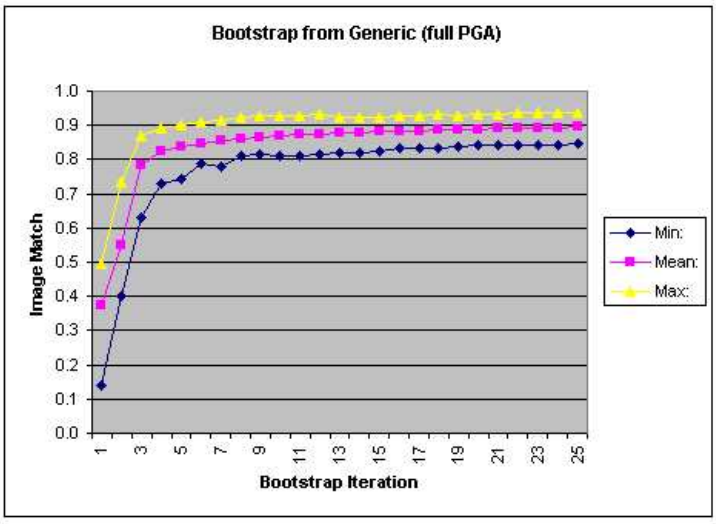

Figure 9.7: *The bootstrap evolution starting from the Generic using global PGA modes in the figural level optimization.*

9.6.4 Convergence

In this work, we address the question of convergence superficially. The image match values above appear to be converging. Convergence of the PGA mean and modes is considered future work. Furthermore, we may want to constrain the evolution of the model such that the converging model has specific correspondence properties. Our expectation is that eventually the method will converge to the same resulting mean model independent of the choice of starting model and optimization options — using different paths through model parameter space.

Here we use the heuristic approach of running the bootstrap until the image match values cease to improve significantly. We therefore use the resulting model from 10 bootstrap iterations starting from the Generic using PGA deformations in the optimization (see figure 9.7).

9.6.5 Exclusion of Outliers

In particular during the early bootstrap iterations, the optimization process reaches poor image match values for some training cases. This is simply due to too much difference between the initial model and these training prostate shapes. As a result, the optimized model then does not correspond to a prostate shape within a reasonable precision.

Therefore, it can be argued that these non-prostate shapes should be excluded from the calculations of the PGA mean and modes. Much like outliers are often identified and excluded during statistical methods in general. This identification could simply be done by defining an image match threshold based on e.g. the variance of these values. In this work we experimented superficially with outlier exclusion and concluded that the PGA bootstrap method performs robustly without this extra step.

9.6.6 Resulting Prostate Shape Model

The resulting Prostate shape model then consists of the mean shape and the deformations illustrated in figure 9.8. The 10 modes of variation include 95% of the variation in the training collection. This ensures little need for atom optimization in the segmentation process. The automatic fitting achieves image matches in the range 0.81–0.93 with mean 0.87. That this is significantly better than the original attempt with the Potato model can be seen in figure 9.9.

The shape model is trained on binary images — therefore the resulting model is only evaluated for segmenting binary images. However, in the full m-rep segmentation framework, the shape model is combined with profile models for the local boundary instead of the just using the Gaussian derivative profile (as done in [Rao, 2003] for instance).

Apart from being directly applicable for segmentation, the shape model — and especially the condensed PGA parameterization — is also directly applicable for shape classification.

Furthermore, the hierarchical m-rep optimization framework also allow easy generalization of the model building method. The work presented here is limited to single-figure m-rep models. The shape model method could easily be extended to multi-object multi-object ensembles due to the modular optimization method where each individual step could be modeled using the PGA bootstrap approach introduced here.

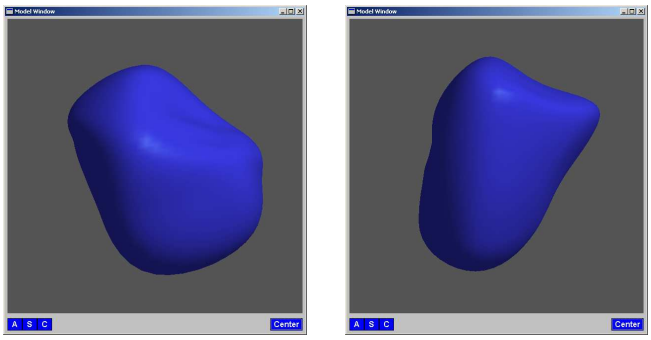

Mode 1: Sagittal view (from the side) possibly showing effects of varying pressure from bladder (is top right).

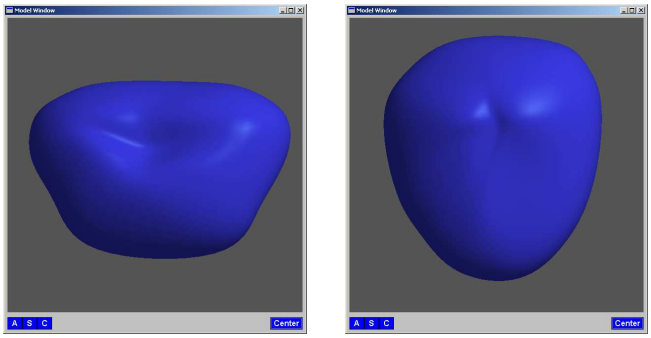

Mode 2: Coronal view: Laurel/Hardy deformation.

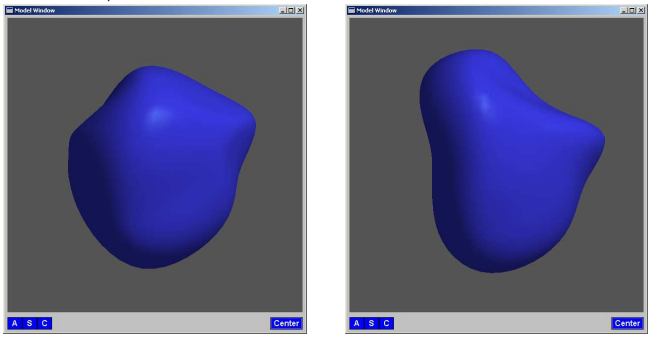

Mode 3: Sagittal: again possibly results of bladder pressure.

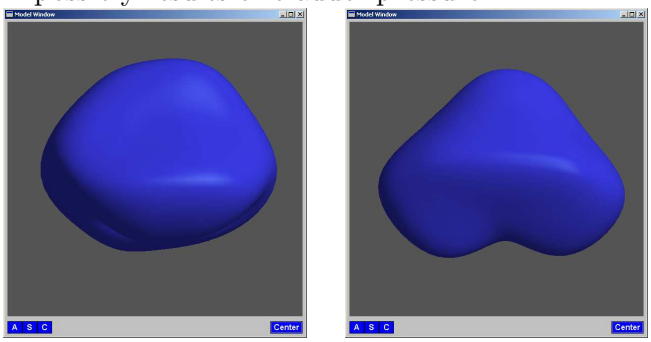

Mode 4: Axial view: Saddle-back where the prostate curves around the rectum — a known problematic behavior that complicates radiation treatment.

Figure 9.8: *The four primary modes of deformation in the Prostate model. Each mode is illustrated by ± 1.5 standard deviation images.*

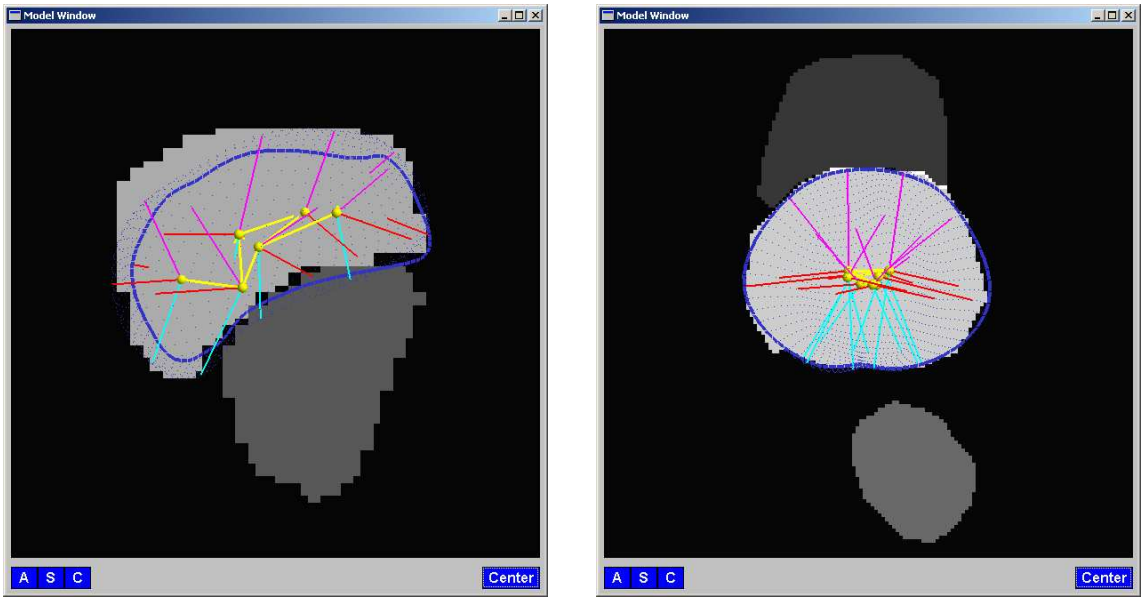

Figure 9.9: *The Prostate model on the cases from figure 9.4 with image matches 0.85 and 0.91 (was 0.60 and 0.84). Since left is a worst case this is highly satisfying.*

9.7 Conclusion

We present a novel shape model construction method using a medial shape representation. The method is essentially automatic based on an iterative bootstrap method that alternates between shape representation optimization and principal geodesic analysis of shape mean and variations. The method constructs an m-rep shape model consisting of a mean and corresponding main modes of variation.

The non-automatic step is the choice of sampling in the medial sheet. We have chosen a 4x4 atom grid that appears to be a suitable compromise between compactness and accuracy.

The method is evaluated through construction of a prostate shape model from a training collection of 46 manually segmented prostates. The resulting model has good quantitative performance. In addition, the modes of variation show deformations that corresponds intuitively well with known prostate behavior. We are especially pleased with the presence of the saddle-back variation in the fourth mode.

Future work is centered on ensuring desirable convergence of the shape model in the bootstrap iterations. Furthermore, we plan to evaluate the method against the method that uses the MDL approach to generate an ASM [Davies et al., 2002]. Central points to evaluate are compactness, correspondence, and legality (how likely are illegal models).

Finally, we look forward to applying the method on other anatomical organs (among others kidneys, hearts and various brain structures).

Acknowledgements

We sincerely thank Per Halverson at Western Wake Radiology (Cary, NC, USA) for supplying pelvis scans for the collection.

Recycling in this Chapter

This chapter is an edited version of [Dam et al., 2003b].

Bibliography

[Abbasi et al., 1999] S. Abbasi, F. Mokhtarian, & J. Kittler. *Curvature Scale Space Image in Shape Similarity Retrieval*. Multimedia Systems, 7:467–476, 1999.

[Blum & Nagel, 1978] H. Blum & R. Nagel. *Shape Description Using Weighted Symmetric Axis Features*. Pattern Recognition, 10(3):167–180, 1978.

[Brox et al., 2003] Brox, Welk, Steidl, & Weickert. *Equivalence Results for TV Diffusion and TV Regularization*. In *Scale Space Methods in Computer Vision*, number 2695 in LNCS, 2003.

[Catté et al., 1992] F. Catté, P.-L. Lions, J.-M. Morel, & T. Coll. *Image Selective Smoothing and Edge Detection by Nonlinear Diffusion*. SIAM Journal of Numerical Analysis, 29:182–193, 1992.

[Cocosco et al., 1997] C.A. Cocosco, V. Kollokian, R.K.-S. Kwan, & A.C. Evans. *BrainWeb: Online Interface to a 3D MRI Simulated Brain Database*. In *Proceedings of 3rd International Conference on Functional Mapping of the Human Brain*, volume 5 of NeuroImage, page 425, May 1997. http://www.bic.mni.mcgill.ca/brainweb/.

[Collins et al., 1998] D.L. Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, & A.C. Evans. *Design and Construction of a Realistic Digital Brain Phantom*. IEEE TMI, 17, June 1998. http://www.bic.mni.mcgill.ca/brainweb/.

[Cootes et al., 1995] T.F. Cootes, C.J. Taylor, D.H. Cooper, & J. Graham. *Active Shape Models: Their Training and Application*. CVIU, (1):38–59, 1995.

[Cremers et al., 2002a] D. Cremers, T. Kohlberger, & C. Schnörr. *Nonlinear Shape Statistics in Mumford-Shah Based Segmentation*. In *7th European Conference on Computer Vision*, volume 2351 of Springer LNCS, 2002.

[Cremers et al., 2002b] D. Cremers, F. Tischhuser, J. Weickert, & C. Schnörr. *Diffusion snakes: introducing statistical shape knowledge into the Mumford–Shah functional*. International Journal of Computer Vision, 50, 2002.

[Crouch et al., 2003] J. Crouch, S.M. Pizer, E.L. Chaney, & M. Zaider. *Medial Techniques to Automate Finite Element Analysis of Prostate Deformation*. Submitted to IEEE Transactions on Medical Imaging, 2003. http://midag.cs.unc.edu/pubs/papers/TMI03-Crouch-FEM.pdf.

[Dam & Letteboer, 2003] Erik Dam & Marloes Letteboer. *Non-linear Diffusion in 3D for Interactive Segmentation of Brain Tissue.* In *Submitted to MICCAI 2003*, 2003.

[Dam & Nielsen, 2000] Erik Dam & Mads Nielsen. *Non-Linear Diffusion for Interactive Multi-scale Watershed Segmentation.* In *Medical Image Computing and Computer-Assisted Intervention — MICCAI 2000*, volume 1935 of Lecture Notes in Computer Science, pages 216—225. Springer, October 2000.

[Dam & Nielsen, 2001] Erik Dam & Mads Nielsen. *Exploring Non-Linear Diffusion: The Diffusion Echo.* In Michael Kerckhove, editor, *Scale-Space Theories in Computer Vision*, Lecture Notes in Computer Science. Springer, 2001.

[Dam & ter Haar Romeny, 2003] Erik Dam & Bart ter Haar Romeny. *Front End Vision and Multi-Scale Image Analysis*, chapter Chapters 13–15: Deep Structure I, II & III. Kluwer, 2003.

[Dam, 2000] Erik Dam. *Evaluation of Diffusion Schemes for Watershed Segmentation.* Master's thesis, University of Copenhagen, 2000. Technical report 2000/1 on http://www.diku.dk/research/techreports/2000.htm.

[Dam, 2003] Erik B. Dam. *Billedanalyse og strålebehandling.* At the science portal of the Danish Radio, 2003. http://www.dr.dk/videnskab/artikler/it/billedanalyse.asp.

[Dam et al., 2000] Erik Dam, Peter Johansen, Ole Fogh Olsen, Andreas Thomsen, Tron Darvann, Andy B. Dobrzeniecki, Nuno V. Hermann, Noriyuki Kitai, Sven Kreiborg, Per Larsen, & Mads Nielsen. *Interactive Multi-Scale Segmentation in Clinical Use.* In *European Congress of Radiology 2000*, March 2000.

[Dam et al., 2002] Erik Dam, Dr. Stephen M. Pizer, Dr. Julian Rosenman, & Gregg Tracton. *M-Rep Based Classification of Organ Conformations for Radiation Treatment Planning of Prostate Cancer.* The 17th Annual UNC Radiology Research Symposium, 2002. Reviewed abstract.

[Dam et al., 2003a] Dam, Fogh Olsen, Johansen, Lillholm, Nielsen, & Thomsen. *Nabla Vision*, 2000–2003. http://www.itu.dk/image/nablavision.

[Dam et al., 2003b] Erik Dam, P. Thomas Fletcher, Stephen M. Pizer, Gregg Tracton, & Julian Rosenman. *Prostate Shape Modeling based on Principal Geodesic Analysis Bootstrapping.* In *Submitted to ICCV workshop VLSM*, 2003.

[Dam et al., 2003c] Erik Dam, Ole Fogh Olsen, & Mads Nielsen. *Approximating Non-linear Diffusion.* In *Scale-space Methods in Computer Vision, Proc. from the 4th int. conference*, volume 2695 of LNCS. Springer, 2003.

[Damon, 1997] James Damon. *Generic Properties of Solutions to Partial Differential Equations.* Arch. Rational Mech. Anal., (140):353–403, 1997.

[Damon, 2003] James Damon. *Smoothness and Geometry of Boundaries Associated to Skeletal Structures.* Technical report, University of North Carolina, Chapel Hill, 2003. http://midag.cs.unc.edu/pubs/papers/Damon_SkelStr_I.pdf, http://midag.cs.unc.edu/pubs/papers/Damon_SkelStr_II.pdf, http://midag.cs.unc.edu/pubs/papers/Damon_SkelStr_III.pdf.

[Davies et al., 2002] R.H. Davies, C.J. Twining, T.F. Cootes, J.C. Waterton, & C.J. Taylor. *A Minimum Description Length Approach to Statistical Shape Modeling.* IEEE Transactions on Medical Imaging, 21(5), 2002.

[Demirci et al., 2003] M. Fatih Demirci, Ali Shokoufandeh, Yakov Keselman, Sven Dickinson, & Lars Bretzner. *Many-to-Many Matching of Scale-Space Hierarchies using Metric Embedding.* In *Scale Space Methods in Computer Vision, Proceedings from 4th Int. Conference*, volume 2695 of LNCS, 2003.

[Duistermaat & Kolk, 2000] J. J. Duistermaat & J. A. C. Kolk. *Lie Groups.* Springer, 2000.

[Duits, 2003] Florack Platel Duits, Felsberg. *α Scale Spaces on a Bounded Domain.* In *Scale Space Methods in Computer Vision*, number 2695 in LNCS, 2003.

[Fischl & Schwartz, 1997] B. Fischl & E.L. Schwartz. *Learning an Integral-Equation Approximation to Nonlinear Anisotropic Diffusion in Image-Processing.* PAMI, 19(4), 1997.

[Fletcher et al., 2002] Fletcher, Gash, Joshi, Liu, Lu, Pizer, Styner, Thall, Tracton, & Yushkevich. *Pablo.* http://midag.cs.unc.edu, 2002. Pablo is a product of MIDAG, University of North Carolina, Chapel Hill.

[Fletcher et al., 2003a] P. T. Fletcher, S. Joshi, C. Lu, & S. M. Pizer. *Gaussian Distributions on Lie Groups and Their Application to Statistical Shape Analysis.* To appear *Information Processing in Medical Imaging*, 2003.

[Fletcher et al., 2003b] P. T. Fletcher, C. Lu, & S. Joshi. *Statistics of Shape via Principal Geodesic Analysis on Lie Groups.* To appear *Computer Vision and Pattern Recognition*, 2003.

[Gage, 1983] M. Gage. *An isoperimetric inequality with applications to curve shortening.* Invent. Math., (76):357–364, 1983.

[Gauch & Pizer, 1993] John M. Gauch & Stephen M. Pizer. *Multiresolution Analysis of Ridges and Valleys in Grey-Scale Images.* IEEE Transactions on Pattern Analysis and Machine Intelligence, 15(6):635–646, June 1993.

[Gauch, 1999] John M. Gauch. *Image Segmentation and Analysis via Multiscale Gradient Watershed Hierachies.* IEEE Transactions on Image Processing, 8(1):69 – 79, January 1999.

[Gomez et al., 2000] G. Gomez, J.L. Marroquin, & L.E. Sucar. *Probabilistic Estimation of Local Scale.* In *Proc. of the ICPR*, 2000.

[Griffin & Colchester, 1995] Lewis D Griffin & Alan C F Colchester. *Superficial and deep structure in linear diffusion scale space: isophotes, critical points and separatrices.* Image and Vision Computing, 13(7):543–557, September 1995.

[Griffin et al., 2003] Griffin, Lillholm, & Nielsen. *Natural Image Profiles are most likely to be Step Edges.* Vision Research, 2003. Submitted.

[ibs, 1999] *The Internet Brain Segmentation Repository*, 1999. MR brain data set 788_6_m and its manual segmentation was provided by the Center for Morphometric Analysis at Massachusetts General Hospital and is available at http://neuro-www.mgh.harvard.edu/cma/ibsr.

[Ijima, 1962] T. Ijima. *Basic theory on normalization of a pattern (in case of typical one-dimensional pattern).* In *Bulletin of Electrotechnical Laboratory.* 1962. In japanese.

[Jackway, 1996] P.T. Jackway. *Gradient watersheds in morphological scale-space.* IEEE Trans. Image Proc., 5:913–921, 1996.

[Jaynes, 2003] E.T. Jaynes. *Probability Theory: The Logic of Science.* http://omega.albany.edu:8008/JaynesBook, 2003.

[Johansen et al., 1999] Peter Johansen, Mads Nielsen, & Sven Kreiborg. *The Computation of Natural Shape*, 1999.
`http://www.diku.dk/research-groups/image/research/NaturalShape/`.

[Joshi et al., 2002] S. Joshi, S. Pizer, P. T. Fletcher, P. Yushkevich, A. Thall, & J. S. Marron. *Multiscale deformable model segmentation and statistical shape analysis using medial descriptions.* Transactions on Medical Imaging, 21(5), 2002.

[Kanters et al., 2003] Frans Kanters, Bram Platel, Luc Florack, & Bart M. ter Haar Romeny. *Content Based Image Retrieval using Multi-scale Top Points.* In *Scale Space Methods in Computer Vision, Proceedings from 4th Int. Conference*, volume 2695 of LNCS, 2003.

[Kass & Terzopoulos, 1988] Michael Kass & Andrew Witkin Demitri Terzopoulos. *Snakes: Active Contour Models.* Int. Journal of Computer Vision, pages 321–331, 1988.

[Koenderink, 1984] Jan J. Koenderink. *The Structure of Images.* Biological Cybernetics, 50:363–370, 1984.

[Koster, 1995] André Koster. *Linking Models for Multi-scale Image Sequences.* PhD thesis, University of Utrecht, 1995.

[Kullback & Leibler, 1951] S. Kullback & R.A. Leibler. *On Information Theory and Sufficiency.* Annals of Mathematical Statistics, 22, 1951.

[Kwan et al., 1996] R.K.-S. Kwan, A.C. Evans, & G.B. Pike. *An Extensible MRI Simulator for Post-Processing Evaluation.* Lecture Notes in Computer Science, 1131:135–140, 1996. http://www.bic.mni.mcgill.ca/brainweb/.

[Letteboer et al., 2001] M. Letteboer, W. Niessen, P. Willems, E. B. Dam, & M. Viergever. *Interactive multi-scale watershed segmentation of tumors in MR brain images.* In *Proc. of the IMIVA workshop of MICCAI*, 2001.

[Letteboer et al., 2003] M. Letteboer, O. F. Olsen, E. B. Dam, P. Willems, M. Viergever, & W. Niessen. *Segmentation of Tumors in MR Brain Images using an Interactive Multi-scale Watershed Algorithm.* IEEE Transactions on Medical Imaging, 2003. Submitted.

[Lifshitz & Pizer, 1990] Lawrence Lifshitz & Stephen Pizer. *A Multiresolution Hierarchical Approach to Image Segmentation Based on Intensity Extrema.* IEEE PAMI, 12(6):529 – 540, June 1990.

[Lindeberg, 1994] Tony Lindeberg. *Scale-Space Theory in Computer Vision.* Kluwer Academic Publishers, 1994.

[Lorensen & Cline, 1987] William E. Lorensen & Harvey E. Cline. *Marching Cubes: A High Resolution 3D Surface Construction Algorithm.* Computer Graphics, 21(4):163–169, 1987.

[Lorenzen et al., 2001] P. Lorenzen, S. Joshi, G. Gerig, & E. Bullitt. *Tumor-Induced Structural Radiometric Asymmetry in Brain Images.* In *Proc. Workshop on Mathematical Methods in Biomedical Image Analysis MMBIA 2001.* IEEE Computer Society, 2001.

[Maxwell, 1870] J.C. Maxwell. *On Hills and Dales.* The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 4th Series, 40(269):421–425, December 1870.

[Mumford & Shah, 1985] D. Mumford & J. Shah. *Boundary Detection by Minimizing Functionals.* In *CVPR85*, pages 22–26, 1985.

[Murakami, 1998] S. Murakami. *Three and four dimensional diagnosis for TMJ.* In *Proc. of CAR 98*, pages 88–91, 1998.

[Niessen et al., 1997] Wiro J. Niessen, Koen L. Vincken, Joachim A. Weickert, & Max A. Viergever. *Nonlinear Multiscale Representations for Image Segmentation.* Computer Vision and Image Understanding, 66(2):233–245, May 1997.

[Nitzberg & Shiota, 1992] Mark Nitzberg & Takahiro Shiota. *Nonlinear Image Filtering with Edge and Corner Enhancement.* IEEE Tr. Pattern Analysis and Machine Intelligence, 14(8):826 – 833, 1992.

[Nordstrom, 1990] K.N. Nordstrom. *Biased Anisotropic Diffusion: A Unified Regularization and Diffusion Approach to Edge Detection.* IVC, 8(4):318–327, 1990.

[Olsen & Nielsen, 1997] Ole Fogh Olsen & Mads Nielsen. *Generic events for the gradient squared with application to multi-scale segmentation.* In *Scale-Space Theory in Computer Vision, Proc. 1st International Conference*, volume 1252 of Lecture Notes in Computer Science, pages 101–112, Utrecht, The Netherlands, July 1997.

[Olsen, 1996] Ole Fogh Olsen. *Multi-Scale Segmentation of Grey-Scale Images.* Technical Report 96/30, Department of Computer Science, University of Copenhagen, 1996.

[Olsen, 1997] Ole Fogh Olsen. *Multi-Scale Watershed Segmentation.* In Jon Sporring, Mads Nielsen, Luc Florack, & Peter Johansen, editors, *Gaussian Scale-Space Theory*, pages 191–200. Kluwer, 1997.

[Olsen et al., 2000] O. F. Olsen, E. B. Dam, M. Lillholm, A. Thomsen, P. Johansen, , & M. Nielsen. *Method and System for Multi-dimensional Segmentation.* Danish patent application filed October 2000, 2000.

[Ott, 1993] Edward Ott. *Chaos in Dynamical Systems.* Cambridge University Press, 1993.

[Perona & Malik, 1990] Pietro Perona & Jitendra Malik. *Scale-Space and Edge Detection Using Anisotropic Diffusion.* IEEE PAMI, 12(7):629 – 639, July 1990.

[Pizer et al., 1994] S.M. Pizer, C.A. Burbeck, J.M. Coggins, D.S. Fritsch, & B.S. Morse. *Object Shape Before Boundary Shape: Scale-Space Medial Axes.* JMIV, 4:303–313, 1994.

[Pizer et al., 1996] S.M. Pizer, D.S. Fritsch, P. Yushkevich, V. Johnson, & E. Chaney. *Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape.* IEEE Transactions on Medical Imaging, 18(10), 1996.

[Pizer et al., 2003a] Pizer, Chen, Fletcher, Fridman, Fritsch, Gash, Glotzer, Jiroutek, Joshi, Muller, Thall, Tracton, Yushkevich, & Chaney. *Deformable M-Reps for 3D Medical Image Segmentation.* IJCV, 2003. To appear, see http://midag.cs.unc.edu/pubs/papers/IJCV01-Pizer-mreps.pdf.

[Pizer et al., 2003b] S.M. Pizer, K. Siddiqi, G. Székely, J. Damon, & S.W. Zucker. *Multiscale Medial Loci and Their Properties.* IJCV, 2003. To appear, see http://midag.cs.unc.edu/pubs/papers/IJCV01-Pizer-medloci.pdf.

[Press et al., 1999] Press, Teukolsky, Vetterling, & Flannery. *Numerical Recipies in C.* Cambridge University Press, 1999.

[Rao, 2003] Manjari Rao. *Analysis of a Locally Varying Intensity Template for Segmentation of Kidneys in CT images.* Master's thesis, University of North Carolina, Chapel Hill, 2003.

[Rudin et al., 1992] L. I. Rudin, S. Osher, & E Fatemi. *Nonlinear total variation based noise removal algorithms.* Physica D, pages 259–268, 1992.

[Samet, 1984] H. Samet. *The Quadtree and Related Hierarchical Data Structures.* Surveys, 16(2):187–260, June 1984.

[Sapiro, 2001] Guillermo Sapiro. *Geometric Partial Differential Equations and Image Analysis.* Cambridge, 2001.

[Scharr & Weickert, 2000] Hanno Scharr & Joachim Weickert. *An Anisotropic Diffusion Algorithm with Optimized Rotation Invariance.* In *Mustererkennung 2000, Proceedings from 22. DAGM-symposium*, pages 460–467. Springer, 2000.

[Shokoufandeh et al., 2002] Ali Shokoufandeh, Sven Dickinson, Clas Jonsson, Lars Bretzner, & Tony Lindeberg. *On the Representation and Matching of Qualitative Shape at Multiple Scales.* In *Proceedings, European Conference on Computer Vision*, LNCS, 2002.

[Siddiqi et al., 1999] K. Siddiqi, A. Shokoufandeh, S. Dickinson, & S. Zucker. *Shock Graphs and Shape Matching.* International Journal of Computer Vision, 30, 1999.

[Sporring et al., 1997] Jon Sporring, Mads Nielsen, Luc Florack, & Peter Johansen. *Gaussian Scale-Space Theory.* Kluwer Academic Publishers, Dordrecht, 1997.

[Styner & Gerig, 2001] Martin Styner & Guido Gerig. *Medial Models Incorporating Object Variability for 3D Shape Analysis.* In *Proc. of Information Processing in Medical Imaging*, 2001.

[ter Haar Romeny, 1994] Bart M. ter Haar Romeny, editor. *Geometry-Driven Diffusion in Computer Vision.* Kluwer Academic Publishers, 1994.

[Thall, 2002] A. Thall. *Fast C^2 interpolating subdivision surfaces using iterative inversion of stationary subdivision rules.* Technical report, University of North Carolina Department of Computer Science, 2002. http://midag.cs.unc.edu/pub/papers/Thall_TR02-001.pdf.

[Tracton et al., 1994] G. Tracton, E. Chaney, J. Rosenman, & S. Pizer. *MASK: combining 2D and 3D segmentation methods to enhance functionality.* In *Proceedings of Mathematical Methods in Medical Imaging III*, volume SPIE Vol 2299, 1994.

[van Hateren & van der Schaaf, 1998] van Hateren & van der Schaaf. *Independent component filters of natural images compared with simple cells in primary visual cortex.* Proceedings of the Royal Society of London Series B - Biological Sciences, 265 (1394):359–366, 1998.

[Vincken, 1995] Koen Vincken. *Probabilistic Multi-scale Image Segmentation by the Hyperstack.* PhD thesis, University of Utrecht, 1995.

[Wang, 2003] Haonan Wang. *Functional Data Analysis of Populations of Tree-structured Objects.* PhD thesis, University of North Carolina, Chapel Hill, 2003.

[Webb, 1988] Steve Webb. *The Physics of Medical Imaging.* Institute of Physics Publishing, London, 1988.

[Weickert, 1998a] Joachim Weickert. *Anisotropic Diffusion in Image Processing.* B. G. Teubner, 1998.

[Weickert, 1998b] Joachim Weickert. *Efficient Image Segmentation Using Partial Differential Equations and Morphology.* Technical Report 98-10, Department of Computer Science, University of Copenhagen, 1998.

[Witkin, 1983] Andrew P. Witkin. *Scale-Space Filtering.* In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 1019–1022, Karlsruhe, Germany, 1983.