Learning and Bayesian Shape Extraction for Object Recognition

Washington Mio, Anuj Srivastava, and Xiuwen Liu

Florida State University, Tallahassee FL 32306, USA

Abstract. We present a novel algorithm for extracting shapes of contours of (possibly partially occluded) objects from noisy or low-contrast images. The approach taken is Bayesian: we adopt a region-based model that incorporates prior knowledge of specific shapes of interest. To quantify this prior knowledge, we address the problem of learning probability models for collections of observed shapes. Our method is based on the geometric representation and algorithmic analysis of planar shapes introduced and developed in [15]. In contrast with the commonly used approach to active contours using partial differential equation methods [12,20,1], we model the dynamics of contours on vector fields on shape manifolds.

1 Introduction

The recognition and classification of objects present in images is an important and difficult problem in image analysis. Applications of shape extraction for object recognition include video surveillance, biometrics, military target recognition, and medical imaging. The problem is particularly challenging when objects of interest are partially obscured in low-contrast or noisy images. Imaged objects can be analyzed in many ways: according to their colors, textures, shapes, and other characteristics. The past decade has seen many advances in the investigation of models of pixel values, however, these methods have only found limited success in the recognition of imaged objects. Variational and level-set methods have been successfully applied to a variety of segmentation, denoising, and inpainting problems (see e.g. [1]), but significant advances are still needed to satisfactorily address recognition and classification problems, especially in applications that require real-time processing.

An emerging viewpoint among vision researchers is that global features such as shapes should be taken into account. The idea is that by incorporating some prior knowledge of shapes of objects of interest to image models, one should be able to devise more robust and efficient image analysis algorithms. Combined with clustering techniques for the hierarchical organization of large databases of shapes [22], this should lead to recognition and classification algorithms with enhanced speed and performance. In this paper, we construct probability models on shape spaces to model a given collection of observed shapes, and integrate these to a region-based image model for Bayesian extractions of shapes from images. Our primary goal is to capture just enough information about shapes

T. Pajdla and J. Matas (Eds.): ECCV 2004, LNCS 3024, pp. 62–73, 2004.

present in images to be able to identify them as belonging to certain categories of objects known *a priori*, not to extract fine details of the contours of imaged objects.

Shape analysis has been an important theme of investigation for many years. Following the seminal work of Kendall [13], large part of the research in quantitative shape analysis has been devoted to "landmark-based" studies, where shapes are represented by finite samplings of contours. One establishes equivalences of representations with respect to shape preserving transformations, and then compares shapes in the resulting quotient space [5,21]. Statistical shape models based on this representation have been developed and applied to image segmentation and shape learning in [7,6]; the literature on applications of this methodology to a variety of problems is quite extensive. A drawback of this approach is that the automatic selection of landmarks is not straightforward and the ensuing shape analysis is heavily dependent on the choices made. Grenander's deformable templates [8] avoids landmarks by treating shapes as points in an infinite-dimensional differentiable manifold, and modeling variations of planar shapes on an action of the diffeomorphism group of $\mathbb{R}^2$ [24,9,18]. However, computational costs associated with this approach are typically very high. A very active line of research in image analysis is based on active contours [12,20] governed by partial differential equations; we refer the reader to [1] for a recent survey on applications of level-set methods to image analysis. Efforts in the direction of studying shape statistics using partial differential equation methods have been undertaken in [17,3,2].

In [15], Klassen et al. introduced a new framework for the representation and algorithmic analysis of continuous planar shapes, without resorting to defining landmarks or diffeomorphisms of $\mathbb{R}^2$. To quantify shape dissimilarities and simulate optimal deformations of shapes, an algorithm was developed for computing geodesic paths in shape spaces. The registration of curves to be compared is automatic, and the treatment suggests a new technique for driving active contours [23]. In this paper, we investigate variants of this model for shape extraction from images. In our formulation, the dynamics of active contours is governed by vector fields on shape manifolds, which can be integrated with classical techniques and reduced computational costs. The basic idea is to create a manifold of shapes, define an appropriate Riemannian structure on it, and exploit its geometry to solve optimization and inference problems.

An important element in this stochastic geometry approach to shape extraction is a model for shape learning. Assuming that a given collection of observed shapes consists of random samples from a common probability model, we wish to learn the model. Examples illustrating the use of landmark-based shape analysis in problems of this nature are presented in [7,6,14,10]. The problem of model construction using the shape analysis methods of [15] presents two main difficulties: the shape manifold is *nonlinear* and *infinite-dimensional*. A most basic notion needed in the study of sample statistics is that of mean shape; Karcher means introduced in [11] are used. As in [5], other issues involving nonlinearity are handled by considering probability densities on the (linear) tangent space at the mean shape. To tackle the infinite dimensionality, we use approximate

finite-dimensional representations of tangent vectors to shape manifolds. We consider multivariate normal models, so that learning reduces to estimations of the relevant parameters. (Other parametric models can be treated with similar techniques.) Implicit in this approach is that large collection of shapes have been pre-clustered and we are modeling clusters of fairly small diameters. Clustering algorithms and hierarchical organizations of large databases of shapes are discussed in [22].

This paper is organized as follows: in Section 2, we briefly review the material of [15], as it provides the foundations of our stochastic geometry approach to shape extraction. Section 3 is devoted to a discussion of shape learning. In Section 4, we present the image model used in the shape extraction algorithm, and applications of the algorithm to imagery involving partial occlusions of objects, low contrast, or noise.

2 Shape Spaces and Geodesic Metrics

In this section, we review the geometric representation of continuous planar shapes, the geodesic metric on shape space, and the algorithmic shape analysis methods introduced and developed in [15].

2.1 Geometric Representation of Shapes

Shapes of outer contours of imaged objects are viewed as closed, planar curves $\alpha\colon I \to \mathbb{R}^2$, where $I = [0, 2\pi]$. To make shape representations invariant to uniform scaling, the length is fixed to be 2π by requiring that curves be parameterized by arc length, i.e., $\|\alpha'(s)\| = 1$, for every $s \in I$. Then, the tangent vector can be written as $\alpha'(s) = e^{j\theta(s)}$, where $j = \sqrt{-1}$. We refer to $\theta\colon I \to \mathbb{R}$ as an *angle function* for α. Angle functions are invariant under translations of $\mathbb{R}^2$, and the effect of a rotation is to add a constant to θ. Thus, to make the representation invariant to rotations of $\mathbb{R}^2$, it suffices to fix the average of θ to be, say, π. In addition, to ensure that θ represents a closed curve, the condition $\int_0^{2\pi} \alpha'(s)\, ds = \int_0^{2\pi} e^{j\theta(s)}\, ds = 0$ is imposed. Thus, angle functions are restricted to the *pre-shape manifold*

$$\mathcal{C} = \left\{ \theta \in \mathbb{L}^2 \,|\, \frac{1}{2\pi} \int_0^{2\pi} \theta(s)\, ds = \pi \text{ and } \int_0^{2\pi} e^{i\theta(s)} ds = 0 \right\} . \tag{1}$$

Here, $\mathbb{L}^2$ denotes the vector space of all square integrable functions on $[0, 2\pi]$, equipped with the standard inner product $\langle f, g \rangle = \int_0^{2\pi} f(s)g(s)\, ds$. For continuous direction functions, the only remaining variability in the representation is due to the action of the reparametrization group S^1 arising from different possible placements of the initial point $s = 0$ on the curve. Hence, the quotient space $\mathcal{S} \equiv \mathcal{C}/S^1$ is defined as the *space of continuous, planar shapes.*

2.2 Geodesic Paths between Shapes

At each point $\theta \in \mathcal{C}$, the tangent space $T_\theta \mathcal{C}$ to the pre-shape manifold $\mathcal{C} \subset \mathbb{L}^2$ naturally inherits an inner product from $\mathbb{L}^2$. Thus, $\mathcal{C}$ is a Riemannian manifold and the distance between points in $\mathcal{C}$ can be defined using minimal length geodesics. The distance between two points (i.e., shapes) θ_1 and θ_2 in $\mathcal{S}$ is defined as the infimum of all pairwise distances between pre-shapes representing θ_1 and θ_2, respectively. Thus, the distance $d(\theta_1, \theta_2)$ in $\mathcal{S}$ is realized by a shortest geodesic in $\mathcal{C}$ between pre-shapes associated with θ_1 and θ_2. We abuse terminology and use the same symbol θ to denote both a pre-shape and its associated shape in $\mathcal{S}$. We also refer to minimal geodesics in $\mathcal{C}$ realizing distances in $\mathcal{S}$ as geodesics in $\mathcal{S}$, and to tangent vectors to these geodesics as tangent vectors to $\mathcal{S}$.

One of the main results of [15] is the derivation of an algorithm to compute geodesics in $\mathcal{C}$ (and $\mathcal{S}$) connecting two given points. An easier problem is the calculation of geodesics satisfying prescribed initial conditions. Given $\theta \in \mathcal{C}$ and $f \in T_\theta \mathcal{C}$, let $\Psi(\theta, f, t)$ denote the geodesic starting at θ with velocity f, where t denotes the time parameter. The geodesic $\Psi(\theta, f, t)$ is constructed with a numerical integration of the differential equation satisfied by geodesics. The correspondence $f \mapsto \Psi(\theta, f, 1)$ defines a map $\exp_\theta \colon T_\theta \mathcal{C} \to \mathcal{C}$ known as the exponential map at θ. The exponential map simply evaluates the position of the geodesic Ψ at time $t = 1$. Consider the exponential map at θ_1. Finding the geodesic from θ_1 to θ_2 is equivalent to finding the direction f such that $\exp_{\theta_1}(f) = \theta_2$. For each $f \in T_{\theta_1} \mathcal{C}$, let $E(f) = \| \exp_{\theta_1}(f) - \theta_2 \|^2$ be the square of the $\mathbb{L}^2$ norm of the residual vector. The goal is to find the vector f that minimizes (i.e., annihilates) E. A gradient search is used in [15] to solve this energy minimization problem. This procedure can be refined to yield geodesics in $\mathcal{S}$ by incorporating the action of the re-parametrization group S^1 into the search.

Figure 1 shows an example of a geodesic path in $\mathcal{S}$ computed with this algorithm. In this paper, we have added the invariance of shapes to reflections in

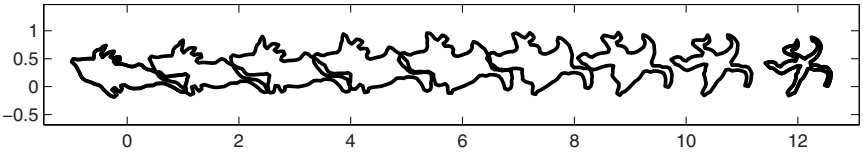

Fig. 1. A geodesic path in shape space

$\mathbb{R}^2$. Essentially, one computes geodesics for both a shape and a reflection, and selects the one with least length.

2.3 Karcher Mean Shapes

The use of *Karcher means* to define mean shapes in $\mathcal{S}$ is suggested in [15]. If $\theta_1, \ldots, \theta_n \in \mathcal{S}$ and $d(\theta, \theta_j)$ is the geodesic distance between θ and θ_j, a Karcher

mean is defined as an element $\mu \in \mathcal{S}$ that minimizes the quantity $\sum_{i=1}^{n} d^2(\theta, \theta_i)$. An iterative algorithm for computing Karcher means in Riemannian manifolds is presented in [16,11] and particularized to the spaces $\mathcal{C}$ and $\mathcal{S}$ in [15]. An example of a Karcher mean shape is shown in Figure 2.

Fig. 2. The Karcher mean shape of eight boots.

2.4 Computational Speeds

To demonstrate the level of performance of the algorithm to compute geodesics between two shapes, Table 2.4 shows the average computation times achieved under several different settings, each estimated by averaging 1,225 calculations on a personal computer with dual Xeon CPUs (at 2.20 GHz) running Linux. Each shape is sampled using T points on the curve and tangent vectors are approximated using $2m$ Fourier terms. Consistent with our analysis, the algorithm for calculating geodesics is linear in T and m. Computational efficiency can be further improved with parallel processing, since the costliest step in the algorithm consists of $2m$ calculations that can be executed independently.

Table 1. Average computation time (in seconds) per geodesic.

T	50	50	100	100	200	200	400	400
m	50	100	100	200	200	400	400	800
Time (secs.)	0.0068	0.0133	0.0268	0.0525	0.1044	0.2066	0.4172	0.8274

3 Shape Learning

An important problem in statistical shape analysis is to "learn" probability models for a collection of observed shapes. Assuming that the given shapes are random samples from the same probability model, we wish to learn the model. These models can then be used as shape priors in Bayesian inferences to recognize or classify newly observed shapes. Implicit in our considerations is the assumption that observed shapes have been pre-clustered, so that we are seeking probability models for clusters of fairly small diameters in $\mathcal{S}$. Clustering techniques on the shape space $\mathcal{S}$ have been studied in [22].

Learning a probability model amounts to estimating a probability density function on shape space, a task that is rather difficult to perform precisely. In this paper, we assume a parametric form for the models so that learning is reduced to an estimation of the relevant parameters. To simplify the discussion of probability models on infinite-dimensional manifolds, the models will be presented in terms of their negative log-likelihood, i.e., the energy of the distribution.

The simplest model is a "uniform Gaussian" on $\mathcal{S}$, whose energy is proportional to $d^2(\theta, \mu)/2$, where μ is the Karcher mean of the sample. The constant of proportionality is related to the variance, as usual. We wish to refine the model to a multivariate normal distribution. Two main difficulties encountered are the *nonlinearity* and the *infinite-dimensionality* of $\mathcal{S}$, which are addressed as follows.

(i) **Nonlinearity.** Since $\mathcal{S}$ is a nonlinear space, we consider probability distributions on the tangent space $T_\mu \mathcal{C}$ at the mean pre-shape $\mu \in \mathcal{C}$, to avoid dealing with the nonlinearity of $\mathcal{S}$ directly. This is similar to the approach taken in [5].

(ii) **Dimensionality.** Our parametric models will require estimations of covariance operators of probability distributions on $T_\mu \mathcal{C} \subset \mathbb{L}^2$. We approximate covariances by an operators defined on finite dimensional subspaces of $T_\mu \mathcal{C}$.

Let $\Theta = \{\theta_1, \ldots, \theta_r\}$ represent a finite collection of shapes. The estimation of the Karcher mean shape μ of Θ is described in [15]. Using μ and the shapes θ_j, $1 \leq j \leq r$, we find tangent vectors $g_j \in T_\mu \mathcal{S}$ such that the geodesic from μ in the direction g_j reaches θ_j in unit time, that is, $\exp_\mu(v_j) = \theta_j$. This lifts the shape representatives to the tangent space at μ.

Let V be the subspace of $T_\mu \mathcal{S}$ spanned by $\{v_1, \ldots, v_r\}$, and $\{e_1, \ldots, e_m\}$ an orthonormal basis of V. Given $v \in V$, write it as $v = x_1 e_1 + \ldots + x_m e_m$. The correspondence $v \mapsto \mathbf{x} = (x_1, \ldots, x_m)$ identifies V with $\mathbb{R}^m$, so we assume that $v_j \in \mathbb{R}^m$. We still have to decide what model to adopt for the probability distribution. We assume a multivariate Gaussian model for $\mathbf{x}$ with mean $\mathbf{0}$ and covariance matrix $K \in \mathbb{R}^{m \times m}$. The estimation of K using sample covariance follows the usual procedures. Depending on the number and the nature of the shape observations, the rank of K may be much smaller than m. Extracting the dominant eigenvectors and eigenvalues of the estimated covariance matrix, one captures the dominant modes of variation and the variances along these principal directions.

To allow small shape variations in directions orthogonal to those determined by the non-zero eigenvalues of K, choose $\varepsilon > 0$ somewhat smaller than the dominant eigenvalues of K. If $K_\varepsilon = K + \varepsilon^2 I_m$, where I_m is the $m \times m$ identity matrix, we adopt the multivariate normal distribution

$$\frac{1}{(2\pi)^{m/2} \det(K_\varepsilon)^{1/2}} \exp\left(-\frac{\mathbf{x}^T \cdot K_\varepsilon^{-1}(\mathbf{x})}{2}\right) \tag{2}$$

on the subspace V of $T_\mu \mathcal{S}$. If $\theta \in \mathcal{S}$, let $g \in T_\mu \mathcal{S}$ satisfy $\Psi(\mu, g, 1) = \theta$, and let $g_V = \sum_{i=1}^m x_i e_i$ be the orthogonal projection of g onto V. We adopt a probability

model on $\mathcal{S}$ whose energy is given up to an additive constant by

$$F(\theta; \mu, K) = \frac{\mathbf{x}^T \cdot K_\varepsilon^{-1}(\mathbf{x})}{2} + \frac{1}{\varepsilon^2} \|g^\perp\|^2 , \tag{3}$$

where $g^\perp \in T_\mu \mathcal{S}$ is the component of g orthogonal to V. Strictly speaking, this definition is only well posed if the exponential map is globally one-to-one. However, for most practical purposes, one can assume that this condition is essentially satisfied because clusters are assumed to be concentrated near the mean and finite-dimensional approximations to θ are used.

The first row of Figure 3 shows eigenshapes associated with the first five eigenvalues (in decreasing order) of the multivariate normal model derived from the shapes in Figure 2. The solid lines show the mean shape, and the dotted lines represent variations about the mean along principal directions. Variations are uniformly sampled on an interval of size proportional to the eigenvalues.

Having obtained a probability model for observed shapes, an important task is to validate it. This can be done in a number of ways. As an illustration, we use the model for random sampling. The second and third rows of Figure 3 show examples of random shapes generated using the Gaussian model learned from the shapes in Figure 2.

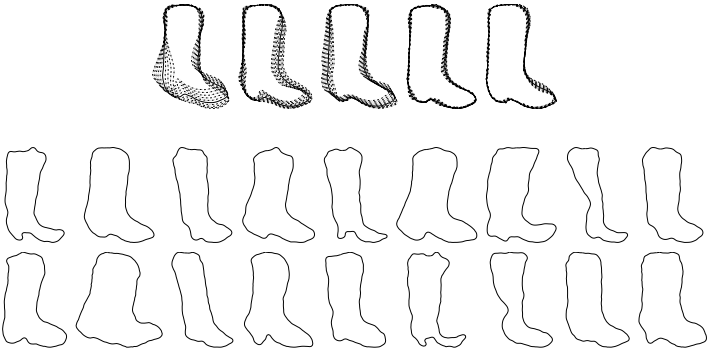

Fig. 3. The first row shows eigenboots in dotted lines, i.e., variations about the mean shape of Figure 2 (displayed in solid lines) along principal directions associated with the five dominant eigenvalues. The second and third rows display 18 random shapes sampled from the proposed multivariate normal model.

Another example is shown in Figure 4. A set of nineteen observed shapes of swimming ducks is analyzed for learning a probability model. We calculated the mean shape – shown on the lower right corner – and estimated the sample covariance matrix K. Figure 5 shows variations of the mean shape along the dominant principal directions and ten random shapes generated using the learned probability model.

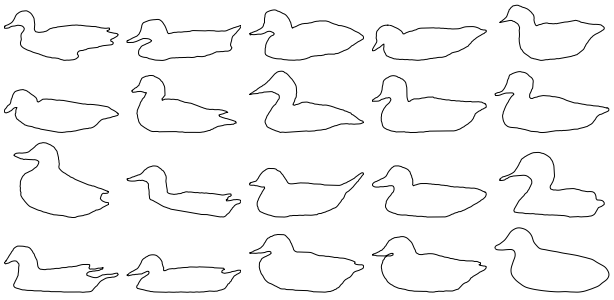

Fig. 4. Nineteen shapes of swimming ducks and their mean shape displayed on the lower right corner.

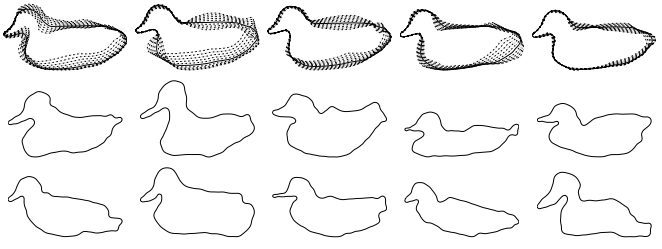

Fig. 5. Eigenducks (first row) and ten random shapes generated by sampling from a multivariate normal tangent-space model.

4 Bayesian Extraction of Shapes

The extraction of shapes of partially occluded objects from noisy or low-contrast images is a difficult problem. Lack of clear data in such problems may severely limit the performance of image segmentation algorithms. Thus, techniques for integrating some additional knowledge about shapes of interest into the inference process are sought. The framework developed in this paper is well suited to the formulation and solution of Bayesian shape inference problems involving this type of imagery. We assume that the shape to be extracted is known *a priori* to be related to a family modeled on a probability distribution of the type discussed in Section 3. The case of several competing models can be treated with a combination of our shape extraction method and hypothesis testing techniques. We emphasize that our goal is to extract just enough features of shapes present in images to be able to recognize objects as belonging to certain known categories, not to capture minute details of shapes. Such low-resolution approach is more robust to noise and allows for greater computational efficiency.

Our analysis thus far has focused on shape, a property that is independent of variables that account for rotations, translations, and scalings of objects. However, shapes appear in images at specific locations and scales, so the process of shape extraction and recognition should involve an estimation of these nuisance variables as well. Hence, in this context, the data likelihood term assumes

knowledge of shape, location, and scale variables, while the prior term depends only on shape. Therefore, we first revisit our shape representation to incorporate these extra variables.

To account for translational effects, we introduce a variable $p \in \mathbb{R}^2$ that identifies the centroid of a constant-speed curve $\alpha\colon I \to \mathbb{R}^2$, which is given by $p = (1/2\pi)\int_0^{2\pi} \alpha(s)\,ds$. We adopt a logarithmic scale for the length by writing $L = e^{\ell}$, $\ell \in \mathbb{R}$. Lastly, to allow arbitrary rotations, we simply relax the constraints on θ used in the description of the pre-shape manifold $\mathcal{C}$ and only require that θ satisfy the closure conditions

$$\int_0^{2\pi} \cos\theta(s)\,ds = 0 \quad \text{and} \quad \int_0^{2\pi} \sin\theta(s)\,ds = 0. \tag{4}$$

Thus, pre-shapes that can change position and are free to shrink or stretch uniformly will be described by triples $(p, \ell, \theta) \in \mathbb{R}^2 \times \mathbb{R} \times \mathbb{L}^2$ satisfying (4). The collection of all such triples will be denoted $\mathcal{F}$. An element $(p, \ell, \theta) \in \mathcal{F}$ represents the curve

$$\alpha(s) = p + e^{\ell}\int_0^s e^{j\theta(x)}\,dx - \frac{e^{\ell}}{2\pi}\int_0^{2\pi}\int_0^s e^{j\theta(x)}\,dx\,ds. \tag{5}$$

For shape extraction, we do not need to further consider the quotient space under the action of the re-parameterization group $\mathbb{S}^1$ on $\mathcal{F}$. The data likelihood and shape prior terms will be invariant under the $\mathbb{S}^1$-action, so the posterior energy will be constant along $\mathbb{S}^1$ orbits. We now describe the posterior energy for our Bayesian inference.

(a) **Data Likelihood.** Let $D \subset \mathbb{R}^2$ be the image domain and $I\colon D \to \mathbb{R}^+$ be an image. A closed curve represented by (p, ℓ, θ) divides the image domain into a region $D_i(p, \ell, \theta)$ inside the curve, and a region $D_o(p, \ell, \theta)$ outside. Let p_i be a probability model for the pixel values inside the curve, and p_o be a model for pixels outside. For simplicity, we assume a binary image model choosing p_i and p_o to be Gaussian distributions with different means. (Alternatively, one can use variants of the Mumford-Shah image model [19]). For a given (p, ℓ, θ), the compatibility of an image I with (p, ℓ, θ) is proportional to

$$H(I|p, \ell, \theta) = -\iint_{D_i(p,\ell,\theta)} \log p_i(I(y))\,dy - \iint_{D_o(p,\ell,\theta)} \log p_o(I(y))\,dy. \tag{6}$$

(b) **Shape Prior.** Let μ_0 and K_0 represent the mean and the covariance matrix associated with the shape prior model. Set the prior energy to be $F([\theta]; \mu, K)$, as in Equation 3, where $[\theta]$ indicates that θ has been normalized to have average π.

Combining the two terms, up to an additive constant, the posterior energy is proportional to

$$P_\lambda\,(p, \ell, \theta|I) = \lambda H(I|p, \ell, \theta) + (1 - \lambda)F([\theta]; \mu, K)\,,$$

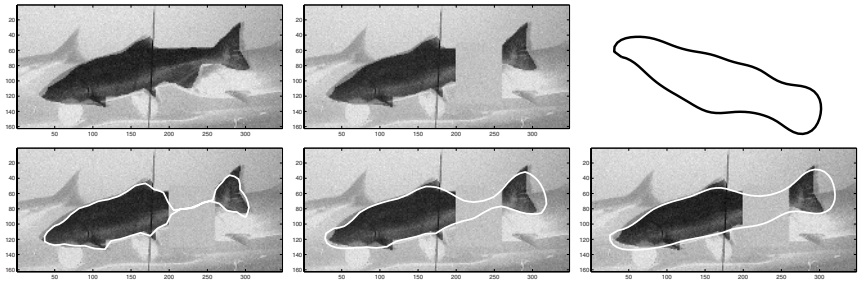

Fig. 6. An example of shape extraction. Top row from left to right: an image, the same image with a partial occlusion, and the prior mean. The bottom row shows MAP shape estimates with an increasing influence of the prior from left to right.

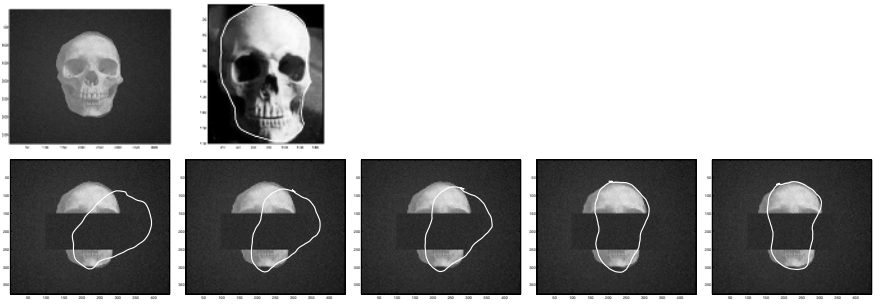

Fig. 7. A shape extraction experiment. The top row shows an image of a skull to be analyzed after being artificially obscured and a second skull whose contour is used as prior mean. The bottom row shows various stages of the curve evolution during the shape extraction.

where $0 < \lambda < 1$. As before, it is convenient to lift θ to the tangent space at the mean μ. Letting $\theta = \Psi(\mu, g, 1) = \exp_\mu(g)$, with g in the tangent space at μ, we rewrite the posterior energy as

$$E_\lambda(p, \ell, g|I) = P_\lambda\left(p, \ell, \exp_\mu(g)|I\right).$$

We use a gradient search for a MAP estimation of (p, ℓ, g), approximating g with a truncated Fourier series.

Shown in Figure 6 are illustrations of this Bayesian shape extraction using a uniform Gaussian prior. The top row shows an object embedded in an image, the same image with the object partially obscured, and the prior mean shape. The bottom row displays MAP estimates of the shape of the object under an increasing influence of the prior. The improvements in discovering hidden shapes despite partial occlusions emphasize the need and power of a Bayesian approach to such problems using shape priors.

Figure 7 depicts the results of another shape extraction experiment. On the top row, the first panel displays the image of a skull that is artificially obscured

for shape extraction. The second panel shows the contour of a different skull used as prior mean. The second row shows various stages of the curve evolution during the gradient search of a MAP estimate of the shape.

5 Summary

We presented an algorithm for the extraction of shapes of partially obscured objects from noisy, low-contrast images for the recognition and classification of imaged objects. The image model adopted involves a data likelihood term based on pixel values and a shape prior term that makes the algorithm robust to image quality and partial occlusions. We discussed learning techniques in shape space in order to construct probability models for clusters of observed shapes using the framework for shape analysis developed in [15] A novel technique that models the dynamics of active contours on vector fields on shape manifolds was employed. Various shape extraction experiments were carried out to demonstrate the performance of the algorithm.

Acknowledgments. This work was partially supported by the National Science Foundation and the Intelligence Technology Innovation Center through the joint "Approaches to Combat Terrorism" Program Solicitation NSF 03-569 (DMS-0345242) and by the grants NSF (FRG) DMS-0101429 and NMA 201-01-2010.

References

1. Chan, T., Shen, J., Vese, L.: Variational PDE Models in Image Processing. Notices Amer. Math. Soc. **50** (2003) 14–26.
2. Cremers, D., Soatto, S.: A Pseudo Distance for Shape Priors in Level Set Segmentation. In: Proc. 2nd IEEE Workshop on Variational, Geometric, and Level-Set Methods in Computer Vision, Nice, France (2003).
3. Cremers, D., F. Tischhäuser, F., Weickert, J., Schnörr, C.: Diffusion Snakes: Introducing Statistical Shape Knowledge into the Mumford-Shah Functional. Internat. Journal Computer Vision, **50** (2002), 295–313.
4. Do Carmo, M.P.: Differential Geometry of Curves and Surfaces. Prentice-Hall (1976).
5. Dryden, I.L., Mardia, K.V.: Statistical Shape Analysis. John Wiley & Sons (1998).
6. Duta, N., Jain, A.K., Dubuisson-Jolly, M.P.: Automatic Construction of 2D Shape Models. IEEE Trans. Pattern Analysis and Machine Intelligence **23** (2001).
7. Duta, N., Sonka, M., Jain, A.K.: Learning Shape Models from Examples Using Automatic Shape Clustering and Procrustes Analysis. In: Proc. Information in Medical Image Processing, Lecture Notes in Computer Science, Vol. 1613. Springer-Verlag, Berlin Heidelberg New York (1999) 370–375.
8. Grenander, U.: General Pattern Theory. Oxford University Press (1993).
9. Grenander, U., Milller, M.I.: Computational Anatomy: an Emerging Discipline. Quarterly of Applied Mathematics, **LVI** (1998) 617–694.
10. Holbolth, A., Kent, J.T., Dryden, I.L.: On the Relation Between Edge and Vertex Modeling in Shape Analysis. Scandinavian Journal of Statistics **29** (2002) 355–374.

11. Karcher, H.: Riemann Center of Mass and Mollifier Smoothing. Comm. Pure Appl. Math. **30** (1977) 509–541.
12. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active Contour Models. International Journal of Computer Vision **1** (1988) 321–331.
13. Kendall, D.G.: Shape Manifolds, Procrustean Metrics and Complex Projective Spaces. Bull. London Math. Society **16** (1984) 81–121.
14. Kent, J.T., Dryden, I.L., Anderson, C.R.: Using Circulant Symmetry to Model Featureless Objects. Biometrika **87** (2000) 527–544.
15. Klassen, E., Srivastava, A., Mio, W., Joshi, S.: Analysis of Planar Shapes Using Geodesic Paths on Shape Spaces. IEEE Trans. Pattern Analysis and Machine Intelligence **26** (2004).
16. Le, H.L., Kendall, D.G.: The Riemannian Structure of Euclidean Shape Spaces: a Novel Environment for Statistics. Annals of Statistics **21** (1993) 1225–1271.
17. Leventon, M., Grimson, W., Faugeras, O.: Statistical Shape Infuence in Geodesic Active Contours. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (2000).
18. Miller, M.I., Younes, L.: Group Actions, Homeomorphisms, and Matching: A General Framework. International Journal of Computer Vision **41** (2002) 61–84.
19. Mumford, D., Shah, J.: Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems. Comm. Pure Appl. Math. **42** (1989) 577–685.
20. Sethian, J.: Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Science. Cambridge University Press (1996).
21. Small, C.G.: The Statistical Theory of Shape. Springer-Verlag (1996).
22. Joshi, S., Srivastava, A., Mio, W., Liu, X.: Hierarchical Organization of Shapes for Efficient Retrieval. In: Proc. 8th European Conference on Computer Vision. Prague, Czech Republic (2004).
23. Srivastava, A., Mio, W., Klassen, E., Liu, X.: Geometric Analysis of Constrained Curves for Image Understanding. In: Proc. 2nd Workshop on Variational, Geometric and Level-Set Methods in Computer Vision. Nice, France (2003).
24. Younes, L.: Computable Elastic Distance Between Shapes. SIAM Journal of Applied Mathematics **58** (1998) 565–586.