USING SHAPE DISTRIBUTIONS AS PRIORS IN A CURVE EVOLUTION FRAMEWORK

Andrew Litvin *William Clem Karl*

Electrical and Computer Engineering Department
Boston University
Email: {litvin,wckarl}@bu.edu

ABSTRACT

In this paper we propose a novel framework for constructing and using a shape prior in a curve evolution framework. The prior shape information is captured through shape distributions, which are histograms of features derived from the shape boundary. The resulting prior captures perceptual shape similarity, is robust to small sample size, and flexible. We further derive a curve evolution force that corresponds to this prior. This enables us to use this prior to perform tasks such as mean shape calculation and image segmentation within a curve evolution framework.

1. INTRODUCTION

The use of prior shape information is indispensable in many image processing applications, such as image segmentation, tracking, classification, etc. Different approaches have been taken to capturing such prior information. For example, one common approach uses principal component analysis of training data to obtain a set of shape representation functions [1] describing allowable shape variations. A recent different approach uses an angle function parameterization coupled with a distance measure on a shape manifold to construct deformable shape models [2]. Another approach to the inclusion of prior shape information is based on explicit modeling and extraction of component parts [3]. Unfortunately, these methods can be sensitive to the chosen boundary parameterization and the resulting models may be difficult to generalize to unseen shapes.

Typical curve evolution-based segmentation methods incorporate a penalty on boundary length, which can be viewed as a shape prior favoring objects with shorter boundaries. Such a prior is intrinsic and generic, but is overly simplistic, resulting in a trade-off between boundary smoothing and suppression of salient shape structure through, for example, the rounding off of corners and shrinking of object size [4].

This work was partially supported by The Air Force Office of Scientific Research under Grant F49620-03-1-0257, The National Institutes of Health under Grant NINDS 1 R01 NS34189, and The Engineering research centers program of the NSF under award EEC-9986821

There have been several attempts to extend this type of prior. An alternative data-driven prior shape model was proposed in [5]. A model of the distribution of curvature and intensity with respect to a segmented curve was found from training data. This spatially stationary model was then used in a maximum a posteriori (MAP) formulation to segment an image. Although giving better results than generic curve length penalty priors, this approach still tends to suppress salient structures. The reason is that the stationary prior coupled with the MAP criterion attempts to drive the curvature at every point on the curve to the same, constant value corresponding to the mode of the distribution.

Our goal is to overcome the drawbacks of existing methods of constructing shape priors for use in a curve evolution context. We aim for a compromise between the focusing ability of the prior model, its generalizability, and its cost to find and implement. Motivated by the ideas in [6, 7], we construct a shape boundary prior as an energy which penalizes the difference between the entire feature *distributions* of a given curve and those of the prior. This new formulation is our first contribution.

Unlike the conventional curve length penalty, our new energy depends on the segmenting curve in a non-local way, making the calculation of the minimizing flow challenging. We propose an efficient solution of this problem by constructing a distribution matching PDE as our second contribution. The overall result is a new flexible and tractable approach to the inclusion of prior shape information into curve-evolution-based segmentation. We present preliminary results of applying our method to two problems of interest, namely mean shape calculation and image segmentation.

2. PRIOR FORMULATION

Let Φ denote a continuously valued feature (for example: curvature) defined on the space Ω (for example: arc length along the curve) and let λ be the value of the feature (for example: the value of curvature). For simplicity, we may assume that $\lambda \in [0, \lambda_{max}]$. We define the "distribution function" $H(\lambda)$ be the fraction of the space Ω with values greater than λ: $H(\lambda) = \frac{\int_\Omega \{\Phi(\Omega)>\lambda\} d\omega}{\int_\Omega d\omega}$. Note that $H(\lambda)$ is related to cumulative distribution function (CDF): $H(\lambda) = 1 - CDF(\lambda)$. We then define our prior energy as:

$$E_C = \sum_{i=1}^{M} w_i \int_{0}^{\lambda_{max}} \left[H^i(C,\lambda) - H_*^i(\lambda)\right]^2 d\lambda \qquad (1)$$

where M is the number of features taken into account, $H^i(C,\lambda)$ and $H_*^i(\lambda)$ are the distribution functions of the i^{th} feature of the curve and prior, respectively, and w_i weighs the importance of each feature. The prior feature distribution functions $H_*^i(\lambda)$ can be obtained in a variety of ways, including prior specification based on intuition or estimation from training data, as we elaborate below. In general, we do not require that $H_*^i(\lambda)$ correspond to the feature distribution of any particular curve.

2.1. Image segmentation

The energy (1) can be used in a number of interesting ways. First, we consider the problem of curve-evolution-based image segmentation. We define the curve C segmenting an object as the minimizer of the following functional:

$$C^* = \underset{C}{\operatorname{argmin}} \left[E_d + \alpha E_C\right] \qquad (2)$$

where E_d is a data fidelity term and term $E_{\mathbf{C}}$ represents a prior term on the curve boundary. While a number of choices are possible, here we take the data fidelity term E_d to be an image log-likelihood, capturing the expected intensity profile in the direction normal to the boundary, as in [5].

In most existing curve evolution based approaches the shape prior term E_C simply penalizes the total length of the curve $E_C = ||C||$, resulting in smooth segmentation boundaries. Instead, we will use the energy (1) with the prior feature distribution functions $H_*^i(\lambda)$ obtained as the average of N training shapes C_j: $H_*^i(\lambda) = \sum_{j=1}^{N} H^i(C_j,\lambda)$. The use of this prior will drive the segmenting curve to one whose feature distributions match the training data.

2.2. Mean shape calculation

As a second application, the proposed framework provides a natural and interesting approach to the determination of the average shape $\overline{C}$ of a collection of N shapes C_i. To this end, let us use the energy (1) to define a distance between curves as follows:

$$d(C,C_j) = \sqrt{\sum_{i=1}^{M} w_i \int_{0}^{\lambda_{max}} \left[H^i(C,\lambda) - H^i(C_j,\lambda)\right]^2 d\lambda} \qquad (3)$$

This distance can be shown to satisfy the triangle inequality and is thus a metric. Given this distance we can now define the mean shape using the Karcher mean formula

$$\overline{C} = \underset{C}{\operatorname{argmin}} \sum_{j=1}^{N} d^2(C,C_j) \qquad (4)$$

The resulting mean curve $\overline{C}$ can also be found as the curve which minimizes the distance to the mean of training data feature distributions (versus minimizing the mean of the distances to the individual feature distributions).

3. FEATURES

Any feature representing measures uniformly sampled along the shape boundary or volume can be used in our framework. The factors governing the choice of features for a particular problem include: the ability of a chosen feature set to capture overall shape topology as well as boundary features and computational tractablity of the resulting optimization problem. In this work we use two features: the set of all inter-point distances for a uniformly sampled boundary and the set of curvatures at multiple scales also sampled along the boundary. The first feature captures global shape topology while the second captures boundary features.

4. CURVE EVOLUTION FRAMEWORK

We minimize (2) or (4) by moving the curve in a direction derived from the gradient of the corresponding functional. The details on minimization of the E_d term in (2) can be found in [5]. We will focus on the challenging part - minimization with respect to our prior term E_C in (1).

For simplicity, we consider the case of a single feature. Since our prior term is additive in the different features, minimizing flows for individual features can be added with corresponding weights to obtain the overall minimizing flow. The gradient decent with respect to (1) (and equivalently with respect to (4)) is naturally expressed as histogram modification PDE (flow) on the space of features. However, this flow (i.e. the steepest descent direction) will, in general, not yield valid curve updates. As a result, we propose a 2 step approach to computation of the overall gradient curve flow.

Let the single feature values at time t be $\Phi_t(C)$ (continuously defined). In our first step, we compute the gradient descent flow for the feature $\Phi_t(C)$ corresponding to (1).

$$\frac{\partial \Phi}{\partial t} = -\nabla_\Phi E = \left[H_*(\Phi(t)) - H_t(\Phi(t))\right] H'_\lambda(\Phi_t(t)) \qquad (5)$$

where $H'_\lambda(\Phi_t(t))$ is the derivative of the distribution with respect to λ evaluated at $\Phi_t(t)$. The stationary point of this flow corresponds to the case when $H_*(\Phi(t)) = H_t(\Phi(t))$, i.e. the distribution function for given curve matches the target prior distribution function. The discrete approximation of eq. (5) is straightforward and can be efficiently computed.

Our flow in (5) is similar to the histogram equalization flow introduced in [8]. To our knowledge, however, an energy minimization interpretation of the flow in [8] has not been presented. While the flow in [8] is a minimizing flow of (1), it does not constitutes the gradient decent evolution equation for the energy term (1).

The flow in (5) is defined on the curve *features* without consideration of consistency. Let this flow in (5) on the set of features be denoted by dF^n. The evolution of the features is constrained by their connection to a valid contour C, which induces a manifold of valid feature sets. The second step of our minimization approach aims to project the unconstrained feature flow onto this manifold. Define the displacement in the normal direction of the underlying contour at node k as $dx[k]$. We desire the projection $dF^n_\perp$ of dF^n that corresponds to a valid shape deformation $dx_\perp[k]$. We take 2 approaches to finding deformations $dx_\perp[k]$ corresponding to dF^n. The first is to perform numerical minimization of the distance between dF^n and $dF^n_\perp$. The second approach is to perform analytical projection of the flow in those cases where such a calculation is tractable, such as when the features are the set of all inter-point distances.

5. RESULTS

5.1. Segmentation

In this experiment we apply our segmentation framework to a synthetic bimodal image with very low signal-to-noise ratio (SNR=-17.5 db). We build our prior model shape distributions on a collection of four triangular shapes similar to the one used to generate the noisy image.

Figure 1 presents the result obtained when using our new prior as well as the results obtained by using the approach in [5] and a generic curve length penalization term $E_C = ||C||$. While the difference between the later two methods (B, C) is not significant, our method (A) gives a shape visually similar to the true shape. As a measure of segmentation error, the symmetric distance (in pixels) between the true boundary and the final result is shown on the top of each panel. While giving some improvement of this error measure, our result is significantly superior visually. The majority of the symmetric distance error for our prior is attributed to the error in location of the shape, while the error of other methods is due to intrinsic shape fluctuations caused by noise.

5.2. Mean shape calculation

We apply our framework to the problem of finding the mean of 2 triangles according to (4) for different choices of the distance measure. The aim is to derive a unique perceptual mean shape without using high level knowledge about topology or landmarks.

We compare our result with two traditionally used shape distance metrics. The first metric is the asymmetric inter-point shape distance defined as

$$d(C_1, C_2) = \int_{C_1} D(x, C_2) ds \tag{6}$$

where the integration is carried out along C_1 and $D(x, C_2)$ is the Euclidean distance between the point x on C_1 and the closest point on the contour C_2. In Figure 2 (A) we show the resulting mean shape corresponding to the two prior shapes, and as can be seen the result is not a triangle.

Another often used global shape difference measure is based on the total area between shapes:

$$d(C_1, C_2) = \int_{A:\ \mathrm{sign}(DT(C_1)) \neq \mathrm{sign}(DT(C_2))} dS \tag{7}$$

where $DT(C_1)$ and $DT(C_2)$ are signed distance transforms for shapes C_1 and C_2 respectively. When used in (4), this shape difference measure yields an infinite number of solutions for the mean shape. These solutions are located in the areas shaded in red in Figure 2 (B). In this case the solution of (4) is not unique.

The result of using our metric (1) is given in Figure 2 (C). We use the 2 features described in Section 3. We initialize the evolution process using the result obtained from (6). The contour produced by our iterative process is shown by the solid red line. The size of this shape is smaller than that of original shapes, which is consistent with the invariance of our measure to the size, location and rotation of the shapes under comparison. We manually scale and shift the contour to match the position of original shapes for visualization purpose. One can see that the scaled result corresponds well to the expected mean shape.

6. CONCLUSIONS

In this paper we present a novel distribution-based shape prior. Our shape prior is based on matching distributions of features calculated along the boundary of a curve. We embed our shape prior into a level set framework and derive a tractable evolution process corresponding to the prior. Our experiments show that our global distribution difference metric can solve the problem of finding the perceptual mean shape in situations when other methods using globally defined similarity criteria fail to yield an acceptable solution. Preliminary experiments on using our measure as a prior in image segmentation yields superior solution comparing to existing methods.

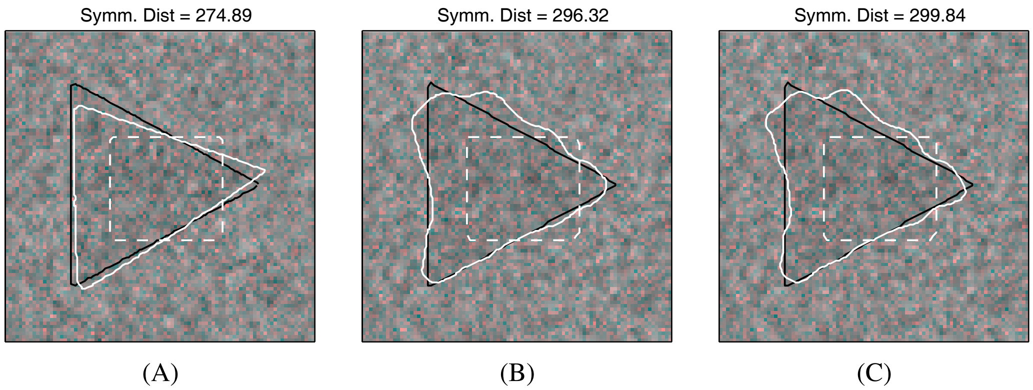

Fig. 1. Segmentation experiment. A: Our method; B: Method in [5]; C: Curve length penalty prior; White - final result; Black - true shape boundary; Dashed, white - initial curve. Symmetric distance (in pixels) between true boundary and final result is shown on the top of each panel.

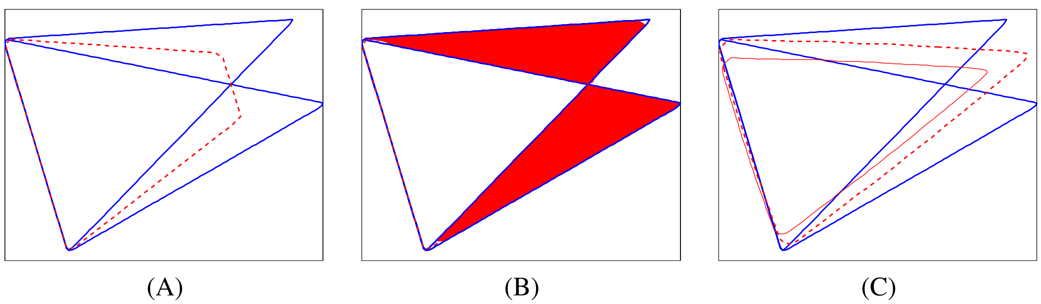

Fig. 2. Mean shape calculation using three shape difference measures. Two blue, solid contours correspond to prior shapes; red dashed line represents the mean shape; filled areas correspond to non-unique solution for the mean shape. (A) - asymmetric distance based measure; (B) - area based measure; (C) - our distribution difference measure (thin solid red line - evolution result; dashed red line - manually scaled result).

Acknowledgments

The authors would like to thank Jayant Shah of Northeastern University, for helpful discussions and advice.

7. REFERENCES

[1] A. Tsai, Yezzi A, Wells W, Tempany C, Tucker D, Fan A, Grimson W, and Willsky A, "Model-based curve evolution technique for image segmentation," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'01)*, 2001.

[2] E. Klassen, A.Srivastava, and W. Mio, "Analysis of planar shapes using geodesic paths on shape spaces," *IEEE Trans. on PAMI, in press.*

[3] Stephen M. Pizer, Daniel S. Fritsch, Paul A. Yushkevich, Valen E. Johnson, and Edward L. Chaney, "Segmentation, registration, and measurement of shape variation via image object shape," *IEEE Transactions on Medical Imaging*, vol. 18(10), pp. 851–865, 1996.

[4] D. Mumford and J. Shah, "Boundary detection by minimizing functionals," in *Proceedings of CVPR*, San Francisco, June 1985, pp. 22–26.

[5] Michael E. Leventon, W. Eric L. Grimson, Olivier Faugeras, and William M. Wells III, "Level set based segmentation with intensity and curvature priors," *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis Proceedings (MMBIA'00)*, pp. 4–11, 2000.

[6] S.-C. Zhu, "Embedding gestault laws in markov random fields," *IEEE trans. on PAMI*, vol. 21, no. 11, 1999.

[7] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions," *ACM transactions on graphics*, vol. 21(4), pp. 807–832, 2002.

[8] G. Sapiro and V. Caselles, "Histogram modification via partial differential equations," *International Conference on Image Processing Proceedings(ICIP'95)*, 1995.