

# Shape Representation via Harmonic Embedding

Alessandro Duci\*    Anthony J. Yezzi†    Sanjoy K. Mitter‡    Stefano Soatto\*

\* Computer Science Department, University of California at Los Angeles, Los Angeles – CA 90095

† Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta – 30332

‡ Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge - MA 02139

alessandro.duci@sns.it, ayezzi@ece.gatech.edu, mitter@mit.edu, soatto@ucla.edu

## Abstract

*We present a novel representation of shape for closed planar contours explicitly designed to possess a linear structure. This greatly simplifies linear operations such as averaging, principal component analysis or differentiation in the space of shapes. The representation relies upon embedding the contour on a subset of the space of harmonic functions of which the original contour is the zero level set.*

## 1. Introduction

The analysis and representation of shape is at the basis of many visual perception tasks, from classification and recognition to visual servoing. This is a vast and complex problem, which we have no intention of addressing in its full generality here. Instead, we concentrate on a specific issue that relates to the representation of closed, planar contours. Even this issue has received considerable attention in the literature. In particular, in their work on statistical shape influence in segmentation [19], Leventon et al. have proposed representing a closed planar contour as the zero level set of a function in order to perform linear operations such as averaging or principal component analysis. The contour is represented by the embedding function, and all operations are then performed on the embedded representation. They choose as their embedding function the signed distance from the contour (whose differential structure is described by the *non-linear* Eikonal Equation) and implement its evolution in the numerical framework of level sets pioneered by Osher and Sethian [26].

While this general program has proven effective in various applications, the particular choice of embedding function presents several difficulties, because signed distance functions are not a closed set under

linear operations: the sum or difference of two signed distance functions is not a signed distance function (an immediate consequence of their nonlinear differential structures). Consequently, the space cannot be endowed with a probabilistic structure in a straightforward manner, and repeated linear operations, including increments and differentiation, eventually lead to computational difficulties that are not easily addressed within this representation. Alternative methods that possess a linear structure rely on parametric representations, for instance various forms of splines [7], that do not guarantee that the general structure of the original shape (geometry and topology) is preserved in a neighborhood. Furthermore, such representations are not geometric, for they depend on the chosen parameterization.

In this paper, we present a novel representation of shape for closed planar contours that is geometric *and* explicitly designed to possess a (locally) linear structure. This allows linear operations such as principal component analysis or differentiation to be naturally defined and easily carried out. The basic idea consists of, again, representing the contour as the zero level set of a function, but this time the function belongs to a linear (or quasi-linear<sup>1</sup>) space. While previous methods relied on the (non-linear) Eikonal equation, ours relies on Laplace equation, which is linear. Our representation allows exploring the neighborhood of a given shape while guaranteeing that the topology and the geometry of the original shape is preserved.

We introduce the simplest form of this idea in Sect. 2, where we point out some of its difficulties. We then extend the representation to an anisotropic operator in Sect. 3, and discuss its finite-dimensional implementation in Sect. 4. Finally, we illustrate some of the properties of this representation in Sect. 5.

---

<sup>1</sup>The representation is restricted to a convex subset of a linear vector space.

## 1.1. Relation to previous work

The literature on shape modeling and representation is too vast to review in the limited scope of this paper. It spans at least a hundred years of research in different communities from mathematical morphology to statistics, geology, neuroanatomy, paleontology, astronomy etc. Some of the earlier attempts to formalize a notion of shape include D'Arcy Thompson's treatise "Growth and Form" [30], the work of Matheron on "Stochastic Sets" [23] as well as that of Thom, Giblin and others [29, 8]. The most common representations of shape rely on a finite collection of points, possibly defined up to equivalence classes of group actions [13, 18, 5, 22]. These tools have proven useful in contexts where distinct "landmarks" are available, for instance in comparing biological shapes with distinct "parts." However, comparing objects that have a different number of parts, or objects that do not have any distinct landmark, is elusive under the aegis of statistical shape spaces. Koenderink [17] is credited with providing some of the key ideas involved in formalizing a notion of shape that matches our intuition. However, Mumford has critiqued current theories of shape on the grounds that they fail to capture the essential features of perception [25].

"Deformable Templates," pioneered by Grenander [11], do not rely on "features" or "landmarks;" rather, images are directly deformed by a (possibly infinite-dimensional) group action and compared for the best match in an "image-based" approach [35]. Another line of work uses variational methods and the solution of partial differential equations (PDEs) to model shape and to compute distances and similarity. In this framework, not only can the notion of alignment or distance be made precise [3, 34, 24, 15, 27], but quite sophisticated theories that encompass perceptually relevant aspects can be formalized in terms of the properties of the evolution of PDEs (e.g. [16]). The work of Kimia et al. [14] describes a scale-space that corresponds to various stages of evolution of a diffusing PDE, and a "reacting" PDE that splits "salient parts" of planar contours by generating singularities. [14] also contains a nice taxonomy of existing work on shape and deformation and a review of the state of the art as of 1994.

The variational framework has also proven very effective in the analysis of medical images [21, 31, 20, 33]. Although most of the ideas are developed in a deterministic setting, many can be transposed to a probabilistic context. Scale-space is a very active research area, and some of the key contributions as they relate to the material of this paper can be found in [12, 28, 2, 1] and references therein. Leventon et al. [19] perform principal component analysis in the aligned

frames to regularize the segmentation of regions with low contrast in brain images. Similarly, [32] performs the joint segmentation of a number of images by assuming that their registration (stereo calibration) is given.

We present a novel representation of shape that supports linear operations. We only consider closed planar contours, and even within this set our representation cannot capture any shape; it does not include a notion of hierarchy or compositionality, which are crucial in a complete theory of shape. Despite its limitations that restrict the class of shapes and the analysis to their global properties, our representation has desirable features when it comes to linear analysis. In fact, it allows us to naturally take linear combinations of shapes; performing principal component analysis (PCA) on the embedding function results in a natural notion of deformation (we call it "principal deformation analysis", or PDA) on the underlying shapes. Endowing the space with a probabilistic structure, although not addressed in this paper, is greatly facilitated by the (quasi-)linear nature of the representation.

We note that, although we represent a contour as the zero level set of an embedding function, our approach is not a level set method in the traditional sense [26]: in fact, in local shape analysis we are interested in guaranteeing that changes of topology do *not* occur. In this sense, our approach is far less general, but in ways that are desirable for the specific problem we address, that of representing a neighborhood of a given shape.

As we will see in Sect. 4, our approach relies on a finite-dimensional set of boundary values at specified locations. In this sense, therefore, our technique could be thought of as an implicit version of splines [7, 6], in the sense that changing the location of the control points results in an evolution of the contour

## 2. Harmonic Embedding

The basic idea is to represent a closed planar contour,  $\gamma$ , as the zero level set of a function  $u$  that inherits the linear structure of its embedding space. This linear space is chosen to be the set of harmonic functions, which naturally leads to the contour being represented as the solution of certain Laplace equations. More formally, consider the domain  $\Omega \doteq \{x \in \mathbb{R}^2 : r < |x| < R\}$ ; we will restrict our attention to contours that are contained in such a domain, for some  $r, R \in \mathbb{R}$ . In particular, we will consider smooth contours  $\gamma \in \mathcal{C}_\Omega \doteq C^\infty(S^1, \Omega)$  from the unit circle to  $\Omega$ . We call the inner and outer boundaries of  $\Omega$ ,  $\partial_0\Omega \doteq \{x \in \mathbb{R}^2 : |x| = r\}$  and  $\partial_1\Omega \doteq \{x \in \mathbb{R}^2 : |x| = R\}$  respectively. Each contour is then represented by a function  $u : \Omega \rightarrow \mathbb{R}$ , and in particular by the value of this function in the inner and outer boundaries

$\mathfrak{h}_\Omega \doteq C^0(S^1) \times C^0(S^1)$ . Among all possible  $u$ , we restrict our attention to *harmonic functions*, i.e. to the set  $\mathbb{H}_\Omega \doteq \{u \in C^\infty(\Omega) \cap C^0(\bar{\Omega}) : \Delta u = 0\}$ .

**Definition 2.1 (Harmonic Embedding).** We say that a contour  $\gamma \in \mathcal{C}_\Omega$  has an *harmonic embedding* if there exists a function  $u \in \mathbb{H}_\Omega$  such that

$$\begin{cases} u(x) = 0 \text{ for } x \in \gamma, \\ u(x) \neq 0 \text{ for } x \in \bar{\Omega} \setminus \gamma. \end{cases} \quad (1)$$

We say that the function  $u$  is an harmonic representation associated to  $\gamma$ . The set of contours that admit an harmonic embedding will be indicated by  $\mathcal{C}_\Omega^*$ . We say that  $u \in \mathbb{H}_\Omega$  represents an harmonic contour if the zero level set of  $u$  is an element of  $\mathcal{C}_\Omega^*$ . The set of all the harmonic representations will be indicated by  $\mathbb{H}_\Omega^*$ .

In plain words, we plan to represent a contour  $\gamma$  by the values that a harmonic function  $u$ , that is zero on the contour, takes on the inner and outer boundaries  $\partial_0\Omega, \partial_1\Omega$ . Naturally, not all contours admit a harmonic embedding.

**Proposition 2.1.** *If  $\gamma \in \mathcal{C}_\Omega^*$  and  $u \in \mathbb{H}_\Omega^*$  is an harmonic function associated to  $\gamma$ , then  $u$  has a constant sign on each connected component of the boundary,  $\partial_0\Omega$  and  $\partial_1\Omega$ , where it takes opposite signs.*

*Proof.* The function  $u$  cannot be zero on the boundary of  $\Omega$ , because  $\gamma$  does not belong to  $\partial_0\Omega \cup \partial_1\Omega$  and the function  $u$  is continuous. This implies that it has constant sign on each connected component of the boundary. If  $u$  has the same sign on both the connected components of the boundary, then the zero level set must be empty as a consequence of the maximum principle.  $\square$

The relevance of this proposition is that, if we assign negative values to  $u$  on the inner boundary and positive values on the outer boundary, the maximum principle guarantees that the resulting zero level set is always simply connected. This is desirable in local shape analysis since we do not want small perturbations of a contour to result in changes of topology. Note that this feature is quite different than traditional level set methods that address more general deformations where changes of topology are desirable. Due to the uniqueness of solution to the Laplace equation, knowing  $u$  is equivalent to knowing its values at the boundaries of the set  $\Omega$ . Therefore, we can use as a representative of  $\gamma$  not the entire  $u$ , but the values  $f_0$  and  $f_1$  that  $u$  takes at the boundaries.

**Definition 2.2 (Boundary representation).** Let  $\mathfrak{h}_\Omega^* \subset \mathfrak{h}_\Omega$  such that

$$\mathfrak{h}_\Omega^* \doteq \{(f_0, f_1) \in \mathfrak{h}_\Omega : f_0 < 0, f_1 > 0\}. \quad (2)$$

The elements of  $\mathfrak{h}_\Omega^*$  are the harmonic shapes and we call the set  $\mathfrak{h}_\Omega^*$  the harmonic shape space.

The map  $\pi_\Omega : \mathbb{H}_\Omega^* \rightarrow \mathfrak{h}_\Omega$  associates to a harmonic function the rescaled boundary values. The map  $\psi_\Omega : \mathfrak{h}_\Omega^* \rightarrow \mathcal{C}_\Omega^*$  associates to a rescaled boundary condition the zero level set of the harmonic function with that boundary conditions. The map  $h_\Omega : \mathbb{H}_\Omega^* \rightarrow \mathcal{C}_\Omega^*$  associates to a harmonic function in  $\mathbb{H}_\Omega^*$  its zero level set in  $\mathcal{C}_\Omega^*$ .

*Remark 2.2.* The harmonic representation is invariant to scale factors; in fact, if  $u$  is the harmonic representation of a contour  $\gamma$  and  $\lambda \neq 0$  then also  $\lambda u$  is an harmonic representation of  $\gamma$ . Therefore, every contour  $\gamma$  admits an entire equivalence class of embedding functions  $u$ , and therefore  $f_0, f_1$ . In order to fix this ambiguity, one can rescale the energy of  $u$ , for instance by fixing the value of the integral of  $|f_0|, |f_1|$ . This will be illustrated in the implementation section (Sect. 4).

The following proposition guarantees that the representation we have introduced is linear when restricted to  $\mathfrak{h}_\Omega^*$ . It should also be noted that the set  $\mathfrak{h}_\Omega^*$  is convex.

**Proposition 2.3 (Linearity).** *Let  $u_1, u_2 \in \mathfrak{h}_\Omega^*$ , then there exists a positive number  $\lambda$  such that  $u_1 \pm \lambda u_2 \in \mathfrak{h}_\Omega^*$ .*

*Proof.* If we take  $\epsilon = \min_{\partial\Omega} |u_1|$ ,  $\mu = \max_{\partial\Omega} |u_2|$ , then  $\lambda = \frac{\epsilon}{2\mu}$  satisfies the requirement.  $\square$

### 3. Anisotropic extension

The definitions above indicate that  $u$  can be used as a linear embedded representation of  $\gamma$ . However, as we have pointed out, not all contours admit a harmonic embedding. In the experimental section 5 we give some examples of contours that are not easily captured by the straightforward harmonic representation. In this section, therefore, we extend the representation to a more general elliptic operator (see [9, 10]) instead of the Laplacian in eq. (1). In particular we use an operator of the form

$$\nabla \cdot (A(x, y) \nabla u) \quad (3)$$

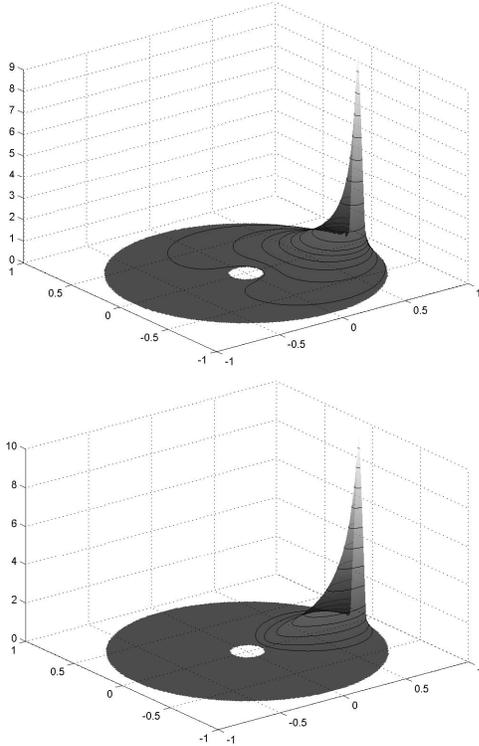
where  $A : \Omega \rightarrow \mathbb{R}^4$  is a differentiable function that associates to each point of the domain a  $2 \times 2$  matrix. It is immediate to see that if  $A(x, y) = I$ , the operator (3) reduces to the Laplacian in (1). Naturally, there are infinitely many choices of  $A(x, y)$ , and this is a power of the representation that is at the disposal of the designer. We choose the following operator:

$$A(x, y) \doteq \frac{1}{\rho^2} \begin{pmatrix} \lambda_1 x^2 + \lambda_2 y^2 & (\lambda_1 - \lambda_2)xy \\ (\lambda_1 - \lambda_2)xy & \lambda_2 x^2 + \lambda_1 y^2 \end{pmatrix} \quad (4)$$

where  $\rho = \sqrt{x^2 + y^2}$ ,  $\lambda_1$  and  $\lambda_2$  are two positive constants. The matrix  $A$  has the following property:

$$(x, y) \mapsto \lambda_1(x, y), \quad (y, -x) \mapsto \lambda_2(y, -x).$$

This choice may seem to come from “black magic,” but its motivation becomes clear by looking at Fig. 1. In our finite-dimensional implementation, described in Sect. 4, if we use the harmonic function  $\Delta u = 0$ , the basis elements are not well localized<sup>2</sup>, and therefore many coefficients are needed to represent even simple contours. The choice of the operator  $A$  above, on the other hand, helps localize and direct the energy of each basis element from the outer to the inner boundary.



**Figure 1.** Basis elements according to the model (1) (top), and the more general model with the operator (3) replacing the Laplacian, for  $\lambda_1 = 20$ ,  $\lambda_2 = 1$  (bottom).

#### 4. Finite-dimensional implementation

So far we have managed to avoid addressing the problem of how to compute the representation  $u$  (or its boundary values  $f_0, f_1$ ) from the contour  $\gamma$  from either (1) or (3). This is not an easy task, since it is not a standard boundary value problem: one could think of

<sup>2</sup>The meaning of the word “basis” will be made precise in the next section.

fixing the boundary value at the inner contour to, say,  $-1$ , then solve the Dirichlet problem between  $\partial_0\Omega$  and  $\gamma$ , and then attempt to extend  $u$  outside  $\gamma$  towards  $\partial_1\Omega$ . Unfortunately, the analytic continuation of  $u$  outside  $\gamma$  is severely ill-posed, and one cannot hope in general to reach the outer boundary for any given contour. A way to fix this would be to carry out the analytic continuation as far as possible (until a singularity develops) toward the outer boundary, then smooth the resulting analytic boundary values. Next, reverse the procedure to solve the Dirichlet problem between the smoothed outer analytic boundary and the contour  $\gamma$ , and then extend it inward toward the inner boundary. One could iterate the procedure back and forth, which would result in a very laborious and numerically ill-conditioned algorithm to find the representation of a given contour.

In this section, therefore, we seek to approximate  $\gamma$  by the zero level set of a function  $u$  that is guaranteed, by construction, to be analytical in  $\Omega$ . The linear structure of the representation comes handy at this point: in fact, we can seek for functions  $u$  of the form

$$u = \sum_{i=1}^n \alpha_{0i} u_{0i} + \alpha_{1i} u_{1i} \quad (5)$$

where  $u_{0i}$  and  $u_{1i}$  are a set of “basis” functions that are constructed to satisfy eq. (1) with the operator (3) replacing the Laplacian, and including the constraints of Remark 2.2 on the inner and outer boundary respectively. More in detail, in order to compute the basis elements, for a chosen positive integer  $n$ , we indicate with  $\{s_{0j}\}_{1 \leq j \leq n}$  and  $\{s_{1j}\}_{1 \leq j \leq n}$  respectively a partition of the inner and outer boundary of  $\Omega$ , with  $|s_{ij}|$  the length of the segment  $s_{i,j}$ , and seek for  $\{u_{ij}\}_{i=0,1; j=1 \dots n}$  that solve the following partial differential equation (PDE)

$$\begin{cases} \nabla \cdot (A(x, y) \nabla u_{ij}(x, y)) = 0 & \text{for } (x, y) \in \Omega \\ u_{ij}(x, y) = 0 & \text{for } (x, y) \in \partial\Omega \setminus s_{ij}, \\ u_{ij}(x, y) = \frac{1}{|s_{ij}|} & \text{for } (x, y) \in s_{ij}. \end{cases} \quad (6)$$

We have integrated this PDE numerically using standard finite-element methods [4], and in Fig. 1 we show sample basis elements for the case of the simple Laplacian  $A(x, y) = I$  (top), and the more general anisotropic operator (4) (bottom).

Note that this PDE needs only to be solved *once* and off-line. Once the basis elements are known, each contour will be represented by the coefficients  $\{\alpha_{ij}\}$ , which can be found following the procedure that we describe next.

First, we need to set up a cost functional that measures the discrepancy between the target contour that we want to represent,  $\gamma$ , and the model contour that

admits a harmonic embedding. A simple cost function is simply the Lebesgue measure of the set symmetric difference<sup>3</sup> between the inside of the contour  $\gamma$ , which we indicate with  $\overset{\circ}{\gamma}$ , and the set of points that correspond to a negative value of  $u$ , which we indicate by  $u^-$ . Other choices of cost functionals are possible, for instance the Hausdorff distance, and we choose the set symmetric difference only for simplicity. Now, ideally one would want to write this functional explicitly in terms of  $\{\alpha_{ij}\}$ , and differentiate it to yield a gradient descent algorithm. Unfortunately, it is not easy to write the zero level set of  $u$  as a function of the parameters  $\alpha$ . However, one can achieve the same results by first computing the gradient flow of the cost functional with respect to an unconstrained  $u$ , then projecting this flow onto the basis  $\{u_{ij}\}$ . More formally, consider a general infinitesimal contour variation  $C : [0, 1] \times [-\epsilon, \epsilon] \rightarrow \Omega$ ; then the gradient descent evolutions of the symmetric difference the target set  $\overset{\circ}{\gamma}$  and the approximating set  $\overset{\circ}{C}$  is given by

$$C_t = \left( \mathcal{X}_{\overset{\circ}{\gamma}} - \frac{1}{2} \right) N$$

where  $N$  is the outward unit normal at  $C(\cdot, 0)$  and  $\mathcal{X}_{\overset{\circ}{\gamma}}$  is the characteristic function of the set  $\overset{\circ}{\gamma}$ . After projecting onto the basis elements  $u_{ij}$  via  $\int_C \langle u_t, u_{ij} \rangle ds$ , one obtains the update formula for the coefficients

$$\alpha'_{ij} = \alpha_{ij} + \mu \int_{C(\cdot, 0)} \frac{u_{ij}}{|\nabla u|} \left( \mathcal{X}_{\overset{\circ}{\gamma}} - \frac{1}{2} \right) ds, \quad (7)$$

where  $\mu$  is a positive constant to be chosen as a design parameter. Finally, one has to enforce the conditions on the sign at the boundary value, described in Proposition 2.1; in this case we simply have:

$$\alpha_{0j} < 0, \quad \alpha_{1j} > 0 \quad \forall j \in \{1, \dots, n\}.$$

We obtain the normalization, with respect to the scale factor of the remark 2.2, dividing by

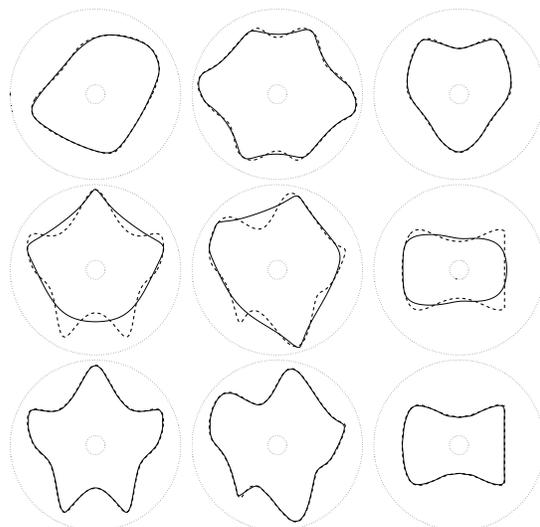
$$\sum_{i=1,2; j=1, \dots, n} |\alpha_{i,j}|.$$

Given a target contour  $\gamma$ , this procedure allows us to estimate its representation  $\alpha_{i,j}$ . In the next section we present some experiments that illustrate the features of this representation.

## 5. Experiments

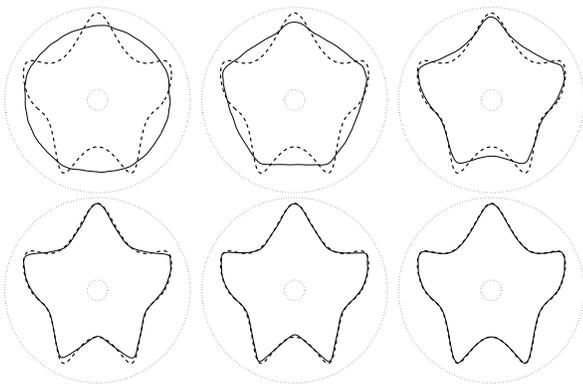
In this section we report some experiments that illustrate the power and limitations of the representation

<sup>3</sup>The *symmetric difference* between two set  $A$  and  $B$  is  $A \Delta B \doteq (A \setminus B) \cup (B \setminus A)$ .



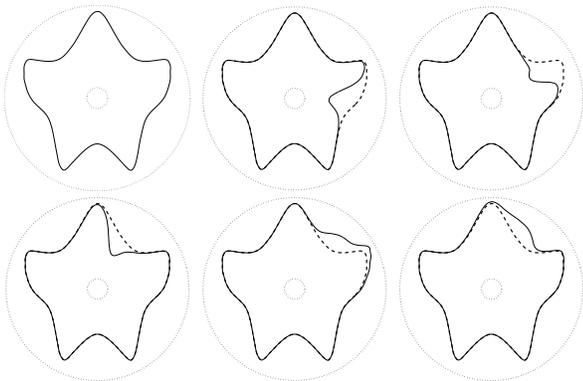
**Figure 2.** (Top) Original contours (dashed lines) in the domain  $\Omega$  enclosed between the inner circle  $\partial_0\Omega$  and the outer circle  $\partial_1\Omega$ . The solid line indicates the best approximation based on harmonic embedding (1). The contours in the second row, however, are not as well captured by the simplest representation. Substituting the anisotropic operator (3) instead of the Laplacian results in improved representational power of our scheme (bottom).

proposed. In Fig. 2 (top) we show a few planar contours (dashed) together with the domain  $\Omega$ , delimited by the inner and outer boundary circles, and - superimposed - the best approximation of the target contour based on the model (1). While these contours are faithfully captured in the representation, the contours in the second row are clearly not well approximated. In the following row we show the same target contours and their approximation using the anisotropic operator (3) instead of the Laplacian. As it can be seen, the target contours are much better approximated. Figure 3 shows the evolution of the contour from an initial estimate (the outer boundary) to the minimum of the set-symmetric difference as described in Sect. 4. One of the strengths of our representation is that it admits a linear structure. Therefore, the computation of increments of a shape (Fig. 4), the average of two or more shapes (Fig. 5) and the analysis of the principal modes of deformation of a shape (Fig. 6) are straightforward to compute. We would like to stress that what is important here is not the result, that is the final *value* of the principal components shown in Fig. 6, but the *procedure* followed to get such principal components. Since our representation is linear by construction, one can simply perform PCA on the coefficients, and automatically get a geometric representation of the de-

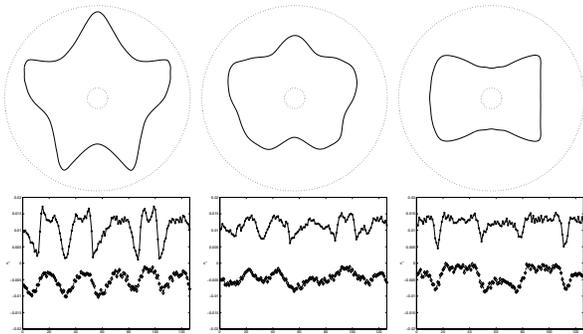


**Figure 3.** Evolution of an initial estimate (the outer contour) to the minimum of the set-symmetric difference that is used to find the best harmonic embedding.

formation that guarantees that no changes of topology occur. Figure 4 is perhaps the most crucial one to il-



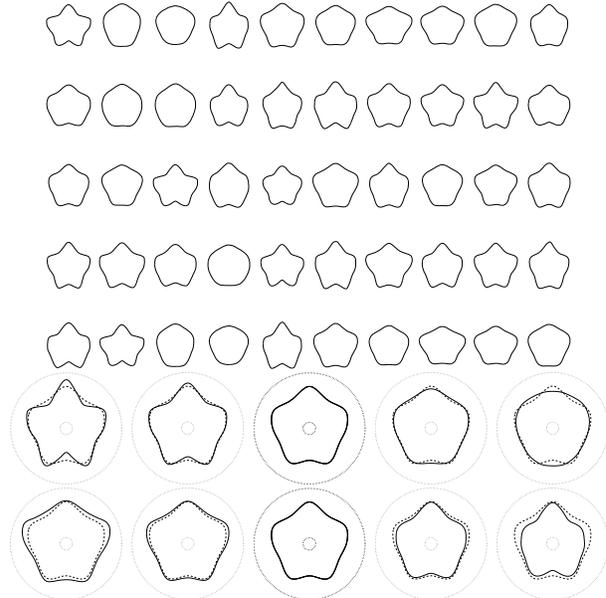
**Figure 4.** Directional derivatives of a shape can be easily computed in the framework of (anisotropic) harmonic embedding. For the figure in (top-left), we show local perturbations along the direction of five basis elements.



**Figure 5.** (Top) The average shape between (left) and (right) can be computed by simply averaging (center) the corresponding coefficients (bottom)  $\{\alpha_{ij}\}$ .

lustrate the features of our method. In fact, if we were

to use parametric representations of a curve, for instance splines, by perturbing the control points we are not guaranteed that the contour does not develop singularities or self-intersects. Our framework enjoys the benefits of splines in having a finite-dimensional linear representation, but is geometric and perturbations are guaranteed not to result in changes of topology.



**Figure 6.** Principal modes of deformation can be easily computed by principal component analysis of the embedded representation. On the top five rows we show a few samples of a collection of shapes. On the bottom in the first row we show the average shape (middle), plus (right) and minus (left) once and twice the first principal component. On the second row we see the effect of the second principal component (left-right) on the mean shape (middle).

## 6. Summary and Conclusions

We have presented a novel representation for the shape of closed planar contours. It does not rely on features or landmark points, and it is characterized by an underlying linear structure that makes linear operations straightforward. We showed examples of computation of the average, increments and principal deformations. Our representation has many shortcomings and it should not be taken as a general tool for shape analysis. It is not invariant with respect to the action of a group (and therefore one has to assume that the shapes are “registered”), it does not possess a notion of compositionality, and it cannot capture many complex contours of practical interest. However, to the best of our knowledge, it is the first and only variational shape

representation to possess a linear structure, which is extremely desirable when it comes to numerical implementation of linear operations.

## Acknowledgments

This research is supported by NSF IIS-0208197/9876145, ONR N00014-02-1-0720, Intel 8029, NIH Pre-NPEBC, and AFOSR F49620-03-1-0095.

## References

- [1] L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel. Axioms and fundamental equations of image processing. *Arch. Rational Mechanics*, 123, 1993.
- [2] L. Alvarez, J. Weickert, and J. Sanchez. A scale-space approach to nonlocal optical flow calculations. In *In ScaleSpace '99*, pages 235–246, 1999.
- [3] R. Azencott, F. Coldefy, and L. Younes. A distance for elastic matching in object recognition. *Proc. 13th Intl. Conf. on Patt. Recog.*, 1:687–691, 1996.
- [4] S. Brenner and L. Scott. *The Mathematical Theory of Finite Element Methods*, volume 15 of *Texts in Applied Mathematics*. Springer, 2002.
- [5] T. K. Carne. The geometry of shape spaces. *Proc. of the London Math. Soc. (3)* 61, 3(61):407–432, 1990.
- [6] D. Cremers, F. Tischhäuser, J. Weickert, and C. Schnörr. Diffusion Snakes: Introducing statistical shape knowledge into the Mumford–Shah functional. *Int. J. of Computer Vision*, 50(3):295–313, 2002.
- [7] C. de Boor. *A practical guide to splines*. Springer Verlag, 1978.
- [8] P. Giblin. *Graphs, Surfaces and Homology*. Chapman and Hall, 1977.
- [9] E. Giusti. *Metodi diretti nel calcolo delle variazioni*. Unione Matematica Italiana, 1994.
- [10] E. Giusti. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 1998.
- [11] U. Grenander. *General Pattern Theory*. Oxford University Press, 1993.
- [12] P. Jackway and R. Deriche. Scale-space properties of the multiscale morphological dilation/erosion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(1):38–51, 1996.
- [13] D. G. Kendall. Shape manifolds, procrustean metrics and complex projective spaces. *Bull. London Math. Soc.*, 16, 1984.
- [14] B. Kimia, A. Tannebaum, and S. Zucker. Shapes, shocks, and deformations I: the components of two-dimensional shape and the reaction-diffusion space. *Int'l J. Computer Vision*, 15:189–224, 1995.
- [15] R. Kimmel and A. Bruckstein. Tracking level sets by level sets: a method for solving the shape from shading problem. *Computer Vision, Graphics and Image Understanding*, (62)1:47–58, 1995.
- [16] R. Kimmel, N. Kiryati, and A. M. Bruckstein. Multi-valued distance maps for motion planning on surfaces with moving obstacles. *IEEE Trans. Robot. & Autom.*, 14(3):427–435, 1998.
- [17] J. J. Koenderink. *Solid Shape*. MIT Press, 1990.
- [18] H. Le and D. G. Kendall. The riemannian structure of euclidean shape spaces: a novel environment for statistics. *The Annals of Statistics*, 21(3):1225–1271, 1993.
- [19] M. Leventon, E. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours, 2000.
- [20] R. Malladi, R. Kimmel, D. Adalsteinsson, V. C. G. Sapiro, and J. A. Sethian. A geometric approach to segmentation and analysis of 3d medical images. In *Proc. Mathematical Methods in Biomedical Image Analysis Workshop*, pages 21–22, 1996.
- [21] R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(2):158–175, 1995.
- [22] K. V. Mardia and I. L. Dryden. Shape distributions for landmark data. *Adv. appl. prob.*, 21(4):742–755, 1989.
- [23] G. Matheron. *Random Sets and Integral Geometry*. Wiley, 1975.
- [24] M. I. Miller and L. Younes. Group action, diffeomorphism and matching: a general framework. In *Proc. of SCTV*, 1999.
- [25] D. Mumford. Mathematical theories of shape: do they model perception? In *In Geometric methods in computer vision, volume 1570*, pages 2–10, 1991.
- [26] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi equations. *J. of Comp. Physics*, 79:12–49, 1988.
- [27] C. Samson, L. Blanc-Feraud, G. Aubert, and J. Zerubia. A level set model for image classification. In *International Conference on Scale-Space Theories in Computer Vision*, pages 306–317, 1999.
- [28] B. ter Haar Romeny, L. Florack, J. Koenderink, and M. V. (Eds.). Scale-space theory in computer vision. In *Lecture Notes in Computer Science*, volume Vol 1252. Springer Verlag, 1997.
- [29] R. Thom. *Structural Stability and Morphogenesis*. Benjamin; Reading, 1975.
- [30] D. W. Thompson. *On Growth and Form*. Dover, 1917.
- [31] P. Thompson and A. W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imaging*, 15(4):402–417, 1996.
- [32] A. Yezzi and S. Soatto. Stereoscopic segmentation. In *Proc. of the Intl. Conf. on Computer Vision*, pages 59–66, 2001.
- [33] A. Yezzi and S. Soatto. Deformotion: deforming motion, shape average and the joint segmentation and registration of images. *Intl. J. of Comp. Vis.*, accepted, 2003.
- [34] L. Younes. Computable elastic distances between shapes. *SIAM J. of Appl. Math.*, 1998.
- [35] A. Yuille. Deformable templates for face recognition. *J. of Cognitive Neurosci.*, 3(1):59–70, 1991.