Shape from Phase: An Integrated Level Set and Probability Density Shape Representation

John Corring
Computer and Information Science and Engineering
University of Florida
Email: corring@cise.ufl.edu

Anand Rangarajan
Computer and Information Science and Engineering
University of Florida
Email: anand@cise.ufl.edu

***Abstract*—The past twenty years has seen the explosion of the "shape zoo": myriad shape representations, each with pros and cons. Of the varied denizens, distance transforms and density function shape representations have proven to be the most utile. Distance transforms inherit the numerous geometric advantages of implicit curve representations while density functions are unmatched in their approach to the modeling of uncertainty and noise in shape features. We have not seen much rapprochement between these two representations in general. In this work, we introduce a complex wave representation (CWR) of shape which has the ability to simultaneously carry probabilistic information via its magnitude and geometric information via its phase, achieving an integration of distance transforms and density function shape representations. The CWR is a parametric representation with cluster centers akin to a mixture model and curve normal information akin to signed distance functions. We demonstrate the perceptual gains of the CWR, highlight the advantages of the probabilistic aspect for noisy shape alignment by a likelihood approach, and fusing both aspects we show that the CWR leads to a feature space in which kernel PCA yields approximate closed curves and probability density functions.**

I. Introduction

The construction of shape representations from unorganized point-sets has long been a staple of pattern recognition, computer vision, medical image analysis and related fields. A variety of shape representations ranging from annotated, sparse landmarks to full blown graph and mesh representations have been deployed over the past twenty years. The myriad representations in this "shape zoo" have in common the fundamental goal of imposing a specific structure on unorganized point-sets with each successful representation accruing gains in different shape matching, indexing and retrieval metrics.

In the shape zoo, the distance transform [1] and the probability density function [2], or shape distribution model, representations have marked out disjoint territories. Their lack of overlap has hindered co-development. Distance transforms, or distance functions, construct a shape representation wherein a bounded domain is parsed into a set of closed, nonintersecting curves with each location marked by the distance to the closest point on a curve. Allowing the field to be signed imbues the field with information that can be used to distinguish shape regions from their surroundings. This scalar field is subsequently used in shape applications such as registration and matching. The hallmark of distance transforms is the implicit representation of a set of curves embedded in 2D via a scalar field. Probability density functions also construct scalar field shape representations but focus instead on assigning a "probability mass" at each location based on the number of points within a neighborhood. This scalar field is subsequently used in shape applications such as registration and indexing. The hallmark of density functions is the implicit grouping of a set of points via an integrable and differentiable scalar field. An orienting generalization that can be made at this juncture is that distance transforms emphasize curve geometry whereas density functions embody location uncertainty.

The lack of rapprochement between distance transforms and density functions stems from their very different perspectives on shape representation: the main advantage of the distance transform is its implicit curve representation whereas that of the density function is its representation of uncertainty (and noise). The principal goal of this paper is *the integration of these two perspectives*: we seek a unified representation which simultaneously embodies curve information while continuing to handle uncertainty. Such a representation would be immensely beneficial to shape analysis provided the respective advantages of both the distance and density functions are preserved. To set the stage, we first turn to the mathematical underpinnings of the distance function since they hold the key to subsequent integration.

Distance transforms satisfy the static Hamilton-Jacobi equation $\|\nabla S(x)\| = 1$ where $S(x)$ is the distance function. If the *signed* distance function is sought, the zero level set of $S(x)$ is a set of curves embedded in 2D (or a set of surfaces embedded in 3D). The difficulty of computing the signed distance function from an unorganized set of points is well known—essentially involving a search over all possible sets of non intersecting curves embedded in 2D (or higher dimensional counterparts in 3D). What is not so well known is the curious fact that the static Hamilton-Jacobi equation—a nonlinear differential equation—is closely related to the static Schrödinger equation—a linear differential equation [3]. It turns out that the distance function Hamilton-Jacobi scalar field $S(x)$ is approximately the *phase* of the complex Schrödinger wave function $\psi(x)$ which in turn is the solution to the wave equation

$$-\hbar^2 \nabla^2 \psi = \psi \tag{1}$$

with the approximation becoming increasingly more accurate

as $\hbar \to 0$. Note that a relationship has only been established between the Hamilton-Jacobi field $S(x)$ and the *phase* of the wave function $\psi(x)$. Since a wave function magnitude is related to a normalizable density function, it is natural to ask whether probabilistic information concerning a shape can be embedded in the wave function magnitude. To answer this question, we turn to the estimation of density functions next.

Density functions estimated from unorganized point-sets come in both parametric and non-parametric flavors. We find shape densities in the form of histograms, mixtures of Gaussians, wavelets and kernel expansions used in the literature [4], [5]. If a mixture of Gaussians is sought, the density function $p(x)$ is peaked at a set of 2D (or 3D) "cluster centers" with the degree of "peakedness" depending on the variance of the underlying cluster. The difficulty of computing the mixture density function from an unorganized set of points is well known—essentially involving a search for cluster exemplars and associated covariance matrices. The question at hand is whether or not we can associate a density function $p(x)$ with the squared magnitude of the wave function with the phase of the wave function continuing to play the role of the distance function. If the answer is in the affirmative, we have a candidate for an integrated shape representation with the wave function magnitude and phase representing location uncertainty and curve geometry respectively.

Distance functions have been put to very productive use in the literature. On the mathematical applications side, in [6] a distance transform of missing regions is used for transporting inpainting information while in [7] distance transforms fuse point-sets into connected meshes—clever applications of the nonlocality of the distance transform beyond the classical use as a vehicle for level-set frameworks [8]. In [9] a new mathematical approach to computing the distance transform is developed, widening the footprint of shape modeling. Recently, a host of new approaches based on partial differential equations ([10], [11], [12]) have gained traction. In [13], [14], shape skeletons, their close association with distance transforms and new scalar fields possessing related properties are all formulated. Most of the techniques above boast an appreciable perceptual or computational gain over traditional shape computation methods.

The perceptual organization literature has also been a rich source of inspiration for the present work. It is difficult to do justice to the past forty years of work in this area—here we attempt a very brief summary. Perceptual organization draws on well founded Gestalt principles of grouping and suggests local mechanisms for connectedness that can lead to global properties (closed curves, illusory contours etc.). While the mathematical principles underlying the mechanisms can vary widely (differential equations, tensor voting, graph partitioning etc.) [11], [15], [16], this literature has been mainly responsible for stressing the importance of grouping and connectedness in vision. Our work is in the same spirit in introducing a new complex wave representation (CWR) of point-sets wherein closed contours emerge from the wave function phase.

In this work, we show that a complex wave function $\psi(x)$ can indeed integrate both shape distribution and level-set information in the magnitude and phase respectively. Drawing upon previous work with mixture models, we design a complex wave function with unknown "cluster exemplars" as free parameters that have to be estimated using (for example) the maximum likelihood principle. Departing from previous work, we include a complex term with unknown curve normal parameters which also have to be estimated from an unorganized point-set. Provided efficient estimation of location (cluster center) and curve (normal) parameters can be achieved—with initial evidence contained herein (see Section IV. A) being quite encouraging—the CWR of shape portrays a highly accurate near and far field with the benefit of linear updating and an easily computable reproducing kernel. In this paper we assume this aspect of the representation is provided—focusing instead on application of the resulting curve representation.

II. Underpinnings of the Representation

We begin by summarizing previous work which introduced an approximation to the unsigned distance function [9] by solving the static Schrödinger equation corresponding to the static Hamilton-Jacobi equation $\|\nabla S\| = 1$:

$$S_\tau(x) \approx -\tau \log \phi_\tau(x; \boldsymbol{\mu}) = -\tau \log \sum_{k=1}^{N} \exp\left\{-\frac{\|x-\mu_k\|}{\tau}\right\}, \tag{2}$$

where $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^N$ is a collection of locations and the scalar field $\phi_\tau(x; \boldsymbol{\mu})$ is the solution to the linear differential equation

$$-\tau^2 \nabla^2 \phi_\tau(x; \boldsymbol{\mu}) + \phi_\tau(x; \boldsymbol{\mu}) = \sum_{k=1}^{N} \delta(x - \mu_k). \tag{3}$$

In (3), τ is a free parameter and the approximation S_τ becomes increasingly accurate as $\tau \to 0$. Since addition is a permitted operation on ϕ_τ, superpositions of solutions are allowed, in sharp contrast to standard distance transforms which do not permit addition.

In the present work, we seek to go beyond the unsigned distance transform (and linear differential equation approximations thereof). In shape analysis, connectedness is fundamental to applications, but is rarely available explicitly from the representation. Unsigned distance transforms solve a wavefront equation that is not suited to dealing with issues of connectedness. The approximation in (3) does not fare any better since it is based on an isotropic Green's function solution evaluated at a point-set. To handle connectedness, ϕ_τ must be modified. Drawing inspiration from the complex nature of wave functions in physics, we introduce a complex modulation factor to (2) that encodes *normal* information. Intuitively, we can "delocalize" the normal of a shape by propagating the phase as suggested by Huygens' principle. This leads to a phase factor $\exp\{i\frac{\nu_k^T(x-\mu_k)}{\lambda}\}$ modulating the real function ϕ_τ above (please see [18] for a more careful consideration of this term).

The modulation acts as a local curve regularization or control factor for the point-set, borne out in level sets of the

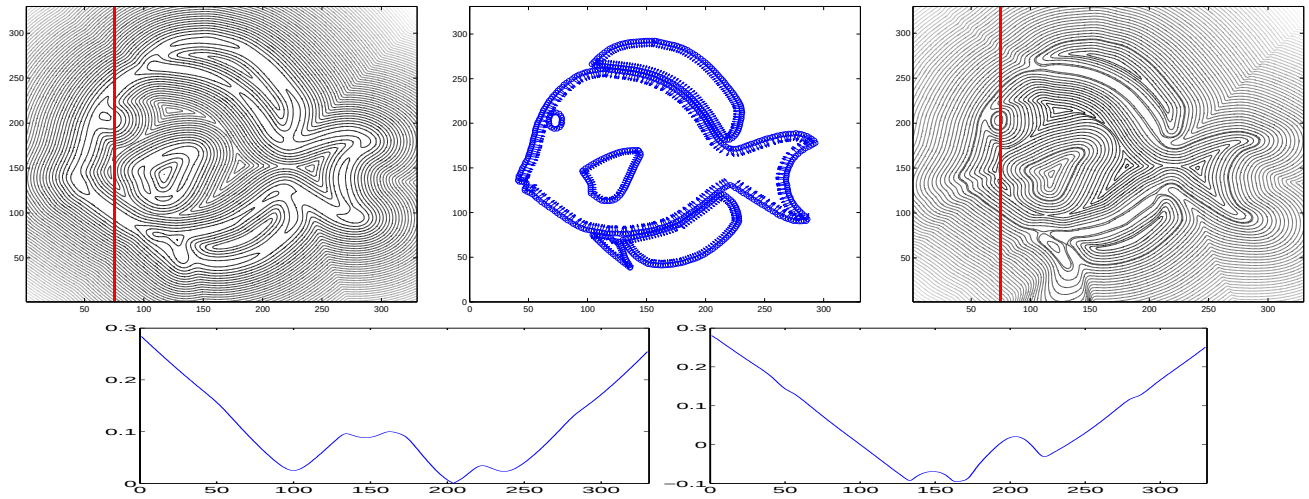

Fig. 1. **Visualization of** ψ. Left: level sets of the unsigned distance transform. Center: oriented point-set. Right: level sets of the modular distance transform. In the second row, the scanlines indicated by the red lines are shown. Near the point locations, the level set of the modular distance transform is much clearer while for the unsigned distance function, the normal information is totally inaccurate near the data points. Conforming to the continuity requirement for perceptual grouping [17] is a primary attribute of ψ. In the unsigned distance transform, the pectoral fin (the small fin below the gill fin) is basically indiscernible whereas in the modular distance transform it is very clear. The point-set used is a sampling from the GatorBait data set[1], specifically *Acanthuridae Acanthurus Chronixis* composed of a linear superposition of shapes that individually form closed curves (refer to Section III. C).

phase of the complex wave. Note that the normal information ν_k is attached to the center μ_k and therefore is an additional set of parameters constraining the wave function ψ:

$$\psi_{\tau,\lambda}(x;\boldsymbol{\mu},\boldsymbol{\nu}) = \sum_{k=1}^{N} \exp\left\{-\frac{\|x-\mu_k\|^2}{\tau} + i\frac{\nu_k^T(x-\mu_k)}{\lambda}\right\}, \tag{4}$$

where $\boldsymbol{\nu} = \{\nu_k\}_{k=1}^N$ is a collection of unit normal vectors. Here we use the Gaussian kernel instead of the exponential kernel since the differential equation aspect is less important. Writing $C = \{\boldsymbol{\mu},\boldsymbol{\nu}\}$ to denote a set of points and the associated unit normal vectors at these points, we call such a collection an *oriented point-set*. The symbol N will be used to represent the number of oriented points in C, with N_i referring to the points in C_i when considering multiple oriented point-sets. Henceforth, oriented point-sets have *unit normals* unless otherwise specified. We will write $\psi(x;C)$ as shorthand for $\psi_{\tau,\lambda}(x;\boldsymbol{\mu},\boldsymbol{\nu})$ with τ and λ suppressed wherever not explicitly needed. The wave function $\psi(x;C)$ contains connectedness information of the curve through the level sets of the phase, probability density information via the squared magnitude, and distance information through the logarithm of the magnitude (as $\lambda \to \infty$ and $\tau \to 0$). (Please see [9] for more details regarding the latter.)

A technical issue arises due to wrapped nature of the 2D wave function phase. At any location x, we obtain the *modular distance* along the normal vector to the zero level set of the phase

$$\tilde{d}(x;C) = \lambda \arctan\left(\frac{\mathrm{Im}[\psi(x;C)]}{\mathrm{Re}[\psi(x;C)]}\right). \tag{5}$$

Note that the phase (carrying orientation data) is now a property of the field, see Fig. 1, and is therefore defined everywhere. The unsigned distance transform obtained from a point-set, despite also being defined everywhere, lacks the crucial connectedness information, causing its zero level-sets to be broken islands marooning the original points. The connectedness component afforded by the phase is critical to shape boundary representation and perceptual grouping.

[1]http://www.cise.ufl.edu/~anand/GatorBait_100.tgz

III. Properties of the Representation

A. Attributes Encoded by ψ

A principal advantage of using distance transforms is the integration of point information via a field. Nonlocal analysis of shapes is enabled by this property. Unfortunately, the tight constraints imposed by the distance transform (such as $\|\nabla S\| = 1$) do not permit averaging, component analysis and the like, thereby limiting the effectiveness of the representation. The wave representation ψ allows for superposition and other operations (see Table I) enabling a richer variety of potential applications than standard distance transforms. With distance information in the magnitude along with connectedness and orientation information provided by the phase, ψ preserves the attractive properties of the signed distance function.

B. Analysis of ψ

There are some similarities between the wave function ψ (with a Gaussian kernel) and a Gabor wavelet. The latter has been extensively studied and used with great success in pattern recognition [19]. Gabor wavelets have primarily been used as *function* approximations and not in the *distributional* sense as we have used them.

$\psi(x;C)$ is continuous almost everywhere, and is defined on all of $\mathbb{R}^2$, so the L_2 norm of $\psi(x)$ can be computed

TABLE I
Technical Layout of the Operations on ψ

Unsigned Distance	$d^2(x;C) \approx -\tau \log(\psi\overline{\psi})$
Modular Distance (MD)	$\tilde{d}(x;C) = \lambda \arctan\left(\frac{\mathrm{Im}[\psi(x;C)]}{\mathrm{Re}[\psi(x;C)]}\right)$
Curve Geometry	$n(x;C) = \lambda\nabla \arctan\left(\frac{\mathrm{Im}[\psi(x;C)]}{\mathrm{Re}[\psi(x;C)]}\right)$
Sampling Probability	$p(x;C) = \lvert\psi\rvert^2/\lVert\psi\rVert_2^2$
Spatial Variance and Frequency Parameters	τ, λ
MD Linearity	$\psi_3(x;C_3) = \psi_1(x;C_1) + \psi_2(x;C_2) \rightarrow \tilde{d}(x;C_3) = \tilde{d}(x;C_1 \bigcup C_2)$
Kernel	$k((\mu_j,\nu_j),(\mu_k,\nu_k)) = \tau\pi \exp\left\{-\frac{\lVert\mu_j-\mu_k\rVert^2}{2\tau} - \frac{\tau\lVert\nu_j-\nu_k\rVert^2}{8\lambda^2} + \frac{i(\nu_j+\nu_k)(\mu_j-\mu_k)}{2\lambda}\right\}$

easily using properties of the Gaussian integral and Fourier transform:

$$\|\psi(x;C)\|_2^2 = 2\tau\pi \sum_{k=1}^{N}\sum_{j\geq k}^{N} \cos\left(\frac{(\nu_j+\nu_k)^T(\mu_k-\mu_j)}{2\lambda}\right) \exp\left\{-\frac{\|\mu_k-\mu_j\|^2}{2\tau} - \tau\frac{\|\nu_k-\nu_j\|^2}{8\lambda^2}\right\}. \quad (6)$$

We can use the inner product (k) from Table I to get

$$\langle\psi(x;C_1),\psi(x;C_2)\rangle = \sum_{k=1}^{N_1}\sum_{j=1}^{N_2} k((\mu_k,\nu_k),(\mu_j,\nu_j)) \quad (7)$$

which we use as a reproducing kernel for comparing shapes with position and normal parameters given by C_1 and C_2. The function ψ at present is unnormalized but can be turned into a square-root density (a probability amplitude) via normalization ($\psi \rightarrow \frac{\psi}{\|\psi\|_2}$) which can then be used to compute shape descriptors dependent on probabilistic features of a shape—such as moment based features. Other interesting applications of this aspect of ψ include maximum likelihood shape matching, wherein the negative log likelihood resulting from the probability density can be used as an objective function to choose an optimal shape correspondence without establishing point to point correspondences.

C. ψ for Oriented Multi-curve Shapes

As discussed above, ψ has unique properties (relative to distance functions) stemming from additivity of the representation, leading to a high level additivity or "superimposability". Depending on the choices of the free parameters τ and λ, modifying a shape with new position and orientation data can be very easy. We briefly justify the viability of this attractive property, and its limitations, below.

When

$$\tilde{d}(x) = \lambda \arctan\left(\frac{\mathrm{Im}[\psi(x;C)]}{\mathrm{Re}[\psi(x;C)]}\right) \quad (8)$$

$$= \lambda \arctan\left(\frac{\sum_{k=1}^{N} \sin(\frac{\nu_k^T(x-\mu_k)}{\lambda}) \exp\{-\frac{\|x-\mu_k\|^2}{\tau}\}}{\sum_{k=1}^{N} \cos(\frac{\nu_k^T(x-\mu_k)}{\lambda}) \exp\{-\frac{\|x-\mu_k\|^2}{\tau}\}}\right) \quad (9)$$

is evaluated, the contribution of each of the cluster centers to the sum decays exponentially; the slow growth of the arctangent yields stability to small contributions. To see what this means for superposition, consider an oriented point-set C_1 and let C_2 be a new oriented point-set to be superimposed. Let q_1 be the zero level set of the unwrapped $\tilde{d}(x;C_1)$ and q_2 be the zero level set of the unwrapped $\tilde{d}(x;C_2)$. Provided that $p(x;C_2) \ll p(x;C_1),\ \forall x \in q_1$ and $p(x;C_1) \ll p(x;C_2),\ \forall x \in q_2$, the superposition of C_1 and C_2 is stable: the resulting zero level sets approximately match $q_1 \cup q_2$.

The takeaway from this is that multiple curves can often be "added" easily: if one has multiple CWRs, then provided that the properties detailed above hold, one can compute the field by simply adding their fields together. Under this operation, the stability of the level sets depends on the distance to the initial set and the free parameters. When fields interact with each other and the above fails, then point discontinuities can arise in the phase field of ψ. However, provided that the abutting shapes have agreeing normal information (as in Fig. 1), the resulting superposition can maintain the desirable features of each of the underlying sets.

In high curvature areas the frequency of the oscillatory part and spatial accuracy of the density play a key role in the level sets of $\tilde{d}$. If the sampling of the curve location or normal data is insufficient, the superposition limitation mentioned above kicks in and the curve may be grouped incorrectly. On the other hand, as superimposed shape boundaries abut, the phase-driven orientation plays a more and more significant role in the perceptual qualities of the field and the level sets as outlined above. The parameters τ and λ act as intrinsic uncertainty parameters between the multi-curve and high curvature paradigms of shapes—a first order extension of their interpretation of uncertainty in spatial and normal information. A route to mitigating the abutment issue (a universal phenomenon in multi-curve representations) is allowing non-uniform frequency and spatial parameters to control the degree of precision of $\tilde{d}$.

IV. Experiments

A. Maximum Likelihood Alignment with $|\psi|^2$ as a Density

The probabilistic interpretation of ψ naturally has the flavor of a mixture of Gaussians. However, ψ has curve normal information and the squared magnitude of ψ is not actually a

mixture of Gaussians. The latter indeed has eigenvector information in the covariance matrix but this *cannot be interpreted as the normal to a curve*: eigenvectors have direction but not orientation since e_k and $-e_k$ are the same eigenvector. In the complex wave ψ, the normal is directly encoded into the representation, and we can solve for it in a number of ways. Our preliminary results for curve normal parameter estimation using maximum likelihood suggest that the interesting (computationally hard) problem of orienting the normals will be an exciting new route to the signed distance function problem. We must emphasize here that the CWR lends itself to multiple avenues of the parameter estimation process: probabilistic, geometric, and data driven (see below).

The unnormalized function $|\psi(x)|^2$ is

$$|\psi(x)|^2 \propto \sum_{j=1,k\geq j}^{N,N} \cos\left(\frac{\nu_j(x-\mu_j)-\nu_k(x-\mu_k)}{\lambda}\right) \exp\left\{-\frac{\|x-\mu_j\|^2}{\tau}-\frac{\|x-\mu_k\|^2}{\tau}\right\}. \tag{10}$$

Note that this is not the L_2 norm but the squared magnitude of ψ at location x. It is not obvious from the expression above, but as $|\psi(x)|^2$ is the magnitude squared of a complex number, it is nonnegative everywhere. When suitably normalized, $|\psi(x)|^2$ can be treated as a probability density function which immediately connects it to the plethora of shape density functions used in the literature.

Here we consider the shape registration problem under a maximum likelihood formulation. $C = \{\mu_k, \nu_k\}_{k=1}^{N_1}$ is given as a template and the task is to find a mapping from $P = \{x_j\}_{j=1}^{N_2}$ to C within a class of admissible maps $\mathcal{H}$. The maximum likelihood optimization problem

$$\max_{f\in\mathcal{H}} \prod_{j=1}^{N_2} |\psi(f(x_j);C)|^2, \tag{11}$$

is robust to Gaussian noise on C (see Fig. 2). Here $\mathcal{H}$ consists of a rotation followed by a shear. Note that this is not the same as maximizing the likelihood of a Gaussian mixture on a test point-set since the cross terms of $|\psi|^2$ interact. Instead, it is uniquely suited to situations in which an *oriented* template is registered to an unoriented point-set. Note that once an alignment is achieved, the problem of signed distance function estimation can be expressed as an extension: extend the organization of the template set to the unorganized set.

B. ψ for kPCA on Curves

In kernel PCA (kPCA) [20], the goal is to build a linear basis of functions, $B = \{e_i\}_{i=1}^{N-1}$, out of the features observed during a training phase. The optimum basis minimizes the error of the original feature set $\{\psi(x;C_i)\}_{i=1}^{N}$. The centered Gram Matrix $K_{ij} = \langle\psi(x;C_i),\psi(x;C_j)\rangle$ corresponding to oriented point-sets i and j is then eigendecomposed following the equation

$$\lambda_i N e_i = K e_i \tag{12}$$

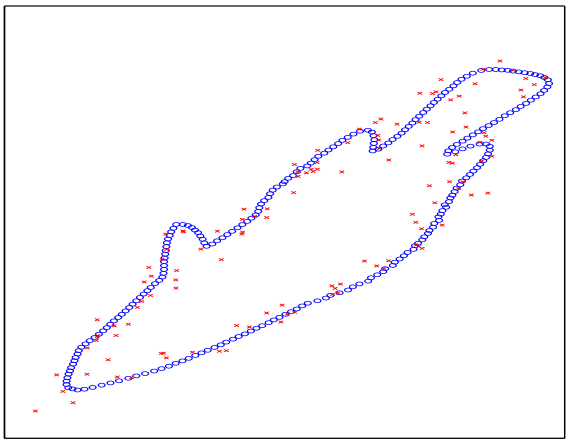

$\sigma_{\text{err}}/\sigma_{\text{data}}$	$\lvert\theta^*-\theta\rvert$	$\lvert s_1^*-s_1\rvert$	$\lvert s_2^*-s_2\rvert$
.045	0.006	0.076	0.046
.06	0.038	0.057	0.094
.075	0.008	0.088	0.027
.09	0.008	0.052	0.052
.105	0.006	0.029	0.042

Fig. 2. **Maximum likelihood alignment using $|\psi(x)|^2$ as a density.** The blue circles (○) are a noiseless template with accurate normal data while the red points (×) are points sampled from the template with Gaussian noise added. For a range of noise parameters, an alignment of the noisy data to the template was found by maximizing the likelihood in (11). The unknown transformation parameters were drawn uniformly with $\theta \in [0, 2\pi)$ and $s_1, s_2 \in [.5, 2]$. θ is the total rotation angle of the template, s_1, s_2 scale in the respective directions of the rotated basis. σ_{data} is the spatial standard deviation of the points in the template, and σ_{err} is the standard deviation of the added Gaussian noise to the test point-set. Root mean squared relative error is reported over 25 trials at each noise threshold.

resulting in $\{e_i\}_{i=1}^{N-1}$ containing the kernel principal components. Then the collection of e_i corresponding to nonzero λ_i are normalized and used as a basis for test patterns.

The first problem we propose to solve with kPCA is to estimate the underlying density function model for a shape, through the magnitude squared of its expansion on a linear space of basis functions provided as training data. As an added and surprising benefit, we show that the expansion itself contains closed curves. The perceptual gains of the CWR are conserved under kPCA encoding. The CWR is therefore uniquely suited to PCA-based compression unlike probability densities which are positive and integrate to one. In summary, we can start with a set of oriented point-sets, go to our feature basis and build a linear subspace, and accurately approximate a closed curve and a probability density corresponding to an unseen, test, oriented point-set in terms of a few basis coefficients. See Fig. 3, where the closed curve estimated by the kernel is shown on the left of the figure. This is a novel aspect of the CWR directly leveraging the properties of linearity and superposition to construct a basis representation. We show that the absolute error of this representation serves as a discriminative measure for classifying oriented point-sets. An additional novel aspect is a framework that also yields a generative approximation corresponding to the classification—the wave function that emerges from the approximation of the

data in the kPCA basis.

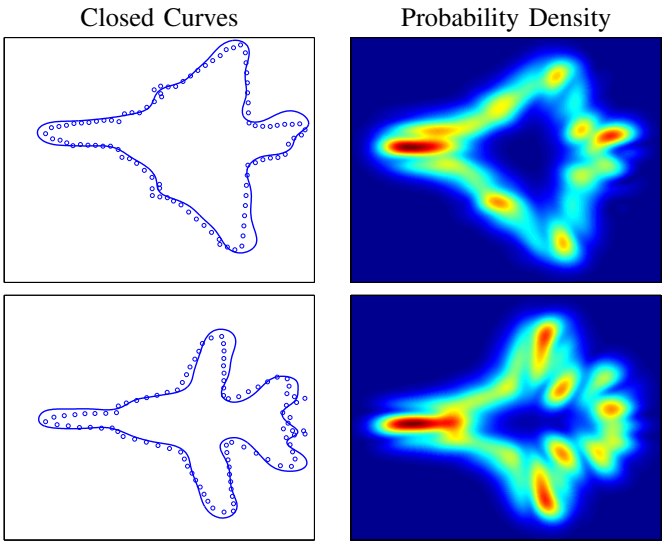

Class 1	err_1	err_2	Class 2	err_1	err_2
1	1.7	2.4	1	2.4	0.72
2	0.73	2.0	2	2.2	0.74
3	0.67	2.9	3	2.0	1.55
4	0.42	2.7	4	2.1	0.34
5	0.54	2.7	5	2.3	1.59

Fig. 3. **kPCA Results: Closed curves from linear combinations of closed curves.** Linear bases $B_1 = \{e_i^1\}_{i=1}^{25}$, $B_2 = \{e_i^2\}_{i=1}^{25}$, each consisting of linear combinations of 25 patterns were used to predict the features for 5 unseen patterns. The first 24 principal components were used as a basis for the point-sets shown. Note that through kPCA, we simultaneously estimated a density function and a closed curve from the span of 24 training patterns with coefficients derived from ≈ 100 elements of the oriented point-set. The data sets used were Mirage and F-14 (wings open) [21].

V. Conclusions and Future Work

The barrier to entry for any potential new member of the shape zoo is high. By integrating two popular shape representations—the signed distance transform and the probability density function—we arrived at a complex wave representation (CWR) for shape, with Huygens' principle acting as a regularizer for point-set organization. With the squared magnitude of the wave related to a probability density and the phase related to a modular signed distance function, we have taken the first steps toward integration.

Our work also suggests that research involving the CWR can proceed on multiple fronts: deformable template matching, dictionary construction, classification and indexing. Deeper understanding of the mathematical properties of the CWR on both a local and global level will elevate the state of the art for dealing with some of the issues that plague the signed distance function estimation community: since the field has geometric and probabilistic aspects, control or regularization of signed distance function estimation can bifurcate. In our opinion, the most exciting aspect of the CWR lies in its potential for advances in perceptual grouping and the computation of signed distance functions from unorganized point-sets.

Acknowledgments

The authors thank Sibel Tari, Sudeep Sarkar, Manu Sethi, and Karthik S. Gurumoorthy for invaluable discussions and encouragement. This work is supported by NSF IIS 1065081.

References

[1] J. A. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proceedings of the National Academy of Sciences*, vol. 93, no. 4, pp. 1591–1595, 1996.

[2] I. L. Dryden and K. Mardia, *Statistical Shape Analysis*, ser. Probability and Statistics. Wiley, 1998.

[3] J. Butterfield, "On Hamilton-Jacobi theory as a classical root of quantum theory," in *Quo Vadis Quantum Mechanics?*, ser. The Frontiers Collection. Springer, 2005, ch. 13, pp. 239–273.

[4] D. Cremers, S. J. Osher, and S. Soatto, "Kernel density estimation and intrinsic alignment for shape priors in level set segmentation," *International Journal of Computer Vision*, vol. 69, no. 3, pp. 335–351, 2006.

[5] A. M. Peter and A. Rangarajan, "Maximum likelihood wavelet density estimation with applications to image and shape matching," *IEEE Transactions on Image Processing*, vol. 17, no. 4, pp. 458–468, 2008.

[6] F. Bornemann and T. März, "Fast image inpainting based on coherence transport," *Journal of Mathematical Imaging and Vision*, vol. 28, no. 3, pp. 259–278, 2007.

[7] A. Sharf, T. Lewiner, A. Shamir, L. Kobbelt, and D. Cohen-Or, "Competing fronts for coarse–to–fine surface reconstruction," in *Computer Graphics Forum*, vol. 25, no. 3. Wiley, 2006, pp. 389–398.

[8] S. Osher and R. Fedkiw, *Level set methods and dynamic implicit surfaces*. Springer, 2003, vol. 153.

[9] M. Sethi, A. Rangarajan, and K. S. Gurumoorthy, "The Schrödinger distance transform (SDT) for point-sets and curves," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 198–205.

[10] M. Breuß, A. Bruckstein, and P. Maragos, *Innovations for Shape Analysis: Models and Algorithms*, ser. Mathematics and Visualization. Springer, 2013.

[11] S. W. Zucker, "Distance images and the enclosure field: Applications in intermediate-level computer and biological vision," in *Innovations for Shape Analysis: Models and Algorithms*, ser. Mathematics and Visualization, M. Breuß, A. Bruckstein, and P. Maragos, Eds. Springer, 2013, ch. 14, pp. 301–323.

[12] S. Tari, "Fluctuating distance fields, parts, three-partite skeletons," in *Innovations for Shape Analysis: Models and Algorithms*, ser. Mathematics and Visualization, M. Breuß, A. Bruckstein, and P. Maragos, Eds. Springer, 2013, ch. 20, pp. 439–466.

[13] S. Tari, J. Shah, and H. Pien, "Extraction of shape skeletons from grayscale images," *Computer Vision and Image Understanding*, vol. 66, no. 2, pp. 133–146, 1997.

[14] G. Aubert and J.-F. Aujol, "Poisson skeleton revisited: A new mathematical perspective." *Journal of Mathematical Imaging and Vision*, vol. 48, pp. 149–159, 2014.

[15] P. Soundararajan and S. Sarkar, "An in-depth study of graph partitioning measures for perceptual organization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 6, pp. 642–660, 2003.

[16] P. Mordohai and G. Medioni, *Tensor Voting: A Perceptual Organization Approach to Computer Vision and Machine Learning*. Morgan and Claypool, 2006.

[17] S. Sarkar and K. L. Boyer, "Perceptual organization in computer vision: A review and a proposal for a classificatory structure," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 23, no. 2, pp. 382–399, 1993.

[18] A. Rangarajan, "Revisioning the unification of syntax, semantics and statistics in shape analysis," *Pattern Recognition Letters*, vol. 43, pp. 39–46, 2014.

[19] T. S. Lee, "Image representation using 2D Gabor wavelets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 10, pp. 959–971, 1996.

[20] B. Schölkopf, A. Smola, and K.-R. Müller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Computation*, vol. 10, no. 5, pp. 1299–1319, 1998.

[21] N. Thakoor, J. Gao, and S. Jung, "Hidden Markov model-based weighted likelihood discriminant for 2D shape classification," *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2707–2719, 2007.