A Schrödinger formalism for simultaneously computing the Euclidean distance transform and its gradient density

Karthik S. Gurumoorthy*
International Center for Theoretical Sciences,
Tata Institute of Fundamental Research,
Bangalore, Karnataka, India
karthik.gurumoorthy@icts.res.in

Anand Rangarajan
Department of Computer and Information
Science and Engineering,
University of Florida,
Gainesville, FL, USA
anand@cise.ufl.edu

ABSTRACT

In this work, we employ the well-known Hamilton-Jacobi to Schrödinger connection to present a unified framework for computing both the Euclidean distance function and its gradient density in two dimensions. Previous work in this direction considered two different formalisms for independently computing these quantities. While the two formalisms are very closely related, their lack of integration is theoretically troubling and practically cumbersome. We introduce a novel Schrödinger wave function for representing the Euclidean distance transform from a discrete set of points. An approximate distance transform is computed from the magnitude of the wave function while the gradient density is estimated from the Fourier transform of the phase of the wave function. In addition to its simplicity and efficient $O(N \log N)$ computation, we prove that the wave function-based density estimator increasingly, closely approximates the distance transform gradient density (as a free parameter approaches zero) with the added benefit of not requiring the true distance function.

Keywords

Euclidean distance functions, Gradient density estimation, Schrödinger equation, Fourier transform

1. INTRODUCTION

One of the most popular representations used in shape analysis is the distance transform (or the Euclidean distance function). The distance transform—closely related to Voronoi diagrams in 2D and in 3D [7]—is a *scalar field* representation of a shape and is widely deployed in computer vision, machine learning, graphics and medical imaging circles [14, 13, 18]. When shapes are parametrized as a point-set, the distance transform at any point on a grid (in 2D or in 3D) is the Euclidean distance to the nearest point in the point-set. For a 2D point-set, the distance transform can be readily visualized as a set of interlocking cones with the tip of each cone at a unique point in the point-set. The interlocking cone geometry dovetails well with the Voronoi diagram representation since one can readily visualize a circular wavefront emanating at constant speed from each Voronoi center (cone tip) and meeting at the Voronoi boundaries—the set of symmetries with each point in this set being equidistant to at least two Voronoi centers. The fast marching and fast sweeping methods leverage this intuition into fast distance transform algorithms [17, 21, 15].

*Corresponding author

ICVGIP '14, December 14-18, 2014, Bangalore, India

http://dx.doi.org/10.1145/2683483.2683486.

2. PREVIOUS WORK ON THE SCHRÖDINGER FORMALISM

The fast marching and the fast sweeping methods—entrenched for the past twenty years as the methods of choice for distance transform computation—are examples of Hamilton-Jacobi solvers. In Hamilton-Jacobi theory, a scalar field $S(X)$ is sought that satisfies a Hamilton-Jacobi equation. Recall from above, the underlying intuition of circular wavefronts (at constant speed) emanating from each Voronoi center. This intuition can be couched (in 2D) in the form of the Hamilton-Jacobi equation $S_x^2 + S_y^2 = 1$ with $S(X) = 0$ at the Voronoi centers. The result is a nonlinear partial differential equation whose solution is the distance transform of a point-set.

Hamilton-Jacobi theory is the pinnacle of classical physics. For almost a century however, we have known of the close correspondence between Hamilton-Jacobi theory and quantum theory—or the correspondence between Hamilton-Jacobi and Schrödinger equations for the same underlying physical events. It is widely known in the theoretical physics literature that a Hamilton-Jacobi equation can be obtained from the corresponding Schrödinger equation in the limit as Planck's constant $\hbar$ is sent to zero [3]. The peculiarity of this relationship is heightened when we take into account the nonlinearity of Hamilton-Jacobi equations as against the linearity of Schrödinger equations.

What then is the Schrödinger equation for the distance transform problem? If the Hamilton-Jacobi equation is $\|\nabla S\|^2 - 1 = 0$ (almost everywhere except at the Voronoi centers), then an appropriate Schrödinger equation is

$$(-i\hbar\nabla)^2 \psi + \psi = \psi_0 \tag{2.1}$$

where ψ_0 takes suitable values over the domain (while being related to the source points for which $S(X) = 0$). If we restrict the solutions to be real (while glossing over normalization issues), a simple correspondence can be discovered between the Hamilton-Jacobi scalar field $S(X)$ and the quantum field $\psi(X)$, specifically $\psi(X) = \exp\left(-\frac{S(X)}{\hbar}\right)$.

At this point, the reader may rightly be questioning the relevance of this material to computer vision. Here's the payoff in the sim-

plest possible terms. Given a set of points $Y = \{Y_k\}_{k=1}^K$, the approximate solution $R(X)$[1] obtained from the Schrödinger equation in (2.1) at a small, non-zero value of $\hbar$ is

$$\phi^{(1)}(X) \approx \sum_{k=1}^{K} \exp\left(-\frac{\|X - Y_k\|}{\hbar}\right) \quad \text{and}$$
$$R(X) = -\hbar \log\left(\phi^{(1)}(X)\right). \tag{2.2}$$

In other words, the above distance transform is strikingly similar to a soft *winner-take-all* competition [20] at each grid point X between the set of Voronoi centers $\{Y_k\}_{k=1}^K$. The degree of competition depends upon both the free parameter $\hbar$ and the distance to the Voronoi centers. We henceforth use τ instead of $\hbar$ to avoid the semantic dissonance created by treating Planck's constant $\hbar$ as a free parameter. The Voronoi boundaries are the set of symmetries at which there is no unique winner among the Voronoi centers. We see that an approximate $R(X)$ is obtained by first solving a linear Schrödinger equation and then passing the result through a non-linearity ($-\log$). The algorithmic payoff: the linear Schrödinger equation can be efficiently solved using a fast Fourier transform (FFT). Previous work describing this approach can be found in [10].

With the histogram of oriented gradients (HOG) emerging as a popular feature descriptor in image analysis [6, 12, 22], it is natural to compute the HOG of distance transforms as a feature descriptor for shape analysis. Since $\|\nabla S\| = 1$ almost everywhere, the distance transform gradients are unit vectors (wherever they exist) and hence, in 2D, the density function of distance transform gradients is one-dimensional, i.e, defined w.r.t. the orientation $\theta = \arctan\left(\frac{S_y}{S_x}\right)$. It turns out that when we define a wave function

$$\phi^{(2)}(X) = \exp\left(\frac{iS(X)}{\tau}\right), \tag{2.3}$$

the squared magnitude of the Fourier transform of $\phi^{(2)}(X)$ is approximately proportional to the density function of distance transform gradient vectors with the approximation becoming increasingly exact as $\tau \to 0$. Furthermore, the one dimensional nature of the density function is verified by the fact that the Fourier transform of the wave function takes values mainly on the unit circle (after axis normalization) in the spatial frequency domain. In other words, *spatial frequencies become histogram bins* in this representation and only the "bins" on the unit circle get enough votes. A detailed explanation of this approach can be found in [11]. While this relationship is to be expected from a physics perspective as both ∇S and spatial frequencies are related to momenta, from machine learning and density estimation perspectives, this relationship appears to be *new*.

3. CONTRIBUTIONS

As presented above, an approximate distance transform can be obtained by solving a linear Schrödinger equation and exploiting the relationship $\phi^{(1)}(X) = \exp\left(-\frac{R(X)}{\tau}\right)$. Further, an approximate density function of distance transform gradients can be obtained from the squared magnitude of the Fourier transform of the wave function $\phi^{(2)}(X) = \exp\left(\frac{iS(X)}{\tau}\right)$. This asymmetry between the exponent (in the former) and the phase (in the latter) is theoretically troubling and experimentally cumbersome. It would be of great benefit to be able to *theoretically guarantee and efficiently compute* both an approximate distance transform and a gradient density from a single Schrödinger representation.

[1]Note that $R(X)$ is the approximate solution whereas $S(X)$ is the true Euclidean distance function. The solution is approximate because $\hbar$ is set to a small *non-zero* value.

3.1 A new wave function for joint estimation

In the current work we are primarily interested in extending our previous formulations described in [10] and [11]. The main thrust of the present work is to provide an answer to the following question: *Can we synthesize a single wave function capable of representing both the distance transform and its gradient density?* In [16] we hinted at one such representation but did not provide any theoretical underpinnings to substantiate our claim. In contrast, we not only introduce a wave function $\phi(X)$ capable of integrating both the distance transform and its gradient density in a single representation but also provide theoretical and experimental support for the new representation.

Consider a wave function

$$\phi(X) = \sum_{k=1}^{K} \exp\left(\frac{(-1+i)\|X - Y_k\|}{\tau}\right). \tag{3.1}$$

Its magnitude $\phi_m(X)$ equals

$$(\phi_m(X))^2 = \left[\sum_{k=1}^{K} \exp\left(\frac{-\|X - Y_k\|}{\tau}\right) \cos\left(\frac{\|X - Y_k\|}{\tau}\right)\right]^2 + \left[\sum_{k=1}^{K} \exp\left(\frac{-\|X - Y_k\|}{\tau}\right) \sin\left(\frac{\|X - Y_k\|}{\tau}\right)\right]^2. \tag{3.2}$$

Let the wave function $\phi_p(X)$ represent the phase component of $\phi(X)$ given by

$$\phi_p(X) \equiv \frac{\phi(X)}{\phi_m(X)} = \sum_{k=1}^{K} \frac{\exp\left(\frac{-\|X-Y_k\|}{\tau}\right) \exp\left(\frac{i\|X-Y_k\|}{\tau}\right)}{\phi_m(X)}. \tag{3.3}$$

Though $\phi_p(X)$ is somewhat similar to $\phi^{(2)}(X)$ in that it also has unit magnitude, the phase component of $\phi_p(X)$ is not the true Euclidean distance function. The definition of $\phi_p(X)$ involves a summation over all the source points $\{Y_k\}_{k=1}^K$ and not just the nearest neighbor of X, as in the case of $\phi^{(2)}(X)$ which is defined using the true distance function $S(X)$. Similarly, although the functions $\phi_m(X)$ and $\phi^{(1)}(X)$ are closely related (especially when τ tends to zero), they are not identical as $\phi_m(X)$ also includes sinusoidal functions.

We establish the following results in this paper.

THEOREM 3.1. *The power spectrum of the phase component of $\phi(X)$, denoted by $\phi_p(X)$, is approximately equal to the density function of the distance transform gradients, with the approximation becoming increasingly exact as $\tau \to 0$.*

THEOREM 3.2. *The approximate distance function $R_m(X)$ recovered from the exponent of the magnitude of $\phi(X)$, specifically $R_m(X) = -\tau \log(\phi_m(X))$ converges to the true distance function $S(X)$ as $\tau \to 0$.*

3.2 Novel Aspects of the Present Work

- The unified framework presented here allows us to compute both the distance function and its gradient density from the same wave function.

- The unified representation bypasses the need for acquiring the true distance function in order to determine its gradient density.
- Gradients of the distance function are generally obtained by initially solving for the distance function and later computing its discrete derivatives. Our work showcases an approach wherein neither the distance function nor its derivatives are explicitly computed.

We would like to emphasize that the main focus of the paper is the aforementioned new mathematical result and hence is of theoretical interest. Efficient computation of the distance function and the gradient density follow from the theoretical results.

4. THE DISTANCE TRANSFORM GRADIENT DENSITY FUNCTION

In [11], we furnished a closed-form expression for the density function of the distance transform gradients which we briefly state here for the sake of completion. As mentioned above, the geometry of the distance transform corresponds to a set of intersecting cones with the origins at the Voronoi centers. The gradients of the distance transform (which exist globally except at the cone intersections and origins) are unit vectors. Therefore the gradient density function is one-dimensional and defined over the space of *orientations*. The orientations are constant and unique along each ray of each cone. The probability distribution function of the true distance function gradients is given by

$$\mathcal{F}(\theta \leq \Theta \leq \theta + \Delta) \equiv \frac{1}{L} \int \int_{\theta \leq \arctan\left(\frac{S_y}{S_x}\right) \leq \theta + \Delta} dx\, dy \quad (4.1)$$

where we have expressed the orientation random variable as $\Theta = \arctan\left(\frac{S_y}{S_x}\right)$. Here, L is the area of the 2D domain Ω whose expression is given below. The probability distribution function induces a *closed-form expression* for its density function $P(\theta)$ and is given by

$$P(\theta) \equiv \lim_{\Delta \to 0} \frac{\mathcal{F}(\theta \leq \Theta \leq \theta + \Delta)}{\Delta} = \frac{1}{L} \sum_{k=1}^{K} \frac{Q_k^2(\theta)}{2}. \quad (4.2)$$

Here $Q_k(\theta)$ is the length of the ray of the k^{th} cone at an orientation θ and the area L can be written as

$$L = \sum_{k=1}^{K} \int_0^{2\pi} \int_0^{Q_k(\theta)} r\, dr\, d\theta = \sum_{k=1}^{K} \int_0^{2\pi} \frac{Q_k^2(\theta)}{2}\, d\theta. \quad (4.3)$$

Though a closed-form expression is attractive, computing the density function $P(\theta)$ via (4.2) is an onerous task as we need to first determine the Voronoi region corresponding to each Voronoi center (source point) Y_k and then for each orientation direction θ, compute the ray length $Q_k(\theta)$ from Y_k to its Voronoi boundary along θ. These involve unwieldy manipulations of complicated data structures. On the other hand, our mathematical result provides an easy mechanism to achieve the same task as explained in Section 7.

5. EQUIVALENCE OF THE POWER SPECTRUM AND THE GRADIENT DENSITY

In this section, we claim that the power spectrum [2] of $\phi_p(X)$—defined in (3.3)— can increasingly, closely approximate the distance transform gradient density as τ tends to zero. We will demonstrate that as $\tau \to 0$, the power spectrum converges to the closed-form expression stated in (4.2).

5.1 Scaled Fourier Transform of the Wave Function Phase

The scaled Fourier transform of $\phi_p(X)$ for a given value of τ equals

$$F_\tau(u, v) \equiv \frac{1}{2\pi\tau\sqrt{L}} \iint_\Omega \phi_p(X) \exp\left(\frac{-i(ux + vy)}{\tau}\right) dx\, dy. \quad (5.1)$$

Here the grid point $X = (x, y)$. The scale factor $\frac{1}{2\pi\tau\sqrt{L}}$ is the normalizing term such that the L_2 norm of F_τ is 1. Let $\mathcal{D}_j$ denote the j^{th} Voronoi region corresponding to the input point Y_j, i.e. if $X = (x, y) \in \mathcal{D}_j$, then $\|X - Y_j\| \leq \|X - Y_k\|, \forall k$ and the inequality is strict for points not lying on the Voronoi boundary. $\mathcal{D}_j$ can be represented by a Cartesian product $[0, 2\pi) \times [0, Q_j(\theta)]$. It follows that $\Omega = \bigcup_{j=1}^{K} \mathcal{D}_j$.

The Voronoi boundary of $\mathcal{D}_j$ is composed of a sequence of straight line segments. As it is one-dimensional—a *measure zero* set in 2D—it doesn't contribute to the integration in (5.1). Hence, we let $\mathcal{D}_j$ correspond *only* to the interior points of the Voronoi region. To be mathematically precise, we consider the Voronoi region $\mathcal{D}_j^\epsilon$ for any arbitrarily small $\epsilon > 0$ such that if $X \in \mathcal{D}_j^\epsilon$, then $\|X - Y_j\| < \|X - Y_k\| - \epsilon, \forall k$ (please note the strict inequality). Since the distance function is not differentiable at the source points, we exclude them too. In other words, the region $\mathcal{D}_j^\epsilon$ is *strictly away* from the Voronoi boundaries of $\mathcal{D}_j$ and the source points and is represented by the Cartesian product $[0, 2\pi) \times [\epsilon Q_j(\theta), (1-\epsilon) Q_j(\theta)]$. We then let $\epsilon \to 0$ to include all the grid locations except the Voronoi boundaries and the source points. The reason for introducing the ϵ factor will become clearer as we proceed.

The Fourier transform for a given $\epsilon > 0$ can be similarly expressed as

$$2\pi\tau\sqrt{L^\epsilon} F_\tau^\epsilon(u, v) \equiv \sum_{j=1}^{K} \iint_{\mathcal{D}_j^\epsilon} \phi_p(X) \exp\left(\frac{-i(ux + vy)}{\tau}\right) dx\, dy. \quad (5.2)$$

From the definition of L in (4.3), the area L^ϵ of the domain $\Omega^\epsilon = \bigcup_{j=1}^{K} \mathcal{D}_j^\epsilon$ equals

$$L^\epsilon = \sum_{j=1}^{K} \int_0^{2\pi} \int_{\epsilon Q_j(\theta)}^{(1-\epsilon) Q_j(\theta)} r\, dr\, d\theta = (1 - 2\epsilon) \sum_{j=1}^{K} \int_0^{2\pi} \frac{Q_j^2(\theta)}{2}\, d\theta = (1 - 2\epsilon) L. \quad (5.3)$$

Note that $\lim_{\epsilon \to 0} L^\epsilon = L$ and $\lim_{\epsilon \to 0} F_\tau^\epsilon(u, v) = F_\tau(u, v)$. For $j \in \{1, 2, \ldots, K\}$, let $\phi_j(X) = \phi_p(X)$ when $X \in \mathcal{D}_j^\epsilon$. Recall that in $\mathcal{D}_j^\epsilon$, $S(X) = \|X - Y_j\|$. Multiplying and dividing $\phi_j(X)$ by $\exp\left(\frac{S(X)}{\tau}\right)$ we get

$$\phi_j(X) = \frac{\exp\left(\frac{iS(X)}{\tau}\right)}{\alpha_j(X)} + \eta_j(X) \quad (5.4)$$

where

$$\eta_j(X) = \sum_{k=1, k\neq j}^{K} \frac{\exp\left(\frac{i\|X-Y_k\|}{\tau}\right)\exp\left(-\frac{\beta_{jk}(X)}{\tau}\right)}{\alpha_j(X)}, \quad (5.5)$$

$$\alpha_j(X) = \exp\left(\frac{S(X)}{\tau}\right)\phi_m(X) \text{ and} \quad (5.6)$$

$$\beta_{jk}(X) = \|X-Y_k\| - \|X-Y_j\|. \quad (5.7)$$

We have the following Lemma whose proof is given in Appendix A.1.

LEMMA 5.1. *If $\alpha_j(X)$ is defined as in (5.6), then*

$$\lim_{\tau\to 0}\alpha_j(X) = 1, \quad (5.8)$$

and furthermore the convergence is uniform $\forall X \in \mathcal{D}_j^\epsilon$.

The stronger uniform convergence result is required by the following Lemma 5.2. By choosing an arbitrarily small but non-zero ϵ, we deliberately *exclude* the points on the Voronoi boundary by only considering the region $\mathcal{D}_j^\epsilon$ in lieu of the true Voronoi region $\mathcal{D}_j$ and obtain uniform convergence.

LEMMA 5.2. *If $\eta_j(X)$ is defined as in (5.5), then*

$$\lim_{\tau\to 0}\frac{1}{2\pi\tau\sqrt{L^\epsilon}}\sum_{j=1}^{K}\iint_{\mathcal{D}_j^\epsilon}\eta_j(X)\exp\left(\frac{-i(ux+vy)}{\tau}\right)dx\,dy = 0. \quad (5.9)$$

The proof is available in Appendix A.2.

5.2 Proof of Theorem 3.1

PROOF. Using Lemmas 5.1 and 5.2, we can conclude that $\phi_j(X) \approx \exp\left(\frac{iS(X)}{\tau}\right)$ at small values of τ. Hence the Fourier transform for a fixed $\epsilon > 0$ stated in (5.2) can be approximated by

$$2\pi\tau\sqrt{L^\epsilon}F_\tau^\epsilon(u,v) \approx \sum_{j=1}^{K}\iint_{\mathcal{D}_j^\epsilon}\exp\left(\frac{iS(X)}{\tau}\right)\exp\left(\frac{-i(ux+vy)}{\tau}\right)dx\,dy, \quad (5.10)$$

as $\tau \to 0$. This lead us to an expression which is similar to the one considered in [11].

The power spectrum defined as the squared magnitude of the Fourier transform [2] is given by $P_\tau^\epsilon(u,v) \equiv |F_\tau^\epsilon(u,v)|^2$. Using the higher-order stationary phase approximation [19], we can establish the convergence of $P_\tau^\epsilon(u,v)$ to the true gradient density of the distance function, $P(\theta)$, given in (4.2) as $\tau \to 0$. As the gradients are unit vectors, the Fourier transform values of $\phi_p(X)$ lie mainly on the unit circle $\tilde{r} = \sqrt{u^2+v^2} = 1$ in the frequency domain and this behavior tightens as $\tau \to 0$. When we poll the power spectrum at locations close to the unit circle, it progressively resembles the density function of the gradients of the distance transform as τ tends to zero and hence can serve as a gradient density estimator at small, non-zero values of τ. For any orientation $\omega_0 \in [0, 2\pi)$, this result can be mathematically expressed as follows:

$$\lim_{\epsilon\to 0}\lim_{\tau\to 0}\int_{\omega_0}^{\omega_0+\Delta}\left\{\int_{1-\delta}^{1+\delta}P_\tau^\epsilon(\tilde{r},\omega)\,d\tilde{r}\right\}d\omega = \int_{\omega_0}^{\omega_0+\Delta}P(\omega)\,d\omega. \quad (5.11)$$

Here $P_\tau^\epsilon(\tilde{r},\omega)$ is the representation of the power spectrum in the polar coordinate system $0 < \delta < 1$ for polling close to the unit circle (for which $\tilde{r} = 1$). The parameter Δ is an *arbitrarily* small interval on the orientation required to cancel out the cross phase factors which arise due to multiple locations indexing the same frequency bin. Detailed exposition of the proof can be found in [11]. We would like to accentuate that our work in [11] focused only on estimating the orientation density and *did not* leverage the unified representation for the distance transform and gradient density as considered in this paper. □

6. RECOVERING THE TRUE EUCLIDEAN DISTANCE FUNCTION

We now show that the function $R_m(X) = -\tau\log(\phi_m(X))$ converges to the true distance function $S(X)$ as $\tau \to 0$.

6.1 Proof of Theorem 3.2

PROOF. Let Y_{k_0} be the closest source-point to X, i.e, $S(X) = \|X-Y_{k_0}\| \leq \|X-Y_k\|, \forall k$. Multiplying and dividing $\phi_m(X)$ in (3.2) by $\exp\left(\frac{S(X)}{\tau}\right)$ we get

$$\begin{aligned} R_m(X) &= S(X) - \tau\log\left(\alpha_{k_0}(X)\right) \\ &= S(X) - \tau\log\left(\sqrt{1+\gamma_{k_0}(X)}\right) \end{aligned} \quad (6.1)$$

where $\gamma_{k_0}(X)$ is defined in (A.3). If X doesn't lie on the Voronoi boundary, then from Lemma 5.1 we have $\alpha_{k_0}(X) \to 1$ as $\tau \to 0$. On the flip side, if X is a point on the Voronoi boundary then the worst scenario would be when X is equidistant from all the point-sources in which case $\beta_{jk}(X)$ defined in (5.7) will equal zero for any pair of j and k. Hence, we get $1+\gamma_{k_0}(X) \leq K^2$. In either case we arrive at the desired result, namely $\lim_{\tau\to 0} R_m(X) = S(X)$. □

Observe that the above inequality also gives an error bound of $\tau\log(K)$ for a fixed value of τ which is tight as:

- It only scales as the logarithm of K.
- It can be made arbitrarily small by appropriately choosing τ.

7. COMPUTATIONAL COMPLEXITY

7.1 Algorithms based on Discrete Convolution

A salient feature of ϕ is its simplicity and efficient computability. The wave function $\phi(X)$ can be expressed as the convolution between the functions $f(X) = \exp\left(\frac{(-1+i)\|X\|}{\tau}\right)$ and $g(X) = \sum_{k=1}^{K}\delta(X-Y_k)$ where $g(X)$ is the sum of spatially shifted delta functions such that it is non-zero only at the point-set locations $\{Y_k\}_{k=1}^K$ and 0 elsewhere. Computation of ϕ at the given N regular grid locations can be performed in $O(N\log N)$ time using the fast Fourier transform (FFT) based discrete convolution algorithm [4]. Once we have the knowledge of $\phi(X)$, both the approximate distance function and its gradient density can also be easily recovered. The distance function $R_m(X)$ can be obtained in $O(N)$ as it only takes the equivalent amount of time to compute the magnitude function $\phi_m(X)$ and its logarithm from $\phi(X)$. Similarly, the phase function $\phi_p(X)$ can be determined at the N grid locations in $O(N)$ time. The discrete Fourier transform of $\phi_p(X)$ can be performed in $O(N\log N)$ and the subsequent squaring operation to calculate the power spectrum requires $O(N)$. Hence the overall time complexity (inclusive of the time required to obtain $\phi(X)$) to compute

both the distance function and orientation density function is only $O(N \log N)$, *independent* of K and holds specifically for the case where $K = O(N)$.

7.2 Numerical stability

In principle we should be able to apply our technique at a very small value of τ and see good convergence results. But we noticed that a naïve double precision-based implementation may be erroneous for values of τ very close to zero. At small values of τ, the magnitude of $f(X)$ drops off very quickly and hence for grid locations which are far away from the point-set, the convolution performed using FFT may not be accurate due to floating point errors. We avoid numerical issues by using the GNU MPFR multiple-precision arithmetic library which provides arbitrary precision arithmetic with correct rounding [8, 9]. Advanpix [1] is another MATLAB® software with arbitrary precision capabilities.

8. EMPIRICAL VALIDATION

We now provide empirical evidence to buttress our theoretical results by implementing our unified framework on a set of 2D shapes.[2] The silhouettes of the shapes were uniformly sampled to obtain the point-set locations $\{Y_k\}_{k=1}^{K}$. We considered the 2D grid $-0.123 \leq x \leq 0.123$ and $-0.123 \leq y \leq 0.123$ with a grid width of $\frac{1}{2^{10}}$. The number of grid locations equals $N = 253 \times 253 = 64009$. For a fixed value of τ, we computed the wave function $\phi(X)$ using the FFT based discrete convolution algorithm and then recovered the distance and the density functions using its magnitude and the phase respectively. For the sake of comparison, we also determined the distance function using the fast sweeping method [21]. The percentage errors incurred by our Schrödinger distance function and the fast sweeping technique in comparison to the true Euclidean distance function were calculated as:

$$\mathcal{E} = \frac{100}{N} \sum_{i=1}^{N} \frac{\Delta_i}{R_i} \tag{8.1}$$

where R_i and Δ_i are respectively the actual distance and the absolute difference of the computed distance to the actual distance at the i^{th} grid point. We list these error values for $\tau = 0.001$ and $\tau = 0.0004$ in Table 1. As one would expect, the error decreases as τ is reduced. The true Euclidean distance function contour plots and those obtained from our Schrödinger distance function at $\tau = 0.0004$ and fast sweeping are shown in Figure 4.

Table 1: Percentage errors incurred in distance function computation

Shape	% error in distance function		
	Schrödinger		Fast sweeping
	τ=0.001	τ=0.0004	
Hand	10.540 %	2.277%	2.359%
Horse	12.464%	2.707%	2.379%
Bird	12.193 %	2.404%	2.082%

In Figure 1, we plot the Fourier transform of the phase of the wave function $\phi(X)$ at $\tau = 0.0004$. Observe that the Fourier transform values mainly lie on the unit circle as the gradients of the distance function are unit vectors. In Figure 2 we compare the density function obtained from the power spectrum with the true density function computed using (4.2). Barring some differences at a few histogram bins resulting mainly from the binning method used for constructing these plots, we note from visual inspection that the two density functions exhibit similar profiles. To examine the convergence of the density function as $\tau \to 0$, we also computed the L_1 norm of the difference between the true and estimated densities for various τ values and summarized them in Table 2. Note that as we decrease τ, the L_1 error also drops. To press this point home, we also plot the density function for various values of τ in Figure 3 which clearly shows convergence to the true orientation density as $\tau \to 0$.

Table 2: L_1 errors for different τ values

τ value	L_1 error in gradient density		
	Hand	Horse	Bird
0.01	0.711	0.675	0.748
0.006	0.424	0.414	0.424
0.001	0.257	0.298	0.253
0.0004	0.247	0.183	0.167

9. DISCUSSION

We have exploited the relationships between Hamilton-Jacobi theory and Schrödinger equations to derive a new wave function capable of approximately representing both the Euclidean distance transform and its gradient density. For the first time, we have proven that an approximate distance transform can be computed from the magnitude of the wave function while an approximate gradient density can be estimated from the Fourier transform of the phase component of the wave function. Furthermore, these computations can be efficiently carried out using fast Fourier transforms in $O(N \log N)$ on a regular grid. The distance transform approximation and gradient density estimator increasingly, closely approximate the true distance and density functions respectively as a free parameter τ approaches zero. A novel aspect of the density estimator is that distance transform gradients need not be *a priori* computed.

Many extensions of the present work are possible. Signed distance function computation is one such extension wherein the domain is parsed into sets of closed curves. In general, signed distance function computation from point-sets is an ill-posed problem and is closely related to perceptual organization. It would be interesting to consider complex wave functions wherein the wave function phase encodes signed distance information. An initial foray with applications to shape analysis is presented in [5]. Finally, we are also interested in studying the performance of the gradient density estimator at small values of τ—in a manner akin to the study of the width parameter in Parzen window density estimators—in order to determine the relationship between the underlying discrete grid and the parameter τ.

Acknowledgments

This research work benefited from the support of the AIRBUS Group Corporate Foundation Chair in Mathematics of Complex Systems established in ICTS-TIFR to Karthik S. Gurumoorthy and NSF IIS 1065081 to Anand Rangarajan.

[2]We thank Kaleem Siddiqi for providing the 2D shapes used in this paper.

10. REFERENCES

[1] Multiprecision computing toolbox for MATLAB (Advanpix LLC.), 2014. url:http://www.advanpix.com.

[2] R. N. Bracewell. *The Fourier transform and its applications*. McGraw-Hill, New York, NY, 3rd edition, 1999.

[3] J. Butterfield. On Hamilton-Jacobi theory as a classical root of quantum theory. In A. Elitzur, S. Dolev, and N. Kolenda, editors, *Quo-Vadis Quantum Mechanics*, chapter 13, pages 239–274. Springer, New York, NY, 2005.

[4] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.

[5] J. Corring and A. Rangarajan. Shape from phase: An integrated level set and probability density shape representation. In *International Conference on Pattern Recognition (ICPR)*, pages 46–51, 2014.

[6] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 886–893. IEEE Computer Society, 2005.

[7] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, New York, NY, 3rd edition, 2008.

[8] L. Fousse, G. Hanrot, V. Lefévre, P. Pélissier, and P. Zimmermann. MPFR: A multiple-precision binary floating-point library with correct rounding. *ACM Transactions on Mathematical Software*, 33:1–15, 2007.

[9] T. Granlund *et al.* GNU MP: The GNU Multiple Precision Arithmetic Library, 2014. url:http://gmplib.org.

[10] K. S. Gurumoorthy and A. Rangarajan. A Schrödinger equation for the fast computation of approximate Euclidean distance functions. In *Scale Space and Variational Methods in Computer Vision (SSVM)*, volume 5567 of *Lecture Notes in Computer Science (LNCS)*, pages 100–111. Springer, 2009.

[11] K. S. Gurumoorthy and A. Rangarajan. Distance transform gradient density estimation using the stationary phase approximation. *SIAM Journal on Mathematical Analysis*, 44(6):4250–4273, 2012.

[12] R. Hu and J. Collomosse. A performance evaluation of gradient field HOG descriptor for sketch based image retrieval. *Computer Vision and Image Understanding (CVIU)*, 117(7):790–806, July 2013.

[13] R. Kimmel. *Numerical geometry of images: Theory, algorithms, and applications*. Springer-Verlag, New York, NY, 2004.

[14] S. J. Osher and R. P. Fedkiw. *Level set methods and dynamic implicit surfaces*. Springer-Verlag, New York, NY, October 2003.

[15] S. J. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.

[16] 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 (CVPR)*, pages 198–205. IEEE Computer Society, 2012.

[17] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

[18] K. Siddiqi and S. Pizer, editors. *Medial Representations: Mathematics, Algorithms and Applications*, volume 37 of *Computational Imaging and Vision*. Springer, 2008.

[19] R. Wong. *Asymptotic approximations of integrals*. Academic Press, New York, NY, 1989.

[20] A. L. Yuille and D. Geiger. Winner-Take-All networks. In M. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, pages 1228–1231. The MIT Press, Boston, MA, 2nd edition, 2002.

[21] H. K. Zhao. A fast sweeping method for eikonal equations. *Mathematics of Computation*, 74(250):603–627, 2005.

[22] Q. Zhu, M.-C. Yeh, K.-T. Cheng, and S. Avidan. Fast human detection using a cascade of histograms of oriented gradients. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1491–1498. IEEE Computer Society, 2006.

APPENDIX

A. PROOF OF LEMMAS

A.1 Proof of Lemma 5.1

PROOF. Based on the expression for $\phi_m(X)$ from (3.2) and using the trigonometric equalities

$$\cos^2(A) + \sin^2(A) = 1,$$
$$\cos(A - B) = \cos(A)\cos(B) + \sin(A)\sin(B), \tag{A.1}$$

$\alpha_j(X)$ can be expressed as

$$\alpha_j(X) = \sqrt{1 + \gamma_j(X)} \tag{A.2}$$

where

$$\gamma_j(X) = \sum_{k=1, k\neq j}^{K} \exp\left(-\frac{2\beta_{jk}(X)}{\tau}\right) + 2\sum_{k=1}^{K}\sum_{m=k+1}^{K} \exp\left(-\frac{\beta_{jk}(X) + \beta_{jm}(X)}{\tau}\right)\cos\left(\frac{\beta_{km}(X)}{\tau}\right). \tag{A.3}$$

If $k \neq j$, then $\beta_{jk}(X) > 0$ and hence

$$\lim_{\tau\to 0} \exp\left(-\frac{\beta_{jk}(X)}{\tau}\right) = 0. \tag{A.4}$$

As the cosine term in the expression for $\gamma_j(X)$ can be bounded independent of τ, we have

$$\lim_{\tau\to 0} \gamma_j(X) = 0 \tag{A.5}$$

and the result follows. As the points $X \in \mathcal{D}_j^\epsilon$ are *strictly* away from the Voronoi boundary of $\mathcal{D}_j$, the convergence is also uniform. □

A.2 Proof of Lemma 5.2

PROOF. In comparison to (A.4) we can deduce a stricter form of convergence, namely

$$\lim_{\tau\to 0} \frac{\exp\left(-\frac{\beta_{jk}(X)}{\tau}\right)}{\tau} = 0 \tag{A.6}$$

uniformly $\forall X \in \mathcal{D}_j^\epsilon$. By applying Lemma 5.1 it then follows that $\frac{\eta_j(X)}{\tau}$ converges *uniformly* to zero as $\tau \to 0$. This allows us to swap the double integral over $\mathcal{D}_j^\epsilon$ and the limit for τ leading to the desired result. □

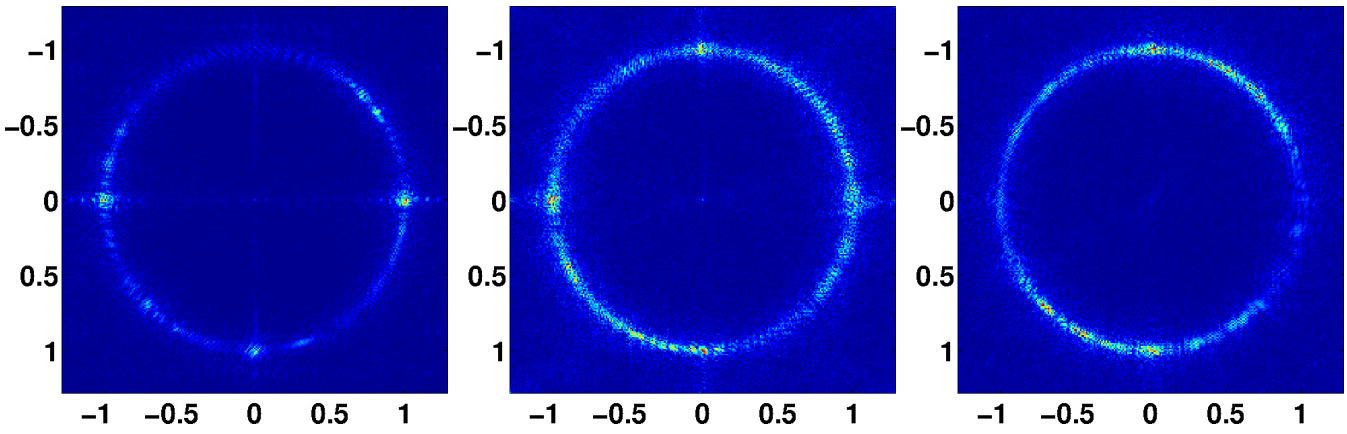

Figure 1: Fourier transform of the phase displayed close to the unit circle

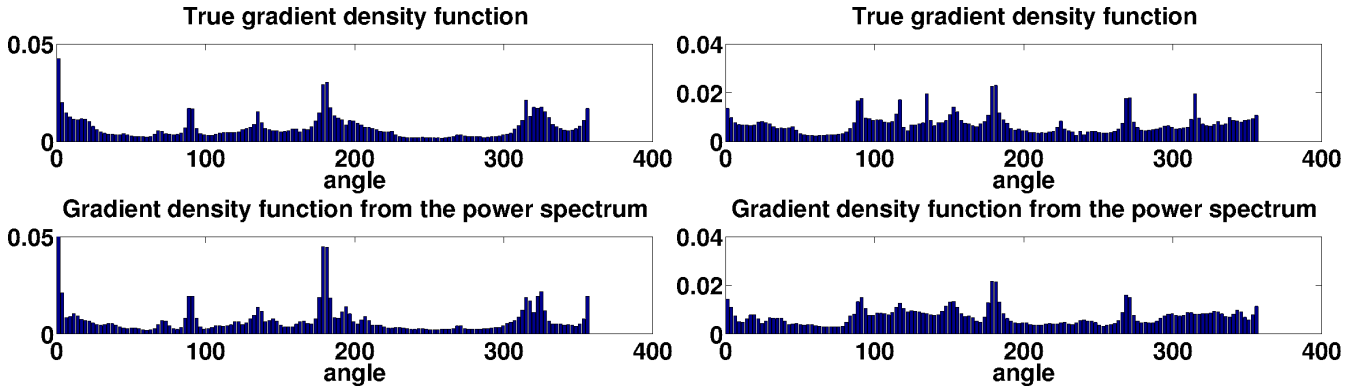

Figure 2: Comparison of the density functions. (i) Top: True orientation density, (ii) Bottom: Density function obtained from the power spectrum.

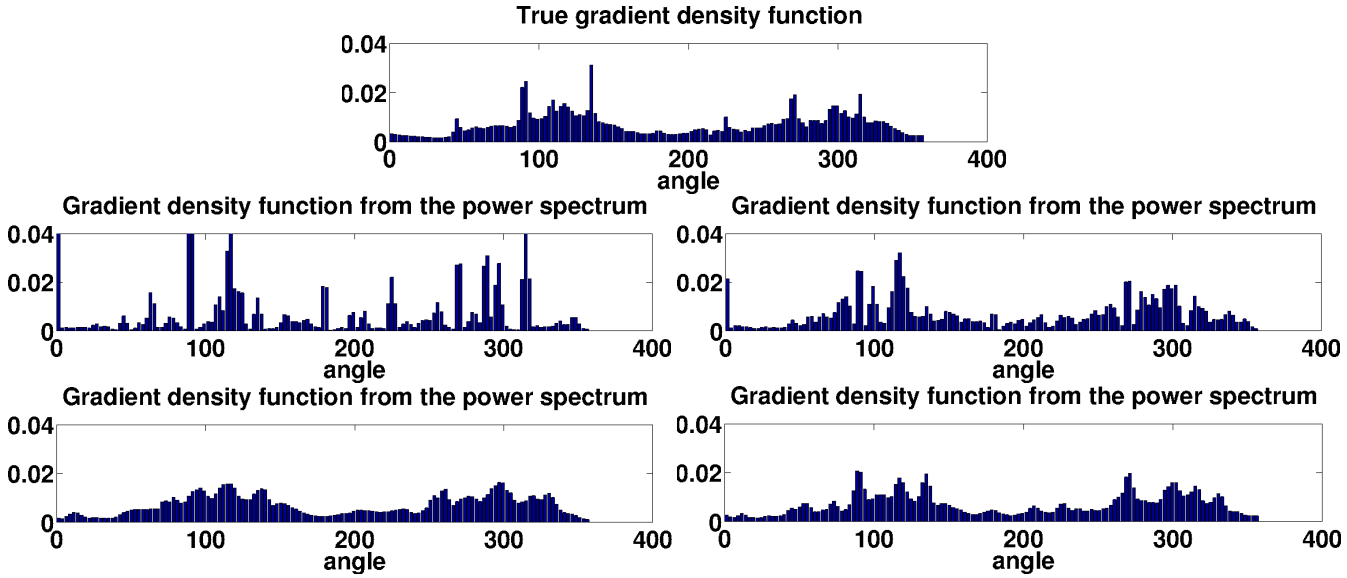

Figure 3: Convergence of the density function for $\tau \in \{0.01, 0.006, 0.001, 0.0004\}$. Read left to right and top to bottom.

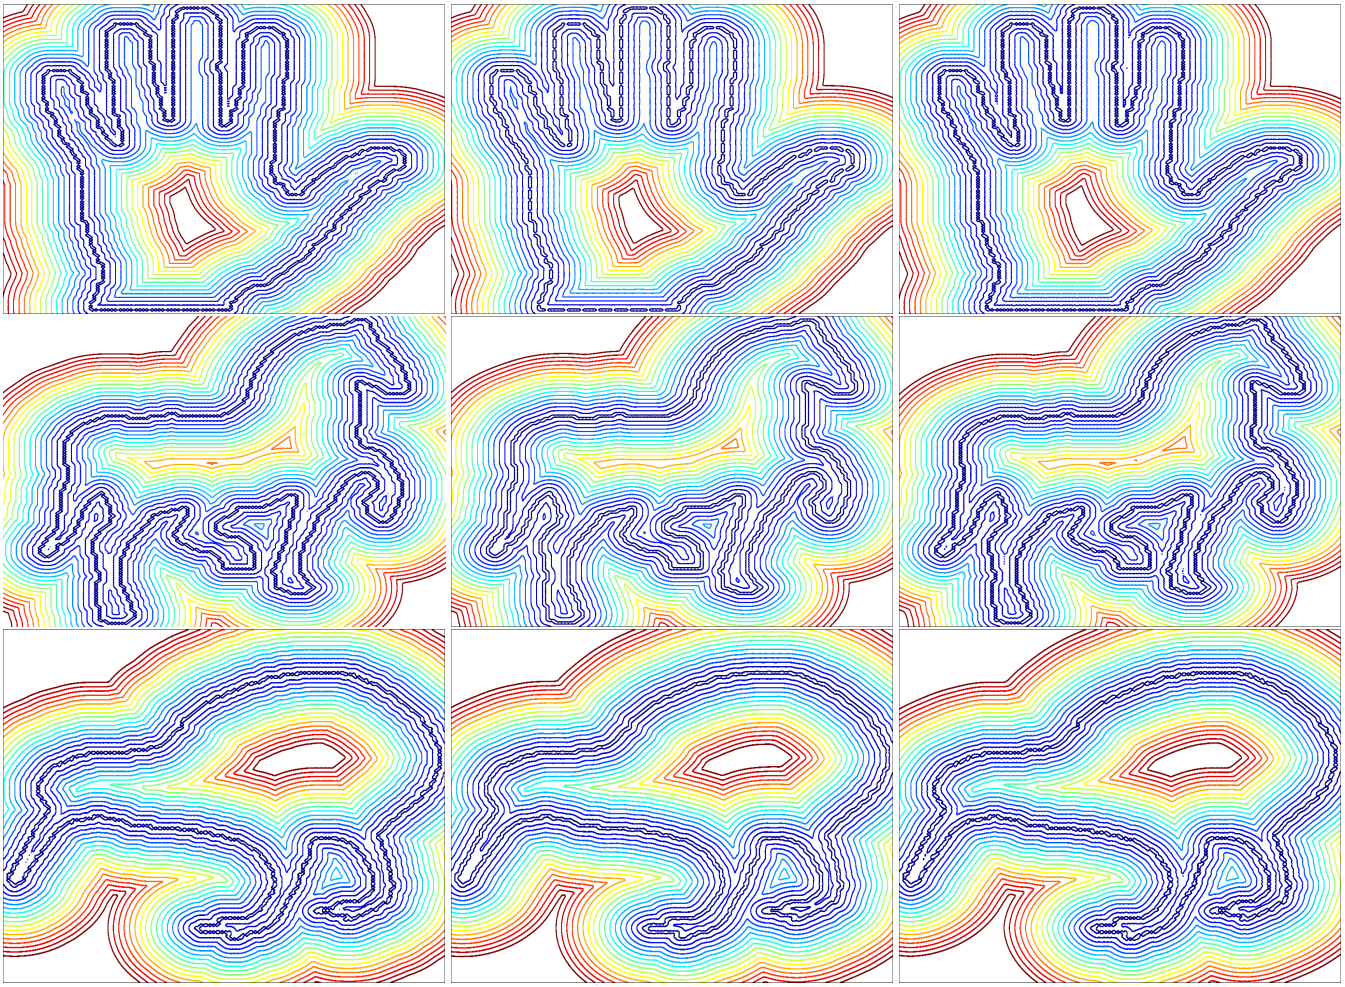

Figure 4: Contour plots: (i) Left: True Euclidean distance function, (ii) Center: Schrödinger distance transform at $\tau = 0.0004$, (iii) Right: Fast sweeping