The Schrödinger Distance Transform (SDT) for point-sets and curves

Manu Sethi, Anand Rangarajan and Karthik Gurumoorthy
Department of Computer & Information Science & Engineering
University of Florida
msethi@cise.ufl.edu, anand@cise.ufl.edu, sgk@ufl.edu

Abstract

Despite the ubiquitous use of distance transforms in the shape analysis literature and the popularity of fast marching and fast sweeping methods—essentially Hamilton-Jacobi solvers, there is very little recent work leveraging the Hamilton-Jacobi to Schrödinger connection for representational and computational purposes. In this work, we exploit the linearity of the Schrödinger equation to (i) design fast discrete convolution methods using the FFT to compute the distance transform, (ii) derive the histogram of oriented gradients (HOG) via the squared magnitude of the Fourier transform of the wave function, (iii) extend the Schrödinger formalism to cover the case of curves parametrized as line segments as opposed to point-sets, (iv) demonstrate that the Schrödinger formalism permits the addition of wave functions—an operation that is not allowed for distance transforms, and finally (v) construct a fundamentally new Schrödinger equation and show that it can represent both the distance transform and its gradient density—not possible in earlier efforts.

1. Introduction

Euclidean distance functions—more popularly (and henceforth) referred to as distance transforms—are widely used in shape and image analysis [15], graphics and medical imaging [19]. Since distance transforms allow us to transition from shapes (usually parametrized as a set of points or a set of curves) to a scalar field, problems such as shape registration can be couched in terms of (rigid, affine, non-rigid) alignment of distance transform fields. Signed and unsigned distance transforms are deployed in 3D as well with their zero level-sets corresponding to surfaces. Furthermore, medial axis methods and skeletonization often draw upon distance transform representations. Efficient algorithms that compute signed and unsigned distance transforms have proliferated in the past three decades [21, 23, 6] with Voronoi approaches on the one hand and fast marching (and fast sweeping methods) on the other cemented as

standard bearers at the moment.

Unsigned distance transforms on point-sets can be visualized as a set of intersecting cones with a point located at the tip of each cone. Signed distance transforms—which are meaningful only in the context of shapes parametrized as non-interlocking closed curves or surfaces—are continuous across the shape boundary but discontinuous at the Voronoi boundaries. Consequently, distance transforms are highly constrained geometric entities. These geometric constraints act as an impediment to shape analysis: for example, it is a non-trivial task to compute shape averages and/or atlases of distance transforms. Despite the fact that distance transform gradients are unit vectors and therefore the histogram of oriented gradients (HOG) of distance transforms is one-dimensional, these properties have not been widely leveraged in shape recognition and indexing. In addition, despite notable exceptions such as the generalized Voronoi diagrams which work on line segments and curves, we have not seen explicit representations of distance transforms of curves (as opposed to point-sets) in the literature.

The above problems with distance transforms help set up the Schrödinger distance transform (SDT) representation introduced here. We first briefly describe the theoretical development leading to the SDT and then highlight its properties while contrasting them with the problems mentioned above. The unsigned distance transform on point-sets is a scalar field $S(X)$ satisfying $\|\nabla S(X)\| = 1$ everywhere (except at the Voronoi centers where $S(X) = 0$ and at the Voronoi boundaries). This is a nonlinear differential equation, but more importantly, is an example of a static, *nonlinear* Hamilton-Jacobi (HJ) equation. It is widely known in the theoretical physics literature [3] that the static, nonlinear Hamilton-Jacobi equation can be embedded in a static, *linear* Schrödinger equation: $-\tau^2 \nabla^2 \psi(X) + \psi(X) = \psi_0(X)$ where $\psi_0(X)$ is non-zero at the point-set locations and τ a free parameter. When we set $\psi(X) = \exp\left\{-\frac{S(X)}{\tau}\right\}$ (and not $\exp\left\{i\frac{S(X)}{\tau}\right\}$ with the distinction elaborated upon below), it is easy to show that we obtain the HJ equation in the limit as $\tau \to 0$. The linearity of the Schrödinger equa-

tion enables the development of new $O(N \log N)$ (where N is the number of grid points) Fast Fourier Transform (FFT) based discrete convolution algorithms for computing unsigned distance transforms. When we set $\psi(X)$ to $\exp\left\{i\frac{S(X)}{\tau}\right\}$, the squared magnitude of the Fourier transform $\Psi(u, v; \tau)$ yields an increasingly tight approximation to the HOG of the distance transform as $\tau \to 0$. To summarize, the transition from nonlinear HJ to linear Schrödinger enables a wave function representation of distance transforms with concrete payoffs in terms of fast algorithms for distance transform and HOG computation.

The motivation for this paper is three-fold. First, we note that most previous work concentrates on distance transform computation on point-sets. Even when the underlying data are closed curves in 2D, fast marching methods (and related fast sweeping methods) essentially solve for distance transforms from point-sets since the curves are discretized as points on a grid. Other than the generalized Voronoi diagram, there is not much previous work on computing distance transforms for curves (parametrized as line segments). The first goal of this paper is to extend the Schrödinger distance transform approach to computing unsigned distance transforms for a set of line segments. The payoff is in the efficient representation of long line segments as seen in figure 1—an aerial image from Google Sketchup which consists of long line segments for which we compute the SDT.
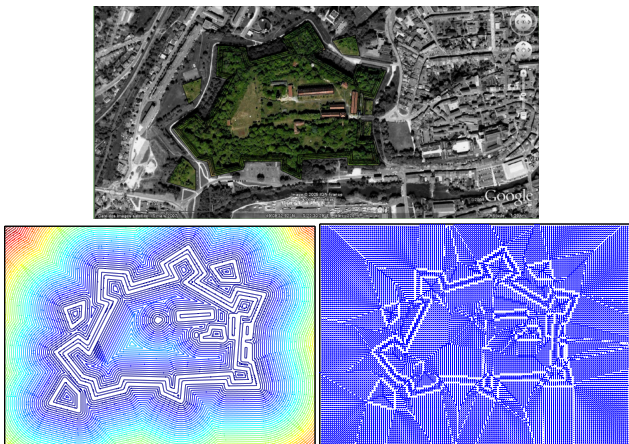

Figure 1: Top: Aerial image of Citadelle de Verde with its wireframe model superimposed. Bottom left: Schrödinger Distance Transform (SDT) computed from line segments. Bottom right: quiver plot of SDT gradients.

Second, please recall from above that a major stumbling block in the computation of shape statistics of distance transforms is the fact that the geometric constraints do not permit the addition of distance transforms. When we switch to the SDT representation, we show that we can add wave functions and obtain meaningful results. Consequently, we can now envisage the computation of SDT atlases and the like which is a step forward for distance transform shape analysis. Finally, we resolve a technical issue alluded to above. Fast SDT computation requires us to use the representation $\psi(X) = \exp\left\{-\frac{S(X)}{\tau}\right\}$ whereas the computation of the HOG of distance transforms forces us to use the representation $\psi(X) = \exp\left\{i\frac{S(X)}{\tau}\right\}$. This asymmetry is theoretically troubling and practically cumbersome. A single representation capable of representing the distance transform *and* its gradient density would be quite useful. To this end, we formulate a new static Schrödinger equation and show that it is capable of representing both the distance transform and the gradient density. Results are presented for all three aspects of this paper.

2. Relationship to Previous Work

As previously mentioned, distance transforms have seen wide deployment in the literature. We restrict focus on previous work that impacts on the three aspects of this paper: (i) Schrödinger distance transforms on curves, (ii) addition of SDTs and (iii) a unified representation for both distance transforms and gradient densities.

Voronoi and Delaunay data structures (in 2D and 3D) are the traditional methods in computer science for efficiently computing distance transforms [6]. In the past thirty years, static Hamilton-Jacobi solvers such as fast marching and fast sweeping have entered the picture. The fast marching method in [16] is a well known algorithm and has a complexity of $O(N \log N)$, where N is the number of grid points. Recently, clever use of an *untidy priority queue* data structure reduces the complexity of fast marching to $O(N)$. Fast sweeping methods are also popular [23] and also (empirically) solve the problem in $O(N)$. In the past two years, static Schrödinger equation solvers with a complexity of $O(N \log N)$ have been used for distance transform computation [11]. It should be mentioned that all of the above approaches essentially convert curves in 2D into points on a grid and then solve for the distance transform. As such, other than the Generalized Voronoi Diagrams (GVD) [9], there is no previous work on computing distance transforms for a set of curves (where the curves are not merely represented as points on a grid). GPU based implementations of GVDs, for example, in [13] and the jump flooding algorithm in [18] are prevalent in the area of graphics. To our knowledge, even these approaches, with the only exception of [22], discretize the curves by rasterizing them into their constituent sets of pixels. In recent work [2], Voronoi diagrams for points were generalized to regions, where the GVD was obtained by taking the union of Voronoi polygons of pixels lying on the boundary of a shape. This again highlights the need to obtain a distance transform for curves and/or line segments.

While distance transforms have been used before for shape alignment, there is very little previous work on using distance transform representations for shape atlas computation. In [4], the distance transform is converted into a density function for shape representation followed by atlas computation in the space of probability densities. The linearity of the SDT is not leveraged here. In [17], distance transforms are converted into density functions using log-odds followed by computations in the space of density functions. In contrast to these approaches, we directly use the linearity of the SDT representation to add wave functions— an operation that is not permitted for distance transforms since addition of two distance transforms does not result in a distance transform. To our knowledge this is new.

The third contribution of this paper is somewhat technical since it focuses on a unified representation of distance transforms and gradient densities. Since the advent of HOG [5] a few years ago, gradient density estimation has risen in prominence. As the distance transform gradients are unit vectors, the gradient density is one dimensional and defined over the space of orientations. As previously mentioned, in [11], the SDT was computed using the representation $\psi(X) = \exp\left\{-\frac{S(X)}{\tau}\right\}$ whereas the gradient density of the SDT was computed using the representation $\psi(X) = \exp\left\{i\frac{S(X)}{\tau}\right\}$ [12]. In this paper, these two disparate representations are unified using a new static Schrödinger equation with a solution $\psi_C(X)$ efficiently computed using an FFT. The novelty here is in its simplicity. The SDT is obtained from the magnitude $|\psi_C(X)|$ and the HOG is obtained from the phase $\frac{\psi_C(X)}{|\psi_C(X)|}$.

The original SDT computed using $\psi(X) = \exp\left\{-\frac{S(X)}{\tau}\right\}$ emerged from an inhomogeneous *screened Poisson* partial differential equation [11]. In [10], a solution to the Poisson equation was used for shape representation where the authors point out that distance transform-based solutions have contributions from only one site point at a grid point and are discontinuous near the Voronoi boundaries. Their representation is smoother than Voronoi since the solution at each grid point depends on many centers. In our work, the static Schrödinger equation is an inhomogeneous screened Poisson equation. We have a free parameter τ which can be increased to give a smoother solution at each site point (by taking contributions from Voronoi centers other than the closest) or reduced to get as close an approximation to the true distance transform as desired.

3. The Schrödinger distance transform (SDT) and its HOG for point-sets

In this section, in an effort to make the paper self contained, we briefly summarize the SDT [11] and its associ-

ated gradient density (HOG) [12] for point-sets. We begin by demonstrating the well known relationship between the nonlinear static Hamilton-Jacobi equation and the static, linear Schrödinger equation.

The static HJ equation for distance transform computation from a 2D point-set $\left\{Y_k \in \mathbb{R}^2, k = 1, \ldots, K\right\}$ is

$$\|\nabla S^*(X)\| = 1 \qquad (1)$$

almost everywhere (except at the point-set locations) with $S^*(X) = 0$ at the point-set Y. Here X is a 2D grid of N locations and the solution is the distance transform $S^*(X)$. The corresponding static Schrödinger equation is

$$-\tau^2 \nabla^2 \psi(X) + \psi(X) = \sum_{k=1}^{K} \psi_0(X - Y_k) \qquad (2)$$

where τ is a free parameter (but usually Planck's constant $\hbar$ in physics) and $\psi(X)$ a wave function. $\psi_0(X)$ is a square integrable function which is typically non-zero only in a small neighborhood around the origin. When we substitute $\psi(X) = \exp\left\{-\frac{S(X)}{\tau}\right\}$ (and $\psi_0(X) = \exp\left\{-\frac{S_0(X)}{\tau}\right\}$ with $S_0(X)$ chosen to ensure that $\psi_0(X)$ is equal to zero almost everywhere), we get

$$\|\nabla S(X)\|^2 - \tau \nabla^2 S(X) = 1 - \sum_{k=1}^{K} \exp\left\{-\frac{S_0(X - Y_k) - S(X)}{\tau}\right\}. \qquad (3)$$

This can be shown to approach (1) almost everywhere (as $\tau \to 0$) except at the point-set Y.

A Green's function approach [1] can be used to obtain the solution to (2) which in 2D [and with a suitable choice of $\psi_0(X)$] is approximately

$$\psi(X) = \sum_{k=1}^{K} \exp\left(-\frac{\|X - Y_k\|}{\tau}\right). \qquad (4)$$

The Schrödinger distance transform (SDT) is obtained using $S(X) = -\tau \log\{\psi(X)\}$ and has been shown in [11] to converge to the true Euclidean distance function as $\tau \to 0$. An FFT-based fast convolution method is used to compute (4) in $O(N \log N)$ time. The SDT for curves developed in the next section follows this basic treatment.

In figure 2 (left), we use a fish image from our Gator-Bait dataset[1] and plot the SDT contours computed using (4). We also pitted the SDT against the fast sweeping algorithm [23] on a 2D grid consisting of points randomly chosen between $(-0.123, -0.123)$ and $(0.123, 0.123)$ with a grid width of $\frac{1}{2^{10}}$. The number of grid points equals $N = 253 \times 253 = 64,009$. We ran 100 experiments, each time randomly choosing $10,000$ grid points as data points. We set $\hbar = 0.0001$ for the SDT and ran the fast sweeping method for 10 iterations—sufficient for it to converge.

---

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

The plot in figure 2 (right) shows the average percentage error calculated according to (4) for both these techniques in comparison to the true Euclidean distance function. From the plot, it is clear that while the fast sweeping method has a percentage error of around 7%, the SDT gave a percentage error of less than **1.5%** providing much better accuracy.
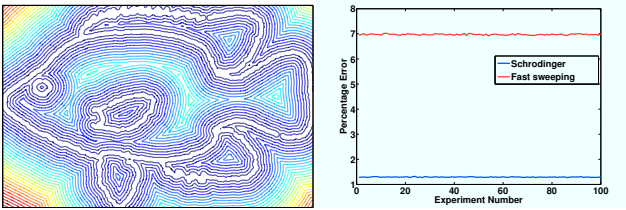

Figure 2: Left: SDT on a fish point-set (from the GatorBait database). Right: Comparison to fast sweeping on random point-sets.

We now turn to distance transform gradient density estimation. Since distance transforms satisfy $\|\nabla S^*(X)\| = 1$ almost everywhere, distance transform gradients are unit vectors. The HOG for distance transforms is one dimensional and defined over the set of orientations. Consider a slightly different static Schrödinger equation from (2) above:

$$-\tau^2 \nabla^2 \psi(X) - \psi(X) = 0 \qquad (5)$$

which is a Helmholtz equation in contrast to the inhomogeneous, screened Poisson equation in (2). When we substitute $\psi(X) = \exp\left(i\frac{S(X)}{\tau}\right)$ in (5), where we have set $|\psi(X)| = 1$ (a vast simplification), we get

$$\|\nabla S(X)\|^2 - i\tau \nabla^2 S(X) = 1 \qquad (6)$$

which again approaches (1) as $\tau \to 0$. It can be shown [12] that the squared magnitude of the Fourier transform of $\psi(X)$, namely $|\Psi(u, v; \tau)|^2$ approaches the true density of the distance transform gradients as $\tau \to 0$. Since the gradients are unit vectors, the Fourier transform values mainly lie on the unit circle (as seen on the right in figure 3) and the squared magnitude of the FFT can be reparametrized as $|\Psi(u(\theta), v(\theta); \tau)|^2$. Consequently, *spatial frequencies on the unit circle become HOG bins*. The proof uses the stationary phase approximation [14] and standard analysis. Note the discrepancy between the two representations with the distance transform relying on a screened Poisson formulation but with its gradient density using a Helmholtz formulation. These two disparate formulations are unified in this paper. In figure 3, we use a fish image to compare the true gradient density with the SDT HOG.
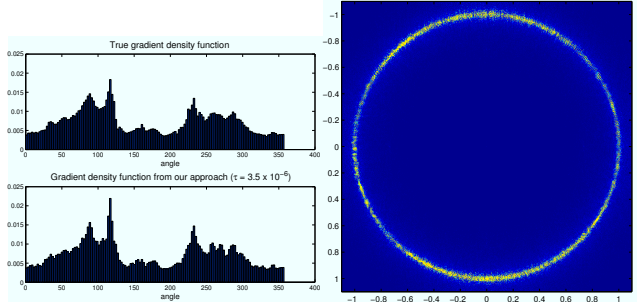

Figure 3: Left: HOG for the distance transform (fish). Right: $|\Psi(u, v; \tau)|^2$—the squared magnitude of the Fourier transform of $\psi(X)$ showing values lying mainly on the unit circle in the spatial frequency domain.

4. The Schrödinger distance transform for a set of curves

4.1. Solving the Schrödinger equation using Laplace's method

Assume now that we are given a set of curves parametrized as line segments. Let $\{Y_k, k \in 1, \ldots, K\}$ denote a set of given line segments. We now seek the Schrödinger distance transform for this set of line segments. We follow the approach of the previous section and construct a Schrödinger solver for this problem. Let $Y_1^{(k)}$ and $Y_2^{(k)}$ be the end points of the kth line segment in 2D with $Y_1^{(k)}, Y_2^{(k)} \in \mathbb{R}^2$. Let $Y_0^{(k)} \in \mathbb{R}^2$ be any point on the line segment. Then $Y_0^{(k)}$ can be represented as $Y_0^{(k)}(\alpha) = Y_1^{(k)}(1 - \alpha) + Y_2^{(k)}\alpha$, where $\alpha \in [0, 1]$. We seek the solution to

$$-\tau^2 \nabla^2 \psi(X) + \psi(X) = \sum_{k=1}^{K} \int_0^1 \psi_0(X - Y_0^{(k)}(\alpha)) d\alpha. \quad (7)$$

As before $\psi_0(X)$ can be chosen to be a square integrable function which is equal to zero almost everywhere except in a neighborhood around the origin. When we choose an appropriate "box" function for $\psi_0(X)$, the solution to the above equation can be obtained using the Green's function approach:

$$\psi(X) = \sum_{k=1}^{K} \int_0^1 K_0\left(\frac{\|X - Y_0^{(k)}(\alpha)\|}{\tau}\right) d\alpha \qquad (8)$$

where K_0 is the zeroth order modified Bessel function of the second kind. This is a difficult integral to approximate but following [11], we replace K_0 by the exponential function (which is justified since the two functions become nearly identical almost everywhere except at the origin as $\tau \to 0$). With this approximation in place, we obtain a series expansion of the integral in (8) using the Laplace method for in-

tegral asymptotics [20]. We briefly describe the application of this method below for a single line segment denoted by Y (with the line segment superscript index suppressed). The final solution is a summation over all line segments.

$$\psi(X) = \int_0^1 \exp\left(-\frac{\sqrt{A\alpha^2 + B\alpha + C}}{\tau}\right) d\alpha \quad (9)$$

where $A = \|Y_1 - Y_2\|^2$, $B = -2\langle(X - Y_1), (Y_1 - Y_2)\rangle$ and $C = \|X - Y_1\|^2$. Now if we substitute $\sqrt{A\alpha^2 + B\alpha + C} = y$ and then further substitute $4Ay^2 = (4AC - B^2)\cosh^2\theta$, then the above integral in (9) becomes

$$\psi(X) = \frac{\sqrt{4AC - B^2}}{2A} \int_a^b \cosh(\theta) \exp\left(-\frac{h(\theta)}{\tau}\right) d\theta \quad (10)$$

where $h(\theta) \equiv \sqrt{\frac{4AC - B^2}{4A}} \cosh(\theta)$ and a and b the new limits.

Now we can split this integral into two parts at the minimizer $\theta_m \in [a, b]$ of $h(\theta)$ and then integrate both the parts separately and add them to arrive at an expression for $\psi(X)$ in place of (10). To integrate each part we use the Laplace method [20] which gives expansions in powers of τ. Denoting the first part as $\psi_1(X)$, we obtain

$$\psi_1(X) = \frac{\sqrt{4AC - B^2}}{2A} \exp\left\{-\frac{1}{\tau}h(\theta_m)\right\} \sum_{s=0}^{\infty} \Gamma\left(\frac{s+1}{2}\right) c_s \tau^{\frac{s+1}{2}} \quad (11)$$

with expressions available for the coefficients c_s. For both parts, the solution only depends on the minimizer θ_m and does not depend upon either a or b, unless the function $h(\theta)$ is monotonic in $[a, b]$. So, the solution to (10) turns out to be $\psi(X) = 2\psi_1(X)$ if $\theta_m \in (a, b)$ or $\psi(X) = \psi_1(X)$ if $\theta_m \in \{a, b\}$ when $h(\theta)$ is monotonic. As $\tau \to 0$, we only require a few terms to get a good approximation to $\psi(X)$. For the case where $\theta_m \in (a, b)$, θ_m always turns out to be 0 because of the substitutions that we have made in going from (9) to (10). We restricted the expansion of (11) to 5 terms in our experiments which is sufficient to get a good approximation. The values of c_s obtained are: $c_0 = \frac{1}{\sqrt{2}}\sqrt{\beta}$, $c_1 = 0$, $c_2 = \frac{3}{8}\sqrt{2}\beta^{3/2}$, $c_3 = \frac{5}{32}\beta^2$ and $c_4 = -\frac{7\sqrt{2}}{32}\beta^{5/2}$ where $\beta = \sqrt{\frac{4A}{4AC - B^2}}$. For the case where θ_m is either a or b we restricted the expansion to only one term in our experiments for which we get $c_0 = \frac{C^{1/4}}{\sqrt{2}}$ or $\frac{(A+B+C)^{1/4}}{\sqrt{2}}$ respectively. After adding the contributions from all the line segments, the approximate distance transform can be obtained from the relation $S(X) = -\tau \log\{\psi(X)\}$. It can be seen from (11) that as $\tau \to 0$, the solution to $S(X)$ approaches the true distance transform. The SDT for a set of curves parametrized as line segments is therefore available as a truncated series expansion. Furthermore, we can analytically differentiate this series expansion to obtain the distance transform gradients as well. We show the SDT con-

tours and its gradients in figures 5 and 6 respectively.
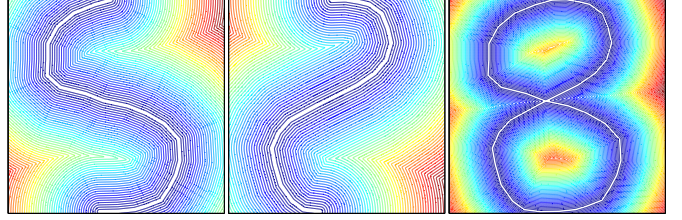
(a) Left and middle: Contour plots of the two SDTs separately computed. Right: Contour plot of the SDT computed by adding the wave functions obtained from the S and reversed S shapes.
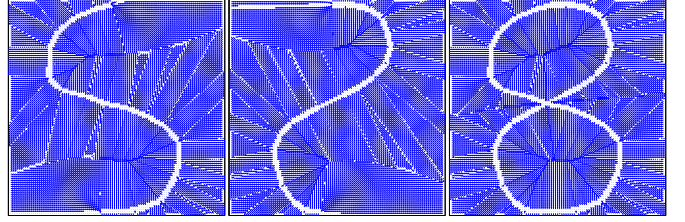
(b) Gradients corresponding to the SDTs computed above.

Figure 4: A figure "8" obtained by adding wave functions of an S and reversed S shapes.

4.2. Exploiting the linearity of the SDT formalism: Adding wave functions

We have mentioned above that the linearity of the Schrödinger equation permits us to add solutions and create a composite distance transform. Assume that we have two curves with each comprising a set of line segments. Further assume that we have successfully obtained solutions $\psi^{(1)}(X)$ and $\psi^{(2)}(X)$ to the two differential equations which differ only in the $\sum_{k=1}^{K} \int_0^1 \psi_0(X - Y_0^{(k)}(\alpha))d\alpha$ terms on the right hand side of (7).

Given the two distance transforms $S^{(1)}(X)$ and $S^{(2)}(X)$ for the two curves, we are not permitted to add them in order to produce a distance transform for the union of the two curves. However, we can add the two wave functions since both are solutions to linear differential equations. Consequently, we can obtain a new distance transform via $S(X) = -\tau \log(\psi^{(1)}(X) + \psi^{(2)}(X))$. We now give an intuitive explanation for this operation. Since we are only interested in solutions at small values of τ, at each location X, the negative exponentials in the solution to $\psi(X)$ essentially compete for that point. Another way of seeing this is to note that the same mechanism is present in Voronoi diagrams where centroids compete to gain membership of grid points. When we have line segments instead of points, despite the fact that the membership mechanism is more complicated, the behavior is essentially the same. The series expansions in $\psi^{(1)}(X)$ and $\psi^{(2)}(X)$ compete for membership of the location X. As $\tau \to 0$, the competition intensifies
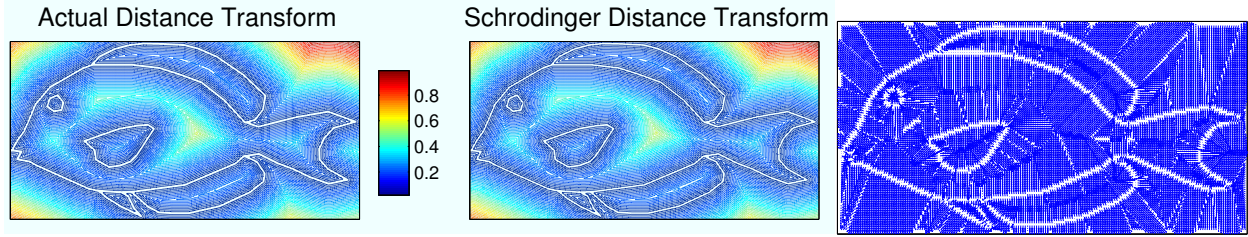
Figure 5: Left and middle: True distance transform and SDT for a fish (line segments). Right: Gradients of the fish SDT (line segments).
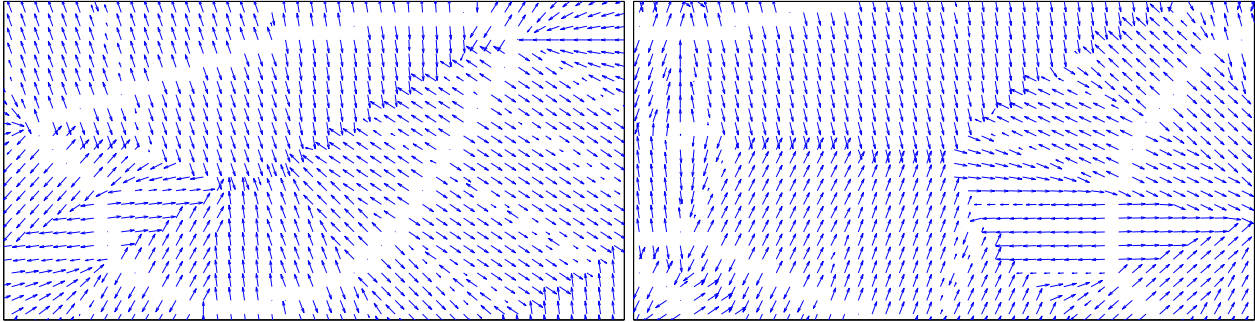

Figure 6: Zoomed quiver plots of the gradients at the middle fin (left) and the tail (right)

and the exact distance transform achieved via the nonlinearity $S(X) = -\tau \log(\psi^{(1)}(X) + \psi^{(2)}(X))$. From this intuitive description, we see that wave functions can indeed be summed but distance transforms cannot. When wave functions are summed, the range of competition at each location increases. Subsequently, when the logarithmic non-linearity is invoked, we obtain good approximations to the true distance transforms at small values of τ. In figure 4, we showcase a simple example of this overall approach. We take two shapes: an S and a reversed S, both composed of line segments and show that the distance transform of the figure 8 (present in the union of the two curves) is obtained by adding the wave functions *separately obtained* from each S shape.

4.3. Results and implementation details

We evaluated the SDT of lines for a 2D image from our fish database and based our results against the actual distance transform computed using line segments. On an image of size 903×460, the grid size chosen was $0 \leq x \leq 4.6053$ and $0 \leq y \leq 2.3460$ and the the grid width was kept as 0.0051. The value of τ was set to 6×10^{-4}. We obtained 92 line segments for the fish image from which the SDT is computed using $S(X) = -\tau \log \{\psi(X)\}$. We resort to the GNU MPFR[2] and GMP[3] libraries to compute our results because for our grid size and the comparatively very small

value of τ, double precision values did not suffice. Figure 5 displays the contours and the gradients of the SDT for line segments and figure 6 displays the zoomed quiver plot for the gradients in figure 5. The ℓ_1 norm of the error between the SDT computed from line segments and the ground truth was 7%.

5. A unified representation for the SDT and its HOG

In Section 3, we briefly summarized previous work on the SDT and its gradient density (HOG). A noteworthy point there was the use of two different Schrödinger equations—an inhomogeneous, screened Poisson for the SDT and a Helmholtz for the HOG. We now integrate the two disparate representations.

5.1. A new complex, static Schrödinger equation

Consider the following Schrödinger equation:

$$-\tau^2 \nabla^2 \psi_C(X) - 2i\psi_C(X) = \sum_{k=1}^{K} \psi_0(X - Y_k). \quad (12)$$

The notation $\psi_C(X)$ is used to emphasize that this is a new, complex static Schrödinger equation. (Static Schrödinger equation for ordinary particles are never complex in the physics literature.) Note that this is neither Poisson nor Helmholtz. The closed-form solution for this equation (in 2D and for an unbounded domain) is
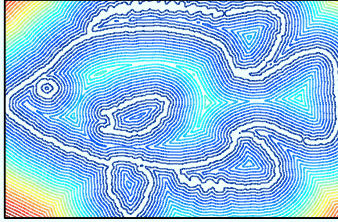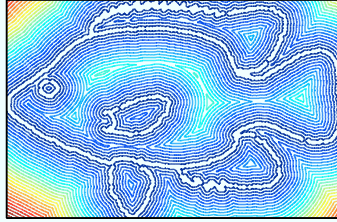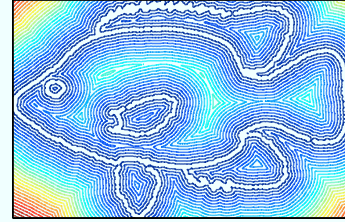

Figure 7: Left: Unified SDT (point-set). Middle: Fast sweeping. Right: Brute force (ground truth).

$$\psi_C(X) = \sum_{k=1}^{K} K_0 \left\{ (1-i) \frac{\|X - Y_k\|}{\tau} \right\}$$

$$\approx \sum_{k=1}^{K} \exp \left\{ (-1+i) \frac{\|X - Y_k\|}{\tau} \right\} \quad (13)$$

with the transition from K_0 to the exponential mirroring the earlier transition in Section 3.

While we have obtained the above solution using a Green's function approach and implemented it using the FFT-based fast convolution method described in [11], it is essential to point out that other multigrid based iterative solutions [7, 8] exist for solving heterogenous Helmholtz equations like the one we have in (12).

We now detail how this wave function contains both the SDT and its HOG in a single representation—something that was not possible in previous work. The wave function in (13) is complex and we can compute its magnitude and phase. The empirical results will show that the SDT computed from the magnitude and the HOG computed from the phase are good approximations to the true distance transform and gradient density. We have not yet formally proved this at the present time but think that the proof will follow the lines of previous work [11, 12].

Consequently, a new SDT can be obtained as

$$S_C(X) = -\tau \log \left| \sum_{k=1}^{K} \exp \left\{ (-1+i) \frac{\|X - Y_k\|}{\tau} \right\} \right| \quad (14)$$

and the new HOG as

$$P_C(u(\theta), v(\theta); \tau) = |\Psi_C(u(\theta), v(\theta); \tau)|^2 \quad (15)$$

where $\Psi_C(u(\theta), v(\theta); \tau)$ is the 2D Fourier transform of the phase $\frac{\psi_C(X)}{|\psi_C(X)|}$ evaluated on the unit circle. Note that we can continue to use fast convolution methods (via the FFT) to compute the SDT from (14) and the HOG from (15).

5.2. Experimental Results for the unified SDT

We implemented the above unified representation for the SDT of point sets on a fish image from our data set and compared it with the results obtained from brute force and fast sweeping approaches. As in the curves case, we use GNU MPFR and GMP libraries to avoid precision issues. The grid size for the fish image is $-0.450 \leq x \leq 0.450$ and $-0.293 \leq y \leq 0.293$. The grid width is 10^{-2} and the number of grid points is 587×900. The value of τ chosen is 9×10^{-3}. Figure 7 shows the contours plots obtained using (14) and also compares it with the fast sweeping method and the ground truth (brute force) result. Figure 8 shows the gradient density obtained using (15) and compares it with the true gradient density. Figure 9 displays the frequency spectrum showing that the FFT values are non-zero mainly on the unit circle. We calculate the ℓ_1 norm of the error between the the two gradient densities by subtracting the frequency values for corresponding bins and then summing their magnitudes. We experimented with various values of τ and found that the error decreased with a decrease in the value of τ. This is shown in figure 9 (right).

6. Contributions

Despite the well known relationship between the Hamilton-Jacobi and Schrödinger equations, only recently do we see shape analysis begin to use Schrödinger distance transform (SDT) representations. This paper builds on earlier work in this space.

We summarize the contributions of this paper in an itemized format.

- Briefly described earlier work on the Schrödinger distance transform (SDT) [11] and its gradient density (HOG) [12] using a static Schrödinger equation formalism.

- Explained that the SDT can be efficiently computed using an FFT in $O(N \log N)$ time for N grid points.

- Developed the SDT for sets of curves (parametrized as line segments) as opposed to merely point-sets. The Laplace approximation is the major technical device used here.

- Exploited the linearity of the SDT formalism by demonstrating that wave functions can be added whereas distance transforms cannot. This will facilitate the creation of atlases in the future.

- Constructed a unified representation for both the SDT and its HOG via a fundamentally new static Schrödinger equation. This is a major departure from previous work which could not represent both quantities in a single differential equation.

- Computed the SDT and its HOG in the new unified representation using an FFT in $O(N \log N)$ time for N grid points.

Immediate future work will focus on the creation of atlases in the SDT framework. We also plan to investigate the SDT in shape registration and indexing applications.
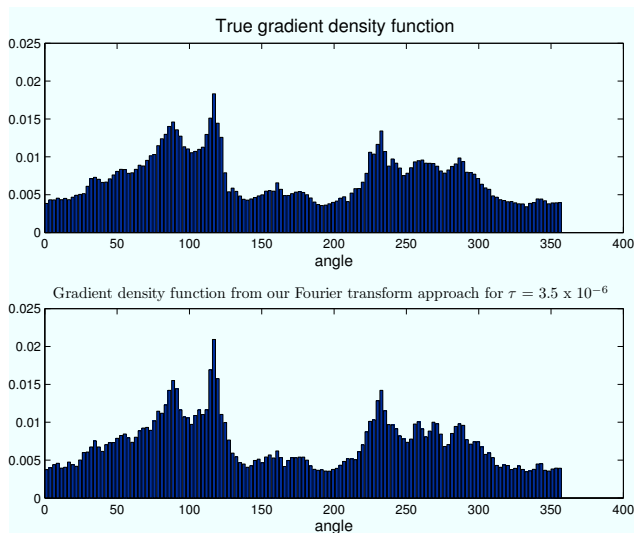

True gradient density function

Gradient density function from our Fourier transform approach for $\tau = 3.5 \times 10^{-6}$

Figure 8: HOG comparison: the true and the SDT gradient densities.

Acknowledgments

References

[1] M. Abramowitz and I. Stegun. *Handbook of mathematical functions with formulas, graphs and mathematical tables*. Dover, New York, NY, 1964. 3

[2] N. Ahuja and S. Todorovic. Connected segmentation tree - a joint representation of region layout and hierarchy. In *IEEE CVPR*, 2008. 2

[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. 1

[4] T. Chen, A. Rangarajan, S. J. Eisenschenk, and B. C. Vemuri. Construction of neuroanatomical shape complex atlas from 3D brain MRI. In *MICCAI (3)*, pages 65–72, 2010. 2
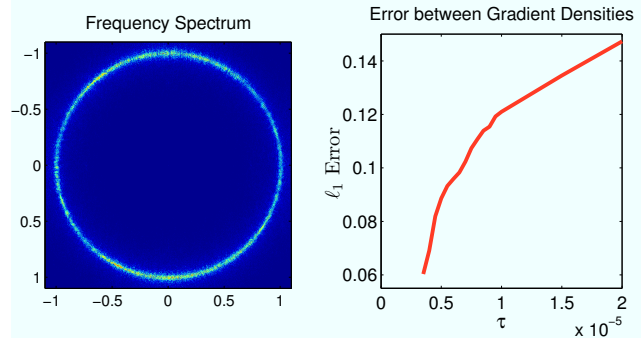

Figure 9: Unified SDT: Left: $\Psi(u, v; \tau)$—the Fourier transform of the *phase* of $\psi(X)$ showing values lying mainly on the unit circle in the spatial frequency domain. Right: Decrease in error of the unified SDT HOG as $\tau \to 0$.

[5] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE CVPR*, pages 886–893, 2005. 2

[6] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2010. 1, 2

[7] Y. A. Erlangga, C. W. Oosterlee, and C. Vuik. A novel multigrid based preconditioner for heterogeneous Helmholtz problems. *SIAM J. Scientific Computing*, 27(4):1471–1492, 2006. 5.1

[8] Y. A. Erlangga, C. Vuik, and C. Oosterlee. On a class of preconditioners for solving the Helmholtz equation. *Appl. Num. Math.*, 50:409–425, 2004. 5.1

[9] M. L. Gavrilova, editor. *Generalized Voronoi Diagram: A Geometry-Based Approach to Computational Intelligence*. Springer, 2008. 2

[10] L. Gorelick, M. Galun, E. Sharon, R. Basri, and A. Brandt. Shape representation and classification using the Poisson equation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 28(12):1991–2005, 2006. 2

[11] K. S. Gurumoorthy and A. Rangarajan. A Schrödinger equation for the fast computation of approximate Euclidean distance functions. In *SSVM*, pages 100–111. Springer, 2009. 2, 3, 3, 4.1, 5.1, 6

[12] K. S. Gurumoorthy, A. Rangarajan, and A. Banerjee. The complex wave representation of distance transforms. In *EMMCVPR*, pages 413–427, 2011. 2, 3, 3, 5.1, 6

[13] K. E. Hoff III, J. Keyser, M. C. Lin, D. Manocha, and T. Culver. Fast computation of generalized Voronoi diagrams using graphics hardware. In *SIGGRAPH*, pages 277–286, 1999. 2

[14] F. Olver. *Asymptotics and special functions*. Academic Press, New York, NY, 1974. 3

[15] S. Osher and R. Fedkiw. *Level set methods and dynamic implicit surfaces*. Springer-Verlag, October 2003. 1

[16] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *J. Comp. Phys.*, 79(1):12–49, 1988. 2

[17] K. M. Pohl, J. W. Fisher III, S. Bouix, M. E. Shenton, R. W. McCarley, W. E. L. Grimson, R. Kikinis, and W. M. Wells III. Using the logarithm of odds to define a vector space on probabilistic atlases. *Medical Image Analysis*, 11(5):465–477, 2007. 2

[18] G. Rong and T. S. Tan. Jump flooding in GPU with applications to Voronoi diagram and distance transform. In *SI3D*, pages 109–116, 2006. 2

[19] K. Siddiqi and S. Pizer, editors. *Medial Representations: Mathematics, Algorithms and Applications*. Computational Imaging and Vision. Springer, 2008. 1

[20] R. Wong. *Asymptotic approximations of integrals*. Academic Press, New York, NY, 1989. 4.1, 4.1

[21] L. Yatziv, A. Bartesaghi, and G. Sapiro. $O(N)$ implementation of the fast marching algorithm. *J. Comp. Phys.*, 212(2):393–399, 2006. 1

[22] Z. Yuan, G. Rong, X. Guo, and W. Wang. Generalized Voronoi diagram computation on GPU. ISVD, pages 75–82, 2011. 2

[23] H. Zhao. A fast sweeping method for eikonal equations. *Math. Comp.*, 74(250):603–627, 2005. 1, 2, 3