A Schrödinger wave equation approach to the eikonal equation: Application to image analysis

Anand Rangarajan and Karthik S. Gurumoorthy*

Department of Computer and Information Science and Engineering
University of Florida, Gainesville, FL, USA

Abstract. As Planck's constant $\hbar$ (treated as a free parameter) tends to zero, the solution to the eikonal equation $|\nabla S(X)| = f(X)$ can be increasingly closely approximated by the solution to the corresponding Schrödinger equation. When the forcing function $f(X)$ is set to one, we get the Euclidean distance function problem. We show that the corresponding Schrödinger equation has a closed form solution which can be expressed as a discrete convolution and efficiently computed using a Fast Fourier Transform (FFT). The eikonal equation has several applications in image analysis, viz. signed distance functions for shape silhouettes, surface reconstruction from point clouds and image segmentation being a few. We show that the sign of the distance function, its gradients and curvature can all be written in closed form, expressed as discrete convolutions and efficiently computed using FFTs. Of note here is that the sign of the distance function in 2D is expressed as a winding number computation. For the general eikonal problem, we present a perturbation series approach which results in a sequence of discrete convolutions once again efficiently computed using FFTs. We compare the results of our approach with those obtained using the fast sweeping method, closed-form solutions (when available) and Dijkstra's shortest path algorithm.

1 Introduction

While image analysis borrows liberally from classical mechanics—with variational principles [1], Euler-Lagrange equations, Hamiltonians [2] and Hamilton-Jacobi theory [3] all in widespread use at the present time—other than a few pioneering works [4,5], there isn't a concomitant borrowing from quantum mechanics. Given the very close relationship between Hamilton-Jacobi theory and Schrödinger wave mechanics, this is somewhat surprising. In this work, we begin with a brief overview of the classical mechanics sequence of i) variational principles and Euler-Lagrange equations [6], ii) Legendre transformations leading to the Hamiltonian [7], iii) canonical transformation of the Hamiltonian which yields the Hamilton-Jacobi theory [7], and finally iv) first quantization to obtain the Schrödinger wave equation [8]. This well known path of development needs an image analysis payoff which we next describe.

* The authors can be contacted at {anand,ksg}@cise.ufl.edu.

It has been a decade since EMMCVPR 1999 at which event we saw [2] the advent of Hamiltonian mechanics to solve the eikonal equation. More specifically, the Hamiltonian approach was also used to analyze the Euclidean distance function problem—an important special case of the eikonal problem wherein $|\nabla S(X)| = 1$ and X a regular grid. The Euclidean distance function problem in its image analysis incarnation can be stated as follows: Given a set of shape silhouettes whose boundaries are parameterized as piecewise smooth curves, compute the *signed* distance at every location on a grid w.r.t. the boundary points. Furthermore, we often seek the gradient, divergence, curvature and medial axes of the signed distance function which are not easy to obtain by other approaches such as the fast marching [9,10] and fast sweeping methods [11] due to the lack of differentiability of the signed distance function. In sharp contrast, we show—using our previous work on this topic [12]—that the Schrödinger wave equation approach to the eikonal results in a closed-form solution which can be expressed as a discrete convolution and computed in $O(N \log N)$ time using a Fast Fourier Transform (FFT) [13] where N is the number of grid points. While the fast marching method is also $O(N \log N)$ (and even $O(N)$ with cleverly chosen data structures [14]), these methods are based on spatial discretizations of the derivative operator (in $|\nabla S(X)| = 1$) whereas the Schrödinger approach does not require derivative discretization. A caveat is that our Euclidean distance function is an approximation since it is obtained for a small but non-zero value of Planck's constant $\hbar$.

The Schrödinger equation approach to the eikonal gives us an unsigned distance function. We complement this by independently finding the sign of the distance function in $O(N \log N)$ time on a regular grid in 2D. We achieve this by efficiently computing the *winding number* for each location in the 2D grid. The winding number is the number of times a closed curve winds around a point. We show that just as in the case of the Schrödinger equation, the winding number can also be written in closed-form, expressed as a discrete convolution and efficiently computed using an FFT. The fact that the winding number can be expressed as a discrete convolution for every location in a 2D grid appears to be a new contribution.

We also leverage the closed-form solution for the unsigned distance function obtained from Schrödinger. Since our distance function is differentiable everywhere, we can once again write down closed-form expressions for the gradients and curvature, express them as discrete convolutions and efficiently compute these quantities in $O(N \log N)$ using FFTs. We visualize the gradients and the maximum curvature using 2D shape silhouettes as the source. The maximum curvature—despite its fundamental drawback of being extrinsic—has a haunting similarity to the medial axis. To our knowledge, the fast computation of the derivatives of the distance function on a regular grid using discrete convolutions is new.

Next, we present a general eikonal solver using a Born expansion-based perturbation method [15]. We cannot solve for $S(X)$ in closed form here, and instead we express the solution as a sequence of discrete convolutions, with each term

efficiently computed using an FFT. We apply this method to image segmentation by first seeding a set of points on the interior and the exterior of the segmentation regions and then solving the eikonal using a forcing function $f(X)$ derived from the image gradients. These results are compared to those obtained using fast sweeping and Dijkstra's shortest path algorithm [16] (since the ground truth is not available). Since all results are obtained for very low values of $\hbar$, some numerical instability issues arise in the FFT-based convolutions. Consequently, higher precision numerical support for fast, discrete convolution is a fundamental requirement and one that we plan to address in future work.

2 A Schrödinger equation for the eikonal problem

In this section, we briefly review the Schrödinger equation approach to (unsigned) Euclidean distance functions. We begin with a Lagrangian variational principle, derive the Hamiltonian via a Legendre transformation, use a canonical transformation to obtain the Hamilton-Jacobi equation and finally quantize Hamilton-Jacobi to obtain the Schrödinger equation.

The Lagrangian variational principle for Euclidean distance functions is an objective function whose solution is the shortest distance between two points in $\mathbb{R}^d$—the Euclidean distance. While we use $d = 2$ for illustration purposes, the approach is general and not restricted to a particular choice of dimension:

$$I[q] = \int_{t_0}^{t_1} L(q_1, q_2, \dot{q}_1, \dot{q}_2, t)dt \tag{1}$$

where $q(t) = \{q_1(t), q_2(t)\}$ is a C^2 path between two points in time t_0 and t_1 with

$$L(q_1, q_2, \dot{q}_1, \dot{q}_2, t) = \frac{1}{2}\left(\dot{q}_1^2 + \dot{q}_2^2\right). \tag{2}$$

The corresponding Euler-Lagrange equations are

$$\ddot{q}_1(t) = 0, \text{ and } \ddot{q}_2(t) = 0 \tag{3}$$

which are tantamount to a straight line in 2D. This choice of L actually yields the *squared* Euclidean distance between two points $q(t_0)$ and $q(t_1)$. If we used the square root of this quantity in the Lagrangian, it becomes homogeneous of degree one in $(\dot{q}_1, \dot{q}_2)$ [as in $L(q_1, q_2, \lambda\dot{q}_1, \lambda\dot{q}_2, t) = \lambda L(q_1, q_2, \dot{q}_1, \dot{q}_2, t)$] and this creates problems for the Legendre transform. Note that the Lagrangian is independent of time t and this fact will later allow us to derive a static Schrödinger equation.

The Hamiltonian is obtained via a Legendre transform [6] applied to the Lagrangian:

$$H(q_1, q_2, p_1, p_2, t) = \sum_{i=1}^{2} p_i\dot{q}_i(p_1, p_2, t) - L(q_1, q_2, \dot{q}_1(p_1, p_2), \dot{q}_2(p_1, p_2), t) \tag{4}$$

where the momenta p_1, p_2 are defined as

$$p_i \equiv \frac{\partial L}{\partial \dot{q}_i},\ i = 1, 2. \tag{5}$$

Equation (5) can be inverted to obtain $\dot{q}_i = \dot{q}_i(p_1, p_2, t)$ [and this fails if the Lagrangian is homogeneous of degree one in $(\dot{q}_1, \dot{q}_2)$].

The Hamilton-Jacobi equation is obtained via a canonical transformation [7] of the Hamiltonian. In classical mechanics, a canonical transformation is defined as a change of variables which leaves the form of the Hamilonian unchanged. For a type 2 canonical transformation, we have

$$\sum_{i=1}^{2} p_i \dot{q}_i - H(q_1, q_2, p_1, p_2, t) = \sum_{i=1}^{2} P_i \dot{Q}_i - K(Q_1, Q_2, P_1, P_2, t) + \frac{dF}{dt} \tag{6}$$

where $F \equiv -\sum_{i=1}^{2} Q_i P_i + F_2(q, P, t)$ which gives

$$\frac{dF}{dt} = -\sum_{i=1}^{2} \left(\dot{Q}_i P_i + Q_i \dot{P}_i \right) + \frac{\partial F_2}{dt} + \sum_{i=1}^{2} \left(\frac{\partial F_2}{\partial q_i} \dot{q}_i + \frac{\partial F_2}{\partial P_i} \dot{P}_i \right). \tag{7}$$

When we pick a *particular* type 2 canonical transformation wherein $\dot{P}_i = 0,\ i = 1, 2$ *and* $K(Q_1, Q_2, P_1, P_2, t) = 0$, we get

$$\frac{\partial F_2}{\partial t} + H(q_1, q_2, \frac{\partial F_2}{\partial q_1}, \frac{\partial F_2}{\partial q_2}, t) = 0 \tag{8}$$

where we are forced to make the identification $p_i = \frac{\partial F_2}{\partial q_i},\ i = 1, 2$. Note that the new momenta P_i are constants of the motion (usually denoted by $\alpha_i,\ i = 1, 2$). Changing F_2 to S as in common practice, we have the standard Hamilton-Jacobi equation for the function $S(q_1, q_2, \alpha_1, \alpha_2, t)$. To complete the circle back to the Lagrangian, we take the total time derivative of the Hamilton-Jacobi function S to get

$$\begin{aligned} \frac{dS(q_1, q_2, \alpha_1, \alpha_2, t)}{dt} &= \sum_{i=1}^{2} \frac{\partial S}{\partial q_i} \dot{q}_i + \frac{\partial S}{\partial t} \\ &= \sum_{i=1}^{2} p_i \dot{q}_i - H(q_1, q_2, \frac{\partial S}{\partial q_1}, \frac{\partial S}{\partial q_2}, t) = L(q_1, q_2, \dot{q}_1, \dot{q}_2, t). \end{aligned} \tag{9}$$

Consequently $S(q_1, q_2, \alpha_1, \alpha_2, t) = \int_{t_0}^{t} L dt$ and the constants $\{\alpha_1, \alpha_2\}$ can now be interpreted as integration constants. For a more accessible treatment of the relationship between Lagrangians and the Hamilton-Jacobi field $S(q_1, q_2, \alpha_1, \alpha_2, t)$, please see [17].

For the Euclidean distance function problem, following (4) and (8), we get $H(q_1, q_2, p_1, p_2, t) = \frac{1}{2}\left(p_1^2 + p_2^2\right)$ and the Hamilton-Jacobi equation is

$$\frac{\partial S}{\partial t} + \frac{1}{2}\left[\left(\frac{\partial S}{\partial q_1} \right)^2 + \left(\frac{\partial S}{\partial q_2} \right)^2 \right] = 0. \tag{10}$$

The Schrödinger equation can sometimes be "derived" using a Feynman path integral approach. The more common approach—termed first quantization[1]—is to convert the relation $p_i = \frac{\partial S}{\partial q_i}$, $i = 1, 2$ into an operator relation $p_i = i\hbar\frac{\partial}{\partial q_i}$, $i = 1, 2$. In a similar fashion, the time operator is $i\hbar\frac{\partial}{\partial t}$. When we quantize the Euclidean distance function problem, we get

$$i\hbar\frac{\partial\psi}{\partial t} + \frac{\hbar^2}{2}\left(\frac{\partial^2\psi}{\partial x^2} + \frac{\partial^2\psi}{\partial y^2}\right) = 0. \tag{11}$$

At first glance, there appear to be some similarities between the Hamilton-Jacobi equation in (10) and the Schrödinger equation in (11). Due to first quantization, the squared first derivatives w.r.t. space in the former have morphed into second derivative operators in the latter. Both equations have first derivatives w.r.t. time.

We now show that the time independence of the Lagrangian in (2) allows us to simplify the former into the static Hamilton-Jacobi equation and the latter into the static Schrödinger equation.

If the Lagrangian is not an explicit function of time, we can seek solutions for the Hamilton-Jacobi equation that are time independent. Setting $S(q_1, q_2, \alpha_1, \alpha_2, t) = S^*(q_1, q_2, \alpha_1, \alpha_2) - Et$ where $E = \frac{1}{2}$ is the total energy for the Euclidean distance function problem, we get

$$\left(\frac{\partial S^*}{\partial q_1}\right)^2 + \left(\frac{\partial S^*}{\partial q_2}\right)^2 = 1 \tag{12}$$

which is the eikonal equation with the forcing term set to one—a nonlinear, first-order differential equation. In a similar fashion, when we set $\psi(x, t) = \phi(x)\exp\left(\frac{it}{2\hbar}\right)$ and use $E = \frac{1}{2}$, we see that $\phi(x)$ satisfies the screened Poisson equation

$$\hbar^2\left(\frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial y^2}\right) = \phi \tag{13}$$

which is a linear, second-order differential equation. A close relationship between ϕ and S^* can be shown by setting $\phi(x) = \exp\left\{-\frac{\hat{S}(x)}{\hbar}\right\}$ and rewriting (13) to get

$$\left(\frac{\partial\hat{S}}{\partial x_1}\right)^2 + \left(\frac{\partial\hat{S}}{\partial x_2}\right)^2 - \hbar\left(\frac{\partial^2\hat{S}}{\partial x_1^2} + \frac{\partial^2\hat{S}}{\partial x_2^2}\right) = 1 \tag{14}$$

which is strikingly similar to the eikonal equation in (12) with the important difference being a viscosity regularization term [18] modulated by the free parameter $\hbar$. [Note that the viscosity term arises naturally from (13)—an intriguing result.] As $\hbar \to 0$, $\hat{S} \to S^*$ which implies that we can solve the static Schrödinger equation in (13) instead of the eikonal equation in (12). In the next section, we describe fast algorithms for solving (14) and also present fast methods for computing the signed distance function and the derivatives of the Euclidean distance function.

[1] First quantization is still mysterious. For an informal but illuminating treatment, please see http://math.ucr.edu/home/baez/categories.html.

3 Fast computation of the signed Euclidean distance function and its derivatives

In our previous work [12], we showed that the static Schrödinger equation in (13) can be efficiently solved using a Fast Fourier Transform (FFT) approach in arbitrary dimensions. The complexity of the FFT is $O(N \log N)$ where N is the total number of grid points (in any dimension). We briefly summarize the FFT-based Euclidean distance function algorithm.

3.1 Unsigned Euclidean distance functions

In the Euclidean distance function problem, we begin by considering the forced version of (13) in 2D:

$$-\hbar^2 \bigtriangledown^2 \phi + \phi = \sum_{k=1}^{K} \delta(X - Y_k). \tag{15}$$

The points $Y_k, k \in \{1, \ldots, K\}$ are a set of seed locations at which $S^*(Y_k) = 0, \forall Y_k, k \in \{1, \ldots, K\}$ with the set X being the locations at which we wish to compute the Euclidean distance function. A Green's function approach can be pursued since the above differential equation is homogeneous except at the seed locations Y. The Green's functions [19] (for an unbounded domain with Dirichlet boundary conditions) are

$$G_{1D}(X - Y) = \frac{1}{2\hbar} \exp\left(\frac{-|X - Y|}{\hbar}\right), \tag{16}$$

$$G_{2D}(X - Y) = = \frac{1}{2\pi\hbar^2} K_0\left(\frac{\|X - Y\|}{\hbar}\right), \tag{17}$$

and

$$G_{3D}(X - Y) = \frac{1}{4\pi\hbar^2} \frac{\exp\left(\frac{-\|X-Y\|}{\hbar}\right)}{\|X - Y\|} \tag{18}$$

in 1D, 2D and 3D respectively where $K_0(r)$ is the modified Bessel function of the second kind. We avoid the singularity at the origin in 2D and 3D by replacing their Green's functions with the exponential function (similar to the 1D Green's function). This is a very good approximation as $\hbar \to 0$ since the 2D and 3D Green's functions converge *uniformly* to the exponential function everywhere away from the origin. With this in place, we write the solution for $\phi(X)$ as

$$\phi(X) = \sum_{k=1}^{K} G(X) * \delta(X - Y_k) = \sum_{k=1}^{K} G(X - Y_k) \tag{19}$$

and the corresponding approximate solution to the eikonal equation (after removing terms independent of X) is

$$\hat{S}(X) = -\hbar \log \sum_{k=1}^{K} \exp\left\{-\frac{\|X - Y_k\|}{\hbar}\right\} \tag{20}$$

with the caveat being that we are using an approximate, unbounded domain Green's function $G(X)$ here. We have shown that an approximate solution for the eikonal (with the forcing term set to one) can be obtained in closed-form as in (20) and efficiently computed using an FFT since equation (19) expresses a *discrete* convolution [13] between the functions

$$G(X) = \exp\left\{-\frac{\|X\|}{\hbar}\right\} \tag{21}$$

and

$$Y_{\text{kron}}(X) \equiv \sum_{k=1}^{K} \delta_{\text{kron}}(X - Y_k). \tag{22}$$

(Here $\delta_{\text{kron}}(X)$ is a Kronecker delta function.) This is a significant result since the time complexity of the discrete convolution is $O(N \log N)$ and the expression $\hat{S}(X)$ in (20) for the Euclidean distance function is continuous and differentiable everywhere (except in 1D at the seed locations).

3.2 Winding numbers for the signed distance function in 2D

The solution for the approximate Euclidean distance function in (20) is lacking in one respect: there is no information on the sign of the distance. This is to be expected since the distance function was obtained only from a set of *points* Y and not a curve or surface. We now describe a new method for computing the signed distance in 2D using winding numbers [20]. (The equivalent concept in 3D and higher dimensions is the topological degree which appears to be a straightforward extension but with possible unexpected pitfalls.)

Assume that we have a closed, parametric curve $\left\{x^{(1)}(t), x^{(2)}(t)\right\}$, $t \in [0, 1]$. We seek to determine if a grid location in the set $\left\{X_i \in \mathbb{R}^2,\, i \in \{1, \ldots, N\}\right\}$ is inside the closed curve. The winding number is the number of times the curve winds around the point X_i (if at all) and if the curve is oriented, counterclockwise turns are counted as positive and clockwise turns as negative. If a point is inside the curve, the winding number is a non-zero integer. If the point is outside the curve, the winding number is zero. If we can efficiently compute the winding number for all points on a grid w.r.t. to a curve, then we would have the sign information (inside/outside) for all the points. We now describe a fast algorithm to achieve this goal.

If the curve is C^1, then the angle $\theta(t)$ of the curve is continuous and differentiable and $d\theta(t) = \left(\frac{x^{(1)}\dot{x}^{(2)} - x^{(2)}\dot{x}^{(1)}}{r^2}\right) dt$ where $r(t) = \sqrt{\left[x^{(1)}\right]^2 + \left[x^{(2)}\right]^2}$. Since we need to determine whether the curve winds around each of the points $X_i, i \in \{1, \ldots, N\}$, define $(\hat{x}_i^{(1)}, \hat{x}_i^{(2)}) \equiv (x^{(1)} - X_i^{(1)}, x^{(2)} - X_i^{(2)})$, $\forall i$. Then the winding numbers for all grid points in the set X are

$$\mu_i = \frac{1}{2\pi} \oint_C \left(\frac{\hat{x}_i^{(1)} \dot{\hat{x}}_i^{(2)} - \hat{x}_i^{(2)} \dot{\hat{x}}_i^{(2)}}{\left[\hat{x}_i^{(1)}\right]^2 + \left[\hat{x}_i^{(2)}\right]^2} \right) dt,\, \forall i \in \{1, \ldots, N\}. \tag{23}$$

As it stands, we cannot actually compute the winding numbers without performing the integral in (23). To this end, we discretize the curve and produce a sequence of points $\left\{Y_k \in \mathbb{R}^2, k \in \{1,\ldots,K\}\right\}$ with the understanding that the curve is closed and therefore the "next" point after Y_K is Y_1. (The winding number property holds for piecewise continuous curves as well.) The integral in (23) becomes a discrete summation and we get

$$\mu_i = \frac{1}{2\pi}\sum_{k=1}^{K}\left(\frac{\left[Y_k^{(1)} - X_i^{(1)}\right]\left[Y_{k\oplus 1}^{(2)} - Y_k^{(2)}\right] - \left[Y_k^{(2)} - X_i^{(2)}\right]\left[Y_{k\oplus 1}^{(1)} - Y_k^{(1)}\right]}{\left[\left(Y_k^{(1)} - X_i^{(1)}\right)^2 + \left(Y_k^{(2)} - X_i^{(2)}\right)^2\right]}\right) \tag{24}$$

$\forall i \in \{1,\ldots,N\}$, where the notation $Y_{k\oplus 1}^{(\cdot)}$ denotes that $Y_{k\oplus 1}^{(\cdot)} = Y_{k+1}^{(\cdot)}$ for $k \in \{1,\ldots,K-1\}$ and $Y_{K\oplus 1}^{(\cdot)} = Y_1^{(\cdot)}$. We can simplify the notation in (24) (and obtain a measure of conceptual clarity as well) by defining the "tangent" vector $\{Z_k, k = \{1,\ldots,K\}\}$ as $Z_k^{(\cdot)} = Y_{k\oplus 1}^{(\cdot)} - Y_k^{(\cdot)}, k \in \{1,\ldots,K\}$ with the $(\cdot)$ symbol indicating either coordinate. Using the tangent vector Z, we rewrite (24) as

$$\mu_i = \frac{1}{2\pi}\sum_{k=1}^{K}\left(\frac{\left[Y_k^{(1)} - X_i^{(1)}\right]Z_k^{(2)} - \left[Y_k^{(2)} - X_i^{(2)}\right]Z_k^{(1)}}{\left[\left(Y_k^{(1)} - X_i^{(1)}\right)^2 + \left(Y_k^{(2)} - X_i^{(2)}\right)^2\right]}\right), \forall i \in \{1,\ldots,N\} \tag{25}$$

We now make the somewhat surprising observation (to us at any rate) that μ in (25) is a sum of two discrete convolutions. The first convolution is between two functions $f_{cr}(X) \equiv f_c(X)f_r(X)$ and $g_2(X) = \sum_{k=1}^{K} Z_k^{(2)}\delta_{\text{kron}}(X - Y_k)$ where the Kronecker delta function is a product of two Kronecker delta functions, one for each coordinate. The second convolution is between two functions $f_{sr}(X) \equiv f_s(X)f_r(X)$ and $g_1(X) \equiv \sum_{k=1}^{K} Z_k^{(1)}\delta_{\text{kron}}(X - Y_k)$. The functions $f_c(X)$, $f_s(X)$ and $f_r(X)$ are defined as

$$f_c(X) \equiv \frac{X^{(1)}}{\sqrt{\left[X^{(1)}\right]^2 + \left[X^{(2)}\right]^2}},\ f_s(X) \equiv \frac{X^{(2)}}{\sqrt{\left[X^{(1)}\right]^2 + \left[X^{(2)}\right]^2}}, \text{ and} \tag{26}$$

$$f_r(X) \equiv \frac{1}{\sqrt{\left[X^{(1)}\right]^2 + \left[X^{(2)}\right]^2}}. \tag{27}$$

where we have abused notation somewhat and let $X^{(1)}$ ($X^{(2)}$) denote the x (y)-coordinate of all the points in the grid set X. Armed with these relationships, we rewrite (25) to get

$$\mu(X) = \frac{1}{2\pi}\left[-f_{cr}(X) * g_2(X) + f_{sr}(X) * g_1(X)\right] \tag{28}$$

which can be computed in $O(N \log N)$ time using two FFTs. We have shown that the sign component of the Euclidean distance function can be separately

computed (without knowledge of the distance) in parallel in $O(N \log N)$ on a regular 2D grid.

3.3 Fast computation of the derivatives of the distance function

Just as the approximate Euclidean distance function $\hat{S}(X)$ can be efficiently computed in $O(N \log N)$, so can the derivatives. This is important because fast computation of the derivatives of $\hat{S}(X)$ on a regular grid can be very useful in medial axes and curvature computations. Below, we detail how this can be achieved. We begin with the gradients and for illustration purposes, the derivations are performed in 2D:

$$\hat{S}_x(X) = \frac{\sum_{k=1}^{K} \frac{\left(X^{(1)} - Y_k^{(1)}\right)}{\sqrt{\left(X^{(1)} - Y_k^{(1)}\right)^2 + \left(X^{(2)} - Y_k^{(2)}\right)^2}} \exp\left\{-\frac{\|X - Y_k\|}{\hbar}\right\}}{\sum_{k=1}^{K} \exp\left\{-\frac{\|X - Y_k\|}{\hbar}\right\}}. \tag{29}$$

A similar expression can be obtained for $\hat{S}_y(X)$. These first derivatives can be rewritten as discrete convolutions:

$$\hat{S}_x(X) = \frac{f_c(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)},\ \hat{S}_y(X) = \frac{f_s(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)}, \tag{30}$$

where $f_c(X)$ and $f_s(X)$ are as defined in (26) and $Y_{\text{kron}}(X)$ is as defined in (22).

The second derivative formulae are somewhat involved. Rather than hammer out the algebra in a turgid manner, we merely present the final expressions—all discrete convolutions—for the three second derivatives in 2D:

$$\begin{aligned}\hat{S}_{xx}(X) = &-(1 + \frac{1}{\hbar})\frac{f_c^2(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)} + \frac{1}{\hbar}\hat{S}_x^2(X) \\ &+\frac{f_r(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)},\end{aligned} \tag{31}$$

$$\begin{aligned}\hat{S}_{yy}(X) = &-(1 + \frac{1}{\hbar})\frac{f_s^2(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)} + \frac{1}{\hbar}\hat{S}_y^2(X) \\ &+\frac{f_r(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)}, \text{ and}\end{aligned} \tag{32}$$

$$\hat{S}_{xy}(X) = -(1 + \frac{1}{\hbar})\frac{f_c(X) f_s(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)} + \frac{1}{\hbar}\hat{S}_x(X)\hat{S}_y(X) \tag{33}$$

where $f_r(X)$ is as defined in (27). We also see that

$$\hat{S}_x^2(X) + \hat{S}_y^2(X) - \hbar\left[\hat{S}_{xx}(X) + \hat{S}_{yy}(X)\right] = (1+\hbar) - 2\hbar\frac{f_r(X) \exp\left\{-\frac{X}{\hbar}\right\} * Y_{\text{kron}}(X)}{\hat{S}(X)} \tag{34}$$

[since $f_c^2(X) + f_s^2(X) = 1$] with the right side going to one as $\hbar \to 0$ for points in X away from points in the seed point-set Y. This is in accordance with (14) and vindicates our choice of the replacement Green's function in (21).

Since we can efficiently compute the first and second derivatives of the approximate Euclidean distance function $\hat{S}(X)$ everywhere on a regular grid, we can also compute derived quantities such as curvature (Gaussian, mean and principal curvatures for a two-dimensional surface). In the next section, we visualize the derivatives and maximum curvature for shape silhouettes.

4 Euclidean distance function experiments

We executed the Schrödinger Euclidean distance function algorithm on a set of 2D shape silhouettes[2]. The grid size is $-20 \leq x \leq 20$ and $-20 \leq y \leq 20$ with a grid spacing of 0.25 and $\hbar = 0.3$. The winding number discrete convolution algorithm is used to mark points as either inside or outside each shape. We visualize the vector fields $(\hat{S}_x, \hat{S}_y)$ in Figure 1 for the 8 shapes and the maximum curvature for a subset of the shapes in Figure 2. We chose the maximum curvature (defined as $H + \sqrt{H^2 - K}$ where H and K are the mean and Gaussian curvatures respectively of the Monge patch given by $\left\{x, y, \hat{S}(x, y)\right\}$) as the vehicle to visualize the medial axes of each shape after first considering the divergence of the unit gradient $\left[\nabla \cdot g = \nabla \cdot \left(\frac{\nabla \hat{S}(X)}{|\nabla \hat{S}(X)|}\right)\right]$ and the entropy $\left(-\frac{\partial \hat{S}(X)}{\partial \hbar}\right)$. The divergence is a good choice for the medial axes provided we update an adaptive grid whereas the entropy requires very high precision numerical computation (which we plan to pursue in the future).

Next, we ran a comparison of the Schrödinger Euclidean distance function algorithm with the fast sweeping method [11] and the exact Euclidean distance. We used a "Dragon" point-set obtained from the Stanford 3D Scanning Repository[3] in 3D and executed the three approaches to construct isosurfaces which are visualized in Figure 3. The common grid was $-2 \leq x \leq 2$, $-2 \leq y \leq 2$ and $-2 \leq z \leq 2$ with a grid spacing of 0.125. Numerical underflow errors in the FFT forced us to run the Schrödinger Euclidean distance function algorithm at four values of $\hbar$, namely, 0.025, 0.045, 0.06, and 0.08. We used the following decision criterion for $\hat{S}(X)$: $\hat{S} = \hat{S}|_{\hbar=0.08}$ if $\hat{S} \geq 2$, $\hat{S} = \hat{S}|_{\hbar=0.06}$ if $1.5 \leq \hat{S} < 2$, $\hat{S} = \hat{S}|_{\hbar=0.045}$ if $0.75 \leq \hat{S} < 1.5$ and $\hat{S} = \hat{S}|_{\hbar=0.025}$ if $\hat{S} < 0.75$. The initial conditions $\hat{S}(Y_k) = 0$, $\forall k \in \{1, \ldots, K\}$ were used to translate the $\hat{S}$ values (upwards or downwards) such that the minimum value was zero. The average percentage error in the Schrödinger approach was 3.89% whereas the average percentage error in the fast sweeping method (where the Gauss-Seidel iterates were run until convergence) was 6.35%. Our FFT-based approach does not begin by discretizing the spatial differential operator as is the case with the fast marching and fast sweeping methods and this could help account for the increased accuracy.

[2] We thank Kaleem Siddiqi for providing us the set of 2D shape silhouettes used in this paper.

[3] This dataset is available at http://graphics.stanford.edu/data/3Dscanrep/.

5 A perturbation approach for the general eikonal problem

In this section, we briefly summarize our perturbation approach (using the well known Born expansion) [15,12] for the general eikonal equation (with forcing functions $f(X)$ bounded away from zero). We consider the static Schrödinger equation (in 2D) with a forcing function $f(X)$:

$$-\hbar^2 \bigtriangledown^2 \phi + f^2 \phi = \sum_{k=1}^{K} \delta(X - Y_k). \tag{35}$$

Equation (35) can be rewritten as

$$(-\hbar^2 \bigtriangledown^2 + \tilde{f}^2) \left[1 + (-\hbar^2 \bigtriangledown^2 + \tilde{f}^2)^{-1} \circ (f^2 - \tilde{f}^2)\right] \phi = \sum_{k=1}^{K} \delta(X - Y_k) \tag{36}$$

with $\tilde{f}(X)$ a constant forcing function. Now, defining the operator A as $A \equiv (-\hbar^2 \bigtriangledown^2 + \tilde{f}^2)^{-1} \circ (f^2 - \tilde{f}^2)$ and ϕ_0 as $\phi_0 \equiv (1 + A)\phi$, we see that ϕ_0 satisfies

$$(-\hbar^2 \bigtriangledown^2 + \tilde{f}^2)\phi_0 = \sum_{k=1}^{K} \delta(X - Y_k) \tag{37}$$

and

$$\phi = (1 + A)^{-1} \phi_0. \tag{38}$$

Using a geometric series approximation for $(1 + A)^{-1}$, we obtain the solution for ϕ as

$$\phi \approx \phi_0 - \phi_1 + \phi_2 - \phi_3 + \ldots + (-1)^T \phi_T \tag{39}$$

where ϕ_i satisfies the recurrence relation

$$(-\hbar^2 \bigtriangledown^2 + \tilde{f}^2)\phi_i = (f^2 - \tilde{f}^2)\phi_{i-1}, \; \forall i \in \{1, 2, \ldots, T\}. \tag{40}$$

The solutions for ϕ_i can then be obtained by convolution

$$\phi_0(X) = \sum_{k=1}^{K} G(X) * \delta(X - Y_k) = \sum_{k=1}^{K} G(X - Y_k), \tag{41}$$

$$\phi_i(X) = G(X) * \left[(f^2 - \tilde{f}^2)\phi_{i-1}\right], \; \forall i \in \{1, 2, \ldots, T\} \tag{42}$$

and an approximate solution to the eikonal equation can be obtained from $\hat{S}(X) = -\hbar \log \phi(X)$. The discrete convolutions in (41) and in (42) can be efficiently computed via FFTs. The number of terms (T) used in the geometric series approximation of $(1 + A)^{-1}$ is *independent* of the grid size N. We set $\tilde{f} = \sqrt{\frac{[\min_X f(X)]^2 + [\max_X f(X)]^2}{2}}$ which turns out to be the optimal value [12].

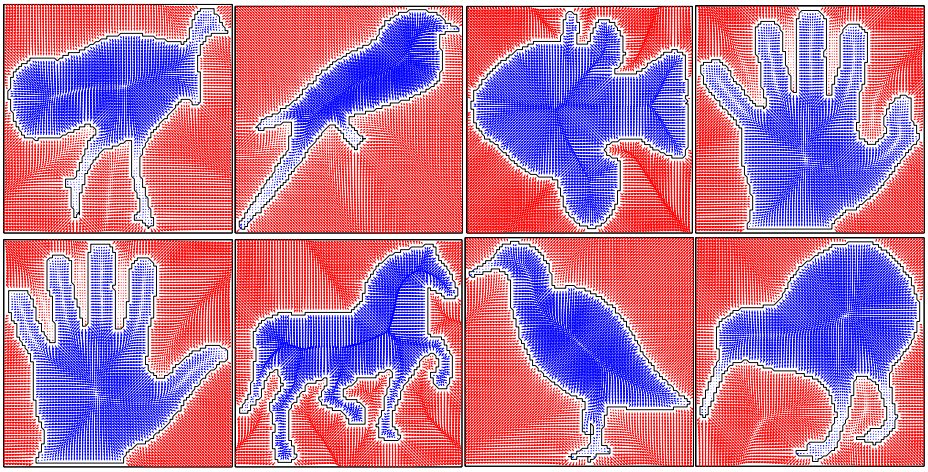

Fig. 1. A quiver plot of $\nabla \hat{S} = (\hat{S}_x, \hat{S}_y)$ for a set of silhouette shapes

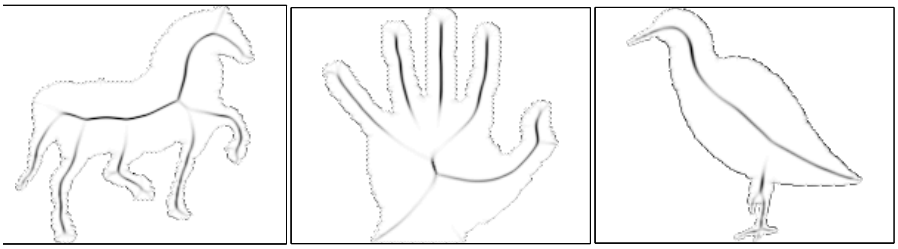

Fig. 2. Maximum curvature plots: i) Horse, ii) Hand, and iii) Bird

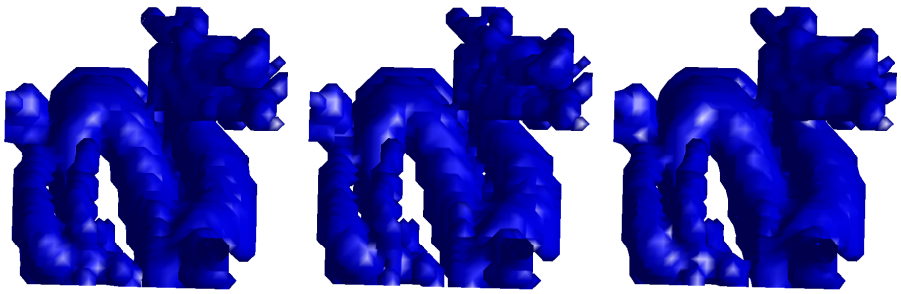

Fig. 3. Dragon surface reconstructed using i) Schrödinger, ii) Exact Euclidean distance and iii) Fast sweeping

6 Image segmentation results using the eikonal solver

To test the eikonal solver, we obtained two images (3096 and 101085) from the Berkeley Segmentation Dataset and Benchmark[4]. After first smoothing them using a 5×5 Gaussian with standard deviation 0.25, we used the following sets of parameters for the two images. For image 3096: we ran the perturbation method for $T = 5$ iterations at $\hbar$ values of 0.05, 0.15, 0.25 and 0.35 with thresholds of 1.5, 2.5 and 4, grid size $-8 \leq x \leq 8$, $-6 \leq y \leq 6$ with grid spacing 0.0313 and the forcing function $f(X) = \frac{|\nabla I(X)|}{\max_X |\nabla I(X)|} + 0.5$. For image 101085, the only changes were: grid size $-5 \leq x \leq 5$, $-8 \leq y \leq 8$ with grid spacing 0.0313 and the forcing function $f(X) = \frac{|\nabla I(X)|}{\max_X |\nabla I(X)|} + 0.01$. Both images were seeded with a set of interior/exterior points, the eikonal algorithms were run twice and we displayed the boundaries in Figure 4 using the eikonal "winner"—the one with the smaller distance. The same approach was used for the fast sweeping method and Dijkstra's shortest path algorithm [16] (since closed form solutions are not available). The results are obviously anecdotal but serve to illustrate the correspondence between the Schrödinger (quantum) and the fast sweeping (classical) approaches. Note the larger scale boundaries in the Schrödinger segmentation of image 101085 which we attribute to the viscosity term in (14).

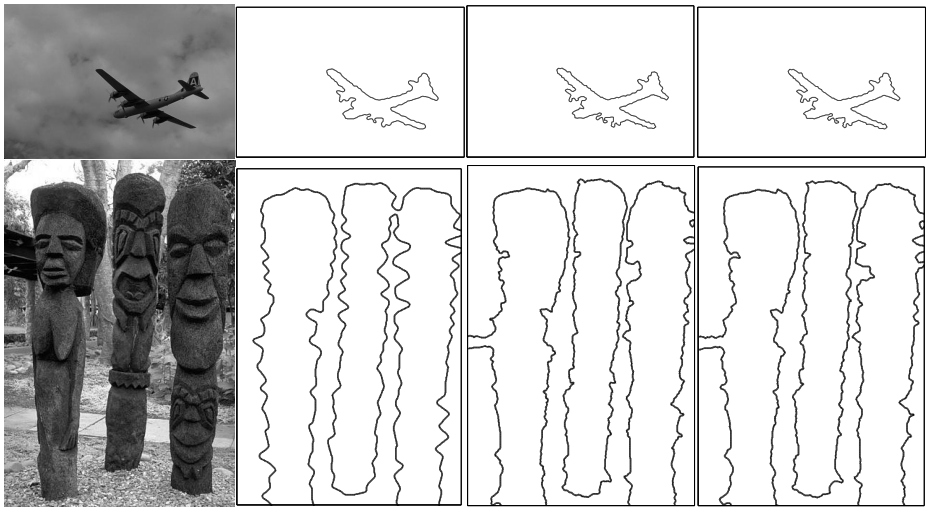

Fig. 4. Image segmentation results. Top from left to right: i) Image 3096, ii) Schrödinger, iii) Dijkstra, iv) Fast sweeping. Bottom from left to right: i) Image 101085, ii) Schrödinger, iii) Dijkstra, iv) Fast sweeping.

[4] This dataset is available at
http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/

7 Discussion

While energy minimization methods have permeated image analysis in the past two decades, one overarching generalization we can make is that the formulations are inspired by classical and not quantum mechanics. Despite the fact that a close correspondence exists between Hamilton-Jacobi theory and Schrödinger wave mechanics, we have not seen image analysis leverage this relationship. When we solve the Schrödinger equation at small values of $\hbar$, we obtain closed-form solutions for the Euclidean distance function problem that can be efficiently computed in $O(N \log N)$. This has immediate applications in image analysis as the sign of the distance function, its gradients and curvature can all be written in closed-form and efficiently computed via FFTs. When applied to shape silhouettes, the gradients and curvature information can aid in medical axes computation. We also show that a perturbation series approach leads to a fast eikonal solver which can be used in image segmentation. Future work will focus on fast, multipole methods [21] for solving the time independent Schrödinger equation since these methods have very attractive computational properties.

Acknowledgments

We thank Arunava Banerjee, Rama Chellappa, Alireza Entezari, Chiu-Yen Kao, Ajit Rajwade, Kaleem Siddiqi, Baba Vemuri and Alan Yuille for helpful discussions.

References

1. Horn, B.: Robot Vision. The MIT Press (1986)
2. Siddiqi, K., Tannenbaum, A., Zucker, S.: A Hamiltonian approach to the eikonal equation. In: EMMCVPR. Volume LNCS 1654., Springer-Verlag (1999) 1–13
3. Osher, S., Fedkiw, R.: Level Set Methods and Dynamic Implicit Surfaces. Springer (2002)
4. Gilboa, G., Sochen, N., Zeevi, Y.: Image enhancement and denoising by complex diffusion processes. IEEE PAMI **26**(8) (2004) 1020–1036
5. Emms, D., Severini, S., Wilson, R., Hancock, E.: Coined quantum walks lift the cospectrality of graphs and trees. In: EMMCVPR. Volume LNCS 3757., Springer (2005) 332–345
6. Arnold, V.: Mathematical Methods of Classical Mechanics. Springer (1989)
7. Goldstein, H., Poole, C., Safko, J.: Classical Mechanics. Addison-Wesley (2001)
8. Griffiths, D.: Introduction to Quantum Mechanics. Benjamin-Cummings (2004)
9. Sethian, J.: A fast marching level set method for monotonically advancing fronts. In: Proc. Nat. Acad. Sci. (1996) 1591–1595
10. Osher, S., Sethian, J.: Fronts propagating with curvature dependent speed: algorithms based on Hamilton-Jacobi formulations. Journal of Computational Physics **79**(1) (1988) 12–49
11. Zhao, H.: A fast sweeping method for eikonal equations. Mathematics of Computation (2005) 603–627

12. Gurumoorthy, K., Rangarajan, A.: A fast eikonal equation solver using the Schrödinger wave equation. Technical Report CVGMI-09-05, Center for Computer Vision, Graphics and Medical Imaging (CVGMI), University of Florida (2009)
13. Bracewell, R.: The Fourier Transform and its Applications. 3rd edn. McGraw-Hill Science and Engineering (1999)
14. Yatziv, L., Bartesaghi, A., Sapiro, G.: O(N) implementation of the fast marching algorithm. Journal of Computational Physics **212**(2) (2006) 393–399
15. Fernandez, F.: Introduction to Perturbation Theory in Quantum Mechanics. CRC Press (2000)
16. Cormen, T., Leiserson, C., Rivest, R., Stein, C.: Introduction to Algorithms. 2nd edn. The MIT Press (September 2001)
17. Butterfield, J.: On Hamilton-Jacobi theory as a classical root of quantum theory. In A. Elitzur, S. Dolev, N. Kolenda, eds.: Quo-Vadis Quantum Mechanics. Springer (2005) 239–274
18. Crandall, M., Ishii, H., Lions, P.: User's guide to viscosity solutions of second order partial differential equations. Bulletin of the AMS **27**(1) (1992) 1–67
19. Abramowitz, M., Stegun, I.: Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. Government Printing Office, USA (1964)
20. Aberth, O.: Precise Numerical Methods Using C++. Academic Press (1998)
21. Gumerov, N., Duraiswami, R.: Fast multipole methods for the Helmholtz equaton in three dimensions. Elseveier series in electromagnetism. Elsevier (2004)