

A new energy minimization framework and sparse linear system for path planning and shape from shading

Adrian M. Peter
Dept. of Eng. Systems
Florida Institute of Technology
Melbourne, FL USA
apeter@fit.edu

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

Mark Moyou
Dept. of Eng. Systems
Florida Institute of Technology
Melbourne, FL USA
mmoyou@my.fit.edu

Anand Rangarajan
Dept. of CISE
University of Florida
Gainesville, FL, USA
anand@cise.ufl.edu

ABSTRACT

For over 30 years, the static Hamilton-Jacobi (HJ) equation, specifically its incarnation as the eikonal equation, has been a bedrock for a plethora of computer vision models, including popular applications such as shape-from-shading, medial axis representations, level-set segmentation, and geodesic processing (i.e. path planning). Numerical solutions to this nonlinear partial differential equation have long relied on staples like fast marching and fast sweeping algorithms—approaches which rely on intricate convergence analysis, approximations, and specialized implementations. Here, we present a new variational functional on a scalar field comprising a spatially varying quadratic term and a standard regularization term. The Euler-Lagrange equation corresponding to the new functional is a linear differential equation which when discretized results in a linear system of equations. This approach leads to many algorithm choices since there are myriad efficient sparse linear solvers. The limiting behavior, for a particular case, of this linear differential equation can be shown to converge to the nonlinear eikonal. In addition, our approach eliminates the need to explicitly construct viscosity solutions as customary with direct solutions to the eikonal. Though our solution framework is applicable to the general class of eikonal problems, we detail specifics for the popular vision applications of shape-from-shading, vessel segmentation, and path planning. We showcase experimental results on a variety of images and complex mazes, in which we hold our own against state-of-the-art fast marching and fast sweeping techniques, while retaining the considerable advantages of a linear systems approach.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ICVGIP '14, December 14-18, 2014, Bangalore, India
Copyright 2014 ACM 978-1-4503-3061-9/14/12 ...\$15.00
<http://dx.doi.org/10.1145/2683483.2683498>

Keywords

Path planning, Quantum, Schrödinger equation, Eikonal equation, Hamilton-Jacobi equation, Screened Poisson equation, Sparse linear systems

1. INTRODUCTION

Variational methods have been a mainstay of computer vision since its inception. Surface reconstruction, optical flow, edge detection and shape from shading (SFS) were given variational treatments prior to the introduction of other frameworks such as Gibbs-Markov random fields and combinatorial optimization. In most of the above applications, variational methods have begun with a functional defined on the field of interest (scalar fields in the case of surface reconstruction and SFS and vector fields in the case of optical flow) followed by the familiar trope of the Euler-Lagrange differential equation and corresponding discretized system of equations (on regular or irregular grids as the case may be). Regardless of the application, variational methods tended to incorporate both data and regularization costs with the terms ranging from straightforward quadratic energies to more complex formulations. In the present work, we bring a fresh perspective on the variational origins of the eikonal equation. We show how to circumvent directly working with the nonlinear eikonal, instead solving the resulting linear Euler-Lagrange equation from our functional. We exhibit how adoption of this framework will lead to very efficient numerical algorithms, allowing us to simply employ sparse linear solvers. In addition, in our framework, *viscosity solutions* [5, 1] are generated as a natural by-product of the linear solution, which is in sharp contrast to contemporary eikonal solvers that superfluously append a viscosity term. Our variational formulation also opens the door for new and innovative ways to incorporate problem dependent constraints for eikonal-related applications.

Level-set methods have seen wide application in computer vision over the past twenty years. Level-sets and related Hamilton-Jacobi (usually eikonal) solvers have permeated segmentation, path planning, mesh processing and active contours. While the Hamilton-Jacobi differential equation

is typically derived in physics via a canonical transformation of a variational Lagrangian, in the above applications, Hamilton-Jacobi and eikonal solvers have been directly used without taking recourse to the physics route. In the process of adapting to an application domain such as vision, the origins of the eikonal equation in optics and light propagation are de-emphasized with the algorithmic aspect showcased instead. Today, fast marching and fast sweeping eikonal solvers have become the staple of level-set methods in general and vision applications in particular.

Given this history of variational and level-set methods in computer vision, the bar has been raised considerably for the introduction of new frameworks and algorithms. The motivation and need for new methodologies must be first addressed, followed (hopefully) by a payoff via new algorithms. We motivate the new energy minimization framework by starting with a simple, variational problem, widely known in computer vision. Given a set of discrete locations $x_k, k \in \{1, \dots, K\}$, we create a scalar field $\phi_0(x)$ which is highly peaked at the given set of locations (typically referred to as source points) and is close to zero elsewhere. (The field $\phi_0(x)$ can be thought of roughly as a set of delta functions with the twist of being square integrable.) Now, consider the reconstruction of a new field $\phi(x)$ which is close to $\phi_0(x)$ in the ℓ_2 sense while also being smooth in terms of an appropriate Sobolev norm. Perhaps the simplest energy functional (variational principle) is

$$I[\phi] = \int_{\Omega} |\phi(x) - \phi_0(x)|^2 dx + \lambda^2 \int_{\Omega} \|\nabla\phi(x)\|^2 dx \quad (1)$$

where Ω is the domain and λ a regularization parameter. Assuming certain technical conditions hold, the solution to the above energy function can be written in terms of the sum of isotropic Green's functions at the source points. At small values of λ , the solution $\phi(x) \approx \gamma \sum_{k=1}^K G(\frac{\|x-x_k\|}{\lambda})$, with γ being a normalization constant, rapidly *and isotropically* decays to zero away from the source points. Isotropic decay implies that the field $\phi(x)$ carries information related to the distance from a source. In the limit as $\lambda \rightarrow 0$, at a point $x \in \Omega$, $\phi(x)$ carries information regarding the distance to the closest source location. This observation permits us to assign $\phi(x)$ a *distance function*-like role.

One of the hallmarks of the eikonal equation $\|\nabla S(x)\| = c(x)$ is its versatility. The interesting applications of the eikonal—shape from shading, path planning etc.—correspond to different choices of the forcing function. When $\|\nabla S\| = 1$, the scalar field $S(x)$ is the actual distance function (given a set of source locations at which $S(x)$ is zero). As we later show, the variational principle in (1) results in a scalar field which is similar to $S(x)$. Given this relationship of the eikonal to the energy function in (1) for the setting of the forcing function to one, is there an energy functional incorporating $c(x)$ with properties similar to the eikonal?

In this work, we design a new energy functional—similar to (1) above—whose solution has many of the properties of the eikonal. The new variational form, with a data term modulated by $c(x)$, results in Euler-Lagrange differential equation which is *linear* in $\phi(x)$. Due to the presence of $c(x)$ though, we can no longer use a Green's function approach to solve the resulting linear differential equation—an *inhomogeneous, screened Poisson* equation. Instead, we directly discretize the screened Poisson on a standard grid and use *sparse, linear system* solvers to obtain the solution. Despite

the eikonal being a nonlinear differential equation, and our screened Poisson a linear one, we show a very close relationship between the two approaches and demonstrate the effect of the free parameter λ in obtaining viscosity-like solutions. We showcase the application of the new energy minimization framework and sparse linear solver in path planning, shape from shading and vessel segmentation.

2. A NEW VARIATIONAL PRINCIPLE AKIN TO THE EIKONAL

We return to the “distance function” variational principle in (1). Given a set of discrete locations $x_k, k \in \{1, \dots, K\}$ and a scalar field $\phi_0(x)$ highly peaked at the source locations, the desired scalar field $\phi(x)$ can be recovered by solving the Euler-Lagrange equation corresponding to the variational problem

$$-\lambda^2 \nabla^2 \phi + \phi(x) = \phi_0(x), \quad (2)$$

a linear (*inhomogeneous, screened Poisson*) differential equation. Recall that $\phi_0(x)$ is somewhat similar to a set of delta functions located at the source points [with the difference being the square integrability of $\phi_0(x)$]. To tease out the similarity to distance functions, consider the following pair of transformations: $\phi(x) = \exp\left\{-\frac{S(x)}{\lambda}\right\}$ and $\phi_0(x) = \exp\left\{-\frac{S_0(x)}{\lambda}\right\}$ with $S(x)$ and $S_0(x)$ arbitrary. (We have abused notation somewhat by using $S(x)$ in the transformation in order to highlight the close relationship of $\phi(x)$ to distance transforms.) While these transformations require mathematical justification—especially so since they imply the non-negativity of $\phi(x)$ and $\phi_0(x)$ —we bypass these issues for now. Taking appropriate derivatives and substituting the results in (2), we get

$$\|\nabla S\|^2 - \lambda \nabla^2 S = 1 - \exp\left\{-\frac{S_0(x) - S(x)}{\lambda}\right\}. \quad (3)$$

Assuming bounded derivatives and provided $S(x) < S_0(x)$ almost everywhere, the above nonlinear differential equation can be approximated at small values of λ as

$$\|\nabla S(x)\| \approx 1, \quad (4)$$

the constant forcing function version of the eikonal equation. The solution $S(x)$ to eq. (4) is a distance transform over the set of discrete source locations and satisfies

$$S(x) = \min_k \|x - x_k\| \quad (5)$$

over the domain Ω .

Since the above derivation has been heuristic (in the extreme), we now solve the differential equation in (1) assuming the existence of a Green's function $G(\frac{\|x\|}{\lambda})$ defined on the domain Ω . The solution for an arbitrary $\phi_0(x)$ is

$$\phi(x) = \int_{\Omega} G\left(\frac{\|x-y\|}{\lambda}\right) \phi_0(y) dy. \quad (6)$$

As $\phi_0(x)$ is closely related to a set of delta functions centered at the source locations (with the important difference of being square integrable), we may approximate $\phi(x)$ as

$$\phi(x) \approx \gamma \sum_{k=1}^K G\left(\frac{\|x-x_k\|}{\lambda}\right) \quad (7)$$

where γ is a constant related to the square integrability of $\phi_0(x)$. For positive valued Green's functions, we see that $\phi(x) > 0$ everywhere satisfying the assumption above that $\phi(x)$ can be written as $\exp\left\{-\frac{S(x)}{\lambda}\right\}$. Given this transformation, the approximate distance function can be expressed as

$$S(x) \approx -\lambda \log \left[\gamma \sum_{k=1}^K G\left(\frac{\|x - x_k\|}{\lambda}\right) \right]. \quad (8)$$

For rapid and isotropic decay (at small values of λ) of the Green's function away from the source points, $S(x)$ can be shown to be a good approximation to the true (unsigned) distance function with the approximation becoming increasingly accurate as $\lambda \rightarrow 0$. Green's function expressions vary with dimension requiring separate analyses on a case by case basis. For dimensions 1 to 3—the cases usually of interest—the analysis goes through and the above observations hold [16].

Can the above distance function variational principle be extended to the general eikonal? More precisely, given the eikonal equation

$$\|\nabla S\| = c(x), \quad (9)$$

where $c(x)$ is the forcing function, does there exist an energy functional (variational principle) similar to (1) above. We take a closer look at the physical meaning of the eikonal to aid us in reverse engineering a new variational principle.

In the eikonal, $S(x)$ is related to the time taken to reach the location x . The forcing function is proportional to the time spent at that location. The source locations represent time at the origin with $S(x)$ equal to zero. Then at a given location x , $S(x)$ is the total time taken to reach that location with large values of the forcing function $c(x)$ acting as impediments, thereby increasing the time taken and vice versa. In our setup, $\phi(x)$ is inversely proportional to $S(x)$. Therefore, large values of $S(x)$ (time taken) correspond to small values of $\phi(x)$ and vice versa. The input $\phi_0(x)$ can be taken to be peaked at the source locations and close to zero in regions known to be inaccessible [with large values of $S(x)$]. Consequently, we see an inverse relationship between $c(x)$ and $\phi_0(x)$ as well. From the above, we see that $c(x)\phi(x)$ should be small, except at the source locations where $\phi(x)$ should be close to $\phi_0(x)$.

Given these observations, consider the following variational principle—a potential candidate to be akin to the eikonal:

$$I[\phi] = \int_{\Omega} c^2(x) |\phi(x) - \phi_0(x)|^2 dx + \lambda^2 \int_{\Omega} \|\nabla \phi(x)\|^2 dx. \quad (10)$$

The principal difference between (10) and (1) is in the data term. The forcing function $c(x)$ now modulates the penalty on the deviation of $\phi(x)$ from $\phi_0(x)$. In particular, note that large values of $c(x)$ force a closer adherence of $\phi(x)$ to $\phi_0(x)$ and this tracks our earlier discussion focused on locations where $c(x)$ is large and $\phi_0(x)$ close to zero. And, we emphasize that $\phi_0(x)$ is a modeling tool at our disposal with no counterpart in the eikonal.

Without further ado, the Euler-Lagrange equation corresponding to the variational problem in (10) is

$$-\lambda^2 \nabla^2 \phi + c^2(x) \phi(x) = c^2(x) \phi_0(x), \quad (11)$$

which is similar to the inhomogeneous, screened Poisson equation in (2). Once again, consider the following pair of

transformations: $\phi(x) = \exp\left\{-\frac{S(x)}{\lambda}\right\}$ and $\phi_0(x) = \exp\left\{-\frac{S_0(x)}{\lambda}\right\}$ with $S(x)$ and $S_0(x)$ arbitrary. While we bypass justification of these transformations, note that $\phi(x)$ and $\phi_0(x)$ are related to time (taken to reach a point) and can therefore be expected to be positive. Taking appropriate derivatives and substituting the results in (11), we get (as shown in §8),

$$\|\nabla S\|^2 - \lambda \nabla^2 S = c^2(x) - c^2(x) \exp\left\{-\frac{S_0(x) - S(x)}{\lambda}\right\}. \quad (12)$$

Assuming bounded derivatives and provided $S(x) < S_0(x)$ almost everywhere, the above nonlinear differential equation can be approximated at small values of λ as

$$\|\nabla S\| \approx c(x), \quad (13)$$

the eikonal equation. We have shown that a linear differential equation—an Euler-Lagrange equation corresponding to a very simple but new variational principle in (10)—can play a role similar to the eikonal. The scalar field $\phi(x)$ is the counterpart to $S(x)$ in the eikonal. When we apply a nonlinear, pointwise transformation on $\phi(x)$, we obtain a nonlinear differential equation, very similar to the eikonal, but containing a viscosity term modulated by a free parameter λ (which also controls the degree of regularization in the variational principle). The linear differential equation in (2) can be discretized to obtain a sparse linear system of equations. Algorithm design is less of an issue in this space with myriad efficient linear solvers at our disposal. To our knowledge, this ours is the first derivation of a formal energy functional that leads to a linear PDE Euler-Lagrange equation (11), which also corresponds, in the limit as $\lambda \rightarrow 0$, to the nonlinear, general eikonal equation (9).

In light of the above discussion, the contributions of the present work include:

- Introduction of a new variational principle, very closely related to standard energy functionals used in vision over the past thirty years, resulting in a linear Euler-Lagrange differential equation that corresponds, in the limit, to the nonlinear general eikonal equation.
- Viscosity terms appear naturally in our setup as opposed to the original eikonal.
- A standard discretization approach resulting in a sparse linear solver.
- Demonstration of our framework in several application domains, including path planning, shape from shading and vessel segmentation.

As a preview, Figure 1 illustrates a solution for $S(x)$ obtained using our linear system approach and the resulting optimal paths from three different locations.

3. RELATED WORK

The eikonal equation (9) has applications in numerous fields, with variants of it having been successfully applied in computer vision (shape from shading [12], distance transforms [14], skeletons [19]), path planning [20], and optimal control, with numerous accounts in theoretical physics. The work in path planning and general geodesic processing is of particular interest, since we will later develop the proper




Figure 1: Left: $S(x)$ obtained via the linear differential equation as opposed to directly solving the non-linear eikonal equation. Middle: Contours of $S(x)$, red dot indicates source point. Right: Optimal paths from various locations to common source. (Note: constraint areas are in white and traversable regions in black.)

linear analog to all previous nonlinear approaches. Mitchell [13] demonstrated the use of the eikonal equation for constrained path planning with multiple constraints, and more recently, the work in [20] has illustrated its use for simple car navigation. With the use of structured meshes, path planning using the eikonal equation can be extended to geodesic processing on curved surfaces [3]. For completeness, we also include [21], which recently demonstrated the use of the classical-quantum connection in the context of reinforcement learning (but made no contact with the eikonal).

Regardless of the application domain, what is common among all these approaches are the numerical algorithms used to solve eq. (9). Almost all default to the use of fast marching and fast sweeping. The fast marching method [14] is a well-known algorithm and has a complexity of $O(N \log N)$, where N is the number of grid points in Ω (domain). The $\log N$ factor is due to the overhead of administering a priority queue data structure. More recently, the work in [23] suggested a clever solution using an untidy priority queue, which improved the complexity to $O(N)$. However, this comes at the price of approximation errors in the computed solution. Analysis in [17] has shown that the explicit estimate of the error introduced by using an untidy priority queue is tied to the speed function $F(x)$ [inverse of cost $c(x)$ in eq. (9)]; the complexity bound is really $O\left(\frac{F_{\max}}{F_{\min}} N\right)$. Given a large ratio $\frac{F_{\max}}{F_{\min}}$, there can be robustness problems. This is a disconcerting characteristic for path planning—our focus in the present work—where very large and small values are assigned to $c(x)$ to designate areas of admissible traversability.

Fast sweeping [24] essentially executes Gauss-Seidel iterations through the grid. It is generally easier to implement than fast marching and is computationally nicer, having predictable memory access patterns. For a single pass over the grid, the complexity is $O(N)$. However, this is an iterative algorithm; hence, the number of passes over the whole grid really depends on the problem at hand—heuristic rules suggest 2^n sweeps for $\Omega \subset \mathbb{R}^n$.

A logical precursor to the work presented here is the Poisson equation approach in [9]. In this previous work, Poisson equations are used in wave function-like shape representations (but not for the eikonal). The use of the Schrödinger equation in [10] leads to an eikonal solver but the authors do not demonstrate a variational principle underlying the differential equation and do not solve a sparse linear system as we have done. From a broader perspective, one can view our work as a further step in the logical progression of the application of linear differential equations in nonlinear

domains.

In our framework, we circuitously avoid direct solution of the nonlinear eikonal, and instead only have to solve a linear wave equation (second-order PDE). We can leverage a plethora of direct and iterative solutions to solve the resulting sparse, linear system, with complexity measures ranging from optimal $O(N)$ for multigrid [18, 22], logarithmic spectral $O(N \log N)$ techniques [7], and polynomial $O(N^{1.5})$ direct solvers [6]. In addition to these favorable complexity measures, sparse linear solvers have less complex implementation requirements than fast marching or fast sweeping, exhibit better numerical stability, and are typically included as part of many numerical libraries, e.g. MATLAB[®].

The remainder of this paper is organized as follows. In the next section, we detail the physical motivations and theory that connects the variational approach used here with standard variational principles used in physics. Here, we also make the connections between our framework and viscosity solutions to the eikonal equation. We then discuss the specifics as they relate to developing our path planning formalism, §5, followed by a brief exposition shape from shading, §6. Experimental results are showcased in §7.1, demonstrating optimal path recovery on complicated constraint domains, segmentation of vessels and surface reconstruction. The last section concludes by summarizing our effort and proposing directions for future work.

4. RELATIONSHIP TO PHYSICS-BASED VARIATIONAL PRINCIPLES AND VISCOSITY SOLUTIONS

In the theoretical physics literature, we encounter the following variational principle [2] on a wave function $\psi(x)$:

$$H[\psi] = \int_{\Omega} \left\{ \frac{\hbar^2}{2m} \|\nabla\psi\|^2 + V(x)|\psi|^2 \right\} dx \quad (14)$$

in which \hbar is the reduced Planck's constant, m is the mass of a free particle moving under a potential $V(x)$. The Euler-Lagrange equation derived from this variational principle is the static Schrödinger equation

$$-\frac{\hbar^2}{2m} \nabla^2 \psi + V(x)\psi = 0. \quad (15)$$

When we set $\psi(x) = R(x) \exp\left\{i\frac{S(x)}{\hbar}\right\}$ in (15), we get two

equations, one for $R(x)$ and the other for $S(x)$:

$$\begin{aligned} \nabla \cdot (R^2 \nabla S) &= 0, \\ \frac{(\nabla S)^2}{2m} + V(x) - \frac{\hbar^2}{2m} \frac{\nabla^2 R}{R} &= 0. \end{aligned} \quad (16)$$

In our setup, the “wave function” is not complex and we set $\phi(x) = \exp\left\{-\frac{S(x)}{\lambda}\right\}$ with λ playing a role similar to \hbar in (14). The change in the definition of ϕ relative to ψ above allows us to obtain a single static Hamilton-Jacobi equation with a viscosity term in (12). Otherwise, there are broad similarities. It remains to be seen if there is any benefit to using the complex wave function formalism in image analysis problems.

There is a remarkable connection between the framework detailed here and the theory of viscosity solutions for eikonal equations [5]. Viscosity solutions arise from the need to mathematically characterize solutions to the eikonal equation. More specifically, we desire the ability to prove properties such as uniqueness of the solution, its regularity, consistency, etc. For a nonlinear PDE such as the general eikonal, eq. (9), we mathematically cannot define a point-wise convergent solution that strictly satisfies the differential equation. Hence, the basic approach in seeking viscosity solutions for eikonal problems is to introduce an additional, parameter-influenced term in the equation to get

$$\|\nabla S\| - \tau \nabla^2 S = c(x), \quad (17)$$

where τ is a free parameter. The solution S is realized in the limiting behavior of $\tau \rightarrow 0$ [5, 1]. The artificial incorporation of this second-order term endows the approximate solution desirable mathematical properties such as uniqueness.

The elegance of the present framework, where we solve the linear PDE corresponding to nonlinear eikonal, is that this viscosity-type behavior is a natural outcome of the Hamilton-Jacobi formalism that connects the classical eikonal equation with its quantum, wave equation counterpart. Referring back to eq. (11), we see that the second-order Laplacian term precisely exhibits this behavior, with λ^2 playing the role of τ . All desirable mathematical niceties are obtained for free—properties of uniqueness, stability, etc. are all well established for such linear, Poisson-type PDEs. *We consider this a novel contribution*; as, to our knowledge, this is the first time that the eikonal equation is solved using a sparse, linear system approach (as we will later detail); thus, circumventing the entire discipline of nonlinear viscosity solutions which aim at directly solving the eikonal equation.

5. APPLICATIONS TO PATH PLANNING

In the pioneering contribution of [12], Kimmel and Sethian developed one of the earliest applications of the eikonal equation to path planning. By finding a solution $S(x)$ (referred to as the value function in the path planning context) to the eikonal equation in eq. (9), we immediately recover the minimum cost to go from a source location x_0 in the state space to any other point x (in the state space). Here, we impose the boundary condition $S(x_0) = 0$, and consider $c(x)$ as the cost to travel through location x (higher the value, the more costly) and prescribe it as a strictly positive function, and $S : \mathbb{R}^d \rightarrow \mathbb{R}^+[0, \infty)$. In comparison to other popular path planning techniques like potential field methods [11], the value function is an example of a navigation function—a potential field free of local minima.

Given a scalar field solution, the optimal paths are determined by gradient descent on $S(x)$, and is typically referred to as backtracking. The backtracking procedure can be formulated as an ordinary differential equation

$$\dot{x} = -\frac{\nabla S(x(t))}{\|\nabla S(x(t))\|}, \quad (18)$$

whose solution $x(t)$ is the reconstructed path from a fixed target location x_T . We typically terminate the gradient backtracking procedure at some t value such that $\|x(t) - x_0\| < \epsilon$, i.e. we get arbitrarily close the source point, for some small $\epsilon > 0$. By construction, the backtracking on $S(x)$ cannot get stuck in local minima—an obvious proof by contradiction validates this claim if one considers S to be differentiable and have local minima $\nabla S = 0$ at some x , but $c(x) > 0$, and contradicts the eikonal equation $\|\nabla S\| = c(x)$. Although, in theory, $S(x)$ can contain saddle points, but usually this is not an issue in practice.

As we have espoused in §2, the linear differential equation—our counterpart to the general nonlinear eikonal in eq. (9)—is the inhomogeneous, screened Poisson equation in (11). Here, we choose to directly solve this linear differential equation and get increasingly better approximations to the true $S(x)$ by solving the system at small values of λ .

Solving the distributional form of the screened Poisson equation is far less complicated to implement than the fast marching or fast sweeping methods needed for directly solving (9). In addition, the computational complexity for implementing (11) can match these algorithms since $O(N)$ sparse, linear system solvers are available [18]. This comes from the fact that a finite difference approximation to (11), using a standard five-point Laplacian stencil, simply results in a sparse linear system of the form [7]:

$$\underbrace{\begin{bmatrix} T_N + 2I_N & -I_N & \cdots & 0 \\ -I_N & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & -I_N \\ 0 & & -I_N & T_N + 2I_N \end{bmatrix}}_A \underbrace{\begin{bmatrix} \phi_1 \\ \phi_2 \\ \vdots \\ \phi_N \end{bmatrix}}_x = \underbrace{\begin{bmatrix} c_1^2 \phi_0(1) \\ \vdots \\ c_{x_0}^2 \phi_0(x_0) \\ \vdots \\ c_N^2 \phi_0(N) \end{bmatrix}}_b, \quad (19)$$

where N represents the number of grid points and T_N is a tri-diagonal block of the form

$$T_N = \begin{bmatrix} 2 & -1 & & 0 \\ -1 & \ddots & \ddots & \\ & \ddots & \ddots & -1 \\ 0 & & -1 & 2 \end{bmatrix}.$$

In eq. (19), there are three minor implementation details hidden in the sparse structure. First, when the corresponding grid locations of the main diagonal terms incur cost from $c(x)$, these will have to be added to that diagonal location. Next, the right-hand side notation [with $\phi_0(x_0)$] highlights

the fact that only the grid location(s) of the source point(s) should have a high value (e.g. one) while all others are close to zero. Finally, the $c(x)$ function should be set to high values for undesirable travel regions for the optimal path, and very small values for favorable travel areas (and set to one on the source point). Though better complexities are achievable using multigrid [18] solvers, one can just as well address many problems in a satisfactory manner using direct, sparse solvers, like MATLAB’s $A \setminus b$ —an approach we adopt for the experiments in the present paper.

6. APPLICATIONS TO SHAPE-FROM-SHADING

Shape-from-shading has long been a popular problem domain for computer vision, having the primary objective of recovering the scalar height field from a single image. Solution approaches utilizing the eikonal equation have been known since the early 80’s [4], and have continually improved upon through the advent of fast sweeping and fast marching methods [12, 15]. The standard forward image model assumes a Lambertian reflectance model generates the luminance via inner product of the surface normal, $N(x)$, with the light source direction, d , i.e. $L(x) = \langle N(x), d \rangle$. For example, if we assume a vertical lighting direction $d = [0, 0, 1]^T$, we get the imaging operator

$$L(x) = \frac{1}{\sqrt{\|\nabla S(x)\|^2 + 1}},$$

where $S(x)$ is the desired scalar height field we wish to recover. This can be obtained by solving the standard eikonal equation, eq. (9), with

$$c(x) = \sqrt{\frac{1}{L(x)^2} - 1} \quad (20)$$

and boundary conditions $S(x_i) = h_i$, i.e. we seed the boundary conditions with the known heights h_i at select grid locations x_i .

As we have detailed in previous sections, our formalism allows one to address any general (non-linear) eikonal equation by solving the linear screened Poisson equation in (11). One simply needs to create the forcing function in (20), and then solve the discretized sparse system as in (19). This immediately yields the recovered height field.

7. EXPERIMENTAL VALIDATION

7.1 Path Planning on Complex Mazes and Extensions to Vessel Segmentation

We applied our path planning approach to a variety of complex maze images. We explicitly chose the maze grid sizes to be much larger than the norm for recent publications that apply the eikonal equation for path planning; these typical sizes are usually smaller than 100×100 . An objective juxtaposition of contemporary fast sweeping and marching techniques, which require special discretization schemes, data structures, sweep orders, etc., versus our approach presented here, clearly illustrates the efficiency and simplicity of the later. Our framework reduces path planning (a.k.a. all-pairs, shortest path or geodesic processing) implementation to four straightforward steps:

2D Grid Dims. (No. of Points)	$A \setminus b$ (sec.)
450×450 ($N = 202,500$), Fig. 1	0.79
434×493 ($N = 213,962$), Fig. 2(a)	1.02
621×473 ($N = 293,733$), Fig. 2(b)	1.22
419×496 ($N = 207,824$), Fig. 2(c)	0.77

Table 1: Maze grids sizes and time to solve sparse system for path planning. Our approach simply uses MATLAB’s \setminus operator to solve the shortest path problem, avoiding complex nuances of discretization schemes and specialized data structures required for fast marching.

1. Define $c(x)$, which assigns a high cost to untraversable areas in the grid and low cost to traversable locations. In the experiments here, we simply let appropriately scaled versions of the maze images be $c(x)$, with *white pixels representing boundaries* and *black pixels the possible solution paths*.
2. Select a source point on the grid and a small value for λ . The solution to (11) will simultaneously recover all shortest paths back to this source from any non-constrained region in the grid.
3. Use standard finite differencing techniques to evaluate (11). This leads to a sparse, block tri-diagonal system, as described in §5, which can be solved by a multitude of linear system numerical packages. We simply use MATLAB’s \setminus operator. Recover approximate solution to eikonal by letting $S(x) = -\lambda \log \phi(x)$.
4. Backtrack to find the shortest path from any allowable grid location to the source, i.e. use eq. (18) to perform standard backtracking on $S(x)$.

The grid sizes and execution times for several mazes are provided in Table 1. Notice that even for larger grids, our time to solve for $S(x)$ is on the order of a few seconds, and this is simply using the basic sparse solver in MATLAB. The time complexity of MATLAB’s direct solver is $O(N^{1.5})$, which makes our approach here slightly slower than the optimal runtime. $O(N)$ is achievable for sparse systems, such as ours, using multigrid methods, but we have opted to showcase the simplicity of our implementation versus pure speed. Figure 2 illustrates our path planning approach on a variety of mazes: (a) demonstrates path planning while paying homage to Schrödinger’s cat, (b) is a traditional maze, and (c) is a whimsical result on a skull maze. Notice in all these mazes there are multiple solution paths back to the source, but only the shortest path is chosen.

The above discussed application of path planning can be readily extended to centerline extraction from medical imagery of blood vessels. One can view the image $I(x)$ as a “maze” where we only want to travel on the vessels in the image. Figure 3, column (a) illustrates three example medical images: eye, brain, and hand. Columns (b) showcases our results, while (c) provides comparative analysis against fast sweeping. Notice that our solution naturally generates smoother centerline segmentations, which is a natural consequence of having a built-in, viscosity-like term in (12). Whereas, the fast marching and fast sweeping methods tend to have sharper transitions in the paths and deviate from the center—viscosity solutions can be used to alleviate this,




Figure 2: (a) Schrödinger’s cat maze, (b) standard maze, (c) skull maze. Multiple solutions are possible but only the shortest path is chosen to be optimal. (Note: constraint areas are in white and traversable regions in black.)




Figure 3: Medical image vessel centerline extraction, top row original images and bottom row our linear solution approach. Paths under our linear systems approach are smooth due to inherent viscosity-like behavior. They also produce segmentations where the extracted vessels are centered on the blood vessels. FM approach requires additional constraints to achieve centerline extraction [8].

but are not organic to the formulation like ours. In fact, it has been shown that additional constraints have to be incorporated to ensure fast marching approaches extract the centerline [8]. Again, we stress the simplicity and ease-of-use of this approach, with only one free parameter λ , making it a viable option for many path planning related applications, such as robotic navigation, optimal manipulation, and vessel extraction in medical images.

7.2 Surface Reconstruction via Shape-from-Shading

For shape-from-shading, we validated height recovery on two common images that often used in the literature: Mozart and a vase. Figure 4 illustrates the recovered surfaces using our method, (a), fast marching, (b), and fast sweeping, (c). Under each image we also list the error of the reconstruction from the known ground truth height field. The error was computed by comparing the true mean gradient magnitudes versus those estimated from the recovered $S(x)$.

The validation shows that our method is competitive with both fast marching and fast sweeping, all the while retaining the efficiency and simplicity of obtain a solution through a

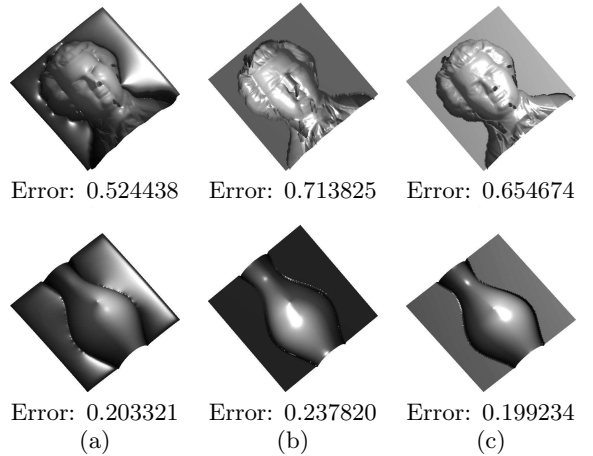


Figure 4: Shape-from-shading surface reconstruction, per column: (a) our linear solution approach, (b) fast marching, (c) fast sweeping. Based on the gradient magnitude error (from the true surface), our approach linear systems approach is better or at least highly competitive.

sparse linear system. Going beyond the present work, our general framework can be adapted to all previous application areas of the eikonal equation, and, as alluded to earlier, the variational objective can be readily modified for to incorporate other constraints that may lead to better reconstructions.

8. CONCLUSION

The Hamilton-Jacobi equation, particularly its specialized form as the eikonal equation, is at the heart of numerous applications in vision (shape-from-shading, path planning, medial axis, etc.), and spurred the rapid development of several innovative computational techniques to directly solve this *nonlinear* PDE, including fast marching, and fast sweeping. However, lost in this flurry of advancing nonlinear solvers was a completely alternative approach, one which allows you to rigorously approximate solutions to the nonlinear eikonal by instead solving a *linear* differential equation which can be discretized and solved using a *sparse, linear solver*. We formally derived new variational principles that lead to the Poisson-class equation and its connection to the classical eikonal equation. This opens the door for myriad future explorations that can be realized by simply modifying the variational objective. We applied our framework to develop a novel solution to the classical all-pairs, shortest path problem (a.k.a. path planning)—for the first time demonstrating that the solution is achievable by simply solving a sparse, linear system. We also illustrated results on shape-from-shading and vessel centerline extraction. Our approach is straightforward to implement (by deploying any sparse linear solver) and holds its own against contemporary fast marching and fast sweeping methods while possessing the considerable advantage of linearity. In addition, a direct consequence of our mathematical formulation is that viscosity solutions are naturally incorporated and obtained when solving the linear differential equation—allowing one to circumvent explicit viscosity constructions required for

any method that tries to directly solve the nonlinear eikonal. In future work, we plan to revisit past uses of the eikonal equation and examine improvements gained by the adoption of our framework. We are also investigating extensions of this approach to other areas such as control theory.

Acknowledgments

This work is supported by NSF IIS 1065081 and the support of the AIRBUS Group Corporate Foundation Chair in Mathematics of Complex Systems established in ICTS-TIFR.

9. REFERENCES

- [1] M. Bardi and I. Capuzzo-Dolcetta. *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Birkhäuser, 2008.
- [2] D. Bohm. A suggested interpretation of the quantum theory in terms of “hidden variables”, i and ii. In *Quantum Theory and Measurement*, Princeton Series in Physics, pages 369–402. Princeton University Press, 1984.
- [3] A. Bronstein, M. Bronstein, and R. Kimmel. *Numerical geometry of non-rigid shapes*. Springer, 2010.
- [4] A. Bruss. The eikonal equation: some results applicable to computer vision. *Journal of Mathematical Physics*, 23(5):890–896, 1982.
- [5] M. Crandall and L. Pierre-Louis. Viscosity solutions of Hamilton-Jacobi equations. *Transactions of the American Mathematical Society*, 277(1):1–42, 1983.
- [6] T. Davis. *Direct Methods for Sparse Linear Systems*. SIAM, 2006.
- [7] J. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [8] T. Deschamps and L. Cohen. Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *Medical Image Analysis*, 5(4):281–299, 2001.
- [9] L. Gorelick, M. Galun, E. Sharon, R. Basri, and A. Brandt. Shape representation and classification using the Poisson equation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12):1991–2005, 2006.
- [10] K. S. Gurumoorthy and A. Rangarajan. A fast eikonal equation solver using the Schrödinger wave equation. *CoRR*, abs/1403.1937, 2014.
- [11] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5(1):90–98, 1986.
- [12] R. Kimmel and J. Sethian. Optimal algorithm for shape from shading and path planning. *Journal of Mathematical Imaging and Vision*, 14:237–244, 2001.
- [13] I. Mitchell. Continuous path planning with multiple constraints. In *IEEE Conference on Decision and Control*, pages 5502–5507, 2003.
- [14] S. Osher and R. Fedkiw. *Level set methods and dynamic implicit surfaces*. Springer-Verlag, 2003.
- [15] E. Pados and O. Faugeras. *Handbook of Mathematical Models in Computer Vision*, chapter Shape from Shading, pages 275–388. Springer, 2006.
- [16] A. Rangarajan and K. S. Gurumoorthy. A Schrödinger wave equation approach to the eikonal equation:

Application to image analysis. In *Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, pages 140–153, 2009.

- [17] C. Rasch and T. Satzger. Remarks on the $O(N)$ implementation of the fast marching method. *IMA Journal of Numerical Analysis*, 29:806–813, 2009.
- [18] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, 2nd edition, 2003.
- [19] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. Zucker. Hamilton-Jacobi skeletons. *International Journal of Computer Vision*, 48(3):215–231, 2002.
- [20] R. Takei, R. Tsai, H. Shen, and Y. Landa. A practical path-planning algorithm for a simple car: a Hamilton-Jacobi approach. In *American Control Conference (ACC)*, pages 6175–6180, July 2010.
- [21] E. Theodorou, J. Buchli, and S. Schaal. A generalized path integral control approach to reinforcement learning. *Journal of Machine Learning Research*, 11:3137–3181, 2010.
- [22] J. Xia, S. Chandrasekaran, M. Gu, and X. Li. Superfast multifrontal method for large structured linear systems of equations. *SIAM J. Matrix Analysis and Applications*, 31:1382–1411, 2009.
- [23] L. Yatziv, A. Bartesaghi, and G. Sapiro. A fast $O(N)$ implementation of the fast marching algorithm. *Journal of Computational Physics*, 212:393–399, 2006.
- [24] H. Zhao. A fast sweeping method for eikonal equations. *Mathematics of Computation*, 74:603–627, 2005.

Appendix

When $\phi(x) = \exp\left\{-\frac{S(x)}{\lambda}\right\}$, the gradient of ϕ is

$$\nabla\phi = -\frac{1}{\lambda} \exp\left\{-\frac{S(x)}{\lambda}\right\} \nabla S. \quad (21)$$

From this, the Laplacian can be written as

$$\nabla^2\phi = \exp\left\{-\frac{S(x)}{\lambda}\right\} \left(\frac{1}{\lambda^2} \|\nabla S\|^2 - \frac{1}{\lambda} \nabla^2 S\right). \quad (22)$$

Substituting (22) into (11), we obtain

$$-\|\nabla S\|^2 + \lambda \nabla^2 S + c^2(x) = c^2(x) \exp\left\{-\frac{S_0(x) - S(x)}{\lambda}\right\}. \quad (23)$$

Recall that $\phi_0(x) = \exp\left\{-\frac{S_0(x)}{\lambda}\right\}$ behaves like a delta function. As $S(x)$ is bounded within the domain under consideration, $S_0(x)$ can be chosen much larger than $S(x)$, except at $x = 0$, so that the right side approaches zero as $\lambda \rightarrow 0, \forall x \neq 0$. The additional $\lambda \nabla^2 S$ term [relative to (9)] is referred to as the *viscosity term* which emerges naturally from our Euler-Lagrange equation—an intriguing result. Since $|\nabla^2 S|$ is bounded, as $\lambda \rightarrow 0$, (23) tends to $\|\nabla S\| = c(x)$ which is the original eikonal equation (9).