

Fast Computation of Weighted Distance Functions and Geodesics on Implicit Hyper-Surfaces

Facundo Mémoli* and Guillermo Sapiro†

**Instituto de Ingeniería Eléctrica, Universidad de la República, Montevideo, Uruguay*; †*Electrical and Computer Engineering, University of Minnesota, Minneapolis, Minnesota 55455*

E-mail: memoli@iie.edu.uy; guille@ece.umn.edu

Received March 26, 2001; revised August 6, 2001

An algorithm for the computationally optimal construction of intrinsic weighted distance functions on implicit hyper-surfaces is introduced in this paper. The basic idea is to approximate the intrinsic weighted distance by the Euclidean weighted distance computed in a band surrounding the implicit hyper-surface in the embedding space, thereby performing all the computations in a Cartesian grid with classical and efficient numerics. Based on work on geodesics on Riemannian manifolds with boundaries, we bound the error between the two distance functions. We show that this error is of the same order as the theoretical numerical error in computationally optimal, Hamilton–Jacobi-based, algorithms for computing distance functions in Cartesian grids. Therefore, we can use these algorithms, modified to deal with spaces with boundaries, and obtain also for the case of intrinsic distance functions on implicit hyper-surfaces a computationally efficient technique. The approach can be extended to solve a more general class of Hamilton–Jacobi equations defined on the implicit surface, following the same idea of approximating their solutions by the solutions in the embedding Euclidean space. The framework here introduced thereby allows for the computations to be performed on a Cartesian grid with computationally optimal algorithms, in spite of the fact that the distance and Hamilton–Jacobi equations are intrinsic to the implicit hyper-surface. © 2001 Academic Press

Key Words: implicit hyper-surfaces; distance functions; geodesics; Hamilton–Jacobi equations; fast computations.

1. INTRODUCTION

Computing distance functions has many applications in numerous areas including mathematical physics, image processing, computer vision, robotics, computer graphics, computational geometry, optimal control, and brain research. In addition, having the distance

function from a seed to a target point, it is straightforward to compute the corresponding geodesic path, since this is given by the gradient direction of the distance function, back propagating from the target toward the seed (see for example [18]). Geodesics are used for example for path planning in robotics [35]; brain flattening and brain warping in computational neuroscience [56, 57, 60, 61, 68]; crests, valleys, and silhouettes computations in computer graphics and brain research [7, 30, 62]; mesh generation [64]; and many applications in mathematical physics. Distance functions are also very important in optimal control [59] and computational geometry for computations such as Voronoi diagrams and skeletons [48]. It is then of extreme importance to develop efficient techniques for the accurate and fast computations of distance functions. It is the goal of this paper to present a computationally optimal technique for the computation of intrinsic weighted distance functions on implicit hyper-surfaces.¹ It is well known already, and it will be further detailed below, that these weighted distances can be obtained as the solution of simple Hamilton–Jacobi equations. We will also show that the framework here presented can be applied to a larger class of Hamilton–Jacobi equations defined on implicit surfaces. We also discuss the application of our proposed framework to other, nonimplicit, surface representations.

1.1. Distance Function Computation and Its Hamilton–Jacobi Formulation

Before proceeding, let us first formally define the concept of *intrinsic* weighted distances on *implicit* hyper-surfaces. Let \mathcal{S} be a (codimension 1) hyper-surface in \mathbb{R}^d defined as the zero level set of a function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$. That is, \mathcal{S} is given by $\{x \in \mathbb{R}^d : \psi(x) = 0\}$. We assume from now on that ψ is a signed distance function to the surface \mathcal{S} . (This is not a limitation, since as we will discuss later, both explicit and implicit representations can be transformed into this form.) Our goal is, for a given point $p \in \mathcal{S}$, to compute the intrinsic g -weighted distance function $d_{\mathcal{S}}^g(p, x)$ for all desired points $x \in \mathcal{S}$.² Note that we are referring to the intrinsic g -distance, that is, the geodesic distance on the Riemannian manifold $(\mathcal{S}, g^2 \mathbb{I})$ (\mathbb{I} stands for the $(d-1) \times (d-1)$ identity matrix) and not on the embedding Euclidean space. For a given positive weight g defined on the surface (we are considering only isotropic metrics for now), the g -distance on \mathcal{S} (that coincides with the geodesic distance of the Riemannian manifold $(\mathcal{S}, g^2 \mathbb{I})$) is given by

$$d_{\mathcal{S}}^g(p, x) \triangleq \inf_{C_{px}[\mathcal{S}]} \{\mathbf{L}_g(C)\}, \quad (1)$$

where

$$\mathbf{L}_g\{C\} \triangleq \int_a^b g(C(l)) \|\dot{C}(l)\| dl \quad (2)$$

is the weighted *length functional* defined for piecewise C^1 curves $C : [a, b] \rightarrow \mathcal{S}$, and $C_{px}[\mathcal{S}]$ denotes the set of curves piecewise C^1 joining p to x , traveling on \mathcal{S} . In general,

¹Although all the examples in this paper are going to be reported for two-dimensional surfaces in 3D (here denoted as 3D surfaces), the theory is valid for general dimension hyper-surfaces, and it will be presented in this generality. A number of applications deal with higher dimensions. For example, for the general theory of harmonic maps, in order to deal with maps onto general open surfaces, it is necessary to have this notion of intrinsic distance [41]. In addition, higher dimensions might appear in motion planning, when explicitly assuming that the robot is not modeled by a point, thereby adding additional constraints to its movements.

²This can certainly be extended to any subset of \mathcal{S} .

we will consider the definition to be valid for any \tilde{g} defined over the domain that the curve may travel through.

We need to compute this distance when all the concerning objects are represented in discrete form in the computer. Computing minimal weighted distances and paths in graph representations is an old problem that has been optimally solved by Dijkstra [21]. Dijkstra showed an algorithm for computing the path in $O(n \log n)$ operations, where n is the number of nodes in the graph. The weights are given on the edges connecting between the graph nodes, and the algorithm is computationally optimal. In theory, we could use this algorithm to compute the weighted distance and corresponding path on polygonal (not implicit) surfaces, with the vertices as the graph nodes and the edges the connections between them (see [33]). The problem is that the optimal paths computed by this algorithm are limited to travel on the graph edges, giving only a first approximation of the true distance. Moreover, Dijkstra's algorithm is not a consistent one: it will not converge to the true desired distance when the graph and grid is refined [42, 43]. The solution to this problem, limited to Cartesian grids, was developed in [27, 51, 52, 59] (and recently extended by Osher and Helmsen, see [45]). Tsitsiklis [59] first described an optimal-control type of approach, while independently Sethian [51, 52] and Helmsen [27] both developed techniques based on upwind numerical schemes. The solution presented by these authors is consistent and converges to the true distance [49, 59], while keeping the same optimal complexity of $O(n \log n)$. Later this work was extended in [32] for triangulated surfaces (see also [7, 36] for related works on numerics on non-Cartesian grids). We should note that the algorithm developed in [32] is currently developed only for triangulated surfaces with acute triangles. Therefore, before the algorithm can be applied, as an initialization step the surfaces have to be preprocessed to remove all obtuse triangles or other polygons present in the representation [31]. Following [52], we call these *fast marching algorithms*.

The basic idea behind the computationally optimal techniques for finding weighted distances, meaning these fast marching algorithms, is to note that the distance function satisfies a Hamilton–Jacobi partial differential equation (PDE) in the viscosity sense; see, for example, [38, 50] for the general topic of distance functions on Riemannian manifolds (and a nice mathematical treatment), and [12, 23, 31, 44, 46, 52] for the planar (and more intuitive) case. This Hamilton–Jacobi equation is given by

$$\|\nabla_S d_S^g\| = g, \quad (3)$$

where ∇_S is the gradient intrinsic to the surface, and d_S^g is the g -distance from a given seed point to the rest of the manifold.³

That is, we can transform the problem of optimal distance computation into the problem of solving a Hamilton–Jacobi equation (recall that g is known, it is the given weight), also known as the Eikonal equation. In order to solve this equation, the current state of knowledge permits us to accurately and optimally (in a computational sense) find (weighted) distances on Cartesian grids as well as on particular triangulated surfaces (after some preprocessing, namely the elimination of obtuse triangles, see [6, 32]). The goal of this paper is to extend this to implicit hyper-surfaces. In other words, we will show how to solve the above Eikonal equation for implicit hyper-surfaces \mathcal{S} .

³ Note that ∇_S and d_S^g become the classical gradient and distance respectively for Euclidean spaces.

Recall that although all the applications in this paper will be presented for 3D surfaces, the theory is valid for any d -dimensional hyper-surfaces, and will then be presented in this generality.

1.2. Distance Function and Geodesics on Implicit Surfaces

The motivations behind extending the distance map calculation to implicit surfaces are numerous: (a) In many applications, surfaces are already given in implicit form, e.g., [10, 13, 16, 25, 45, 46, 53, 66, 62], and there is then a need to extend to this important representation the previously mentioned fast techniques. We could of course triangulate the implicit surface, eliminate obtuse triangles, and then use the important algorithm proposed in [32]. This is not a desirable process in general when the original data is in implicit form, since it affects the distance computation accuracy because of errors from the triangulation, and also adds the computational cost of the triangulation itself, triangulation that might not be needed by the specific application. If, for example, what it is needed is to compute the distance between a series of points on the surface, the computational cost added by the triangulation is unnecessary. Note that finding a triangulated representation of the implicit surface is of course dimensionality dependent, and adds the errors of the triangulation process. Moreover, accurate triangulations that ensure correctness in the topology are computationally expensive, and once again there is no reason to perform a full triangulation when we might be interested just in the intrinsic distance between a few points on the implicit surface. (b) It is a general agreement that the work on implicit representations and Cartesian grids is more robust when dealing with differential characteristics of the surface and partial differential equations on this surface. Numerical analysis on Cartesian grids is much more studied and supported by fundamental results than the work on polygonal surfaces. It is even recognized that there is no consensus on how to compute basic differential quantities over a triangulated surface (see, for example, [22]), although there is quite an agreement for implicit surfaces. Moreover, representing an hyper-surface with structured elements such as triangles is certainly difficult for dimensions other than 2 or 3. (c) If the computation of the distance function is just a part of a general algorithm for solving a given problem, it is not elegant, accurate, nor computationally efficient to go back and forth from different representations of the surface.

Before proceeding, we should note that although the whole framework and theory is here developed for implicit surfaces, it is valid for other surface representations as well after preprocessing. This will be explained and discussed later in the paper (Section 5). Moreover, we will later assume that the embedding is a distance function. This is not a limitation, since many algorithms exist to transform a generic embedding function into a distance one; see also Section 5. Therefore, the framework here presented can be applied both to implicit (naturally) and other surface representations such as triangulated ones.

In order to compute intrinsic distances on surfaces, a small but important number of techniques have been reported in the literature. As mentioned before, in a very interesting work Kimmel and Sethian [32] extended the fast marching algorithm to work on triangulated surfaces. In its current version, this approach can only be used when dealing with 3D triangulated surfaces and its extension to deal with higher dimensions seems very involved. Moreover, it can only correctly handle acute triangulations (thereby requiring a preprocessing step). And of course, it doesn't apply to implicit surfaces without some preprocessing (a triangulation).

Another very interesting approach to computing intrinsic distances, this time working with implicit surfaces, was introduced in [16]. This will be further described below, but first let's make some comments on it. First, this is an evolutionary/iterative approach, whose steady state gives the solution to the corresponding Hamilton–Jacobi equation. Therefore, this approach is not computationally optimal for the class of Hamilton–Jacobi equations discussed in this paper. (Although when properly implemented, the computational complexity of this iterative scheme is the same as in the fast marching method here proposed, the inner loop is more complex, making the iterative algorithm slower.)⁴ Second, very careful discretization must be done to the equation proposed in [16] because of the presence of intrinsic jump functions that might change the zero level-set (i.e., the surface). On the other hand, the numerical implementation is not necessarily done via the utilization of *monotone schemes*, as required by our approach and all the fast marching techniques previously mentioned (thereby having a theoretical error $\Theta(\sqrt{\Delta x})$ [19]), and better accuracy might then be obtained.

In order to compute the intrinsic distance on an implicit surface, we must then solve the corresponding Hamilton–Jacobi equation presented before. In order to do this in a computationally efficient way, we need to extend the fast marching ideas in [27, 45, 51, 52, 59], which assume a Cartesian grid, to work in our case. Since an implicit surface is represented in a Cartesian grid, corresponding to the embedding function, the first and most intuitive idea is then to attempt to solve the *intrinsic Eikonal equation* using the *fast marching* technique. The first step toward our goal is to express all the quantities in the intrinsic Eikonal equation by its *implicit-extended* representations. What we mean is that the *intrinsic* problem (we consider $g = 1$ for simplicity of exposition),

$$\begin{cases} \|\nabla_{\mathcal{S}} d_{\mathcal{S}}(x)\| = 1 & \text{for } p \in \mathcal{S} \\ d_{\mathcal{S}}(p) = 0, \end{cases} \quad (4)$$

with $p \in \mathcal{S}$ the seed point, is to be extended to all \mathbb{R}^d (or at least to a band surrounding \mathcal{S}), and the derivatives are to be taken tangentially to $\{\psi = 0\}$. Considering then the projection of the Euclidean gradient onto the tangent space of \mathcal{S} to obtain the intrinsic one, and denoting by \hat{d} the Euclidean extension to the intrinsic distance $d_{\mathcal{S}}$, we have to numerically solve, in the embedding Cartesian grid, the equation

$$\begin{cases} \|\nabla \hat{d}(x)\|^2 - |\nabla \hat{d}(x) \cdot \nabla \psi(x)|^2 = 1 & \text{for } x \in \mathbb{R}^d \\ \hat{d}(l(p)) = 0, \end{cases} \quad (5)$$

where $l(p)$ is the ray through p normal to the level sets of ψ .

This is exactly the approach introduced in [16], as discussed above, to build-up the evolutionary approach, given by the PDE

$$\phi_t + \text{sgn}(\phi_0) (\sqrt{\|\nabla \phi\|^2 - |\nabla \phi \cdot \nabla \psi|^2} - 1) = 0, \quad (6)$$

where $\phi_0(x) = \phi(x, 0)$ is the initial value of the evolving function, generally a step-like

⁴ The general framework introduced in [16] is applicable beyond the Hamilton–Jacobi equations discussed in this paper (see also [8, 17]). Here we limit the comparison between the techniques to the equations where both approaches are applicable.

function (convolved with the signum) that tells inside from outside of the zero level-set. One then finds $\hat{d}(\cdot) = \phi(\cdot, \infty)$.

Of course, in order to obtain a computationally optimal approach, we want to solve the stationary problem (5), and not its iterative counterpart (6). It turns out that the basic requirements for the construction of a fast marching method, even with the recent extensions in [45], do not hold for this equation. This can be explicitly shown, and has also been hinted by Kimmel and Sethian in their work on geodesics on surfaces given as graphs of functions.⁵

To recap, the fast marching approach cannot be directly applied to the computation of intrinsic distances on implicit surfaces defined on a Cartesian grid (Eq. (5)), and the state of the art in numerical analysis for this problem says that in order to compute intrinsic distances one either has to work with triangulated surfaces or has to use the iterative approach mentioned above. The problems with both techniques were reported before, and it is the goal of this paper to present a third approach that addresses all these problems.

1.3. Our Contribution

The basic idea here presented is conceptually very simple. We first consider a small h offset of \mathcal{S} . That is, since the embedding function ψ is a distance function, with \mathcal{S} as its zero level set, we consider all points x in \mathbb{R}^3 for which $|\psi(x)| \leq h$. This gives a region in \mathbb{R}^d with boundaries. We then modify the (Cartesian) fast marching algorithm mentioned above for computing the distance transform inside this h -band surrounding \mathcal{S} . Note that here all the computations are, as in the works in [27, 51, 52, 59], in a Cartesian grid. We then use this Euclidean distance function as an approximation of the intrinsic distance on \mathcal{S} . In Section 2 we show that the error between these two distances, under reasonable assumptions on the surface \mathcal{S} , is of the same order as the numerical error introduced by the fast marching algorithms in [27, 51, 52, 59].⁶ Therefore, when adapting these algorithms to work on Euclidean spaces with boundary adaptation described in Section 3, we obtain an algorithm for the computation of intrinsic distances on implicit surfaces with the same simplicity, computational complexity, and accuracy as the optimal fast marching techniques for computing Euclidean distances on Cartesian grids.⁷ In Section 3 we also explicitly discuss the numerical error of our proposed technique. Examples of the algorithm here proposed are given in Section 4. Since Osher and Helmsen have recently shown that the fast marching algorithm can be used to solve additional Hamilton–Jacobi equations, we show that the framework here proposed can be applied to equations from that class as well; this is done in Section 5. This section also discusses the use of the framework here presented for nonimplicit surfaces. Finally, some concluding remarks are given in Section 6.

⁵ We have also benefited from private conversations with Stan Osher and Ron Kimmel to confirm this claim.

⁶ In contrast with works such as [1, 47], where an offset of this form is just used to improve the complexity of the level-sets method, in our case the offset is needed to obtain a small error between the computed distance transform and the real intrinsic distance function; see next section.

⁷ Although in this paper we deal with the fast marching techniques, other techniques for computing distance functions on Cartesian grids, e.g., the fast technique reported in [11] for uniform weights, could be used as well, since the basis of our approach is the approximation of the intrinsic distance by an extrinsic one.

2. DISTANCE FUNCTIONS: INTRINSIC VS. EXTRINSIC

The goal of this section is to present the connection between the intrinsic distance function and the Euclidean function computed inside a band surrounding the (implicit) surface. We will completely characterize the difference between these two functions, mainly based on results on shortest paths on manifolds with boundary. The results here presented will also justify the use of the Cartesian fast marching algorithms for the computation of intrinsic weighted distances on implicit surfaces.

Recall that we are dealing with a closed hyper-surface \mathcal{S} in \mathbb{R}^d represented as the zero level-set of a distance function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$. That is, $\mathcal{S} = \{\psi = 0\}$. Our goal is to compute a g -weighted distance map on this surface from a seed point $q \in \mathcal{S}$.

Let

$$\Omega_h \triangleq \bigcup_{x \in \mathcal{S}} B(x, h) = \{x \in \mathbb{R}^d : |\psi(x)| \leq h\}$$

be the h -offset of \mathcal{S} (here $B(x, h)$ is the ball centered at x with radius h). It is well known that for a smooth \mathcal{S} , $\partial\Omega_h$ is also smooth if h is sufficiently small, see Appendix A for references. Ω_h is then a *manifold with smooth boundary*.

Our computational approach is based on approximating the solution of the *intrinsic* problem ($d_{\mathcal{S}}^g(p)$ is the intrinsic g -weighted distance on \mathcal{S})

$$\begin{cases} \|\nabla_{\mathcal{S}} d_{\mathcal{S}}^g(p)\| = g & \text{for } p \in \mathcal{S} \\ d_{\mathcal{S}}^g(q) = 0. \end{cases} \quad (7)$$

by that of the *Euclidean* (or *extrinsic*) one,

$$\begin{cases} \|\nabla d_{\Omega_h}^{\tilde{g}}(p)\| = \tilde{g} & \text{for } p \in \Omega_h \\ d_{\Omega_h}^{\tilde{g}}(q) = 0, \end{cases} \quad (8)$$

where \tilde{g} is a smooth *extension* of g in a domain containing Ω_h , and $d_{\Omega_h}^{\tilde{g}}(p)$ is the Euclidean \tilde{g} -weighted distance in Ω_h . Our goal is to be able to control, for points on \mathcal{S} , $\|d_{\mathcal{S}}^g - d_{\Omega_h}^{\tilde{g}}\|_{L_{\infty}(\mathcal{S})}$ with h . Note that we have replaced the intrinsic gradient $\nabla_{\mathcal{S}}$ by the Euclidean gradient and the intrinsic distance $d_{\mathcal{S}}^g(p)$ on the surface by the Euclidean distance $d_{\Omega_h}^{\tilde{g}}(p)$ in Ω_h . We have then transformed the problem of computing an intrinsic distance into the problem of computing a distance in an Euclidean manifold with boundary.

We will show that under suitable (and likely) geometric conditions on \mathcal{S} we can indeed control $\|d_{\mathcal{S}}^g - d_{\Omega_h}^{\tilde{g}}\|_{L_{\infty}(\mathcal{S})}$ with h . In order to materialize this, we first need to briefly discuss the extension \tilde{g} and to review some basic background material on Riemannian manifolds with boundary.

2.1. The Extension of the Weight g

We require that $\tilde{g}|_{\mathcal{S}} = g$, and that \tilde{g} is *smooth* and nonnegative within Ω_h . There are situations when one has a readily available extension, and others in which the extension has to be “invented.” We call the former *natural extension* and the latter *general extension*. Both cases, as argued below, will provide smooth functions \tilde{g} .

In many applications, the weight $g : \mathcal{S} \rightarrow \mathbb{R}$ depends on the curvature structure of the hyper-surface. Denoting $\mathbf{B}_{\mathcal{S}}(\cdot) : \mathcal{S} \rightarrow {}_s\mathbb{R}^{d \times d}$ the second fundamental form of \mathcal{S} , and $\Lambda(\mathbf{B}_{\mathcal{S}}(x))$ the set of its eigenvalues, this means that

$$g(x) = F(\Lambda(\mathbf{B}_{\mathcal{S}}(x))),$$

where F is a given function. In this case it is utterly *natural* to take advantage of the implicit representation by noting that $\mathbf{B}_{\mathcal{S}}(x) = \mathbf{H}_{\psi|_{T_x\mathcal{S}}}(x)$ for $x \in \mathcal{S}$, where \mathbf{H}_{ψ} is the Hessian of ψ and $T_x\mathcal{S}$ is the tangent space to \mathcal{S} at x (see [37]). The *natural* extension then becomes

$$\tilde{g}(x) = F(\Lambda(\mathbf{H}_{\psi|_{T_x\mathcal{S}(x)}}(x))), \quad x \in \Omega_h, \tag{9}$$

where $\mathcal{S}(x) \triangleq \{y \in \mathbb{R}^d : \psi(y) = \psi(x)\}$.

This extension is valid for $\{x \in \mathbb{R}^d : |\psi(x)| < 1/\mathcal{M}_{\mathcal{S}}\}$, where $\mathcal{M}_{\mathcal{S}}$ absolutely bounds all principal curvatures of \mathcal{S} ; see Appendix A.

When the weight g cannot be directly extended to be valid for a tubular neighborhood of the hyper-surface, one has to do that in a pedestrian way. One such extension comes from propagating the values of g along the normals of \mathcal{S} in a constant fashion; i.e.,

$$\tilde{g}(x) = g(\Pi_{\mathcal{S}}(x)), \quad x \in \Omega_h, \tag{10}$$

where $\Pi_{\mathcal{S}}(\cdot) : \mathbb{R}^d \rightarrow \mathcal{S}$ stands for the normal projection onto \mathcal{S} . This extension is well defined and smooth as long as there is a unique *foot* in \mathcal{S} for every x in the domain of the desired extension Ω . Taking h sufficiently small we can guarantee that $\Omega \supset \Omega_h$ if \mathcal{S} is smooth. See Appendix A for some details.

In practice, this extension can be accomplished solving the equation [14]

$$\phi_t + \text{sgn}(\psi)\nabla\psi \cdot \nabla\phi = 0$$

with initial conditions given by any $\phi(\cdot, 0)$, such that $\phi(\cdot, 0)|_{\mathcal{S}} = g$. Then $\tilde{g}(\cdot) \triangleq \phi(\cdot, \infty)$.

2.2. Shortest Paths and Distance Functions in Manifolds with Boundary

Since we want to approximate the problem of intrinsic distance functions by a problem of distance functions in manifolds with boundary, and to prove that the latter converges to the former, we need to review basic concepts on this subject. We will mainly include results from [2, 3, 63]. We are interested in the existence and smoothness of the geodesic curves on manifolds with boundary, since our convergence arguments below depend on these properties. We will assume throughout this section that (\mathcal{M}, m) is a *connected* and *complete* Riemannian manifold with boundary (this will later become the h -offset Ω_h with the metric $\tilde{g}^2\mathbb{I}$, where \mathbb{I} now stands for the $d \times d$ identity matrix).

DEFINITION 2.1 Let $p, q \in \mathcal{M}$, then if $d_{\mathcal{M}}(\cdot, \cdot) : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ is the distance function in \mathcal{M} (with its metric m), a shortest path between p and q is a path joining them such that its Riemannian length equals $d_{\mathcal{M}}(p, q)$.

Now, since \mathcal{M} is *complete*, for every pair of points p and q there exists a *shortest path* joining them; see [2]. The following results deal with the regularity of this shortest path.




FIG. 1. The minimal path is C^1 , but not C^2 .

THEOREM 2.1 *Let (\mathcal{M}, m) be a C^3 manifold with C^1 boundary B . Then any shortest path of \mathcal{M} is C^1 .*

When (\mathcal{M}, m) is a flat manifold (i.e., \mathcal{M} is a codimension 0 subset of \mathbb{R}^d and the metric m is isotropic and constant), it is easy to see that any *shortest path* must be a straight line whenever it is in the interior of \mathcal{M} , and a shortest path of the boundary B when it is there. This will be the situation for us from now on.

It might seem a bit awkward that one cannot achieve a higher regularity class than C^1 for the shortest paths, even by increasing the regularity of $\mathcal{M} \cup B$, but a simple counterexample will convince the reader. Think of \mathcal{M} as \mathbb{R}^2 with the open unit disc removed (see Fig. 1), and its Euclidean metric. The acceleration in all the open segment (AP) is $\vec{0}$, and in all the open arc (PQ) is $-\vec{e}_r$; that is, it points inward and has modulus 1. That is, even in most simple examples, C^2 regularity is not achievable. It is, however, very easy to check that in this case $\dot{\gamma}$ is actually *Lipschitz*.

For the general situation, in [3, 39] the authors proved that shortest paths do have *Lipschitz continuous* first derivatives, which means that in fact shortest paths are twice differentiable *almost everywhere* by Rademacher's Theorem. This fact will be of great importance below.

For a more comprehensive understanding of the theory of shortest paths and distance functions in Riemannian manifolds with boundary, see [2, 3, 39, 63] and references therein.

2.3. Convergence Result for the Extrinsic Distance Function

We now show the relation between the Euclidean distance in the band Ω_h and the intrinsic distance in the surface \mathcal{S} . Below we will denote $d_{\mathcal{S}} \triangleq d_{\mathcal{S}}^1$, and $d_{\Omega_h} \triangleq d_{\Omega_h}^1$.

Observation 2.1. Since we assume the implicit surface \mathcal{S} to be compact, the continuous function $d_{\mathcal{S}} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ attains its maximum. Therefore, we can define the diameter of

the set as

$$\text{diam}(\mathcal{S}) \triangleq \sup_{p,q \in \mathcal{S}} d_{\mathcal{S}}(p, q) < \infty.$$

Observation 2.2. Since $\mathcal{S} \subset \Omega_h$ we have that for every pair of points p and q in \mathcal{S} , $d_{\Omega_h}(p, q) \leq d_{\mathcal{S}}(p, q)$, so in view of the previous observation we have

$$d_{\Omega_h}(p, q) \leq \text{diam}(\mathcal{S}) \quad \forall p, q \in \mathcal{S}.$$

Observation 2.3. Since we are assuming \tilde{g} to be a smooth extension of g to all $\Omega \supset \Omega_h$ (we stress the fact that the extension does not depend on h), \tilde{g} will be Lipschitz in Ω , and we call $K_{\tilde{g}}$ its associated constant. Further, we will denote $M_g \triangleq \max_{\{x \in \mathcal{S}\}} g(x)$ and $M_{\tilde{g}} \triangleq \sup_{\{x \in \Omega\}} \tilde{g}(x)$.

We need the following *Lemma* whose simple proof we omit (see, for example, [18]).

LEMMA 2.1. *When a \tilde{g} -shortest path travels through an interior region, its curvature is absolutely bounded by*

$$B_{\tilde{g}} \triangleq \sup_{\{x \in \Omega\}} \left(\frac{\|\nabla \tilde{g}(x)\|}{\tilde{g}(x)} \right).$$

The following Lemma will be needed in the proof of the theorem below. Its proof can be found in Appendix B.

LEMMA 2.2. *Let $f : [a, b] \rightarrow \mathbb{R}$ be a $C^1([a, b])$ function such that f' is Lipschitz. Let $\varphi \in L^\infty([a, b])$ denote (one of) f' 's weak derivative(s). Then*

$$\int_a^b f'^2(x) dx = f f' \Big|_a^b - \int_a^b f(x) \varphi(x) dx.$$

We are now ready to present one of the main results of this section. We bound the error between the intrinsic distance on \mathcal{S} and the Euclidean one in the offset Ω_h . As we will see below, in the most general case, the error is of the order $h^{1/2}$ (h being half the offset width). We will later discuss that this is also the order of the theoretical error for the numerical approximation in fast marching methods. That will lead us to conclude that our algorithm does keep the convergence rate within the theoretically proven order for fast marching methods' numerical approximation. However, for all practical purposes, the order of convergence in the numerical schemes used by fast marching methods is that of h ; see [49]. We will also argue that for all practical purposes we can guarantee no decay in the overall rate of convergence. We defer the detailed discussion on this to after the presentation of the general bound below.

THEOREM 2.2. *Let A and B be two points on the smooth hyper-surface \mathcal{S} (see Fig. 2). Let $d_h^{\tilde{g}} = d_{\Omega_h}^{\tilde{g}}(A, B)$ and $d_S^g = d_S^g(A, B)$. Then, for points on the surface \mathcal{S} , we have that for sufficiently small h*

$$|d_S^g - d_h^{\tilde{g}}| \leq h^{\frac{1}{2}} C(h) \text{diam}(\mathcal{S}),$$

where $C(h)$ depends on the global curvature structure of \mathcal{S} and on \tilde{g} , and approaches a constant when $h \downarrow 0$ (it does not depend on A nor B , we give a precise form of $C(h)$ in the proof).




FIG. 2. Tubular neighborhood.

Proof. Let $d_h = d_{\Omega_h}(A, B)$; $d_S = d_S(A, B)$; and let $\gamma : [0, d_h]$ denote a Ω_h \tilde{g} -distance minimizing arc-length parameterized path between $A = \gamma(0)$ and $B = \gamma(d_h)$, such that $\|\dot{\gamma}\| = 1$. Let $\delta = \Pi_{\psi}(\gamma) = \gamma - \psi(\gamma)\nabla\psi(\gamma)$ be the orthogonal projection of γ onto S . This curve will be as smooth as γ for small enough h ; see Appendix A. For sufficiently small h , the boundary of Ω_h will be smooth, since S is smooth and no shocks will be generated (see next section and Appendix A). So we can assume that γ is C^1 and that $\dot{\gamma}$ is Lipschitz, since it is a shortest path within a smooth Riemannian manifold with boundary; see Section 2.2 above.

It is evident that (this is a simple but key observation)

$$\mathbf{L}_{\tilde{g}}\{\gamma\} = d_h^{\tilde{g}} \stackrel{(1)}{\leq} d_S^g \stackrel{(2)}{\leq} \mathbf{L}_g\{\delta\},$$

since

- (1) $S \subset \Omega_h$ and $\tilde{g}|_S = g$
- (2) δ need not be a g -shortest path between A and B on S .

We then have

$$\begin{aligned} |d_S^g - d_h^{\tilde{g}}| &\leq |\mathbf{L}_g\{\delta\} - \mathbf{L}_{\tilde{g}}\{\gamma\}| = |\mathbf{L}_{\tilde{g}}\{\delta\} - \mathbf{L}_{\tilde{g}}\{\gamma\}| \\ &\leq \int_0^{d_h} |\tilde{g}(\delta)\|\dot{\delta}\| - \tilde{g}(\gamma)\|\dot{\gamma}\|| dt \\ &\leq \int_0^{d_h} |\tilde{g}(\delta)\|\dot{\delta}\| - \tilde{g}(\delta)\|\dot{\gamma}\|| dt + \int_0^{d_h} |\tilde{g}(\delta)\|\dot{\gamma}\| - \tilde{g}(\gamma)\|\dot{\gamma}\|| dt \\ &= \int_0^{d_h} |g(\delta)\|\dot{\delta}\| - \|\dot{\gamma}\|| dt + \int_0^{d_h} |\tilde{g}(\delta) - \tilde{g}(\gamma)| dt \\ &\leq M_g \int_0^{d_h} \|\dot{\gamma} - \dot{\delta}\| dt + K_{\tilde{g}} \int_0^{d_h} \|\gamma - \delta\| dt \\ &= M_g \int_0^{d_h} \|\nabla\psi(\gamma) \cdot \dot{\gamma} \nabla\psi(\gamma) + \psi(\gamma)\mathbf{H}_{\psi}(\gamma)\dot{\gamma}\| dt + K_{\tilde{g}} \int_0^{d_h} \|\psi(\gamma)\nabla\psi(\gamma)\| dt \\ &\leq M_g \int_0^{d_h} |\nabla\psi(\gamma) \cdot \dot{\gamma}| dt + h M_g \int_0^{d_h} \|\mathbf{H}_{\psi}(\gamma)\dot{\gamma}\| dt + K_{\tilde{g}} h d_h. \end{aligned}$$

We now bound the first two terms at the end of the preceding expression.

1. We first bound the second term in the preceding expression; this will be an ingredient to the bounding of the first term as well. We have

$$\|\mathbf{H}_\psi(\gamma)\dot{\gamma}\| \leq \sup_{\{v:\|v\|=1;p:d(p,S)\leq h\}} \|\mathbf{H}_\psi(p)v\| = \sup_{\{p:d(p,S)\leq h\}} \max(|\lambda(p)|, |\mu(p)|),$$

where $\lambda(p)$ and $\mu(p)$ denote the largest and the smallest eigenvalue of $\mathbf{H}_\psi(p)$, respectively. Now, as we know from Appendix A, the maximum absolute eigenvalue of $\mathbf{H}_\psi(p)$, $K(p)$, is bounded by

$$K(p) \leq \frac{\mathcal{M}_S}{1 - |\psi(p)|\mathcal{M}_S},$$

where \mathcal{M}_S is the maximum absolute eigenvalue of $\mathbf{H}_\psi|_S$; that is,

$$\mathcal{M}_S = \sup_{\{x \in S\}} \max_{\{1 \leq i \leq d\}} |\lambda_i(\mathbf{H}_\psi(x))|,$$

where $\lambda_i(\cdot)$ stands for the i -th eigenvalue of a symmetric matrix.

Then

$$\int_0^{d_h} \|\mathbf{H}_\psi(\gamma(s))\dot{\gamma}(s)\| ds \leq d_h \frac{\mathcal{M}_S}{1 - h\mathcal{M}_S}.$$

2. Let us define the function $f : [0, d_h] \rightarrow \mathbb{R}$, $f(t) = \psi(\gamma(t))$. Then formally $\dot{f}(t) = \nabla\psi(\gamma(t)) \cdot \dot{\gamma}(t)$ and $\ddot{f}(t) = \mathbf{H}_\psi(\gamma(t))[\dot{\gamma}(t), \dot{\gamma}(t)] + \nabla\psi(\gamma(t)) \cdot \ddot{\gamma}(t)$. Since $\dot{\gamma}(\cdot)$ is Lipschitz, and ψ is regular we can guarantee that $\dot{f}(\cdot)$ is also Lipschitz, so $\ddot{f}(\cdot)$ exists almost everywhere. We want to bound

$$\int_0^{d_h} |\dot{f}(t)| dt.$$

We note first that $f(0) = f(d_h) = 0$, and $|f(t)| \leq h$, $|\ddot{f}(t)| \leq \frac{\mathcal{M}_S}{1-h\mathcal{M}_S} + B_{\tilde{g}}$ for almost every $t \in [0, d_h]$. In fact, we have that for those subintervals of $[0, d_h]$ in which the shortest path travels through $\partial\Omega_h$, either $f(t) = h$, or $f(t) = -h$ for the whole subinterval, and therefore $f(t)$ is constant for each subinterval, so $\ddot{f}(t) = 0$ there. On the other hand, when γ is in the interior of Ω_h , it is a \tilde{g} -geodesic, so its acceleration is bounded by $B_{\tilde{g}}$, as we have seen in Lemma 2.1. Therefore, we conclude that $|\ddot{f}(t)| \leq |\mathbf{H}_\psi(\gamma(t))[\dot{\gamma}(t), \dot{\gamma}(t)]| + B_{\tilde{g}}$. Combining all this we have that for almost every $t \in [0, d_h]$,

$$|\dot{f}(t)| - B_{\tilde{g}} \leq \sup_{\{v:\|v\|=1;d(p,S)\leq h\}} |\mathbf{H}_\psi(p)[v, v]| \leq \sup_{\{p:d(p,S)\leq h\}} \max(|\lambda(p)|, |\mu(p)|)$$

and the given bound follows as before.

Applying Cauchy–Schwartz inequality we obtain

$$\int_0^{d_h} |\dot{f}(t)| dt \leq \sqrt{(d_h) \int_0^{d_h} \dot{f}^2(t) dt}.$$

Now using Lemma 2.2:

$$\begin{aligned} \int_0^{d_h} \dot{f}^2(t) dt &= \dot{f} f|_0^{d_h} - \int_0^{d_h} f \ddot{f} dt = - \int_0^{d_h} f \ddot{f} dt \\ &\leq \int_0^{d_h} |f| |\ddot{f}| dt \leq (d_h)h \left(\frac{\mathcal{M}_S}{1 - h\mathcal{M}_S} + B_{\tilde{g}} \right). \end{aligned}$$

Finally,

$$\int_0^{d_h} |\nabla \psi(\gamma) \cdot \dot{\gamma}| dt \leq (d_h)h \sqrt{h \left(\frac{\mathcal{M}_S}{1 - h\mathcal{M}_S} + B_{\tilde{g}} \right)}.$$

Using both computed bounds, we find that

$$|d_S^g - d_h^{\tilde{g}}| \leq \text{diam}(S)\sqrt{h} \left[M_g \sqrt{\frac{\mathcal{M}_S}{1 - h\mathcal{M}_S} + B_{\tilde{g}}} + M_g \sqrt{h} \frac{\mathcal{M}_S}{1 - h\mathcal{M}_S} + K_{\tilde{g}} \sqrt{h} \right]. \quad \blacksquare \tag{11}$$

From the preceding Lemma we obtain:

COROLLARY 2.1 For a given point $q \in S$

$$\|d_{\Omega_h}^{\tilde{g}}|_S(q, \cdot) - d_S^g(q, \cdot)\|_{L^\infty(S)} \xrightarrow{h \downarrow 0} 0.$$

Remark. The rate of convergence obtained with the techniques shown above is of order \sqrt{h} . A quick look over the proof of convergence shows that the term responsible for the $h^{1/2}$ rate is $\int_0^{d_h} |\dot{f}(t) dt$. All other terms have the higher order of h . Suppose we can find a finite collection of (disjoint) intervals $I_i = (a_i, b_i)$ such that $\text{sgn}(\dot{f})$ is constant (f is monotonic) within each I_i , $\cup_{i=1}^N I_i \subseteq [0, d_h]$, where N is the cardinality of that collection of intervals, and $\dot{f}(t) = 0$ for $t \in [0, d_h] \setminus \cup_{i=1}^N I_i$. Then, we could write

$$\begin{aligned} \int_0^{d_h} |\dot{f}(t)| dt &= \sum_{i=1}^N \text{sgn}(\dot{f})|_{(a_i, b_i)} \int_{a_i}^{b_i} \dot{f}(t) dt = \sum_{i=1}^N \text{sgn}(\dot{f})|_{(a_i, b_i)} (f(b_i) - f(a_i)) \\ &= \sum_{i=1}^N |f(b_i) - f(a_i)| \leq \sum_{i=1}^N (|f(b_i)| + |f(a_i)|) \\ &\leq 2Nh \text{ since } f(t) = \psi(\gamma(t)) \text{ and } \gamma(\cdot) \text{ travels through } \Omega_h, \end{aligned}$$

obtaining a higher rate of convergence, h . It is quite convincing that cases where $N = \infty$ can be considered pathological. We then argue that for all practical purposes the rate of convergence achieved is indeed h (at least). Moreover, for simple cases like a sphere (or other convex surfaces), it is very easy to show explicitly that the error is (at least) of order h .⁸ Notwithstanding, we are currently studying the space of surfaces and metrics g for which we can guarantee that $N < \infty$, and advances in this subject will be reported elsewhere.

⁸ In this case, as in the case of convex surfaces, the geodesic is composed of two straight lines inside the band, tangent to its inner boundary, and a geodesic on the inner boundary of the band; see Fig. 1.

This shows that we can approximate the intrinsic distance with the Euclidean one on the offset band Ω_h . Moreover, as we will detail below, the approximation error is of the same order as the theoretical numerical error in fast marching algorithms. Thereby, we can use fast algorithms in Cartesian grids to compute intrinsic distances (on implicit/implicitized surfaces), enjoying their computational complexity without affecting the convergence rate given by the underlying numerical approximation scheme.

3. NUMERICAL IMPLEMENTATION AND ITS THEORETICAL ERROR

In this section we first discuss the simple modification that needs to be incorporated into the (Cartesian) fast marching algorithm in order to deal with Euclidean spaces with manifolds with boundary. We then propose a way of estimating the (now discrete) offset h , and bound the total numerical error of our algorithm, thereby showing our assertion that the error with our algorithm is of the same order as the one obtained with the fast marching algorithm for Cartesian grids (or triangulated 3D surfaces).

As stated before, we are dealing with the numerical implementation of the Eikonal equation inside an open, bounded, and connected domain Ω (this will later become the offset Ω_h). The general equation, when $P(x)$ is the weight (it becomes \tilde{g} for our particular case), is given by

$$\begin{cases} \|\nabla f(x)\| = P(x) \quad \forall x \in \Omega \\ f(r) = 0, \end{cases} \tag{12}$$

with r the seed point. Note that following the results in the previous section, we are now dealing with the Eikonal equation in Euclidean space, and so the Euclidean gradient is used above.

The upwind numerical scheme to be used for this equation is of the form $(\Delta x_1 = \Delta x_2 = \dots = \Delta x_d = \Delta x)$ [49],

$$\begin{cases} \sum_{j=1}^d \max^2(\hat{f}(p) - m_j, 0) = (\Delta x)^2 P^2(p) \\ m_j = \min(\hat{f}(p + \Delta x \vec{e}_j), \hat{f}(p - \Delta x \vec{e}_j)), \end{cases} \tag{13}$$

where \hat{f} is the numerically computed value of f for every point p in the discrete domain

$$\mathcal{D}(\Omega, \Delta x) \triangleq \Omega \cap (Z\Delta x).$$

Here, \vec{e}_j with $j = 1, 2, \dots, d$, are the elements of the canonical basis of \mathbb{R}^d .

We now describe the fast marching algorithm for solving the above equation. For this we follow the presentation in [52]. For clarity we write down the algorithm in *pseudo-code* form. Details on the original fast marching method on Cartesian grids can be found in the mentioned references.

At all times there are three kinds of points under consideration:

- **NarrowBand.** These points have to them associated an already guessed value for \hat{f} , and are immediate neighbors to those points whose value has already been “frozen.”
- **Alive.** These are the points whose \hat{f} value has already been frozen.

- **Far Away.** These are points that haven't been processed yet, so no tentative value has been associated to them. For that reason, they have $\hat{f} = \infty$, forcing them not to be considered as part of the up-winding stencil in the Gudunov's Hamiltonian.

The steps of the algorithm include the following

- *Initialization:*

1. Set $\hat{f} = 0$ for every point belonging to the set [**Alive**]. These are the seed point(s) if they lie on the grid. If the seed is not a grid point, its corresponding **Neighbors**⁹ are set **Alive** and are given an initial value \hat{f} simply computed via interpolation (taking into account the distance from the neighbor grid points to the seed point).

2. Find a tentative value of \hat{f} for every **Neighbor** of an **Alive** point and tag each **NarrowBand**.

3. Set $\hat{f} = \infty$ for all the remaining points in the discrete domain.

- *Advance:*

1. Beginning of loop: Let (p_{min}) be the point \in [**NarrowBand**], which takes the least value of \hat{f} .

2. Insert the point p_{min} to the set [**Alive**] and remove it from [**NarrowBand**].

3. Tag as **Neighbors** all those points in the discrete domain that can be written in the form $p_{min} \pm \Delta x \vec{e}_j$, and belong to [**NarrowBand**] \cup [**FarAway**]. If a **Neighbor** is in [**FarAway**], remove it from that set ([**FarAway**]) and insert it to [**NarrowBand**].

4. Recalculate \hat{f} for all **Neighbors** using Eq. (13)

5. Set [**Neighbor**] = *empty set*.

6. Back to the beginning (step 1).

The boundary conditions are taken such that points beyond the discrete domain have $\hat{f} = \infty$.

The condition that is checked all the time, and that really defines the domain the algorithm is working within, is the one that determines if a certain point q is **Neighbor** of a given point p that belongs to the domain. The only thing one has to do in order to make the algorithm work in the domain Ω_h specified by $\{x \in \mathbb{R}^d : |\psi(x)| \leq h\}$ is change the way the **Neighbor** checking is done. More precisely, we should check

$$q \in \mathbf{Neighbor}(p) \text{ iff } \{(|\psi(q)| \leq h) \ \&\& \ (q \text{ can be written like } p \pm \Delta x \vec{e}_j)\},$$

the emphasis here being on the test “ $|\psi(q)| \leq h$.” We could also achieve the same effect by giving an infinite weight to all points outside Ω_h ; that is, we treat the outside of Ω_h as an obstacle. Therefore, with an extremely simple modification to the fast marching algorithm, we make it work as well for distances on manifolds with boundary, and therefore, for intrinsic distances on implicit surfaces. This is of course supported by the convergence results in the previous section and the analysis on the numerical error presented below.

3.1. Bounding the Offset h

We now present a technique to estimate h , the size of the offset of the hyper-surface \mathcal{S} that actually defines the computational domain Ω_h . The bounds on h are very simple. On one hand, we need h to be large enough so that the upwind scheme can be implemented, meaning that h has to be large enough to include the stencil used in the numerical implementation.

⁹ For a grid point p , any of its 2D-neighboring points can be written like $p \pm \Delta x_i \vec{e}_i$.




FIG. 3. We depict the situation that leads to the lower bound for h in the 2D case. In red: the curve. In black: the centers of $\bar{B}(x \in \mathcal{S}, d^{1/2} \Delta x)$. In green: the points of $\mathcal{D}(\Omega_h, \Delta x)$ that fall inside $\bar{B}(x, d^{1/2} \Delta x)$ for some $x \in \mathcal{S}$, and in blue those that don't.

On the other hand, h has to be small enough to guarantee that Ω_h remains simply connected with smooth boundaries and that \tilde{g} remains smooth inside Ω_h .

Let \mathcal{M}_S be as before a bound for the absolute sectional curvature of \mathcal{S} , and let Δx be the grid size. In addition, let W be the maximal offsetting of the surface \mathcal{S} that guarantees that the resulting set remains connected and different parts of the boundary of that set do not touch each other. We show below that a suitable bounding of h is (recall that d is the dimension of the space),

$$\Delta x \sqrt{d} < h < \min \left\{ \frac{1}{\mathcal{M}_S}, W \right\}. \tag{14}$$

Let us introduce some additional notation. We denote by *cell* the unit cell of the computational grid. Let x be a point in Ω_h ; we denote by $n(x)$ the number of cells $C_1(x), \dots, C_{n(x)}(x)$ that contain x . It is clear that if $x \in \mathcal{D}(\Omega_h, \Delta x)$ (it is a grid point), then x is contained in 2^d cells having x as a vertex. It is also clear that $n(x) \leq 2^d$. For a given cell C we call $\mathcal{P}(C)$ the set of points of $\mathcal{D}(\Omega_h, \Delta x)$ that compose C (i.e., its vertices). We will denote by $\mathbf{C}(x)$ the set $\bigcup_{i=1}^{n(x)} C_i(x)$, and by $\mathcal{P}(x)$ the set $\bigcup_{i=1}^{n(x)} \mathcal{P}(C_i(x))$.

The lower bound comes from forcing that, for every $x \in \mathcal{S}$, all points in $\mathbf{C}(x)$ lie within Ω_h (note of course that we want h to be as small as possible); see Fig. 3. That is,

$$\bigcup_{x \in \mathcal{S}} \mathbf{C}(x) \subset \Omega_h.$$

Once again, this constraint comes in order to guarantee that there are “enough” points to make the discrete calculations. We try to make $\mathbf{C}(x) \subset \bar{C}(x, l)$, where $C(x, l)$ stands for the hypercube centered in x , with side length $2l$, and sides parallel to the gridding directions. The worst scenario is when x is a point in the discrete domain, and we must have $l \geq \Delta x$. Finally, we observe that $C(x, l) \subset B(x, l\sqrt{d})$. The condition then becomes

$$\bigcup_{x \in \mathcal{S}} B(x, \Delta x \sqrt{d}) \subset \Omega_h = \bigcup_{x \in \mathcal{S}} B(x, h),$$

which provides the lower bound, $h > \Delta x \sqrt{d}$.

The upper bound includes two parts. First, we shouldn't go beyond W , since if we do, different parts of the offset surface might touch each other, a situation that can even create a nonsimply connected band Ω_h . The second part of the upper bound comes from seeking that when traveling on a characteristic line of ψ at a point p of \mathcal{S} , no shocks occur inside Ω_h . It is a simple fact that this won't happen if $h < \frac{1}{M_S}$; see Appendix A. It is extremely important to guarantee this both to obtain smooth boundaries for Ω_h and to obtain smooth extensions of the metric g (\tilde{g}).

Note of course that in general, h and also Δx can be position dependent. We can use an adaptive grid, and in places where the curvature of \mathcal{S} is high, or places where high accuracy is desired, we can make Δx small.

3.2. The Numerical Error

It is time now to explicitly bound the numerical error of our proposed method. As stated above, it is our goal to formally show that we are within the same order as the computationally optimal (fast marching) algorithms for computing distance functions on Cartesian grids. Note that the numerical error for the fast marching algorithm on triangulated surfaces has not been reported, although it is of course bounded by the Cartesian one (since this provides a particular "triangulation").

3.2.1. Numerical Error Bound of the Cartesian Fast Marching Algorithm

The aim of this section is to bound a quantity that measures the difference between the numerically computed value $\hat{d}_S^g(p, \cdot)$ and the real value $d_S^g(p, \cdot)$. Any such quantity will compare both functions on \mathcal{S} , but in principle the numerically computed value will not be defined all over the hyper-surface. So we will be dealing with an interpolation stage, that we comment further below in Section 3.2.2.

Let us fix a point $p \in \mathcal{S}$, and let $\hat{f}(\cdot)$ be the numerically computed solution (according to (13)), and $f(\cdot)$ the *real* viscosity solution of the problem (12). The approximation error is then bounded by (see [49])

$$\max_{p \in \mathcal{D}(\Omega, \Delta x)} |\hat{f}(p) - f(p)| \leq C_L (\Delta x)^{\frac{1}{2}}, \quad (15)$$

where C_L is a constant. In practice, however, the authors of [49] observed first-order accuracy. As we have seen, we also find an error of order $h^{1/2}$ for the general approximation of the weighted intrinsic distance on \mathcal{S} with the distance in the band Ω_h , and a practical order of h (see Remark and Theorem 2.2).

Before proceeding with the presentation of the whole numerical error of our proposed algorithm, we need the following simple lemma whose proof we omit.

LEMMA 3.1. *For a convex set $D \subset \Omega$, and $y, z \in D$, f satisfies*

$$|f(z) - f(y)| \leq \|P\|_{L^\infty(\Omega)} \|z - y\|.$$

Remark. Using the preceding Lemma and (15), it is easy to see that for x such that $\mathbf{C}(x) \subset \Omega$,

$$|\hat{f}(p) - \hat{f}(q)| \leq 2C_L (\Delta x)^{\frac{1}{2}} + \|P\|_{L^\infty(\Omega)} \sqrt{d} \Delta x, \quad \forall p, q \in \mathcal{P}(x) \quad (16)$$

a relation we will shortly use.

3.2.2. The Interpolation Error

Since following our approach we are now computing the distance function in the band Ω_h , in the corresponding discrete Cartesian grid, we have to interpolate this to obtain the distance on the zero level-set \mathcal{S} . This interpolation produces a numerical error which we now proceed to bound.

Given the function $\zeta : \mathcal{D}(\Omega, \Delta x) \rightarrow \mathbb{R}$ (Ω being a generic domain, which becomes the band Ω_h for our particular case), we define the function $\mathcal{I}(\zeta) : \Omega \rightarrow \mathbb{R}$ through an interpolation scheme. We will assume that the interpolation error is bounded in the following way:¹⁰

$$\sup_{y \in \mathcal{P}(x)} |\zeta(y) - \mathcal{I}(\zeta)(x)| \leq \max_{z \in \mathcal{P}(x)} \zeta(z) - \min_{z \in \mathcal{P}(x)} \zeta(z)$$

for every $x \in \Omega$.

3.2.3. The Total Error

We now present the complete error (numerical plus interpolation) introduced by our algorithm, without considering the possible error in the computation of \tilde{g} (or in other words, we assume that the weight was already given in the whole band Ω_h).

Let p be a point in \mathcal{S} . We denote by

- $d_{\mathcal{S}}^g(p, \cdot) : \mathcal{S} \rightarrow \mathbb{R}$ the *intrinsic* g -distance function from p to any point in \mathcal{S} .
- $d_h^{\tilde{g}}(p, \cdot) : \Omega_h \rightarrow \mathbb{R}$ the \tilde{g} -distance function from p to any other point in Ω_h .
- $\widehat{d}_h^{\tilde{g}}(p, \cdot) : \mathcal{D}(\Omega_h, \Delta x) \rightarrow \mathbb{R}$ the *numerically computed* value of $d_h^{\tilde{g}}(p, \cdot)$ to any point in the discrete domain.
- $\mathcal{I}(\widehat{d}_h^{\tilde{g}})(p, \cdot) : \mathcal{S} \rightarrow \mathbb{R}$ the result of interpolating $\widehat{d}_h^{\tilde{g}}$ (that's only specified for points in $\mathcal{D}(\Omega_h, \Delta x)$) to points in $\mathbb{R}^d \supset \mathcal{S}$.

The goal is then to bound $\|d_{\mathcal{S}}^g(p, \cdot) - \mathcal{I}(\widehat{d}_h^{\tilde{g}})(p, \cdot)\|_{L^\infty(\mathcal{S})}$, and we proceed to do so now.

Let x be in \mathcal{S} and y in $\mathcal{P}(x)$, then

$$\begin{aligned} |d_{\mathcal{S}}^g(p, x) - \mathcal{I}(\widehat{d}_h^{\tilde{g}})(p, x)| &\leq |d_{\mathcal{S}}^g(p, x) - d_h^{\tilde{g}}(p, x)| + |d_h^{\tilde{g}}(p, x) - d_h^{\tilde{g}}(p, y)| \\ &\quad + |d_h^{\tilde{g}}(p, y) - \widehat{d}_h^{\tilde{g}}(p, y)| + |\widehat{d}_h^{\tilde{g}}(p, y) - \mathcal{I}(\widehat{d}_h^{\tilde{g}})(p, x)|, \end{aligned} \quad (17)$$

and using Proposition 2.2, Lemma 3.1 (15), and simple manipulations (in that order) we obtain

$$\begin{aligned} &|d_{\mathcal{S}}^g(p, x) - \mathcal{I}(\widehat{d}_h^{\tilde{g}})(p, x)| \\ &\leq C(h) \text{diam}(\mathcal{S}) h^{\frac{1}{2}} + M_{\tilde{g}} \|x - y\| + C_L(\Delta x)^{\frac{1}{2}} + \left(\max_{y \in \mathcal{P}(x)} \widehat{d}_h^{\tilde{g}}(p, y) - \min_{y \in \mathcal{P}(x)} \widehat{d}_h^{\tilde{g}}(p, y) \right). \end{aligned}$$

The last term can be dealt with using (16). Since we want both $h \downarrow 0$ and $\frac{h}{\Delta x} \uparrow \infty$, in order to have increasing fidelity in the approximation of $d_h^{\tilde{g}}$ by its numeric counterpart $\widehat{d}_h^{\tilde{g}}$,¹¹ we can choose (for instance) $h = C_x (\Delta x)^\gamma$ for some constant $C_x > \sqrt{d}$ and $\gamma \in (0, 1)$. We

¹⁰ One may imagine several interpolation schemes satisfying this not-stringent-at-all condition.

¹¹ This way, we will have an increasing number of points in Ω_h .

then obtain

$$\|d_S^g(p, \cdot) - \mathcal{I}(\widehat{d}_h^g)(p, \cdot)\|_{L^\infty(\mathcal{S})} \leq (\Delta x)^{\frac{\gamma}{2}} \mathcal{C}(\Delta x; \mathcal{S}), \quad (18)$$

where $\mathcal{C}(\Delta x; \mathcal{S})$ goes to a constant (that depends on \mathcal{S}) as $\Delta x \downarrow 0$, and this provides the desired bound.

To recap, we have obtained that the use of an Euclidean approximation in the band Ω_h to the intrinsic distance function on the level-set \mathcal{S} doesn't (meaningfully) change the order of the whole numerical approximation, in the worst case scenario. While in the most general case the theoretical bound for the error of our method is of order $h^{1/2}$ and the general order of the error of the underlying numerical scheme is $(\Delta x)^{1/2}$, for all practical purposes the approximation error (over \mathcal{S}) between both distances (d_S^g and \widehat{d}_h^g) is of order h (see remark after Corollary 2.1), and the practical numerical error between \widehat{d}_h^g and \widetilde{d}_h^g is of order $(\Delta x)^\beta$ (for some $\beta \in [\frac{1}{2}, 1)$ for our first order schemes). Then, the practical bound for the total error becomes something of order $(\Delta x)^{\min(\beta, \gamma)}$. Therefore, choosing a big enough γ (< 1) dispels any concerns about worsening the overall error rate when doing Cartesian computations on the band.¹²

To conclude, let's point out that since we are working within a narrow band (Ω_h) of the surface \mathcal{S} , we are actually not increasing the dimensionality of the problem. We can then work with a Cartesian grid while keeping the same dimensionality as if we were working on the surface.¹³

4. EXPERIMENTS

We now present a number of 3D examples of our algorithm. Recall that although all the examples are given in 3D, the theory presented above is valid for any dimension $d \geq 3$.

Two classes of experimental results are presented. We show a number of intrinsic distance functions for implicit surfaces, as well as geodesics computed using these functions. In order to compute interesting geodesics, we use also nonuniform weights, permitting the computation of crest/valleys, and optimal paths with obstacles. We also experimentally compare our results with those obtained on triangulated surfaces using the fast marching technique developed by Kimmel and Sethian (for this we use the results and software reported in [6]).

Figures 4, 5, and 6 show the intrinsic distance function for implicit surfaces computed with the method here proposed ($g = 1$). An arbitrary seed point on the implicit surface has been chosen, and pseudo colors are used to improve the visualization. Red corresponds to low values of the distance and blue to the high ones. We observe that, as expected, the distance (colors) vary smoothly, and that close points have similar colors and far points have very different colors (close and far measured on the surface of course).

In Fig. 7 we compare the result of our approach with that of fast marching on a triangulated surface (all triangulated-surface computations were done with the package reported in [6]). We also show absolute differences (error) between distances obtained through both

¹² Note that the numerical scheme used by the fast marching algorithm decreases its accuracy when nondifferentiable points of the distance appear, this can happen for instance when the domain contains the cut locus of the initial set [15]. In any case, $(\Delta x)^{\frac{1}{2}}$ is the slowest error rate achievable.

¹³ The number of points in the band can be roughly estimated by the quantity $2h \text{ area}[\mathcal{S}]$ when $\Delta x = 1$.




FIG. 4. Distance maps from a point on the sphere, torus, and teapot (three views are presented for each model).

approaches. The particular pattern of the error is the subject of future research. In Fig. 8 we show level lines of the intrinsic distance function computed with the technique here proposed.

Before concluding this part of the experiments, let's give some technical details on the implementation. The code for the examples in this paper was written in C++. For visualization purposes, VTK was used. Most of the “hard code” was done taking advantage of Blitz++'s double templated arrays and related routines, see [9]. The implicit models used in this paper were obtained from [67] (other techniques, e.g., [40], could be used as well). All the code was compiled and run in a 450 Mhz Pentium III, with 256 Mb of RAM, working under Linux (RedHat 6.2). The compiler used was `egcs-2.91.66` and the level of optimization was 3. In Table 1 we show running times of the intrinsic distance map algorithm for some of the implicit models we used, along with the corresponding *offset*-value (h) and size and number of grid points in Ω_h for each model.

4.1. Geodesics on Implicit Surfaces

To find geodesic curves on the implicit surface, we backtrack starting from a specified target point toward the seed point, while traveling on the surface in the direction given




FIG. 5. Distance map from a point on a portion of white/gray matter boundary of the cortex.




FIG. 6. Distance map from one seed point on a knot. In this picture we evidence that the algorithm works well for quite convoluted geometries (as long as h is properly chosen). Note how points close in the Euclidean sense but far away in the intrinsic sense receive very different colors, indicating their large (intrinsic) distance.




FIG. 7. Top: Distance map from a single seed point (situated at the nose) on an implicit bunny ($g = 1$). The figure on the top-left was obtained with the implicit approach (d_h) here presented, while the one on the top-right was derived with the fast marching on triangulated surfaces (d_Δ) technique. Bottom: Three views of the absolute difference between both distance functions ($d_h - d_\Delta$). The maximal difference (error between the distances) is 4.1439, being 91.5599 the maximal computed distance in the band.

by the (negative) intrinsic-distance gradient. This means that after we have computed the intrinsic distance function as explained above, we have to solve the following ODE (which obviously keeps the curve on \mathcal{S}):

$$\begin{cases} \dot{\gamma} = -\nabla_\psi d_{\Omega_h}^g(\gamma) \\ \gamma(0) = p \in \mathcal{S}, \end{cases}$$

where

$$\nabla_\psi d_{\Omega_h}^g(p) \triangleq \nabla d_{\Omega_h}^g(p) - (\nabla d_{\Omega_h}^g(p) \cdot \nabla \psi(p)) \nabla \psi(p)$$

TABLE 1

Model	Size	$\#\mathcal{D}(\Omega_h, \Delta x)$	h	Running Time (secs)
Brain	$122 \times 142 \times 124$	168,603	1.75	9.4
Bunny	$81 \times 80 \times 65$	38,107	1.75	1.99
Knot	$80 \times 81 \times 44$	16,095	1.0	0.76
Sphere	$70 \times 70 \times 70$	11,800	1.75	0.65
Torus	$64 \times 64 \times 64$	21,704	1.75	1.16
Teapot	$80 \times 55 \times 46$	24,325	1.75	1.22




FIG. 8. Top: Level lines for the intrinsic distance function depicted in Fig. 7 (left). Bottom: Level lines for the intrinsic distance function depicted in Fig. 4 (second row). In both rows, the (22) levels shown are 0.03, 0.05, 0.1, . . . , 0.95, 0.97 percent of the maximum value of the intrinsic distance, and the coloring of the surface corresponds to the intrinsic distance function. Three views are presented. Note the correct separation between adjacent level lines. Note also how these lines are “parallel.”

is the gradient of $d_{\Omega_h}^g$ at $p \in \mathcal{S}$ projected onto the tangent space to $\mathcal{S} = \{\psi = 0\}$ at p . Since we must discretize the above equation, one can no longer assume that at every instant the geodesic path γ will lie on the surface, so a projection step must be added. In addition, since all quantities are known only at grid points, an interpolation scheme must be used to perform all evaluations at positions given by γ . We have used a simple Runge–Kutta integration procedure, with adaptive step, namely an ODE23 procedure.

Before presenting examples of geodesic curves, we should note that we are assuming that $\nabla_{\psi} d_{\Omega_h}^g$, the extrinsic gradient of the distance in the band, is a good approximation of $\nabla_{\mathcal{S}} d_{\mathcal{S}}^g$, the intrinsic gradient of the intrinsic distance (and not just $d_{\Omega_h}^g$ a good approximation of $d_{\mathcal{S}}^g$ as we have previously proved). Bounding the error between these two gradients, e.g., using the framework of viscosity solutions (since intrinsic distances are not necessarily smooth), is the subject of current work (see also next section for a numerical experiment).

The figures described next illustrate the computation of geodesic curves on implicit surfaces for different weights g . In all the figures the geodesic curve is drawn on top of the surface, which is colored as before, colors indicating the intrinsic weighted distance.

In Fig. 9 we present both the geodesic curve computed with our technique and the one computed with the fast marching algorithm on triangulated surfaces following the implementation reported in [6].

In Fig. 10 we show the computation of sulci (valleys) on an implicit surface representing the boundary between the white and gray matter in a portion of the human cortex (data obtained from MRI). Here the (extended) weight \tilde{g} is a function of the mean curvature given by [6]

$$\tilde{g}_{\text{valley}}(x) = \omega + \left(\mathbf{M}(x) - \min_{y \in \Omega_h} \mathbf{M}(y) \right)^p,$$

where \mathbf{M} stands for the mean curvature of the level sets of ψ , so it is computed simply as $\mathbf{M}(x) = \Delta\psi(x)$. In the example presented we used $\omega = 100$ and $p = 3$. More details on the use of this approach for detecting valleys (and creases) can be found in [6] and in the references therein.

In Fig. 11 we show the computation of geodesic curves with obstacles on implicit surfaces. This is an important computation for topics such as motion planning on surfaces.

4.2. Simple Numerical Accuracy Validations

We conclude the examples with some simple numerical validations. Since for a sphere, for instance, the real distances can be computed, we compare these with those numerically computed with our algorithm. As previously explained, for this case the error of our proposed band-based approximation of the intrinsic (continuous) distance is of order h (actually, it can be shown that the order is slightly superlinear). We have tested the computed distance between given seed points in the sphere for a number of different grid sizes (resolutions) in the cube $[0, 1]^3$, obtaining the errors given in Table 2. In obtaining the data we have used $h = 2(\Delta x)^{0.7}$. It can be observed an overall error

$$\left(\|d_S - \mathcal{I}(\hat{d}_h)\|_{L^\infty(S)}\right)$$

rate given approximately by $(\Delta x)^{0.653}$.

Although a thorough study of the approximation of $\nabla_S d_S^g$ by $\nabla_\psi d_{\Omega_h}^g$ will be the subject of future research, we will present below some numeric evidence. Note of course that the main concern, as previously explained, is at the cut locus (singularities on the gradient of the distance function). To the best of our knowledge, complete analysis of the accuracy of the gradient of the intrinsic distance has not yet been performed for triangulated surfaces.

TABLE 2

Size	h	Overall Numerical Error
100	0.079621	0.111035
120	0.070081	0.101189
140	0.062912	0.090596
160	0.057298	0.084766
180	0.052764	0.077357
200	0.049012	0.072119
220	0.045849	0.069262
240	0.043140	0.064085
260	0.040789	0.061003
280	0.038727	0.057780
300	0.036901	0.056024
320	0.035271	0.053569
340	0.033806	0.051469
360	0.032480	0.048853
380	0.031273	0.047292
400	0.030170	0.046195
420	0.029157	0.044747
440	0.028223	0.043254
460	0.027358	0.041999
480	0.026555	0.040501
500	0.025807	0.039396




FIG. 9. Top: Distance map (weight = 1) and geodesic curve between two points on an implicit bunny. We show two geodesics superimposed, the black one is the one obtained via the implicit back propagation described in the text, while the white one is obtained when performing the back propagation computation in the triangulated surface. It is important to note that in both cases the distance function used is the one computed with our implicit approach; to feed the data to the triangulated surfaces back-propagation algorithm, we first interpolated the intrinsic distance to points onto the triangulated surface. We can clearly see that both geodesics overlap almost entirely, justifying the proposed implicit back propagation approach when compared to the one on the triangulated surface. Bottom: We repeat the top figure, but now for the white curve (\mathcal{C}_Δ) the distance used was also computed on the triangulated surface. In other words, the black curve (\mathcal{C}_h) corresponds to complete implicit computations, both distance and back propagation, while the white one corresponds to complete computations on the triangulated domain. For this particular example, the geodesic obtained with the computations on the implicit surface is actually shorter than the one obtained with computations on the triangulated representation.




FIG. 10. These four figures show the detection of valleys over implicit surfaces representing a portion of the human cortex. We use a mean curvature based weighted distance. In the left-upper corner we show the mean curvature of the brain surface (clipped to improve visualization). It is quite convincing that this quantity can be of great help to detect valleys. In the remaining figures, we show two curves over the surface, whose coloring correspond to the mean curvature (not clipped, from red, yellow, green to blue, as the value increases). The red curve corresponds to the *natural* geodesic ($g = 1$), while the white curve is the weighted-geodesic that should travel through “nether” regions. Indeed, a very clear difference exists between both trajectories, since the white curve makes its way through regions where the mean curvature attains low values. The figure in the right-low quadrant is a zoomed view of the same situation.

We make all our computations again for simplicity, over a sphere, taking $g = 1$ (we will discard the superscripts g and \tilde{g} for the remains of this section). As an indicator of how well $\nabla_{\psi} d_{\Omega_h}$ approximates $\nabla_S d_S$ (over \mathcal{S}) we look at how much the quantity $\|\nabla_{\psi} d_{\Omega_h}\|$ differs from 1. In Fig. 12 we show nine histograms of the aforementioned quantity, for 1000 points on the sphere and for nine (increasing) values of h . It can be observed that the values of $\|\nabla_{\psi} d_{\Omega_h}\|$ spread more and more as h grows.

5. EXTENSIONS

5.1. General Metrics: Solving Hamilton–Jacobi Equations on Implicit Surfaces

Since the very beginning of our exposition we have restricted ourselves to *isotropic* metrics. As stated in the introduction, this already has many applications, and just a few




FIG. 11. Distance map and geodesic curve between two points on an implicit bunny surface with an intrinsic obstacle on it. We now use a binary weight, $g = \{1, \infty\}$, being infinity at the obstacle. This permits, as illustrated in the figure, the computation of optimal paths with obstacles on implicit surfaces. The blue path corresponds to the obstacle-weighted distance function, and the white one to the natural ($g = 1$) distance function. Both geodesics are shown over the surface of the bunny, the pseudocolor representing the weighted distance for the surface with obstacle. The obstacle is also shown in blue. Note that the geodesic is not touching the obstacle due to the low grid resolution used to define it in this example (low resolution which makes it actually not a binary but a multivalued obstacle).

were shown in the previous section. Since the fast marching approach has been recently extended to more general Hamilton–Jacobi equations by Osher and Helmsen [45], we are immediately tempted to extend our framework to these equations as well. These equations have applications in important areas such as adaptive mesh generation on manifolds, [28], and semiconductors manufacturing.

Then, we are led to investigate the extension of our algorithm to general metrics of the form, $\mathbf{G} : \mathcal{S} \rightarrow \mathbb{R}^{d \times d}$, that is, a positive definite 2-tensor. Our new definition of weighted length becomes

$$\mathbf{L}_{\mathbf{G}}\{C\} \triangleq \int_a^b \sqrt{\mathbf{G}(C(t))[\dot{C}(t), \dot{C}(t)]} dt,$$

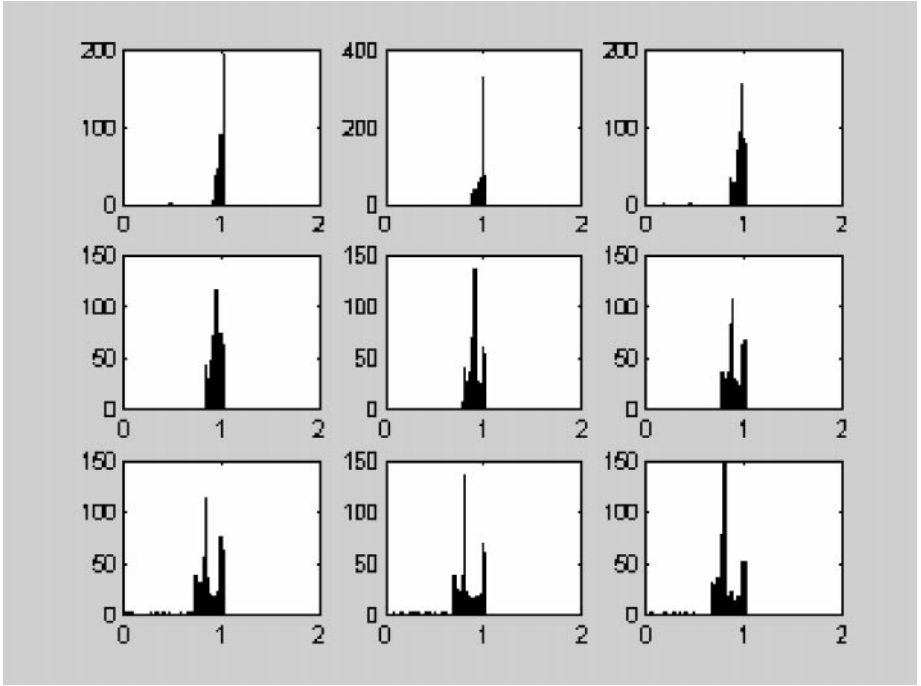


FIG. 12. Histograms of $\|\nabla_{\psi} d_{\Omega_h}\|$ for several (increasing) values of h , for 1000 points uniformly distributed on a sphere. From left to right and top to bottom, the histograms are plotted for increasing values of h .

and the problem is to find for every $x \in S$ (for a fixed $p \in S$),

$$d_S^G(x, p) \triangleq \inf_{C_{px}[S]} \{\mathbf{L}_G(C)\}. \tag{19}$$

As before, we attempt to solve the approximate problem in the band Ω_h , with an extrinsic distance

$$d_{\Omega_h}^{\tilde{G}}(x, p) \triangleq \inf_{C_{px}[\Omega_h]} \{\mathbf{L}_{\tilde{G}}(C)\}, \tag{20}$$

where

$$\mathbf{L}_{\tilde{G}}\{C\} \triangleq \int_a^b \sqrt{\tilde{G}(C(t))[\dot{C}(t), \dot{C}(t)]} dt$$

for an adequate extension \tilde{G} of G . The solution of the extrinsic problem satisfies (in the viscosity sense) the Eikonal equation

$$(\tilde{G}^{-1})(x) [\nabla d_{\Omega_h}^{\tilde{G}}, \nabla d_{\Omega_h}^{\tilde{G}}] = 1. \tag{21}$$

The first issue now is the numerical solvability of the preceding equation using a fast marching type of approach. Osher and Helmsen [45] have extended the capabilities of the fast marching to deal with Hamilton–Jacobi equations of the form

$$H(x, \nabla f) = a(x)$$

for geometrically based Hamiltonians $H(x, \vec{p}) : \Omega(\subset \mathbb{R}^d) \times \mathbb{R}^d \rightarrow \mathbb{R}$ that satisfy

$$\begin{aligned} H(x, \vec{p}) &> 0 && \text{if } \vec{p} \neq \vec{0} \\ H(x, \vec{p}) &&& \text{is homogeneous of degree 1 in } \vec{p} \\ p_i H_{p_i}(x, \vec{p}) &\geq 0 && \text{for } 1 \leq i \leq d \quad \forall x \in \Omega, \forall \vec{p}. \end{aligned} \quad (22)$$

It easily follows that these conditions hold for (21) considering

$$H(x, \vec{p}) \triangleq \sqrt{(\tilde{\mathbf{G}}^{-1})(x)[\vec{p}, \vec{p}]},$$

when the matrix $\tilde{\mathbf{G}}^{-1}(x)$ is diagonal. Therefore, we can solve this kind of Hamilton–Jacobi equations (the extrinsic problem) with the extended fast marching algorithm.

In order to show that our framework is valid for these equations as well, all what we basically need to do is to prove that the extrinsic distance (20) on the offset Ω_h converges to the intrinsic one on the implicit surface \mathcal{S} , i.e., (19). This can be done repeating the steps in the convergence proof previously reported in Section 2.3 for isotropic metrics. Combining this with the results of Osher and Helmsen we then obtain that our framework can be applied to a larger class of Hamilton–Jacobi equations: general intrinsic Eikonal equations. The extension of these ideas to even more general intrinsic Hamilton–Jacobi equations of the form $\hat{H}(x, \nabla_S u) = \hat{a}(x)$ $x \in \mathcal{S}$ remains to be studied, and eventual advances will be reported elsewhere.

5.2. Nonimplicit Surfaces

The framework we presented was here developed for implicit surfaces, although it applies to other surface representations as well. First, if the surface is originally given in polygonal or triangulated form, or even as a set of unconnected points and curves, we can use a number of available techniques, e.g., [34, 40, 47, 55, 58, 65, 67] (and some very nice public domain software [40]), to first implicitize the surface and then apply the technique here proposed.¹⁴ Note that the implicitation needs to be done only once per surface as a preprocessing step and will remain valid for all subsequent uses of the surface. This is important, since many applications have been shown to benefit from an implicit surface representation. Moreover, as we have seen, all what we need is to have a Cartesian grid in a small band around the surface \mathcal{S} . Therefore, there is no explicit need to perform an implicitation of the given surface representation. For example, if the surface is given by a cloud of unconnected points, we can compute distances intrinsic to the surface defined by this cloud, as well as intrinsic geodesic curves, without explicitly computing the underlying surface. All that is needed is to embed this cloud of points in a Cartesian grid and consider only those points in the grid at a distance h or less from the points in the cloud. The computations are then done on this band.

6. CONCLUDING REMARKS

In this article we have presented a novel computationally optimal algorithm for the computation of intrinsic distance functions and geodesics on implicit hyper-surfaces. The

¹⁴ The same techniques can be applied to transform any given implicit function into a distance one.

underlying idea is based on using the classical Cartesian fast marching algorithm in an offset bound around the given surface. We have provided theoretical results justifying this approach and presented a number of experimental examples. The technique can also be applied to 3D triangulated surfaces, or even surfaces represented by clouds of unconnected points, after these have been embedded in a Cartesian grid with proper boundaries. We have also discussed that the approach is valid for more general Hamilton–Jacobi equations as well.

Many questions remain open. Recently, T. Barth (and independently D. Chopp) have shown techniques to improve the order of accuracy of fast marching methods. It will be interesting to see how the method proposed here can be extended to match such accuracy. Related to this, we are currently working on tighter bounds for the error between $d_{\Omega_n}^g$ and d_S^g , as well as bounds for the error between their corresponding derivatives. We are interested in extending the framework presented here to the computation of distance functions on high codimension surfaces and general embeddings. More generally, it remains to be seen what class of intrinsic Hamilton–Jacobi (or in general, what class of intrinsic PDEs) can be approximated with equations in the offset band Ω_n . In an even more general approach, what kind of intrinsic equations can be approximated by equations in other domains, with offsets being just a particular and important example. Even if fast marching techniques do not exist for these equations, it might be simpler and even more accurate to solve the approximating equations in these domains than in the original surface \mathcal{S} . The framework here presented then not only offers a solution to a fundamental problem, but also opens the doors to a new area of research.

APPENDIX A: DISTANCE MAPS IN EUCLIDEAN SPACE

We now present a few important results on distance maps. These have been mainly adapted (and adopted) from [4, 5, 26, 54].

Wherever ψ is smooth we know that it satisfies the *Eikonal* equation

$$\|\nabla\psi\| = 1. \tag{A.1}$$

The distance function satisfies this PDE everywhere in the *viscosity sense* [29, 20]. It is also well known that within a sufficiently small neighborhood of $\mathcal{S} = \{\psi = 0\}$, $\psi(\cdot)$ is *smooth*, or at least as smooth as \mathcal{S} . These assertions can be made precise through the following Lemma from [24]:

LEMMA A.1. *Let \mathcal{S} be a C^k ($k \geq 2$) codimension 1 closed hyper-surface of \mathbb{R}^d . Then, the signed distance function to \mathcal{S} is $C^k(U)$ for a certain neighborhood U of \mathcal{S} .*

Differentiating $\|\nabla\psi\|^2 = 1$, we obtain

$$D(\nabla\psi)\nabla\psi = 0.$$

Therefore,

$$\mathbf{H}_\psi \nabla\psi = 0 \tag{A.2}$$

meaning that the normal to \mathcal{S} at p is an eigenvector of the Hessian, associated to the null eigenvalue. Differentiating again we obtain

$$D^3\psi \nabla\psi + (D^2\psi)^2 = 0. \tag{A.3}$$

The next Lemma, whose detailed proof can be found in [4], is mainly based in the relations (A.2) and (A.3), and it is used to verify that the function $\mu : (-\varepsilon, \varepsilon) \rightarrow \mathbb{R}^{d \times d}$ defined by $\mu(t) = \mathbf{H}_\psi(p_0 + t\nabla\psi(p_0))$ (p_0 is any point in the manifold $\{\psi = 0\}$) satisfies the following **ODE**:

$$\dot{\mu}(t) + \mu^2(t) = 0 \quad t \in (-\varepsilon, \varepsilon)$$

LEMMA A.2. *The eigenvectors of \mathbf{H}_ψ are constant along the characteristic lines $x(s) = x_0 + s\nabla\psi(x(s))$ (arc length parametrized, x_0 is a point onto \mathcal{S}) of ψ within any neighborhood where it is smooth, and the eigenvalues vary according to*

$$\lambda_i(s) = \frac{\lambda_i(0)}{s \lambda_i(0) + 1}.$$

We use the above formula to bound the maximum offset $|\epsilon|$ of $\{\psi = 0\}$ that keeps $\{\psi = \epsilon\}$ smooth, we just take $|\epsilon|(\max_{1 \leq i \leq d-1} |\lambda_i(0)|) < 1$.

We now obtain bounds on the eigenvalues of the Hessian of the distance function:

COROLLARY A.1. *The eigenvalues $\lambda_i(p)$ of $\mathbf{H}_\psi(p)$ (principal curvatures of $\{x : \psi(x) = \psi(p)\}$) are absolutely bounded by*

$$|\lambda_i(p)| \leq \frac{\mathcal{M}_\mathcal{S}}{1 - |\psi(p)|\mathcal{M}_\mathcal{S}},$$

where $\mathcal{M}_\mathcal{S}$ absolutely bounds all eigenvalues of $\mathbf{H}_\psi(p)$, $p \in \mathcal{S}$; and $|\psi(p)|$ is sufficiently small.

To conclude, let's present some concepts on projections onto the implicit surface \mathcal{S} , zero level-set of the distance function ψ . It is clear that the projection of a point $p \in \mathbb{R}^d$ onto \mathcal{S} is given by

$$\Pi_\mathcal{S}(p) = p - \psi(p)\nabla\psi(p).$$

This projection is well defined as long as there is only one $x \in \mathcal{S}$ such that $\Pi_\mathcal{S}(p) = x$. This can be guaranteed when working within a small tubular neighborhood of a smooth surface \mathcal{S} . Moreover, this map is smooth within a certain tubular neighborhood of \mathcal{S} [54]:

THEOREM A.1. *If $\mathcal{S} \subset \mathbb{R}^d$ is a compact C^k ($k \geq 2$) codimension 1 hyper-surface, then there is a $h(\mathcal{S}) > 0$ such that the map $\Pi_\mathcal{S}$ is well defined and belongs to $C^{k-1}(\{x : d(x, \mathcal{S}) < h\}, \mathbb{R}^d)$.*

APPENDIX B: TECHNICAL LEMMA

LEMMA B.1. *Let $f : [a, b] \rightarrow \mathbb{R}$ be a $C^1([a, b])$ function such that f' is Lipschitz. Let $\varphi \in L^\infty([a, b])$ denote (one of) f'' 's weak derivative. Then one has*

$$\int_a^b f'^2(x) dx = ff'|_a^b - \int_a^b f(x)\varphi(x) dx$$

Proof. Let $\text{ext}(f')$ denote the Lipschitz extension of f' to all \mathbb{R} given by

$$\text{ext}(f')(x) = \begin{cases} f'(a) & \text{for } x < a \\ f'(x) & \text{for } x \in [a, b] \\ f'(b) & \text{for } x > b \end{cases}$$

Then let $\text{ext}(f)$ be given by any (bounded) primitive of $\text{ext}(f')$, that is, $\text{ext}(f) = \int \text{ext}(f')$. Let $\hat{\varphi} \in L^\infty \mathbb{R}$ denote $\text{ext}(f')$'s weak derivative, and we have that $\hat{\varphi}|_{[a,b]}$ and φ coincide as weak derivatives of f . Let $\{\eta_\epsilon(\cdot)\}_{\{\epsilon>0\}}$ be a family of bounded support mollifiers. Then we define the function

$$f_\epsilon = \text{ext}(f) * \eta_\epsilon.$$

It is clear that we will have ($\xrightarrow{\epsilon \downarrow 0}$ means uniform convergence)

(a)

$$f_\epsilon \xrightarrow{\epsilon \downarrow 0} \text{ext}(f) \text{ over compact sets of } \mathbb{R}$$

(b)

$$f'_\epsilon \xrightarrow{\epsilon \downarrow 0} \text{ext}(f') \text{ over compact sets of } \mathbb{R}$$

(c)

$$f''_\epsilon \xrightarrow{\epsilon \downarrow 0} \hat{\varphi} \text{ locally in } L^2(\mathbb{R})$$

Since $f'_\epsilon \in C^\infty(\mathbb{R})$, we can use integration by parts to conclude that

$$\int_a^b f'_\epsilon{}^2(x) dx = f'_\epsilon f_\epsilon \Big|_a^b - \int_a^b f_\epsilon(x) f''_\epsilon(x) dx.$$

Now the left-hand side will converge to $\int_a^b f'^2(x) dx$ in view of (b); the first term in the right-hand side will obviously converge to $f f'|_a^b$. For the remaining term we observe the following, using Cauchy–Schwartz inequality (let $\langle \cdot, \cdot \rangle : L^2([a, b]) \times L^2([a, b]) \rightarrow \mathbb{R}$ denote $L^2([a, b])$'s internal product):

$$\begin{aligned} & \left| \int_a^b f_\epsilon(x) f''_\epsilon(x) dx - \int_a^b f(x)\varphi(x) dx \right| \\ &= |\langle f''_\epsilon, f_\epsilon \rangle - \langle \varphi, f \rangle| = |\langle f''_\epsilon, f_\epsilon - f \rangle + \langle f, f''_\epsilon - \varphi \rangle| \\ &\leq (b - a) \left(\left(\max_{\{x \in [a,b]\}} |f''_\epsilon(x)| \right) \|f_\epsilon - f\|_{L^2([a,b])} \right. \\ &\quad \left. + \left(\max_{\{x \in [a,b]\}} |f(x)| \right) \|f''_\epsilon - \varphi\|_{L^2([a,b])} \right). \end{aligned}$$

Now, everything is under control since

$$\max_{\{x \in [a,b]\}} |f''_\epsilon(x)| \leq \|\varphi\|_{L^\infty([a,b])}.$$

Hence, we have proved

$$\int_a^b f_\epsilon(x) f_\epsilon''(x) dx \xrightarrow{\epsilon \downarrow 0} \int_a^b f(x) \varphi(x) dx$$

the last step of the proof. ■

ACKNOWLEDGMENTS

Many colleagues and friends helped us during the course of this work and they deserve our most sincere acknowledgment. Prof. Miguel Paternain, in the early stages of this work, pointed out very relevant references. Prof. Stephanie Alexander, Prof. David Berg, and Prof. Richard Bishop helped us with our questions about geodesics on manifolds with boundaries. Prof. Stanley Osher helped us with issues regarding the fast marching algorithm in general, and his recent extensions in particular, while Prof. Ron Kimmel discussed with us topics related to his work on triangulated surfaces and provided general comments on the paper. Prof. Osher also helped with the literature on rate of convergence of Hamilton–Jacobi numerical approximations. Prof. Omar Gil helped us with some deep insights into technical aspects of the work. This work benefited from rich conversations with Alberto Bartesaghi (who also provided us with some of his code to extract geodesics in triangulated surfaces), Cesar Perciante, Luis Vásquez, Federico Lecumberry, and Prof. Gregory Randall. We also enjoyed important discussion with Prof. Vicent Caselles on the density of zeros of general functions. We also thank Prof. Luigi Ambrosio for his comments on the regularity of level-sets of distance functions. We are indebted to Alvaro Pardo for his careful reading of the manuscript and for his suggestions. We thank Prof. L. T. Cheng and Richard Tsai for their comments. Prof. H. Zhao provided some of the implicit surfaces. FM in particular wants to thank Omar Gil, who has been a great and patient teacher, and Gregory Randall for his constant support and encouragement. FM performed part of this work while visiting the University of Minnesota. This work was partially supported by a grant from the Office of Naval Research ONR-N00014-97-1-0509, the Office of Naval Research Young Investigator Award, the Presidential Early Career Awards for Scientists and Engineers (PECASE), a National Science Foundation CAREER Award, by the National Science Foundation Learning and Intelligent Systems Program (LIS), by the Comisión Sectorial de Investigación Científica (CSIC), the Comisión Académica de Posgrado (CAP), and Tecocom through a scholarship to FM.

REFERENCES

1. D. Adalsteinsson and J. A. Sethian, A fast level set method for propagating interfaces, *J. Comput. Phys.* **118**, 269 (1995).
2. R. Alexander and S. Alexander, Geodesics in Riemannian manifolds with boundary, *Indiana Univ. Math. J.* **30**, 481 (1981).
3. S. Alexander, I. D. Berg, and R. L. Bishop, The Riemannian obstacle problem, *Illinois J. Math.* **31**, 167 (1987).
4. L. Ambrosio and N. Dancer, *Calculus of Variations and Partial Differential Equations*, Topics on Geometrical Evolution Problems and Degree Theory (Springer-Verlag, New York, 2000).
5. L. Ambrosio and H. M. Soner, Level set approach to mean curvature flow in arbitrary codimension, *J. Diff. Geom.* **43**, 693 (1996).
6. A. Bartesaghi and G. Sapiro, A system for the generation of curves on 3D brain images, *Hum. Brain Mapping* **14**, 1 (2001).
7. T. J. Barth and J. A. Sethian, Numerical schemes for the Hamilton–Jacobi and level set equations on triangulated domains, *J. Comput. Phys.* **145**, 1 (1998).
8. M. Bertalmio, L. T. Cheng, S. Osher, and G. Sapiro, Variational problems and partial differential equations on implicit surfaces: The framework and examples in image processing and pattern formation, IMA University of Minnesota, preprint, June 2000.
9. Blitz++ website: available at www.oonumerics.org/blitz.
10. J. Bloomenthal, *Introduction to Implicit Surfaces* (Kaufmann, San Francisco, California, 1997).
11. M. Boue and P. Dupuis, Markov chain approximation for deterministic control problems with affine dynamics and quadratic cost in the control, *SIAM J. Numer. Anal.* **36**, 667 (1999).

12. A. M. Bruckstein, On shape from shading, *Comp. Vision Graph. Image Proc.* **44**, 139 (1988).
13. V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert, Minimal surfaces based object segmentation, *IEEE-PAMI* **19**, 394 (April 1997).
14. S. Chen, B. Merriman, S. Osher, and P. Smereka, "A simple level set method for solving Stefan problems," *J. Comput. Phys.* **135**, 8 (1995).
15. D. L. Chopp, Some improvements of the fast marching method, *SIAM J. Sci. Comput.*, in press.
16. L. T. Cheng, *The Level Set Method Applied to Geometrically Based Motion, Material Science, and Image Processing*, Ph.D. Thesis Dissertation (UCLA, June 2000).
17. L. T. Cheng, P. Burchard, B. Merriman, and S. Osher, *Motion of Curves Constrained on Surfaces Using a Level Set Approach*, UCLA CAM Report 00-32 (September 2000).
18. L. D. Cohen and R. Kimmel, Global minimum for active contours models: A minimal path approach, *Int. J. Comput. Vis.* **24**, 57 (1997).
19. M. G. Crandall and P. L. Lions, Two approximations of solutions of Hamilton–Jacobi equations, *Math. Comp.* **43**, 1 (1984).
20. M. G. Crandall and P. L. Lions, Viscosity solutions of Hamilton–Jacobi equations, *Trans. Am. Math. Soc.* **277** (1983).
21. E. Dijkstra, A note on two problems in connection with graphs, *Numerische Math.* **1**, 269 (1959).
22. M. Desbrun, M. Meyer, P. Schröder, and A. Barr, Discrete differential-geometry operators in nD , Caltech, USC Report (July 22, 2000).
23. S. Fomel, A variational formulation of the fast marching Eikonal solver, Stanford Exploration Project, Report SERGEY, May 1 (May 2000), pp. 357–376.
24. R. L. Foote, Regularity of the distance function, *Proc. Am. Math. Soc.* **92** (September 1984).
25. S. F. Frisken, R. N. Perry, A. Rockwood, and T. Jones, Adaptively sampled fields: A general representation of shape for computer graphics, in *Computer Graphics (SIGGRAPH)* (New Orleans, July 2000).
26. J. Gomes and O. Faugeras, Representing and evolving smooth manifolds of arbitrary codimension embedded in \mathbb{R}^n as the intersection of n hypersurfaces: The vector distance functions, RR-4012 INRIA (Oct. 2000).
27. J. Helmsen, E. G. Puckett, P. Collela, and M. Dorr, Two new methods for simulating photolithography development in 3D, *Proc. SPIE Microlithography* **IX**, 253 (1996).
28. P. Hoch and M. Rasclé, Hamilton–Jacobi equations on a manifold and applications to grid generation or refinement, available at <http://www-math.unice.fr/~hoch/psfiles/hamja.ps>.
29. H. Ishii, A simple direct proof of uniqueness for solutions of the Hamilton–Jacobi equations of Eikonal type, *Proc. Amer. Math. Soc.* **100**, 247 (1987).
30. N. Khaneja, M. I. Miller, and U. Grenander, Dynamic programming generation of geodesics and sulci on brain surfaces, *IEEE Trans. Pattern Anal. Mach. Intelligence* **20**, 1260 (1998).
31. R. Kimmel, *Numerical Geometry of Images: Theory, Algorithms, and Applications*, Technion CIS Report 9910 (October 1999).
32. R. Kimmel and J. A. Sethian, Computing geodesic paths on manifolds, *Proc. Nat. Acad. Sci.* **95**, 8431 (1998).
33. N. Kiryati and G. Székely, Estimating shortest paths and minimal distances on digitized three dimensional surfaces, *Pattern Recognition* **26**, 1623 (1993).
34. V. Krishnamurthy and M. Levoy, Fitting smooth surfaces to dense polygon meshes, *Comput. Graph.* 313 (1996).
35. J. C. Latombe, *Robot Motion Planning* (Kluwer Academic, Boston, 1991).
36. F. Lafon and S. Osher, High order two dimensional nonoscillatory methods for solving Hamilton–Jacobi scalar equations, *J. Comput. Phys.* **123**, 235 (1996).
37. J. M. Lee, *Riemannian Manifolds, An Introduction to Curvature* (Springer-Verlag, New York, 1997).
38. C. Mantegazza and A. C. Menucci, Hamilton–Jacobi equations and distance functions on Riemannian manifolds, available at <http://cvgmt.sns.it/papers/manmen99/>.
39. A. Marino and D. Scolorzi, Geodetiche con ostacolo, *Boll. Un. Mat. Ital.* **6** (1983).
40. S. Mauch, Closest point transform, available at www.ama.caltech.edu/~seanm/software/cpt/cpt.html.

41. F. Memoli, G. Sapiro, and S. Osher, Harmonic maps onto implicit manifolds, preprint, March 2001.
42. J. S. B. Mitchell, An algorithmic approach to some problems in terrain navigation, *Artif. Intelligence* **37**, 171 (1988).
43. J. S. B. Mitchell, D. Payton, and D. Keirse, Planning and reasoning for autonomous vehicle control, *Int. J. Intelligent Syst.* **2**, 129 (1987).
44. S. Osher, A level-set formulation for the solution of the Dirichlet problem for Hamilton–Jacobi equations, *SIMA J. Numer. Anal.* **24**, 1145 (1993).
45. S. J. Osher and R. P. Fedkiw, Level set methods, UCLA CAM Report 00-07 (February 2000).
46. S. J. Osher and J. A. Sethian, Fronts propagation with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations, *J. Comput. Phys.* **79**, 12 (1988).
47. D. Peng, B. Merriman, S. Osher, H. Zhao, and M. Kang, A PDE-based fast local level set method, *J. Comput. Phys.* **155**, 410 (1999).
48. F. P. Preparata and M. I. Shamos, *Computational Geometry*, Texts and Monographs in Computer Science (Springer-Verlag, New York, 1990).
49. E. Rouy and A. Tourin, A viscosity solutions approach to shape-from-shading, *SIAM. J. Numer. Anal.* **29**, 867 (1992).
50. T. Sakai, *Riemannian Geometry*, AMS Translations of Mathematical Monographs **149** (Providence, 1996).
51. J. Sethian, Fast marching level set methods for three-dimensional photolithography development, in Proc. SPIE International Symposium on Microlithography, *Santa Clara, California, March 1996*.
52. J. A. Sethian, A fast marching level-set method for monotonically advancing fronts, *Proc. Nat. Acad. Sci.* **93**, 1591 (1996).
53. J. A. Sethian, *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Materials Sciences* (Cambridge Univ. Press, Cambridge, UK, 1996).
54. L. Simon, *Theorems on Regularity and Singularity of Energy Minimizing Maps* (Birkhäuser, Boston, 1996).
55. G. Taubin, Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation, *IEEE Trans. PAMI* **13**, 1115 (1991).
56. A. W. Toga, *Brain Warping* (Academic Press, New York, 1998).
57. P. Thompson, R. Woods, M. Mega, and A. Toga, Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain, *Human Brain Mapping* **9**, 81 (2000).
58. Y. Tsai, “Rapid and accurate computation of the distance function using grids,” UCLA CAM Report 00-36, October 2000, available at www.math.ucla.edu.
59. J. N. Tsitsiklis, Efficient algorithms for globally optimal trajectories, *IEEE Trans. Autom. Control* **40**, 1528 (1995).
60. D. C. Van Essen, H. Drury, S. Joshi, and M. I. Miller, Functional and structural mapping of human cerebral cortex: Solutions are in the surfaces, *Proc. Nat. Acad. Sci.* **95**, 788 (February 1998).
61. B. Wandell, S. Chial, and B. Backus, Cortical visualization, *J. Cognit. Neurosci.* **12**(5), 739 (2000).
62. A. Witkin and P. Heckbert, Using particles to sample and control implicit surfaces, *Computer Graphics (SIGGRAPH)*, pp. 269–278, 1994.
63. F. E. Wolter, *Cut Loci in Bordered and Unbordered Riemannian manifolds*, Doctoral Dissertation (Technische Universität Berlin, 1985).
64. Z. Wood, M. Desburn, P. Schröder, and D. Breen, Semi-Regular mesh extraction from volume, in *Proc. IEEE Visualization 2000*.
65. G. Yngve and G. Turk, “Creating smooth implicit surfaces from polygonal meshes, Technical Report GIT-GVU-99-42 (Graphics, Visualization, and Usability Center. Georgia Institute of Technology, 1999), available at www.cc.gatech.edu/gvu/geometry/publications.html.
66. A. Yezzi, S. Kichenassamy, P. Olver, and A. Tannenbaum, Geometric active contours for segmentation of medical imagery, *IEEE Trans. Medical Imaging* **16**, 199 (1997).
67. H. Zhao, S. Osher, B. Merriman, and M. Kang, Implicit, non-parametric shape reconstruction from unorganized points using a variational level set method, *Comp. Vision Image Understanding* **80**, 295 (2000).
68. G. Zigelman, R. Kimmel, and N. Kiryati, *Texture Mapping using Surface Flattening via Multidimensional Scaling*, Technion-CIS Technical Report 2000-01 (2000).