

**ON GEOMETRIC VARIATIONAL MODELS FOR
INPAINTING SURFACE HOLES**

By

Vicent Caselles

Gloria Haro


Guillermo Sapiro

and

Joan Verdera

IMA Preprint Series # 2089

(January 2006)



INSTITUTE FOR MATHEMATICS AND ITS APPLICATIONS

UNIVERSITY OF MINNESOTA
400 Lind Hall
207 Church Street S.E.
Minneapolis, Minnesota 55455-0436

Phone: 612/624-6066 Fax: 612/626-7370

URL: <http://www.ima.umn.edu>

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JAN 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE On Geometric Variational Models for Inpainting Surface Holes (PREPRINT)				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Minnesota, Institute for Mathematics and its Applications, 207 Church Street SE, Minneapolis, MN, 55455-0436				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Geometric approaches for filling-in surface holes are introduced and studied in this paper. The basic idea is to represent the surface of interest in implicit form, and fill-in the holes with a scalar, or systems of, geometric partial differential equations often derived from optimization principles. These equations include a system for the joint interpolation of scalar and vector fields, a Laplacian-based minimization, a mean curvature diffusion flow, and an absolutely minimizing Lipschitz extension. The theoretical and computational framework, as well as examples with synthetic and real data, are presented in this paper.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 45	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

On Geometric Variational Models for Inpainting Surface Holes

V.Caselles¹, G.Haro,^{1,2} G.Sapiro,³ and J.Verdera¹

Corresponding author: Gloria Haro

Full postal address:

Institute for Mathematics and its Applications

University of Minnesota

400 Lind Hall

207 Church Street S.E.

Minneapolis, MN 55455-0436

email: **haro@ima.umn.edu** Phone: (612) 626-1501 Fax: (612) 626-7370

Affiliations:

¹ Dept. de Tecnologia, University of Pompeu-Fabra, Passeig de Circumvalacio, 8,
08003 Barcelona, Spain, {vicent.caselles, gloria.haro, joan.verdera}@upf.edu

² Institute for Mathematics and its Applications, University of Minnesota, Minneapolis,
MN 55455, USA, haro@ima.umn.edu

³ Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA, guille@ece.umn.edu

Abstract

Geometric approaches for filling-in surface holes are introduced and studied in this paper. The basic idea is to represent the surface of interest in implicit form, and fill-in the holes with a scalar, or systems of, geometric partial differential equations, often derived from optimization principles. These equations include a system for the joint interpolation of scalar and vector fields, a Laplacian-based minimization, a mean curvature diffusion flow, and an absolutely minimizing Lipschitz extension. The theoretical and computational framework, as well as examples with synthetic and real data, are presented in this paper.

Keywords: Inpainting, variational formulations, interpolation, surface holes, scalar and vector fields, Laplacian, mean curvature, absolute minimizing Lipschitz extension.

1 Introduction

Inpainting is a term used in art to denote the modification of images (painting, photographs, etc) in a form that can not be detected by an ordinary observer. It normally refers to the filling-in of regions of missing information or the replacement of regions by a different kind of information. This is a very important topic in image processing, with applications including image coding and wireless image transmission (e.g., recovering lost blocks), special effects (e.g., removal of objects), and image restoration (e.g., scratch removal). The basic idea behind the computer algorithms that have been proposed in the literature is to fill-in these regions with available information from their surroundings. This information can be automatically detected as in [12, 25], or hinted by the user as in more classical texture filling techniques [22, 26, 38]. Several names have been used for this filling-in operation, including *disocclusion* in [8, 32], or *inpainting* in [11, 12, 13]. In the context of this paper, and following [12], we shall refer to it as digital *inpainting*.

It turns out that images are not the only kind of data where there is a need for digital inpainting. Surfaces obtained from range scanners often have holes, regions where the 3D model is incomplete. The main cause of holes are occlusions, but these can also be due to low reflectance, constraints in the scanner placement, or simply lack of sufficient coverage of the object by the scanner. This is frequently observed in the scanning of art pieces [31], and is in part due to the fact that complicated geometry has a lot of self-occlusions and details. Art pieces also impose significant restrictions on the scanner placement. Holes are also observed in common scenarios where LADAR data is collected (e.g., a house behind an occluding tree), and in all the major areas where range scanners are used. With the increasing popularity of range scanners as 3D shape acquisition devices, with applications in geoscience, art (e.g., archival), medicine (e.g., prohestetics), manufacturing (from cars

to clothes), and defense (e.g., LADAR), it is very important to be able to inpaint this missing information. This is often needed for post-processing as well as for presentation. It is the goal of this paper to present a framework for inpainting these surface holes.

Our work is inspired by the one reported in [21], and it is presented as an alternative to this method. This pioneering work addressed the problem of hole filling via isotropic diffusion of volumetric data (that is, iterative Gaussian convolution of some distance function to the known data). The approach proposed by these authors addresses holes with complicated topology, a task very difficult with mesh representations. The reader is directed to this paper for an excellent and detailed description of the nature of holes in scanning statues and for a literature review in the subject. We should only note that most algorithms on reconstructing surfaces from range data are point-cloud reconstruction based and treat holes as regions with low sampling density, thereby interpolating across them [2, 6, 10, 18, 24, 27]. Of course, these algorithms often do not distinguish between a real hole in the data and one due to the lack of sampling, and equally fill or fail to fill both cases in the same fashion. Other point-cloud methods evolve a surface over time until it approximates the data [17, 42, 44], or fit a set of 3D radial basis functions to the data, compute a weighted sum of them and use a level set of this last function as reconstructed surface [23, 14]. Mesh based methods for surface reconstruction [39, 20, 41] can perform hole filling as a post-process or integrate hole filling into surface reconstruction [20]. One of our proposed models is closely related to the one presented in [19] (and of course to our previously introduced 3D surface inpainting model [40]), where the authors use the Willmore flow with a finite element implementation. In contrast with their work, our model works on implicit surfaces, thereby allowing for more complicated hole topologies, and also naturally leads to systems of low order differential equations.

The first algorithm here proposed is an extension of our previous work on image inpainting [7, 8, 12] (see also [11, 13, 16, 32, 34, 37]). In particular, we show how to adapt the variational formulation we presented in [7, 8] to the problem of surface hole filling. As in [21], the use of volumetric data (that is, the surface is represented as the zero level-set of a function) brings us topological freedom. In contrast with [21], we use a system of coupled anisotropic (geometric) partial differential equations designed to smoothly continue the isophotes of the embedding function, and therefore the surface of interest (as the zero level isophote). These equations are based on the geometric characteristics of the known surface (e.g., the curvatures), and as [21], are applied only at the holes and a neighborhood of them (being these equation anisotropic and geometry based, they lead to a slightly slower algorithm than the one reported in [21], as expected with geometric flows). A preliminary version of this (first) model was presented in [40]. We formalize this and improve it here with an automatic initialization method. This initialization is based on the computation of a conical neighborhood $\mathcal{F}$ of the known part of the surface, call it $\mathcal{S}$, where the distance function is uniquely attained. Thereby we can define the signed distance function d_s and then ∇d_s is the extension of the unit normal to $\mathcal{S}$ to a neighborhood of it. This construction also helps us to label both parts of the surface as interior and exterior, and this is useful in this first method.

We also develop additional curvature based hole surface inpainting methods. The first of them is based on a variational model which integrates the Laplacian of a distance function (i.e., a function which satisfies $|\nabla D| = 1$, and $D = d_s$ in the conical neighborhood $\mathcal{F}$), in a open set containing the hole. Recall that the Laplacian of the distance function gives the mean curvature of its level sets. The second method is more heuristic and is based on the diffusion of a function which represents mean curvature of level sets of an

underlying implicit function.

Finally, we also present simpler methods based on the Laplace equation and the so-called AMLE model, which permit to reconstruct a function which is distance-like near the known part of the surface and whose zero level set can be interpreted as the reconstructed surface. If our interest is just to find a smooth reconstruction, this approach may be sufficient. If one wants a reconstruction which is based on minimizing mean curvature, it can serve as an initialization.

These algorithms, except the one based on curvature diffusion which is less reliable, exhibit a similar behavior in reconstructing surface holes for synthetic and real data. As mentioned above, the reconstructions based on the Laplace or AMLE equation can be used as initializations for the curvature based approaches. Describing and studying all these techniques provides a comprehensive understanding of the different possible frameworks for filling-in surface holes.

The remainder of this paper is organized as follows. Section 2 describes our variational approach for image inpainting in any dimension N and its adaptation to the reconstruction of surface holes. In Section 3 we present the two additional curvature based approaches: a variational one minimizing the integral of the absolute value (or a power of it) of the mean curvature, and an heuristic one, based on the diffusion of curvature. Section 4 describes two simple methods for surface reconstruction based on Laplace equation and the so-called AMLE model. Section 5 describes the numerical algorithms used in our computations. In Section 6 we present some numerical experiments on hole filling obtained with the algorithms previously introduced. Finally, in Section 7, we summarize the main conclusions of the paper.

2 Joint interpolation of vector fields and gray levels and its application to surface inpainting

Let us describe the variational approach to filling-in by joint interpolation of vector fields and gray values which was introduced in [7, 8]. Our purpose is to adapt it to the problem of hole filling on surfaces.

Let Q be a hyper-rectangle in $\mathbb{R}^N$ and Ω an open bounded subset of Q with smooth boundary. Suppose that we are given an image $u_0 : Q \setminus \overline{\Omega} \rightarrow \mathbb{R}$, where $\overline{\Omega}$ denotes the closure of Ω . Using the information of u_0 on $Q \setminus \overline{\Omega}$ we want to reconstruct the image u_0 inside the hole of missing information Ω . In our context, the function u_0 is an implicit representation of the known data. In [7, 8] we proposed to fill-in the hole Ω (on images) using both the gray level u and the vector field of normals θ to the level sets of the image outside the hole. This permitted to design energy functionals which minimize a power of (mean) curvature and to write them in terms of the pair of variables (u, θ) . This is the approach that we shall explore next with the purpose of interpolating holes in surfaces.

We denote by $L^p(Q)$, $1 \leq p < \infty$, the space of (measurable) functions $f : Q \rightarrow \mathbb{R}$ such that $\int_Q |f(x)|^p dx < \infty$. By $L^\infty(Q)$ we denote the space of bounded functions $f : Q \rightarrow \mathbb{R}$.

Let $\tilde{\Omega}$ be an open subset of Q with smooth boundary such that $\overline{\Omega} \subset\subset \tilde{\Omega}$. The band around Ω , used to fill it in, is the set $B = \tilde{\Omega} \setminus \overline{\Omega}$. To fill-in the hole Ω we use the information of u_0 contained in B , mainly the gray level u_0 and the vector field of normals (i.e., the gradient directions) to the level sets of u_0 in B , which we denote by θ_0 . We assume that θ_0 is a vector field with values in $\mathbb{R}^N$ satisfying $\theta_0(x) \cdot \nabla u_0(x) = |\nabla u_0(x)|$ and $|\theta_0(x)| \leq 1$. Intuitively, this amounts to say that $\theta_0(x) = \frac{\nabla u_0(x)}{|\nabla u_0(x)|}$ when $\nabla u_0(x) \neq 0$, and we give freedom to $\theta_0(x)$ when $\nabla u_0(x) = 0$ requiring only that $|\theta_0(x)| \leq 1$. The basic

goal then is to extend in a smooth way the pair (u_0, θ_0) from the band $B = \tilde{\Omega} \setminus \overline{\Omega}$ to a pair of functions (u, θ) inside Ω . For that we attempt to continue the isosurfaces of u_0 (i.e., the hypersurfaces $[u_0 = \lambda] := \{x : u_0(x) = \lambda\}$ or, more generally, the boundaries of the level sets $[u_0 \geq \lambda]$, $\lambda \in \mathbb{R}$) in B inside Ω by taking into account the principle of good (interpreted here as smooth) continuation. The energy functional proposed in [7, 8] was based on the following principles:

a) We constrain the solution (u, θ) to coincide with the data on the band B . For that we ask that $u = u_0$ on B . The vector field θ should also satisfy $|\theta| \leq 1$ on Ω and should be related to u by $\theta \cdot \nabla u = |\nabla u|$, i.e., we impose that θ is the vector field of directions of the gradient of u . This implies that $\theta(x) = \theta_0(x)$ on the points $x \in B$ where $\nabla u_0(x) \neq 0$ and we leave freedom to $\theta(x)$ on the points where $x \in B$ where $\nabla u_0(x) = 0$ besides satisfying that $|\theta(x)| \leq 1$.

b) We impose that the vector field θ_0 in the band B is smoothly continued by θ inside Ω . Note that if θ are the directions of the normals to the level hypersurfaces of u , then $\text{div}(\theta)$ (a possible measure of smoothness of the vector field) is the mean curvature. The smooth continuation of the levels sets of u_0 inside Ω is imposed by requiring that $\text{div } \theta \in L^p(\tilde{\Omega})$, i.e., $\int_{\tilde{\Omega}} |\text{div } \theta|^p \, dx < \infty$. For consistency we shall require that θ_0 is such that $\text{div } \theta_0 \in L^p(B)$ and assume that (u_0, θ_0) may be extended to a pair of functions (u, θ) in the class of functions where we minimize our energy (see (1)).

Based on these basic principles, we proposed in [7, 8] to interpolate the pair (θ, u) in

Ω by minimizing the functional

$$\begin{aligned} & \text{Minimize } \int_{\tilde{\Omega}} |\operatorname{div}(\theta)|^p (\gamma + \beta |\nabla k * u|) dx \\ & |\theta| \leq 1, \|u\|_{\infty} \leq M, \end{aligned} \tag{1}$$

$$|\nabla u| - \theta \cdot \nabla u = 0 \text{ in } \tilde{\Omega},$$

$$u = u_0 \text{ in } B, \quad \theta \cdot \nu^{\tilde{\Omega}} = \theta_0 \cdot \nu^{\tilde{\Omega}} \text{ in } \partial\tilde{\Omega},$$

where $p > 1$, $\gamma > 0$, $\beta \geq 0$, k denotes a regularizing kernel of class C^1 such that $k(x) > 0$ a.e., and $M = \|u_0\|_{L^\infty(B)} := \sup_{x \in B} |u_0(x)|$. As usual, $\nu^{\tilde{\Omega}}$ denotes the outer unit normal to $\partial\tilde{\Omega}$. Let us note that if u is the characteristic function of a set $A \subseteq \mathbb{R}^3$ (i.e., $u(x) = \chi_A(x) = 1$ if $x \in A$ and $= 0$ otherwise) with smooth boundary and θ is a smooth extension of the unit normal to ∂A , then $\int_{\tilde{\Omega}} |\operatorname{div}(\theta)|^p |\nabla u| dx = \int_{\partial A} |H|^p dS$, where $H(x)$ is the mean curvature of ∂A and dS denotes the surface area. The convolution of ∇u with the kernel k is done for technical reasons, it permits to prove the existence of a minimum for (1) [7, 8], though we can dismiss it in practice. Finally, let us also note that the constant γ has to be > 0 , it implies an L^p bound on $\operatorname{div} \theta$, and this is useful to prove that the limits of minimizing sequences satisfy the constraint $|\nabla u| - \theta \cdot \nabla u = 0$ in $\tilde{\Omega}$ [7, 8]. We refer to [8] for a detailed theoretical analysis of this formulation and its approximation by smoother functionals. Let us finally note that the use of smooth continuation principles based on powers of mean curvature as smoothness measure was proposed in [33] with the purpose of image segmentation with depth (reconstructing then the occluded boundaries), and used as a base for variational approaches of image disocclusion in [7, 8, 32].

2.1 Surface inpainting

Let us now describe how to adapt the above formulation to inpaint (fill-in) holes (or gaps) on surfaces $\mathcal{S}$, which we assume to be embedded in $\mathbb{R}^3$. To avoid any confusion with our previous use of the word hole, let us use the word *gap* of the surface. Assume, to fix ideas, that $\mathcal{S}$ is a smooth compact connected surface, and $\mathcal{M}$ is a part of $\mathcal{S}$ which is unknown or could not be obtained during scanning (or is damaged and needs to be replaced). Let us identify $\mathcal{S}$ with its known part. Let us choose a bounding box Q in $\mathbb{R}^3$ strictly containing the surface gap $\mathcal{M}$ and part of $\mathcal{S}$ (see, for instance, Figure 8). Let $\partial\mathcal{M}$ be the boundary of the gap. Even if $\mathcal{M}$ is unknown, its relative boundary in $\mathcal{S}$ is known. Let $\mathcal{F}$ be a neighborhood of $\mathcal{S} \cap Q$ such that

$$\mathcal{F} = \{x \in Q : d(x, \mathcal{S} \cap Q) < \alpha d(x, \partial\mathcal{M})\}, \quad 0 < \alpha < 1.$$

We assume that $\mathcal{F} \setminus (\mathcal{S} \cap Q)$ consists of two connected components, which can be identified as the two sides of the surface $\mathcal{S}$ (see Figure 1). The information derived from the region $\mathcal{F}$ is considered reliable and we impose it as a constraint in our reconstruction. Let $d^{\mathcal{F}}(x)$ be the distance of a point $x \in \mathcal{F}$ to $\mathcal{S} \cap Q$. By changing the sign of $d^{\mathcal{F}}$ in one of the sides of the surface we may define the signed distance function to $\mathcal{S} \cap Q$ in $\mathcal{F}$ (take it positive inside and negative outside). We denote it by $d_s^{\mathcal{F}}(x)$, or simply, by d_s . The vector field in $\mathcal{F}$,

$$N(x) = \nabla d_s(x),$$

is an extension of the unit normal vector field on $\mathcal{S} \cap Q$ to its neighborhood $\mathcal{F}$. Again, we consider this information as reliable and it will be used as a constraint.

To adapt functional (1) to surface hole reconstruction we must make explicit the hole




Figure 1: a) Section of the surface $\mathcal{S}$ with the *hole* $\mathcal{M}$ and the neighborhood $\mathcal{F}$. b) Sign assignment to the two faces of $\mathcal{S}$. The hole Ω is defined as the ball minus $\mathcal{F}$. The band B is defined as $Q \setminus \overline{\Omega}$.

Ω and the functions (u_0, θ_0) which are the known data on a neighborhood of Ω . To define the hole we take a ball $\mathcal{G}$ (or any open set homeomorphic to a ball) such that $\overline{\mathcal{G}} \subset \subset Q$ and containing the boundary of the gap $\partial \mathcal{M}$ in its interior. We define the hole Ω by removing from $\mathcal{G}$ the points of $\mathcal{F}$. We take the band $B = Q \setminus \overline{\Omega}$ (notice that Q is identified to the set $\tilde{\Omega}$ of Section 2).




Figure 2: If the ball is not large enough, we may have difficulties in the assignment of labels. Notice that if we extend the ball, the labelling can be done in a consistent way.

We then consider $u_0 : Q \setminus \overline{\Omega} \rightarrow \mathbb{R}$ a characteristic function, that is, a binary function taking values 0 and 1. The values $u_0(x) = 1$ and $u_0(x) = 0$ represent the points which are interior, respectively, exterior, to $\mathcal{S}$. Recall that we are assuming that $\mathcal{F} \setminus (\mathcal{S} \cap Q)$ consists of two connected components which represent the two sides of the surfaces. We then label these two sides with the values $u_0 = 1$, representing the inner part of the surface, and $u_0 = 0$, representing its outer side. The set $[u_0 = 1] \cap \mathcal{F}$ (respectively, the set $[u_0 = 0] \cap \mathcal{F}$) is represented in Figure 1.b with the sign $+$ (resp., the sign $-$). By propagation, we extend this labelling to the rest of $Q \setminus \Omega$, knowing it already in $\mathcal{F}$. For simplicity, we assume that, by eventually extending the ball, this labelling can be done in a consistent way: we assume that $\mathcal{G}$ can be chosen so that $\partial\mathcal{G}$ is divided into two labelled regions consistent with the labels in $\mathcal{F} \setminus (\mathcal{S} \cap Q)$ (see figures 1.b and 2). We call A the set of points x in $Q \setminus \Omega$ such that $u_0(x) = 1$, hence $u_0(x) = \chi_A(x)$. In this case, by minimizing (1), we want to reconstruct the set A inside the hole Ω knowing the set outside Ω .

Recall that we minimize (1) by solving the gradient descent equations (9), (10) presented below, using the numerical approach described in Section 5.1, and we need the initial conditions for u and θ in $Q = \Omega \cup B$, so that $\theta \cdot \nabla u = |\nabla u|$. For that, we define u_{init} in $Q = \Omega \cup B$ as the extension of u_0 inside Ω by a geodesic propagation, so that u_{init} takes values 0 and 1. Then we define D_s as the signed distance to $\partial[u_{\text{init}} = 0]$ (negative in $[u_{\text{init}} = 0]$ and positive in $[u_{\text{init}} = 1]$) so that, by construction, D_s is an extension of d_s (recall that d_s is defined in $\mathcal{F}$). We take the vector field $\theta_{\text{init}} = \nabla D_s$ in $Q = \Omega \cup B$ and we define $\theta_0 = \theta_{\text{init}}$ in B . Observe that $\nabla u_0 = \nu \delta_{\mathcal{S} \cap Q}$, where ν denotes the inner unit normal to $\mathcal{S}$ and $\delta_{\mathcal{S} \cap Q}$ is the Hausdorff measure on $\mathcal{S} \cap Q$. We have $\theta_{\text{init}} \cdot \nabla u_{\text{init}} = |\nabla u_{\text{init}}|$.

3 Alternative curvature based approaches

Following the description in the introduction, we now present alternative filling-in approaches for the problem of filling holes in three dimensions.

3.1 Energy in terms of distance functions

Recall that if $\mathcal{S}'$ is a smooth manifold of class C^2 , then the signed distance function D to $\mathcal{S}'$ is also of class C^2 in a neighborhood of $\mathcal{S}'$. The vector field ∇D is an extension of the unit normal to $\mathcal{S}'$ and satisfies $|\nabla D| = 1$. The operator $\Delta D(x) = \operatorname{div} \nabla D(x)$ represents the sum of the principal curvatures of the isosurface $[D = D(x)] := \{y \in Q : D(y) = D(x)\}$. When we look at this as a function in $\mathbb{R}^3$, the distance function is Lipschitz and it satisfies $|\nabla D| = 1$ in the viscosity sense. The isosurfaces may develop singularities and the only thing we can expect is that the mean curvature is a Radon measure. Indeed, recall that the mean curvature of a polyhedral surface is a Dirac's measure concentrated at the edges, and the signed distance function may have such type of singularities. We shall assume that the signed distance D to the surface $\mathcal{S}$ is such that $\Delta D \in \mathcal{M}(Q)$ where $\mathcal{M}(Q)$ denotes the space of Radon measures in Q [1]. We define

$$W(Q) = \{u \in L^1(Q) : \nabla u \in L^1(Q), \Delta u \in \mathcal{M}(Q)\}.$$

We propose to fill-in the three dimensional holes via the minimization of the functional

$$\operatorname{Min}_{\{D \in W(Q), |\nabla D| = 1, D = d_s \text{ in } \mathcal{F}\}} \int_Q |\Delta D(x)| dx. \quad (2)$$

This energy integrates the mean curvature on the isosurfaces of D . Due to the singularities of the isosurfaces of D , the integral of a power of the mean curvature with an exponent

$p > 1$ may be infinite. Let us observe that problem (2) has a minimizer as soon as the admissible set is nonempty, and we assume that this is the case. Details on the implementation of the proposed approach are presented in Section 5.2.

3.2 Curvature diffusion and distance reconstruction

We also present studies based on diffusion of the mean curvature (see also [43] for related work based on the linear Poisson equation). For convenience, let us write $Q_{\mathcal{F}} := Q \setminus \mathcal{F}$. We propose to diffuse the mean curvature of $\mathcal{S}$ and then reconstruct the surface with the prescribed curvature, that is, we propose to solve the system of PDEs

$$\omega_t = \Delta \omega \quad \text{in } [0, \infty) \times (Q \setminus \mathcal{F}) \quad (3)$$

$$u_t = |\nabla u| \left(\operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right) - \omega \right) \quad \text{in } [0, \infty) \times (Q \setminus \mathcal{F})$$

with the following boundary conditions on $\partial\mathcal{F} \cap Q$:

$$\omega = \Delta d_s \quad \text{on } \partial\mathcal{F} \cap Q, \quad (4)$$

$$\nabla u \cdot \nu = \nabla d_s \cdot \nu \quad \text{on } \partial\mathcal{F} \cap Q,$$

where d_s denotes the signed distance to $\mathcal{S}$ and ν denotes the outer unit normal to $\partial\mathcal{F}$. Observe that we did not write the Dirichlet boundary condition $u = d_s$ because it is not possible, in general, to impose it in a classical sense to the equation for u (mean curvature type equation). Let us comment on the boundary conditions used on $\partial Q_{\mathcal{F}} \setminus \partial\mathcal{F}$. First of all, we observe that the ideal scenario would be to consider $Q = \mathbb{R}^N$ and solve the system of PDEs (3) in $[0, \infty) \times (\mathbb{R}^N \setminus \mathcal{F})$ with the boundary conditions (4), but this is impossible

at the numerical level. For that we modify the boundary conditions on $\partial Q_{\mathcal{F}} \setminus \partial \mathcal{F}$, we do linear extrapolation of u and ω along the normal to the level sets of u . Details on the implementation of the proposed flow are presented in Section 5.2.

4 The Laplace and the Absolute Minimizing Lipschitz extension interpolation

In [15] we studied and classified interpolation algorithms which satisfy a reasonable series of axioms in terms of the solution of a partial differential equation. Two particular examples are the Absolutely Minimizing Lipschitz Extension, denoted as AMLE in the sequel, and the Laplacian interpolation. We now discuss the applicability of AMLE and the Laplace equation to the problem of filling-in surface holes.

The Laplacian interpolation is based on solving the PDE

$$-\Delta u = 0 \quad \text{in } Q_{\mathcal{F}}, \quad (5)$$

with specified boundary data on $\partial Q_{\mathcal{F}}$. Indeed, boundary data is only known in $\partial \mathcal{F} \cap Q$, where we should impose that $u = d_s$. Thus, a reasonable assumption would be to consider

$$\frac{\partial u}{\partial \nu} = 0 \quad \text{in } \partial Q_{\mathcal{F}} \setminus \partial \mathcal{F} \quad (6)$$

where ν denotes the outer unit normal to $\partial Q_{\mathcal{F}} \setminus \partial \mathcal{F}$. In some sense, from the theoretical point of view, the lack of knowledge of boundary conditions for u in $\partial Q_{\mathcal{F}} \setminus \partial \mathcal{F}$ is an obstacle to the possibility of using (5) to reconstruct the surface $\mathcal{S} \cap Q$ (which is defined as $\partial[u > 0]$). In spite of this, we can solve (5) with the boundary condition (6), or use,

instead of (6), a linear extrapolation of u along its gradient direction on $\partial Q_{\mathcal{F}} \setminus \mathcal{F}$. The results obtained with this last boundary condition are presented in Section 6. We should of course mention that this approach is closely related to the work in [21], based on linear diffusion.

The AMLE interpolation ([4, 5]) is based on solving the PDE

$$\langle D^2 u(\nabla u), \nabla u \rangle = 0 \quad \text{in } Q_{\mathcal{F}}. \quad (7)$$

with boundary data on $\partial Q_{\mathcal{F}}$ (here ∇u and $D^2 u$ denote the gradient and the Hessian matrix of u , respectively, so that in Cartesian coordinates, $\langle D^2 u(\nabla u), \nabla u \rangle = \sum_{i,j=1}^3 \frac{\partial^2 u}{\partial x_i \partial x_j} \frac{\partial u}{\partial x_i} \frac{\partial u}{\partial x_j}$). As a short review on (7), let us recall that this equation can be solved on general domains and boundary data, in particular the data can be given in a finite number of surfaces, curves and/or points. Indeed, we may assume that the boundary data $\varphi \in Lip_{\partial}(Q_{\mathcal{F}})$ where

$$Lip_{\partial}(Q_{\mathcal{F}}) = \left\{ g \in C(\partial Q_{\mathcal{F}}) : |||g||| = \sup_{x,y \in \partial Q_{\mathcal{F}}} \frac{|g(x) - g(y)|}{d_{\partial Q_{\mathcal{F}}}(x,y)} < \infty \right\},$$

and $d_{Q_{\mathcal{F}}}(x,y)$ is the geodesic distance between x and y in $Q_{\mathcal{F}}$, i.e., the minimal length of all possible paths joining x and y and contained in $Q_{\mathcal{F}}$ [29].

Let us recall that, if X is an open set or a smooth manifold in $\mathbb{R}^N$, $W^{1,\infty}(X)$ (resp. $W^{1,p}(X)$) denotes the space of functions $u \in L^{\infty}(X)$ (resp. $u \in L^p(X)$) such that $\nabla u \in L^{\infty}(X)$ (resp. $\nabla u \in L^p(X)$). By $W_0^{1,\infty}(X)$ (resp. $W_0^{1,p}(X)$), we denote the closure in $W^{1,\infty}(X)$ (resp. in $W^{1,p}(X)$), of the smooth functions with compact support in X .

Existence and uniqueness of viscosity solutions for the AMLE model (7) with boundary data $\varphi \in Lip_{\partial}(Q_{\mathcal{F}})$ was proved by Jensen [29]. Moreover, he proved that the viscosity solution of (7) is an absolutely minimizing Lipschitz extension of φ , i.e., $u \in W^{1,\infty}(Q_{\mathcal{F}}) \cap$

$C(\overline{Q_{\mathcal{F}}})$ and satisfies

$$\|\nabla u\|_{L^\infty(Q'; R^N)} \leq \|\nabla w\|_{L^\infty(Q'; R^N)} \quad (8)$$

for all $Q' \subseteq Q_{\mathcal{F}}$ and w such that $u - w \in W_0^{1,\infty}(Q')$. Let us add that the AMLE model was introduced by Aronsson in [4, 5] as the Euler-Lagrange equation of the variational problem (8) (which can be interpreted as the limit as $p \rightarrow \infty$ of the variational problems $\|\nabla u\|_{L^p(Q'; R^N)} \leq \|\nabla w\|_{L^p(Q'; R^N)}$ for all $Q' \subseteq Q_{\mathcal{F}}$ and w such that $u - w \in W_0^{1,p}(Q')$ [9, 29]). The above results were extended in [29] to the case of continuous boundary data $\varphi \in C(\partial Q_{\mathcal{F}})$, and Jensen proved that in that case, the *AMLE* is locally Lipschitz continuous in $Q_{\mathcal{F}}$ [29].

Concerning boundary conditions, the same remarks we made for the Laplace equation (5) can be done here, that is, boundary data are only known in $\partial\mathcal{F} \cap Q$ where we should impose that $u = d_s$ (by the results in [30], there exist absolutely minimizing Lipschitz extensions of $d_s|_{\partial\mathcal{F} \cap Q}$ which satisfy (7), but there is no uniqueness result for them). From the theoretical point of view, the lack of knowledge of boundary conditions on $\partial Q_{\mathcal{F}} \setminus \partial\mathcal{F}$ is an obstacle to the possibility of using (7) to reconstruct the surface $\mathcal{S} \cap Q$ (which is defined as $\partial[u > 0]$). To overcome this difficulty, we propose to use either Neumann boundary conditions (6), or to use linear extrapolation of the values of u along the direction ∇u on $\partial Q_{\mathcal{F}} \setminus \partial\mathcal{F}$. The experiments displayed in Section 6 were obtained with the last method. Details on the implementation of both Laplacian and AMLE interpolation methods are presented in Section 5.3.

5 Numerical considerations

We now present some basic concepts related to the numerical implementation of the different filling-in models described above. For later use let us define the backward and forward derivative operators by:

$$\begin{aligned}\delta_{x_1}^+ u(i, j, k) &= u(i+1, j, k) - u(i, j, k) & \delta_{x_1}^- u(i, j, k) &= u(i, j, k) - u(i-1, j, k) \\ \delta_{x_2}^+ u(i, j, k) &= u(i, j+1, k) - u(i, j, k) & \delta_{x_2}^- u(i, j, k) &= u(i, j, k) - u(i, j-1, k) \\ \delta_{x_3}^+ u(i, j, k) &= u(i, j, k+1) - u(i, j, k) & \delta_{x_3}^- u(i, j, k) &= u(i, j, k) - u(i, j, k-1)\end{aligned}$$

where indexes (i, j, k) denote the voxel location. We use the superscript c to denote centered differences in space derivatives, and we define $\nabla^c u = (\delta_{x_1}^c u, \delta_{x_2}^c u, \delta_{x_3}^c u)$ where $\delta_{x_1}^c u(i, j, k) = \frac{1}{2} (u(i+1, j, k) - u(i-1, j, k))$, with analogous expressions for $\delta_{x_2}^c u(i, j, k)$, $\delta_{x_3}^c u(i, j, k)$.

5.1 Joint interpolation of vector fields and gray levels

To minimize the functional (1) we use the steepest descent method. If we denote the energy by $\tilde{E}(\theta, u)$, the steepest descent equations are

$$\theta_t = -\nabla_{\theta} \tilde{E}(\theta, u) \tag{9}$$

and

$$u_t = -\nabla_u \tilde{E}(\theta, u) \tag{10}$$

in $(0, \infty) \times \tilde{\Omega}$, supplemented with the corresponding boundary data and initial conditions.

The constraints on (θ, u) can be incorporated either by penalization or by projecting onto

them after each time step. We tested both methods in an implicit (also in an explicit) in time discretization of (9), (10). Let us explain in some detail the implicit in time implementation of (9), (10) with the constraint $\theta \cdot \nabla u = |\nabla u|$ incorporated by penalization. The constraint $u = u_0$ in B is also incorporated by penalization. Thus we consider

$$\tilde{E}(u, \theta) = \int_{\tilde{\Omega}} |\operatorname{div}(\theta)|^p (\gamma + \beta |\nabla k * u|) dx + \eta \int_{\tilde{\Omega}} (|\nabla u| - \theta \cdot \nabla u) + \lambda \int_B (u - u_0)^2 dx \quad (11)$$

which corresponds to the energy (1) plus two penalization terms for the constraints $\theta \cdot \nabla u = |\nabla u|$, and $u = u_0$ on B , with $\eta, \lambda > 0$. At the discrete level, the integral is replaced by a sum, the gradients are discretized with forward derivatives and the divergence with backward derivatives (or conversely, gradients with backward derivatives, while divergences with forward ones). We follow this structure since at the continuous level the divergence is the dual operator of the gradient, while at the discrete level the divergence discretized with backward (resp., forward) derivatives is the dual of the gradient with forward (resp., backward) derivatives. Thus, we use as discrete approximations: $\nabla u = (\delta_{x_1}^+ u, \delta_{x_2}^+ u, \delta_{x_3}^+ u)$, and similarly for $\nabla k * u$, while $\operatorname{div} \theta = \sum_{i=1}^3 \delta_{x_i}^- \theta_i$ if $\theta = (\theta_1, \theta_2, \theta_3)$. To simplify our notation, let us write $g(\theta) := \beta |\operatorname{div}(\theta)|^p$, $h(u) := \gamma + \beta |\nabla k * u|$. Then

$$\nabla_{\theta} \tilde{E}(\theta, u) = -p \nabla [h(u) |\operatorname{div}(\theta)|^{p-2} \operatorname{div}(\theta)] - \eta \nabla u = 0 \quad (12)$$

and

$$\nabla_u \tilde{E}(\theta, u) = -\operatorname{div} \left(k * \left(g(\theta) \frac{\nabla k * u}{|\nabla k * u|} \right) \right) - \eta \operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right) + \eta \operatorname{div} \theta + 2\lambda (u - u_0) \chi_B = 0, \quad (13)$$

both equations holding in $\tilde{\Omega}$. As we mentioned above, the gradients are discretized with forward derivatives, while backward ones are used to discretize the divergences. To solve equations (9) and (10), we use an implicit discretization in time. To be precise, we write

$$\nabla_{\theta} \tilde{E}(\theta, \theta', u, v) = -p \nabla [h(u)(\epsilon + |\operatorname{div}(\theta')|^{p-2}) \operatorname{div}(\theta)] - \eta \nabla \mathbf{u}, \quad (14)$$

and

$$\begin{aligned} \nabla_u \tilde{E}(\theta, \theta', u, v) = & -\operatorname{div} \left(\mathbf{k} * \left(g(\theta) \frac{\nabla \mathbf{k} * \mathbf{u}}{\sqrt{\epsilon + |\nabla \mathbf{k} * \mathbf{v}|^2}} \right) \right) - \eta \operatorname{div} \left(\frac{\nabla \mathbf{u}}{\sqrt{\epsilon + |\nabla \mathbf{v}|^2}} \right) \\ & + \eta \operatorname{div}(\theta) + 2\lambda(\mathbf{u} - \mathbf{u}_0) \chi_B. \end{aligned} \quad (15)$$

Then, we use the discretization in time given by

$$\theta^{n+1} - \theta^n = -\Delta t \nabla_{\theta} \tilde{E}(\theta^{n+1}, \theta^n, u^n, u^n), \quad (16)$$

and

$$u^{n+1} - u^n = -\Delta t \nabla_u \tilde{E}(\theta^{n+1}, \theta^{n+1}, u^{n+1}, u^n). \quad (17)$$

Finally, we make the change of variables $\xi^{n+1} := \theta^{n+1} - \theta^n$, $v^{n+1} := u^{n+1} - u^n$ and we have

$$\xi^{n+1} = -\Delta t \nabla_{\theta} \tilde{E}(\xi^{n+1} + \theta^n, \theta^n, u^n, u^n), \quad (18)$$

$$v^{n+1} = -\Delta t \nabla_u \tilde{E}(\theta^{n+1}, \theta^{n+1}, v^{n+1} + u^n, u^n). \quad (19)$$

In practice we solve (18) in $\tilde{\Omega}$ with the boundary condition

$$\theta^{n+1} \cdot \nu^{\tilde{\Omega}} = \theta_0 \cdot \nu^{\tilde{\Omega}} \quad \text{on } \partial \tilde{\Omega},$$

Now, since $\theta^{n+1} \cdot \nu^{\tilde{\Omega}}|_{\partial\tilde{\Omega}} = \theta^n \cdot \nu^{\tilde{\Omega}}|_{\partial\tilde{\Omega}}$, we may write this boundary condition as

$$\xi^{n+1} \cdot \nu^{\tilde{\Omega}} = 0 \quad \text{on } \partial\tilde{\Omega}.$$

The boundary condition for (19) is (see [8])

$$k * \left(g(\theta^{n+1}) \frac{\nabla k * u^{n+1}}{\sqrt{\epsilon + |\nabla k * u^n|^2}} \right) \cdot \nu^{\tilde{\Omega}} + \eta \left(\frac{\nabla u^{n+1}}{\sqrt{\epsilon + |\nabla u^n|^2}} \right) \cdot \nu^{\tilde{\Omega}} = \eta \theta_0 \cdot \nu^{\tilde{\Omega}} \quad \text{on } \partial\tilde{\Omega}.$$

which is written in terms of $v^{n+1} = u^{n+1} - u^n$ as

$$B_k(\theta^{n+1}, v^{n+1}, u^n) = \eta \theta_0 \cdot \nu^{\tilde{\Omega}} - B_k(\theta^{n+1}, u^n, u^n),$$

where $B_k(\theta^{n+1}, u, u^n) = k * \left(g(\theta^{n+1}) \frac{\nabla k * u}{\sqrt{\epsilon + |\nabla k * u^n|^2}} \right) \cdot \nu^{\tilde{\Omega}} + \eta \left(\frac{\nabla u}{\sqrt{\epsilon + |\nabla u^n|^2}} \right) \cdot \nu^{\tilde{\Omega}}$. Then we use a conjugate gradient method to solve (18) and (19). The constraint $|\theta| \leq 1$ is incorporated by renormalizing θ^n (when $|\theta^n| > 1$) after each time step. The constraints on $\|u\|_\infty$ can be also introduced after each time step. In spite of the penalization term, the relationship $\theta \cdot \nabla u = |\nabla u|$ is (numerically) lost and we reinforce it after a certain number of time steps.

We can also set $\eta = 0$ and incorporate the constraint that $|\nabla u| = \theta \cdot \nabla u$ by projecting onto it after each time step. We also tested this both in a time implicit and explicit discretization of equations (9) and (10). After each time step of θ and u we redefine

$$\theta(i, j, l) = \frac{\theta(i, j, l) + \alpha \nabla u(i, j, l)}{\max(1, |\theta(i, j, l) + \alpha \nabla u(i, j, l)|)}$$

for some $\alpha > 0$. As it has been shown in [28] this is a good way of imposing that $|\theta| \leq 1$ and $\theta \cdot \nabla u = |\nabla u|$. We have found quite similar results using both described methods for imposing the constraint.

In our experiments, we take k a Gaussian kernel with small variance, say one or two pixels. In practice, one can also dismiss the kernel k . The initial conditions for u and θ are taken as we explained at the end of Section 2.1 so that $\theta \cdot \nabla u = |\nabla u|$ in $Q := \Omega \cup B$. We could also construct an initialization (u, θ) using the solution Laplace equation, or AMLE (see Section 4). As regards the parameters in functional (11), in our experiments we choose: $\beta = 1$, $\lambda = 100$, $\gamma = 1$ and $\eta = 10$ (although we have experimented with $\gamma = 0, \gamma = 0.5$ and $\eta = 1$ and the results are always good). The parameter ϵ in equations (14) and (15) is set to 0.1.

Finally, let us observe that, since both constraints $u = u_0$ in B and $\theta \cdot \nabla u = |\nabla u|$ in $\tilde{\Omega}$ have been imposed by penalization, this gives some robustness to noise, but this can be also achieved with a previous smoothing of the surface near the hole.

5.2 Curvature-base approaches

We use the steepest descent method also to minimize the functional (2), and we solve

$$D_t = -\Delta \left(\frac{\Delta D}{|\Delta D|} \right) \quad (20)$$

In order to satisfy the constraint $|\nabla D| = 1$ we make use of the PDE that computes the signed distance function [35]:

$$D_t = -\text{sign}(D)(|\nabla D| - 1) \quad (21)$$

Then, as a numerical approach to minimize (2) we combine (20) and (21) at each time step. To go from D^n to D^{n+1} we first solve

$$D^* = D^n - \Delta t \Delta \left(\frac{\Delta D^n}{|\Delta D^n| + \epsilon} \right)$$

using the standard discretization for the Laplacian ($\Delta w = \sum_{i=1}^3 \delta_{x_i}^+ \delta_{x_i}^- w$ for any function w) and then solve

$$D^{n+1} = D^* - \Delta t \text{sign}(D^*) (|\nabla D^*| - 1)$$

using an upwind scheme for the gradient magnitude [35, 36]. The precise numerical scheme is:

$$D^{n+1} = D^* - \Delta t [\max(\text{sign}(D^*), 0)(|\nabla^+ D^*| - 1) + \min(\text{sign}(D^*), 0)(|\nabla^- D^*| - 1)]$$

where

$$|\nabla^+ D^*| = \left[\sum_{i=1}^3 \max(\delta_{x_i}^- D^*, 0)^2 + \min(\delta_{x_i}^+ D^*, 0)^2 \right]^{1/2}, \quad (22)$$

$$|\nabla^- D^*| = \left[\sum_{i=1}^3 \max(\delta_{x_i}^+ D^*, 0)^2 + \min(\delta_{x_i}^- D^*, 0)^2 \right]^{1/2}. \quad (23)$$

As boundary conditions, we use $D^n = D^* = d_s$ on $\partial \mathcal{F} \cap Q$ and linear extrapolation of D^n and D^* along its gradient direction on $\partial Q_{\mathcal{F}} \setminus \partial \mathcal{F}$.

Even if not fully theoretically justified, unless we work in a small neighborhood of the surface $D = 0$, a similar scheme may be used to minimize the functional

$$\text{Min} \left\{ D: |\nabla D| = 1, \Delta D \in L^2(Q), D = d_s \text{ in } \mathcal{F} \right\} \int_Q |\Delta D(x)|^2 dx. \quad (24)$$

In order to solve the system (3), we solve first the equation $\omega_t = \Delta \omega$ using an explicit

Euler scheme and the standard discretization for the Laplacian. Then, we reconstruct the distance function by solving $u_t = |\nabla u| \left(\operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right) - \omega \right)$. For that, we use the following explicit scheme ([36]):

$$u^{n+1} = u^n - \Delta t \left[|\nabla^+ u^n| \max(\omega - \operatorname{div}^- \nabla_n^+ u^n, 0) + |\nabla^- u^n| \min(\omega - \operatorname{div}^- \nabla_n^+ u^n, 0) \right]$$

where

$$\operatorname{div}^- \nabla_n^+ u^n = (\delta_{x_1}^-, \delta_{x_2}^-, \delta_{x_3}^-) \cdot \left(\frac{(\delta_{x_1}^+ u^n, \delta_{x_2}^+ u^n, \delta_{x_3}^+ u^n)}{\sqrt{\epsilon^2 + |(\delta_{x_1}^+ u, \delta_{x_2}^+ u, \delta_{x_3}^+ u)|^2}} \right),$$

where $\epsilon > 0$ is a small parameter, and $|\nabla^+ u^n|$, $|\nabla^- u^n|$ are given as in (22), (23).

5.3 The Laplace equation and the AMLE

The Laplace equation (5) is solved by computing the steady state of equation $u_t = \Delta u$. This PDE is discretized with an implicit (or explicit) Euler scheme in time and the standard discretization for the Laplacian. In case of the implicit Euler scheme, the corresponding linear system is solved using the conjugate gradient method. As boundary conditions we use $u = d_s$ on $\partial \mathcal{F} \cap Q$ and linear extrapolation of u along its gradient direction on $\partial Q_{\mathcal{F}} \setminus \partial \mathcal{F}$ (one can also use the Neumann boundary conditions (6)). The same boundary conditions are used for the AMLE equation whose numerical scheme we describe now.

The AMLE equation (7) is also solved computing the steady state of its associated evolution problem. Following [15], we use a nonlinear overrelaxation method (with an implicit Euler scheme in time and centered differences for space derivatives). Thus, we compute u^{n+1} at (i, j, k) by:

$$u^{n+1} = u^n - w \frac{(1 + 2\Delta t)|\nabla^c u^*|^2 u^* - u^n |\nabla^c u^*|^2 - \Delta t \sum_{i,j=1}^3 \bar{u}_{x_i x_j}^* \delta_{x_i}^c u^* \delta_{x_j}^c u^*}{(1 + 2\Delta t)|\nabla^c u^*|^2 + \epsilon} \quad (25)$$

where the relaxation parameter is $0 < w < 2$ (we use, in practice, $w = 1.5$).

When solving (25), voxels are visited in a lexicographic order. Let us denote by $\preceq$ the lexicographic order for indexes in Z^3 . Now, when computing the value of $u^{n+1}(i, j, k)$ we use u^* , defined below, which is a function that takes the same values as u^n or u^{n+1} depending at the pixel location (p, q, r) where we evaluate it:

$$u^*(p, q, r) = \begin{cases} u^n(p, q, r) & \text{if } (i, j, k) \preceq (p, q, r), \\ u^{n+1}(p, q, r) & \text{otherwise,} \end{cases} \quad (26)$$

Finally, we define $\bar{u}_{x_i x_j}^*$ as

$$\bar{u}_{x_i x_j}^* := \begin{cases} \delta_{x_i}^- \delta_{x_i}^+ u^* - 2u^* & \text{if } i = j, \\ \delta_{x_i}^c \delta_{x_j}^c u^* & \text{otherwise.} \end{cases}$$

This iterative scheme (25) gives the solution of (7).

6 Experimental results

We now show experimental results illustrating the filling-in techniques here proposed.

6.1 Simple geometric objects

First, we present experiments of geometric objects done with the different methods discussed above: the joint interpolation of vector fields and gray levels (abbreviated JIVFGL), minimization of the absolute value of the Laplacian of the distance function (and also the case of power 2), curvature diffusion, the AMLE, and Laplace equations. The images in our experiments have been rendered using the AMIRA Visualization and Modeling System [3].

Figure 3 shows a pyramid with a hole and its corresponding reconstructions; the basis of the pyramid is a square whose sides have length 61 voxels, the height is of 30 voxels and the hole is parallel to the basis going from voxels of height 8 to voxels of height 12, both included. Figure 3(a) displays an ideal pyramid from which we cut a piece, the pyramid with the hole is displayed in Figure 3(b). Notice that the boundary of the hole looks rounded, this is only due to our display and does not correspond to the data. Figure 3(c) shows its reconstruction by joint interpolation of vector fields and gray levels (JIVFGL), i.e., functional (1) with $N = 3$. Figure 3(d) shows the result obtained minimizing the absolute value of the Laplacian of the distance function, i.e., (2). Figure 3(e) shows the result obtained minimizing the square of the Laplacian of the distance function, i.e., (24). Figure 3(f) shows the result obtained with curvature diffusion (3). Finally, Figures 3(g) and 3(h) display the results obtained solving the AMLE and Laplace equations in 3D, respectively. Note how the reconstructions obtained with the model JIVFGL and the square of the Laplacian for example manage to fill-in a relatively large hole. The others do a decent work that can certainly be used as a very good initial condition for the best models to refine. For a better display of the results we show in Figure 4 an orthoslice of each of the reconstructions, corresponding to a plane parallel to the basis on the middle of the hole (voxels located at height 10).

Figure 5 displays a torus with a hole and its corresponding reconstructions; the torus is generated by the revolution of a circle of radius $r = 10$ voxels along a circle of radius $R = 20$ voxels, the hole extends over an angle of $c = 35^\circ$ (approximately). Figure 5(a) displays the ideal torus from which we cut a piece. Figure 5(b) shows the gap in the torus, notice again that the boundary of the hole looks rounded, this is only due to our display and does not correspond to the data. Figure 5(c) shows its reconstruction by joint

interpolation of vector fields and gray levels (JIVFGL). Figure 5(d) displays the result obtained minimizing the absolute value of the Laplacian of the distance function, i.e., (2). Figure 5(e) displays the result obtained minimizing the square of the Laplacian of the distance function, i.e., (24). Figures 5(f) and 5(g) display the results obtained solving the AMLE and Laplace equations in 3D. Note that the best results are obtained with the JIVFGL model and with (24). For a better display of our results, we show in Figure 6 an orthoslice of each reconstruction through a plane passing by the middle of the torus.


6.2 Experiments with Michelangelo's David

For this real data, with the purpose of adapting it to our algorithm, the data, originally given as a triangulated surface, was converted to an implicit representation in a regularly spaced 3D grid. The result is visualized again as a triangulated surface. Figure 7 shows a rendering of a scanned version of Michelangelo's David which has several holes.


Figures 8(a) and 8(b) show some particular holes with a bounding box isolating them. Figures 8(c) and 8(d) show the triangulated surface (the data) around the hole. The reconstructed surface by the JIVFGL model are shown in Figures 8(e) and 8(f). The reconstructed surfaces look natural.

Figure 9(a) shows the hole in David's left hand. Figures 9(b),(c),(d),(e) and (f) show the corresponding results obtained minimizing the absolute value of the Laplacian, the square of the Laplacian, diffusion of curvature, the AMLE, and Laplace equations, respectively. All reconstructions look again natural, but we observe that the result obtained with curvature diffusion is less smooth. The differences between reconstructions are more clearly seen in the orthoslices of Figure 10.


Figure 11(a) shows the hole in David's hair. Figures 11(b),(c),(d),(e) and (f) show the




(a) An ideal pyramid.




(b) Pyramid with a gap.




(c) Reconstruction using JIVFGL.




(d) Minimizing the integral of $|\Delta D|$.




(e) Minimizing the integral of $|\Delta D|^2$.



(f) Curvature diffusion.





(g) AMLE.






(h) Laplace equation.

Figure 3: Reconstruction of a pyramid with a gap.



(a) Pyramid and orthoslice 10 of
z axis.




(b) Orthoslice of the ideal recon-
struction.



(c) Reconstruction using
JIVFGL.

(d) Minimizing the integral of
 $|\Delta D|$.

(e) Minimizing the integral of
 $|\Delta D|^2$.




(f) Curvature diffusion.


(g) AMLE.

(h) Laplace equation.


Figure 4: Orthoslices of the reconstruction of a pyramid with a gap.




(a) An ideal torus.




(b) Torus with a gap.




(c) Reconstruction using JIVFGL.




(d) Minimizing the integral of $|\Delta D|$.



(e) Minimizing the integral of $|\Delta D|^2$.




(f) AMLE.




(g) Laplace equation.


Figure 5: Reconstruction of a torus with a gap.




(a) Slice of the ideal torus.




(b) The gap of the torus.




(c) Reconstruction using JIVFGL.




(d) Minimizing the integral of $|\Delta D|$.



(e) Minimizing the integral of $|\Delta D|^2$.



(f) AMLE.



(g) Laplace equation.

Figure 6: Orthoslices of the reconstruction of a torus with a gap.




Figure 7: Scanned version of Michelangelo's David (data from the Stanford Michelangelo's project)


corresponding results obtained minimizing the absolute value of the Laplacian, the square of the Laplacian, diffusion of curvature, the AMLE, and Laplace equations, respectively. Again, the comparison between different reconstructions can be done with the orthoslices displayed in Figure 12.

7 Conclusions


In this note we have shown how to extend our previous work on variational image inpainting to fill-in surface holes. The idea, inspired by [21] and [7, 8], is to represent the surface of interest by means of a function u , as an upper level set $[u > 0]$, and minimize an energy functional which integrates a power of the mean curvature of the level sets of u . Then we use a gradient descent method and, thus, we run a system of coupled geometric partial differential equations that permit to geometrically continue the surface into the hole. We have also discussed other curvature based hole surface reconstruction models, one of them




(a) David's left hand.




(b) A detail of its hair.




(c) A zoomed detail of
(a) showing the triangulated
surface with the hole.



(d) A zoomed detail of
(b) showing the triangulated
surface with the hole.




(e) The reconstruction of the
hole in (c).




(f) The reconstruction of the
hole in (d).


Figure 8: Two different parts of David's surface and their corresponding bounding boxes Q containing a hole on the surface. The corresponding results obtained using JIVFGL.




(a) The hole in David's left hand.




(b) Minimizing the absolute value of the Laplacian.




(c) Minimizing the square of the Laplacian.



(d) Curvature diffusion.




(e) AMLE.



(f) Laplace equation.

Figure 9: The hole in David's left hand and some reconstructions using different methods.



(a) Hand surface and coronal orthoslice number 35.



(b) Orthoslice of the gap.



(c) Reconstruction using JIVFGL.



(d) Minimizing the absolute value of the Laplacian.



(e) Minimizing the square of the Laplacian.



(f) Curvature diffusion.




(g) AMLE.




(h) Laplace equation.


Figure 10: Orthoslices of the reconstruction of the hole in David's left hand.




(a) The hole in David's hair.




(b) Minimizing the integral of $|\Delta D|$.




(c) Minimizing the integral of $|\Delta D|^2$.



(d) Curvature diffusion.




(e) AMLE.



(f) Laplace equation.

Figure 11: The hole in David's hair and some reconstructions using different methods.



(a) Hair surface and axial orthoslice number 59.

(b) Orthoslice of the gap.



(c) Reconstruction using JIVFGL.



(d) Minimizing the integral of $|\Delta D|$.



(e) Minimizing the integral of $|\Delta D|^2$.



(f) Curvature diffusion.



(g) AMLE.



(h) Laplace equation.

Figure 12: Orthoslices of the reconstruction of the hole in David's hair.

based on a variational model which integrates the Laplacian of a distance function, the other is heuristic and is based on the diffusion of a function which represents the mean curvature of level sets of an underlying implicit function. In all these cases, we have shown reconstruction of surface holes both for synthetic and real data.

Finally, we have also shown simpler methods based on the Laplace equation and the so-called AMLE model which reconstructs a function which is distance-like near the known part of the surface and whose zero level set can be interpreted as the reconstructed surface. If our interest is just to find a smooth reconstruction, this approach may be sufficient. If one wants a reconstruction which is based on minimizing mean curvature, it can just serve as an initialization stage.

Acknowledgments

This paper is dedicated to Andrés Francisco Solé Martínez, a friend and colleague who departed this world too early.

GS wants to thank Prof. Marc Levoy from Stanford University for providing him an early version of his paper on hole filling. This paper and follow up conversations with Prof. Levoy resulted in challenging ourself to extend our techniques for image inpainting to the problem of hole filling, leading to the work here presented. We also thank the Stanford Michelangelo Project for data provided for this work, Marcelo Bertalmío who constantly provides good feedback to us on inpainting problems, and Coloma Ballester for her help at several instances of this work. The second author was partially supported by a grant from the Spanish Ministerio de Ciencia y Tecnología, reference FP2000-5801 and by the Institute for Mathematics and its Applications (University of Minnesota). The first, second and fourth authors acknowledge partial support by the Departament

d'Universitats, Recerca i Societat de la Informació de la Generalitat de Catalunya and by PNPGC project, reference BFM2003-02125. The third author was partially supported by the Office of Naval Research, the National Science Foundation, the National Geospatial-Intelligence Agency, and DARPA.

References

- [1] L. Ambrosio, N. Fusco and D. Pallara, Functions of Bounded Variation and Free Discontinuity Problems, Oxford Mathematical Monographs, 2000.
- [2] N. Amenta, M. Bern, and M. Kamvysselis, A new Voronoi-based surface reconstruction algorithm, Proc. SIGGRAPH 1998, ACM, (1998) 415–421.
- [3] Amira, Amira Visualization and Modeling System, <http://www.AmiraVis.com>
- [4] G. Aronsson, Extension of functions satisfying Lipschitz conditions, Ark. for Math., 6 (1967) 551–561.
- [5] G. Aronsson, On the partial differential equation $u_x^2 u_{xx} + 2u_x u_y u_{xy} + u_y^2 u_{yy} = 0$, Ark. for Math., 7 (1968) 395-425.
- [6] C.L. Bajaj, F. Bernardini, and G. Xu, Automatic reconstruction of surfaces and scalar fields from 3D scans, Proc. ACM SIGGRAPH 1995, ACM, (1995) 109-118.
- [7] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera, Filling-in by joint interpolation of vector fields and grey levels, IEEE Trans. Image Processing 10 (2001) 1200-1211.
- [8] C. Ballester, V. Caselles, and J. Verdera, Dissocclusion by joint interpolation of vector fields and gray levels, Multiscale Model. Simul. 2 (2003) 80-123.

- [9] T. Battacharya, E. DiBenedetto and J. Manfredi, Limits as $p \rightarrow \infty$ of $\Delta_p u_p = f$ and related extremal problems, *Rendiconti Sem. Mat. Fascicolo Speciale NonLinear PDEs*, Univ. di Torino (1989) 15-68.
- [10] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin, The ball-pivoting algorithm for surface reconstruction, *IEEE Transactions on Visualization and Computer Graphics* 5 (1999) 349–359.
- [11] M. Bertalmio, A. L. Bertozzi, and G. Sapiro, Navier-Stokes, fluid dynamics, and image and video inpainting, *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, Hawaii, (2001) 355-362.
- [12] M. Bertalmío, G. Sapiro, V. Caselles, and C. Ballester, Image inpainting, *Proc. SIGGRAPH 2000*, ACM, (2000) 417–424.
- [13] M. Bertalmío, L. Vese, G. Sapiro, and S. Osher, Simultaneous texture and structure image inpainting, *IEEE Transactions on Image Processing* 12 (2003) 882-889.
- [14] J.C. Carr, R.K. Beatson, J.B. Cherrie, T.J. Mitchell, W.R. Fright, B.C. McCallum, and T.R. Evans, Reconstruction and representation of 3D objects with radial basis functions, *Proc. SIGGRAPH 2001*, ACM, (2001) 67–76.
- [15] V. Caselles, J.M. Morel and C. Sbert, An axiomatic approach to image interpolation, *IEEE Transactions on Image Processing* 7 (1998) 376–386.
- [16] T. Chan and J. Shen, Local inpainting models and TV inpainting, *SIAM J. Appl. Math.* 62 (2001) 1019-1043.

- [17] Y. Chen and G. Medioni, Description of complex objects from multiple range images using an inflating balloon model, *Computer Vision and Image Understanding*, 61 (1995) 325-334.
- [18] U. Clarenz, M. Rumpf, and A. Telea, Finite elements on point based surfaces, *Computers and Graphics*, to appear, 2005.
- [19] U. Clarenz, U. Diewald, G. Dziuk, M. Rumpf, and R. Rusu, A finite element method for surface restoration with smooth boundary conditions, *Computer Aided Geometric Design* 21 (2004) 427-445.
- [20] B. Curless and M. Levoy, A volumetric method for building complex models from range images, *Proc. SIGGRAPH 1996*, ACM, 5 (1996) 303-312.
- [21] J. Davis, S. Marschner, M. Garr, and M. Levoy, Filling holes in complex surfaces using volumetric diffusion, *First International Symposium on 3D Data Processing, Visualization, and Transmission* (2002).
- [22] J. S. De Bonet, Multiresolution sampling procedure for analysis and synthesis of texture images, *Proc. SIGGRAPH 1997*, ACM, (1997) 361-368.
- [23] H. Dinh, G. Turk, and G. Slabaugh, Reconstructing surfaces using anisotropic basis functions , *Proceedings IEEE Int. Conference on Computer Vision* (2001) 606-613.
- [24] H. Edelsbrunner and E.P. Mucke, Three-dimensional alpha shapes, *ACM Transactions on Graphics* 13 (1994) 43-72.
- [25] A. A. Efros and T. K. Leung, Texture synthesis by non-parametric sampling, *IEEE International Conference on Computer Vision*, Corfu, Greece, (1999) 1033-1038.

- [26] D. Heeger and J. Bergen, Pyramid based texture analysis/synthesis, Proc. SIGGRAPH 1995, ACM, (1995) 229–238.
- [27] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, Surface reconstruction from unorganized points, Proc. SIGGRAPH 1992, ACM, (1992) 71-78.
- [28] K. Ito and K. Kunisch, An active set strategy based on the augmented Lagrangian formulation for image restoration, M2AN Math. Model. Numer. Anal. 33 (1999) 1-21.
- [29] R. Jensen, Uniqueness of Lipschitz extensions: Minimizing the sup norm of the gradient, Arch. Rat. Mech. Anal. 123 (1993) 51–74.
- [30] P. Juutinen, Absolutely minimizing Lipschitz extensions on a metric space, Annales Academiæ Scientiarum Fennicæ Mathematica 27 (2002) 57-67.
- [31] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk, The Digital Michelangelo Project: 3D scanning of large statues, Proc. SIGGRAPH 2000, ACM, (2000) 269-276.
- [32] S. Masnou and J. Morel, Level-lines based disocclusion, IEEE Int. Conf. Image Processing (1998) 259-263.
- [33] M. Nitzberg, D. Mumford, and T. Shiotani, Filtering, Segmentation, and Depth, Springer-Verlag, Berlin, 1993.
- [34] J. M. Ogden, E. H. Adelson, J. R. Bergen, and P. J. Burt, Pyramid-based computer graphics, RCA Engineer 30 (1985) 4-15.
- [35] S. Osher and R. Fedkiw, Level set methods and dynamic implicit surfaces, Applied Mathematical Sciences, 153, Springer-Verlag, New York, 2003.

- [36] S. Osher and J.A. Sethian, Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulation, *J. Computational Physics* 79 (1988) 12-49.
- [37] S. Rane, G. Sapiro, and M. Bertalmio, Structure and texture filling-in of missing image blocks in wireless transmission and compression applications, *IEEE Trans. Image Processing* 12 (2003) 269-303.
- [38] E. Simoncelli and J. Portilla, Texture characterization via joint statistics of wavelet coefficient magnitudes, *Proc. 5th IEEE Int. Conf. on Image Processing* (1998) 62-66.
- [39] G. Turk, and M. Levoy, Zippered polygon meshes from range images, *Proc. SIGGRAPH 1994*, ACM, (1994) 311-318.
- [40] J. Verdera, V. Caselles, M. Bertalmío, and G. Sapiro, Inpainting surface holes, *IEEE International Conference on Image Processing, ICIP 2003*, Barcelona, Spain, (2003) 903-906.
- [41] M.D. Wheeler, Y. Sato, K. Ikeuchi, Consensus surfaces for modeling 3D objects from multiple range images, *Proc. IEEE Int. Conference on Computer Vision*, (1998) 917-924.
- [42] R. Whitaker, A level-set approach to 3D reconstruction from range data, *International Journal of Computer Vision* 29 (1998) 203-231.
- [43] Y. Yu, K. Zhou, D. Xu, X. Shi, H. Bao, B. Guo and H.-Y. Shum, Mesh editing with Poisson-based gradient field manipulation, *ACM Transactions on Graphics* 23 (2004) 641-648.

- [44] H.K. Zhao, S. Osher, and R. Fedkiw, Fast Surface Reconstruction using the Level Set Method, Proc. First IEEE Workshop on Variational and Level Set Methods, in conjunction with Proc. IEEE ICCV (2001) 194-202.