

INPAINTING SURFACE HOLES

Joan Verdera * Vicent Caselles † Marcelo Bertalmio ‡ and Guillermo Sapiro §

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

ABSTRACT

An algorithm for filling-in surface holes is introduced in this paper. The basic idea is to represent the surface of interest in implicit form, and fill-in the holes with a system of geometric partial differential equations derived from image inpainting algorithms. The framework and examples with synthetic and real data are presented.

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 [5, 10], or hinted by the user as in more classical texture filling techniques [9, 11, 17]. Several names have been used for this filling-in operation: *disocclusion* in [2, 13], or *inpainting* in [4, 5, 6]. In the context of this paper, and following [5], 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 [12], which in part due to the complicated geometry have 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., prohesthetics), 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 an algorithm for inpainting these surface holes.

Our work is inspired by the one reported in [8], 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.

The algorithm here proposed is an extension of our previous work on image inpainting [2, 3, 5] (see also [4, 6, 7, 13, 15, 16]). In particular, we show how to adapt the variational formulation we presented in [2, 3] to the problem of surface hole filling. As in [8], 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 [8], 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 [8], are applied only at the holes and a neighborhood of them (being these equation anisotropic and geometry based, they lead to

*Dept. de Tecnologia, Universitat Pompeu-Fabra, Passeig de Circumvalació, 8, 08003 Barcelona, Spain, joan.verdera@tecn.upf.es

†Dept. de Tecnologia, Universitat Pompeu-Fabra, Passeig de Circumvalació, 8, 08003 Barcelona, Spain, vicent.caselles@tecn.upf.es

‡Dept. de Tecnologia, Universitat Pompeu-Fabra, Passeig de Circumvalació, 8, 08003 Barcelona, Spain, marcelo.bertalmio@tecn.upf.es

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



Fig. 1. A rendering of a scanned version of Michelangelo's David (Michelangelo's project)

a slightly slower algorithm than the one reported in [8]). We present below both the underlying theory and examples for real data.

2. A VARIATIONAL APPROACH TO INPAINTING IN N DIMENSIONS

We now describe a variational approach to the joint interpolation of vector fields and gray values. For more details, see [2, 3]. In the following section we will see how we adapt this to the problem of hole filling on surfaces.

Let Q be a hyperrectangle in R^N and Ω be an open bounded subset of Q with smooth boundary. Suppose that we are given an image $u_0 : Q \setminus \bar{\Omega} \rightarrow R$, where $\bar{\Omega}$ denotes the closure of Ω . Using the information of u_0 on $Q \setminus \bar{\Omega}$ 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 [2, 3] the authors proposed to fill-in the hole Ω using both the gray level and the vector field of normals to the level sets of the image outside the hole. Let $\tilde{\Omega}$ be an open subset of Q with smooth boundary such that $\bar{\Omega} \subset \tilde{\Omega}$. The band around Ω will be the set $B = \tilde{\Omega} \setminus \bar{\Omega}$. To fill-in the hole Ω we shall 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 shall assume that θ_0 is a vector field with values in R^N satisfying $\theta_0(x) \cdot \nabla u_0(x) = |\nabla u_0(x)|$ and $|\theta_0(x)| \leq 1$. In practice we take $\theta_0(x) = \frac{\nabla u_0(x)}{|\nabla u_0(x)|}$ when $\nabla u_0(x) \neq 0$, and $\theta_0(x) = 0$ if $\nabla u_0(x) = 0$. The basic goal then is to smoothly extend the pair (u_0, θ_0) from the band $B = \tilde{\Omega} \setminus \bar{\Omega}$ to a pair of functions (u, θ) inside Ω . For that we attempt to continue the isosurfaces of u_0 (i.e the hypersurfaces $[u_0 = \lambda]$) or, more

generally, the boundaries of the level sets $[u_0 \geq \lambda]$, $\lambda \in R$) in B inside Ω by taking into account the principle of good (smooth) continuation. As we already said, we should constraint the solution (u, θ) to coincide with the data on B . The vector field θ should satisfy $|\theta| \leq 1$ on Ω and should be related to u by $\theta \cdot \nabla u = |\nabla u|$, i.e., we should impose that θ is the vector field of directions of the gradient of u . Finally, we should impose that the vector field θ_0 in the band B is smoothly continued by θ inside Ω . Note that if θ represents the directions of the normals to the level hypersurfaces of u , then $div(\theta)$ represents its mean curvature. The smooth continuation of the levels sets of u_0 inside Ω is imposed by requiring that $div \theta \in L^p(\tilde{\Omega})$, i.e., $\int_{\tilde{\Omega}} |div \theta|^p dx < \infty$. For consistency we shall require that $div \theta_0 \in L^p(B)$.

We propose to interpolate the pair (θ, u) in Ω by minimizing the functional

$$\text{Minimize } \int_{\tilde{\Omega}} |div(\theta)|^p (\gamma + \beta |\nabla k * u|) dx$$

$$u = u_0, \theta = \theta_0 \text{ in } B, |\nabla u| - \theta \cdot \nabla u = 0 \text{ in } \tilde{\Omega} \quad (2.1)$$

$$|\theta| \leq 1, \|u\| \leq M$$

where $p > 1$, $\gamma > 0$, $\beta \geq 0$, k denotes a regularizing kernel of class C^1 such that $k(x) > 0$ a.e., $M = \|u_0\|_{L^\infty(B)}$. Let us note that if u is the characteristic function of a set $A \subseteq R^3$ (i.e. $u(x) = \chi_A(x) = 1$ if $x \in A$ and $= 0$ otherwise) with smooth boundary and θ is an smooth extension of the unit normal to ∂A , then $\int_{\tilde{\Omega}} |div(\theta)|^p |\nabla u| dx = \int_{\partial A} |H|^p dS$ where $H(x)$ is the mean curvature of ∂A and dS denotes the surface area. Let us say that for mathematical reasons we have convolved the ∇u term of (2.1) to be able to ensure the existence of a minimum for (2.1) [3, 2]. Finally, let us also note that the constant γ has to be > 0 , otherwise we do not get an L^p bound on $div \theta$. We refer to [2] for a detailed theoretical analysis of this formulation and its approximation by smoother functionals. Let us finally note that the use of smooth continuation using powers of the curvature as smoothness method was proposed in [14] with the purpose of image segmentation with depth (reconstructing then the occluded boundaries), and used as a base for variational approaches of disocclusion in [2, 3, 13].

The minimization of (2.1) is given by the gradient descent method, thereby obtaining a system of two coupled PDE's, one for u and one for the vector field θ :

$$\theta_t = -\nabla_{\theta} E(\theta, u) \quad u_t = -\nabla_u E(\theta, u) \quad \text{in } \tilde{\Omega}, \quad (2.2)$$

(where $E(\theta, u)$ denotes the energy defined in (2.1)) supplemented with the corresponding boundary data and initial conditions. The constraints on (θ, u) can be incorporated either by penalization or by projection after each time step. In our experiments, we take k to be a Gaussian kernel with

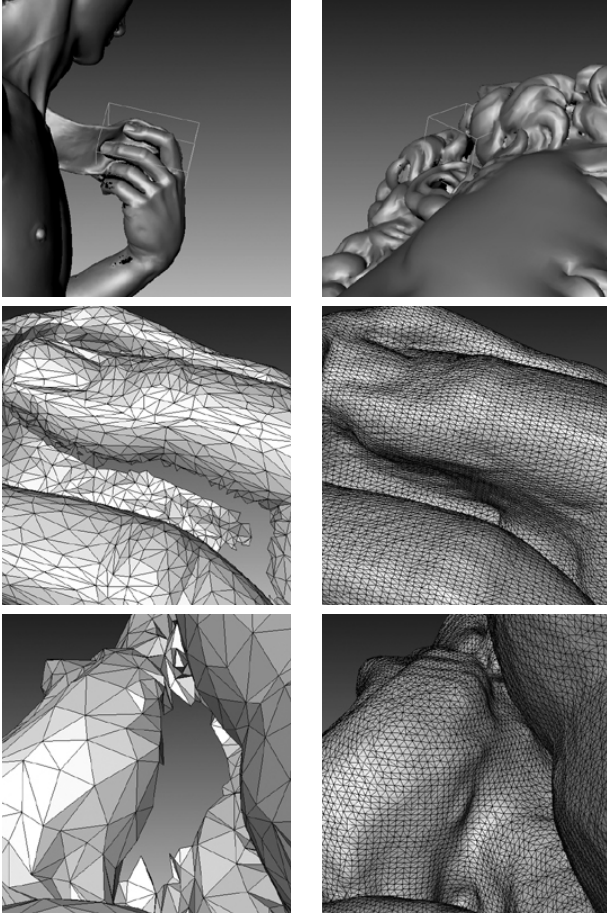


Fig. 2. From top to bottom and left to right: a) David's left hand. b) A detail of its hair. c) A zoomed detail of a) showing the triangulated surface with the hole. d) The reconstruction of the hole in c) displayed as a triangulated surface. e) A zoomed detail of b) showing the triangulated surface with the hole. f) The reconstruction of the hole in e) displayed as a triangulated surface.

small variance, say one or two pixels. In practice, one can also dismiss the kernel k . The initial conditions are ad-hoc interpolations, for instance, we can take u inside Ω as the average value of u_0 in B , θ inside Ω being the direction of the gradient of u .

3. SURFACE INPAINTING

The above algorithm can be used when $u_0 : Q \setminus \overline{\Omega} \rightarrow R$ is the characteristic function of a set $A \subseteq R^3$ which is only known in $Q \setminus \overline{\Omega}$. In this case, by minimizing (2.1), we want to reconstruct the set A inside the hole Ω knowing the set outside Ω , hence knowing also the unit normal θ_0 at the part of its boundary $\partial A \setminus \overline{\Omega}$. As noted in [2, 3], the vector field θ_0 can be extended to all the known part of A , i.e., to $A \setminus \overline{\Omega}$.




Fig. 3. From top to bottom and left to right: a) David's left foot. b) A detail of the fingers with a hole. c) The reconstruction of b) displayed as a triangulated surface. d) A rendering of the same result.

Let us describe how to use the above formulation to inpaint (fill-in) holes (or gaps) on surfaces \mathcal{S} , which we assume to be embedded in 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. Let \mathcal{S}^* be the known part of \mathcal{S} . Let us choose a bounding box Q in R^3 strictly containing the gap \mathcal{M} and part of \mathcal{S}^* (see Figs. 2.a, 2.b, 3.a). If the surface \mathcal{S} had no gap then we would consider as the set A the part of the interior of \mathcal{S} inside Q . If \mathcal{S} has a gap, then we construct a surface \mathcal{T} such that $\mathcal{S}' := \mathcal{S}^* \cup \mathcal{T}$ separates R^3 into its interior and its exterior, and we define A as the part of the interior of \mathcal{S}' inside Q . We take as hole Ω a ball containing the gap \mathcal{M} and the added surface \mathcal{T} , and let $B = Q \setminus \overline{\Omega}$ be the band where the known information of \mathcal{S} is used. This information consists in $u_0 = \chi_{A \cap B}$ and the normal vector field θ_0 to \mathcal{S}^* in B . We constrain $u = u_0$ and $\theta = \theta_0$ in B and we also constrain u to be 1 on the part of \mathcal{S}^* contained in Ω . Then we minimize (2.1) by solving the equations (2.2) where the initial condition for u is χ_A . The surface \mathcal{T} can be taken as part of the surface of the smallest ball containing the boundary of the gap (hence, the gap). The ball Ω can be taken as a slight dilation of the smallest ball containing the gap. We shall assume that this construction is possible for the gaps considered here. Otherwise, we should produce a similar construction.

4. EXPERIMENTAL RESULTS

With the purpose of adapting them to our algorithm, the data, originally given as a triangulated surface, were converted to an implicit representation in a regularly spaced 3D grid. The result was visualized again as a triangulated surface. Figure 1 displays a rendering of a scanned version of Michelangelo's David which has several holes. Figures 2.a, 2.b, 3.a display some particular holes with a bounding box isolating them. Figures 2.c, 2.e, 3.b display the triangulated surface (the data) around the hole. The reconstructed surface is displayed in Figures 2.d, 2.f, 3.c. Figure 3.d displays a rendering of Figure 3.c. These images have been rendered using the AMIRA Visualization and Modeling System [1].

5. CONCLUSIONS AND FUTURE DIRECTIONS

In this note we have shown how to extend our previous work on image inpainting to fill-in surface holes. The idea, inspired by [8] and [2, 3], is to represent the surface of interest as the zero level-set of a function, and then run a system of coupled geometric partial differential equations that smoothly continue the surface into the hole.

As we see it, this work provides a first step into the use of geometric (anisotropic) PDE's for surface hole filling. In a future work, we plan to investigate the possibility of adapting to this problem other techniques that have been shown to be very efficient for image inpainting. As in [8], the line-of-sight constraints should be introduced into our algorithm as well to help us in a more automatic initialization.

Acknowledgments. 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 myself 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. We also thank Oliver Sander for helping us with AMIRA. The first three 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 BFM2000-0962-C02-01. The third author was also partially supported by the Programa Ramon y Cajal (Ministerio de Ciencia y Tecnología, Spain). The fourth author was partially supported by the Office of Naval Research, the Presidential Early Career Awards for Scientists and Engineers (PECASE), and National Science Foundation CAREER Award.

6. REFERENCES

- [1] Amira, "Amira Visualization and Modelling System," <http://www.AmiraVis.com>
- [2] C. Ballester, V. Caselles, and J. Verdera, "Disocclusion by joint interpolation of vector fields and gray levels," pre-print, 2002.
- [3] 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**, pp. 1200-1211, August 2001.
- [4] 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, December 2001.
- [5] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," *Computer Graphics (SIGGRAPH 2000)*, pp. 417-424, July 2000.
- [6] M. Bertalmio, L. Vese, G. Sapiro, S. Osher, "Simultaneous texture and structure image inpainting," pre-print, July 2002 (mountains.ece.umn.edu/guille/inpainting.htm).
- [7] T. Chan and J. Shen, "Local inpainting models and TV inpainting," *SIAM J. Appl. Math.* **62:3**, pp. 1019-1043, 2001.
- [8] 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*, June 2002.
- [9] J. S. De Bonet, "Multiresolution sampling procedure for analysis and synthesis of texture images," *Proceedings of ACM SIGGRAPH*, July 1997.
- [10] A. A. Efros and T. K. Leung, "Texture synthesis by non-parametric sampling," *IEEE International Conference on Computer Vision, Corfu, Greece*, pp. 1033-1038, Sept. 1999.
- [11] D. Heeger and J. Bergen, "Pyramid based texture analysis/synthesis," *Computer Graphics (SIGGRAPH 1995)*, pp. 229-238, July 1995.
- [12] 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," *Computer Graphics (SIGGRAPH 2000 Proceedings)*, July 2000.
- [13] S. Masnou and J. Morel, "Level-lines based disocclusion," *IEEE Int. Conf. Image Processing*, October 1998.
- [14] M. Nitzberg, D. Mumford, and T. Shiota, "Filtering, Segmentation, and Depth," Springer-Verlag, Berlin, 1993.
- [15] J. M. Ogden, E. H. Adelson, J. R. Bergen, and P. J. Burt, "Pyramid-based computer graphics," *RCA Engineer* **30(5)**, pp. 4-15, 1985.
- [16] 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*, to appear.
- [17] E. Simoncelli and J. Portilla, "Texture characterization via joint statistics of wavelet coefficient magnitudes," *Proc. 5th IEEE Int'l Conf. on Image Processing*, 1998.