Dual Domain Extrapolation

Bruno Lévy – ISA - INRIA Lorraine

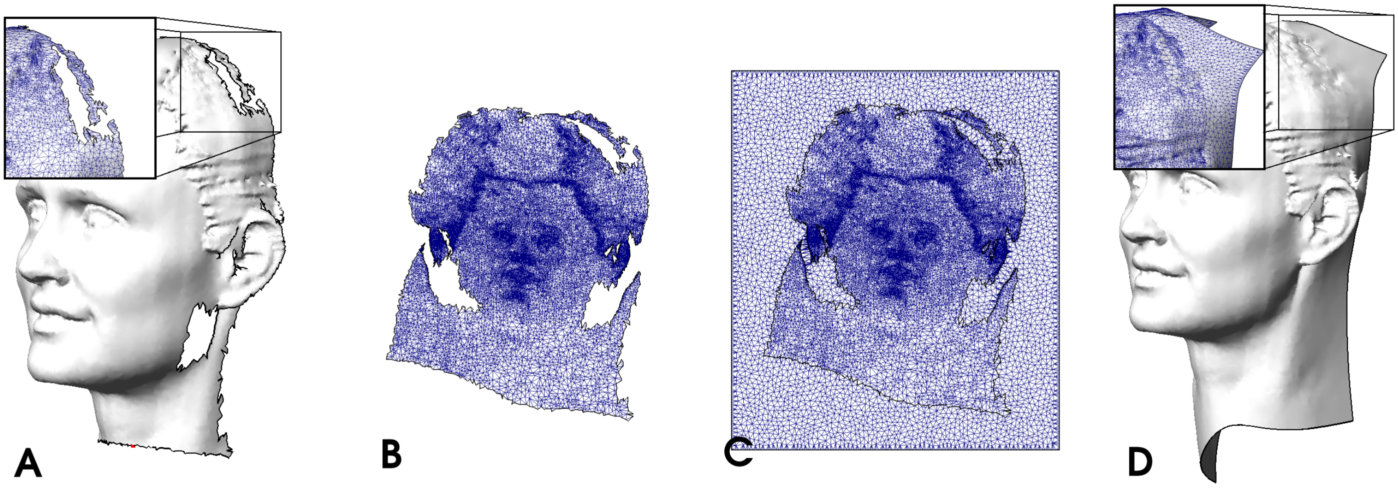

A: Due to shadows, scanned meshes often have complex holes and irregular borders. B: a global parameterization of the surface is computed. C: filling the holes and extrapolating the borders become 2D problems in parameter space. D: The locations of the new vertices in 3D space are computed by approximating a minimal energy surface (MES).

Abstract

Shape optimization and surface fairing for polygon meshes have been active research areas for the last few years. Existing approaches either require the border of the surface to be fixed, or are only applicable to closed surfaces. In this paper, we propose a new approach, that computes natural boundaries. This makes it possible not only to smooth an existing geometry, but also to extrapolate its shape beyond the existing border. Our approach is based on a global parameterization of the surface and on a minimization of the squared curvatures, discretized on the edges of the surface. The so-constructed surface is an approximation of a minimal energy surface (MES). Using a global parameterization makes it possible to completely decouple the outer fairness (surface smoothness) from the inner fairness (mesh quality). In addition, the parameter space provides the user with a new means of controlling the shape of the surface. When used as a geometry filter, our approach computes a smoothed mesh that is discrete conformal to the original one. This allows smoothing textured meshes without introducing distortions.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Boundary representations.

Keywords: shape optimization, polygon meshes, differential geometry, minimal energy surfaces, natural boundaries

1 Introduction

Geometric Design has defined tools to process and create geometry. The following three operators are widely used to optimize the shape of surfaces:

- ◇ **Surface fairing** used, for instance, to remove the high-frequency noise from a scanned surface;
- ◇ **Surface blending**, whose goal is to create a surface that smoothly connects an existing set of patches;
- ◇ **Surface extrapolation**, used to extend the shape of an existing surface.

Using differential geometry, it is possible to define the criteria that should be met by the constructed surfaces. This was initially applied to parametric surfaces (see e.g. [Moreton and Séquin 1992]).

Mesh models are among the most common representations of surface geometry in Computer Graphics, and adapting geometric design tools to this type of representation has been an active research area for the last few years (see e.g. [Kobbelt 1997]). In the specific case of a mesh model, not only the shape of the surface, referred to as the *outer fairness*, should be considered, but also the quality of the triangulation, referred to as the *inner fairness*.

Surface fairing and surface blending have been well studied for mesh models, and have now efficient discrete counterparts. However, existing discrete fairing methods either need the border of the surface to be fixed, or can only be applied to closed surfaces. This prevents these methods from being uses as surface extrapolators. In this paper, we propose a new discrete fairing method that does not suffer from this limitation, which allows implementing new geometric design tools for meshed models.

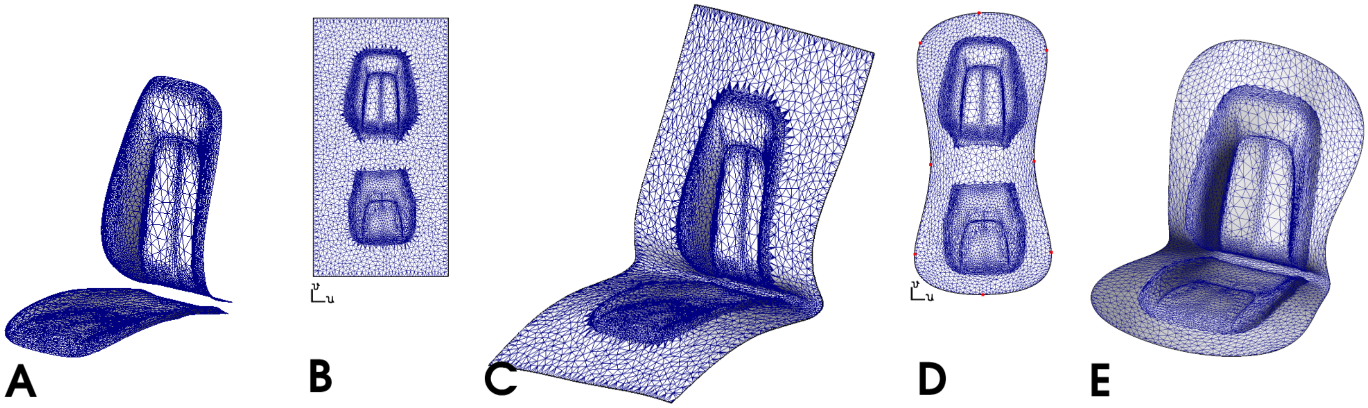

Figure 1: *Extrapolation and border editing in parameter space. A: original surface; B: constrained triangulation of the parameter space; C: extrapolating a minimal energy surface; D: the border has been edited and the constrained triangulation reconstructed; E: resulting surface: changing the border in parameter space means cropping a "virtual" surface, extrapolated from the initial surface.*

1.1 Previous Work

Mesh fairing methods can be classified into two main categories, geometric filtering and geometric optimization:

In **geometric filtering** approaches, the goal is often to remove high frequencies from an initial mesh. *Discrete diffusion* is proposed in [Taubin 1995], in which signal processing provides an elegant mathematical background to study these problems. Discrete diffusion has been then improved in [Desbrun et al. 1999] by introducing a more stable smoothing operator, which allows using larger values for the smoothing parameter. The method proposed in [Ohtake et al. 2000] combines discrete diffusion with an inner fairness criterion, in order to improve the quality of the mesh. All the methods mentioned above are very useful for filtering noisy geometries, but since only G^0 continuity is ensured, it is difficult to use them for surface design. Higher order continuity can be obtained by combining two smoothing steps [Taubin 1995], or using curvature flows of higher orders [Brakke 1992; Hsu et al. 1992; Chopp and Sethian 1999].

In **geometric optimization** approaches (see e.g. [Burchard et al. 1994]), global *fairness* criteria are defined and optimized. Fairness is often defined using notions from differential geometry (mean curvature, Gaussian curvature ...) or approximation of physics (thin-plate energy). This family of methods has been first defined for parametric surfaces, for which optimizing the fairness means solving a Partial Differential Equation [Bloor and Wilson 1990]. The differential operators involved in these PDEs can be approximated on polygonal meshes. In [Pinkall and Polthier 1993], a method is proposed to construct discrete minimal surfaces (i.e. surfaces of minimal area) by minimizing the Dirichlet energy (i.e. the integral of the squared Laplacian) relative to a conformal parameterization. The methods proposed in [Mallet 1992] and in [Kobbelt et al. 1998] also minimize an approximation of the Dirichlet energy. Note that for these two methods, the used local parameterizations are not necessarily conformal, which means the Dirichlet energy does not necessarily correspond to the area of the surface. In [Welch and Witkin 1994], an approximation of curvature is defined for triangulated surfaces, together with a minimization algorithm. In [Schneider and Kobbelt 2000; Schneider and Kobbelt n. d.], they propose to minimize the Laplacian of the mean curvature, an intrinsic geometric value measuring the variation of the curvature. Kobbelt coined the term *Discrete Fairing* in [Kobbelt 1997] to categorize the family of methods based on the optimization of differential operators approximated on meshes.

For both families of methods, the Laplacian operator plays an important role. For instance, the Dirichlet energy is the integral of the squared Laplacian. To approximate the Laplacian on a polygonal mesh, several numerical schemes have been proposed, based on local parameterizations [Taubin 1995; Kobbelt et al. 1998; Desbrun et al. 1999; Guskov et al. 1999].

Note that meshed models can also be obtained from unorganized set of points by using reconstruction methods and level set representations. The methods described in [Boissonnat and Cazals 2002], [Museth et al. 2002] and [Carr et al. 2001] reconstruct a function from $\mathbb{R}^3$ to $\mathbb{R}$ fitted to set of points. The reconstructed surface is obtained as the zero-set of this function, extracted by applying marching-cube-like methods. As in these two latter methods, our approach minimizes an objective function based on differential operators. The main difference is that it operates on the meshed model directly rather than fitting an intermediary representation.

1.2 Overview

In this paper, we present the Dual Domain Extrapolation method, a new mesh fairing approach with the following characteristics:

- ◇ A discrete approximant of the Minimal Energy Surfaces (MES) criterion is defined and efficiently minimized thanks to Newton's multivariate non-linear optimization method;
- ◇ The method does not need any boundary condition. Therefore, the border of the surface is naturally extrapolated;
- ◇ Thanks to a global parameterization, it is possible to completely dissociate the outer fairness (surface quality) from the inner fairness (mesh quality);
- ◇ The global shape of the surface is not affected by the discretization;
- ◇ Using a global parameterization provides the user with new means of designing the shape of the surface. By designing the shape of the border in parameter space, the user trims a virtual surface, extrapolated from the initial surface.
- ◇ Any vertex of the initial surface may be constrained or let free to move. The method can be used as an extrapolator as well as a low pass filter;

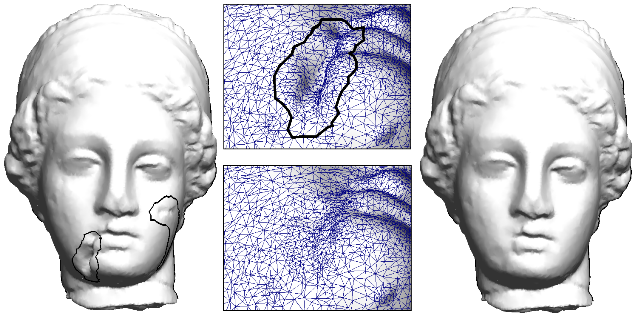

Figure 2: *DDE used to smooth a mesh in user-selected zones. The mesh in the smoothed regions is discrete conformal to the original one.*

- When our method is used as a low pass filter, the smoothed surface is discrete conformal to the original one, which means textured models can be filtered without introducing any texture deformation;
- *Limitations of the method:* since it is based on a global parameterization with natural boundaries, our method is only applicable if such a parameterization can be constructed. The surface needs to be homeomorphic to a disc (however, a user-assisted solution is proposed for surfaces of higher genus), and global overlaps need to be avoided (this will be discussed in Section 2.1).

The paper is organized as follows. The next section introduces the approximation scheme used to compute the curvatures, and the objective function minimized by our method. Section 3 presents some results obtained with our method used as an extrapolator, a low-pass filter, or as a geometric design tool. The paper concludes with possible improvements of the method.

2 Dual Domain Extrapolation

Minimal Energy Surfaces (MES) are defined to be surfaces that minimize the following objective function:

$$E_{MES} = \int_{\Omega} \kappa_{min}^2 + \kappa_{max}^2 \, dudv \qquad (1)$$

Our goal is to define a discrete approximant of this energy with the possibility of optimizing the geometry of a triangulated surface without any boundary condition. To approximate the differential operators involved in this type of energy, existing mesh fairing approaches use local parameterizations. For instance, to approximate the Dirichlet energy, the umbrella operator used in [Kobbelt et al. 1998] is defined at each vertex relative to a parameterization of its 1-ring neighbors. Unfortunately, the shape of the obtained surface is dependent on the expression of these local parameterizations. Moreover, they introduce an unwanted coupling between the outer fairness (surface quality) and the inner fairness (mesh quality). This can be improved by using more elaborated numerical schemes [Schneider and Kobbelt 2000; Schneider and Kobbelt n. d.].

To overcome these problems, instead of using local parameterizations, our approach is based on a global parameterization of the surface. Moreover, expressing the MES criterion as a coupling between global (u,v) coordinates and geometric (x,y,z) coordinates makes it unnecessary to fix the boundaries of the surface. A global parameterization was already used to define very flexible remeshing algorithms in [Alliez et al. 2002], in which re-triangulation is performed in parameter space. Our method is also based on the idea of using the parameter space,

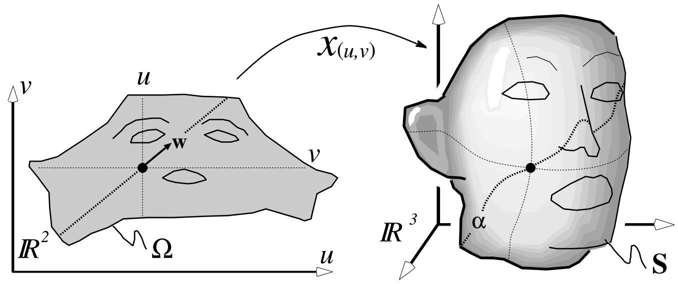

Figure 3: *Computing directional curvatures on a surface.*

with the difference that triangles are created beyond the border of the parameter space, or in zones corresponding to holes of the original surface, where no geometry is known *a priori*. Optimizing the location of these vertices in 3D space is then performed by minimizing our discrete MES criterion.

Our DDE method comprises the following steps:

1. Construct a global parameterization of the surface,
2. (optional) select the zones of the surface that should be smoothed (see Figure 2),
3. (optional) Fill-in the holes and/or extend the border, by adding triangles in parameter space, using a Constrained Delaunay Triangulation (Figure 1-B).
4. Make the mesh regular, by iteratively inserting a new vertex at the circumcenter of the largest triangle (the Delaunay condition and constraints are restored by edge swapping), until the area of the largest triangle is equal to the average area of the initial triangles,
5. Optimize the new vertices and the vertices belonging to a zone selected in step 2, by minimizing our discrete approximant of the MES criterion (Figure 1-C),
6. Remove skinny triangles, by inserting a vertex in the triangle with the highest *area in 3D / area in parameter space* ratio (as in step 4), until the highest ratio reaches the average ratio of the initial triangles,
7. (optional) Interactively edit the boundary curve in parameter space (see Figure 1-D). After each modification, steps 3,4,5,6 (triangulation and optimization) are re-executed (Figure 1-E).

2.1 Global parameterization

In our context, to construct a global parameterization of the surface (step 1), it is necessary to use a method that can compute natural boundaries, such as MIPS [Hormann and Greiner 2000], ABF [Sheffer and de Sturler 2001], LSCM [Lévy et al. 2002] or DCP [Desbrun et al. 2002]. Unfortunately, these methods do not offer the same guarantees as [Floater 1997], which fixes the border on a convex polygon. In our case, depending on the used parameterization method, overlaps and triangle flips may be encountered. ABF guarantees the absence of triangle flips, and provides a way of fixing global overlaps, by adding constraints. However, ABF is numerically expensive, and difficult to apply to large meshes. For this reason, in our experiments, we have used LSCM, which is much easier to solve numerically (note that the same result would be obtained with DCP, since it minimizes the same energy as LSCM). We did not encounter overlaps with our examples, but if they occur, it is easy to detect them and to switch to ABF.

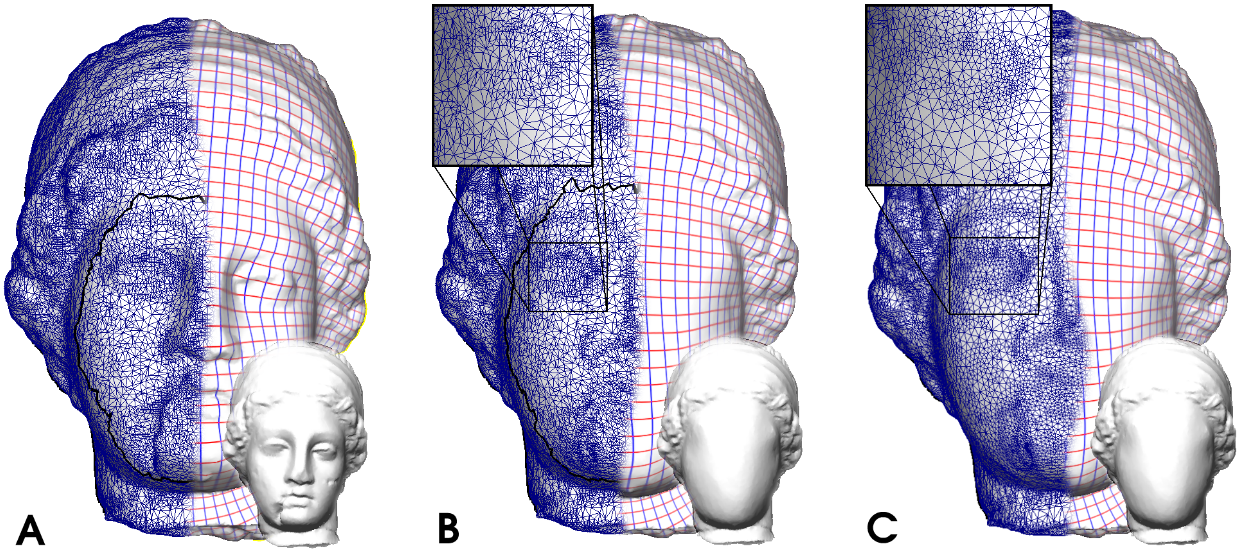

Figure 4: *Outer and inner fairness optimization. A: original mesh with selected zone; parameterization; B: smoothed mesh obtained by minimizing the discrete MES criterion. The triangulation is discrete conformal to the original surface; parameterization of the smoothed mesh; C: result obtained after optimizing the inner fairness (mesh quality) in parameter space and re-optimizing the discrete MES optimization relative to this new parameterization. Note that the outer and the inner fairness are completely decoupled.*

2.2 Curvature approximation

Different approximation schemes for surface curvature were proposed in the literature. Most of them consider the 1-ring neighborhoods of the vertices, and compute the curvatures from quadratic surfaces fitted to the neighborhoods [Welch and Witkin 1994], or directly apply least squares fitting to the eigen structure of the second fundamental form [Moreton and Séquin 1992], which avoids ill-conditioning. In our approach, we consider the neighborhoods of the edges, composed of two triangles, and approximate the second order derivatives thanks to finite differences. The resulting numerical scheme is much simpler to compute, and does not require any boundary condition.

In this section, we consider a surface $\mathbf{S}$ provided with a parameterization $\mathbf{x}(.,.) : \Omega \subset \mathbb{R}^2 \rightarrow \mathbf{S}; (u,v) \mapsto \{x(u,v), y(u,v), z(u,v)\}$. The directional curvature at a point $\mathbf{u} = (u,v)$ relative to a direction $\mathbf{w}$ in parameter space is defined to be the curvature of the curve $\alpha(.) : \alpha(t) = \mathbf{x}(\mathbf{u} + t.\mathbf{w})$ (see Figure 3, previous page). It is well known (see e.g. [do Carmo 1976]) that the curvature is given by the norm of the second order derivatives of $\alpha(.)$, parameterized by arc-length s:

$$\kappa_{\mathbf{w}}(\mathbf{u}) = \left\| \frac{\partial^2 \alpha}{\partial s^2}(0) \right\| \tag{2}$$

In our case, the surface $\mathbf{S}$ is a triangulated mesh, and its parameterization $\mathbf{x}(.,.)$ is a piecewise linear function. For such a surface, we propose to approximate the second order derivatives using finite differences between first order derivatives:

$$\begin{aligned} \kappa_{\mathbf{w}}(\mathbf{u}) &\simeq \frac{1}{2\delta} \left\| \frac{\partial \alpha}{\partial s}(\delta) - \frac{\partial \alpha}{\partial s}(-\delta) \right\| \\ &= \frac{1}{2\delta} \left\| \frac{\partial \alpha}{\partial t} \cdot \frac{\partial t}{\partial s}(\delta) - \frac{\partial \alpha}{\partial t} \cdot \frac{\partial t}{\partial s}(-\delta) \right\| \\ &= \frac{1}{2\delta} \left\| \frac{J(\delta).\mathbf{w}}{\|J(\delta).\mathbf{w}\|} - \frac{J(-\delta).\mathbf{w}}{\|J(-\delta).\mathbf{w}\|} \right\| \end{aligned} \tag{3}$$

where $J(\delta)$ denotes the Jacobian matrix of $\mathbf{x}(.,.)$ (i.e. the matrix of the differential $d\mathbf{x}$) at the point $\mathbf{u} + \delta.\mathbf{w}$, given by:

$$J = \begin{pmatrix} \partial x/\partial u & \partial y/\partial u & \partial z/\partial u \\ \partial x/\partial v & \partial y/\partial v & \partial z/\partial v \end{pmatrix}^t \tag{4}$$

As shown in Figure 5 below, we consider now an edge e shared by two triangles T and T'. In 3D, the points lying on the edge e can be considered as cylindrical points (having a null minimum curvature κ_{min}). In parameter space, the principal directions associated with κ_{min} and κ_{max} are $\mathbf{w}_1$ and $\mathbf{w}_2$ respectively, where $\mathbf{w}_1$ is aligned with the edge e in parameter space, and $\mathbf{w}_2$ is orthogonal to $\mathbf{w}_1$. Note that since $\mathbf{x}(.,.)$ is a piecewise linear parameterization, the Jacobian matrix J is constant over each triangle T , and will be denoted J_T in what follows. The finite difference approximation of the maximum curvature κ_{max} is then given by:

$$\begin{aligned} \kappa_{max} = \kappa_{\mathbf{w}_2}(\mathbf{u}) &\simeq \frac{1}{2\delta} \left\| \frac{J_T.\mathbf{w}_2}{\|J_T.\mathbf{w}_2\|} - \frac{J_{T'}.\mathbf{w}_2}{\|J_{T'}.\mathbf{w}_2\|} \right\| \\ J_T &= \frac{1}{2\mathscr{A}(T)} \begin{pmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \\ z_1 & z_2 & z_3 \end{pmatrix} \begin{pmatrix} (v_2 - v_3) & (u_3 - u_2) \\ (v_3 - v_1) & (u_1 - u_3) \\ (v_1 - v_2) & (u_2 - u_1) \end{pmatrix} \end{aligned} \tag{5}$$

where $\mathscr{A}(T)$ denotes the area of the triangle T in parameter space.

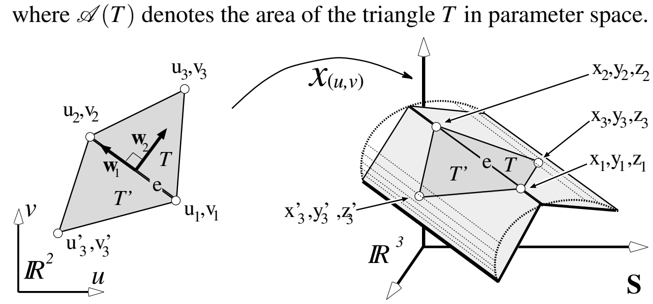

Figure 5: *Approximating the curvature around an edge e. The points of e can be considered to be cylindrical, with a null curvature κ_{min} associated with the direction $\mathbf{w}_1$ of e, and a maximum curvature κ_{max} associated with a direction $\mathbf{w}_2$, orthogonal to $\mathbf{w}_1$*

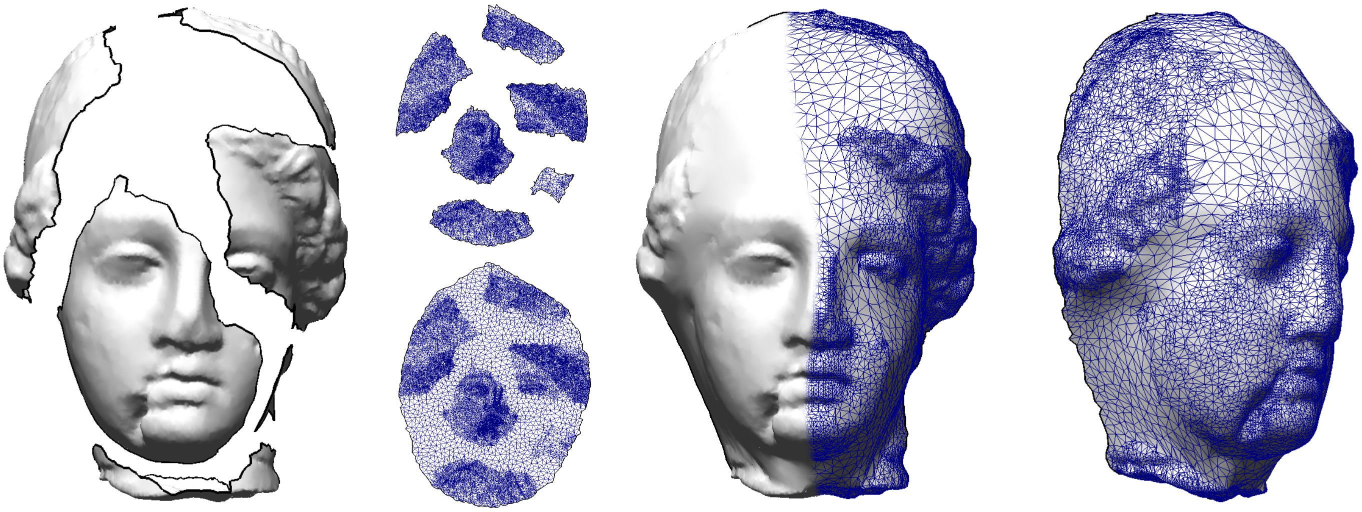

Figure 6: *Generating a complex blending surface. The charts are parameterized, the user arranges them in parameter space, a constrained Delaunay triangulation is computed, and the new vertices are optimized by minimizing the discrete MES criterion.*

2.3 Discrete MES optimization

The expression of our discrete MES energy F is then obtained by summing the approximation of the curvatures (Equation 5) over the edges e of the triangulation, weighted by the areas of the triangles T and T' in parameter space.

$$E_{MES} \simeq F(x) = \frac{1}{6\delta} \sum_{e \in \mathscr{E}} \frac{1}{\mathscr{A}(T)+\mathscr{A}(T')} \left\| \frac{J_T.\mathbf{w}_2(e)}{\|J_T.\mathbf{w}_2(e)\|} - \frac{J_{T'}.\mathbf{w}_2(e)}{\|J_{T'}.\mathbf{w}_2(e)\|} \right\|^2$$

$$e = (i,j) \quad ; \quad \mathbf{w}_2(e) = \begin{bmatrix} v_i - v_j \\ u_j - u_i \end{bmatrix} \tag{6}$$

To minimize the objective function F, we have experimented different strategies (see e.g. [Nocedal and Wright 2000]). Their perfomances are compared in section 3:

◇ *Newton's method:* this requires the gradient ∇F and the Hessian $\nabla^2 F$ of the function. For $\nabla^2 F$, our implementation uses expressions computed in Maple;

◇ *Quasi-Newton BFGS method:* only requires ∇F ($\nabla^2 F$ is approximated). This gives the best performaces. However, BFGS is quite difficult to implement;

◇ *SQP (Sequential Quadratic Programming):* at each iteration k, we consider the terms $||J_T.\mathbf{w}_2(e)||$ and $||J_{T'}.\mathbf{w}_2(e)||$ to be constant. They are initialized to 1 at the first iteration. This approximate of $\nabla^2 F$ is less accurate than BFGS, and therefore less efficient, but much easier to implement.

2.4 Inner fairness optimization

As shown in Figure 1, moving the border in parameter space means letting it 'slide' along a virtual surface, extrapolated from the initial patch. Based on this idea, it is possible to optimize the inner fairness in parameter space (see Figure 4):

1. choose a region of interrest and lock its border;
2. optimize the shape of the triangles in 2D parameter-space, using Laplacian smoothing;
3. optimize the MES criterion, using the smoothed 2D triangulation as the parameter-space

The vertices "slide" along the virtual smooth surface, while the global geometry is not altered. The inner fairness is completely decoupled from the outer fairness.

3 Results and conclusions

Figures 6, 7, 8 show possible applications of DDE. The tool demonstrated in Figure 6 may have applications in archaeological modeling, providing the user with a flexible and intuitive interface to repair surfaces (parameter-space "jigsaw"). Table 1 shows the size of the models, computation times, and number of iterations (the time spent in the constrained Delaunay triangulation is not included since it is negligible).

The principal limitation of our method is that it is unable to directly process a model of arbitrary topology (closed models, donuts, . . .). A user-assisted solution to this problem is proposed in Figure 8. More generally, it would be possible to apply our method to an automatically generated atlas, by adding inter-chart continuity conditions to the objective function. We think it is also possible (and better) to use local parameterization satisfying compatibilities conditions, obtained thanks to a global optimization process.

As compared to the method presented in [Schneider and Kobbelt n. d.], the class of surfaces minimizing the MES criterion is not as general as the one minimizing the Laplacian of the mean curvature.

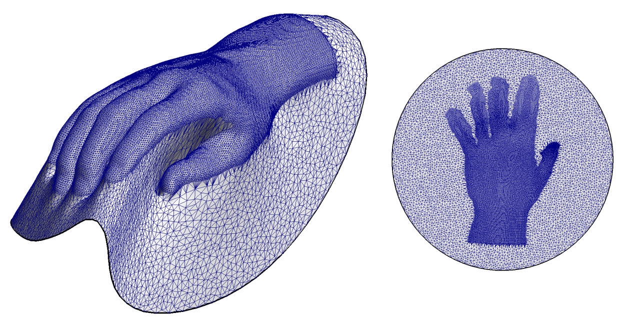

Figure 7: *Dual domain extrapolation and associated parameter space.*

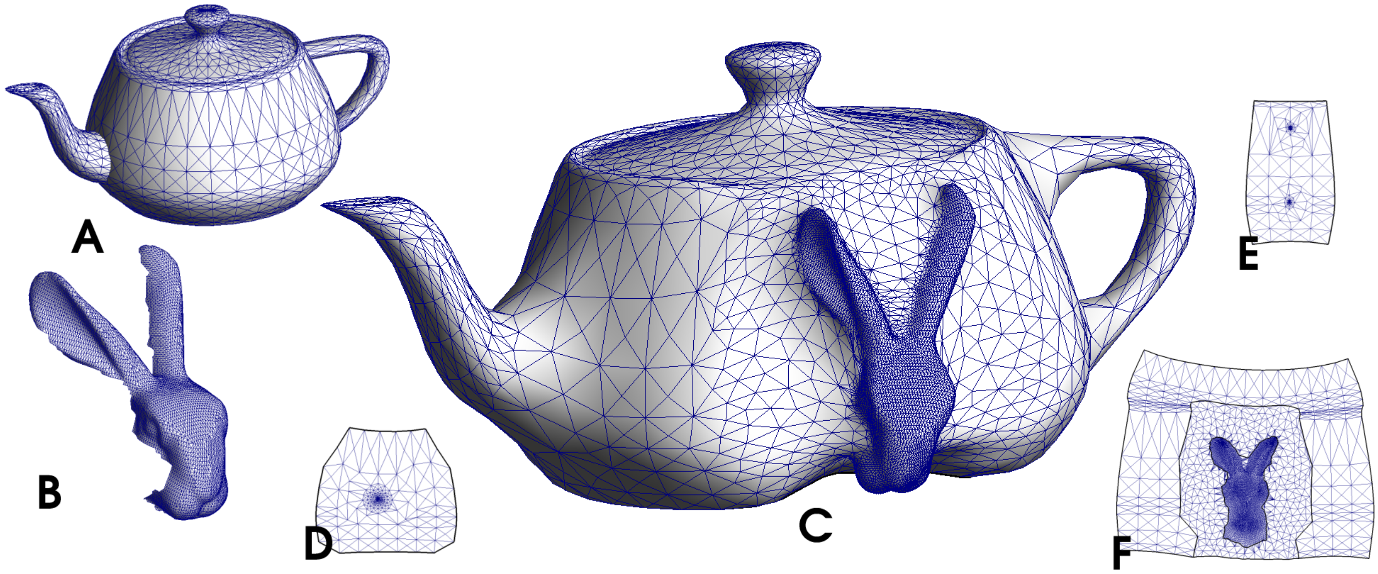

Figure 8: *A,B,C: Would you like a cup of tea, Mr. Rabbit ? (Alice's adventures in Wonderland); D,E,F: associated parameter spaces (Alice's adventures in Flatland). For this model, which is not homeomorphic to a disc, DDE was applied to parameterizations constructed in user-selected zones (total time of the session: 10 min.)*

	woman (1^{st} page)	chair (fig. 1)	statue (fig. 4)	jigsaw (fig. 6)
♯ initial vertices	21778	3102	18842	13078
LSCM time (s)	17	1.5	15.2	8.8
♯ free vertices	3357	1246	4109	1871
♯ Newton iters	4	3	2	3
Newton time (s)	9.2	1.5	16	1.8
♯ BFGS iters	5	4	4	3
BFGS time (s)	8.1	1.2	12	1.1
♯ SQP iters	8	5	5	5
SQP time (s)	17	2.8	31	5.1

Table 1: *Compared timings for Newton, BFGS and SQP*

One of our goals in future research is to find a numerical scheme that enables natural boundaries to be implemented for this latter criterion.

Aknowledgements

Thanks to CGAL [Fabri et al. 2000] for the Constrained Delaunay code, to Cyberware and Stanford University for the models.

References

ALLIEZ, P., MEYER, M., AND DESBRUN, M. 2002. Interactive Geometry Remeshing. *ACM Transactions on Graphics (SIGGRAPH conf. proc.)*, 347–354.

BLOOR, M., AND WILSON, M. 1990. Using partial differential equations to generate free-form surfaces. *Computer-Aided Design*, 22, 202–212.

BOISSONNAT, J.-D., AND CAZALS, F. 2002. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comp. Geometry Theory and Applications*, 185–203.

BRAKKE, K. 1992. The surface evolver. *Experimental Mathematics 1*, 2, 141–165.

BURCHARD, H., AYERS, J., FREY, W., AND SAPIDIS, N. 1994. Approximation with aesthetic constraints. In *Designing Fair Curves and Surfaces*, N. Sapidis, Ed., SIAM.

CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3D objects with radial basis functions. In *SIGGRAPH Conference Proceedings*, 67–76.

CHOPP, D., AND SETHIAN, J. 1999. Motion by intrinsic Laplacian of curvature. *Interfaces and Free Boundaries*, 1, 1–18.

DESBRUN, M., MEYER, M., SCHRÖDER, P., AND BARR, A. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH Conference Proceedings*, ACM, 317–324.

DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic parameterizations of surface meshes. In *Proceedings of Eurographics*, 209–218.

DO CARMO, M. 1976. *Differential geometry of curves and surfaces*. Prentice Hall, Englewood Cliffs, Inc.

FABRI, A., GIEZEMAN, G.-J., KETTNER, L., SCHIRRA, S., AND SCHÖNHERR, S. 2000. On the Design of CGAL, a Computational Geometry Algorithms Library. *Softw. – Pract. Exp. 30*, 11, 1167–1202. www.cgal.org.

FLOATER, M. 1997. Parametrization and smooth approximation of surface triangulations. *Computer Aided Geometric Design 14*, 3 (April), 231–250.

GUSKOV, I., SWELDENS, W., AND SCHRÖDER, P. 1999. Multiresolution signal processing for meshes. In *SIGGRAPH Conference Proceedings*, ACM, 325–334.

HORMANN, K., AND GREINER, G. 2000. MIPS: An efficient global parametrization method. In *Curve and Surface Design: Saint-Malo 1999*, P.-J. Laurent, P. Sablonnire, and L. Schumaker, Eds. Vanderbilt University Press, 153–162.

HSU, L., KUSNED, R., AND SULLIVAN, J. 1992. Minimizing the squared mean curvature integral for surfaces in space forms. *Experimental Mathematics 1*, 3, 191–207.

KOBBELT, L., CAMPAGNA, S., VORSATZ, J., AND SEIDEL, H. 1998. Interactive multi-resolution modeling on arbitrary meshes. In *SIGGRAPH Conference Proceedings*, 105–114.

KOBBELT, L. 1997. Discrete fairing. In *Proceedings of the Seventh IMA Conference on the Mathematics of Surfaces*, 101–131.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. *ACM Transactions on Graphics (SIGGRAPH conf. proc.)*, 362–371.

MALLET, J. 1992. Discrete Smooth Interpolation. *Computer Aided Design 24*, 4, 263–270.

MORETON, H., AND SÉQUIN, C. 1992. Functional minimization for fair surface design. In *SIGGRAPH Conference Proceedings*, ACM.

MUSETH, K., BREEN, D. E., WHITAKER, R. T., AND BARR, A. H. 2002. Level set surface editing operators. *ACM Transactions on Graphics (SIGGRAPH Conf. Proc.) 21*, 3, 330–338.

NOCEDAL, AND WRIGHT. 2000. *Numerical Optimization*. Springer.

OHTAKE, Y., BELYAEV, A., AND BOGAEVSKI, I. 2000. Polyhedral surface smoothing with simultaneous mesh regularization. In *Geometric Modeling and Processing Proceedings*, 229–237.

PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Math. 2*, 15.

SCHNEIDER, R., AND KOBBELT, L. Geometric fairing of irregular meshes for freeform surface design, to appear in the CAGD journal.

SCHNEIDER, R., AND KOBBELT, L. 2000. Generating fair meshes with G^1 boundary conditions. In *Geometric Modeling and Processing Proceedings*, 251–261.

SHEFFER, A., AND DE STURLER, E. 2001. Parameterization of faceted surfaces for meshing using angle based flattening. *Engineering with Computers 17*, 326–337.

TAUBIN, G. 1995. A signal processing approach to fair surface design. In *SIGGRAPH Conference Proceedings*, ACM, 351–358.

WELCH, W., AND WITKIN, A. 1994. Free-form shape design using triangulated surfaces. In *SIGGRAPH Conference Proceedings*, ACM.