INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Flattening of 3D data

Robert Acar — Brent W. Seales

N° 5048

Décembre 2003

THÈME 3

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--5048--FR+ENG

Flattening of 3D data

Robert Acar * , Brent W. Seales †

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet Ariana

Rapport de recherche n° 5048 — Décembre 2003 — 17 pages

Abstract: The digital library project strives to digitise special collections of libraries; this consists in storing as binary data, photographs of the content of ancient or rare manuscripts. The object is typically not in a flat plane. One collects, along with the photograph of the unflattened object (and the inevitably distorted text), a positional reading of its surface using laserometer. It is then a mathematical problem of how to use the latter information to undo the distortion of the photograph before storing the digitised image. We discuss a variational formulation and implementation of this.

Key-words: Digital preservation, Image restoration, Document analysis, Digital library, Isometry, Nonlinear least-squares

The first author thanks the University of Kentucky and INRIA for their hospitality, and was partially supported by NSF grant EIA9977071

* University of Puerto Rico
† University of Kentucky

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex (France)
Téléphone : +33 4 92 38 77 77 — Télécopie : +33 4 92 38 77 65

Applatissement de données 3D

Résumé : Le but du projet de la bibliothèque numérique est de numériser les collections spéciales des bibliothèques; ceci consiste à transformer en données binaires des photographies du contenu de manuscripts rares ou anciens. L'objet, typiquement, n'est pas dans un plan. On enregistre, en même temps que des photographies de l'objet non plat et du texte déformé qui s'y trouve, la forme et la position de sa surface en utilisant un laseromètre. La manière de se servir de cette information pour enlever la distortion de la photographie avant d'enregistrer l'image numérique est alors un problème mathématique. Nous en examinons une formulation variationnelle et l'implantation correspondante.

Mots-clés : Conservation numérique, Restauration de l'image, Analyse de documents, Bibliothèque numérique, Isométrie, Méthode des moindres carrés non-linéaire

1 Introduction

Many manuscripts in the possession of libraries and museums around the globe have already been damaged by historical disasters, and are subject to continuing further deterioration (see [2], [3]). The purpose of the *Digital Athenæum* project is to restore, analyse, and make accessible previously damaged manuscripts and artifacts. In [2] is described a 3-D acquisition technique using light projector and camera. The 3-D data, which includes (but is not restricted to) the space coordinates of a discrete sample of points on the surface to be registered, is used in conjunction with 2-D images (say, photographs). In many instances, the manuscript used to be flat, but underwent some bending and crimpling. In some cases, the bending in question is merely a result of the way the manuscript is stored, such as when sheets are collected in binders, and the repositor does not allow the patron or user to take them out of the binder. An important step preceding the storage in digital form is, then, to perform a virtual flattening of the surface to its original shape. It is this part of the restoration process which concerns us in the present work. In [2], this is achieved by using physics-based simulation: the surface is modelled by a mass-spring system and is "forced to collide with a plane" (or, in other words, dropped to the floor). Sampled points on the surface are given a mass and, to ensure some sort of shape rigidity, neighbouring particles are connected by Hookean springs. Such a sheet falling on a horizontal plane obstacle will tend then to oscillate to a limiting (mostly flat) position. The same technique has been applied to model the shape of cloth draped over solid support: see [1], [6]. We propose to follow a more geometric approach, where the shape of the surface is taken into account, but without any assumptions about local elasticity properties, which are implicit in the physics model.

2 Problem statement

Let us assume for a moment that, rather than having merely discrete data at hand, we have complete information on the position of the surface region, so that we know the coordinates of every point on that surface. We ask then: if this surface is the image under an isometry χ of a plane region, is χ unique (up to affine plane congruence?) If there were two such maps χ and ψ, then, since isometry is a property preserved under composition and taking inverse, the map $\psi^{-1}\chi$ must itself be an isometry. Call it ϕ.

To fix ideas, let us say that the plane region is a rectangle R. We are given an isometry ϕ of R onto itself; let us require that the action of ϕ on part of the boundary be specified — a minimal assumption would be that ϕ fixes 3 points of the boundary, but let us even assume that it leaves the whole boundary fixed. Does that force ϕ to be the identity? If the answer is yes, then we know that our problem has an unique solution. We address this question in appendix A.

3 Algorithm

The algorithm we propose relies on finding the minimum of a certain objective function (which we also refer to as the "criterion") reflecting the fact that the output, a certain triangulation of a plane region, results from an isometric flattening of the observed surface region. The vertices of this triangulation are in bijective correspondence with the vertices of a triangulation of the observed surface. The vertices of the observed region and of the flattened region constitute each a discretization of their respective objects, and, under the flattening, vertex i of the surface goes to vertex i of the flattened region.

An isometry has the property of preserving both local area and local distance. In the flattened region, they are the usual Euclidean area and distance; in the orginal surface, they are to be handled as area and distance on a Riemannian manifold: for instance, the distance between two (nearby) points is the length of the geodesic connecting them. Call ϕ the map performing the flattening. Our first approach was to compare between an area criterion (require that ϕ merely preserve local area), and a metric criterion (require that ϕ simply preserve local metric). Experiments suggested that the metric criterion was definitely the more reliable of the two (see appendix B). In retrospect, we understand why: preserving local area does not characterize an isometry, but preserving local metric does.

We proceed with the mathematical description of the problem. Let v_i, $1 \leq i \leq n$ be the set of observed points in space; each v_i consists of three coordinates x_i^0, y_i^0, z_i^0. We assume that the surface to which the data belongs can be represented in non-parametric form as $z = f(x, y)$; this assumption, up to coordinate change, will hold in general if the surface is not too severely warped. Our goal is to find a corresponding set of n points $u_i = (x_i, y_i)$ so that if $\phi(x_i, y_i) = (x_i^0, y_i^0, z_i^0)$, ϕ^{-1} will be the desired flattening. (See fig 1). Note that, without loss of generality, we choose the flattening plane to be the (x, y) plane itself. Call $u_1^0, \ldots, u_n^0$ the projections of $v_1, \ldots, v_n$ on the (x, y) plane. The algorithm will then be:

1. Perform, in the (x, y) plane, a Delaunay triangulation of $u_1^0, \ldots, u_n^0$. It consists of m triangles $T_1, \ldots T_m$. Call $e_1, \ldots, e_p$ the edges of this triangulation. By innocuous abuse of notation, if e_j joins the vertices u_a and u_b, call $\phi(e_j)$ the edge joining v_a and v_b, so that e_j is the projection of $\phi(e_j)$ on the (x, y) plane. In the same manner, denote by $\phi(T_k)$ that triangle having projection T_k. $\phi(T_1), \ldots \phi(T_m)$ represents a triangulation of the surface data. We will use these edges and triangles in the variational formulation, and we regard $u_1^0, \ldots, u_n^0$ as (almost) the zeroth iterate of our final solution $u_1, \ldots, u_n$.

2. Define the criterion (objective function), or energy:

$$\mathcal{E} = \mathcal{M} + \alpha\mathcal{A}$$

where $\mathcal{M} = \sum(|l(e_j) - l(\phi(e_j))|^2$ and $\mathcal{A} = \sum(|a(T_k) - a(\phi(T_k))|^2)$, and where l means Euclidean length and a means Euclidean area. (Note that the spatial triangulation is piecewise flat, so that each spatial triangle is a plane triangle). α is a control parameter: $\alpha = 0$ means that our criterion is purely metric-preserving, whereas $\alpha \gg 1$ means that our criterion is tilted towards preserving area.

3. Before we start the iteration in view of making the criterion minimum, we must make sure that the problem is well-posed by fixing enough many variables. There is more than one way to do this; so far, we have elicited to choose two adjacent nodes among the u^0s (say, numbered u_1^0 and u_2^0) and replace the distance between them, which is the projection of the distance $d(v_1, v_2)$, by $d(v_1, v_2)$. Typically, u_1^0, u_2^0 are chosen to fall on the boundary of the convex hull of $u_1^0, \ldots, u_n^0$. Holding these two fixed, the only genuine variables will then be $u_3, \ldots u_n$. Let r be the ratio between

$d(v_1, v_2)$ and $d(u_1^0, u_2^0)$. Up to symmetry, the new values of $u_3^0, \ldots u_n^0$ are obtained by a recoordinatisation consisting of a congruence followed by a dilation of ratio r. We may assume, e.g., that the new u_1^0 will be at the origin of the (x, y) plane, u_2^0 has now coordinates $(0, d(v_1, v_2))$, and all the other free vertices have nonnegative abscissæ $x_3^0, \ldots x_n^0$.

4. Having fixed the initial guess in this fashion, we now iterate until either a stopping criterion is met, or a certain number of iterations is exceeded, indicating lack of convergence. We have used the Matlab routine `lsqnonlin`, which we have found to be more robust than methods that use gradient information. Note that the objective function, while smooth, is nonconvex (see appendix B).

Two remarks:

1. When performing the initial Delaunay triangulation, it is necessary, as a preliminary step, to excise those triangles which will appear at the boundary, and which are not projections of genuine members of an approximate triangulation of the surface itself, but rather are "hanging down the sides". Since they tend to be elongated and skinny, one can discard them using a threshold on the largest angle. Here, we simply used a threshold on the area: we discarded those triangles the area of which was smaller than the average area by more than one standard deviation. Other criteria will have to be adopted ad hoc.

2. Since we are not using gradient information, it is not necessary to sum squares in the definition of the energy (or, in mathematical terms, we could use an L^1 instead of an L^2 criterion). The use of the squares is indicated only because we are using `lsqnonlin`, where the algorithm is naturally adapted to a sum of squares.

4 Results

We tested the algorithm on an example involving 25 nodes. The steps are illustrated in figures 2–5. Fig 2 shows the initial Delaunay triangulation of spatial data. Fig 3 shows the position of the initial guess. Figs 4 and 5 show the result of running `lsqnonlin` respectively when $\alpha = 1$ and $\alpha = 0$. Both figures show the placement of the solution, marked by the hollow dots, relative to the initial guess, marked by the fixed triangulation.

The point of this example is merely to check the soundness of the algorithm. On larger examples, matlab runs into trouble: we were unable to carry out the computation on a 15876-node (126×126) example, and we found that it was already considerably slowed down on a 1600-node subpatch of the data, due to matlab's excessive memory requirements. Further testing of the method would require using another high-level scientific language such as C or Fortran.

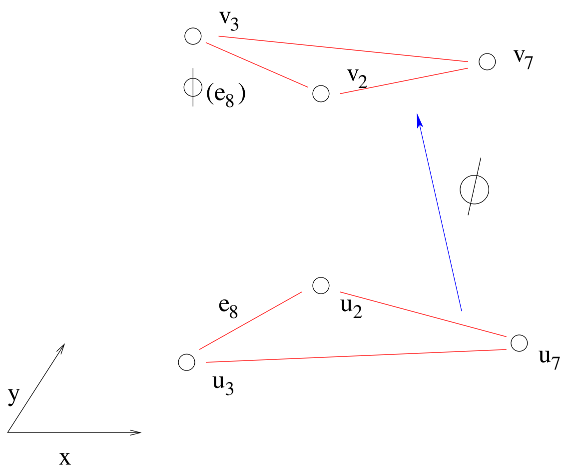

Figure 1: ϕ^{-1} is the flattening map

5 Future work

One omission from the method described, is how to take into account information about curvature. Indeed, a surface isometric to a flat region must have everywhere zero Gaussian curvature. Such information can be embedded in a refinement of our procedure as follows: given the spatial data, use local approximation to the curvature in order to segment the spatial region into one part which can be flattened, and one which cannot. This will arise, for instance, in any situation where an initially flat object has portions which developed,

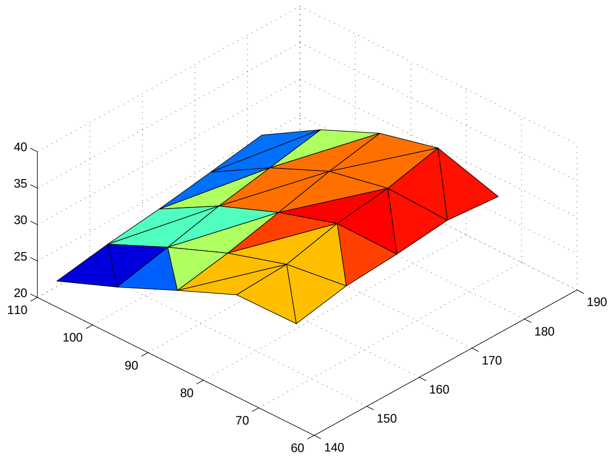

Figure 2: Delaunay triangulation of data

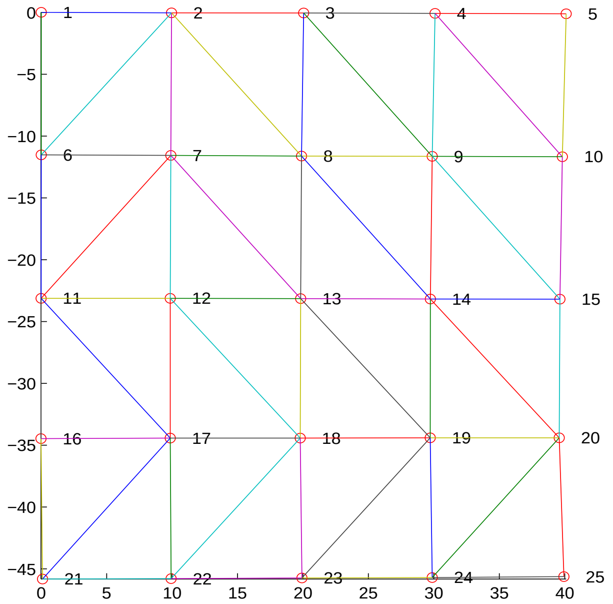

Figure 3: Initial guess

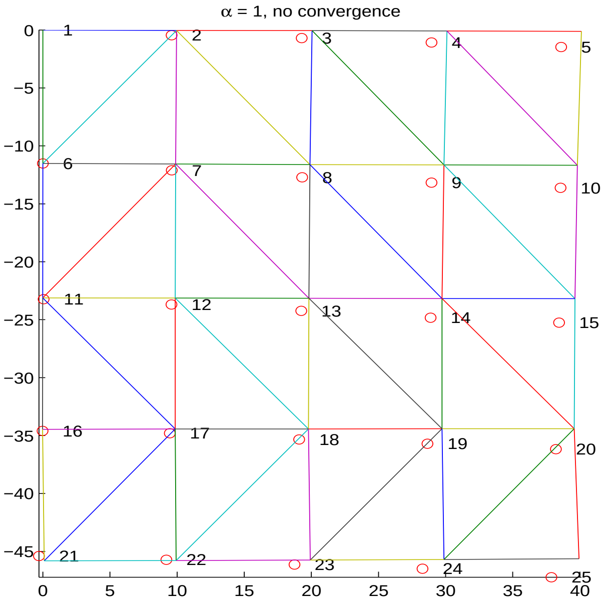

Figure 4: Result when $\alpha = 1$

over time, warping or bubbling. It is then of interest to modify the problem so that, now, the flattened data is replaced by nearly flattened, and one devises a dual way to store and view the curved part, using, say, a graphical user's interface.

Another purpose for using curvature information is to be able to generate data synthetically in view of testing this or other methods; the mathematics of this problem are now of differential-geometric nature. We have started work to this end.

A The continuous problem

In section 2, we showed that uniqueness of the solution of the continuous problem amounts to uniqueness of ϕ, isometric self-map of a plane region R, under the assumption that ϕ leaves part of the boundary fixed (in other words, such a ϕ can only be the identity). It is known that this is indeed the case if one uses the fact that isometry preserves local (Euclidean) distance. Is it also

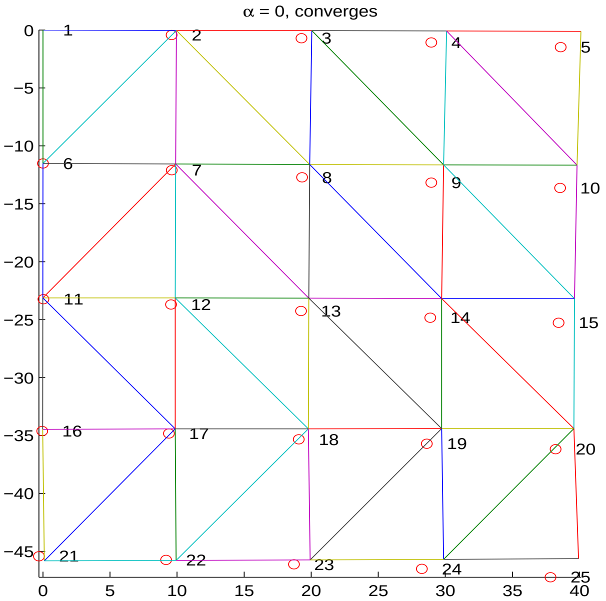

Figure 5: Result when $\alpha = 0$

the case if one requires merely that ϕ preserves local area? The answer turns out to be no, as the following example shows in case where R is the unit disc:

Let $g(r)$ be a function defined on $[0, 1]$, with the property that g sends 0 and 1 to 0, but g is not identically zero. An example of such a g is $g(r) = ar(1-r)$, a any real number. Define then ϕ, in polar coordinates, by: $\phi(r, \theta) = (r, \theta + g(r))$. ϕ can easily be seen to preserve the area measure $r\,\mathrm{d}r\,\mathrm{d}\theta$. ϕ also leaves the boundary of the unit disc fixed (and its center, as it happens), and ϕ is not the identity. One can think of ϕ as a shear with radial symmetry.

This lack of uniqueness can then be extended to any polygonal region P: by the Riemann mapping theorem (see, e.g., [5]), P is conformally equivalent to the unit disc. If ψ is a conformal map from R to P, then $\psi\phi\psi^{-1}$, the conjugate of ϕ, is a self-map of P which is not the identity, which leaves the boundary of P fixed, and which preserves area, since its Jacobian, product of the Jacobian of each factor, is one. Note that here we are not explicitly using the fact that ψ is conformal, only that it is a diffeomorphism. Figures 6 and 7 help visualize this self-map. In these figures, the conformal map from the unit disc to the square was computed using Toby Driscoll's Schwarz-Christoffel Matlab toolbox (see [4]).

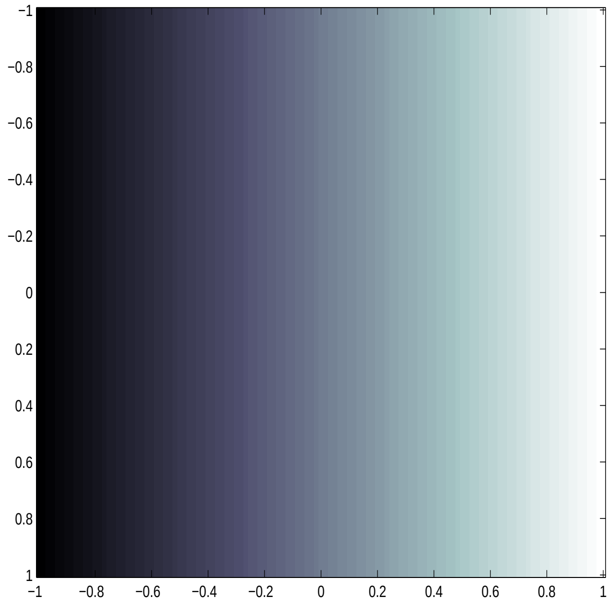

Figure 6: Unit square

B Discretization

The solution of the flattening problem is cast as finding the minimum of a certain criterion E, i.e., as a variational problem. This discretization yields a function of finitely many variables (namely, the positions of the "free" nodes); i.e., the state space is finite-dimensional. Of this variational problem, we need to answer the following questions:

- Does the solution exist?
- Is the solution unique?
- How sensitive is the solution to perturbations in the data?
- Since the variational problem is inherently a discretization, we don't expect it to be the same as the solution of the flattening problem. But does it approximate it? In which sense?

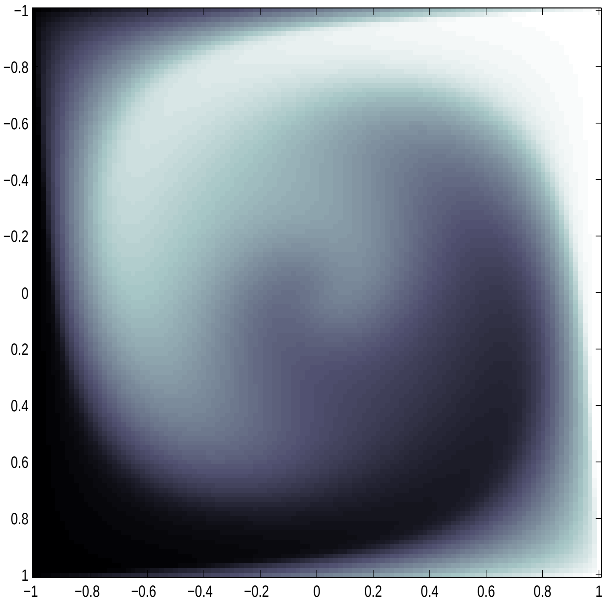

Figure 7: Area-preserving perturbation

We can answer the first question by the affirmative: the growth property of E in the large guarantees that there exist bounded (hence compact) lower level sets, so that a global minimum exists. The answer to the second question would be yes, if E were strictly convex; but E is not even convex (see below). However, our experiments seem to indicate that E is strictly convex in the neighbourhood of the minima, which would imply that minima are isolated, hence locally unique. We do not have yet a proof of the uniqueness of the global minimum. Note that since the data is, in a sense, approximate, since the distances and areas arise from a piecewise flat approximation of the surface, we do not expect the value at the computed minimum to be exactly zero.

We see by now that, lacking convexity, a procedure to solve numerically the variational problem will need to have reasonably robust global convergence properties. The answers to the third and fourth question are both related to the degree of convexity of E in the neighbourhood of the global minimum (which is also a local minimum) and await further analysis.

To visualize the convexity properties of E, we consider an example where the data is already flat (here we have 6 variable nodes, of which 2 are interior), and where we allow only one node to vary, the others being held fixed, equal to the corresponding vs, which are shown in figure 8. Figures 9 and 10 show the plots of $\mathcal{E}$ corresponding to values of α respectively 0 and 10. This indicates that, while neither $\mathcal{M}$ nor $\mathcal{A}$ is everywhere convex (even though they are convex in the neighbourhood of the minimum), the nonconvexity of $\mathcal{M}$ is mild compared to that of $\mathcal{A}$.

Referring to the same v data, we show the result of the numerical minimisation, starting from the initial state shown in fig 11. We find that, while the pure metric criterion yields convergence to the desired solution, the area criterion ($\alpha = 10$) converges to a spurious solution (corresponding to a local minimum of $\mathcal{E}$). The results are shown in Figs 12-13.

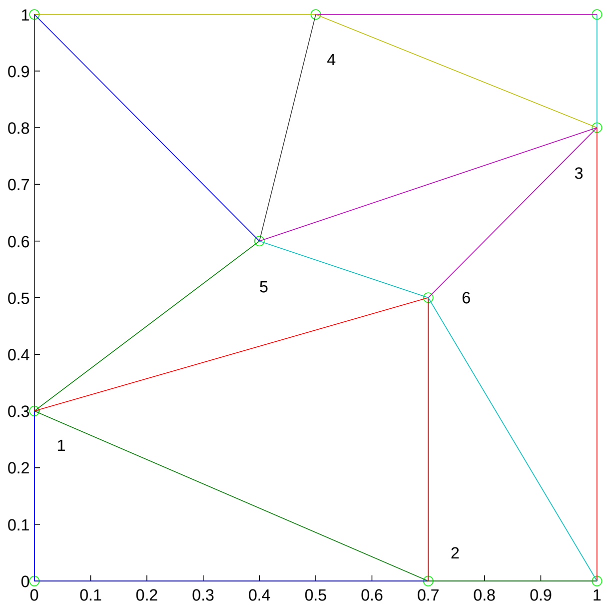

Figure 8: Target configuration

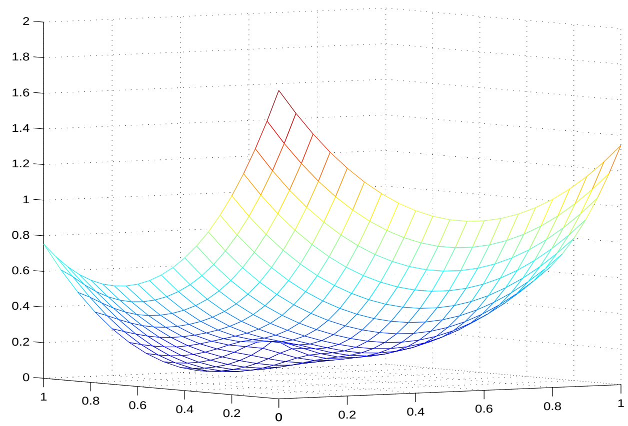

Figure 9: Varying node 4, $\alpha = 0$

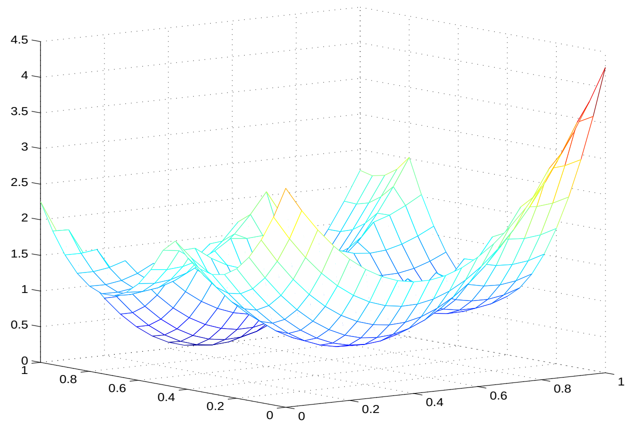

Figure 10: Varying node 4, $\alpha = 10$

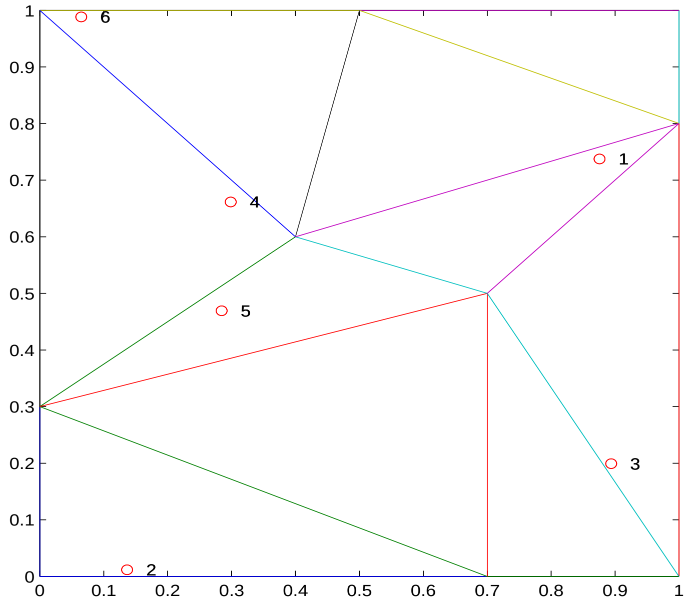

Figure 11: Initial guess

References

[1] D. Baraff and A. Witkin *Large steps in cloth simulation*, Computer Graphics (Proc. SIGGRAPH), pages 43–52, August 1998

[2] M.S. Brown and W.B. Seales *The Digital Atheneum: new approaches for preserving, restoring and analyzing damaged manuscripts*, Proceedings of the first ACM/IEEE-CS joint conference on Digital libraries, 2001

[3] M.S. Brown, W.B. Seales, K. Kiernan and J. Griffioen *3D Acquisition and Restoration of Medieval Manuscripts*, Communications of the ACM Special Issue on Digital Libraries, May 2001

[4] Toby Driscoll *Schwarz-Christoffel toolbox for Matlab* http://www.math.udel.edu/∼driscoll/SC/

[5] S.G. Krantz *The Riemann Mapping Theorem* in Handbook of Complex Analysis. Boston, MA: Birkhäuser, pp. 86-87, 1999

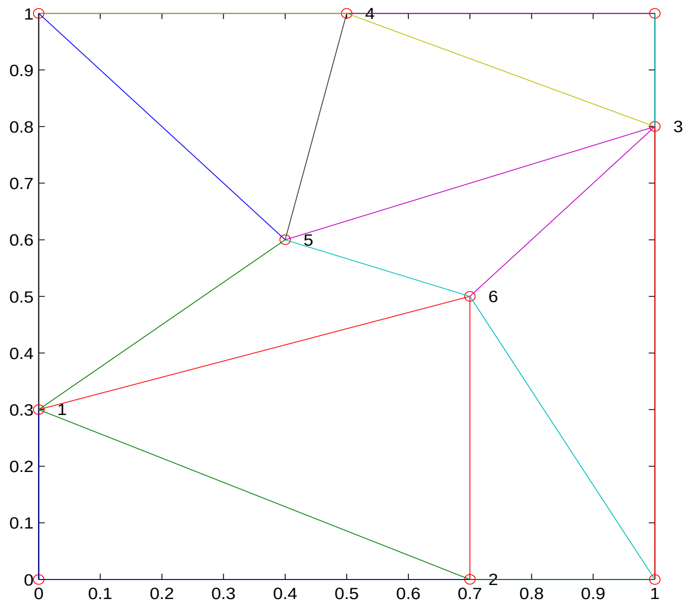

Figure 12: Result when $\alpha = 0$

[6] D. Terzopoulos, J.C. Platt, and A.H. Barr *Elastically deformable models*, Computer Graphics (Proc. SIGGRAPH), pages 21:205–214, 1987.

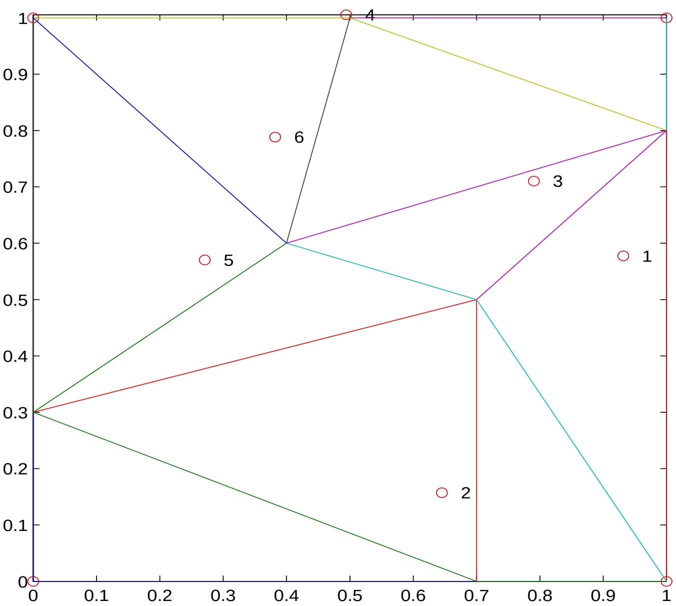

Figure 13: Result when $\alpha = 10$

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)
Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399