On Template-Based Reconstruction from a Single View: Analytical Solutions and Proofs of Well-Posedness for Developable, Isometric and Conformal Surfaces

A. Bartoli, Y. Gérard, F. Chadebecq and T. Collins
ISIT – UMR6284 CNRS / Université d'Auvergne, Clermont-Ferrand, France

Abstract

Recovering a deformable surface's 3D shape from a single view registered to a 3D template requires one to provide additional constraints. A recent approach has been to constrain the surface to deform quasi-isometrically. This is applicable to surfaces of materials such as paper and cloth. Current 'closed-form' solutions solve a convex approximation of the original problem whereby the surface's depth is maximized under the isometry constraints (this is known as the maximum depth heuristic). No such convex approximation has yet been proposed for the conformal case.

We give a unified problem formulation as a system of PDEs for developable, isometric and conformal surfaces that we solve analytically. *This has important consequences. First, it gives the first* analytical algorithms *to solve this type of reconstruction problems. Second, it gives the first algorithms to solve for the* exact constraints. *Third, it allows us to study the* well-posedness *of this type of reconstruction: we establish that isometric surfaces can be reconstructed unambiguously and that conformal surfaces can be reconstructed up to a few discrete ambiguities and a global scale. In the latter case, the candidate solution surfaces are obtained analytically.*

Experimental results on simulated and real data show that our methods generally perform as well as or outperform state of the art approaches in terms of reconstruction accuracy.

1. Introduction

Significant progress has been made over the last decade on the 3D reconstruction of deformable surfaces from monocular video data. In the particular template-based setup, a reference 3D view of the surface to be reconstructed, called *template*, is available. 3D reconstruction is then carried out from 3D to 2D correspondences which are established between the template and an input image of the deformed surface. State-of-the-art has reported extremely impressive results in template-based reconstruction, both in terms of image matching [3, 5, 11, 12, 20] and shape inference [1, 2, 4, 5, 6, 7, 8, 10, 14, 15, 16, 17, 18, 20]. Reprojection constraints are clearly not sufficient to well-constrain the problem since a large number of surfaces project at the same image location. To further constrain the problem, most works assume the surface to deform isometrically and to be developable (isometric to the plane) [1, 2, 4, 5, 7, 10, 14, 15, 18, 20]. Isometry has been combined with shading [7], the problem's dimensionality reduced by learning the admissible shape space [7] (originally proposed in [19]) or stitching multiple local surface patches of low complexity [15]. Some other works reconstruct elastic surfaces using shading [8], quasi-conformal deformations [6] or temporal smoothness [16].

This paper is concerned with reprojection constraints and surface deformation constraints. The latter are non-convex. The literature has established a number of well behaved *numerical* methods solving convex approximations to the isometric problem [1, 2, 7, 10, 14, 15, 18, 20] that cannot be used for elastic surfaces. We propose a different approach. First, we formulate the problem as a system of PDEs which unifies the cases of developable, isometric and conformal surfaces. Conformal surfaces preserve their angles during deformation. They do not physically exist but the conformal model, allowing for local isotropic scaling, forms a simple relaxation of the isometric model which has given good empirical reconstruction results on medical data [6]. Second, we solve the system of PDEs *analytically*. This leads to simple algorithms that solve the original constraints fast, whether developable, isometric or conformal, without optimization. Third, we establish the following two theorems regarding well-posedness:

Theorem 1 *Template-based isometric surface reconstruction from a single view registered to the template generally has a single solution. This holds true for developable and non-developable surfaces.*

Theorem 2 *Template-based conformal surface reconstruction from a single view registered to the template generally has a small, discrete set of solutions (at least two) and a global scale ambiguity.*

Despite the amount of recent literature concerned with solving this type of 3D reconstruction problem algorithmically, very little effort was devoted to analyzing the problem's well-posedness from a theoretical standpoint.

We demonstrate our algorithms and compare them to existing ones using simulated data and three real datasets. These experiments show that our algorithms compete with and often outperform state of the art approaches. They also show that 3D reconstruction for extensible surfaces *is possible* with the conformal deformation model.

Paper organization. §2 reviews the state of the art. §3 sets our general differential problem modeling and §4 introduces the necessary change of variables. §5 derives our analytical solution to the conformal problem. §6 specializes the conformal solution to the isometric and developable cases. §7 gives our algorithmic implementation. §8 reports experimental results. Finally §9 gives our conclusions. Note that theorems 1 and 2 are respectively proved in §§6 and 5.

2. State of the Art

Template-based 3D reconstruction of a deformable surface from a single view has two main steps: *(i)* registration of the input image to the template and *(ii)* shape inference from reprojection and deformation constraints. Image registration has been solved by the pixel-based method [3], the feature-based method [11] and a combination of both [12]. The shape inference step has been solved in many methods by constraining the surface to deform isometrically [1, 2, 4, 7, 10, 14, 15, 18, 20]. Almost all these methods use a convex approximation[1] of the original problem. Empirical results [14, 1] showed that the most accurate results were obtained using a convex approximation that maximizes the surface's depth (the *maximum depth heuristic* [10]) under inextensibility constraints [1, 14, 20, 7]. An analytical solution under orthographic projection has also been proposed [4]. Extensible (non-isometric) surfaces have recently been modeled by quasi-conformal surfaces [6] (surfaces that preserve their angles during deformation) and an iterative reconstruction algorithm was proposed. Conformal transformations are extensively used in Computer Graphics, for instance to find a 2D parameterization of some given 3D surface for texture-mapping [21]. Extensible surfaces were also reconstructed by using the shading cue [8]. Finally, it is worth of note that numerical methods were proposed that solve for both the registration and shape inference steps simultaneously [20, 5].

[1]The term 'closed-form' as used in the current literature is somewhat ambiguous. We make the following clear distinction. *Convex (numerical) methods* solve a convex approximation to the original problem. These include depth maximizing methods. *Analytical methods* solve the original nonconvex problem with no numerical optimization. These include the proposed methods.

To sum up, state of the art shows that significant efforts were put into finding effective numerical methods. In particular, convex numerical methods were proposed to solve the shape inference step by approximating the original deformation constraint using the maximum depth heuristic. This unfortunately cannot be used with extensible surfaces. However, the literature does not bring any analytical solution to the original problem, nor does it give any *theoretical analysis* of the problem's well-posedness (though numerical studies were proposed based on spectral analysis [17, 7, 18]). In this paper we use PDEs as a modeling framework, from which we derive analytical solutions to developable, isometric and conformal 3D surface reconstruction. This allows us to propose simple reconstruction algorithms and carefully analyze the problem's well-posedness from a theoretical standpoint.

3. Modeling

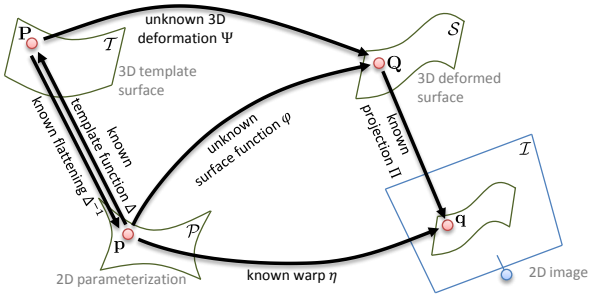

Figure 1. **General modeling of monocular template-based reconstruction.** The surface is assumed to have a disk topology.

Geometric modeling. Our problem modeling is illustrated in figure 1. We consider two 3D surfaces, $\mathcal{T} \subset \mathbb{R}^3$ the known template surface and $\mathcal{S} \subset \mathbb{R}^3$ an unknown deformed surface. These two surfaces are related by some unknown continuously differentiable 3D deformation $\Psi \in \mathbb{C}_1$. Ψ maps a point $\mathbf{P} \in \mathcal{T}$ to the corresponding point $\mathbf{Q} = \Psi(\mathbf{P}) \in \mathcal{S}$. Let $\mathcal{P} \subset \mathbb{R}^2$ be some 2D parameterization of $\mathcal{T}$ and $\mathcal{S}$ obtained by conformally flattening $\mathcal{T}$ [21]. The template surface function $\Delta \in \mathbb{C}_1$ is an invertible function mapping a point $\mathbf{p} \in \mathcal{P}$ to its 3D point $\mathbf{P} = \Delta(\mathbf{p}) \in \mathcal{T}$. The unknown surface function $\varphi \in \mathbb{C}_1$ maps a point $\mathbf{p} \in \mathcal{P}$ to $\mathbf{Q} = \varphi(\mathbf{p}) \in \mathcal{S}$. We write the input image as $\mathcal{I} \subset \mathbb{R}^2$. The known camera projection operator Π maps $\mathbf{Q} \in \mathcal{S}$ to $\mathbf{q} = \Pi(\mathbf{Q}) \in \mathcal{I}$. Let $\mathbf{q}^\top = (x \;\; y)$ and $\mathbf{Q}^\top = (X \;\; Y \;\; Z)$. Assuming that the pin-hole camera model's intrinsics contained in matrix K are known and the effects 'undone' in the image (by applying K^{-1}) the projection operator is simply $\Pi(\mathbf{Q}) = \frac{1}{Z}(X \;\; Y)^\top$. We write $\eta \in \mathbb{C}_1$ as the known warp function that maps $\mathbf{p} \in \mathcal{P}$ to $\mathbf{q} = \eta(\mathbf{p}) \in \mathcal{I}$.

Differential modeling. The *reprojection constraint* is thus simply obtained as $\Pi \circ \varphi = \eta$. The *deformation constraints* are differential constraints, expressed using the surface function's Jacobian matrix (we use the notation $\mathsf{J}_\varphi = \frac{\partial \varphi}{\partial \mathbf{p}}$ for the Jacobian matrix of some function φ). The deformation constraints are found starting from the equation $\varphi = \Psi \circ \Delta$ and considering its partial derivatives $\mathsf{J}_\varphi = (\mathsf{J}_\Psi \circ \Delta)\mathsf{J}_\Delta$ obtained from the chain rule. Let $\mathbb{O}$ be the set of (3×3) orthonormal matrices, $\tilde{\mathbb{O}}$ the set of (3×2) column-orthonormal matrices, $\mathbb{S}$ the set of (3×3) orthogonal matrices (orthonormal up to scale, $\mathbb{O} \subset \mathbb{S}$) and $\tilde{\mathbb{S}}$ the set of (3×2) column-orthogonal matrices (column-orthonormal up to scale, $\tilde{\mathbb{O}} \subset \mathbb{S}$). It is then simple to derive the 3D surface's differential properties [13] which express the deformation constraints on J_φ:

Surface	J_Ψ	J_Δ	J_φ	Comments
Developable	$\mathcal{T} \to \mathbb{O}$	$\mathcal{P} \to \tilde{\mathbb{O}}$	$\mathcal{P} \to \tilde{\mathbb{O}}$	
Isometric	$\mathcal{T} \to \mathbb{O}$	$\mathcal{P} \to \tilde{\mathbb{S}}$	$\mathcal{P} \to \tilde{\mathbb{S}}$	Known scale
Conformal	$\mathcal{T} \to \mathbb{S}$	$\mathcal{P} \to \tilde{\mathbb{S}}$	$\mathcal{P} \to \tilde{\mathbb{S}}$	Unknown scale

Our problem is thus to find $\varphi : \mathcal{P} \to \mathbb{R}^3$ such that:

$$\left\{ \begin{array}{l} \Pi \circ \varphi = \eta \\ \mathsf{J}_\varphi^\top \mathsf{J}_\varphi = \lambda \mathrm{I} \end{array} \right. \quad \lambda^2 = \left\{ \begin{array}{ll} 1 & \text{(developable)} \\ \det(\mathsf{J}_\Delta^\top \mathsf{J}_\Delta) & \text{(isometric)} \\ \text{unknown } \mathcal{P} \to \mathbb{R} & \text{(conformal),} \end{array} \right. \tag{1}$$

where I is the (2×2) identity matrix. The local scale function $\lambda : \mathcal{P} \to \mathbb{R}$, $\lambda > 0$ in the isometric case is obtained from $\mathsf{J}_\varphi^\top \mathsf{J}_\varphi = \lambda \mathrm{I}$. We have this since $\lambda^2 = \det(\mathsf{J}_\varphi^\top \mathsf{J}_\varphi) = \det(\mathsf{J}_\Delta^\top \mathsf{J}_\Delta)$ since $\mathsf{J}_\varphi = (\mathsf{J}_\Psi \circ \Delta)\mathsf{J}_\Delta$ and $\mathbb{O}$ is the image of J_Ψ. Note that previous work concentrated on the developable case where $\lambda = 1$ is a constant function.

Let $\eta = (\eta_x \ \ \eta_y)^\top$ and $\varphi = (\varphi_x \ \ \varphi_y \ \ \varphi_z)^\top$. We now use the two equations from the reprojection constraint to eliminate φ_x and φ_y and define the alias $\gamma = \varphi_z > 0$ (since the surface lies in front of the camera). The reprojection constraint is rewritten $\varphi_x = \eta_x \gamma$ and $\varphi_y = \eta_y \gamma$, and thus:

$$\varphi = \gamma \begin{pmatrix} \eta \\ 1 \end{pmatrix} \ \Rightarrow \ \mathsf{J}_\varphi = \begin{pmatrix} \frac{\partial \eta_x}{\partial x}\gamma + \eta_x \frac{\partial \gamma}{\partial x} & \frac{\partial \eta_x}{\partial y}\gamma + \eta_x \frac{\partial \gamma}{\partial y} \\ \frac{\partial \eta_y}{\partial x}\gamma + \eta_y \frac{\partial \gamma}{\partial x} & \frac{\partial \eta_y}{\partial y}\gamma + \eta_y \frac{\partial \gamma}{\partial y} \\ \frac{\partial \gamma}{\partial x} & \frac{\partial \gamma}{\partial y} \end{pmatrix}. \tag{2}$$

This simply enforces that every point on the surface $\mathcal{S}$ lies on the corresponding line of sight predicted by the warp η.

4. Preliminary Rewriting

We now put the system of PDEs (1) in a form that will be solved. We start from the deformation constraints:

$$\left(\mathsf{J}_\varphi^\top \mathsf{J}_\varphi\right)_{1,1} = \left(\mathsf{J}_\varphi^\top \mathsf{J}_\varphi\right)_{2,2} = \lambda \qquad \text{and} \qquad \left(\mathsf{J}_\varphi^\top \mathsf{J}_\varphi\right)_{1,2} = 0.$$

We define the following four functions A, B_x, B_y and C as:

$$A = 1 + \|\eta\|_2^2 \quad B_x = \left\| \frac{\partial \eta}{\partial x} \right\|_2^2 \quad B_y = \left\| \frac{\partial \eta}{\partial y} \right\|_2^2 \quad C = \left(\frac{\partial \eta}{\partial x} \right)^\top \frac{\partial \eta}{\partial y}.$$

Using equation (2) we rewrite the system of PDEs (1) as:

$$\begin{array}{l} A\left(\frac{\partial \gamma}{\partial x}\right)^2 + \frac{\partial A}{\partial x}\gamma\frac{\partial \gamma}{\partial x} + B_x\gamma^2 = A\left(\frac{\partial \gamma}{\partial y}\right)^2 + \frac{\partial A}{\partial y}\gamma\frac{\partial \gamma}{\partial y} + B_y\gamma^2 = \lambda \\ A\frac{\partial \gamma}{\partial x}\frac{\partial \gamma}{\partial y} + \frac{1}{2}\frac{\partial A}{\partial y}\gamma\frac{\partial \gamma}{\partial x} + \frac{1}{2}\frac{\partial A}{\partial x}\gamma\frac{\partial \gamma}{\partial y} + C\gamma^2 = 0. \end{array}$$

We define $f = \gamma\sqrt{A}$ (by definition $A \geq 1$ and thus $f > 0$) and replace the unknown function γ by $\gamma = fA^{-\frac{1}{2}}$ giving:

$$\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{B_x}{A} - \left(\frac{\partial A}{\partial x}\right)^2 \frac{1}{4A^2}\right) f^2 = \lambda \tag{3}$$

$$\left(\frac{\partial f}{\partial y}\right)^2 + \left(\frac{B_y}{A} - \left(\frac{\partial A}{\partial y}\right)^2 \frac{1}{4A^2}\right) f^2 = \lambda \tag{4}$$

$$\frac{\partial f}{\partial x}\frac{\partial f}{\partial y} = \left(\frac{1}{4A^2}\frac{\partial A}{\partial x}\frac{\partial A}{\partial y} - \frac{C}{A}\right) f^2. \tag{5}$$

Finally, we define two functions U and V as:

$$U = \frac{B_y - B_x}{A} + \left(\left(\frac{\partial A}{\partial x}\right)^2 - \left(\frac{\partial A}{\partial y}\right)^2\right)\frac{1}{4A^2} \tag{6}$$

$$V = \frac{1}{4A^2}\frac{\partial A}{\partial x}\frac{\partial A}{\partial y} - \frac{C}{A}. \tag{7}$$

It can be verified that in the conformal case the original problem (1) is equivalent to finding $f : \mathcal{P} \to \mathbb{R}$, $f > 0$ such that:

$$\left(\frac{\partial f}{\partial x}\right)^2 - \left(\frac{\partial f}{\partial y}\right)^2 = Uf^2 \tag{8}$$

$$\frac{\partial f}{\partial x}\frac{\partial f}{\partial y} = Vf^2. \tag{9}$$

Note that functions A, B_x, B_y and C, as well as U and V, depend only on the known warp η. Every solution f to (8,9) yields a solution φ to (1) as:

$$\varphi = fA^{-\frac{1}{2}}\begin{pmatrix} \eta \\ 1 \end{pmatrix}. \tag{10}$$

5. Conformal Deformation

We now present two propositions. The first gives analytical solutions to the system of PDEs (1) in the conformal deformation case. The second gives an analytical solution to the local isotropic scaling function λ corresponding to a conformal solution.

Proposition 1 (conformal solutions) *The solutions of the system of PDEs (8,9) are $f = \exp g$ with:*

$$\frac{\partial g}{\partial x} = \frac{1}{\sqrt{2}}\delta\sqrt{\sqrt{U^2 + 4V^2} + U} \tag{11}$$

$$\frac{\partial g}{\partial y} = \frac{1}{\sqrt{2}}\delta\,\mathrm{sign}(V)\sqrt{\sqrt{U^2 + 4V^2} - U} \tag{12}$$

where $\delta : \mathcal{P} \to \pm 1$ is a piecewise constant sign function which may vary between the connected components separated by closed curves of equations $U = V = 0$, and represents a convex/concave ambiguity.

Proposition 2 (conformal scaling) *The conformal scaling function λ associated to a solution f is:*

$$\lambda = \frac{1}{2}\left(\frac{B_x + B_y}{A} + \sqrt{U^2 + 4V^2} - \frac{1}{4A^2}\left\|\frac{\partial A}{\partial \mathbf{q}}\right\|_2^2\right) f^2 \qquad (13)$$

Theorem 2 follows from proposition 1, and from the fact that g is recovered up to an additive constant c by integrating its partial derivatives. Thus, $f = \exp(c+g) = (\exp c)(\exp g)$. From equation (10) we then have $\varphi^\top = (\exp c)(\exp g)A^{-\frac{1}{2}}(\eta\ \ 1)^\top$ showing that φ is recovered up to a global scale $\exp c$.

5.1. Proof of Proposition 1

We first prove the forward implication: we show that any solution is of the form given in proposition 1. We then prove the reverse implication: we show that every function of the form given in proposition 1 is a valid solution.

Forward implication, $V \neq 0$**.** This is the general case. We multiply equation (8) by $\left(\frac{\partial f}{\partial x}\right)^2$ and $\left(\frac{\partial f}{\partial y}\right)^2$. This leads to:

$$\left(\frac{\partial f}{\partial x}\right)^4 - \left(\frac{\partial f}{\partial x}\right)^2\left(\frac{\partial f}{\partial y}\right)^2 = Uf^2\left(\frac{\partial f}{\partial x}\right)^2$$

$$\left(\frac{\partial f}{\partial x}\right)^2\left(\frac{\partial f}{\partial y}\right)^2 - \left(\frac{\partial f}{\partial y}\right)^4 = Uf^2\left(\frac{\partial f}{\partial y}\right)^2.$$

By substituting equation (9) we obtain:

$$\left(\frac{\partial f}{\partial x}\right)^4 - Uf^2\left(\frac{\partial f}{\partial x}\right)^2 - V^2f^4 = 0$$

$$\left(\frac{\partial f}{\partial y}\right)^4 + Uf^2\left(\frac{\partial f}{\partial y}\right)^2 - V^2f^4 = 0.$$

The two equations have the same (positive) discriminant $d = (U^2 + 4V^2)f^4$, giving:

$$\left(\frac{\partial f}{\partial x}\right)^2 = \frac{1}{2}\left(\epsilon_x\sqrt{U^2+4V^2} + U\right)f^2 \qquad (14)$$

$$\left(\frac{\partial f}{\partial y}\right)^2 = \frac{1}{2}\left(\epsilon_y\sqrt{U^2+4V^2} - U\right)f^2, \qquad (15)$$

where $\epsilon_x = \pm 1$ and $\epsilon_y = \pm 1$. We notice from the first equation that $\epsilon_x\sqrt{U^2+4V^2} + U$ must be positive. Since $U^2 + 4V^2 \geq U^2$ and $V \neq 0$ we obtain $\epsilon_x = 1$. A similar reasoning for the second equation gives $\epsilon_y = 1$, and thus:

$$\left(\frac{\partial f}{\partial x}\right)^2 = \frac{1}{2}\left(\sqrt{U^2+4V^2} + U\right)f^2 \qquad (16)$$

$$\left(\frac{\partial f}{\partial y}\right)^2 = \frac{1}{2}\left(\sqrt{U^2+4V^2} - U\right)f^2. \qquad (17)$$

Introducing two sign functions $\delta_x : \mathbb{R}^2 \to \pm 1$ and $\delta_y : \mathbb{R}^2 \to \pm 1$, we obtain:

$$\frac{\partial f}{\partial x} = \frac{1}{\sqrt{2}}\delta_x\sqrt{\sqrt{U^2+4V^2} + U}\,f$$

$$\frac{\partial f}{\partial y} = \frac{1}{\sqrt{2}}\delta_y\sqrt{\sqrt{U^2+4V^2} - U}\,f.$$

By substituting the above expressions into equation (9) we obtain $\delta_x\delta_y = \operatorname{sign}(V)$. Defining $\delta_x = \delta$ we get $\delta_y = \delta\operatorname{sign}(V)$. Finally, by setting $g = \ln f$, with $\frac{\partial g}{\partial x} = \frac{1}{f}\frac{\partial f}{\partial x}$ and $\frac{\partial g}{\partial y} = \frac{1}{f}\frac{\partial f}{\partial y}$ and substituting in the two equations directly above, we get equations (11,12). Consequently, any solution will be of the form given in proposition 1 if $V \neq 0$.

Forward implication, $V = 0$**.** In this special case equation (9) implies that at least one of the partial derivatives of f vanishes. Fortunately, equations (11,12) still hold, as we demonstrate below. Assuming $U \neq 0$ for now, equation (8) implies that one and only one partial derivative of f vanishes. If $U > 0$, equation (8) implies $\left(\frac{\partial f}{\partial x}\right)^2 = Uf^2$ and $\frac{\partial f}{\partial y} = 0$, which, substituting f by g leads to:

$$\frac{\partial g}{\partial x} = \zeta U \qquad \frac{\partial g}{\partial y} = 0 \qquad \zeta = \pm 1.$$

This is exactly what one obtains by setting $V = 0$ and $U > 0$ in equations (11,12), where we arbitrarily choose $\operatorname{sign}(0) = 1$, and with ζ given by the δ sign function. If $U < 0$, the same reasoning leads to:

$$\frac{\partial g}{\partial x} = 0 \qquad \frac{\partial g}{\partial y} = \zeta U \qquad \zeta = \pm 1,$$

which again is obtained by setting $V = 0$ and $U < 0$ in equations (11,12). The last case is $U = 0$. It implies from equations (8,9) that both partial derivatives of f vanish. Replacing f by g, and since $f > 0$, this also implies that both partial derivatives of g vanish, which is the result given by setting $U = V = 0$ in equations (11,12). This completes the proof of the forward implication.

Reverse implication. The solution f must satisfy equations (8,9). Using the fact that $f = \exp g$ and so $\frac{\partial f}{\partial x} = \frac{\partial g}{\partial x}f$ and $\frac{\partial f}{\partial y} = \frac{\partial g}{\partial y}f$, by substituting in equations (8,9) we arrive at:

$$\left(\frac{\partial g}{\partial x}\right)^2 - \left(\frac{\partial g}{\partial y}\right)^2 = U \qquad \frac{\partial g}{\partial x}\frac{\partial g}{\partial y} = V.$$

It is straightforward to verify that these two equations are satisfied by the partial derivatives of g given by equations (11,12). This completes the proof of proposition 1.

The sign function δ. Proposition 1 shows how we can obtain g and then f. The former depends on the unknown sign function δ. If δ were not constrained the number of possible solutions would be infinite. It turns out that the number of solutions can be reduced to generally just a few by considering $\varphi \in \mathbb{C}_1$. We first note that there are at least two solutions corresponding to δ and $-\delta$, a convex one and a concave one. Because $\eta \in \mathbb{C}_1$ and $\varphi \in \mathbb{C}_1$, it is straightforward to show that $g \in \mathbb{C}_1$. This is because $A = 1 + \|\eta\|_2^2 \in \mathbb{C}_1$ and thus $f = \gamma\sqrt{A} \in \mathbb{C}_1$. Since $A \geq 1$ and $f > 0$ we have $g = \ln f \in \mathbb{C}_1$. Therefore $\frac{\partial g}{\partial x} \in \mathbb{C}_0$ and $\frac{\partial g}{\partial y} \in \mathbb{C}_0$ and it is clear from equations (11,12) that δ may change only if $\frac{\partial g}{\partial x} = \frac{\partial g}{\partial y} = 0$, which implies $U = V = 0$. Consequently, δ may vary in the special case where two closed curves $U = 0$ and $V = 0$ overlap. This thus defines $\kappa \geq 1$ connected components in $\mathcal{P}$. The sign function δ is of constant value over each connected component and leads to two solutions for g over each component, and thus a total of $2^\kappa \geq 2$ solutions.

5.2. Proof of Proposition 2

Replacing the partial derivatives of f from equations (16,17) in equations (3,4) we obtain:

$$\left(\frac{1}{2}\left(\sqrt{U^2+4V^2}+U\right)+\left(\frac{B_x}{A}-\left(\frac{\partial A}{\partial x}\right)^2\frac{1}{4A^2}\right)\right)f^2 = \lambda$$

$$\left(\frac{1}{2}\left(\sqrt{U^2+4V^2}-U\right)+\left(\frac{B_y}{A}-\left(\frac{\partial A}{\partial y}\right)^2\frac{1}{4A^2}\right)\right)f^2 = \lambda$$

Replacing U by its definition (6) we obtain the local scale function (13). This completes the proof.

6. Isometric Deformation

In the isometric case (whether developable or non-developable), we obtain a very compact, simple and elegant analytical solution given by the following proposition.

Proposition 3 (isometric solution) *The solution of the system of PDEs (1) for an isometric deformation is* $\varphi^\top = \gamma(\eta^\top \ \ 1)$ *(with* η *the known warp) and:*

$$\gamma = \sqrt{\frac{2\lambda}{B_x + B_y + A\sqrt{U^2+4V^2} - \frac{1}{4A}\left\|\frac{\partial A}{\partial \mathbf{q}}\right\|_2^2}}. \qquad (18)$$

$\lambda = 1$ *if the surface is developable and* $\lambda = \sqrt{\det(\mathsf{J}_\Delta^\top \mathsf{J}_\Delta)}$ *otherwise (with* Δ *the known template surface function).*

Theorem 1 follows from proposition 3. This is an important result: it shows that a 3D point can be uniquely reconstructed from the value of the warp and its partial derivatives at this point only, and gives an analytical solution to this. Note also that this result could be used directly in Shape-from-Texture.

The isometric deformation is a special case of conformal deformation, where the amount of local isotropic scaling is given by the known local scale function λ. Therefore, equation (13) holds, and can be used to solve for f instead of λ. Because $\gamma = fA^{-\frac{1}{2}}$ we get equation (18) which completes the proof.

7. Implementation

Here we turn the theory in §§5 and 6 into practical 3D reconstruction algorithms. Template-based 3D reconstruction has two phases. The first phase is to define the template; this phase is done off-line. The second phase is to reconstruct the surface as seen in an input image given the template. This is done on-line.

Off-line 3D template reconstruction. Because the literature mostly addresses the case of developable surfaces, it typically obtains the template from an image showing the object flat, which may be rectified to a fronto-parallel pose. This gives both the template's texture map and parameterization space $\mathcal{P}$; the template surface function is the identity matrix, $\Delta = \mathrm{I}$. In the non-developable case, whether isometric or conformal, the object cannot be physically flattened. We use standard methods to build the template's 3D shape $\mathcal{T}$ of disc topology. This can then be easily conformally flattened, giving the flattening function Δ^{-1}, the template's texture map and its parameterization space $\mathcal{P}$.

Off-line precomputations. In our implementation we express the warp function on the basis of the tensor-product of cubic B-splines with an 8×8 grid of control points. The control points in the template are fixed. Note that with this model, the warp's partial derivatives are obtained analytically. We regularly discretize the template space $\mathcal{P}$ at a large number of points $\{\mathbf{q}_1, \ldots, \mathbf{q}_m\}$, and precompute the value of local scaling at these points in the non-developable isometric and conformal cases as $\lambda_j = \sqrt{\det\left(\mathsf{J}_\Delta^\top(\mathbf{q}_j)\mathsf{J}_\Delta(\mathbf{q}_j)\right)}$, $j = 1, \ldots, m$.

On-line 3D deformable reconstruction. The 3D shape corresponding to a new incoming image is computed in two steps. The first step common to all three types of surface deformations is to estimate the warp η. We use a registration algorithm [12] to robustly compute the control point positions in the input image. The second step then computes the numerical values $\varphi_1, \ldots, \varphi_m$ of the surface function φ at the m sampled points. This computation is different between the isometric and the conformal deformation models.

In the isometric case each point is computed independently. We thus use a simple 'for loop' over $j = 1, \ldots, m$, where we compute $\eta(\mathbf{q}_j)$ then $\gamma(\mathbf{q}_j)$ from equation (18) using the precomputed λ_j, and finally store $\varphi_j^\top = \gamma(\mathbf{q}_j)(\eta(\mathbf{q}_j)^\top \ \ 1)$.

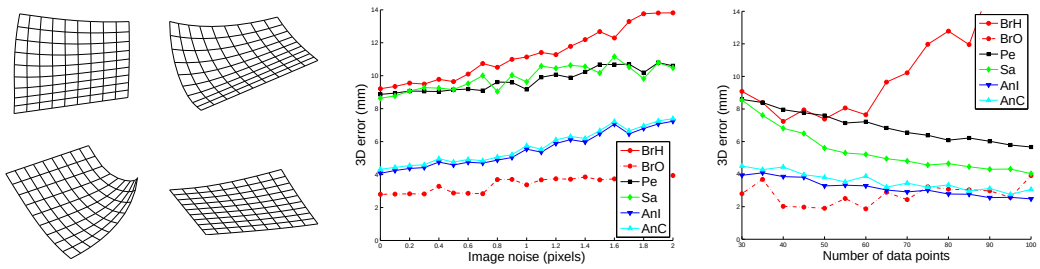

(a) (b) (c)

Figure 2. **Results on synthetic datasets.** (a) Examples of simulated developable shapes. (b) 3D reconstruction error against image noise. (c) 3D reconstruction error against number of data points. **BrO** is a non-convex and non-analytical method that serves as gold standard. the proposed analytical methods **AnI** and **AnC** thus give more accurate results than existing convex methods.

In the conformal case, we first find the connected components $\mathcal{C}_1, \ldots, \mathcal{C}_\kappa$, $\kappa \geq 1$ which cover the parameterization space $\mathcal{P}$ such that $\mathcal{P} = \cup_{s=1}^{\kappa} \mathcal{C}_s$. Each component $\mathcal{C}_s$ has a constant sign function δ, and we thus write $\delta(\mathbf{p}) = \delta_s \in \pm 1$ if $\mathbf{p} \in \mathcal{C}_s$. Following proposition 1, we compute $U_j = U(\mathbf{q}_j)$ and $V_j = V(\mathbf{q}_j)$, $j = 1, \ldots, m$, and find the closed curves $U = V = 0$. We then compute the partial derivatives of g from equations (11,12). At this point, the 2^κ solutions for g can be found by integrating its partial derivatives with all the possible sign combinations for $\delta_1, \ldots, \delta_\kappa$. For efficiency, we however propose to first independently reconstruct each of the κ components. Let $g_s : \mathcal{C}_s \to \mathbb{R}$ and $\varphi_s : \mathcal{C}_s \to \mathbb{R}^3$ be respectively the restrictions of functions g and φ to $\mathcal{C}_s$. The former is obtained up to an unknown additive constant $c_s \in \mathbb{R}$ by integrating the partial derivatives in the least squares sense as $\delta_s g_s + c_s$. The κ corresponding convex and concave surface patches are then parameterized by $\varphi_s = d_s (\exp g_s)^{\delta_s} A^{-\frac{1}{2}} \begin{pmatrix} \eta & 1 \end{pmatrix}^\top$, where $d_s = \exp(c_s)$ is an unknown multiplicative constant. The patches are then stitched together to form a smooth surface for all the 2^κ possible sign combinations for $\{\delta_1, \ldots, \delta_\kappa\} \in [-1; 1]^\kappa$, while computing the patches' arbitrary scales $d_2, \ldots, d_\kappa$ maximizing the agreement where the patches meet (we fix the global scale ambiguity by setting $d_1 = 1$).

Because the surface's continuity constraint is used in the reconstruction process, it cannot disambiguate the conformal solutions, as in [2]. However, other constaints such as surface integrability or shading can be used to select one solution amongst the 2^κ candidates. This might however not resolve the global scale ambiguity. If needed, the values of the 3D surface deformation function Ψ can be obtained easily at the sampled points by composing φ with Δ^{-1}. Also, it is worth of note that many of the calculations in the above algorithms are pointwise and thus highly parallelizable on the GPU. It is perfectly feasible for realtime reconstruction with our methods.

8. Experimental Results

8.1. Compared Algorithms

We compared our algorithms **AnI** (our proposed analytical isometric solution) and **AnC** (our proposed analytical conformal solution) with four other algorithms representing the state of the art in template-based 3D reconstruction using reprojection and deformation constraints. We used the original implementation made available by the authors of these algorithms. These algorithms are all applicable to developable surfaces; they are **Sa** ([14] – a convex numerical solution based on the maximum depth heuristic), **Pe** ([10] – an iterative solution based on the maximum depth heuristic) and **BrH** ([1] – a convex solution based on the maximum depth heuristic with SOCP). In addition, we compare with **BrO** ([1] – a non-analytical and non-convex solution that minimizes the surface's reprojection error with an isometry prior and a smoothness penalty under the developability constraint). **BrO** serves as the gold standard reference.

8.2. Simulated Data

We simulated a bent A4 sheet of paper using a developable surface model [9] as shown in figure 2. We randomly drew m points on the simulated paper's surface and projected them with a calibrated pin-hole camera with intrinsics $\mathsf{K} = \mathrm{diag}(500, 500, 1)$ located 500 mm from the surface to an image of size (244×244). Some gaussian noise with standard deviation σ was then added to the image points. We measured the 3D residual error in mm as the average distance between the simulated and the reconstructed 3D points. For each configuration, we kept the average E of the 3D residual error over 100 trials.

We computed E as a function of σ for $m = 30$ points and as a function of m with $\sigma = 0.5$ pixels. Our results are shown in figure 2. The accuracy of all methods degrades with increasing noise and improves with increasing number of points, except for **BrH** which degrades as the number of points increases. **BrH** almost consistently gives the least accurate result; we believe that this is partly due to

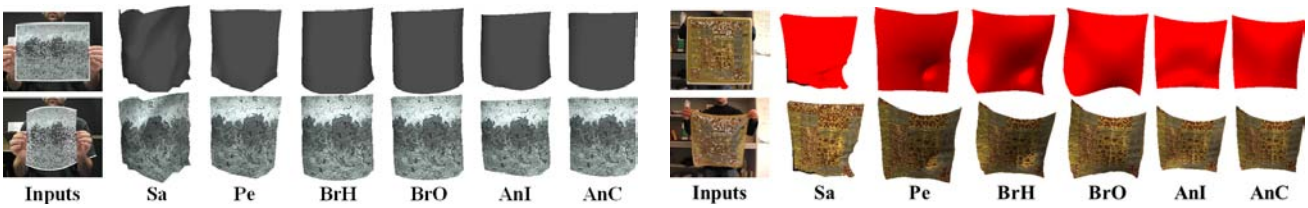

Figure 3. **The paper and cloth datasets.** The inputs and results for the 6 compared algorithms for the paper and cloth datasets, shown on the left and on the right respectively.

the inaccuracy of the warp estimation routine used by this method. **BrO**, the only tested non-convex or non-analytical method, serving here as gold standard, almost consistently gives the most accurate result. **Pe** and **Sa** both outperform **BrH** and have the same behaviour with respect to increasing noise. However, **Sa**'s accuracy improves dramatically with the number of points, and eventually meets the accuracy of **AnC**, **AnI** and **BrO** when the number of points becomes large (greater than 100). Finally, **AnC** and **AnI**, the proposed analytical methods, give very simular results and have the same behaviour with respect to both the amount of noise and number of points, but are slightly more sensitive to noise than the other methods. However, they consistently give better results than the existing convex methods, and meet the accuracy of the gold standard **BrO** for 50 points or more.

8.3. Real Data

We now present results and comparisons of the different methods on three real datasets.

Paper [16]. The 2D parameterization $\mathcal{P}$ is readily available as one image of the dataset showing the surface flat. The registration algorithm [12] returned 87 keypoint matches. They are shown with the reconstructed surfaces computed with all methods in figure 3. Because we do not have ground truth, we report below the average discrepancy in centimeters between surfaces. We clearly see that all methods agree at the order of a centimeter:

	Sa	**Pe**	**BrH**	**BrO**	**AnI**	**AnC**
Sa	0	1.22	1.23	1.23	1.19	1.22
Pe	1.22	0	0.40	0.67	0.13	0.68
BrH	1.23	0.40	0	0.29	0.32	0.31
BrO	1.23	0.67	0.29	0	0.58	0.13
AnI	1.19	0.13	0.32	0.58	0	0.59
AnC	1.22	0.68	0.31	0.13	0.59	0

Cloth [18]. The 2D parameterization $\mathcal{P}$ is readily available as one image of the dataset showing the surface flat. The registration algorithm [12] returned 192 keypoint matches. They are shown with the reconstructed surfaces in figure 3. The average surface discrepancy is large (around 10 centimeters) between any pairs of surfaces. We observe that all methods manage to reconstruct the surface to some extent. **AnC** qualitatively gives a very convincing reconstruction.

Uterus [6]. This dataset shows a human uterus filmed by a laparoscope. The uterus undergoes challenging extensible deformations caused by the motion of a surgery tool. The 2D parameterization $\mathcal{P}$ was obtained by conformal flattening of a mesh with 207 vertices representing the 3D template surface $\mathcal{T}$. Those vertices were reconstructed by rigid Structure-from-Motion from 85 frames showing the uterus without deformations. This is all illustrated in figure 4. For

(a) (b) (c)

Figure 4. **The uterus dataset.** (a) Two input rigid frames. (b) 3D template $\mathcal{T}$. (c) Parameterization space $\mathcal{P}$ (top) and input image $\mathcal{I}$ (bottom), both overlaid with the 57 point matches.

this dataset, as expected, all isometric methods failed to recover a sensible shape. The exception is **AnC** which uses the conformal deformation model and thus tolerates elastic deformations. The reconstructed surfaces are shown in figure 5. Amongst the multiple solutions returned by **AnC**, the correct one can be easily distinguished. The solution obtained from **AnI** is clearly incorrect, as for the other isometric methods.

9. Conclusion

We have presented the first analytical solutions to full perspective template-based 3D reconstruction of developable, isometric and conformal surfaces. While the first case was largely studied in the literature, only iterative convex and non-convex numerical methods were known. Our analysis shows that an isometric surface, whether developable or not, can be uniquely reconstructed, while a conformal surface can be computed up to a set of a few dis-

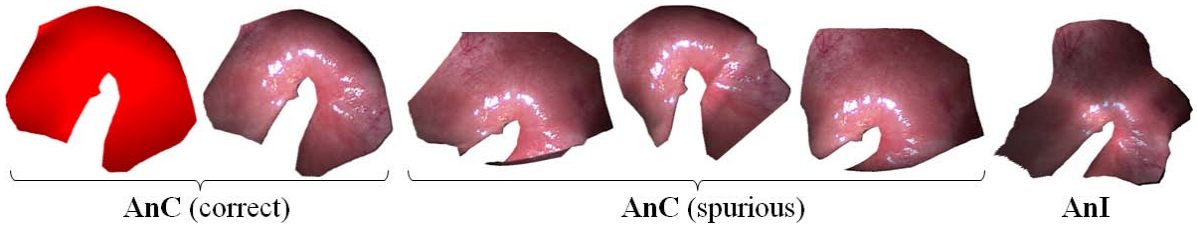

Figure 5. **The uterus dataset.** From left to right: the correct and some of the spurious solutions of **AnC** and the solution of **AnI**.

crete ambiguities and a global scale factor. An additional cue such as shading is needed to resolve these ambiguities. Possible extensions of our work include the study of more flexible models than conformal (such as equiareal deformations and quadratic local models), the study of the uncalibrated camera case and of templateless 3D reconstruction.

Acknowledgments. We are grateful to the authors of [6, 16, 18] for making their datasets available to us and to the authors of [1, 10, 14] for making their algorithms' implementation available to us.

References

[1] F. Brunet, R. Hartley, A. Bartoli, N. Navab, and R. Malgouyres. Monocular template-based reconstruction of smooth and inextensible surfaces. *Asian Conference on Computer Vision*, 2010. 1, 2, 6, 8

[2] A. Ecker, K. Kutulakos, and A. Jepson. Semidefinite programming heuristics for surface reconstruction ambiguities. *European Conference on Computer Vision*, 2008. 1, 2, 6

[3] V. Gay-Bellile, A. Bartoli, and P. Sayd. Direct estimation of non-rigid registrations with image-based self-occlusion reasoning. IEEE *Transactions on Pattern Analysis and Machine Intelligence*, 32(1):87–104, January 2010. 1, 2

[4] N. A. Gumerov, A. Zandifar, R. Duraiswami, and L. S. Davis. 3D structure recovery and unwarping surfaces applicable to planes. *International Journal of Computer Vision*, 66(3):261–281, 2006. 1, 2

[5] A. Malti, A. Bartoli, and T. Collins. A pixel-based approach to template-based monocular 3D reconstruction of deformable surfaces. IEEE *International Workshop on Dynamic Shape Capture and Analysis at ICCV*, 2011. 1, 2

[6] A. Malti, A. Bartoli, and T. Collins. Template-based conformal shape-from-motion from registered laparoscopic images. *Conference on Medical Image Understanding and Analysis*, 2011. 1, 2, 7, 8

[7] F. Moreno-Noguer, J. Porta, and P. Fua. Exploring ambiguities for monocular non-rigid shape estimation. *European Conference on Computer Vision*, 2010. 1, 2

[8] F. Moreno-Noguer, M. Salzmann, V. Lepetit, and P. Fua. Capturing 3D stretchable surfaces from single images in closed form. *International Conference on Computer Vision and Pattern Recognition*, 2009. 1, 2

[9] M. Perriollat and A. Bartoli. A quasi-minimal model for paper-like surfaces. *Workshop "Towards Benmarking Automated Calibration, Orientation and Surface Reconstruction from Images"*, 2007. 6

[10] M. Perriollat, R. Hartley, and A. Bartoli. Monocular template-based reconstruction of inextensible surfaces. *International Journal of Computer Vision*, 95(2):124–137, November 2011. 1, 2, 6, 8

[11] J. Pilet, V. Lepetit, and P. Fua. Fast non-rigid surface detection, registration and realistic augmentation. *International Journal of Computer Vision*, 76(2):109–122, February 2008. 1, 2

[12] D. Pizarro and A. Bartoli. Feature-based non-rigid surface detection with self-occlusion reasoning. *International Journal of Computer Vision*, 97(1):54–70, March 2012. 1, 2, 5, 7

[13] A. Presley. *Elementary Differential Geometry*. Springer-Verlag, 2010. Second Edition. 3

[14] M. Salzmann and P. Fua. Reconstructing sharply folding surfaces: A convex formulation. *International Conference on Computer Vision and Pattern Recognition*, 2009. 1, 2, 6, 8

[15] M. Salzmann and P. Fua. Linear local models for monocular reconstruction of deformable surfaces. IEEE *Transactions on Pattern Analysis and Machine Intelligence*, 33(5), May 2011. 1, 2

[16] M. Salzmann, R. Hartley, and P. Fua. Convex optimization for deformable surface 3-D tracking. *International Conference on Computer Vision*, 2007. 1, 7, 8

[17] M. Salzmann, V. Lepetit, and P. Fua. Deformable surface tracking ambiguities. *International Conference on Computer Vision and Pattern Recognition*, 2007. 1, 2

[18] M. Salzmann, F. Moreno-Noguer, V. Lepetit, and P. Fua. Closed-form solution to non-rigid 3D surface registration. *European Conference on Computer Vision*, 2008. 1, 2, 7, 8

[19] M. Salzmann, J. Pilet, S. Ilic, and P. Fua. Surface deformation models for nonrigid 3D shape recovery. IEEE *Transactions on Pattern Analysis and Machine Intelligence*, 29(8):1–7, August 2007. 1

[20] A. Shaji, A. Varol, L. Torresani, and P. Fua. Simultaneous point matching and 3D deformable surface reconstruction. *International Conference on Computer Vision and Pattern Recognition*, 2010. 1, 2

[21] A. Sheffer, B. Lévy, M. Mogilnitsky, and A. Bogomyakov. ABF++: Fast and robust angle based flattening. SIGGRAPH, 2005. 2