Monocular Template-Based 3D Reconstruction of Extensible Surfaces with Local Linear Elasticity

Abed Malti[1] Richard Hartley[2] Adrien Bartoli[1] Jae-Hak Kim[2]
[1] ALCoV/ISIT UMR 6284 CNRS/Université d'Auvergne, Clermont-Ferrand, France
[2]Australian National University and NICTA, Canberra, Australia

Abstract

We propose a new approach for template-based extensible surface reconstruction from a single view. We extend the method of isometric surface reconstruction and more recent work on conformal surface reconstruction. Our approach relies on the minimization of a proposed stretching energy formalized with respect to the Poisson ratio parameter of the surface. We derive a patch-based formulation of this stretching energy by assuming local linear elasticity. This formulation unifies geometrical and mechanical constraints in a single energy term. We prevent local scale ambiguities by imposing a set of fixed boundary 3D points. We experimentally prove the sufficiency of this set of boundary points and demonstrate the effectiveness of our approach on different developable and non-developable surfaces with a wide range of extensibility.

1. Introduction

Monocular template-based 3D reconstruction and representation of nonrigid objects require models with the ability to assume a wide variety of shapes and to track complex motion. The models must be able to recover a plausible deformed shape (or at least a set of discrete plausible deformed shapes) from noise-corrupted features while making the weakest possible assumptions about the observed shape.

On the one hand, previous approaches based on physical constraints address this problem in the case of isometric [6, 15], and conformal [9, 4] deformations. On the other hand, statistical learning approaches [16, 13] have shown effectiveness only on isometric surfaces. In this paper, we present a new physical-based approach for monocular template-based reconstruction of extensible and compressible surfaces with only one mechanical parameter. As physical priors, our approach relies only on the Poisson ratio of the surface material assumed to be linear Hookean (i.e. constant elasticity). This Poisson parameter models the tendency of a material to be compressed in a transverse direction when it is stretched in the longitudinal direction [5]. Our formulation is based on the principle that any extensible/compressible surface lies in the minimal stretching energy state subject to external applied constraints. Accordingly, we formulate the reconstruction problem as being to estimate the shape that has minimal stretching energy given a set of boundary points, and is consistent with the measured image data. The results are close to the real deformed surface as will be shown later in the experimental results section.

*This research has received funding from the EU's FP7 through ERC grant 307483 FLEXABLE. NICTA is funded by the Australian Government through the Australian Research Council.

2. Related Work and Contribution

Different types of constraints have been proposed and can be categorized as statistical or physical constraints. Statistical constraints often model the deformation as a linear combination of basis vectors, which can be learned online either for human face reconstruction [11] or for generic shapes [15, 13, 18]. Non-linear learning methods were applied in human tracking [14] and then extended for more generic surfaces [16].

Physical constraints include spatial and temporal priors on the surface. In [3] physical constraints are used as priors for a coarse-to-fine shape-basis statistical model. A physical prior that has been studied is the isometry constraint [6], which requires that any surface geodesic distance is preserved after deformation. This approach has proven its accuracy for paper-like surfaces and was recently extended for non-developable surfaces undergoing conformal deformation[9]. More recently, [4] have studied the well-posedness of both isometric and conformal deformations. Their applicability in real world deforming object remains limited (paper like surfaces, deforming balls under isotropic conditions). Thus, the problem of monocular template-based 3D reconstruction of realistic deformations has not yet been tackled. However, a SLAM (Simultaneous Localization And Mapping) method for elastic surfaces was attempted with fixed boundary conditions [1] and later extended to free boundary conditions [2]. This approach re-

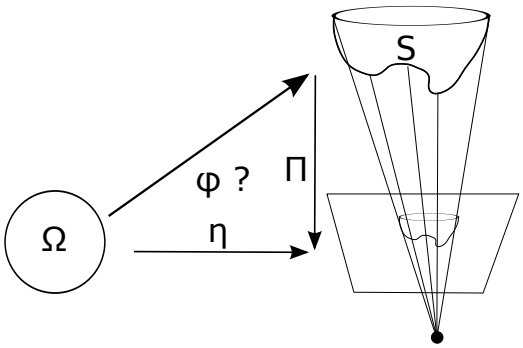

Figure 1. Diagram of monocular template-based 3D reconstruction. The surface is assumed to be homeomorphic to a disc.

lies on the Navier-Stokes fluid-flow equation. It uses an FEM (Finite Element Method) to model the surface and approximate the deformation forces. The surface and the deformation forces are both estimated using an EKF (Extended Kalman Filter). It is important though to distinguish SLAM and monocular template-based reconstruction methods. In the former a time sequential smoother is embedded in a Bayesian a-posteriori estimator where for initialization the scene is required to be static. These assumptions simplify the problem since during tracking small deformations are sequentially integrated and added to the surface. However, in monocular template-based approaches, the rate of deformation of the deformed surface may be very large and multiple solutions may appear.

Contributions. This work goes beyond isometric and conformal deformations to consider elastic deformations for monocular template-based reconstruction. The contributions of the paper are threefold. First, we formalize the reconstruction problem of a generic surface in terms of the minimization of stretching energy. From classic linear elasticity theory [5] we derive a patch-based formulation of the stretching energy of an elastic surface. This formulation unifies geometrical and mechanical constrains in a single energy term. Second, we show that we do not have global scale ambiguities and we prevent local scale ambiguities by imposing a set of fixed boundary 3D points. We experimentally prove that this set of boundary condition sufficiently constrain the solution. Third, we propose an iterative method to solve our extensible surface reconstruction problem.

3. Geometric Priors as Physical Constraints: Isometric and Conformal

A smoothly deforming surface $\mathbf{S}$ homeomorphic to a disc, can be modelled as an embedding of a template Ω into $\mathbb{R}^3$. It is described by a surface function φ of two variables $(u, v) \in \Omega : \varphi : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^3$. The problem of monocular template-based reconstruction can be then summarized in figure 1. It is intuitively constrained by a set of point correspondences, represented by the 2D warp η, between the template and the input image. The nature of the surface deformations are encoded in the differential properties of φ. The Jacobian, denoted J_φ, is a 3×2 matrix which measures and characterizes the local extension of the deformations. The first fundamental form $\mathbf{I}_\varphi$ defined as: $\mathbf{I}_\varphi = \mathsf{J}_\varphi^\top \mathsf{J}_\varphi$, gives the required tensor to measure distances on the deformed shape. It is a 2×2 matrix which maps the local distances from Ω to $\mathbf{S}$. Previous work has focused on two types of mappings: isometric and conformal. If φ is an isometry, then $\mathbf{I}_\varphi$ is the identity. This property has been used as the main constraint in formalizing the problem of monocular template-based 3D reconstruction of isometric surfaces [6]. If φ is conformal, *i.e.* preserves angles, then $\mathbf{I}_\varphi$ is of the form: $\mathbf{I}_\varphi = \lambda\,\mathsf{I}$, where $\lambda : \Omega \rightarrow \mathbb{R}$ controls the amount of local isotropic scaling and I is the identity matrix of dimension 2. This deformation assumption has been used as the main constraint in formalizing the problem of monocular template-based 3D reconstruction of conformal surfaces [9, 4]. In our work, we link this geometric tensor to mechanical priors and then unify geometric and mechanical constraints in one equation, as will be seen in the next section.

4. Mechanical Priors as Physical Constraints: Our Stretching Energy

Our idea is to model the material being deformed as made of some elastic material, and minimizing the deformation energy in trying to fit the surface to the data. For isotropic materials, the deformation energy is computed in terms of the Young's modulus E and the Poisson ratio ν [5].

4.1. Background

Consider a rod of some material with rest length L, which is longitudinally stretched with a force F. Young's modulus expresses the relationship of the force to the extension of the rod:

$$F = E\,dL/L\,, \tag{1}$$

where dL/L is the relative extension of the rod, and E is Young's modulus. It is measured in Pascals (Newtons per square metre – the square metres relating to the cross-section of the rod).

The Poisson ratio ν models the tendency of a material to become thinner when it is stretched. A value of 0 indicates that no thinning takes place. To be exact, for a unit cube of material,

$$-\nu = \frac{\Delta_y}{\Delta_x} = \frac{\Delta_z}{\Delta_x}$$

where Δ_x, Δ_y and Δ_z are small changes in dimension when the cube is stretched (or compressed) in the X direction without constraints in the other directions, as shown in figure 2-a. For most materials, $0 \leq \nu \leq 0.5$, though for some materials, $\nu < 0$ which means that the material will expand laterally when stretched, or equivalently shrink when

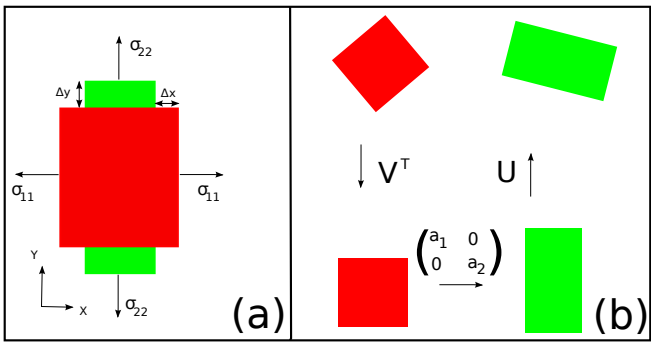

Figure 2. $2D$ deformation of a surface of negligible thickness. (a): stretching the surface in the X direction introduces a compression in the Y direction. (b): any deformation can be seen as stretching/compression in the direction of the singular vectors of the Jacobian of the deformation surface function φ.

compressed. Rubber has a value of ν close to 0.5, which means that it does not significantly change volume when compressed in one direction. Cork has a value close to zero, it does not expand laterally when compressed or stretched. This explains its usefulness in stopping wine bottles.

In the case of two-dimensional sheets, the relationship between the stress σ (infinitesimal force) and strain ϵ (infinitesimal displacement) tensors is given by Hooke's Law [5]:

$$\begin{pmatrix}\epsilon_{xx}\\ \epsilon_{yy}\\ \epsilon_{xy}\end{pmatrix} = \frac{1}{E}\begin{pmatrix}1 & -\nu & \\ -\nu & 1 & \\ & & 1+\nu\end{pmatrix}\begin{pmatrix}\sigma_{xx}\\ \sigma_{yy}\\ \sigma_{xy}\end{pmatrix}, \qquad (2)$$

and hence

$$\begin{pmatrix}\sigma_{xx}\\ \sigma_{yy}\\ \sigma_{xy}\end{pmatrix} = \frac{E}{1-\nu^2}\begin{pmatrix}1 & \nu & \\ \nu & 1 & \\ & & 1-\nu\end{pmatrix}\begin{pmatrix}\epsilon_{xx}\\ \epsilon_{yy}\\ \epsilon_{xy}\end{pmatrix}. \qquad (3)$$

where σ_{xx} is the stress along the x axis, σ_{yy} is the stress along the y axis and σ_{xy} is the mutual stress (shearing). The ϵs' indices represents the strains in the corresponding directions.

4.2. Patch-Based Stretching Energy

One dimensional stretching. Let us return to the 1-dimensional case, and consider the work done in extending a rod (or a spring) with initial length 1 and Young's modulus E. Suppose it is stretched to length L and let us compute the expended energy.

Suppose that we are pulling on the end of the rod, which at some moment has length ℓ. The force being applied at that moment is equal to $E(\ell-1)$. If it is extended to length $\ell+d\ell$, then the energy expended is $E\,(\ell-1)d\ell$. Integrating to get the total work done, we find

$$W = E\int_1^L (\ell-1)d\ell = (L-1)^2/2\,.$$

This makes sense, since the force increases proportionally from 0 to $E(L-1)$ during the stretching process. This shows the necessity of using some infinitesimal process when computing the energy expended.

Two dimensional stretching. Consider a unit square, transformed by a linear transformation A. We could consider an affine transformation, but that would make no real difference. The question is, what is the deformation energy of the material in this deformed state? Clearly it is equal to the expended energy used to deform it. The first thing we want to do is to simplify the computation, by simplifying the deformation. Thus, suppose that the matrix A has a Singular Value Decomposition (SVD) $\mathsf{A} = \mathsf{U}\mathsf{D}\mathsf{V}^\top$ with a_1 and a_2 as real singular values. As shown in figure 2-b, the transformation takes place in three stages, first a rotation by $\mathsf{V}^\top$, then axial stretching by a_1 and a_2, then rotation by U. It is clear that rotation does not change the stored energy, so we have reduced the problem to one of axial stretching.

Clearly, the amount of energy expended depends proportionally on the area of the piece of material being stretched, so we consider a unit square. We want to deform this to dimension $a_1 \times a_2$.

According to (3), when the dimensions of the square are $x \times y$, the force on the faces will be

$$\begin{pmatrix}\sigma_{xx}\\ \sigma_{yy}\end{pmatrix} = \frac{E}{1-\nu^2}\begin{pmatrix}1 & \nu\\ \nu & 1\end{pmatrix}\begin{pmatrix}x-1\\ y-1\end{pmatrix},$$

since the strain $\epsilon_{xy} = 0$ (the SVD decomposition allows us to account only for axial deformations). Ignoring the constant $E/(1-\nu^2)$, the work done in increasing the dimension by (dx, dy) is proportional to

$$\begin{aligned} dW &= (dx,dy)\begin{pmatrix}1 & \nu\\ \nu & 1\end{pmatrix}\begin{pmatrix}x-1\\ y-1\end{pmatrix} \qquad (4)\\ &= (x-1)dx + (y-1)dy + \nu(x-1)dy + \nu(y-1)dx\\ &= \left((x-1)+\nu(y-1)\right)dx + \left((y-1)+\nu(x-1)\right)dy\,. \end{aligned}$$

The complete energy taken to deform from an initial shape 1×1 to $a_1 \times a_2$ is proportional to the integral of this over some path from $(1,1)$ to (a_1, a_2).

We can deform first over x (with $y = 1$) then over y (with $x = a_1$), giving

$$\begin{aligned} &\int_1^{a_1}(x-1)\,dx + \int_1^{a_2}\left((y-1)+\nu(a_1-1)\right)dy\\ &= \left((a_1-1)^2+(a_2-1)^2\right)/2 + \nu(a_1-1)\,(a_2-1)\,. \end{aligned} \qquad (5)$$

From symmetry of this expression, it is obvious that we get the same result if we deform first over y, then over x. From physical considerations, the integral must be the same over any path from $(1,1)$ to (a_1, a_2); this may be verified. This derivation allows us to state the following theorem:

Theorem 1 *Let a planar region Ω be deformed by a differentiable deformation mapping φ. Let $a_1(u,v)$ and $a_2(u,v)$ be the singular values of the Jacobian J_φ defined at point $(u,v) \in \Omega$. Then the deformation energy of the mapping φ is given by*

$$\mathcal{E}_s[\varphi] = \frac{E}{2(1-\nu^2)} \iint_\Omega (a_1(u,v)-1)^2 + (a_2(u,v)-1)^2 + 2\nu(a_1(u,v)-1)(a_2(u,v)-1)\, du\, dv\,. \tag{6}$$

This theorem unifies the geometric and the mechanic properties of a deformable surface; it describes any deformation as an integration of local deformations of patches along the principal axes of stretching/shrinking. These axes correspond to the eigenvectors of the first fundamental form. The amount of deformation is measured by the singular values a_1 and a_2. The Young's modulus is present as a multiplication factor and the minimization of the stretching energy is independent of its values.

Let us compare this formula with the error term used in [6] for isometric deformations. For a linear transformation A the term used there is $\|\mathsf{A}^\top\mathsf{A}-\mathsf{I}\|^2$. (this is applied in [6] to the Jacobian of the transformation, so $\mathsf{A} = \mathsf{J}_\varphi$). In this case, setting $\mathsf{A} = \mathsf{U}\mathsf{D}\mathsf{V}^\top$, we see that $\mathsf{A}^\top\mathsf{A} - \mathsf{I} = \mathsf{U}(\mathsf{D}^2 - \mathsf{I})\mathsf{U}^\top$, and $\|\mathsf{A}^\top\mathsf{A} - \mathsf{I}\|_\mathsf{F}^2 = (\mathsf{a}_1^2 - 1)^2 + (\mathsf{a}_2^2 - 1)^2$ which is similar to, but not quite the same thing as the energy term given by replacing $\nu = 0$ in equation (6). In particular, it involves 4-th degree terms in the deformation values, which would appear to overly penalize large deformations. By contrast, equation (6) involves only quadratic terms.

5. Variational Formulation

Let us assume a set $\mathcal{C}$ of point correspondences between the template Ω and the deformed surface $\mathbf{S}$. Minimizing the stretching energy under this natural boundary condition means (for perspective cameras) sliding the shape along the sight lines until the shape lies in a minimum stretching energy. However, for a set of adjacent local patches it is still possible that local depth variations can compensate for local stretching energy among adjacent patches and provide the same global stretching energy. Characterizing local ambiguities would require one to investigate an analytic solution of this problem. Such a study is a very hard problem since the equation (6) is non-linear. In order to overcome this issue, we experimentally prove in §7 that a set of 3D boundary points are sufficient to effectively prevent local ambiguities. Henceforth, if we assume that a set $\mathcal{F}$ of sufficient boundary $3D$ points are known, then the reconstruction problem can be formalized in the following variational approach:

$$\begin{aligned} &\min_{\varphi} \mathcal{E}_s[\varphi] && \textbf{(Minimal energy)} \\ &\text{s.t.} \\ &\varphi(u_i, v_i) = \mathbf{Q}_i,\ i \in \mathcal{F} && \textbf{(3D fixed boundaries)} \\ &\Pi(\varphi(u_j, v_j)) = \eta(u_j, v_j),\ j \in \mathcal{C} && \textbf{(Image fitting)} \end{aligned} \tag{7}$$

where $\mathbf{Q}_i,\ i \in \mathcal{F}$ is the set of known boundary points, Π is the matrix of camera intrinsics assuming perspective projection, and $(u_i, v_i),\ i \in \mathcal{C}$, are the template points that match points $\eta(u_i, v_i)$ in the input image.

6. Implementation

In order to solve the variational problem stated in equation (7), we propose a two step practical solution: first initialization through a sub-optimal convex formulation, then refinement using non-linear optimization.

Initialization. In the first step, we roughly estimate the 3D coordinates of the correspondence points by maximizing their depths. This convex formulation was used in [17, 6] for isometric surfaces in a formulation using SOCP (Second Order Cone Programming). In our implementation we relax the isometric condition to:

$$\|\mathbf{v}'_i - \mathbf{v}'_j\| \leq \kappa \|\mathbf{v}_i - \mathbf{v}_j\| + \epsilon_\tau, \quad i,j \in \mathcal{C} \tag{8}$$

where $\mathbf{v}_i$ and $\mathbf{v}_j$ are known 3D points from the template Ω and their corresponding unknown 3D points are $\mathbf{v}'_i$ and $\mathbf{v}'_j$ in the deformed surface. Further, ϵ_τ is a small real value which models the tolerance to noise in correspondences, and κ is a real parameter chosen so that edges are able to stretch. As in [17, 6], $\mathbf{v}_i$ and $\mathbf{v}_j$ are chosen to be closest neighbours.

Boundary points. We also add to this classic convex formulation a further constraint in the form of boundary points by setting the following condition:

$$\| \mathbf{v}'_i - \mathbf{Q}_i \| \leq 0, \quad i \in \mathcal{F} \tag{9}$$

In the case where the boundary points are rigid parts of the template, we compute their 3D poses in the camera frame using the EPnP source code from [12][1].

Non-Linear optimization. In order to minimize the stretching energy, we propose to parametrize the surface function φ by fitting the 3D points from the initialization step to a cubic B-spline [7] defined in the parametrization domain $\Omega \subset \mathbb{R}^2$. Then a discrete set of surface points is uniformly sampled to compute a discrete sum of the stretching energy. The minimization problem (7) is then reformulated as:

$$\min_{\mathbf{L}_\varphi} \underbrace{\sum_{i\in\Omega} \left[(a_1(u_i,v_i)-1)^2 + (a_2(u_i,v_i)-1)^2 + 2|\nu(a_1(u_i,v_i)-1)(a_2(u_i,v_i)-1)|\right]}_{\text{Elasticity}} + \lambda_1 \underbrace{\sum_{i\in\mathcal{F}} \|\varphi(u_i,v_i) - \mathbf{Q}_i\|}_{\text{(3D fixed boundaries)}} + \lambda_2 \underbrace{\sum_{i\in\mathcal{C}} \|\Pi(\varphi(u_i,v_i)) - \eta(u_i\, v_i)^\top\|}_{\text{(Image fitting)}} \tag{10}$$

[1] http://cvlab.epfl.ch/software/EPnP/index.php

λ_1 and λ_2 are two real values that tune the importance of the fixed boundary and the image fitting conditions. We experimentally set λ_1 to 0.8 and λ_2 to 0.2.

The singular values $a_1(u_i, v_i)$ and $a_2(u_i, v_i)$ are computed in closed form from the matrix $\mathbf{I}_\varphi = \mathsf{J}_\varphi^\top \mathsf{J}_\varphi$ at each point. This is more efficient that using SVD.

7. Experimental Results

7.1. Compared Methods and Parameters

We compared our algorithm **Po** (minimizing stretching energy expressed with Poisson's ratio) with four algorithms representing methods of monocular template-based 3D reconstruction: **Sa** ([15]- a convex numerical solution based on the maximum depth heuristic), **Br** ([6] – an iterative and non-convex solution that minimizes the surface's reprojection error with an isometry prior and a smoothness penalty), **Ma** ([9] – an iterative and non-convex solution that minimizes the surface's reprojection error with a conformal prior and a smoothness penalty) and **Ba** ([4] – an analytic solution with a conformal prior). We evaluated these methods on both developable and non-developable surfaces by varying the amount of deformation. The amount of extensibility is expressed as a percentage of the relative variation of the stretching energy with respect to the ground-truth template.

7.2. Synthetic Data

In order to obtain physically meaningful elastic deformations we used *3D Studio Max* [10]. We used paper-like material ($\nu = 0$) and rubber material ($\nu = 0.5$). We also used two different template shapes: flat and ball-like surfaces. The template of the **developable-template** (flat) data sets has a size of 200×300 mm^2 and the template of the **non-developable-template** (ball-like) data sets has a size of $70 \times 100 \times 140$ mm^3. We randomly drew 100 points on the simulated deformed surfaces and projected them with a calibrated perspective camera with intrinsics $\Pi = \text{diag}(450, 450, 1)$ located 600 mm from the surface. Gaussian noise with standard deviation $\sigma = 1$ pixel was added to the image points. The set of known boundary points is represented by 20 points randomly drawn from the surface. We measured the 3D residual error in mm as the average distance between the simulated and the reconstructed 3D points. For each extensibility ratio, we kept the average of the 3D residual error over 100 samples. The results are shown in figure 3: figures 3-a and 3-c show some examples of deformed surfaces. Figure 3-b shows the average reconstruction error for the **developable-template** data sets comparing the five algorithms. It can be seen that the error of reconstruction increases when the extensibility ratio increases except for our algorithm **Po** which remains more stable. For isometric deformations (0% of extensibility ratio) the accuracy of the iterative methods **Br**, **Ma** and **Po** is high and similar. We do not display the accuracy result of **Br** for extensibility higher than 15% to keep the error scale of the most accurate methods. The analytic conformal solution **Ba** is less accurate than the iterative conformal method **Ma** because it is more sensitive to noise. From 10% of extensibility ratio the accuracy of the conformal methods **Ba** and **Ma** substantially decreases because the conformal constraint becomes too strong to model such amounts of elastic deformations. Figure 3-d shows the average reconstruction error for the **non-developable-template** data sets comparing the three algorithms. We did not include **Sa** and **Br** because they are designed for developable-template surfaces. Similar observations can be made for the **non-developable-template** data sets. We note that our algorithm **Po** presents similar accuracies for both data sets since it deals with the extensibility locally it can handle both developable-template and non-developable-template surfaces.

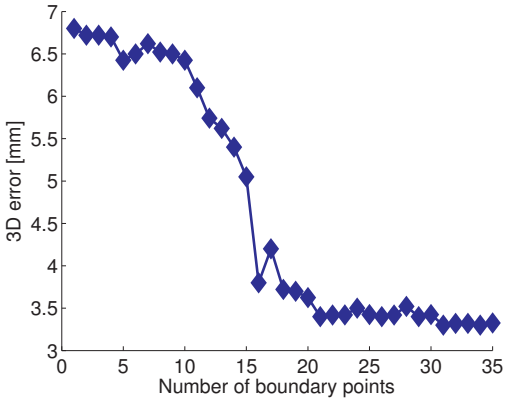

Figure 4. 3D error with respect to the number of boundary points for the synthetic data sets.

Necessary and sufficient condition for boundary points. To experimentally prove that the presence of boundary points are necessary and sufficient, we run under the conditions of our synthetic data set an experiment where we varied the number of boundary points from 1 to 35. The ratio of extensibility ranges from 0% to 30% with a step of approximately 2.5%. The boundary points were randomly sampled on the synthetic deformed surfaces. The average reconstruction error is shown in figure 4. It appears that over 17 points, the error is stabilized around 3.4 mm. This result confirms our belief that there is no global scale ambiguity that would produce huge errors. However, local ambiguities may appear due to local compensation between shrinking and extensions. To reduce these effects it is then sufficient to constrain the minimization of our stretching energy with 3D boundary points.

7.3. Real Data

To acquire real data sets with ground-truth, we used a stereo camera system consisting of two Point Grey Grasshopper cameras. The two cameras have a resolution of 640×480 pixels. A set of 10 deformed shapes was used for each material of the following data sets.

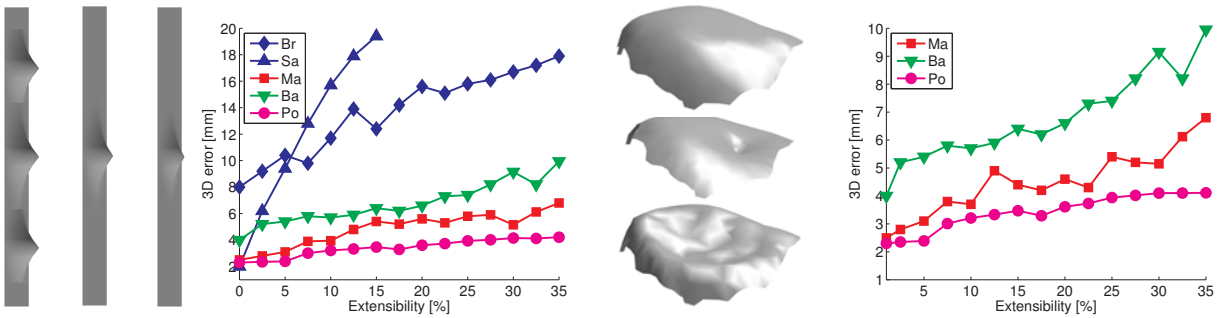

(a) (b) (c) (d)

Figure 3. Results on Synthetic data. **Developable-template** data sets: (a) Examples of simulated shapes. (b) 3D error with respect to the extensibility ratio. **Non-developable-template** data sets: (c) Examples of simulated shapes. (d) 3D error with respect to the extensibility ratio.

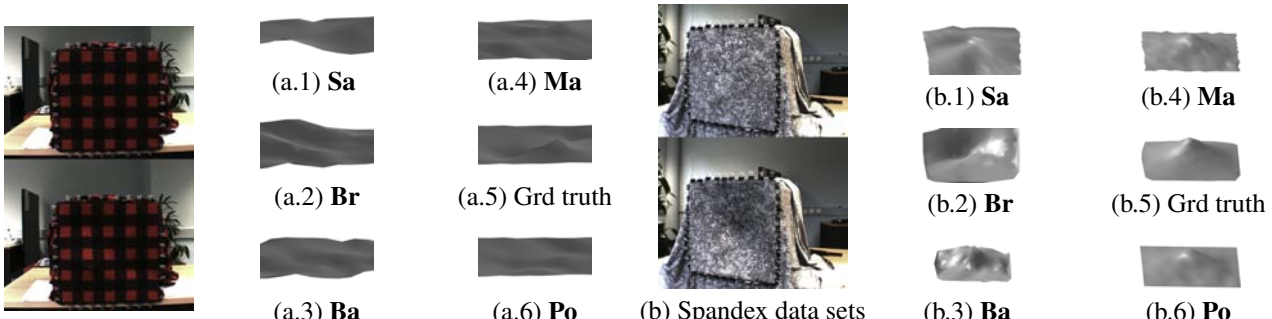

(a) **Redchecker** data sets (b) Spandex data sets

Figure 5. 3D reconstruction of **Developable-template** real data sets.

Developable-template real data sets. As real stretchable and compressible developable surfaces, we used a **redchecker** pattern fabric made of polyester ($\nu = 0.3$) and an irregular pattern **spandex** ($\nu = 0.5$). Both materials can have a large amount of stretching deformation within their tear limit. As demonstrated, 3D boundary points are required in our algorithm, so both materials are clipped on the edge of a box to hold the surface during image acquisition. Features are obtained manually or semi-automatically with SIFT [8]. For the **redchecker** pattern, we took advantage of its corner points for manual-driven point selection. A total of 100 feature points and 36 fixed boundary points were used for this data set. For the **spandex**, features are found by interactive point selection with initial matches from SIFT, which is important in highly-deformed areas, where it is difficult to find point correspondences automatically. A total of 212 feature points and 45 fixed boundary points were used for this data set. The extensibility ratio ranges from 1% to 31%. The 3D error was measured at about 1.3 mm for our **Po** method, 5.3 mm for **Ma** and 6.9 mm for **Ba**. See figures 8-a.1 and 8-a.2 for a summary of 3D errors with **developable-template** real data sets.

Non-developable-template real data sets. As non-developable surfaces, we used a **balloon** pattern fabric made of rubber ($\nu = 0.5$), a **cap** made of quasi-inextensible polyester material ($\nu \approx 0$), an in-vivo piglet **bladder** and a piece of **heart-tissue** ($\nu = 0.35$, experimentally set). See figures 6 and 7 for examples of reconstruction. *(i)* **balloon** data sets: the balloon is clipped to a ring in order to fix its borders. A total of 95 feature points and 30 fixed boundary points were used for this data set. The extensibility ratio ranges from 0% to 22%. The 3D error was measured about 1.6 mm for our **Po** method, 4.3 mm for **Ma** and 8.9 mm for **Ba**.. *(ii)* **cap** data sets: in this data set we deformed the top of the cap and we kept rigid the surrounding part. The boundary points were estimated by EPnP as explained in §6. We used a set of 40 boundary points and a set of 130 point correspondences. The extensibility ratio was almost zero and $\nu \approx 0$. The 3D error was measured at about 2.5 mm for our method **Po**, 4.1 mm for **Ma** and 7.2 mm for **Ba**. *(iii)* **Bladder** and **heart-tissue** data sets. To demonstrate a real application of our proposed method we reconstructed elastic deformations from in-vivo piglet tissue. If tissues are deformed with rigid tools during surgery then we

(a.1) **Balloon** (a.2) **Ba** (a.3) **Ma** (a.4) Grd truth (a.5) **Po** (b.1) **Cap** (b.2) **Ba** (b.3) **Ma** (b.4) Grd truth (b.5) **Po**

Figure 6. 3D reconstructions of **Non-developable-template** real data sets: **balloon** and **cap** data sets.

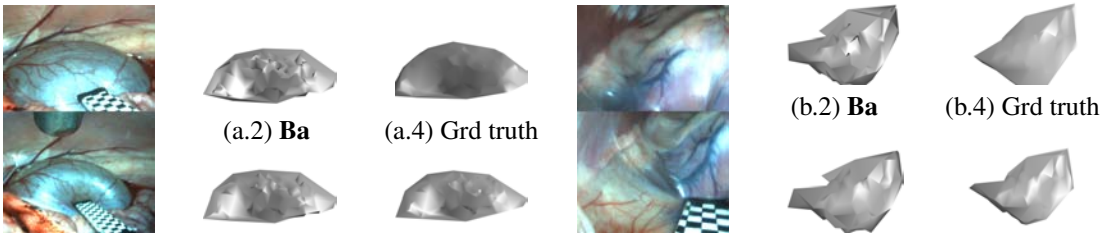

(a.1) **Bladder** (a.2) **Ba** (a.3) **Ma** (a.4) Grd truth (a.5) **Po** (b.1) **Heart-tissue** (b.2) **Ba** (b.3) **Ma** (b.4) Grd truth (b.5) **Po**

Figure 7. 3D reconstructions of **Non-developable-template** real data sets: **bladder** and **heart-tissue** data sets.

can infer boundary conditions on the contact between tool and tissue if we assume that the rigid shape of the tool is known. For this purpose we stitch a checkerboard on a flat tool, see figure 7-a.1-b.1, that we use to deform a bladder and a heart tissue of a piglet in vivo. In this experiment we used the edge contact of the tool consisting of 10 pre-calibrated points as a boundary condition. We also used the stereo setup to infer ground-truth 3D points. A total of 81 and 72 point correspondences were used for the **bladder** and the **heart-tissue** respectively. The extensibility ratio ranges from 5% to 20%. The 3D error was measured for both **Bladder** and **heart-tissue** data sets and were about 2.2 mm for our method **Po**, 7.3 mm for **Ma** and 11.9 mm for **Ba**. The 3D errors with the **Non-developable-template** real data sets are summarized in figure 8-b.

8. Applicability of the Method

The surface reconstruction method developed in this paper is based on the idea of minimizing deformation energy subject to constraints imposed by the template-to-image correspondences. For this method to be effective, a necessary requirement is that the surface to be reconstructed actually should be in a minimum energy state. If the surface is deformed by external forces or constraints, known or unknown, then this may not be the case. An inflated balloon or a stretched piece of rubber will not be in a minimum energy state, except with respect to the space of configurations dictated by the external constraints. This is why we need to include boundary conditions, fixing the position of certain points on the surface. In cases where the surface is deformed by use of an externally applied tool, as in the case of the surgical application, or by fixing the border of the reconstructed object, such constraints may be applied.

The image correspondences provide constraints that will constrain the reconstructed surface to lie in an energy state that is not minimal within the space of fully unconstrained surfaces. The effect of the image constraints is to limit the space of possible configurations, so that the reconstructed surface may assume the minimal energy configuration within a constrained space. This is appropriate for surfaces not under significant tension, such as cloth. It is important to note that for elastic surfaces such as rubber under tension, the image constraints are not sufficient. Image constraints allow points on surfaces to move along rays through the camera centre, defined by the point matches. A surface under tension has a bias to shrink along the image rays towards the camera centre, resulting in a reconstructed image smaller than the ground-truth, but with less stretching energy. Application of further position constraints mitigates this tendency.

Materials such as cloth are not well modelled by an isometric deformation. They deform quite differently from more exactly isometric materials such as paper. A piece of crumpled or ruffled cloth typically takes a shape that is incompatible with the assumption of isometric deformation, which implies a non-curved or developable surface (having zero Gaussian curvature) made up of straight lines. For such deformations, nearly isometric, but not quite, the minimal energy model is appropriate and effective.

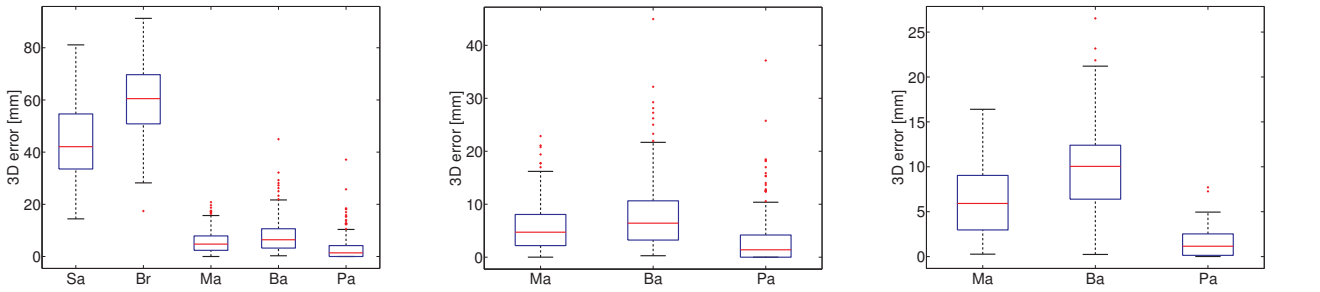

(a.1) **Developable-template** real data sets (a.2) zoom on (a.1) (b) **Non-developable-template** real data sets

Figure 8. 3D error on real data sets.

9. Conclusion

In this paper we presented a new approach to monocular template-based 3D reconstruction of elastic surfaces. Our method relies on a single mechanical parameter, the Poisson ratio. Our analysis defines the deformation energy with respect to this parameter by assuming linear elasticity of local surface patches. Our method is further constrained by a set of boundary points for which we experimentally proved the necessary and sufficient condition. Synthetic and real data experiments have proven the effectiveness of our approach above classic isometric and conformal reconstruction method for elastic surfaces. For future work, it may be fruitful to formally investigate the set of local ambiguities and derive formal properties of the boundary conditions.

Acknowledgment

We gratefully thank the authors of [15], [6], [9] and [4] for providing their source code.

References

[1] A. Agudo, B. Calvo, and J. Montiel. Fem models to code non-rigid ekf monocular slam. *IEEE Workshop on Dynamic Shape Capture and Analysis of ICCV*, 2011. 1

[2] A. Agudo, B. Calvo, and J. Montiel. Finite element based sequential bayesian non-rigid structure from motion. *CVPR*, 2012. 1

[3] A. Bartoli, V. Gay-Bellile, U. Castellani, J. Peyras, S. Olsen, and P. Sayd. Coarse-to-fine low-rank structure-from-motion. *CVPR*, 2008. 1

[4] A. Bartoli, Y. Gerard, F. Chadebecq, and T. Collins. On template-based reconstruction from a single view: Analytical solutions and proofs of well-posedness for developable, isometric and conformal surfaces. *CVPR*, 2012. 1, 2, 5, 8

[5] A. P. Boresi, R. J. Schmidt, and O. M. Sidebottom. *Advanced Mechanics of Materials*. Wiler, 1993. 1, 2, 3

[6] F. Brunet, R. Hartley, A. Bartoli, N. Navab, and R. Malgouyres. Monocular template-based reconstruction of smooth and inextensible surfaces. *ACCV*, 2010. 1, 2, 4, 5, 8

[7] P. Dierckx. *Curve and surface fitting with splines*. Oxford University Press, 1993. 4

[8] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 6

[9] A. Malti, A. Bartoli, and T. Collins. Template-based conformal shape-from-motion from registered laparoscopic images. *MIUA*, 2011. 1, 2, 5, 8

[10] M. Matossian. *3D Studio MAX 3: Visual QuickStart Guide*. Peachpit Press, 1999. 5

[11] I. Matthews and S. Baker. Active appearance models revisited. *International Journal of Computer Vision*, 60(2):135–164, November 2004. 1

[12] F. Moreno-Noguer, V. Lepetit, and P. Fua. Accurate non-iterative o(n) solution to the pnp problem. *ICCV*, 2007. 4

[13] F. Moreno-Noguer and J. Porta. Probabilistic simultaneous pose and non-rigid shape recovery. *CVPR*, 2011. 1

[14] R. Navaratnam, A. W. Fitzgibbon, and R. Cipolla. The joint manifold model for semi-supervised multi-valued regression. *ICCV*, 2007. 1

[15] M. Salzmann and P. Fua. Reconstructing sharply folding surfaces: A convex formulation. *CVPR*, 2009. 1, 5, 8

[16] M. Salzmann and P. Fua. Linear local models for monocular reconstruction of deformable surfaces. IEEE *Transactions on Pattern Analysis and Machine Intelligence*, pages 931–944, 2011. 1

[17] M. Salzmann, R. Hartley, and P. Fua. Convex optimization for deformable surface 3-D tracking. *ICCV*, 2007. 4

[18] M. Salzmann and R. Urtasun. Beyond feature points: Structured prediction for monocular non-rigid 3d reconstruction. *ECCV*, 2012. 1