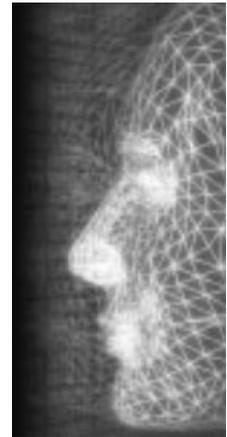


Robust deformation capture from temporal range data for surface rendering

By Umberto Castellani*, Vincent Gay-Bellile and Adrien Bartoli



Imagine an object such as a paper sheet being waved in front of some sensor. Reconstructing the time-varying 3D shape of the object finds direct applications in computer animation. The goal of this paper is to provide such a deformation capture system for surfaces. It uses temporal range data obtained by sensors such as those based on structured light or stereo. So as to deal with many different kinds of material, we do not make the usual assumption that the object surface has textural information. This rules out those techniques based on detecting and matching keypoints or directly minimizing color discrepancy. The proposed method is based on a planar mesh that is deformed so as to fit each of the range images. We show how to achieve this by minimizing a compound cost function combining several data and regularization terms, needed to make the overall system robust so that it can deal with low quality datasets. Carefully examining the parameter to residual relationship shows that this cost function can be minimized very efficiently by coupling nonlinear least squares methods with sparse matrix operators. Experimental results for challenging datasets coming from different kinds of range sensors are reported. The algorithm is reasonably fast and is shown to be robust to missing and erroneous data points. Copyright © 2008 John Wiley & Sons, Ltd.

Received: 9 April 2008; Revised: 11 September 2008; Accepted: 11 September 2008

KEY WORDS: deformable models; 3D registration; Levenberg–Marquardt; bending energy; motion capture

Introduction

Capturing time-varying surface deformations is an important problem which has been recently tackled thanks to the advances in real-time range sensors engineering.^{1–3} A range sensor gives images where the intensity of a pixel is to be interpreted as its depth. Research on deformable models is active in the fields such as computer graphics for 3D morphing and animation,⁴ medical image processing for data alignment and segmentation,^{5,6} and computer vision, for example, contour detection,⁷ face synthesis and expression recognition.⁸

The aim in this paper is to provide an effective system that allows one to capture the deformation of a physical surface being smoothly waved, such as the page of

a book being turned, in front of a range sensor. As opposed to most of the existing methods,^{9–13} we do not use keypoint tracks^{9,10,12} or optical flow.^{11,13} Instead, we only use the range images. This has the advantage that the method is independent of the appearance of the surface. In particular, those textureless surfaces that defeat existing methods are handled. We have to face the two problems of reconstructing the observed surface at each time instant, and registering the reconstructed surfaces. While most of the existing approaches bypasses the first problem,^{14–17} the one we propose solves both simultaneously. We target applications such as 3D data compression, augmented reality and computer animation, all requiring an accurate registration of the range point clouds over time, as well as a reconstruction of the underlying surface. As an example, videos can be synthesized from the captured deformations and a user-provided texture map. The advantage of using range data is that real 3D deformations are obtained, as opposed to a 2D flow field in the monocular image registration case,

*Correspondence to: U. Castellani, Department of Computer Science, University of Verona, Italy. E-mail: umberto.castellani@univr.it

see for example,¹⁸ The input data are a sequence of range images and a coarse boundary of the surface of interest. The whole process is highly robust, filling in possible holes in the range data and detecting erroneous points, while establishing reliable registration even for flat areas that usually defeat shape-based registration methods. We model the surface with a mesh as^{18,19} do for 2D image registration.

Our framework is implemented through two main lines of contributions extending our previous work.²⁰ First, we show that the problem is well modeled by using a mesh that is deformed to fit each point cloud. This model allows us to write a cost function whose global minimum is the sought after solution. This cost function has several data and penalty terms. The data terms incorporate surface to data points distances and boundary information. Furthermore, it explicitly embeds a min operator, thus avoiding the traditional two steps in iterated closed point (ICP)-like algorithms through distance transform. The penalty terms include spatial, that is, surface-related and temporal smoothness as well as inextensibility of the surface, if applicable. The data terms are robustified in order to deal with missing and erroneous points. Second, following,²¹ we use the Levenberg–Marquardt (LM) algorithm to minimize the error function. A careful analysis reveals that the Jacobian matrix involved in the normal equations to be solved at each iteration is highly sparse, for all the data and penalty terms we use. This makes tractable and fast the estimation of dense deformation fields.

Roadmap. The Section ‘Previous Work’ describes the state-of-the-art. The problem statement is given in Section ‘Problem Statement’, and the minimization strategy is described in Section ‘Minimization Procedure’. Experimental results are reported in ‘Experiments’ Section. Finally, conclusions are drawn in ‘Discussion’ Section.

Previous Work

We review existing work on surface motion capture and non-rigid shape registration.

Surface Motion Capture

Several works have been proposed in computer graphics on surface deformation modeling, especially for cloth

motion capture.^{9–13} In Reference [10] a calibrated multi-camera setup is proposed to observe real-time cloth deformation. SIFT keypoints are used to identify a pattern printed on the cloth. The keypoints are tracked in the video thanks to a seed-and-grow approach adapted to the deforming geometry. In Reference [9], quad-marked surfaces are tracked by using a multi-baseline stereo system. Markov Random Fields models are used to introduce additional assumptions on the surface appearance and neighborhood consistency. In Reference [12] the authors use a color-coded cloth texture for reliably matching circular features between different camera views. In Reference [13], a direct estimation of the deformable motion parameters is proposed for range image sequences. The range flow is estimated by introducing depth constraints to the 2D displacements. Similarly in Reference [11], the optical flow is computed from frame to frame using the depth.

Nonrigid Shape Registration

The registration of 3D point clouds is a challenging topic mainly tackled in the framework of ICP for rigid scenarios.^{21,22} However, researchers have recently addressed the case of deformable objects. Roughly speaking the literature on nonrigid registration can be divided into two main categories. The first one directly uses the point clouds. The second one abstracts the point clouds with some probabilistic model.

In Reference [14] the authors propose to jointly compute the correspondences and the nonrigid transformation parameters between two point clouds. The algorithm uses the expectation-maximization (EM) paradigm. It combines the soft-assign and deterministic annealing within a robust framework. Thin-plate splines (TPS) are used for representing the spatial mapping. Nonrigid alignment is proposed in Reference [23] to account for errors in the point clouds, obtained by scanning a rigid object. The authors use TPS to represent the nonrigid warp between a pair of views, that they estimate through hierarchical ICP.²² Medical applications are proposed in References [5,6]. In Reference [6], MR brain scan registration is performed by a modified Newton method over a hierarchical spline-based optical flow representation. In Reference [5], a localized radial basis function (RBF) is proposed, making a point to depend only on its neighboring centers.

Probabilistic approaches^{15–17} are based on modeling each of the point sets by a kernel density function.²⁴ The

(dis)similarity among such densities is computed by introducing appropriate distance functions. Registration is carried out without explicitly establishing correspondences. In Reference [15], the authors propose a correlation-based approach²⁴ to point set registration by representing the point sets as Gaussian mixture models (GMMs). A closed-form solution for the L_2 norm distance between two Gaussian mixtures makes fast computation possible. In Reference [17], registration is carried out simultaneously for several 3D range datasets. The method proposes an information-theoretic approach based on the Jensen–Shannon divergence measure. In Reference [16], nonrigid registration is treated as a maximum likelihood (ML) estimation problem by introducing the coherent point drift (CPD) paradigm. Smoothness constraints are introduced based on the assumption that points close to one another tend to move coherently over the velocity field. The proposed energy function is minimized with the EM algorithm.

The Proposed Approach

Our method draws on the strengths of some of the above mentioned approaches. It combines a deformable surface represented by a mesh with an ICP-like registration method that takes spatial and temporal smoothness into account, as well as the range image data and boundary information, required to prevent the computed surface to shrink or slide arbitrarily. The optimization is performed very efficiently using a distance transform of the range image.

Problem Statement

Surface Representation

The range sensor provides a sequence of range images that we interpret as 3D point clouds D_i with l_i points each:

$$D_i = \begin{pmatrix} d_{i,1}^x & d_{i,1}^y & d_{i,1}^z \\ \vdots & \vdots & \vdots \\ d_{i,l_i}^x & d_{i,l_i}^y & d_{i,l_i}^z \end{pmatrix}$$

The reconstructed surface at time i is represented by a *geometry image*.²⁵ The model M is organized as three $R \times C$ matrices, representing the deformation of a regular flat grid. Each matrix is reshaped in a single vector of size

$\mu = RC$, giving M_i as

$$M_i = \begin{pmatrix} m_{i,1}^x & m_{i,1}^y & m_{i,1}^z \\ \vdots & \vdots & \vdots \\ m_{i,\mu}^x & m_{i,\mu}^y & m_{i,\mu}^z \end{pmatrix}$$

In practice, the number of data points is much larger than the number of model points, that is, $l_i \gg \mu$. Upon convergence, our algorithm determines for each model point if there is a corresponding point in the current point cloud. Points may be missing because of occlusions or corrupted sensor output. This approach has the advantage that it naturally gives the reconstructed surface by interpolating the mesh points. Point cloud registration is obtained by composing the deformation fields.

Cost Function

Our cost function combines two data and three penalty terms:

$$e(M) = e_g(M) + \lambda_b e_b(M) + \lambda_s e_s(M) + \lambda_t e_t(M) + \lambda_x e_x(M) \tag{1}$$

where λ_b , λ_s , λ_x , and λ_t are weighting parameters. Note that we drop the frame index i for clarity purposes, and denote M_i as M and M_{i-1} as \tilde{M} .

The data terms are used to attract the estimated surface to the actual point cloud. The first term e_g is for global attraction, while the second one e_b deals with the boundary. These terms must account for possible erroneous points by using robust statistics. The penalty terms are e_s , e_t , and e_x . The two first ones respectively account for *spatial smoothness* and *temporal smoothness* e_s . The third one penalizes the *surface stress* and is related to the non-extensibility of the surface, and therefore to material properties of the surface.

This cost function is minimized in an ICP-like manner, as described in the next section. All the five terms are explained below in details.

Data Term: Global Surface Attraction. This term globally attracts the model to the data points in a closest point manner.²² In order to avoid the traditional two steps arising in ICP-like algorithms, we explicitly embed the min operator in this data term, as suggested in Reference [21]. Denoting E_M and E_D the sets of boundary points

in the model and in the data, we get the following data term, integrating the model to data points matching step:

$$\sum_{m \in M \setminus E_M} \min_{d \in D \setminus E_D} \|d - m\|^2 \quad (2)$$

where d and m are 3D vectors respectively representing a data and a model point. It is worth noting that, as opposed to Reference [21], our unknowns are not the rigid motion parameters (i.e., the classical roto-translations) but correspond to the whole *nonrigid motion field* in M .

An *outliers rejection* strategy is introduced by defining a robust function w . Following the X84 rule,²⁶ function w discards (i.e., it puts their residual to 0) those correspondences which residual error differs by more than 5.2 median absolute deviation (MAD) from the median. The value 5.2 corresponds to about 3.5 standard deviations, which includes more than 99.9% of a Gaussian distribution. Therefore, Equation (2) is modified so as to get the following robustified data term:

$$e_g(M) = \sum_{m \in M \setminus E_M} w \left(\min_{d \in D \setminus E_D} \|d - m\|^2 \right) \quad (3)$$

Data Term: Boundary Attraction. This term attracts boundary model points to boundary data points. It is defined in a similar manner to the global attraction term (3) except that the sum and min operators are over the boundary points:

$$e_b(M) = \sum_{m \in E_M} w \left(\min_{d \in E_D} \|d - m\|^2 \right) \quad (4)$$

Note that the detection of boundaries points in the range image is out of scope of the paper. In the ‘Experiments’ Section we describe some *ad hoc* boundary detection methods for the proposed applicative scenarios.

Penalty Term: Spatial Smoothness. This term discourages surface discontinuities by penalizing its second derivatives, as an approximation to its curvature. According to the definition of the geometry image,²⁵ the model M is a displacement field parameterized[†] by (u, v) with $u = [1 \cdots R]$ and $v = [1 \cdots C]$, that is, $M(u, v) =$

[†]Recall that the model points lie on a grid.

$[M^x(u, v), M^y(u, v), M^z(u, v)]$. The spatial smoothness term can thus be taken as the surface bending energy:

$$e_s(M) = \int_{\mathcal{R}} \int_{\mathcal{R}} \left(\frac{\partial M^2}{\partial^2 u} \right)^2 + 2 \left(\frac{\partial M^2}{\partial u \partial v} \right)^2 + \left(\frac{\partial M^2}{\partial^2 v} \right)^2 du dv$$

Using a finite difference approximation for the first and second derivatives,¹⁹ the bending energy can be expressed in discrete form as a quadratic function of M :

$$e_s(M) = \text{vect}(M)^\top \mathcal{K} \text{vect}(M) \quad (5)$$

where \mathcal{K} is a $3\mu \times 3\mu$, highly sparse matrix, and $\text{vect}(M)$ is the vectorization operator which rearranges matrix M to a vector.

Penalty Term: Temporal Smoothness. This defines a dependency between the current and the previous point clouds, M and \tilde{M} :

$$e_t(M) = \|M - \tilde{M}\|^2 \quad (6)$$

This makes the surface deformation smooth over time and can be used within a sequential processing approach. It is obviously not used on the first frame of the sequence.

Penalty Term: Non-Extensibility. This term discourages surface stretching. It favors the mesh vertices to preserve their distance with their local neighborhood:²⁷

$$e_x(M) = \sum_{m \in M} \sum_{k \in \mathcal{N}(m)} (\|m - k\|^2 - L_{m,k}^2)^2 \quad (7)$$

where $L_{m,k}$ are constants which are computed at the first frame after robust initialization and $\mathcal{N}(m)$ is the neighborhood of the mesh vertex m , with $\#\mathcal{N}(m) = 8$.

Minimization Procedure

The above described cost function (1) is a sum of squared residuals nonlinearly depending on the unknowns in M . In order to minimize this cost function, we use a nonlinear least squares optimization algorithm, namely the LM algorithm. We extend the LM-ICP approach proposed in Reference [21] to deformable objects. LM requires one to provide the partial derivatives of the residuals through a Jacobian matrix.

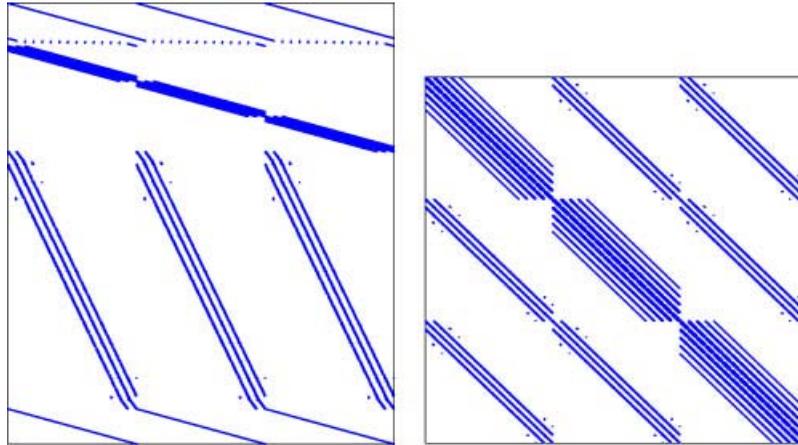


Figure 1. Example of Jacobian (left) and Hessian (right) matrix patterns: dark means non-zero. The size of these matrices is 3760×900 and 900×900 , so they respectively have 3 384 000 and 810 000 entries. They, however, have only 25 626 and 46 716 non-zero entries.

The Hessian matrix[‡] $H = J^T J + \lambda I$ must be inverted at each LM iteration, the problem is not tractable if the number of model points is too high (if the deformation field is too dense). The Jacobian matrix stacks the individual Jacobian matrices for all the five terms in the cost function:

$$J^T = \begin{pmatrix} J_d^T & J_b^T & J_s^T & J_t^T & J_x^T \end{pmatrix}$$

where $J_d^{\mu \times 3\mu}$, $J_b^{E_B \times 3\mu}$, $J_s^{3\mu \times 3\mu}$, $J_x^{\xi \times 3\mu}$, $J_t^{\mu \times 3\mu}$ are related to the global attraction, boundary attraction, spatial smoothness, temporal smoothness, and non-extensibility terms, respectively, and $\xi = \#\mathcal{N}(M)$. For instance, a grid with 15×20 points has a Jacobian matrix with 3760×900 elements ($\mu = 300$, $E_B = 66$, $\xi = 2194$). It is worth noting that the most expensive Jacobian computation is J_d since it requires the estimation of closest points at each iteration. As suggested in Reference [21], we solve this step with the *distance transform*: the distance between data points and a set of points lying on a discrete volumetric grid is pre-computed by defining a proper *distance field* data structure. In this fashion, the closest point computation is carried out in one step for all points.²¹ Therefore, the Jacobian matrix J_d is estimated by finite differences on such *distance fields* which remain constant through the minimization.

[‡]We use ‘Hessian matrix’ for the damped Gauss–Newton approximation to the true Hessian matrix.

One advantage of the proposed approach is that the Jacobian matrix J is very sparse. We thus use the sparsity to speed up each iteration using the technique in Reference [28]. In particular, we use the sparse Cholesky factorization package²⁹ included in the Matlab routines[§]. Figure 1 shows a plot of the Jacobian matrix $J^{3760 \times 900}$ (left), and the corresponding Hessian matrix H (right). Dark points show non-zero entries. The sparsity of the Jacobian and Hessian matrices is clearly evidenced.

Experiments

Two kinds of experiments have been set up. In the first one, a structured-light 3D scanner^{||} is used for scanning different deformations of a paper sheet. In the second experiment, the sensor is a passive-stereo system[¶] which allows us to acquire a sequence of 3D point clouds in realtime. The deformation of a portion of a blanket is modeled.

Initial conditions determine an estimate for both the model position, and the grid size. In practice, a correct starting grid allows LM to converge, as well as to determine the parameters $L_{m,k}$ in Equation (7). In the absence of any texture information, that is,

[§]More precisely, we use the ‘mldivide’ matlab function.

^{||}Data courtesy of Johnny Park from Purdue University (<http://web.ics.purdue.edu>).

[¶]Data courtesy of eVS (<http://www.evsys.net>).

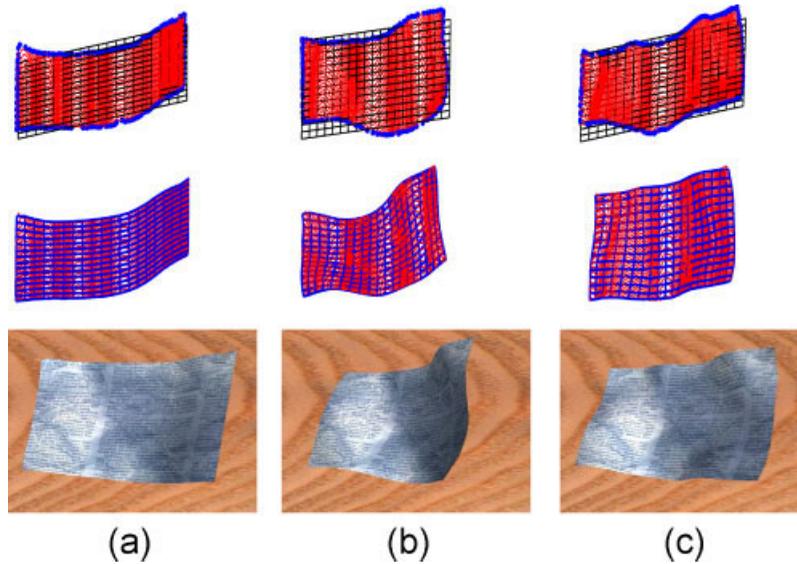


Figure 2. The paper sequence. Three examples of point clouds with the grid mesh superimposed at the starting position (top) and after model fitting (middle). Boundary points are evidenced. The reconstructed model is shown with a new texture (bottom).

in experiment 1, we detect them as strong depth discontinuities in the range image. In experiment 2, we use the intensity information to segment the boundary. The result is however quite noisy, but our system handles it thanks to its robustified data terms.

Experiment 1: Paper Sheet from a Structured-light Scanner

Several scans have been carried out while bending the paper. The sensor provides accurate and high-resolution 3D point clouds. The initial orientation of the grid is estimated by fitting a plane to the data. By projecting the points to the plane, both the grid size and boundaries are easily computed. There is no temporal dependency between the scans since they show totally different paper deformations. The temporal smoothness constraint e_t is thus inhibited.

Figure 2 shows three examples. Images on the top row visualize the model and data before surface fitting. Boundary points are highlighted. In the first example (Figure 2(a)), the deformation is mainly on the horizontal boundary. In the second one (Figure 2(b)), the paper is bent from the top-right to the bottom-left corner. In the third one (Figure 2(c)), the deformation is basically spread to the whole paper. Images on the central row show the result of our robust fitting. The

registration is accurate for both the interior points and the boundary. The recovered meshes are smooth as expected. Finally, three synthetic reconstructions are shown on the bottom row. Any texture can be overlaid on the model, for rendering realistic paper deformation from arbitrary points of view. Figure 3 shows the paper sheet reconstruction with different textures.

Finally, since the same grid is fit on all the images, the correspondences between different frames are recovered as a by-product (i.e., the corresponding points are those lying on the same grid position). In this manner, it is possible to synthetically interpolate the intermediate frames between two or more reference images. Figure 4 shows some intermediate frames between the actual scanned images shown in Figure 2(a), (b), and (c), respectively. Here, intermediate grids are recovered by simple linear interpolation between corresponding mesh vertices. More sophisticated interpolation schemes incorporating the non-extensibility penalty e_x could be used so as to make the surface behave similarly to a real paper sheet.

Experiment 2: Blanket from a Stereo System

A long sequence of point clouds is acquired for the second experiment. The sensor acquires the images at



Figure 3. Textured views. Once the surface deformation is captured, any image can be used as a texture.

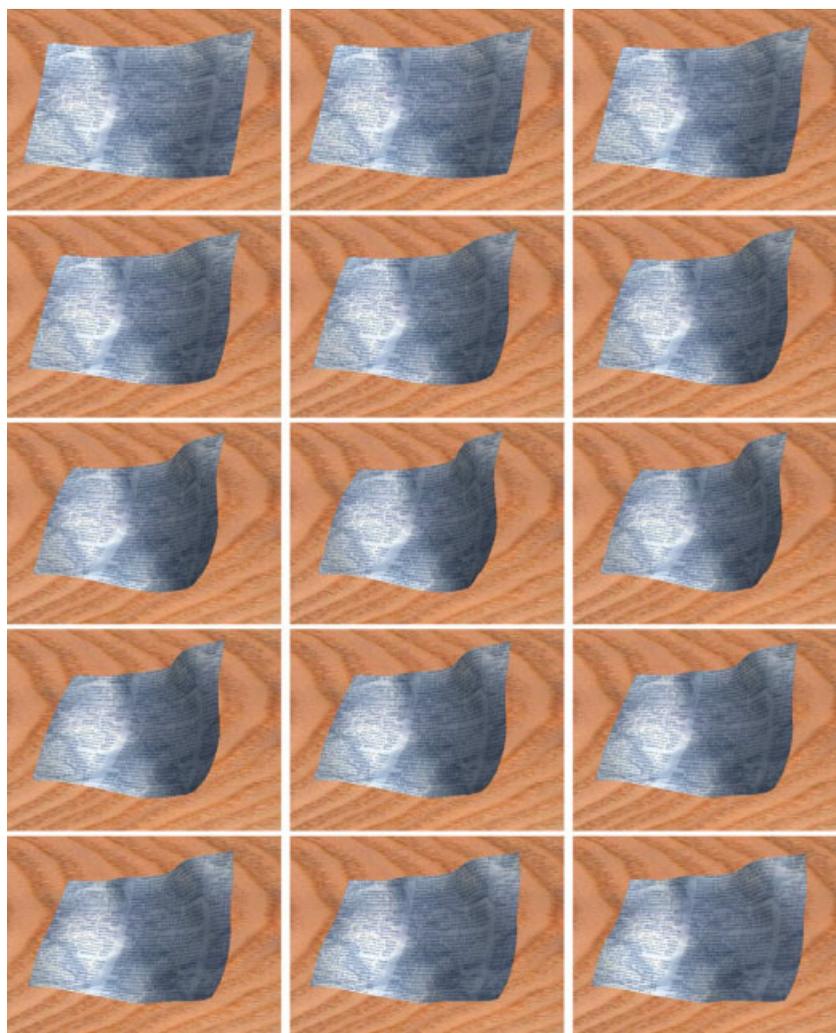


Figure 4. Synthetically interpolated views. Images are generated by linear interpolation between Figure 2(a), (b), and (c), respectively.

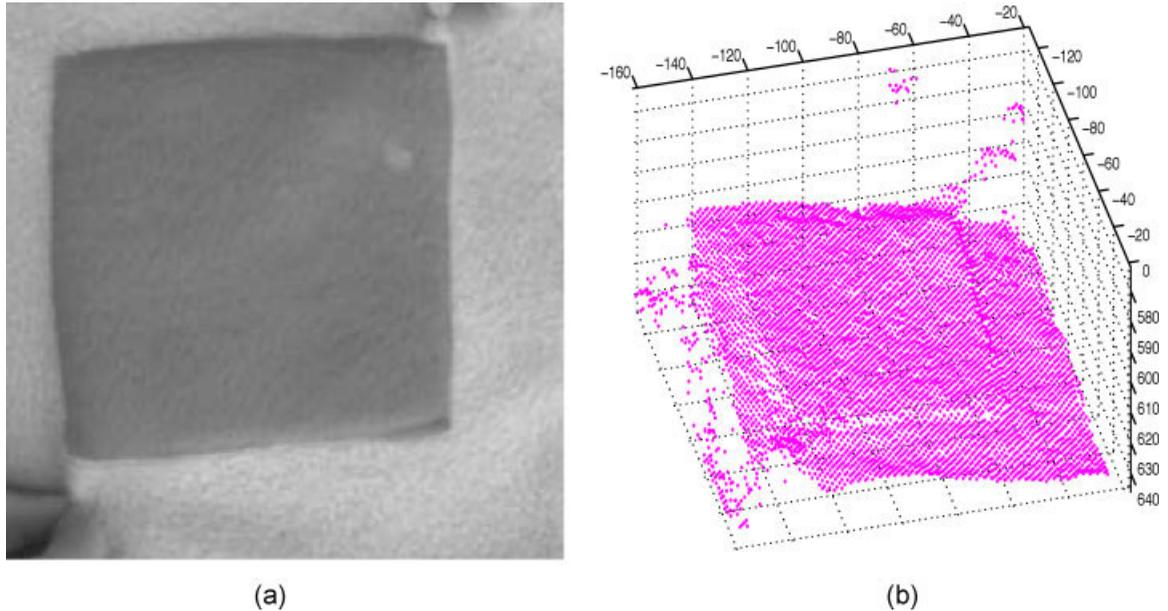


Figure 5. The blanket sequence: intensity image of the blanket (a) and the 3D point cloud (b).

25 frames-per-second (FPS), and provides both intensity (i.e., 2D) and 3D information. The quality of the 2D images is very low, and the range data are noisy. Moreover, the sensor can operate only on a very limited field of view (i.e., 30 cm³). We use a blanket as the object of interest. Figure 5(a) shows a picture of the blanket. Note that although 2D correspondences are recovered by the stereo system between the left and right images, the 2D correspondences between the frames contiguous in time are not reliable due to the distortions introduced by the surface deformation. In this experiment, we aim at observing the blanket deformation only on the portion delimited by the dark square. Figure 5(b) shows the 3D point cloud. There are many spurious points especially on the boundaries, and the scene is not easily recognizable. We use the intensity image** for selecting automatically our region of interest (i.e., the dark square), from which we recovered both the 3D data and the boundary. Figure 6(a) shows the image-boundary extracted by standard image processing techniques, while Figure 6(b) depicts the 3D data (i.e., the selected point cloud and 3D boundary). The sequence is made of 100 point clouds. Model initialization is

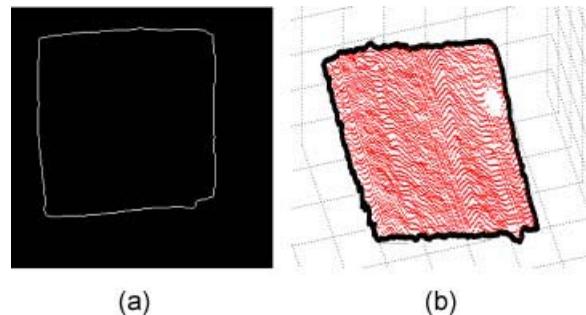


Figure 6. Data extraction: 2D boundary (a) and selected 3D data (b). 3D boundary is highlighted with dark color.

carried out for the first cloud only. Each iteration uses the output of the previous one as an initial condition. Figure 7 shows a selection of the output sequence. For each frame, is visualized: (1) the intensity image, with the extracted 2D boundary and the 2D projection of the estimated model and (2) the point cloud—after the region of interest selection, evidencing both the 3D boundary and the grid. The blanket is handled from the bottom-left and upper-right corners, respectively. On the early frames, the blanket is gradually bent toward the square center, then it is strongly stretched, moving the corners far from each other. Finally, in the late frames, random deformations are generated especially around the corners. Some frames are particularly challenging. In

**The intensity is the left image of the stereo-pair, which is associated to the disparity map. Indeed, there is a mapping between the 2D and 3D information. Note that we do not use intensity information for fitting.

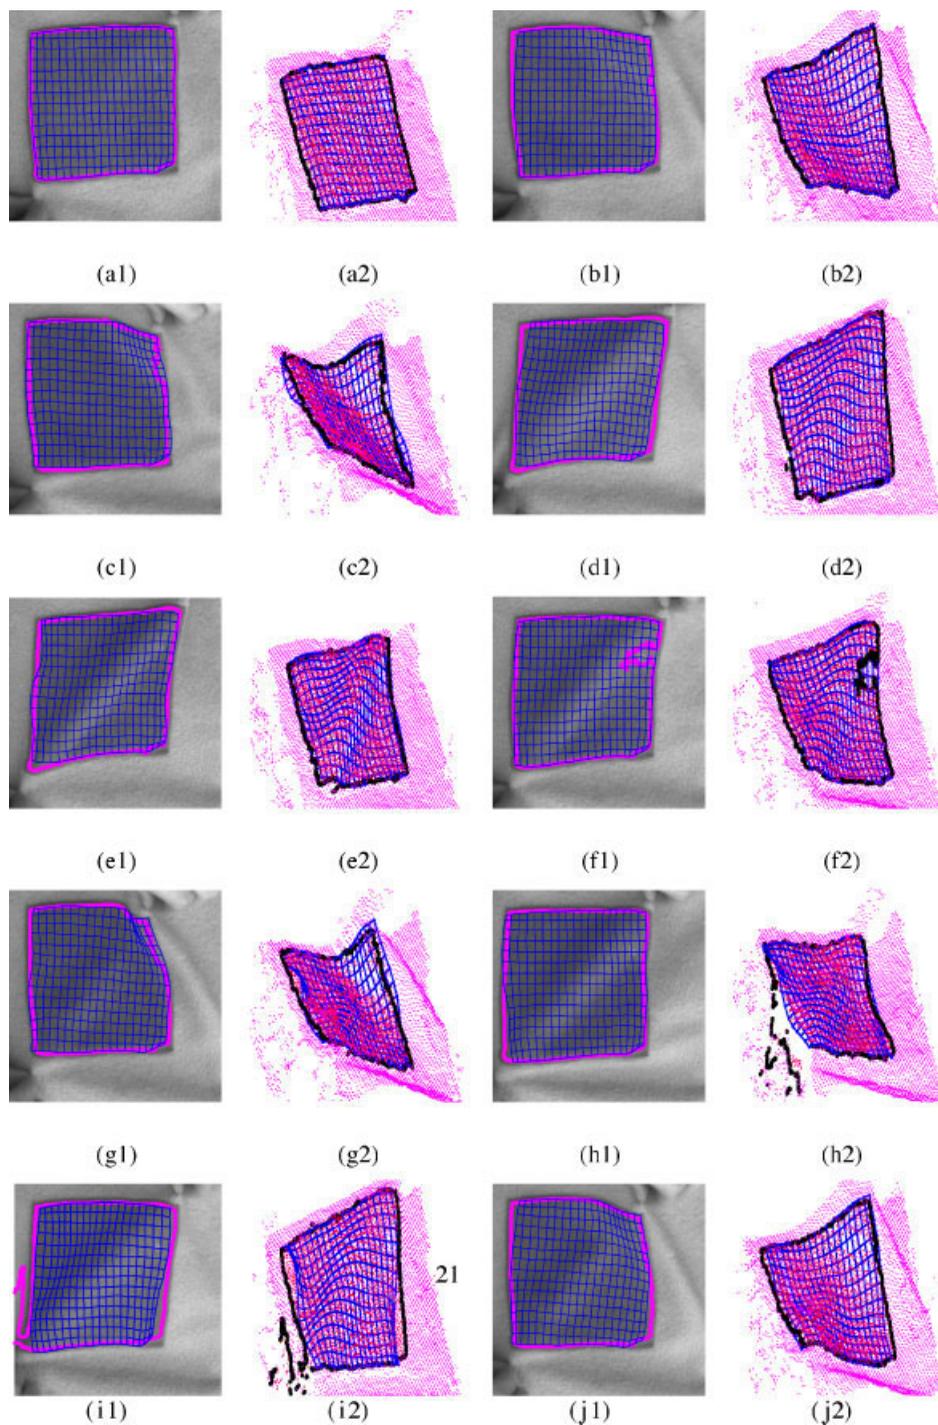


Figure 7. Blanket sequence: 10 selected frames. For each frame the 2D intensity (·,1) and the 3D data (·,2) is visualized. The grid models are shown in the 3D space as well as their projection in the 2D image.



Figure 8. Synthesized movie: some selected frames. Each frame of the movie is projected to the reconstructed model by simulating a deforming video screen.

frame (c) a strong shrinking is evidenced on the top-right corner. In frame (f) a wide hole appears on the top-right side. In frames (h) and (i) data boundaries are clearly wrong on the bottom-left side. Results are satisfying since the fitting is correct for the whole sequence. The mesh grids are well superimposed on data points maintaining a smooth shape. Nevertheless, the projections of the grids to the 2D images confirm the accuracy of the registration. Finally, after the mesh fitting, a dense set of accurate deformations is available. We used them to synthesize

a movie rendered as if it was projected to a deforming screen. We project every frame of the movie to a model of our sequence^{††}. Some frames are shown in Figure 8.

Performance Evaluation. For both experiments, a model of size 15×20 is used. We have seen that a higher value of λ_b is necessary (i.e., $\lambda_b = 1.5$) for a correct

^{††}We loop over the extracted 100 models.

Experiment	Pre-process (s)	Energy (s)	# LM iter	Time per frame (second)	# Frames	Total time (second)
Exp. 1	19.36	2.22	11	24.42	1	43.78
Exp. 2	15.24	2.34	7	16.34	100	1649.34

Table 1. Performances of the main steps of the proposed framework. Running times (second) and iterations are mean values

convergence of the algorithm to the optimal solution. The other terms are set almost equally to 1. The distance transform parameters are important: the size of the voxels trades off speed and result accuracy. Here, we have divided the volume into $36 \times 36 \times 18$ voxels.

Table 1 shows the running times of both the proposed experiments. For each experiment we highlight the time spent for pre-processing (i.e., grid pre-alignment and the computation of constants $L_{m,k}$), the total energy computation, the mean number of LM iterations, the total time for each frame and the total time of the whole experiment. Note that the computational cost of one LM iteration is very similar for both experiments. This is due to the fact that it depends only on the chosen grid size, being independent on the amount of input data. The computational cost of pre-processing instead depends of the total number of points and thus the first experiment takes a longer time.

The method has been implemented in Matlab on a Pentium M 1.86 GHz. Note that since the pre-processing is carried out only for the first frame of the sequence it strongly affects the total time of the first experiment (for which only single frames are processed) while being less relevant for the second experiment.

Discussion

We proposed a new approach for capturing the deformation of 3D surfaces from range image or 3D scans. As a surface model, we deform a generic geometry image, which is aligned with the observed data points. A cost function is devised that combines *a priori* information, such as spatio-temporal smoothness, and observations. Both non-extensibility and boundary attraction terms are crucial for disambiguating this intrinsically ill-posed problem. The optimization phase is solved with the LM algorithm, while taking advantage of the sparsity of the Jacobian and Hessian matrices.

Results are promising since the performances are satisfying for the analyzed cases. The method has been tested onto two kinds of datasets, thereby evidencing

the versatility in dealing with different sensors. In the first experiment, the source data were accurate and the estimated models was according to what we expected. In the second experiment, a whole sequence of 3D point clouds has been processed. This has allowed us to capture cloth deformations without using any markers nor special projected patterns. Although data were very noisy, especially around the boundary, the method performed robustly. We discussed the behavior of our algorithm in the presence of holes and broken boundary. Moreover, since the captured deformation depends only on the size of the chosen grid, the proposed method allows one to use an arbitrary level of deformable details in the captured surfaces. We also highlight that the method is easy to implement and that it is reasonably fast. Finally, some graphical results were shown for generating new synthetic deformations, and for changing the appearance of the original surface, as well as for synthesizing a deforming video. This shows the usefulness of the method for computer animation. Several applications are made possible by the proposed method, including (i) cloth motion capture, with no marker or pattern, (ii) dense augmented reality for special effects such as a liquid spreading on a deformable surface, and (iii) surface completion and interpolation in the presence of missing or erroneous frames. The main limitation regards the hypothesis to work on planar-like surfaces. The extension to shapes with more complex topology will be addressed in future work.

References

1. Koninckx TP, Griesser A, Van Gool L. Real-time range scanning of deformable surfaces by adaptively coded structured light. In *Proceedings of the International Conference on 3D Digital Imaging and Modeling*, 2003; 293–300.
2. May S, Werner B, Surmann H, Pervolz K. 3d time-of-flight cameras for mobile robotics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006; 790–795.

3. Rusinkiewicz S, Holt OH, Levoy M. Real-time 3D model acquisition. *ACM Transactions on Graphics (Proc. SIGGRAPH) 2002*; **21(3)**: 438–446.
4. Alexa M. Recent advances in mesh morphing. *Computer Graphics Forum 2002*; **21(2)**: 173–196.
5. Fornefett M, Rohr K, Stiehl HS. Radial basis functions with compact support for elastic registration of medical images. *Image and Vision Computing 2001*; **19**: 87–96.
6. Vemuri BC, Huang S, Sahni S, et al. An efficient motion estimator with application to medical image registration. *Medical Image Analysis 1998*; **2(1)**: 79–98.
7. Kass M, Witkin A, Terzopoulos D. Snakes: active contour models. *International Journal of Computer Vision 1987*; **1(4)**: 321–331.
8. Blanz V, Vetter T. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques, 1999*; 187–194.
9. Guskov I, Klibanov S, Bryant B. Trackable surfaces. In *Proceedings of the Eurographics/SIGGRAPH Symposium on Computer Animation, 2003*; 251–257.
10. Pritchard D, Heidrich W. Cloth motion capture. *Computer Graphics Forum 2003*; **22(3)**: 263–272.
11. Schols V, Magnor MA. Cloth motion from optical flow. In *Vision, Modelling and Visualization, 2004*; 117–124.
12. Scholz V, Stich T, Keckeisen M, Wacker M, Magnor M. Garment motion capture using color-coded patterns. *Computer Graphics Forum 2005*; **24(3)**: 439–448.
13. Yamamoto M, Boulanger P, Beraldin JA, Rioux M, Domey J. Direct estimation of deformable motion parameters from range image sequence. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 1990*; 460–464.
14. Chui H, Rangarajan A. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding 2003*; **89(2–3)**: 114–141.
15. Jian B, Vemuri BC. A robust algorithm for point set registration using mixture of Gaussians. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2005*; 1246–1251.
16. Myronenko A, Song X, Carreira-Perpinan M. Non-rigid point set registration: coherent point drift. In *Proceedings of the Neural Information Processing Systems, 2007*; 1009–1016.
17. Wang F, Vemuri BC, Rangarajan A. Groupwise point pattern registration using a novel CDF-based Jensen-Shannon Divergence. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2006*; 1283–1288.
18. Pilet J, Lepetit V, Fua P. Real-time non-rigid surface detection. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2005*; 822–828.
19. Prasad M, Zisserman A, Fitzgibbon AW. Single view reconstruction of curved surfaces. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2006*; 1345–1354.
20. Castellani U, Gay-Bellile V, Bartoli A. Joint reconstruction and registration of a deformable planar surface observed by a 3d sensor. In *Proceedings of the International Conference on 3D Digital Imaging and Modeling, 2007*; 1138–1139.
21. Fitzgibbon AW. Robust registration of 2D and 3D point sets. *Image Vision Computing 2003*; **21(13–14)**: 1145–1153.
22. Rusinkiewicz S, Levoy M. Efficient variants of the ICP algorithm. In *Proceedings of the International Conference on 3D Digital Imaging and Modeling, 2001*; 145–152.
23. Brown BJ, Rusinkiewicz S. Non-rigid range-scan alignment using thin-plate splines. In *Proceedings of the International Symposium on 3D Data Processing, Visualization and Transmission, 2004*; 822–828.
24. Tsin Y, Kanade T. A correlation-based approach to robust point set registration. In *Proceedings of the European Conference on Computer Vision, 2004*; 558–569.
25. Gu X, Gortler SJ, Hoppe H. Geometry images. *ACM Transaction on Graphics 2002*; **21(3)**: 355–361.
26. Castellani U, Fusiello A, Murino V. Registration of multiple acoustic range views for underwater scene reconstruction. *Computer Vision and Image Understanding 2002*; **87(3)**: 78–89.
27. Salzmann M, Ilic S, Fua P. Physically valid shape parameterization for monocular 3-D deformable surface tracking. In *British Machine Vision Conference, 2005*; 1138–1139.
28. Pissanetzky S. *Sparse Matrix Technology*. Academic Press: London, 1984.
29. <http://www.cise.ufl.edu/research/sparse/cholmod>

Authors' biographies:



Umberto Castellani is *Ricercatore* (i.e., Research Assistant) of Department of Computer Science at University of Verona. He received his Dottorato di Ricerca (PhD) in Computer Science from the University of Verona in 2003 working on 3D data modelling and reconstruction. During his PhD, he had been Visiting Research Fellow at the Machine Vision Unit of the Edinburgh University, in 2001. In 2007 he has been an Invited Professor for two months at the LASMEA laboratory in Clermont-Ferrand, France. In 2008, he has been Visiting Researcher for two months at the PRIP laboratory of the Michigan State University (USA). His main research interests concern the processing of 3D data coming from different acquisition systems such as 3D models from 3D scanners, acoustic images for the vision in underwater environment, and MRI scans for biomedical applications. The addressed methodologies are focused on the intersections among Machine Learning, Computer Vision and Computer Graphics.



Vincent Gay-Bellile received the MS degree in computer science from the University Blaise Pascal, Clermont-Ferrand, France, in 2005. He is currently a PhD candidate at Blaise Pascal University and a member of CEA (Commissariat à l'Energie Atomique), Saclay. He is working in partnership with the LASMEA and LSVE laboratory under the supervision of Professors Jean-Thierry Lapresté, and Doctors Adrien Bartoli and Patrick Sayd. His main research interests are in image registration and structure from motion for nonrigid scenarios in computer vision. He received the Best Student Paper Award at CORESA '07, together with Adrien Bartoli and Patrick Sayd. In December 2008, he will be a permanent researcher at CEA.



Adrien Bartoli received the PhD degree from INPG/INRIA, Grenoble, France, where he did his research with the Perception Group (ex MOVI) under the supervision of Professors Peter Sturm and Radu Horaud. He was a postdoctoral researcher at the University of Oxford, United Kingdom, in the Visual Geometry Group, under the supervision of Professor Andrew Zisserman. He has been a permanent CNRS research scientist (Charge de Recherche) at the LASMEA Laboratory, Clermont-Ferrand, France, since October 2004 and a visiting professor at DIKU in Copenhagen, Denmark, from 2006 to 2009. His main research interests are in image registration and structure from motion for rigid and nonrigid scenarios in computer vision. He is the author or coauthor of more than 75 scientific papers. He received the 2004 INPG PhD thesis prize and the Best Student Paper Award at CORESA '07, together with his PhD student Vincent Gay-Bellile and colleague Patrick Sayd. Since September 2006, he has been a co-head of the ComSee research team at LASMEA.