Non-rigid Face Modelling Using Shape Priors

Alessio Del Bue, Xavier Lladó, and Lourdes Agapito

Queen Mary University of London, London E1 4NS, UK
{alessio,llado,lourdes}@dcs.qmul.ac.uk

Abstract. Non-rigid 3D shape recovery is an inherently ambiguous problem. Given a specific rigid motion, different non-rigid shapes can be found that fit the measurements. To solve this ambiguity prior knowledge on the shape and motion should be used to constrain the solution. This paper is based on the observation that often not all the points on a moving and deforming surface such as a human face are undergoing non-rigid motion. Some of the points are frequently on rigid parts of the structure – for instance the nose – while others lie on deformable areas. First we develop a segmentation algorithm to separate rigid and non-rigid motion. Once this segmentation is available, the rigid points can be used to estimate the overall rigid motion and to constrain the underlying mean shape. We propose two reconstruction algorithms and show that improved reconstructions can be obtained when the priors on the shape are used on synthetic and real data.

1 Introduction

In this paper we focus on the estimation of the 3D shape and motion of a deformable object such as a human face which is moving rigidly while performing different facial expressions. The face can be thought of as an underlying rigid body undergoing a global rotation and translation while suffering some local non-rigid deformations. Our aim is the simultaneous recovery of motion and 3D non-rigid shape from multiple images exploiting prior knowledge on the structure such as the rigidity of some of the observed points.

In the past years numerous techniques have been proposed to solve the *structure from motion* problem in the case of rigid objects and more recently the framework has also been extended to deal with non-rigid objects. The main challenge in non-rigid structure from motion is to disambiguate the contribution to the image motion given by the shape deformation and that caused by the rigid motion. Bregler et al [3] introduced a representation for non-rigid 3D shape where any configuration can be expressed as a linear combination of basis shapes that define the principal modes of deformation of the object. They proposed a factorization method that exploits the rank constraint on the measurement matrix and enforces orthonormality constraints on camera rotations to recover the motion and the non-rigid 3D shape. Their work can be seen as an extension of Tomasi and Kanade's factorization framework [12] to the case of deformable objects. Torresani et al. [13] extended the method of Bregler et al. to a trilinear optimization problem by minimizing 2D image reprojection error using Alternating Least Squares. Brand [2] proposed an alternative optimization method and added an extra constraint on the basis shapes: the deformations should be as small as possible relative to the rigid shape.

W. Zhao, S. Gong, and X. Tang (Eds.): AMFG 2005, LNCS 3723, pp. 97–108, 2005.

The main problem with these approaches stems from the fact that deformation and motion are ambiguous. Given a specific configuration of points on the image plane, different 3D non-rigid shapes and camera motions can be found that fit the measurements. To solve this ambiguity prior knowledge on the shape and motion should be used to constrain the solution. Recently, Xiao et al. [17] proved that the orthogonality constraints were insufficient to disambiguate rigid motion and deformations. They identified a new set of constraints on the shape bases which, when used in addition to the rotation constraints, provide a closed form solution to the problem of non-rigid structure from motion. However, their solution requires that there be K frames (where K is the number of basis shapes) in which the shapes are known to be independent. Non-linear optimization schemes that minimize image reprojection error have also been proposed to refine an initial solution [1,4]. The advantage of these methods is that they provide a maximum likelihood estimate in the presence of Gaussian noise and prior knowledge on any of the model parameters can be easily incorporated to the cost function in the form of penalty terms. The need for incorporating prior information on the motion or shape parameters to avoid the ambiguities inherent to non-rigid shape estimation is also recognised by Torresani et.al. [15,14] who propose an algorithm that learns the time-varying shape of a non-rigid 3D object from uncalibrated 2D tracking data. Temporal smoothness in the object shape can be imposed within their framework which can also deal with missing data.

In this paper we focus on the observation that often not all the points on a moving and deforming surface – such as a human face – are undergoing non-rigid motion. Some of the points are frequently on rigid parts of the structure – for instance the nose – while others lie on deformable areas. Intuitively, if a segmentation is available, the rigid points can be used to estimate the overall rigid motion and to constrain the underlying mean shape by estimating the local deformations exclusively with the parameters associated to the non-rigid component of the 3D model. Our observation is also supported by recent studies on the notion of shape average by Yezzi and Soatto introduced in [18] where the authors precisely separate motion and deformation components for robustly matching, registering and tracking deformable objects.

Our approach first performs rigid and non-rigid motion segmentation on the fully observed image data to separate both types of motion using an automatic measure of deformability of shapes [10]. Once the points have been segmented into the rigid and non-rigid sets we recover the overall rigid motion from the rigid set and we formalise the problem of non-rigid shape estimation as a constrained minimization adding priors on the degree of deformability of each point. We perform experiments on synthetic and real data which validate the approach and show that the addition of priors on the rigidity of some of the points improves the 3D reconstruction.

The paper is organised as follows. In section 2 we describe the non-rigid factorization framework. In section 3 we propose a rigid and non-rigid motion segmentation algorithm. Section 4 presents two alternative algorithms to recover the 3D shape using rigidity constraints on the non-deforming segmented points. Finally, in section 5 we show experiments on synthetic and real data to validate the segmentation algorithm and the 3D reconstruction methods.

2 Non-rigid 3D Modelling Using Factorization

Tomasi and Kanade's factorization algorithm has recently been extended to the case of non-rigid 3D structure, assuming affine viewing conditions [3,2,13,4]. The model used to express the deformations is point-wise and the 3D shape of any specific configuration $\mathtt{S}$ is approximated by a linear combination of a set of K basis shapes $\mathtt{S}_k$ which represent the principal modes of deformation of the object:

$$\mathtt{S} = \sum_{k=1}^{K} l_k \mathtt{S}_k \qquad \mathtt{S}, \mathtt{S}_k \in \Re^{3\times P} \quad l_k \in \Re \tag{1}$$

where each basis shape $\mathtt{S}_k$ is a $3 \times P$ matrix which contains the 3D locations of P object points for that particular mode of deformation. Assuming an orthographic camera model the shape is then projected onto an image frame i giving P image points:

$$\left[\, \mathbf{x}_{i1} \; \ldots \; \mathbf{x}_{iP} \,\right] = \mathtt{R}_i \left(\sum_{k=1}^{P} l_{ik} \mathtt{S}_k \right) \tag{2}$$

where each $\mathbf{x}_{ij} = [u_{ij} v_{ij}]^T$ contains the horizontal and vertical image coordinates of the point – referred to the centroid of the object – and $\mathtt{R}_i$ encodes the first two rows of the rotation matrix for a specific frame i. If all P points are tracked in F image frames we may construct the measurement matrix $\mathtt{W}$ which can be expressed as:

$$\mathtt{W} = \begin{bmatrix} \mathbf{x}_{11} & \ldots & \mathbf{x}_{1P} \\ \vdots & & \vdots \\ \mathbf{x}_{F1} & \ldots & \mathbf{x}_{FP} \end{bmatrix} = \begin{bmatrix} l_{11}\mathtt{R}_1 & \ldots & l_{1K}\mathtt{R}_1 \\ \vdots & & \vdots \\ l_{F1}\mathtt{R}_F & \ldots & l_{FK}\mathtt{R}_F \end{bmatrix} \begin{bmatrix} \mathtt{S}_1 \\ \vdots \\ \mathtt{S}_K \end{bmatrix} \tag{3}$$

Clearly, the rank of the measurement matrix is constrained to be at most $3K$, where K is the number of deformations. This rank constraint can be exploited to factorize the measurement matrix into a motion matrix $\hat{\mathtt{M}}$ and a shape matrix $\hat{\mathtt{S}}$ by truncating the SVD of $\mathtt{W}$ to rank $3K$. However, this factorization is not unique since any invertible $3K \times 3K$ matrix $\mathtt{Q}$ can be inserted in the decomposition leading to the alternative factorization: $\mathtt{W} = (\hat{\mathtt{M}}\mathtt{Q})(\mathtt{Q}^{-1}\hat{\mathtt{S}})$. The problem is to find a transformation matrix $\mathtt{Q}$ that renders the appropriate replicated block structure of the motion matrix shown in Equation (3) and that removes the affine ambiguity, upgrading the reconstruction to a metric one.

In this paper we address the problem of non-rigid shape estimation and we propose a new solution which incorporates information on the degree of deformability of the 3D points. First we identify the points whose motion can be explained purely by a rigid transformation. This knowledge provides some constraints or priors on the values of the 3D shape which will allow to solve the inherent problem of ambiguity present in the motion and 3D shape in non-rigid factorization.

3 Automatic Rigid and Non-rigid Motion Segmentation

As a first step to 3D structure recovery we propose to separate points in the sequence that exhibit a purely rigid motion from those which are also suffering some non-rigid

deformations. To do this we apply a subset selection technique on the non-rigid component of the point trajectories encoded in the measurement matrix $\mathtt{W}$. Subset selection is a technique commonly used in feature selection problems where a group of features is extracted to obtain a robust solution to a particular estimation problem [7].

Under the factorization framework, features are represented by their image point trajectories stored in $\mathtt{W}$. Our goal is to find the set of features whose motion can be modelled exactly as a rigid motion. In this case we formulate the segmentation problem as finding a subset of trajectories $\mathtt{W}_{rigid}$ within the measurement matrix such that the following condition is satisfied:

$$rank(\mathtt{W}_{rigid}) = 3. \tag{4}$$

The segmentation algorithm follows a *sequential backward selection strategy* [9] by initially considering all the trajectories in the measurement matrix and iteratively deleting one by one those which are contributing most to the rank of the matrix, i.e. the points that exhibit the most non-rigid motion. As the stop criterion for the classification task, we compute the rank of the measurement matrix of the remaining points which will become 3 when only the rigid trajectories are left.

Obviously the rank of the rigid points will not be exactly equal to 3 in the presence of noise. Instead, we have used an automatic method to determine the deformability index of a set of trajectories [10]. This method estimates the value of K – the number of independent basis shapes needed to describe the non-rigid motion – automatically in a non-iterative way. It provides a fixed threshold for comparing the eigenvalues of the matrix to determine the rank. For the case of a 3D rigid body the deformability index K is equal to 1 while in the case of a non-rigid body the index is $3K$ therefore this provides a good selection criterion to separate both sets of trajectories. The complete algorithm is detailed below:

- Initialize $\mathtt{W}_{rigid} = \mathtt{W}$
- Determine the deformability index K for $\mathtt{W}_{rigid}$

1. Compute $\mathtt{W}_{rigid} \simeq \mathtt{U}\mathtt{D}\mathtt{V}^T$ with SVD.
2. Define $\mathtt{S} = \mathtt{D}^{1/2}\mathtt{V}^T$
3. Extract the non-rigid component of the shape matrix $\tilde{\mathtt{S}}_{3(\mathtt{K}-1)\times\mathtt{P}} = \left[\tilde{\mathbf{S}}_1 \ ... \ \tilde{\mathbf{S}}_P\right]$ where each $\tilde{\mathbf{S}}_j$ is a $3(K-1)\times 1$ vector which contains the 3D coordinates of the j^{th} 3D point associated to the $K-1$ non-rigid bases.
4. Determine the maximum norm vector: $\tilde{\mathbf{S}}_t = \max\{\parallel \tilde{\mathbf{S}}_1 \parallel, \ ... \ , \parallel \tilde{\mathbf{S}}_P \parallel\}$.
5. Remove the selected trajectory t from $\mathtt{W}_{rigid}$ and determine the new deformability index K.
6. If $K = 1$ stop the iteration.
7. Else, go to step 1.

We have obtained successful rigid and non-rigid motion segmentations on synthetic sequences using this algorithm. The results will be discussed in the experimental section. Note that the method converges to the right solution only if there is a unique set of rigid points such that $K = 1$. In the case where different groups of features satisfy the rank condition the algorithm could converge to the wrong set.

4 Non-rigid Shape and Motion Estimation Using Rigidity Constraints

Once we have segmented the scene into rigid and non-rigid points, we can use the information on the rigidity of the points to constrain the shape estimation. First we define the constraints that arise based on the observation that a generic shape is composed by points with different degrees of deformation. Kim and Hong [8] defined the *degree of non-rigidity* of a point as its degree of deviation from the average shape to classify points into three classes: rigid, near-rigid and non-rigid. Based on this measure they proposed a method to estimate average shape using the degree of non-rigidity to weight the contribution of each point in an iterative certainty re-weighted factorization scheme. In contrast, we use the knowledge that some points of the scene are rigid to construct specific linear constraints which will in turn eliminate the inherent ambiguities present in non-rigid shape estimation.

4.1 Rigidity Constraint

Definition (rigid point). *If the motion of a point p is completely rigid for the entire sequence, the structure referring to the point can be expressed entirely by the first basis ($K = 1$) called the rigid basis.*

It follows from this definition that a completely rigid point p is entirely parameterized by:

$$\mathbf{S}_p = \begin{bmatrix} \mathbf{S}_{p1} \\ \mathbf{0} \end{bmatrix} \tag{5}$$

where $\mathbf{S}_{p1}$ is a 3-vector which contains 3D coordinates of the rigid component and $\mathbf{0}$ is a $3(K-1)$-vector of zeros. It is possible to reorder the measurement matrix after the detection of all the rigid points by defining the permutation matrix $\mathtt{P}$ such that:

$$\mathtt{WP} = \left[\mathtt{W}_{rigid} \middle| \mathtt{W}_{nonrigid} \right] = \begin{bmatrix} l_{11}\mathtt{R}_1 & \ldots & l_{1K}\mathtt{R}_1 \\ \vdots & & \vdots \\ l_{F1}\mathtt{R}_F & \ldots & l_{FK}\mathtt{R}_F \end{bmatrix} \left[\begin{array}{c|c} \mathtt{S}_{rigid} & \mathtt{S}_{nonrigid} \\ \mathtt{0} & \end{array} \right] \tag{6}$$

where $\mathtt{S}_{rigid}$ is a $(3 \times r)$ matrix containing the 3D coordinates of the r rigid points, $\mathtt{S}_{nonrigid}$ is a $(3K \times d)$ matrix containing the 3D coordinates of the K basis shapes for the d deformable points and $\mathtt{0}$ is a $3(K-1) \times r$ matrix of zeros. Notice that it is now possible to apply Tomasi and Kanade's rigid factorization on the measurement matrix containing the image trajectories of the rigid points $\mathtt{W}_{rigid}$ and decompose it into the motion and rigid structure components as:

$$\mathtt{W}_{rigid} = \begin{bmatrix} \mathtt{R}_1 \\ \vdots \\ \mathtt{R}_f \end{bmatrix} \mathtt{S}_{rigid} \tag{7}$$

4.2 Non-rigid Shape and Motion Estimation

In this section we solve for the non-rigid shape and motion given the 2D image tracks and incorporating the above constraint on the automatically segmented rigid points. Our approach is to minimize image reprojection error subject to the rigidity of the non-deforming points. The cost function being minimised is:

$$\chi = \sum_{i,j} \| \mathbf{x}_{ij} - \hat{\mathbf{x}}_{ij} \|^2 = \sum_{i,j} \| \mathbf{x}_{ij} - (\mathtt{R}_i \sum_k l_{ik} \mathbf{S}_k) \|^2 \tag{8}$$

where $\mathbf{x}_{ij}$ are the measured image points and $\hat{\mathbf{x}}_{ij}$ the estimated image points. We propose two alternative solutions to this constrained minimization: a linear alternate least squares approach which incorporates the rigidity constraints using Generalised Singular Value Decomposition and a fully non-linear minimization scheme using priors on the rigid shape parameters in a Maximum A Posteriori estimation.

Linear Equality-Constrained Least Squares Using GSVD. First we propose an alternating least squares scheme to minimize the cost function described in equation (8). The algorithm alternates between solving for the basis shapes $\mathtt{S}$ and for the configuration weights l_{ik}. The configuration weights are initialised to random values. The scheme can be summarised as follows:

1. Given $\mathtt{R}_i$ and l_{ik} equation (3) can be used to estimate $\mathtt{S}$ linearly subject to the constraint $\tilde{\mathbf{S}}_p = \mathbf{0}$ for $p \in \Omega$ with Ω being the set of r points considered to be rigid throughout the sequence.
2. Given $\mathtt{R}_i$ and $\mathtt{S}$ solve for all l_{ik} using linear least-squares.
3. Iterate the above two steps until convergence.

Note that the algorithm does not solve for the overall rigid motion encoded in the rotation matrices $\mathtt{R}$ since these are calculated before hand by running the rigid factorization algorithm of Tomasi and Kanade on the rigid points. Rearranging equation (3) the problem of solving for $\mathtt{S}$ subject to the rigidity constraint can be expressed as an unconstrained least squares system of the form:

$$\min \| \begin{bmatrix} \mathtt{A} \\ \lambda \mathtt{C} \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{b} \\ \lambda \mathbf{d} \end{bmatrix} \|_2 \tag{9}$$

where $\mathtt{A}$ represents the linear equations, $\mathtt{C}$ the linear constraints and $\mathbf{b}$ and $\mathbf{d}$ are the known observations. It is shown [5] that for $\lambda \to \infty$ the final solution lies on the surface defined by $\mathtt{C}\mathbf{x} = \mathbf{d}$ and thus we obtain a linear equality-constrained least squares (LSE) problem:

$$\min \| \mathtt{A}\mathbf{x} - \mathbf{b} \|_2 \tag{10}$$

subject to:

$$\mathtt{C}\mathbf{x} = \mathbf{d} \tag{11}$$

A method to solve the above LSE problem is directly to factorize both $\mathtt{A}$ and $\mathtt{C}$ using Generalized Singular Value Decomposition (GSVD) (see [6] for details).

Bundle Adjustment Using Priors. The correct approach to non-rigid factorization is to formulate the problem as a non-linear least square estimation minimizing the distance of the reprojection error in the model parameters:

$$\arg\min_{\mathtt{R}_i\mathbf{S}_k l_{ik}} \sum_{i,j} \| \mathbf{x}_{ij} - \hat{\mathbf{x}}_{ij} \|^2 = \arg\min_{\mathtt{R}_i\mathbf{S}_k l_{ik}} \sum_{i,j} \| \mathbf{x}_{ij} - (\mathtt{R}_i \sum_k l_{ik}\mathbf{S}_k) \|^2 \tag{12}$$

where $\mathbf{x}_{ij}$ are the measured image points and $\hat{\mathbf{x}}_{ij}$ the estimated image points.

This method has the advantage of providing a true maximum likelihood estimate, provided the noise distribution is Gaussian. Besides, prior knowledge on any of the model parameters can be easily incorporated to the error cost function in the form of penalty terms. However, it suffers from the fact that it requires an initialization that is close to the global minimum. Therefore these methods are generally used as a final refinement step.

One of the main advantages of performing a prior segmentation of rigid and non-rigid motion is firstly that the rigid motion (estimates of the rotation matrices $\mathtt{R}$) can be pre-computed by performing rigid factorization on the rigid points. This provides a very good initial estimate for the rotation parameters, which coupled with the priors on the 3D shape help solve the ambiguities.

Our prior expectation is that the points detected as being rigid have a zero non-rigid component and can therefore be modelled entirely by the first basis shape:

$$\mathbf{S}_p = \begin{bmatrix} \mathbf{S}_{p1} \\ \tilde{\mathbf{S}}_p \end{bmatrix} = \begin{bmatrix} \mathbf{S}_{p1} \\ \mathbf{0} \end{bmatrix} \tag{13}$$

Therefore our expected prior value of the coordinates of the non-rigid bases $\tilde{\mathbf{S}}_p$ is zero in this case. For every rigid point in the scene we model the distribution of $\tilde{\mathbf{S}}_p$ as a Gaussian with a small variance and solve the problem as a Maximum A Posteriori estimation (MAP).

5 Results

5.1 Synthetic Data

The synthetic 3D data consisted of a set of random points sampled inside a cube of size $50 \times 50 \times 50$ units. Five sequences were generated with 8, 16, 32, 64 and 128 non-rigid points sampled inside the cube. Each sequence also included 8 rigid points (the vertices of the cube). Figure 1 shows the 3D data used in each of the five sequences with the rigid points joined up for display purposes. Our aim is to show the performance of our approach under different degrees of non-rigidity. The deformations for the non-rigid points were generated using random basis shapes as well as random deformation weights. Two basis shapes were used and the first basis shape had the largest weight equal to 1. The data was then rotated and translated over 25 frames and projected onto the images using an orthographic camera model and Gaussian noise was added to the image coordinates. The overall rotation about any axis was 90 degrees at most.

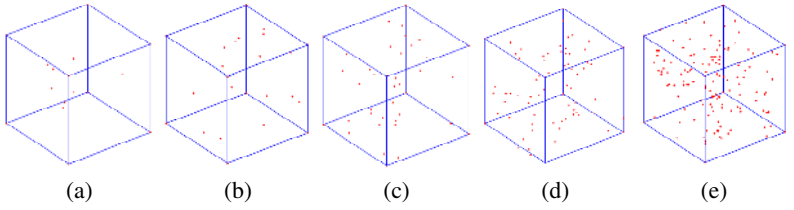

Fig. 1. Synthetic sequence. Example of ground truth of the 3D shape with 8 rigid points (vertices of the cube) and (a) 8, (b) 16, (c) 32, (d) 64 and (e) 128 non-rigid points.

Rigid and Non-rigid Motion Segmentation. Figure 4(a) shows results of the motion segmentation algorithm on a sequence using 8 rigid and 32 non-rigid points. The noise level for this particular experiment was set to be $\sigma = 1.5$ pixels. The $-y$ axis of the graph shows the current value of the deformation index (K) and the $-x$ axis represents the total number of total points left at each iteration. The algorithm classifies points according to the current value of K. The first 32 points are selected as non-rigid as their deformability index K is consistently close to 2. When the 33^{rd} point is selected one can observe a sudden drop in the value of K to 1.5 which then tends to 1. This is the cut-off point and the 8 remaining points are correctly classified as rigid.

3D Reconstruction. We have tested 3 reconstruction algorithms: the linear GSVD method, bundle adjustment without priors (MLE) and bundle-adjustment incorporating priors on the 3D structure (MAP). Figure 2 shows the 2D image reprojection error, relative 3D reconstruction error and absolute rotation error using each of the 3 algorithms, for varying ratios of rigid/non-rigid scene points and different levels of measurement noise. It becomes clear that GSVD and MAP outperform MLE thus showing the improved performance when prior information on the shape is incorporated. In fact the GSVD and MAP error curves appear superimposed which shows that they converge to

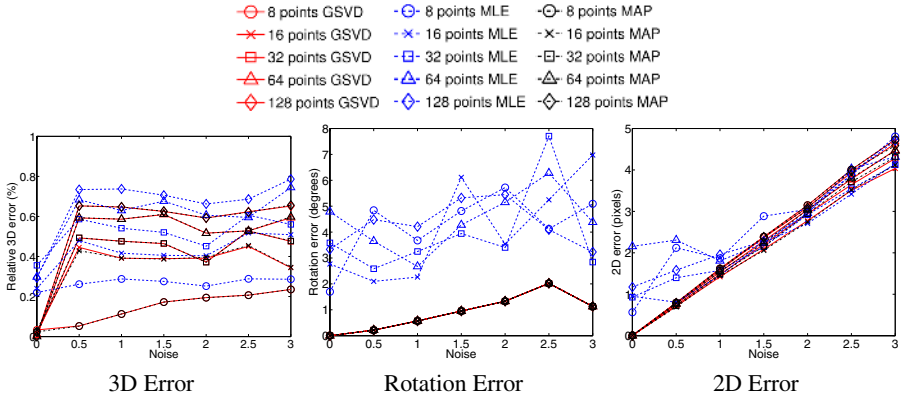

Fig. 2. Relative 3D error (%), rms rotation error (deg) and 2D reprojection error for the synthetic experiments for different ratios of rigid/non-rigid points and increasing levels of noise

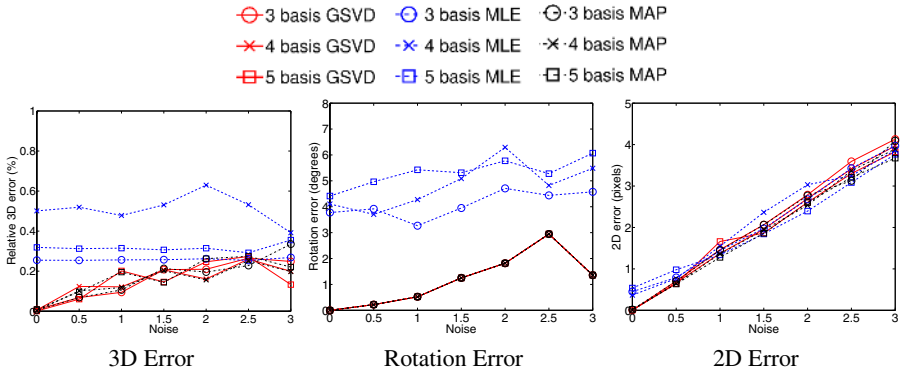

Fig. 3. Relative 3D error (%), rms rotation error (deg) and 2D reprojection error for the synthetic experiments for different numbers of basis shapes and increasing levels of noise

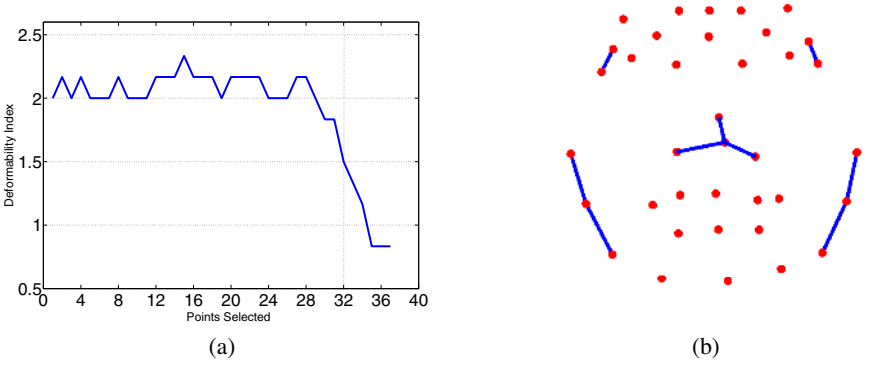

Fig. 4. (a) Deformability index for the automatic segmentation experiment. The graph shows its sudden decrease upon selection of point #33 (the first rigid point). (b) Face data used in the real experiment. Points connected with wireframes show the selected rigid points located on the nose, temples and side of the face.

the same solution, with the main observable difference being the higher speed of convergence for the MAP approach. Note that the MLE approach is not able to compute a correct 3D reconstruction even for the noiseless case showing that the added priors are fundamental to avoid local minima given by ambiguous configurations of motion and deformation parameters.

The number of basis shapes were then varied ($d = 3$, 4 and 5) to test the performance of the algorithm with respect to this parameter. Figure 3 shows the 2D image reprojection error, relative 3D reconstruction error and absolute rotation error obtained with GSVD, MLE and MAP. As expected, the error increases with the number of basis shapes for all 3 algorithms. Once more GSVD and MAP have almost identical performance and provide better results than MLE.

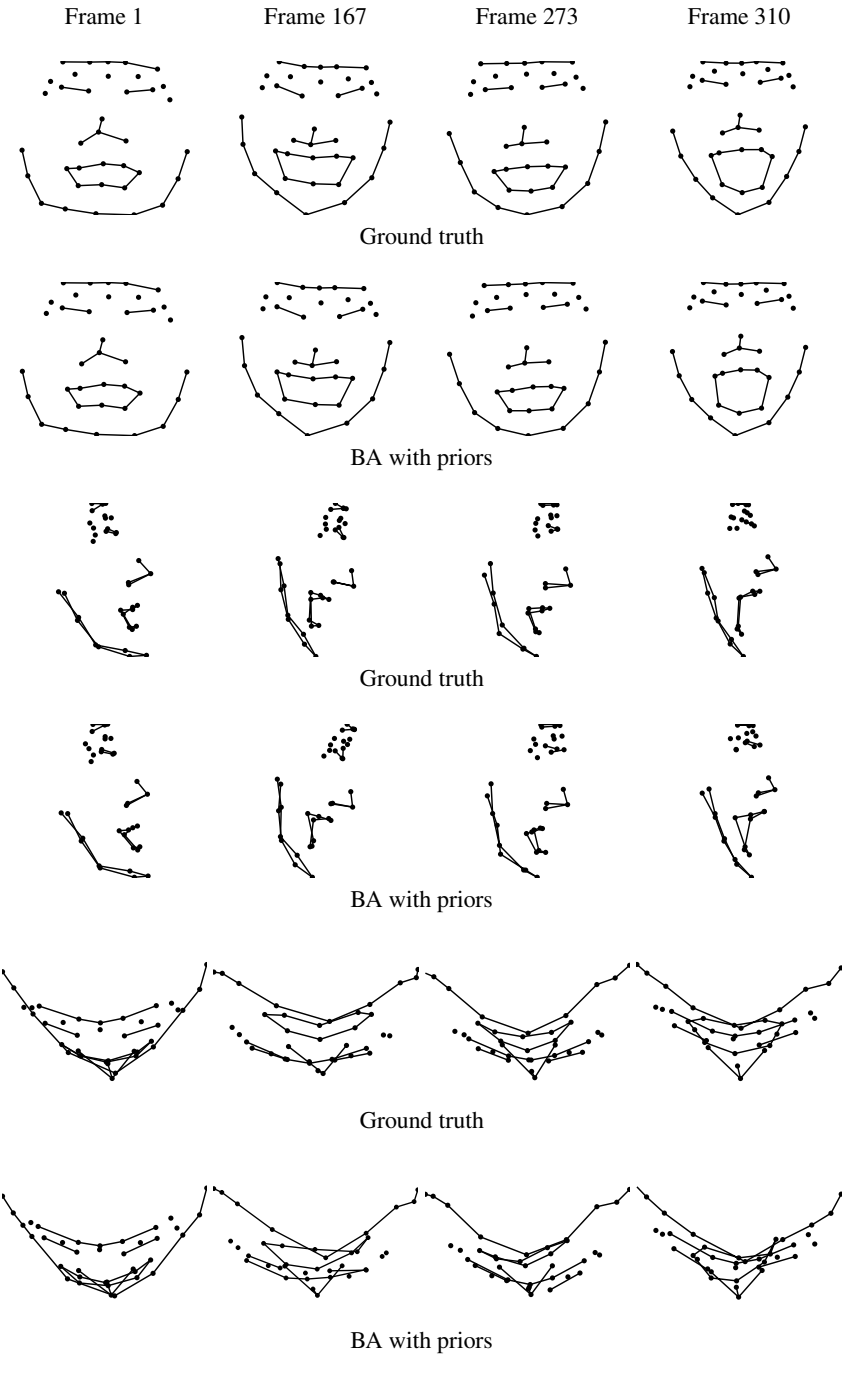

Fig. 5. Front, side and top views of the ground truth and reconstructed face with priors. Reconstructions are shown for frames 1, 167, 273 and 310.

5.2 Real Data

In this experiment we use real 3D data of a human face undergoing rigid motion – mainly rotation – while performing different facial expressions. The 3D data was captured using a VICON motion capture system by tracking the subject wearing 37 markers on the face. The 3D points were then projected synthetically onto an image sequence 310 frames long using an orthographic camera model and noise of variance $\sigma = 0.5$ pixels was added to the image coordinates. In this case the segmentation of points into rigid and non-rigid sets was done manually. Figure 4(b) shows a frontal view of the face where the 14 rigid points – situated on the nose, temples and the side of the face – are marked with circles.

Figure 5 shows the ground truth and reconstructed shape from front, side and top views using the bundle adjustment algorithm incorporating rigidity priors on the non-deforming points. The deformations are very well captured by the model even for the frames in which the facial expressions are more exaggerated. Crucially, the addition of the priors helps remove the ambiguity between the rotational and deformation components.

6 Future Work

We are currently investigating other solutions for rigid and non-rigid motion segmentation using alternative approaches [11,16]. Torresani et al.'s [15] method for learning non-rigid 3D shape from 2D motion using the expectation maximization algorithm could also be used to perform the segmentation.

We also plan to exploit looser rigidity constraints on the shape such as the 3D points behaving rigidly only for a set of frames and to include other priors such as the symmetry of the structure.

Acknowledgements

We acknowledge Dr. Amit K. Roy-Chowdhury for fruitful discussions on the estimation of the deformability index. This work has been partially supported by EPSRC grant GR/S61539/01. Alessio Del Bue holds a Queen Mary Studentship award.

References

1. H. Aanæs and F. Kahl. Estimation of deformable structure and motion. In *Workshop on Vision and Modelling of Dynamic Scenes, ECCV'02, Copenhagen, Denmark*, 2002.
2. M. Brand. Morphable models from video. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, Kauai, Hawaii*, December 2001.
3. C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3d shape from image streams. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, Hilton Head, South Carolina*, pages 690–696, June 2000.
4. A. Del Bue, F. Smeraldi, and L. Agapito. Non-rigid structure from motion using non-parametric tracking and non-linear optimization. In *Proceedings of the 2004 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'04) Volume 1*, Washington, DC, USA, 2004.

5. G. H. Golub and C. F. Van Loan. *Matrix Computation*. John Hopkins University Press, 1991. Second Edition.
6. P. C. Hansen. Regularization, gsvd and truncated gsvd. *BIT*, 29(3):491–504, 1989.
7. A.K. Jain and D. Zongker. Feature selection: Evaluation, application, and small sample performace. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):153–158, February 1997.
8. T. Kim and K-S Hong. Estimating approximate shape and motion of deformable objects with a monocular view. In *Proc. Asian Conference on Computer Vision*, Jeju Island, Korea, January 2004.
9. J. Kittler. Feature selection and extraction. In T. Y. Young and K. S. Fu, editors, *HPRIP*, pages 59–83, Orlando, FL, 1986. Academic Press.
10. A. Roy-Chowdhury. A measure of deformability of shapes with applications to human motion analysis. In *IEEE Conference in Computer Vision and Pattern Recognition*, volume 1, pages 398–404, June 2005.
11. Y. Sugaya and K. Kanatani. Multi-stage optimization for multi-body motion segmentation. *IEICE Transactions on Information and Systems*, E87-D(7):1935–1942, 2004.
12. C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: A factorization approach. *International Journal in Computer Vision*, 9(2):137–154, 1992.
13. L. Torresani, D. Yang, E. Alexander, and C. Bregler. Tracking and modeling non-rigid objects with rank constraints. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, Kauai, Hawaii*, 2001.
14. Lorenzo Torresani and Aaron Hertzmann. Automatic non-rigid 3d modeling from video. In *Proc. 8th European Conference on Computer Vision, Prague, Czech Republic*, pages 299–312, May 2004.
15. Lorenzo Torresani, Aaron Hertzmann, and Christoph Bregler. Learning non-rigid 3d shape from 2d motion. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.
16. R. Vidal and R. Hartley. Motion segmentation with missing data using powerfactorization and gpca. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 310–316, Washington D.C., June 2004.
17. J. Xiao, J. Chai and T. Kanade. A closed-form solution to non-rigid shape and motion recovery. In *Proc. 8th European Conference on Computer Vision, Prague, Czech Republic*, May 2004.
18. Anthony J. Yezzi and Stefano Soatto. Deformotion: Deforming motion, shape average and the joint registration and approximation of structures in images. *International Journal of Computer Vision*, 53(2):153–167, 2003.