

Extracting Average Shapes from Occluded Non-rigid Motion

Alessio Del Bue

Institute for Systems and Robotics, Instituto Superior Tecnico, Lisbon, Portugal

Abstract. This paper presents a method to efficiently estimate average 3-D shapes from non-rigid motion in the case of missing data. Such a shape can be further used to accomplish full reconstruction of deformable objects and registration of non-rigid shapes. The approach is based firstly on a power method which linearly provides an initial estimate of the 3-D structure and motion components of the object shape. Secondly, non-linear optimisation is used to refine the initial linear estimation. Tests on both real and synthetic sequences show the procedure effectiveness in dealing with different degrees of occlusions in the measurements.

1 Introduction

Recently the inference of the 3-D structure of a deforming body viewed by an uncalibrated camera has attracted increasing interest. In a Structure from Motion (SfM) domain, non-rigid shapes have posed new problems since they violate the rigidity constraints on which previous SfM methods strongly rely. Most of the model-free approaches to non-rigid SfM available nowadays are based either on closed-form solutions [12], assuming pre-specified shape priors, or iterative non-linear optimisation techniques [5, 1, 11], requiring an appropriate initialisation in order to converge. In the latter case, average shape and motion [9] have experimentally proven to be a successful initialisation to such tasks and they can be easily computed when the full trajectory of a point lying on the deforming body is available.

However, in the case of missing data affecting the trajectories (i.e. a point being occluded for some frames) a solution for the average shape is not currently available. Estimation of structure and motion from occluded data (see [3] for a review) is an essential task for most practical applications given the difficulty to obtain complete trajectories. At this end, the solution proposed here is an iterative power method which can estimate average shapes in the case of missing data and its reliability is assessed in a full 3-D reconstruction task for deforming objects. The approach is based on the notion of average shape introduced in [9] which penalizes in a certainty-reweighted scheme the non-rigidity of trajectories. This method can be extended by reformulating power methods for SfM [7] to include the notion of non-rigidity of a trajectory and extend it to the case of missing data.

In detail, the paper firstly introduces the non-rigid factorization framework and the definition of average shape (Section 2). Then, power methods for SfM are presented in Section 3 for the case of rigidly moving objects. The new approach with missing entry is explained in Section 4 and experiments (Section 5) show its effectiveness on synthetic test and on a face modelling task.

2 Non-rigid Structure from Motion

2.1 A Factorization Approach to Deformable Modelling

Tomasi and Kanade’s factorization algorithm [10] has been reformulated to the case of non-rigid 3-D structure [2]. A linear approximation of a set of K basis shapes S_k is used to describe a 3-D time varying shape X such that:

$$X = \sum_{k=1}^K l_k S_k \quad X, S_k \in \mathbb{R}^{3 \times P} \quad l_k \in \mathbb{R} \quad (1)$$

Each basis shapes S_k represent the mode of deformations of the deforming body and they are parameterised as a $3 \times P$ matrix which contains the 3-D 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:

$$W_i = [w_{i1} \dots w_{iP}] = R_i \left(\sum_{k=1}^K l_{ik} S_k \right) \quad (2)$$

where each $w_{ij} = [u_{ij} v_{ij}]^T$ with $j = 1 \dots P$ contains the horizontal and vertical image coordinates of the point – referred to the centroid of the object – and 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 W which can be expressed as:

$$W = \begin{bmatrix} w_{11} & \dots & w_{1P} \\ \vdots & & \vdots \\ w_{F1} & \dots & w_{FP} \end{bmatrix} = \begin{bmatrix} l_{11}R_1 & \dots & l_{1K}R_1 \\ \vdots & & \vdots \\ l_{F1}R_F & \dots & l_{FK}R_F \end{bmatrix} \begin{bmatrix} S_1 \\ \vdots \\ S_K \end{bmatrix} = MS. \quad (3)$$

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 \tilde{M} and a shape matrix \tilde{S} by truncating the SVD of W to rank $3K$. However, this factorization is not unique since any invertible $3K \times 3K$ matrix Q can be inserted in the decomposition leading to the alternative factorization: $W = (\tilde{M}Q)(Q^{-1}\tilde{S})$. The focal problem to solve in non-rigid factorization schemes is to find the Q that renders the appropriate replicated block structure of the motion matrix and that removes the affine ambiguity, upgrading the reconstruction to a metric one.

2.2 Extracting Average Shapes from Deformations

Based on the framework described in the previous section, Kim & Hong [9] recently introduced a measure called the Degree of Non-rigidity (*DoN*) to estimate the deviation of a deformable point from its average position. This measure can in turn be used to extract an average shape using an iterative certainty reweighted scheme. If the average 3-D shape of a time varying shape $X_i = [X_{i1} \dots X_{iP}]$ is given by $\hat{X} = [\hat{X}_1 \dots \hat{X}_P]$ the *DoN* for point j is defined as:

$$DoN_j = \sum_{i=1}^F (X_{ij} - \hat{X}_j)(X_{ij} - \hat{X}_j)^T. \quad (4)$$

The 2-D projection C_j of the DoN will be thus given by:

$$C_j = \sum_{i=1}^F R_i (X_{ij} - \hat{X}_j)(X_{ij} - \hat{X}_j)^T R_i^T = \sum_{i=1}^F (w_{ij} - \hat{w}_{ij})(w_{ij} - \hat{w}_{ij})^T \quad (5)$$

where w_{ij} are the image coordinates of point j at frame i and \hat{w}_{ij} are the coordinates of its projected mean shape. While the DoN cannot be computed without an estimation of the mean 3-D shape, the value of its projection can be estimated directly from image measurements.

An initial estimate of the projected 2-D mean shapes \hat{w}_{ij} could be given simply by the first basis shape S_1 (as in equation (3)) which could be computed with a rank-3 approximation $SV D_3(W) = \hat{M}\hat{S}$. The projected deviation from the mean for all the points would then be defined by $\{w_{ij} - \hat{w}_{ij}\} = W - \hat{M}\hat{S}$. However, a straight application of a rank-3 factorization over the first basis component does not produce an accurate measure of C_j as showed in [9]. To adjust the covariances, the average shape and C_j are iteratively estimated until convergence. However, the procedure is unable to deal with the case of missing data affecting the measurements. We will show in the next section how power methods can efficiently solve this issue.

3 Power Methods for Structure from Motion

SfM algorithms based on factorization require an initial decomposition of the motion M and structure matrix S given the data W . In this context, power methods were introduced with the name *powerfactorization* by Schaffalitzky and Hartley [7] to efficiently factorise rank-constrained image measurements. This approach is an alternation method which iteratively estimates M and S by simply executing multiplications and matrices inverse. The update rules at iteration q are given by [7]:

$$\begin{aligned} M_q &= WS_{q-1}^T (S_{q-1}^T S_{q-1})^{-1} \\ S_q &= (M_q^T M_q)^{-1} M_q^T W \end{aligned} \quad (6)$$

They are a straightforward derivation from the orthogonal power method [6] which convergence rate depends on the ratio of the dominant singular values of W . In the case of an affine camera viewing a moving rigid body, the update rules (6) can be modified to account for the geometrical properties of the measurements. For each frame $i = 1 \dots F$, the projection of a point $j = 1 \dots P$ can be expressed as:

$$w_{ij} = A_i X_j + a_i \quad (7)$$

where A_i is a 2×3 camera projection matrix, X_j a 3-vector of the 3-D coordinates and a_i a 2-vector of the affine camera translation. In a more compact form, equation (7) can be rewritten for every point at each frame as:

$$W_i = [A_i | a_i] \begin{bmatrix} X_1 & \dots & X_P \\ 1 & \dots & 1 \end{bmatrix} = M_i \begin{bmatrix} X \\ \mathbf{1}^T \end{bmatrix} \quad (8)$$

where $\mathbf{1}$ is a P -vector of ones. Finally, the global expression for each frame can be written as:

$$\mathbf{W} = \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_F \end{bmatrix} = \begin{bmatrix} [\mathbf{A}_1 | \mathbf{a}_1] \\ \vdots \\ [\mathbf{A}_F | \mathbf{a}_F] \end{bmatrix} \begin{bmatrix} \mathbf{X} \\ \mathbf{1}^T \end{bmatrix} = [\mathbf{A} | \mathbf{a}] \begin{bmatrix} \mathbf{X} \\ \mathbf{1}^T \end{bmatrix} = \mathbf{M}\mathbf{S} \quad (9)$$

The algorithm for extracting the affine motion and structure of a rigid object can be summarized as follows:

- Initialize \mathbf{X}_0 with random entries.
- Compute the $2F \times 4$ update of \mathbf{M}_q given equation (6).
- Extract the $2F \times 1$ measurements centroid \mathbf{a}_q such that $\mathbf{M}_q = [\mathbf{A}_q | \mathbf{a}_q]$.
- Compute the $3 \times P$ update of \mathbf{X}_q such that: $\mathbf{X}_q = (\mathbf{A}_q^T \mathbf{A}_q)^{-1} \mathbf{A}_q^T (\mathbf{W} - \mathbf{T}_q)$ where $(\mathbf{W} - \mathbf{T}_q)$ are the centered coordinates and $\mathbf{T}_q = \mathbf{a}_q \mathbf{1}_{1 \times P}$

4 Average Shape Estimation with Missing Data

4.1 Power Iterations and Degree of Non-rigidity

In the case of affine estimation of average shape $\hat{\mathbf{S}}$ and motion $\hat{\mathbf{M}}$, strongly non-rigid trajectories (which in turn provide high covariances \mathbf{C}_j) are penalized in the estimation of the average components. The estimation task can be recast in the minimisation of a cost function χ such that:

$$\chi = \sum_{i,j} (\mathbf{w}_{ij} - \hat{\mathbf{M}}_i \bar{\mathbf{X}}_j)^T \mathbf{C}_j^{-1} (\mathbf{w}_{ij} - \hat{\mathbf{M}}_i \bar{\mathbf{X}}_j) \quad (10)$$

where $\bar{\mathbf{X}}_j$ contains the the homogeneous coordinate for the average point such that $\bar{\mathbf{X}}_j = [\hat{\mathbf{X}}_j^T \ 1]^T$. Minimizing χ can be carried out with a minor reformulation of the power method [7]. In brief, each matrix \mathbf{C}_j^{-1} can be factored as $\mathbf{C}_j^{-1} = \mathbf{B}_j^T \mathbf{B}_j$ giving:

$$\chi = \sum_{i,j} (\mathbf{w}_{ij} - \hat{\mathbf{M}}_i \bar{\mathbf{X}}_j)^T \mathbf{B}_j^T \mathbf{B}_j (\mathbf{w}_{ij} - \hat{\mathbf{M}}_i \bar{\mathbf{X}}_j) = \sum_{i,j} \|\mathbf{B}_j \mathbf{w}_{ij} - \mathbf{B}_j \hat{\mathbf{M}}_i \bar{\mathbf{X}}_j\|^2 \quad (11)$$

Notice the similarity with equation (6) which hints to a solution of the minimization of (10) with a power approach. In order to obtain the two updates rules for motion and structure we can rewrite (11) such that:

$$\chi = \sum_{i,j} \|\mathbf{B}_j \mathbf{w}_{ij} - \mathbf{B}_j \tilde{\mathbf{X}}_j \tilde{\mathbf{m}}_i\|^2 \quad (12)$$

with:

$$\tilde{\mathbf{X}}_j = \begin{bmatrix} \bar{\mathbf{X}}_j^T & 0 \\ 0 & \bar{\mathbf{X}}_j^T \end{bmatrix} \text{ and } \tilde{\mathbf{m}}_i = [\mathbf{m}_{1i}^T \ \mathbf{m}_{2i}^T]^T \quad (13)$$

where \mathbf{m}_{1i}^T and \mathbf{m}_{2i}^T are respectively the first and second 4×1 rows of $\hat{\mathbf{M}}_i$. Given the quadratic costs (11) and (12) we can express the power updates for the motion as:

$$\tilde{\mathbf{m}}_i = \left(\sum_j \tilde{\mathbf{X}}_j^T \mathbf{B}_j^T \mathbf{B}_j \tilde{\mathbf{X}}_j \right)^{-1} \sum_j \tilde{\mathbf{X}}_j^T \mathbf{B}_j^T \mathbf{B}_j \mathbf{w}_{ij} = \left(\sum_j \tilde{\mathbf{X}}_j^T \mathbf{C}_j^{-1} \tilde{\mathbf{X}}_j \right)^{-1} \sum_j \tilde{\mathbf{X}}_j^T \mathbf{C}_j^{-1} \mathbf{w}_{ij} \quad (14)$$

After rearranging $\tilde{\mathbf{m}}_i \mapsto \hat{\mathbf{M}}_i = [\mathbf{A}_i | \mathbf{a}_i]$ we obtain:

$$\hat{\mathbf{X}}_j = \left(\sum_i \mathbf{A}_i^T \mathbf{C}_j^{-1} \mathbf{A}_i \right)^{-1} \sum_i \mathbf{A}_i^T \mathbf{C}_j^{-1} (\mathbf{w}_{ij} - \mathbf{a}_i) \quad (15)$$

where \mathbf{a}_i is the overall translation component as defined in (9). Once the estimates for the average $\hat{\mathbf{M}}$ and $\hat{\mathbf{S}}$ are available, \mathbf{C}_j is update by equation (5).

4.2 The Missing Data Case

We can now assume that some points are not visible in some frames due to occlusion. In order to include missing data, we can modify the power update equations in (14) and (15) to simply not include the equations regarding the missing entries giving:

$$\tilde{\mathbf{m}}_i = \left(\sum_j \tilde{\mathbf{X}}_j^T \mathbf{C}_j^{-1} \tilde{\mathbf{X}}_j \right)^{-1} \sum_j \tilde{\mathbf{X}}_j^T \mathbf{C}_j^{-1} Z_{ij} \mathbf{w}_{ij} \quad (16)$$

$$\hat{\mathbf{X}}_j = \left(\sum_i \mathbf{A}_i^T \mathbf{C}_j^{-1} \mathbf{A}_i \right)^{-1} \sum_i \mathbf{A}_i^T \mathbf{C}_j^{-1} Z_{ij} (\mathbf{w}_{ij} - \mathbf{a}_i) \quad (17)$$

where Z_{ij} is a scalar which is zero whenever a point is missing and one otherwise. The updates have the property of efficiently estimating the centroid at each frame \mathbf{a}_i since the measurement matrix of missing data may be not mean-centered. Schematically, the algorithm can be outlined as follows¹:

- Initialize \mathbf{X} with random entries.
- Compute the $2F \times 4$ update of $\hat{\mathbf{M}}_i$ for $i = 1 \dots F$ given equation (16).
- Given $\hat{\mathbf{M}}_i = [\mathbf{A}_i | \mathbf{a}_i]$, extract the measurements centroid \mathbf{a}_i .
- Compute the update of the average 3-D structure with (17).
- Recompute $\mathbf{C}_j = \sum_{i=1}^F Z_{ij} (\mathbf{w}_{ij} - \mathbf{a}_i - \mathbf{A}_i \hat{\mathbf{X}}_j) (\mathbf{w}_{ij} - \mathbf{a}_i - \mathbf{A}_i \hat{\mathbf{X}}_j)^T$.
- Iterate until convergence.

A metric update of the average shape can be then obtained in the case of orthographic [10] and weak or para-perspective cameras [8] by computing the correct 3×3 transformation \mathbf{Q} for the average shape.

4.3 Non-linear Optimisation and Non-rigid SfM

A full deformable 3-D reconstruction as presented in Section 2.1 can be successfully computed linearly only when particular assumptions over the data are given and when the full trajectories are available. For instance, in [12] the authors proved the existence of a unique solution and a closed form algorithm when K independent 3-D shapes can be identified in the measured data. On the other hand, a more general solution consists in performing non-linear optimisation [5, 1, 11] by minimizing a cost function which reflects the full deformable model as presented in equation (3) giving:

$$\min_{\mathbf{R}_i \mathbf{S}_{kj} l_{ik}} \sum_{i,j} Z_{ij} \|\mathbf{w}_{ij} - \hat{\mathbf{x}}_{ij}\|^2 = \min_{\mathbf{R}_i \mathbf{S}_{kj} l_{ik}} \sum_{i,j} Z_{ij} \|\mathbf{w}_{ij} - (\mathbf{R}_i \sum_k l_{ik} \mathbf{S}_{kj})\|^2 \quad (18)$$

¹ for clarity we drop the iteration subscript q

where \mathbf{S}_{kj} is the 3×1 basis for the point j such that $\mathbf{S}_k = [\mathbf{S}_{k1} \cdots \mathbf{S}_{kP}]$. Again, the least-squares entries for the missing data are omitted. Initialisation of the model parameters are provided by the average shape computed with our power approach.

5 Experiments

5.1 Synthetic Data

The proposed power approach is first validated using randomly generated synthetic data of a deforming shape. The 3-D bodies are generated by firstly sampling the first basis shape \mathbf{S}_1 over the surface of a sphere. The following basis $\mathbf{S}_2 \dots \mathbf{S}_K$, which express the modes of deformation of the body, are generated randomly. In order to obtain a given deformation at frame i , the configuration weights $l_{i1} \dots l_{iK}$ are computed by fitting 4-order polynomials to random samples, this gives more regular deformation rather than erratic motion. The computed 3-D shapes are then normalized to obtain a specific ratio of deformation defined as $\frac{\sum_{i=1}^F \|\sum_{k=2}^K l_{ik} \mathbf{S}_k\|^2}{\sum_{i=1}^F \|\mathbf{S}_1\|^2}$ which is fixed to 0.25. The final measurement matrix \mathbf{W} is obtained by projecting each 3-D non-rigid shape onto the image plane by means of random orthographic cameras. Finally, points are eliminated given different ratios of missing data.

We test the algorithm performances in providing a meaningful initialisation to the optimisation problem as defined in section 4.3. Firstly, we noticed problems in the convergence if the algorithm for computing the average shape does not include the iterative re-weighting with \mathbf{C}_j . The overall results are showed in table 1 with different levels of image noise affecting the data. A decrease in the algorithm's performances is given for ratios of 30% and more missing data. Regarding the mean shape computation, convergence is generally achieved after 15 iterations with 10% of missing data, higher ratios increase this number however, in the worst case, the algorithm was not performing more than 50 steps.

5.2 Real Data

The real experiments are focused on extracting a mean shape from a deforming face² exhibiting a light rotation and non-rigid motion especially in the mouth region. We selected a 700 frames long sequence from the overall 5000 frames and 56 points are collected to form the measurement matrix \mathbf{W} . Occluded points appear with an overall 20% ratio of missing entries. The recovered mean shape (see figure 1) is then used to initialize a full deformable reconstruction and some frames presenting the recovered 3-D depth and deformations are presented (front and side view). The approach is able to successfully recover a reasonable estimates of the depth and deformations even if the subject is not performing strong rigid motion. The final number of iterations for the power method was of 50 followed by 40 iteration of non-linear optimisation.

² sequence available at: www-prima.inrialpes.fr/FGnet/data/01-TalkingFace/talking_face.html

6 Conclusions

We presented a power approach to estimate average shapes from non-rigid motion in the case of missing data. Experimentally we have shown the effectiveness of the method in a deformable 3-D reconstruction task with affine cameras. The extracted average shape and motion have been shown to provide a reliable initialisation for SfM optimisation tasks in the tested cases. As a further study, the method may be extended to more general camera models (i.e. full perspective), however initial results have shown increased instability in the convergence given by the difficulty in decoupling deformations from perspective distortions. In such cases, an approach using shape priors as presented in [4] may help to successfully compute a reliable average shape.

Acknowledgments. This work was supported by Fundação para a Ciência e a Tecnologia (ISR/IST plurianual funding) through the POS.Conhecimento Program (include FEDER funds). Thanks to João Paulo Costeira for support and for a careful reading of this paper.

References

1. M. Brand. A direct method for 3d factorization of nonrigid motion observed in 2d. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, San Diego, California*, pages 122–128, 2005.
2. 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.
3. A. M. Buchanan and A. Fitzgibbon. Damped newton algorithms for matrix factorization with missing data. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, San Diego, California*, volume 2, pages 316–322, 2005.
4. A. Del Bue, X. Lladó, and L. Agapito. Non-rigid face modelling using shape priors. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, New York, NY*, 2006.
5. A. Del Bue, F. Smeraldi, and L. Agapito. Non-rigid structure from motion using ranklet-based tracking and non-linear optimization. *Image and Vision Computing*, 25(3):297–310, March 2007.
6. G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins, 1989.
7. R. I. Hartley and F. Schaffalitzky. Powerfactorization: an approach to affine reconstruction with missing and uncertain data. In *Australia-Japan Advanced Workshop on Computer Vision*, Adelaide, Australia, September 2003.
8. K. Kanatani and Y. Sugaya. Factorization without factorization: complete recipe. *Memories of the Faculty of Engineering, Okayama University*, 38(1–2):61–72, 2004.
9. Taeone Kim and Ki-Sang Hong. Estimating approximate average shape and motion of deforming objects with a monocular view. *International Journal of Pattern Recognition and Artificial Intelligence*, 19(4):585–601, 2005.
10. C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: A factorization approach. *International Journal of Computer Vision*, 9(2):137–154, 1992.
11. 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.
12. J. Xiao, J. Chai, and T. Kanade. A closed-form solution to non-rigid shape and motion recovery. *International Journal of Computer Vision*, 67(2):233–246, April 2006.

Missing %	Noise				
	0	0.5	1	1.5	2
10%	1, 32	1, 47	1, 89	2, 11	2, 13
20%	2, 85	3, 69	3, 45	3, 69	4, 05
30%	3, 75	4, 74	4, 76	5, 03	5, 78
40%	3, 99	4, 64	5, 18	5, 47	6, 87

Rotation Error

Missing %	Noise				
	0	0.5	1	1.5	2
10%	0.84	1.10	1.02	1.38	1.94
20%	1.26	1.38	2.05	1.26	2.55
30%	1.41	1.62	2.19	2.21	2.18
40%	1.78	1.86	1.96	2.39	2.40

3-D Structure error

Table 1. Left: Mean of the the absolute rotation error expressed in degrees. Right: 3-D reconstruction error expressed in percentage relative to the scene size. The variance of the added noise is expressed in terms of image pixel. The value are computed on 10 trials for each configuration of noise and missing data ratios.

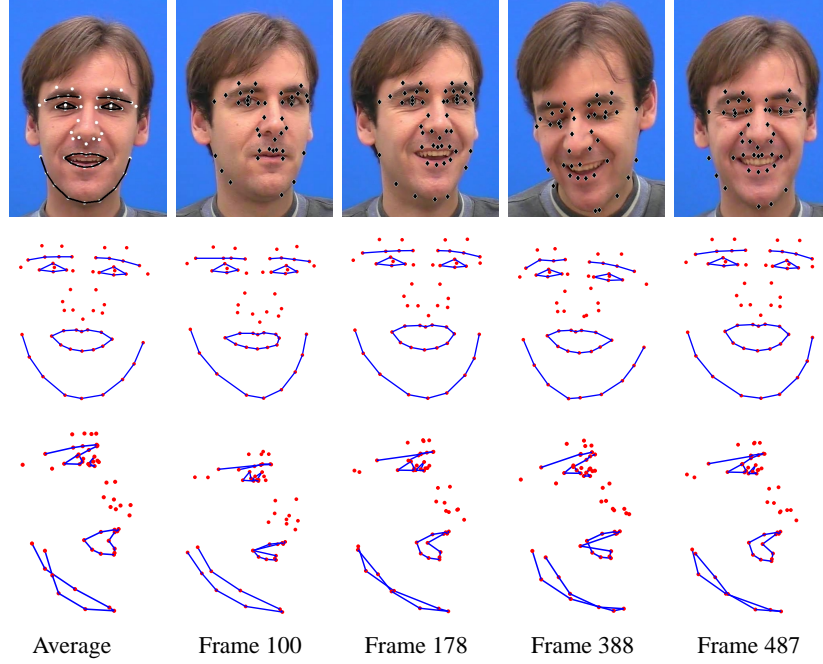


Fig. 1. The first column shows the complete set of 56 points used for reconstruction (first row) and the recovered 3-D average shape (front and side views). The remaining columns present 4 key frames of the sequence with the available points at each frame. The second and third rows present respectively the front and side views of the reconstructed 3-D structure after non-linear optimisation. The number of basis shapes was fixed to $K = 6$.