

Articulated Structure From Motion by Factorization

P. Tresadern and I. Reid

Dept. of Engineering Science, University of Oxford, Oxford, England OX1 3PJ
{pat,ian}@robots.ox.ac.uk

Abstract

Multibody affine Structure From Motion (SFM) methods commonly assume independent motion between objects such that the ‘measurement matrix’ has rank $4k$. When multiple views are available, each object is then independently calibrated to a metric co-ordinate frame.

However, articulated motion results in a further decrease in rank – a fact that we exploit to detect articulated objects and determine their degrees of freedom using simple linear methods. Furthermore, these objects cannot be recovered and calibrated independently since this violates articulation constraints.

We show that articulation constraints can be imposed during factorization and self-calibration to recover consistent 3D structure and motion, from which link lengths and joint angles can be computed. The stability of the method is evaluated using synthetic data for comparison with ground truth and results are also presented for real image sequences.

1. Introduction

Structure From Motion (SFM) methods commonly assume a static scene is imaged by a moving camera. When perspective effects are sufficiently small, affine projection may also be assumed. As a result, affine structure and motion can be recovered from feature tracks using the Singular Value Decomposition (SVD), as elegantly demonstrated in the seminal ‘factorization method’ of Tomasi and Kanade [14].

Subsequent studies have investigated *independently moving objects* for motion segmentation [4], SFM with motion priors and missing data [8], and perspective camera calibration [6]. Other work has used the factorization method to represent deformable objects as a linear combination of “basis shapes” [1, 2] for small changes in shape.

In this paper, however, we look at *articulated* objects (e.g. furniture, the human body) that cannot be represented by a single statistical shape model. In this case, the relative

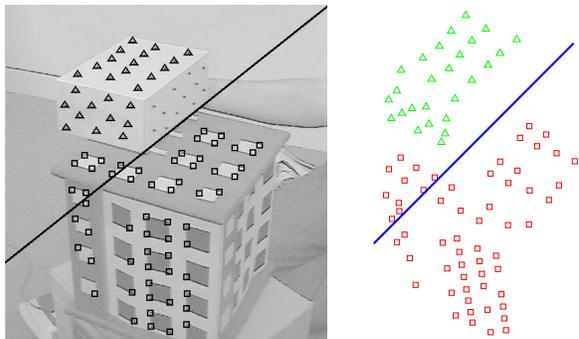


Figure 1. (left) Tracked features and recovered rotation axis. (right) Recovered 3D structure and axis of rotation.

motions of the objects are *dependent* and can be shown [19] to result in the failure of algorithms that assume independent motion (e.g. subspace-based motion segmentation [4]). If we are to recover accurate structure and motion *that satisfies articulation constraints*, this dependency should be incorporated from the beginning (i.e. during factorization).

Specifically, in this work we show how to detect articulated motion and determine the type of coupling between two objects from feature tracks alone. In particular, we show that articulation constraints can be applied as part of the factorization and self-calibration processes themselves. The axis (or centre) of rotation can then be projected into the image with ease whilst link lengths and joint angles are recovered in a metric co-ordinate frame (Figure 1).

Other than human motion analysis (see [7] for a review), single frame pose estimation [13] and model-based tracking systems [5], little attention has been paid to articulated objects. The only directly related work we know of is that of Sinclair *et al* [11] who recover metric articulated structure and motion using perspective cameras. They assume that articulation is about a hinge and the axis of rotation is approximately vertical in the image. Non-linear minimization is used to find points on the axis and they assume that some planar structure is visible.

In contrast, we exploit an affine projection model since the two objects are connected and it is sensible to assume that their relative depth is small compared to their distance from the camera. This greatly simplifies the recovery of articulated structure from motion since we can use linear methods rather than expensive search and non-linear optimization techniques. Furthermore, we do not assume to know the type of articulation *a priori* (we recover this from image data), nor that the axis of rotation is visible in the image, nor do we require any planar structure (visible or otherwise). Although we use fixed cameras in these experiments, this is not a requirement and the method is equally applicable for independently moving cameras.

We outline factorization for articulated objects in Section 2 and discuss self-calibration in Section 3. Results are shown in Section 4 and we conclude in Section 5.

2. Factorization with articulation

In [14] it was shown that a matrix, \mathbf{W} , of feature tracks from a static scene imaged by a moving orthographic camera is at most rank 4 since:

$$\mathbf{W} = \begin{bmatrix} \mathbf{x}_1^1 & \cdots & \mathbf{x}_n^1 \\ \vdots & & \vdots \\ \mathbf{x}_1^F & \cdots & \mathbf{x}_n^F \end{bmatrix} = [\mathbf{R} \quad \mathbf{t}] \begin{bmatrix} \mathbf{S} \\ \mathbf{1} \end{bmatrix} \quad (1)$$

where \mathbf{R} is a $2F \times 3$ ‘motion’ matrix, \mathbf{t} is the $2F \times 1$ projected centroid of the features over the sequence and \mathbf{S} is the $3 \times N$ matrix of 3D feature locations. Tracks from a given object therefore lie in a 4D subspace such that affine structure and motion are recovered by factorization using the SVD and calibrated to a metric coordinate frame by imposing constraints on the rows of \mathbf{R} .

For two independent motions, the ‘motion space’ scales accordingly such that $\text{rank}(\mathbf{W}) = 8$. However, for *dependent* motions there is a decrease in $\text{rank}(\mathbf{W})$ that we use both to detect articulated motion and to estimate the parameters of the joint. For the remainder of the paper, quantities associated with a second object are primed (e.g. \mathbf{R}' , \mathbf{t}' , etc).

2.1. Robust motion segmentation

It is necessary to segment the objects in order to group feature tracks according to the object that generated them. However, many existing methods are prone to failure in the presence of dependent motions [4] and gross outliers [17]. We therefore implement a RanSaC strategy for motion segmentation and outlier rejection [15].

Since four points in general position are sufficient to define an object’s motion, we use samples of four tracks to find consensus among the rest. Having recovered the first object’s motion, we remove all corresponding features and

repeat for the second. All remaining feature tracks are rejected as outliers since the factorization method uses the SVD (a linear least squares operation) and gross outliers severely degrade performance.

Having segmented the motions, we group the columns of \mathbf{W} accordingly and project each object’s features onto its closest $\text{rank} = 4$ matrix to reduce noise. It is then necessary to compute the SVD again – this time on the combined matrix of *both* sets of tracks – in order to estimate the parameters of the coupling between them.

2.2. Objects coupled by a universal joint

When two objects are linked by a two or three degree of freedom ‘universal’ joint (Figure 2), the position of one body with respect to the other is constrained but *their rotations are independent*.

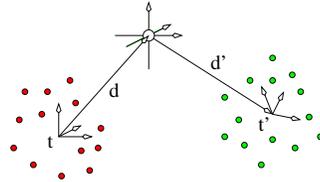


Figure 2. Schematic of a universal joint.

In particular, the centre of rotation (or ‘joint centre’) satisfies both motions such that the two 4D subspaces have a 1D intersection and $\text{rank}(\mathbf{W}) = 7$. Geometrically:

$$\mathbf{t} + \mathbf{R}\mathbf{d} = \mathbf{t}' - \mathbf{R}'\mathbf{d}' \quad (2)$$

and we can show that:

$$\mathbf{W} = [\mathbf{R} \quad \mathbf{R}' \quad \mathbf{t}] \begin{bmatrix} \mathbf{S} & \mathbf{d} \\ \mathbf{S}' + \mathbf{d}' \\ \mathbf{1} & \mathbf{1} \end{bmatrix} \quad (3)$$

where the quantities of interest are $\mathbf{d} = [u, v, w]^T$ and $-\mathbf{d}' = [u', v', w']^T$, representing the position vectors of the joint centre in the affine coordinate frames of the first and second object, respectively.

From Equation (2), it follows that $[\mathbf{d}^T, \mathbf{d}'^T, -1]^T$ lies in the right (column) nullspace of $[\mathbf{R}, \mathbf{R}', \mathbf{t}' - \mathbf{t}]$. As a result, \mathbf{d} and \mathbf{d}' can be recovered once \mathbf{R} , \mathbf{R}' , \mathbf{t} and \mathbf{t}' are known.

Since \mathbf{t} and \mathbf{t}' are the 2D centroids of the two point clouds, they are simply the row means of the matrix of feature tracks for the first and second object, respectively. Following [14] we translate each object to the origin, giving the ‘normalized’ $\text{rank} = 6$ system:

$$\widetilde{\mathbf{W}} = [\mathbf{R} \quad \mathbf{R}'] \begin{bmatrix} \mathbf{S} & \\ & \mathbf{S}' \end{bmatrix}. \quad (4)$$

This is effectively “full rank” since the rotations are independent and have been decoupled from the translations (where the dependency resides). From Equation (4), we can recover \mathbf{R} and \mathbf{R}' by factorization using the SVD. In practice, however, taking the SVD of $\widetilde{\mathbf{W}}$ recovers a full structure matrix, $[\mathbf{V}, \mathbf{V}']$, rather than the block diagonal form seen in Equation (4). We therefore separate the objects by premultiplying $[\mathbf{V}, \mathbf{V}']$ with a matrix, \mathbf{A}_U :

$$\begin{aligned} \mathbf{A}_U[\mathbf{V}, \mathbf{V}'] &= \begin{bmatrix} N_L(\mathbf{V}') \\ N_L(\mathbf{V}) \end{bmatrix} [\mathbf{V}, \mathbf{V}'] \\ &= \begin{bmatrix} N_L(\mathbf{V}')\mathbf{V} & \mathbf{0} \\ \mathbf{0} & N_L(\mathbf{V})\mathbf{V}' \end{bmatrix} \end{aligned} \quad (5)$$

where $N_L(\cdot)$ is an operator that returns the left (row) nullspace of its (matrix) argument. Finally, we transform the recovered motion matrix, $[\mathbf{U}, \mathbf{U}']$, accordingly: $[\mathbf{U}, \mathbf{U}']\mathbf{A}_U^{-1} \rightarrow [\mathbf{R}, \mathbf{R}']$. Having recovered \mathbf{R} , \mathbf{R}' , \mathbf{t} and \mathbf{t}' we can now compute \mathbf{d} and \mathbf{d}' . The reprojected joint centre is then simply $\mathbf{t} + \mathbf{R}\mathbf{d}$ (or $\mathbf{t}' - \mathbf{R}'\mathbf{d}'$).

Although in this case we could recover \mathbf{R} and \mathbf{R}' by factorization of each object independently, here we use a method that deals with both objects simultaneously for consistency with the hinge case where independent factorization is not so straightforward.

2.3. Objects coupled by a hinge

For bodies linked by a hinge (Figure 3), *their relative orientation is also constrained* since their co-ordinate frames have a common axis that is parallel to the axis of rotation.

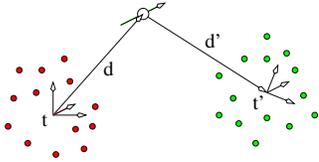


Figure 3. Schematic of a hinge joint.

In this case, *all* points on the rotation axis satisfy both motions such that the subspaces have a 2D intersection and $rank(\mathbf{W}) = 6$. Aligning the rotation axis with the x -axis by choosing an appropriate global co-ordinate frame, we denote the motion matrices by $\mathbf{R} = [\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3]$ and $\mathbf{R}' = [\mathbf{c}_1, \mathbf{c}'_2, \mathbf{c}'_3]$ to give the ‘normalized’ system:

$$\widetilde{\mathbf{W}} = [\mathbf{c}_1 \ \mathbf{c}_2 \ \mathbf{c}_3 \ \mathbf{c}'_2 \ \mathbf{c}'_3] \begin{bmatrix} x_1 \cdots x_{n_1} & x'_1 \cdots x'_{n_2} \\ y_1 \cdots y_{n_1} & \\ z_1 \cdots z_{n_1} & \\ & y'_1 \cdots y'_{n_2} \\ & z'_1 \cdots z'_{n_2} \end{bmatrix}. \quad (6)$$

Due to the dependency in rotation, factorizing the objects independently is not straightforward. Using the form in Equation (6) ensures that both objects have the same x -axis and respect the “common axis” constraint such that rotations are *not* independent. To zero out entries of the recovered $[\mathbf{V}, \mathbf{V}']$ we premultiply with a matrix, \mathbf{A}_H :

$$\mathbf{A}_H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ & N_L(\mathbf{V}') & & & \\ & & N_L(\mathbf{V}) & & \end{bmatrix}, \quad (7)$$

and transform $[\mathbf{U}, \mathbf{U}']$ accordingly.

Note that the ‘joint centre’ may lie anywhere on the axis of rotation, provided that $u + u' = k$ where k is the distance between object centroids parallel to the rotation axis. As a result, $[u + u', v, w, v'w', -1]^T$ lies in the nullspace of $[\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}'_2, \mathbf{c}'_3, \mathbf{t}' - \mathbf{t}]$ and can be recovered with ease. The reprojected axis of rotation is then given by the line $l(\alpha) = \mathbf{t} + [\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3][\alpha, v, w]^T$ where α is any real number.

In this case, we can show that:

$$\mathbf{W} = [\mathbf{c}_1 \ \mathbf{c}_2 \ \mathbf{c}_3 \ \mathbf{c}'_2 \ \mathbf{c}'_3 \ \mathbf{t}] \begin{bmatrix} \mathbf{x} & \mathbf{x}' + \mathbf{u} + \mathbf{u}' \\ \mathbf{y} & \mathbf{v} \\ \mathbf{z} & \mathbf{w} \\ & \mathbf{y}' + \mathbf{v}' \\ & \mathbf{z}' + \mathbf{w}' \\ \mathbf{1} & \mathbf{1} \end{bmatrix} \quad (8)$$

where $\mathbf{x} = [x_1, \dots, x_{n_1}]$, $\mathbf{u} = [u, \dots, u]$, etc.

2.4. Longer kinematic chains

Although we demonstrate this method for two links, it is equally applicable to longer kinematic chains – each additional link increases $rank(\mathbf{W})$ by $4 - m$ where $m=1$ for a universal joint and $m=2$ for a hinge.

3. Self-calibration

For many applications, it is desirable to recover joint angles and distances between joints. However, lengths and angles are not preserved in an affine coordinate frame so we self-calibrate structure and motion to a metric space.

Conventionally, constraints are imposed on the rows of \mathbf{R} so that each 2×3 block corresponding to a given frame (denoted by \mathbf{R}_f) is close to the first two rows of a scaled rotation matrix. This is achieved by postmultiplying \mathbf{R} with a 3×3 upper triangular matrix, \mathbf{B} . Specifically, if we denote \mathbf{R}_f by:

$$\mathbf{R}_f = \begin{bmatrix} \mathbf{i}^T \\ \mathbf{j}^T \end{bmatrix}, \quad (9)$$

the constraints of unit aspect ratio and zero skew are expressed algebraically as:

$$\mathbf{i}^T \mathbf{B} \mathbf{B}^T \mathbf{i} - \mathbf{j}^T \mathbf{B} \mathbf{B}^T \mathbf{j} = 0, \quad (10)$$

$$\mathbf{i}^T \mathbf{B} \mathbf{B}^T \mathbf{j} = 0. \quad (11)$$

These constraints are linear in the elements of the matrix $\mathbf{\Omega} = \mathbf{B} \mathbf{B}^T$, which are recovered by linear least squares [9].

For two motions, \mathbf{B} is a 6×6 upper triangular matrix that takes the following form:

$$\mathbf{B} = \begin{bmatrix} a & b & c & & & \\ & d & e & & & \\ & & f & & & \\ & & & a' & b' & c' \\ & & & & d' & e' \\ & & & & & 1 \end{bmatrix} \quad (12)$$

to prevent interaction between the objects. For independent motions, $f=1$ (the same as calibrating the objects independently) since we cannot constrain the scale factors to be equal due to the depth/scale ambiguity (unless projection is known to be truly orthographic).

In contrast, for articulated objects it is sensible to assume that there is no significant difference in depth between two bodies as they are attached to each other. As a result, we constrain the two motion matrices to be equal in scale such that f is a free parameter.

Furthermore, for objects joined by a hinge, we have the additional constraint that the motions share a common axis whereby \mathbf{B} takes the form:

$$\mathbf{B} = \begin{bmatrix} a & b & c & b' & c' \\ & d & e & & \\ & & f & & \\ & & & d' & e' \\ & & & & 1 \end{bmatrix}. \quad (13)$$

However, this means that the constraints on $\mathbf{\Omega}$ are no longer linear. Therefore, we perform self-calibration on the motion matrix $[\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}'_1, \mathbf{c}'_2, \mathbf{c}'_3]$ using a matrix of the form given in Equation (12). We then rescale the upper 3×3 submatrix such that $a = a'$ and rearrange the elements to give the form shown in Equation (13).

Note that this is a sub-optimal solution that should be refined further using non-linear optimization techniques. We defer this for future work in order to demonstrate the accuracy that is achieved using purely linear methods.

3.1. Recovering lengths and angles

Premultiplying \mathbf{d} and \mathbf{d}' by \mathbf{B}^{-1} gives their metric equivalent such that underlying link lengths are recovered within the kinematic structure. For two bodies joined at a hinge, we choose the x -axis as the axis of rotation such that (with a slight abuse of notation) at a given frame:

$$[\mathbf{c}'_2 \quad \mathbf{c}'_3]_{2 \times 2} = [\mathbf{c}_2 \quad \mathbf{c}_3]_{2 \times 2} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (14)$$

QR decomposition of $[\mathbf{c}_2 \quad \mathbf{c}_3]^{-1} [\mathbf{c}'_2 \quad \mathbf{c}'_3]$ then gives a rotation matrix from which the angle at the joint, θ , can be recovered.

4. Results

4.1. Universal joint

Figure 4(left) shows a frame from the ‘Head’ sequence where a model head was coupled to a box by a ball and socket joint. Both the box and the head were rotated about the joint centre to recover structure and motion.

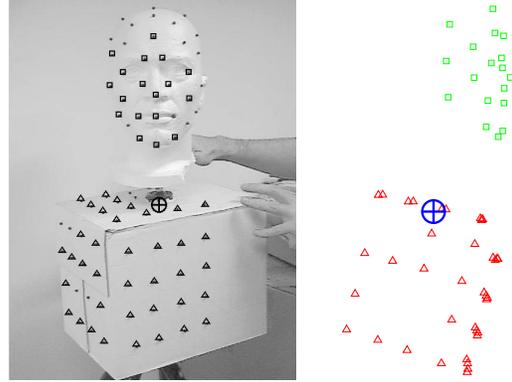


Figure 4. (left) ‘Head’ image with reprojected features and joint centre. (right) Recovered 3D structure and joint centre.

We see that the recovered joint centre, reprojected into the image, lies within a few pixels of the true joint centre. Visual examination of the recovered 3D structure in Figure 4(right) suggests that the location of the joint centre is indeed accurate.

4.2. Hinge joint

Similarly, Figure 1(left) shows a frame from the ‘Hinge’ sequence where two boxes were coupled by a hinge joint. Analysis of the recovered 3D structure (Figure 5) shows that the recovered axis lies close to the intersection of the two planes. We stress that edge information is not used in this method, nor do we compute homographies between planes in the scene.

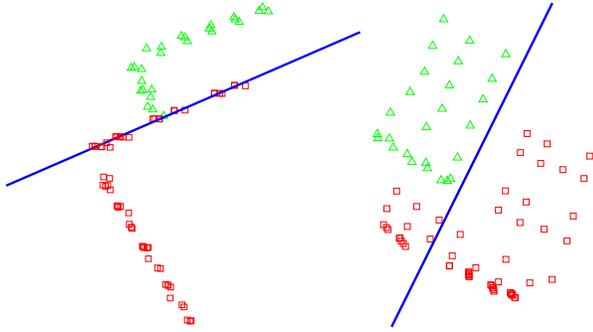


Figure 5. Recovered 3D features and rotation axis from ‘Hinge’ sequence.

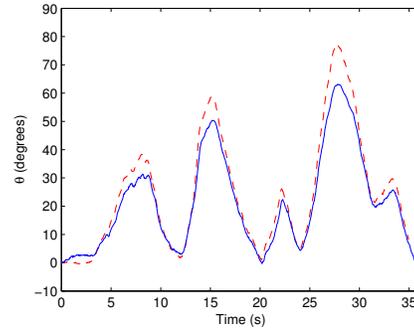


Figure 6. Recovered joint trajectories for two sequences.

4.3. Detecting articulated motion

Since articulated motion results in a drop in $rank(\mathbf{W})$, the singular values indicate the nature of any dependency. We used real image sequences of two bodies undergoing (i) independent motion, (ii) articulated motion at a universal joint and (iii) articulated motion at a hinge to compose \mathbf{W} and recover its singular values. Table 1 shows σ_6 , σ_7 and σ_8 (scaled such that $\sum \sigma = 1$) plus their ratios where we see that the type of articulation can be readily observed as a sharp drop in “effective rank”.

Table 1. Comparison of singular values for different motions

Dependency	$\sigma \times 10^3$			σ_6/σ_7	σ_7/σ_8
	σ_6	σ_7	σ_8		
None	4.9	4.4	3.0	1.11	1.46
Universal joint	6.1	4.4	0.7	1.39	6.28
Hinge	4.5	0.4	0.3	11.25	1.33

4.4. Recovering joint angles

We now demonstrate the recovery of the joint angle for the ‘hinge’ sequence. As we did not have ground truth, we computed the angle independently for two synchronized views of the same motion and compared the values recovered from each view (Figure 6). We see that there is a error in angle of up to 15° as a result of poorly constrained self-calibration due to limited motion of the base object. Despite this error, the recovered joint angles exhibit a clear correlation (a more rigorous evaluation of sensitivity follows in Section 4.5).

As an aside, we note that the signals in Figure 6 could potentially be used to synchronize two image sequences of

the same motion. However, specific synchronization methods exist that may be more appropriate [3, 16, 18].

4.5. Stability and Noise Sensitivity

The factorization algorithm is well known for its robustness, recovering the Maximum Likelihood solution in the presence of isotropic Gaussian noise [10]. We now show that the same applies in the articulated case.

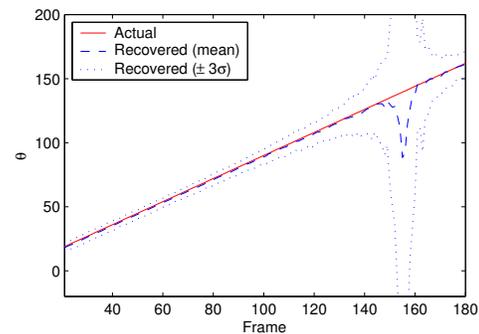


Figure 7. Distribution of joint angle error, over 100 trials, for noise level of $\sigma_n = 3$ pixels.

We generated a synthetic sequence of two bodies undergoing articulated motion and added zero-mean Gaussian noise of $\sigma_n \approx 3$ pixels (typical noise levels were measured as $\sigma_n \approx 1$ pixel for real sequences of a similar image size). From this sequence, the joint angle was recovered using the described methods and compared with ground truth.

Figure 7 shows that the distribution of error in the joint angle is typically small, increasing towards frame 155. At this point, a kinematic singularity occurred such that the affine depth ambiguity often resulted in the joint angle being under-estimated. This ambiguity has been a common cause of complaint in other studies of articulated motion [12].

4.6. Recovering link length

In a similar experiment, we generated a synthetic sequence of a three-link chain where both axes of rotation were parallel. For this sequence, we applied a modified version of the method for longer kinematic chains with parallel axes of rotation. The length of the middle link was then computed as the distance between the two recovered axes. Since affine projection means that structure and motion can only be recovered up to a global scale, we assume orthographic projection to compare the recovered length with its ground truth value of 134.2 units.

Table 2. Distribution of link length error over 100 trials.

	σ_n				
	0	1	2	3	4
Mean length	134.2	134.4	134.8	136.1	138.0
Std. dev.	0.000	0.989	2.420	3.327	4.412

Table 2 shows the distribution of error (over 100 trials) for varying levels of image noise. We see that the recovered length is typically close to the correct value, even at relatively high levels of noise (typical noise levels for a real sequence were measured as $\sigma_n \approx 1$ pixel).

5. Discussion

We have developed the factorization method [14] for articulated objects, showing that the rank of a matrix of feature tracks indicates the type of joint present and that joint centres and axes of rotation can be recovered using straightforward linear methods. Self-calibration was discussed and results presented for synthetic data (for comparison with ground truth) and real image sequences. Several avenues for further development remain:

- Investigation of other dependencies (e.g. prismatic joints) and degeneracies (e.g. planar structure);
- Generalization to direct methods (i.e. optic flow) so that the method may be applied to more general sequences (e.g. video of humans, although obtaining sufficient flow on a person may prove very difficult with current camera technology);
- Using the recovered axes and joint centres to define an object-centred coordinate frame for constrained 3D registration of non-overlapping feature sets;
- Comparison between statistical shape models [1, 2] and articulated SFM.

Acknowledgments

The authors would like to thank Fred Schaffalitzky for a fruitful discussion on self-calibration and the reviewers for their constructive comments. PT is supported by the Engineering and Physical Sciences Research Council (EPSRC).

References

- [1] M. Brand. Morphable 3D models from video. In *CVPR*, 2001.
- [2] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3D shape from image streams. In *CVPR*, 2000.
- [3] Y. Caspi and M. Irani. A step towards sequence-to-sequence alignment. In *CVPR*, 2000.
- [4] J. Costeira and T. Kanade. A multi-body factorization method for motion analysis. In *ICCV*, 1995.
- [5] T. Drummond and R. Cipolla. Real-time visual tracking of complex structures. *TPAMI*, 24(7):932–946, 2002.
- [6] A. Fitzgibbon and A. Zisserman. Multibody structure and motion : 3-D reconstruction of independently moving objects. In *ECCV*, 2000.
- [7] D. M. Gavrila. The visual analysis of human movement: A survey. *CVIU*, 73(1):82–98, 1999.
- [8] A. Gruber and Y. Weiss. Multibody factorization with uncertainty and missing data using the EM algorithm. In *CVPR*, 2004.
- [9] L. Quan. Self-calibration of an affine camera from multiple views. *IJCV*, 19(1):93–110, 1996.
- [10] I. D. Reid and D. W. Murray. Active tracking of foveated feature clusters using affine structure. *International Journal of Computer Vision*, 18(1):41–60, Apr. 1996.
- [11] D. Sinclair, L. Paletta, and A. Pinz. Euclidean structure recovery through articulated motion. In *Proc. 10th Scandinavian Conference on Image Analysis, Lappeenranta, Finland*, pages 991–998, 1997.
- [12] C. Sminchisescu and B. Triggs. Kinematic jump processes for monocular 3D human tracking. In *CVPR*, 2003.
- [13] C. J. Taylor. Reconstruction of articulated objects from point correspondences in a single uncalibrated image. *CVIU*, 2000.
- [14] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: A factorization approach. *IJCV*, 9(2):137–154, 1992.
- [15] P. H. S. Torr and D. W. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *IJCV*, 24(3):271–300, 1997.
- [16] P. Tresadern and I. Reid. Synchronizing image sequences of non-rigid objects. In *BMVC*, 2003.
- [17] R. Vidal and R. Hartley. Motion segmentation with missing data using PowerFactorization and GPCA. In *CVPR*, 2004.
- [18] L. Wolf and A. Zomet. Correspondence-free synchronization and reconstruction in a non-rigid scene. In *Proc. Workshop on Vision and Modelling of Dynamic Scenes*, 2002.
- [19] L. Zelnik-Manor and M. Irani. Degeneracies, dependencies and their implications in multi-body and multi-sequence factorization. In *CVPR*, 2003.