

Graph Rigidity for Near-Coplanar Structure from Motion

Jack Valmadre^{*†} Ben Upcroft[†] Sridha Sridharan[†] Simon Lucey^{*†}

^{*}ICT Centre

CSIRO

1 Technology Ct, Brisbane, Australia 4069

{jack.valmadre, simon.lucey}@csiro.au

[†]School of Engineering Systems

Queensland University of Technology

2 George St, Brisbane, Australia 4000

{ben.upcroft, s.sridharan}@qut.edu.au

Abstract—Recent algorithms for monocular motion capture (MoCap) estimate weak-perspective camera matrices between images using a small subset of approximately-rigid points on the human body (i.e. the torso and hip). A problem with this approach, however, is that these points are often close to coplanar, causing canonical linear factorisation algorithms for rigid structure from motion (SFM) to become extremely sensitive to noise. In this paper, we propose an alternative solution to weak-perspective SFM based on a convex relaxation of graph rigidity. We demonstrate the success of our algorithm on both synthetic and real world data, allowing for much improved solutions to markerless MoCap problems on human bodies. Finally, we propose an approach for solving the two-fold ambiguity over bone direction using a k -nearest neighbour kernel density estimator.

I. INTRODUCTION

An estimate of 3D human pose obtained from a monocular image can be used to initialise an articulated tracker (e.g. [1]–[3]), inject rigid keyframes into non-rigid SFM algorithms (e.g. [4]) or train improved parts-based detectors (e.g. [5]). In this paper, we consider the problem of estimating 3D human pose given the projected 2D position of each joint in several non-consecutive images. These point correspondences may be manually annotated or automatically generated using methods such as [5], [6]. We do not assume any prior on the skeletal dimensions, except that they do not vary between images.

Wei and Chai [7] recently proposed an approach for solving non-rigid structure from motion (NRSFM) specifically for human bodies, taking advantage of empirically validated rigid constraints in the torso and hip. Their approach is notable in comparison to previous literature in the area [9]–[11] as it i) makes no assumptions about bone lengths, ii) is not limited/constrained to modeling shapes previously seen in a training set, and iii) can handle missing body points.

Subsequent work by Valmadre and Lucey [12] demonstrated that the non-linear minimisation in the original algorithm could be replaced by the factorisation algorithm of Tomasi and Kanade [8], which is much faster and not prone to local optima, since it only requires an SVD and a small quadratic program. Unfortunately, when the observed structure is coplanar the decomposition is not unique even up to a 3×3 affine transform, supporting Ullman’s theorem

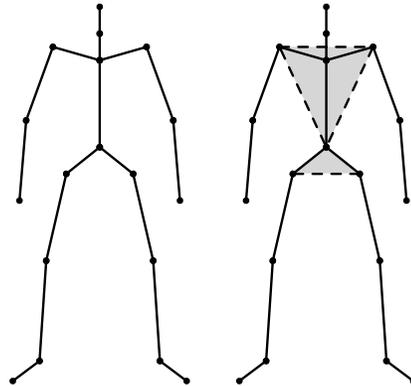


Figure 1: Wei and Chai [7] recognised that the human torso and hip are approximately rigid. These rigid substructures are essential for estimating the camera matrices of 3D non-rigid human structure from 2D point projections. Unfortunately, in practice these rigid substructures are often close to coplanar making them sensitive under the canonical Tomasi and Kanade factorisation method [8].

that at least three views and four non-coplanar points are required [13]. In practice, the vector corresponding to the third singular value will be determined by noise and an affine metric upgrade will not be sufficient.

In the area of perspective geometry, Li [14] recently presented a seminal algorithm for reconstructing a 3D scene from 2D point correspondences without having to estimate camera matrices. The approach treats the structure as a rigid 3D graph, first finding the lengths of the edges, then solving for the vertex positions using graph realisation, also known as graph embedding or multi-dimensional scaling. Both stages are expressed as a convex optimisation using the trace heuristic for matrix rank. A drawback of their algorithm, however, is a requirement that the cameras are already calibrated.

This paper contributes an algorithm which extends the work of [14] to a weak-perspective camera. Our work differs fundamentally to Li’s, however, in a number of ways. Firstly, in Li’s work the camera centres themselves are nodes in the graph, whereas in our weak-perspective extension, the graph

consists solely of the observed structure. As a direct result, we do not require camera calibration to be known *a priori*. Secondly, ours is the first work to characterise theoretically and empirically the robustness that graph rigidity SFM exhibits to coplanarity. We show that our algorithm degrades gracefully rather than catastrophically in the limit of coplanarity. However, in situations where the structure has considerable variance in three dimensions, canonical SFM remains a more accurate and computationally attractive option.

II. TOMASI AND KANADE'S SFM

Tomasi and Kanade's algorithm involves the rank-3 factorisation of a measurement matrix $\mathbf{W} \in \mathcal{R}^{2F \times N}$ into a rotation matrix $\mathbf{R} \in \mathcal{R}^{2F \times 3}$ and a structure matrix $\mathbf{S} \in \mathcal{R}^{3 \times N}$. Using an SVD to factorise $\mathbf{W} = \hat{\mathbf{R}}\hat{\mathbf{S}}$ minimises

$$\|\mathbf{W} - \hat{\mathbf{R}}\hat{\mathbf{S}}\|_F^2, \quad (1)$$

corresponding to the maximum likelihood estimator under an assumption of Gaussian projection noise. Metric reconstruction (up to a similarity transform) is obtained by finding $\mathbf{G} \in \mathcal{R}^{3 \times 3}$,

$$\mathbf{W} = \hat{\mathbf{R}}\hat{\mathbf{S}} = \hat{\mathbf{R}}\mathbf{G}\mathbf{G}^{-1}\hat{\mathbf{S}} = \mathbf{R}\mathbf{S}, \quad (2)$$

that optimally restores the orthogonality and aspect ratio of the rotation matrices.

When $\text{rank}(\mathbf{W}) < 3$, the factorisation is not unique, even up to an affine transformation. For a set of coplanar points, the structure matrix \mathbf{S} , and therefore the measurement matrix \mathbf{W} , is at most rank-2. In this case, the true rotation and structure are not necessarily given by an affine transformation of $\hat{\mathbf{R}}$ and $\hat{\mathbf{S}}$.

Poelman and Kanade [15] extended the factorisation algorithm to weak perspective projection. They also proposed a linear least-squares solution for the affine correction matrix by instead solving for $\mathbf{Q} = \mathbf{G}\mathbf{G}^T$. If the resulting \mathbf{Q} matrix is not positive semidefinite, then an imaginary \mathbf{G} matrix results and the method fails. Novins and Arvo [16] suggested that semidefinite programming be used to constrain $\mathbf{Q} \succeq 0$.

III. WEI AND CHAI'S SFM

Wei and Chai [7] recently presented a method for solving for the 3D pose of a human body given point correspondences between uncalibrated, monocular images. Their approach was originally employed to solve non-rigid SFM. Recently, however, it was demonstrated [12] that their approach is only applicable to rigid SFM problems. Wei and Chai effectively reformulated the task of SFM using Pythagoras' theorem. Assume we have a set of B edges (or "bones") \mathcal{B} which define a rigid graph. Let ℓ_i be the length of edge i , q_i^t be the projected length of edge i in frame t , s_t be the weak perspective scale of frame t and z_i^t be the relative out-of-plane depth across edge i in frame t ,

$$\ell_i^2 - (q_i^t)^2 s_t^{-2} - (z_i^t)^2 = 0, \quad \forall t \in \mathcal{F}, i \in \mathcal{B}. \quad (3)$$

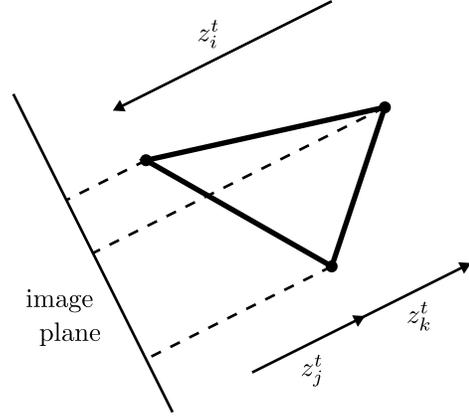


Figure 2: Wei and Chai [7] enforce rigidity by constraining the sum of the depths around a closed rigid triangle to be zero for each frame. Since the sign of the depths is unknown, they seek an expression in the squared depths.

This is an under-constrained system of homogeneous linear equations $\mathbf{A}\mathbf{x} = 0$, with

$$\mathbf{x} = \begin{bmatrix} \ell \\ \mathbf{s} \\ \mathbf{z} \end{bmatrix}, \quad \ell = \begin{bmatrix} \ell_1^2 \\ \vdots \\ \ell_B^2 \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} s_1^{-2} \\ \vdots \\ s_F^{-2} \end{bmatrix},$$

$$\mathbf{z} = \begin{bmatrix} \mathbf{z}_1 \\ \vdots \\ \mathbf{z}_F \end{bmatrix}, \quad \mathbf{z}_t = \begin{bmatrix} (z_1^t)^2 \\ \vdots \\ (z_B^t)^2 \end{bmatrix}.$$

Defining a set of rigid triangles \mathcal{T} , rigidity is enforced by requiring that the sum of the depths around a closed triangle must be zero (see Figure 2),

$$z_i^t + z_j^t + z_k^t = 0 \quad (4)$$

$$(z_i^t)^2 = (z_j^t)^2 + 2z_j^t z_k^t + (z_k^t)^2 \quad (5)$$

$$[(z_i^t)^2 - (z_j^t)^2 - (z_k^t)^2]^2 = 4(z_j^t)^2 (z_k^t)^2. \quad (6)$$

Each rigid triangle $(i, j, k) \in \mathcal{T}$ comprises an extra length e_i and the lengths of two existing bones ℓ_j and ℓ_k . Substituting (3) into (6) gives a system of quadratic equations in the squared variables, with \mathbf{x} augmented to include the extra lengths e_i ,

$$[e_i^2 - (q_i^t)^2 s_t^{-2} - (z_j^t)^2 - (z_k^t)^2]^2 = 4(z_j^t)^2 (z_k^t)^2 \quad (7)$$

$$\mathbf{x}^T \mathbf{C}_{ijk}^t \mathbf{x} = 0, \quad \forall t \in \mathcal{F}, (i, j, k) \in \mathcal{T}. \quad (8)$$

The overall problem can therefore be expressed as,

$$\underset{\ell, \mathbf{e}, \mathbf{s}, \mathbf{z}}{\text{minimise}} \quad \|\mathbf{A}\mathbf{x}\|_2^2 + \lambda \sum_{t \in \mathcal{F}} \sum_{(i, j, k) \in \mathcal{T}} (\mathbf{x}^T \mathbf{C}_{ijk}^t \mathbf{x})^2 \quad (9)$$

subject to $\mathbf{x} \succeq 0, \mathbf{1}^T \ell = 1$.

The parameter λ trades off the penalty for non-rigidity against projection error and must be chosen heuristically.

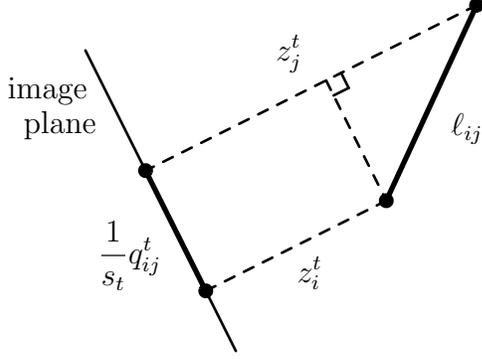


Figure 3: Weak perspective projection defines a right-angled triangle. This simple insight can be employed to pose rigid SFM as a graph rigidity problem that unlike Li [14] does not require calibrated cameras.

Unfortunately this optimisation is not convex due to the introduction of the quartic rigidity term, and in practice is prone to local minima and lacks a deterministic solution, unlike traditional Tomasi and Kanade factorisation.

IV. GRAPH RIGIDITY SFM

Although it was never presented as such, Wei and Chai’s SFM method is an implicit application of graph rigidity theory. A graph is rigid if the position of its vertices can be determined uniquely (up to a Euclidean transform) from the length of its edges. For a 3D graph with N vertices to be rigid, it must have at least $3N - 6$ “properly-distributed” edges [14], which corresponds to a maximally connected graph for a triangle or tetrahedron. In this section we will present a convex relaxation for weak-perspective rigid SFM based on graph rigidity.

A. Posing SFM as Graph Embedding

Assume we have a set of edges defining a rigid graph, \mathcal{E} . This time let ℓ_{ij} be the length of edge (i, j) , s_t be the scale of frame t , q_{ij}^t be the projected length of edge (i, j) in frame t and z_i^t be the distance of point i from the image plane, relative to an arbitrary origin. The expression for projection error is constructed using Pythagoras’ theorem (see Figure 3). This formulation is similar to that of [7], with the exception that the z components are expressed in terms of point position not edge depth, implicitly encoding the rigidity in a single objective function,

$$E = \sum_{t \in \mathcal{F}} \sum_{(i,j) \in \mathcal{E}} [\ell_{ij}^2 - (q_{ij}^t)^2 s_t^{-2} - (z_i^t - z_j^t)^2]^2 \quad (10)$$

$$E = \sum_{t \in \mathcal{F}} \sum_{(i,j) \in \mathcal{E}} [(\mathbf{b}_{ij})^T \boldsymbol{\ell} - (\mathbf{c}_{ij}^t)^T \mathbf{s} - \mathbf{z}_t^T \mathbf{D}_{ij} \mathbf{z}_t]^2 \quad (11)$$

$$\boldsymbol{\ell} = \begin{bmatrix} \ell_{(1)}^2 \\ \vdots \\ \ell_{(E)}^2 \end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix} s_1^{-2} \\ \vdots \\ s_F^{-2} \end{bmatrix}, \quad \mathbf{z}_t = \begin{bmatrix} z_1^t \\ \vdots \\ z_N^t \end{bmatrix}.$$

Substituting $\mathbf{Z}_t = \mathbf{z}_t \mathbf{z}_t^T$ gives a linear least-squares objective and a non-convex rank constraint.

$$E = \sum_{t \in \mathcal{F}} \sum_{(i,j) \in \mathcal{E}} [(\mathbf{b}_{ij})^T \boldsymbol{\ell} - (\mathbf{c}_{ij}^t)^T \mathbf{s} - \text{tr}(\mathbf{D}_{ij} \mathbf{Z}_t)]^2 \quad (12)$$

$$E = \|\mathbf{A}\mathbf{x}\|_2^2, \quad \mathbf{x} = \begin{bmatrix} \boldsymbol{\ell} \\ \mathbf{s} \\ \mathbf{z} \end{bmatrix}, \quad \mathbf{z} = \begin{bmatrix} \text{vec}(\mathbf{Z}_1) \\ \vdots \\ \text{vec}(\mathbf{Z}_F) \end{bmatrix} \quad (13)$$

$$\begin{aligned} & \text{minimise}_{\boldsymbol{\ell}, \mathbf{s}, \mathbf{Z}_t} \|\mathbf{A}\mathbf{x}\|_2^2 \\ & \text{subject to} \quad \text{rank}(\mathbf{Z}_t) \leq 1, \\ & \quad \boldsymbol{\ell} \succeq 0, \mathbf{s} \succeq 0, \mathbf{1}^T \boldsymbol{\ell} = 1. \end{aligned} \quad (14)$$

B. Convex Relaxation

Fazel [17] presents the result that the nuclear norm of a matrix $\|\cdot\|_*$ is the convex envelope of its rank. For a positive semidefinite matrix,

$$\|\mathbf{A}\|_* = \text{tr}(\mathbf{A}) = \sum \lambda(\mathbf{A}) = \|\lambda(\mathbf{A})\|_1. \quad (15)$$

Therefore minimising the trace of a positive semidefinite matrix minimises the ℓ_1 -norm of its eigenvalues, which is likely to find a matrix which has a sparse vector of eigenvalues (i.e. is low rank). We propose the following convex optimisation, inspired by the work of Li [14],

$$\begin{aligned} & \text{minimise}_{\boldsymbol{\ell}, \mathbf{s}, \{\mathbf{Z}_t\}_{t \in \mathcal{F}}} \sum_{t \in \mathcal{F}} \text{tr}(\mathbf{Z}_t) \\ & \text{subject to} \quad \mathbf{Z}_t \succeq 0, \quad \forall t \in \mathcal{F}, \\ & \quad \mathbf{A}\mathbf{x} = 0, \quad \forall t \in \mathcal{F}, (i, j) \in \mathcal{E}, \\ & \quad \boldsymbol{\ell} \succeq 0, \mathbf{s} \succeq 0, \mathbf{1}^T \boldsymbol{\ell} = 1. \end{aligned} \quad (16)$$

Each vector of depths \mathbf{z}_t is recovered up to a sign ambiguity using the SVD to find a rank-1 factorisation of \mathbf{Z}_t .

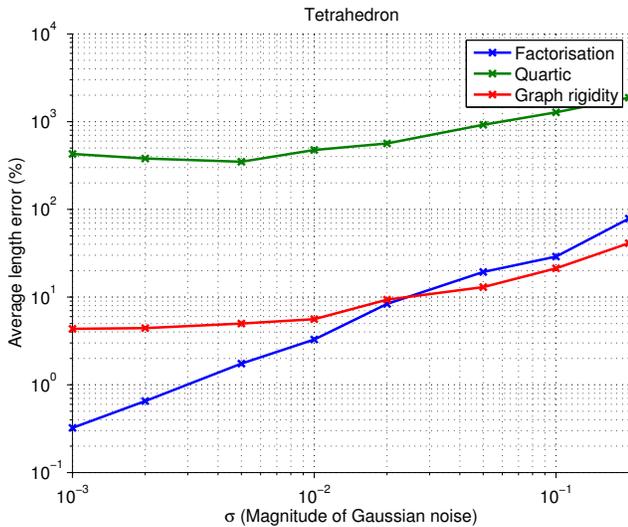
Note that the projections are now enforced by an equality constraint. While it is possible to keep this term in the objective and minimise the residual of the projection equations, we have found that the relaxation is too weak, and the solution tends towards \mathbf{Z}_t matrices with small, non-sparse eigenvalues.

V. SYNTHETIC EVALUATION OF RIGID SFM

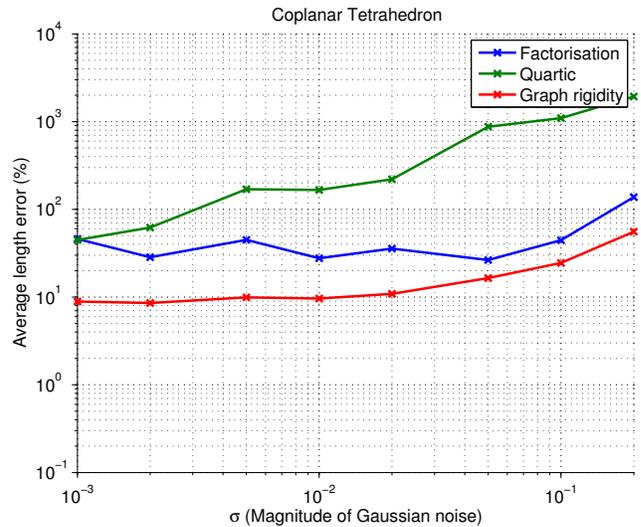
A. Experiment

To test the validity of our claims concerning rigid SFM when dealing with a small number of points that are close to coplanar we generated a series of synthetic experiments.

Random scenes containing four points and five weak-perspective cameras (with scale varying by a factor of 3) were constructed, then Gaussian noise of increasing magnitude was added to each of the projections. A random



(a) Experiment with non-coplanar structure



(b) Experiment with coplanar structure

Figure 4: Median reconstruction error for the three algorithms described. When non-coplanar structure is observed, the reconstruction error using Tomasi-Kanade factorisation decreases steadily with decreasing noise, whereas our proposed algorithm does not recover a perfect reconstruction even for non-coplanar data with noiseless projections due to the relaxation. However, it is much more robust to coplanarity than factorisation. Although Wei and Chai’s quartic objective sometimes achieves accurate reconstructions, the median error is extremely large due to its high susceptibility to local minima.

point in each of the scenes was then projected on to the plane defined by the other three, simulating a coplanar tetrahedron characteristic of the human torso, and identical noise was added to the projections. Comparing the behaviour of reconstruction error with respect to increasing noise between the coplanar and non-coplanar scenes shows how robust (or otherwise) an SFM algorithm is to coplanarity. Random points were generated uniformly in $(-0.5, 0.5)$ and random rotation matrices were generated using [18].

Reconstruction error was measured by the average edge length error, giving a rotation-invariant metric. To compare lengths to ground truth, they had to be scaled into and back out of the camera frame. For example, if a reconstruction algorithm estimates the scales to be \hat{s}_t and the ground truth scales are s_t , then the error between the estimated length \hat{l}_{ij} and true length l_{ij} is

$$\varepsilon_{ij} = \frac{1}{F} \sum_{t \in \mathcal{F}} \left| \frac{1}{s_t} (\hat{s}_t \hat{l}_{ij}) - l_{ij} \right|. \quad (17)$$

B. Results

Experimental results are presented in Figure 4.

Tomasi and Kanade factorisation achieves perfect reconstruction for non-coplanar points when there is no noise. Performance deteriorates relatively quickly for increasing noise because when there are only four points, \mathbf{W} is exactly rank-3 and the SVD does not remove any noise. When faced with coplanar data, the metric reconstruction can only find

a poor solution due to the non-unique decomposition of \mathbf{W} into $\hat{\mathbf{R}}$ and $\hat{\mathbf{S}}$.

Wei and Chai’s SFM occasionally finds a good solution but generally seems to end up in a local minima. It is probably possible to achieve better convergence by tweaking the regularisation constant and the initialisation value, but this sensitivity is exactly what makes the quartic optimisation algorithm unattractive anyway.

Our convex graph rigidity method does not achieve zero error even when there is no noise. This is because of the “softness” of the relaxation. However, it is clearly more robust to coplanarity than Tomasi and Kanade’s SFM, while still providing a deterministic solution. The loss in accuracy in using our algorithm when there is low projection noise and non-coplanar structure suggests that an adaptive strategy could be employed, where the magnitudes of the first three singular values are compared to decide which algorithm to use.

VI. APPLICATION TO HUMAN BODIES

A. Free Bones

After solving for the weak-perspective scales and the structure of the torso using our graph-based algorithm, we adopt the approach of Valmadre and Lucey [12], approximating the lengths of the remaining “free” bones (arms, legs

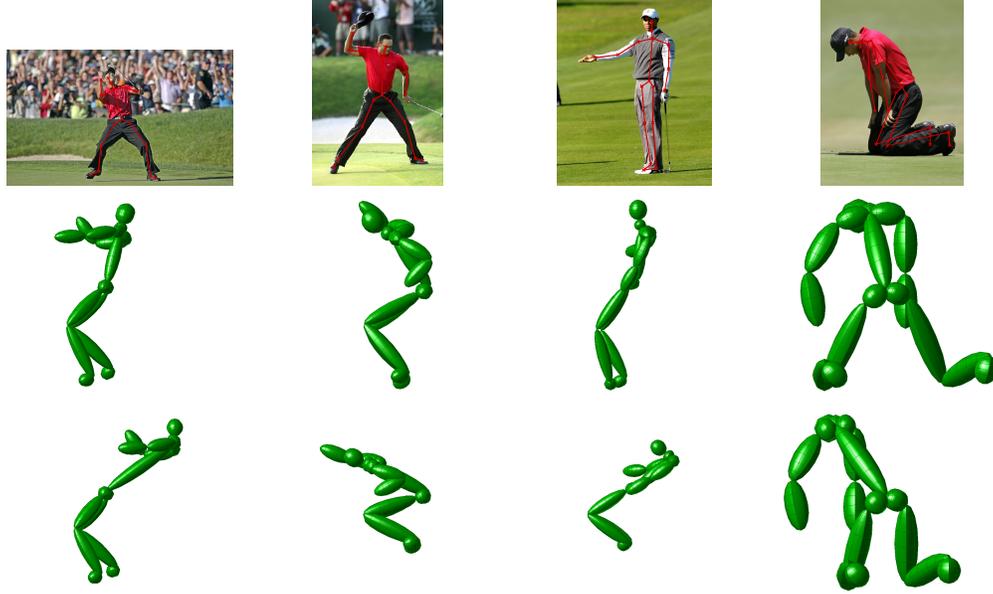


Figure 5: Four labeled images of golfer Tiger Woods from a set of six (top). Solution for human pose using our relaxed graph rigidity method (middle). The same solution applied using Tomasi and Kanade’s SFM to solve for the rigid sub-problem (bottom). The weak-perspective scale critically affects the length of the free bones due to Equation 18.

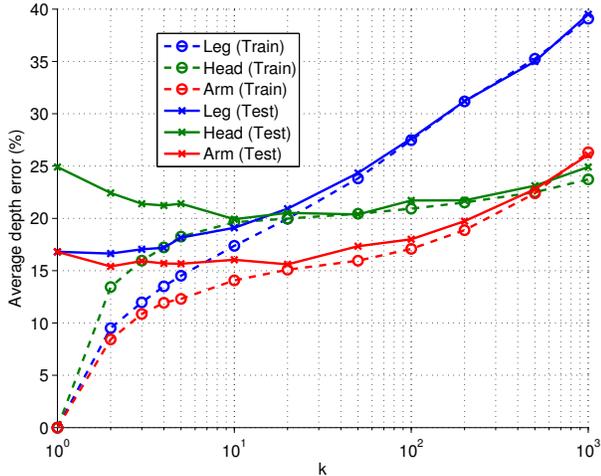


Figure 6: Testing and training error for different values of k . Given all possible discrete depth solutions for an example pose, the classifier chooses the most likely one. It is possible to achieve an error rate of 15-25% using small values of k .

and neck) by their maximum observed projection.

$$\ell_{ij} \approx \max_{t \in \mathcal{F}} \left(\frac{1}{s_t} q_{ij}^t \right) \quad (18)$$

However, there still exists a two-fold sign ambiguity for each free bone with non-zero depth.

B. Discrete Sign Ambiguities

Wei and Chai [7] used a second local optimisation stage to minimise the projection error of the forward kinematic chain subject to joint angle constraints drawn from biomechanics. Two disadvantages of this method are that it is not globally optimal, being highly dependent on initialisation, and that it will only present the user with a single random solution, while there may be several which satisfy the constraints.

Valmadre and Lucey [12] pointed out that by treating each rigid linkage independently, the full discrete solution space could be easily enumerated. They evaluated the inverse kinematics for each solution and eliminated those which violate joint angle limits. A weakness of this method is that when there are several valid solutions, it can be time consuming to choose the correct one.

A critical downfall of both of these methods is that the joint angle limits are too hard. If a solution is only marginally outside the limits, it may be mistakenly eliminated. This is exasperated by the non-linear nature of the inverse kinematics solution. For example, when an elbow is almost straight, the angle of twist at the shoulder becomes very sensitive to small reconstruction errors.

C. Most Likely Pose

To avoid these issues, we propose a method to find the *most likely pose* given a set of example poses. This is posed as a k -nearest neighbour problem, where distance is measured as the sum of Euclidean distances between each joint in a limb. Rotation is normalised with respect to the

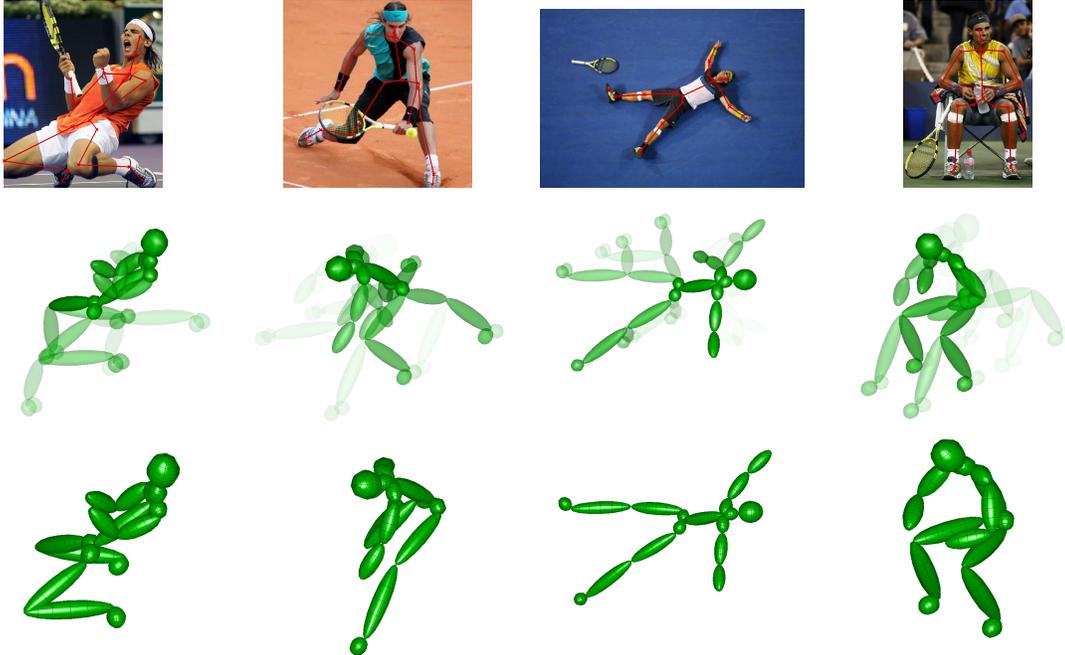


Figure 7: Four labeled images of tennis player Rafael Nadal from a set of nine (top). The discrete solution space after fixing the direction of the torso and hips, with more likely configurations rendered more opaque (middle). The solution chosen by a human user (bottom).

torso and bone lengths are re-scaled to their average. Of the discrete set of solutions, the pose with the minimum-distance k -nearest neighbour is determined to be the most likely. Separate models were trained for arms, legs and the neck. Left and right training examples were mirrored to leverage the known symmetry of the human body.

Our k NN model was trained on 10^4 examples from the CMU MOCAP database,¹ using an 80/20 split to choose k by cross-validation (see Figure 6). Qualitative results of application to a real-world problem are shown in Figure 7, using $k = 3$.

VII. CONCLUSION

We have presented a novel solution to rigid SFM using a convex relaxation of graph rigidity. The algorithm has been empirically shown to be more robust to situations involving a small number of coplanar points, both quantitatively on synthetic datasets and qualitatively on real world problems using the human body. We acknowledge and experimentally confirm that Tomasi and Kanade’s factorisation algorithm is more accurate when the observed structure is not close to coplanar. Additionally, we demonstrate a simple method for choosing the most likely of a set of poses, which greatly increases the efficiency of manually labelling frames for markerless monocular MoCap.

¹<http://mocap.cs.cmu.edu/>

REFERENCES

- [1] M. Yamamoto and K. Koshikawa, “Human motion analysis based on a robot arm model,” in *CVPR*, no. 6. IEEE, 1991, pp. 664–665.
- [2] J. M. Rehg and T. Kanade, “Visual tracking of high DOF articulated structures: An application to human hand tracking,” in *ECCV*. Springer, 1994, pp. 35–46.
- [3] C. Bregler and J. Malik, “Tracking people with twists and exponential maps,” in *CVPR*. IEEE, 2002, pp. 8–15.
- [4] Y. Zhu, M. Cox, and S. Lucey, “3D Motion Reconstruction for Real-World Camera Motion,” in *CVPR*. IEEE, 2011.
- [5] L. Bourdev and J. Malik, “Poselets: Body part detectors trained using 3D human pose annotations,” in *ICCV*. IEEE, Sep. 2009, pp. 1365–1372.
- [6] G. Mori and J. Malik, “Estimating human body configurations using shape context matching,” in *ECCV*, vol. 2352. Springer, 2002, pp. 150–180.
- [7] X. Wei and J. Chai, “Modeling 3D human poses from uncalibrated monocular images,” in *ICCV*. IEEE, 2009.
- [8] C. Tomasi and T. Kanade, “Shape and Motion from Image Streams under Orthography: A Factorization Method,” *IJCV*, vol. 9, no. 2, pp. 137–154, Nov. 1992.
- [9] C. J. Taylor, “Reconstruction of articulated objects from point correspondences in a single uncalibrated image,” in *CVPR*, vol. 80, no. 3. IEEE, Dec. 2000, pp. 677–684.

- [10] C. Barrón and I. A. Kakadiaris, “Estimating Anthropometry and Pose from a Single Uncalibrated Image,” *Computer Vision and Image Understanding*, vol. 81, no. 3, pp. 269–284, Mar. 2001.
- [11] C. Bregler, A. Hertzmann, and H. Biermann, “Recovering non-rigid 3D shape from image streams,” in *CVPR*, vol. 2. IEEE, 2000, pp. 690–696.
- [12] J. Valmadre and S. Lucey, “Deterministic 3D Human Pose Estimation Using Rigid Structure,” in *ECCV*. Springer, 2010, pp. 467–480.
- [13] S. Ullman, “The interpretation of structure from motion,” *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 203, no. 1153, pp. 405–426, 1979.
- [14] H. Li, “Multi-View Structure Computation without Explicitly Estimating Motion,” in *CVPR*. IEEE, 2010.
- [15] C. Poelman and T. Kanade, “A paraperspective factorization method for shape and motion recovery,” *PAMI*, vol. 19, no. 3, pp. 206–218, Mar. 1997.
- [16] K. Novins and J. Arvo, “An Exploration of Semidefinite Programming Applied to Rigid Motion Factorization,” in *Image and Vision Computing New Zealand*, 2005.
- [17] M. Fazel, “Matrix rank minimization with applications,” Ph.D. dissertation, Stanford University, 2002.
- [18] J. Arvo, “Fast Random Rotation Matrices,” 1992.