

Structure from Recurrent Motion: From Rigidity to Recurrency

Xiu Li^{1,2} Hongdong Li^{2,3} Hanbyul Joo² Yebin Liu¹ Yaser Sheikh²

Tsinghua University¹ Carnegie Mellon University² Australian National University³

Abstract

This paper proposes a new method for Non-Rigid Structure-from-Motion (NRSfM) from a long monocular video sequence observing a non-rigid object performing recurrent and possibly repetitive dynamic action. Departing from the traditional idea of using linear low-order or low-rank shape model for the task of NRSfM, our method exploits the property of shape recurrency (i.e., many deforming shapes tend to repeat themselves in time). We show that recurrency is in fact a generalized rigidity. Based on this, we reduce NRSfM problems to rigid ones provided that certain recurrency condition is satisfied. Given such a reduction, standard rigid-SfM techniques are directly applicable (without any change) to the reconstruction of non-rigid dynamic shapes. To implement this idea as a practical approach, this paper develops efficient algorithms for automatic recurrency detection, as well as camera view clustering via a rigidity-check. Experiments on both simulated sequences and real data demonstrate the effectiveness of the method. Since this paper offers a novel perspective on rethinking structure-from-motion, we hope it will inspire other new problems in the field.

1. Introduction

Structure-from-Motion (SfM) has been a success story in computer vision. Given multiple images of a rigidly moving object, one is able to recover the 3D shape of the object as well as camera locations by using geometrical multi-view constraints. Recent research focus in SfM has been extended to the reconstruction of non-rigid dynamic objects or scenes from multiple images, leading to “Non-Rigid Structure from Motion” (or NRSfM in short).

Despite remarkable progresses made in NRSfM, existing methods suffer from serious limitations. Notably, they often assume simple linear models, either over the non-rigid shape [7] or over motion trajectories [1], or both [9]. These linear models, while they are useful for characterizing certain classes of deforming objects (e.g, face, human pose, or clothing), are unable to capture a variety of dynamic objects

in rapid deformation, which are however common in reality.

This paper presents a new method for non-rigid structure from motion. Contrary to the traditional wisdom for NRSfM, we do not make a linear model assumption. Instead, we describe how to exploit shape *recurrency* for the task of non-rigid reconstruction. Specifically, we observe that in our physical world many deforming objects (and their shapes) tend to repeat themselves from time to time, or even only occasionally. In the context of SfM if a shape reoccurs in the video we say it is *recurrent*.

This observation of recurrency enables us to use the existing knowledge of multi-view geometry to reconstruct a shape. Given a video sequence, if one is able to recognize a shape that was seen before, then these two instances of images can be used as a virtual stereo pair of the same *rigid* object in a space. Therefore one can simply apply standard rigid SfM techniques to reconstruct a non-rigid object, without developing new methods. For instance, the techniques used in rigid SfM [12] such as the use of the fundamental matrix, computing camera poses by the Perspective-n-Point (PnP) algorithm, triangulating 3D points, bundle adjustment, and rigid factorization can be used without modifications. We conducted experiments on both synthetic and real data, showing the efficacy of our method.

2. The Key Insight

Rigidity is a fundamental property that underpins almost all work in rigid Structure-from-Motion (SfM). We say an object is *rigid* if its shape remains constant over time. For this reason, multiple images of the same object, taken from different viewpoints, can be viewed as redundant observations of the same target, making the task of rigid SfM mathematically well-posed and solvable. In contrast, the shape of a non-rigid object changes over time, violating the rigidity assumption and rendering NRSfM ill-posed.

In this paper, we show that *shape recurrency* is in fact a *generalized rigidity*. At first glance, finding rigid-pairs can be thought as a restrictive condition; however, satisfying this condition is far easier. Recurrent motions are ubiquitous in our surroundings, including human’s walking, animal’s running, leaves’ waving, clock’s pendulum sway-




Figure 1. This composite slow-motion picture of ‘figure-skating’ clearly illustrates the basic idea of our non-rigid SfM method. Despite the skater’s body poses kept changing dynamically over time, there were moments when she struck (nearly) identical posture, e.g., as indicated by the two red arrows and two blue arrows. Using a pair of such recurrent observations- albeit distant in time, one can reconstruct the 3D pose (shape) of the skater at that time instants, by using only standard rigid-SfM techniques.

ing, car wheels’ rotating, and so on. Many human motions such as martial arts, dance, and sport games contain various repetitive motions and patterns. Even dramatic or non-periodic motions can be included, as long as a visual observation is long enough, making it highly probable to revisit a previously-seen scene again. If we are given multiple sequences for the similar human motions, although they are not exactly the same scenes, it can increase the chances of finding recurrent motions.

This is the key insight of this paper. To further illustrate this idea, consider the example of figure skating in Fig. 1, showing a composite (strobe-type) photograph made by fusing multiple frames of slow-motion photos, which vividly captures the dynamic performance of the skater. Examine carefully each of the individual postures of the skater at different time steps; it is not difficult for one to recognize several (nearly) repeated poses.

To apply our idea of reconstructing non-rigid shapes from recurrence, we propose a novel method to formulate NRSfM problem as a graph-clustering problem, which can be solved by a Normalized-Cut [35] framework. In particular, we build a method to compute the probability representing the rigidity of a shape from images at two different time instances. The final recurrence relations are globally solved considering all connections in the constructed graph.

3. Problem Formulation and Main Algorithm

Consider a non-rigid dynamic 3D object observed by a moving pinhole camera, capturing N images at time steps of $t = 1, 2, \dots, N$. Our task is then to recover all the N temporal shapes of the object, $S(1), S(2), \dots, S(N)$. To be precise, the shape of the object at time t , $S(t)$, is defined by a set of M feature points (landmarks) on the object: $S(t) = [X_{t1}, X_{t2}, \dots, X_{tM}]$, where X_{ti} denotes the homogeneous coordinates of the i -th feature point of the object at

time t . Clearly the $S(t)$ is a $4 \times M$ matrix.

Given a pinhole camera with a projection matrix P , a single 3D point X is projected on the image at position x by a homogeneous equation $x \simeq PX$. For the shape of a temporally deforming object at time t , we have $x(t) \simeq P(t)S(t)$, where $x(t)$ denotes the image measurement of the shape $S(t)$ at time t , and $P(t)$ defines the camera matrix of the t -th frame.

By collecting all N frames of observations of the non-rigid object at time $t = 1, \dots, N$, we obtain the basic equation system for N -view M -point NRSfM problem:

$$\begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(N) \end{bmatrix} \simeq \begin{bmatrix} P(1) & & & \\ & P(2) & & \\ & & \ddots & \\ & & & P(N) \end{bmatrix} \cdot \begin{bmatrix} S(1) \\ S(2) \\ \vdots \\ S(N) \end{bmatrix}. \quad (1)$$

Definition 3.1 (Rigidity). Given two 3D shapes S and S' with correspondences in a space, we can say that they form a *rigid pair* if they are related by a rigid transformation T . Note that a rigid transformation can be compactly represented by a 4×4 matrix T , hence we have: $S' = TS$, $\exists T \in \mathbb{SE}(3)$.

We use $S \approx S'$ to denote that S and S' form a *rigid pair*.

Example 3.1 (Rigid Object). The shape of a rigid object remains constant all the time: $S(t) \approx S(t'), \forall t \neq t'$.

Example 3.2 (Periodic Deformation). A non-rigid object undergoing periodic deformation with period p will return to its previous shape after a multiplicity of periods, leading to $S(t) \approx S(t + kp), \forall k \in \mathbb{N}$.

Example 3.3 (Recurrent Object). A shape at time t re-occurs after some δ -time lapse: $S(t) \approx S(t + \delta)$.

3.1. Rigidity Check via Epipolar Geometry

If two 3D shapes (represented by point clouds) and their exact correspondences are given, checking whether they are rigidly related or not is a trivial task. However, this is not possible in the case of NRSfM where the shapes are not known *a priori*. All we have are two corresponding images of the shapes, and the rigidity-test has to be conducted based on the input images only.

In this paper, we use *epipolar-test* for this purpose. It is based on the well-known result of epipolar geometry: if two 3D shapes differ by only rigid Euclidean transformations, then, their two images must satisfy the *epipolar relationship*. Put it mathematically, we have $S \approx S' \Rightarrow \mathbf{x}'_i^\top \mathbf{F} \mathbf{x}_i = 0, \forall i$, where \mathbf{F} is the *fundamental matrix* between the two images for S and S' , respectively. Note that the RHS equation must be verified over all pairs of correspondences of $(\mathbf{x}_i, \mathbf{x}'_i), \forall i$.

Also note that satisfying epipolar relationship is only a *necessary condition* for two shapes S and S' to be rigid. This is because the epipolar relationship is invariant to any 4×4 projective transformation in 3-space. As a result, it is a weaker condition than the rigidity test, suggesting that even if two images pass the epipolar-test they still possibly be non-rigidly related. Fortunately, in practice, this is not a serious issue, because the *odds* that a generic dynamic object (with more than 5 landmark points) changes its shape precisely following a 15-DoF 3D projectivity is negligible. In other words, there is virtually no risk of mistaking.

The above idea of epipolar-test looks very simple. As such, one might be tempted to rush to implementing the following simple and straightforward algorithm:

1. Estimate a fundamental matrix from the correspondences using the linear 8-point algorithm;
2. Compute the mean residual error computed by averaging all the point-to-epipolar-line distances evaluated on key points in the image;
3. If this mean residual error is less than a pre-defined tolerance, return 'rigid', else return 'non-rigid'.

Unfortunately, despite the simplicity of the above algorithm, it is however not useful in practice, because of the following two reasons. (1) Ill-posed estimation: It is well known that linear methods for epipolar geometry estimation are very sensitive to outliers; a single outlier may destroy the fundamental matrix estimation. However, in our context, the situation is much worse (than merely having a few outliers). This is because, whenever the two feature point sets are in fact *not* rigidly related, forcing them to fit to a single fundamental matrix by using any linear algorithm can only yield a *meaningless* estimation, subsequently leading to a *meaningless* residual errors and unreliable decision. In

short, fitting all feature points to a single epipolar geometry is ill-posed. Instead, in order to do a proper rigidity-test one must consider the underlying 3D rigid-reconstructability of all these image points. (2) Degenerate cases: Even if two sets of points are indeed connected by a valid and meaningful fundamental matrix, there is no guarantee that a valid 3D reconstruction can be computed from the epipolar geometry. For example, when the camera is doing a pure rotation, there will not be enough disparity (parallax) in the correspondences to allow for a proper reconstruction—because the two cameras have only one center of projection—depth can not be observed. In such cases, the two sets of images can be mapped to each other by a planar homography, and the fundamental matrix estimations are non-unique.

Our solution: We propose a new algorithm for rigidity-test, named “Modified Epipolar Test”, which resolves both of the above issues. First, it uses (minimum) sub-set sampling mechanism to ensure that the estimated two-view epipolar geometries (e.g., fundamental matrices) are meaningful. Second, it adopts model-selection to exclude degenerate cases associate with planar homography. Detailed Epipolar-Test algorithm will be presented in Section-4.

3.2. Main Algorithm

Given the above rigidity-test is in place, we are now ready to present the main algorithm of the paper, namely *Structure-from-Recurrent-Motion (SfRM)*.

Algorithm 1: A high-level sketch of our Structure-From-Recurrent-Motion algorithm

Input: N perspective views of a non-rigid shape $S(t), t = 1, \dots, N$. Choose K , i.e., the desired number of clusters.

Output: The reconstructed 3D shapes of $S(t), \forall t \in \{1, \dots, N\}$ up to non-rigid transformations.

- 1 **for** $(i = 1, \dots, N, j = 1, \dots, N)$ **do**
 - 2 Call **Algorithm 2** (i.e., modified-epipolar-test) to get A matrix whose (i, j) -th entry $A(i, j)$ gives the probability that the two images i, j are rigidly related.
 - 3 **end**
 - 4 [Clustering] Form a view-graph $G(V, E, A)$ connecting all N views, and the A matrix is used as the affinity matrix. Run a suitable graph clustering algorithm to cluster the N views into K clusters.
 - 5 [Reconstruction] Apply any rigid SfM-reconstruction method to each of the K clusters.
-

Note that the core steps of the algorithm are A -matrix computation and graph clustering. It should be also noted

that our algorithm only makes use of rigid SfM routines to achieve non-rigid shape reconstruction.

4. Modified Epipolar Test

In this section, we describe our modified epipolar-test algorithm. The output of this algorithm is the probability measuring whether these two images can be the projections of a same rigid shape. As discussed in the previous section, we implement this by checking *whether or not these two sets of correspondences are related by a certain fundamental matrix, and at the same time not related by any planar homography*. The latter condition (i.e., excluding homography) is to ensure 3D reconstruction is possible. Our algorithm is inspired by an early work of McReynolds and Lowe for the same task of rigidity-checking [27], however ours is much simpler—without involving complicated parameter tuning and non-linear refinement. Rigidity-checking was also applied for solving multi-view geometry problems without via camera motion [23].

We will proceed by presenting our algorithm description first, followed by necessary explanations and comments.

Algorithm 2: Modified Epipolar Test algorithm

Input: Two input images, with M feature correspondences $\{(\mathbf{x}_i, \mathbf{x}'_i) | i = 1..M\}$
Output: The probability P that the two images are rigidly related.

1. **(Initialization):** Set parameters $\sigma_F, \sigma_H, \tau_F, \tau_H$.
2. **(Estimate fundamental matrices):** Sample all possible 8-point subsets from the M points; Totally there are $\binom{M}{8}$ such subsets. Store them in a list, and index its entries by k .
for $k = 1, \dots, \binom{M}{8}$ **do**
 - Pick the k -th 8-point subset, estimate fund-matrix F_k with the linear 8-point algorithm.
 - Given F_k , compute the geometric (point to epipolar-line) distances for all the M points by F_k , i.e., $d_F(\mathbf{x}'_i, F_k \mathbf{x}_i)$.
 - Convert the distances to probability measures by applying Gaussian kernel. Compute the product of all probability measures by:

$$P_F(k) = \prod_{i=1..M} \exp\left(-\frac{d_F^2(\mathbf{x}'_i, F_k \mathbf{x}_i)}{\sigma_F^2}\right), \quad (2)$$

end

Find the minimum of all the $\binom{M}{8}$ probabilities: i.e.,

$$P_F = \min_{k \in \binom{M}{8}} P_F(k). \quad (3)$$

3. (Estimate homography)

Run a similar procedure as above, for homography estimation, via sampling all 4-point subsets $l \in \binom{M}{4}$. The overall homography probability can be computed by:

$$P_H = \min_{l \in \binom{M}{4}} \prod_{i=1..M} \exp\left(-\frac{d_H^2(\mathbf{x}'_i, H_l \mathbf{x}_i)}{\sigma_H^2}\right). \quad (4)$$

4. (Compute overall probability)

By now we have both P_F , and P_H . Compare them with their respective tolerances δ_F , and δ_H .

if $(P_F \geq \tau_F)$ **and** $(P_H < \tau_H)$, **then**

Set $P = P_F(1 - P_H)$, **return** P .

else

Set $P = 0$, **return** P .

end

4.1. Why does the algorithm work?

In Step 3 of the algorithm 2, we sample subsets of the data points, each consists of 8 points, the minimally required number to *linearly* fit a fundamental matrix. This way we avoid forcing too many points to fit a single epipolar geometry. If the cameras are calibrated, one could also sample 5 points and use the non-linear 5-point essential-matrix algorithm for better sampling efficiency (e.g. [24, 25]).

Once a fundamental matrix F_k is estimated from an 8-tuple, we evaluate the probability of how likely every other feature points (not in the 8-tuple) satisfies this fundamental matrix. Assuming this probability is independent for each point, the product of Eq.2 gives the total probability P_k of how well this F_k explains all the M points. Exhausting all $\binom{M}{8}$ subsets, we pick the **least** one (in Eq.3) as a (i.e., conservative) estimate of the rigidity score. In Step 4, we repeat a similar sampling and fitting procedure for homography estimation. The idea is to perform *model-selection* [36] to filter out degenerate cases. Finally in Step 5, we report the overall probability (of rigidity-check for the two images) as the product of P_F and $(1 - P_H)$ when P_F is sufficiently high (i.e., $\geq \tau_F$) and P_H is sufficiently low (i.e., $< \tau_H$); otherwise report '0'. In summary, our algorithm provides a way to estimate the rigidity-score defined by the worst-case goodness-of-fitting achieved for all tentative fundamental matrices for each 8-tuple, while at the same time our algorithm favors the case hardly explained by a homography.

4.2. How to speed up the computation?

Our Algorithm 2 can be computationally expensive due to its exhaustive subset enumeration step (Step 3). For example, when $M = 100$, $\binom{M}{8}$ gives a large number of 186 billions.




Figure 2. Examples of A matrices: (from left to right), periodic, recurrent, and rigid scenarios.

Below we will show that one can almost safely replace the enumeration step with a randomized sampling process with much fewer samples, yet at little loss of accuracy. Specifically, we only need to replace the first line (of “For $k \in [1, \binom{M}{8}]...$ ” in Step-3 with “Randomly sample minimal 8-tuples for $k \leq K$ times..”.

Suppose there are about e proportion of valid subsets (i.e., e is the inlier ratio). By ‘valid’ we mean this 8-tuple gives a rise to a good epipolar geometry which explains all data points well enough. Then the odds (i.e., probability) of picking a valid 8-tuple by only sampling once is e , and the odds of getting an outlier is $1 - e$. If one samples K times, then the total odds of getting all K outliers is $(1 - e)^K$. Finally, the odds of getting at least one valid estimation is $p = 1 - (1 - e)^K$. This predicted odds can be very high in practice, suggesting that even a small number of random samples suffices. Note that this proof is akin to the probability calculation used in RANSAC, one can refer to [37] for details.

5. View Clustering and Block Reconstruction

For a given video sequence containing N views, we construct a complete *view-graph* $G(V, E, A)$ of N nodes where each corresponds to one view. E denotes the set of edges, and A the affinity matrix in which $A(i, j)$ measures the similarity between node- i and node- j .

After Step 3 of Algorithm 1, we have obtained an $N \times N$ matrix A . We will use this A as the affinity matrix of our view-graph.

Fig. 2 shows visualizations of example A s that characterizes different types of dynamic movements of objects in videos, showing periodic motion, recurrent motion, and rigid motion respectively. Fig. 2, bright colors in the matrix indicate at which views a particular shape re-occurred.

5.1. Spectral Clustering

Given a view-graph $G(V, E, A)$ with the rigidity matrix A as its affinity matrix, and choose a suitable number K as the intended number of clusters, we suggest to use spectral clustering technique to perform K -way camera view clustering. If two views are clustered to the same group, it means the two views are related by a rigid transformation.




Figure 3. For periodic motion, with $K = 40$ (i.e., one period). From left to right, original A matrix, rearranged A matrix after clustering and clustering membership.




Figure 4. For general recurrent motion, with $K = 25$. From left to right, original A matrix, rearranged A matrix after clustering and clustering membership.

Specifically, we use Shi-Malik’s Normalized-Cut [35] for its simplicity. The algorithm goes as follows: First, compute a diagonal matrix whose diagonal entries are $D(i, i) = \sum_j A(i, j)$. Then, form a Symmetric normalized Laplacian by $L = D^{-1/2} A D^{-1/2}$. Next, take the least $\log_2 K$ eigenvectors corresponding to the second smallest and higher eigen-values of L and run K-means algorithm on them to achieve K -way clustering. Some examples are given below, in Fig. 3 and 4.

5.2. Block-wise Rigid Reconstruction

After the spectral clustering, the A matrix is rearranged to a block-diagonal structure. Each block represents a cluster of views which are rigidly connected, up to an accuracy about the diameter of the cluster. Therefore, they can be considered as multiple rigid projections of the same shape. Hence any standard rigid-SfM technique can be used to recover the 3D shape. In our experiments specifically, we use incremental bundle adjustment which adds new frames gradually to a local triangulation thread.

5.3. Scale Normalization

As each rigid shape cluster is reconstructed independently, all recovered shapes are up to an ambiguous scale. To achieve globally consistent reconstruction results, we align the shape scale by normalizing distance between two selected landmarks (e.g., by normalizing the maximum limb-length for human body).

6. Results

The input of our method is multi-frame feature correspondences, as in other NRSfM methods (e.g., [11, 1]). Finding correspondences is a difficult task by itself, especially for non-rigid deforming objects where self-occlusions may happen frequently. In our experiments, for the synthetic data, we assume that the correspondences are provided. For the real data, we use the publicly available OpenPose [4] library to detect human poses, faces, and hands in sequences.

6.1. Periodic walking sequence

This first experiment aims to validate that our Algorithm 1 and 2 work for a real sequence with periodic movements — which is a special (and simpler) case of recurrent motion. We use a sequence capturing a walking person at a constant speed, where a moving camera is observing this person from different viewpoints, resulting in a nearly periodic sequence. We apply the OpenPose [4] library to detect 14 landmark points on the person over all 700 frames. Example frames are shown in Fig. 5. For the entire sequence, the rigidity (i.e., affinity) matrix computed by our Algorithm 2 is shown on the left of Fig. 6. This figure shows that there exists a strong periodicity, shown as bright bands along the main diagonals. Moreover the period can be readily read out as $p=40$ frames, although our algorithm does not make use of this result. Instead, frames with repetitive shapes are automatically grouped together via view-graph clustering. The middle and the right figure of Fig. 6 show the re-arranged affinity matrix after spectral clustering and the final clustering membership result respectively, where the evident ‘blocky’ structure clearly reveals the grouping. We then perform a rigid-SfM for all views within each block. Fig. 7 shows example pose reconstruction results; note the poses are in 3D.

So far, our algorithm has only focused on recovering the non-rigid shape itself, ignoring its absolute pose in the world coordinate frame. In practice this can be easily fixed, assuming that the ego-motion of each camera view can be recovered by, for example, standard rigid-SfM/SLAM techniques against a stationary background. We conduct this experiment by first tracking background points, then estimating absolute camera poses relative to the background, followed by Procrustes alignment between the absolute camera poses and each reconstructed human poses. The final reconstruction result, with both background point clouds and human poses and trajectories, is shown in Fig. 8.

6.2. Recurrent dancing sequence

This experiment aims to demonstrate the performance of our method on a general (*non-periodic*) video sequence which is likely to contain recurrent movements.



Figure 5. A (nearly) periodic walk sequence.




Figure 6. Affinity matrices before, and after spectral clustering (i.e. N-Cut). The ‘blocky’ structure becomes evident after N-cut. Right: the final view-clustering result.




Figure 7. 3D reconstruction results of the walking sequence.




Figure 8. Consistent 3D reconstruction of both dynamic foreground object (and temporal trajectories) and a static background scene.

We choose a solo dancing sequence captured by the CMU Panoptic Studio [17]. This dataset contains videos from camera arrays. In order to increase the probability of successfully reconstruction, we do not directly use one specific camera, but, instead, extract a time-consecutive video by randomly ‘hopping’ between different cameras in the dataset, to simulate a video as if captured by a ‘monocular camera randomly roaming in space’.

This dancing sequence is challenging as the motion of the dancer is fast and the dance itself is complicated creating many unnatural body movements. The computed affinity matrix is shown in Fig. 9, showing that there is no obvious




Figure 9. The computed original affinity matrix, and the block-wise pattern after clustering on the CMU dancing sequence. There is no obvious cyclic pattern in the original affinity matrix. After clustering, more clear recurrence patterns are revealed.




Figure 10. 3D reconstruction results on the dance sequence.

structure. However, after applying our graph-clustering, we can see a clear block-wise pattern (albeit noisy), suggesting that the video indeed contains many recurrent (repetitive) body poses. Example reconstruction results (along with the discovered recurrent frames) are shown in Fig. 10.

6.3. Quantitative evaluation

To quantitatively measure the performance of our method, we use the Blender to generate synthetic deformations with recurrence. We use the flying cloth dataset [38] and fold the sequence by several times to mimic recurrency. Camera views are randomly generated. Fig. 11 shows some sample frames of the data.

In this sequence, all ground-truth (object shape, camera poses) are given. Noises of different levels are added to image planes. Our method successfully detects recurrency and reconstructs the shape as shown in second row of Fig. 11.

The reconstruction quality is measured by shape errors after alignment, as well as the portion of successfully reconstructed frames. We evaluate on two criteria at different noise levels. Results are given in Fig. 13.

We compare our method with other the state-of-the-art template-free NRSfM methods[1, 7]. The result is shown in Fig. 12. In terms of overall reconstruction accuracy their performances are comparable, while ours is superior for frames exhibiting strong recurrency.




Figure 11. Cloth waving in the wind and our SfRM reconstructions.




Figure 12. Histograms of reprojection errors by different methods. Here we compare our method with shape-basis based method[7] and trajectory-basis based method[1].




Figure 13. SfRM performance at different noise levels. When noise increases, the reconstruction error increases whereas the success ratio falls. This result shows our method handles increasing amount of noises gracefully.

6.4. Timing

Fig. 14 gives the timing results of our SfRM system (excluding rigid reconstruction), showing a clear linear relationship w.r.t. the number of feature points, as well as w.r.t. the number of random samples (in algorithm-2), but is quadratically related to the number of image frames. In our experiments we chose K —the number of clusters— empirically. For future work we would like to investigate how to automatically determine K .

We also test our method on face and hand data captured by the Panoptic Studio. Sample qualitative results are shown in Fig. 15.




Figure 14. Timing (in seconds) as a function of the number of random samples, the number of points, and the number of frames.




Figure 15. Example images of 3D reconstruction of face and hand data.

7. Related work

The idea of our SfRM method is rather different from conventional NRSfM approaches. For space reason we will not review the NRSfM literature here but refer interested readers to recent publications on this topic and references therein [7, 29, 15, 19, 8, 18]. Below, we focus on previous work with similar ideas.

A cornerstone of our method is the mechanism to detect shape recurrence in a video. Similar ideas were proposed for periodic dynamic motion analysis [2, 32, 34, 33]. Our work is specifically inspired by [2, 14]. However, there are major differences. First, their methods assume strictly periodical motions, and need to estimate the period automatically [6] or manually [2]. This way, their methods can only handle limited periodical motions such as well-controlled walking and running. In contrast, our method extends to more general cases of recurrent motions, which include both *a-periodic*, and re-occurring cases. Moreover, their methods assume a camera to be static, and under the the periodical assumption, the target is not allowed to turn around and has to move (walking or running) on a straight line, capturing only partial surfaces [2] or trajectories [32]. Comparably, our method allows free-form target movements and camera motions. Finally, our method is fully automatic, while their methods rely on significant level of manual interactions.

Our method can be applied for 3D human pose recovery, therefore it is related to many work in this domain,

[13, 31, 30, 28, 20, 21]. In particular, our method is related to the research directions which try to lift 3D pose from 2D images, e.g. [3, 5, 26]. Earlier work in this direction either requires the integration of knowledge of the bone length of the target [22], or human pose and shape space priors [3]. Although in experiments we use 3D human poses, mainly as exemplar recurrent movements, our method does not take advantage of any category-specific priors. Rather, we treat poses as general point clouds in 3D. It can be applied to other objects beyond human body. Another category of work on human pose capture relies on the existence of large-scale pose database for retrieving the most similar pose based on a 3D-2D pose similarity metric [10, 5, 16]. Their performance is heavily depend on the size and quality of the database of specific type of targets, while ours works in general scenarios. A recent deep learning approach by Martinez et al. [26] shows that a well-designed network for directly regressing 3D keypoint positions from 2D joint detection showed good performance. However, they rely on large amount of training data of specific class, while ours works without training.

8. Conclusion

We have presented a new method for solving Non-rigid Structure-from-Motion (NRSfM) for a long video sequence containing recurrent motion. It directly extends the concept of rigidity to recurrency as well as periodicity. With this new method at hand, one is able to directly use traditional rigid SfM techniques for non-rigid problems. Key contributions of this work include a randomized algorithm for robust two-view rigidity check, and a view-graph clustering mechanism which automatically discovers recurrent shape enabling the subsequent rigid reconstructions. Finite but adequate experiments have demonstrated the usefulness of the proposed method. The method is practically relevant, thanks to the ubiquity of recurrent motions in reality. One may criticize our method will not work if a shape was only seen for one times. We admit this is a fair criticism, but we argue that if that happened it would be of little real practical value to reconstruct any shape with such a fleeting nature. Our proposed view-graph and shape-clustering algorithms are examples of unsupervised machine-learning techniques. In this regard, we hope this paper may offer insights that bridge SfM research with learning methods.

Acknowledgement. We would like to thank reviewers and ACs for their valuable comments. This work was completed when XL was a visiting PhD student to CMU under the CSC Scholarship (201706210160). HL is grateful for YS for his very generous hosting. HL’s work is funded in part by Australia ARC Centre of Excellence for Robotic Vision (CE140100016). YL’s research is funded by the National Key Foundation for Exploring Scientific Instrument (2013YQ140517) and NSFC grant (No. 61522111).

References

- [1] I. Akhter, Y. Sheikh, S. Khan, and T. Kanade. Nonrigid structure from motion in trajectory space. In *NIPS*, 2009. 1, 6, 7
- [2] S. Belongie and J. Wills. Structure from periodic motion. In *Spatial Coherence for Visual Motion Analysis*. Springer, 2006. 8
- [3] F. Bogo, A. Kanazawa, C. Lassner, P. V. Gehler, J. Romero, and M. J. Black. Keep it SMPL: automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016. 8
- [4] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 6
- [5] C. Chen and D. Ramanan. 3d human pose estimation = 2d pose estimation + matching. In *CVPR*, 2017. 8
- [6] R. Cutler and L. S. Davis. Robust real-time periodic motion detection, analysis, and applications. *TPAMI*, 2000. 8
- [7] Y. Dai, H. Li, and M. He. A simple prior-free method for non-rigid structure-from-motion factorization. *IJCV*, 2014. 1, 7, 8
- [8] M. Gallardo, T. Collins, A. Bartoli, and F. Mathias. Dense non-rigid structure-from-motion and shading with unknown albedos. In *CVPR*, 2017. 8
- [9] P. F. Gotardo and A. M. Martinez. Non-rigid structure from motion with complementary rank-3 spaces. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3065–3072. IEEE, 2011. 1
- [10] A. Gupta, J. Martinez, J. J. Little, and R. J. Woodham. 3d pose from motion for cross-view action recognition via non-linear circulant temporal encoding. In *CVPR*, 2014. 8
- [11] R. Hartley and R. Vidal. Perspective nonrigid shape and motion recovery. In *ECCV*, 2008. 6
- [12] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 1
- [13] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 2014. 8
- [14] P. P. Ivan Laptev, Serge Belongie and J. Wills. Periodic motion detection and segmentation via approximate sequence alignment. In *ICCV*, 2013. 8
- [15] P. Ji, H. Li, Y. Dai, and I. Reid. Maximizing rigidity revisited: A convex programming approach for generic 3d shape reconstruction from multiple perspective views. In *ICCV*, 2017. 8
- [16] H. Jiang. 3d human pose reconstruction using millions of exemplars. In *ICPR*, 2010. 8
- [17] H. Joo, H. Liu, L. Tan, L. Gui, B. Nabbe, I. Matthews, T. Kanade, S. Nobuhara, and Y. Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *ICCV*, 2015. 6
- [18] S. Kumar, A. Cherian, Y. Dai, and H. Li. Scalable dense non-rigid structure-from-motion: A grassmannian perspective. In *CVPR*, 2018. 8
- [19] S. Kumar, Y. Dai, and H. Li. Multi-body non-rigid structure-from-motion. In *3DV*, 2016. 8
- [20] W. Kusakunniran, Q. Wu, H. Li, and J. Zhang. Automatic gait recognition using weighted binary pattern on video. In *IEEE AVSS*, 2009. 8
- [21] W. Kusakunniran, Q. Wu, J. Zhang, H. Li, and L. Wang. Recognizing gaits across views through correlated motion co-clustering. *TIP*, 2014. 8
- [22] H. Lee and Z. Chen. Determination of 3d human body postures from a single view. *Computer Vision, Graphics, and Image Processing*, 1985. 8
- [23] H. Li. Multi-view structure computation without explicitly estimating motion. In *CVPR*, 2010. 4
- [24] H. Li and R. Hartley. Five-point motion estimation made easy. In *ICPR*, 2006. 4
- [25] H. li and R. Hartley. A simple solution to the six-point two-view focal-length problem. In *ECCV*, 2006. 4
- [26] J. Martinez, R. Hossain, J. Romero, and J. J. Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017. 8
- [27] D. P. McReynolds and D. G. Lowe. Rigidity checking of 3d point correspondences under perspective projection. *TPAMI*, 1996. 4
- [28] D. Mehta, S. Sridhar, O. Sotnychenko, H. Rhodin, M. Shafiei, H. Seidel, W. Xu, D. Casas, and C. Theobalt. Vnect: real-time 3d human pose estimation with a single RGB camera. *TOG*, 2017. 8
- [29] H. S. Park, T. Shiratori, I. A. Matthews, and Y. Sheikh. 3d reconstruction of a moving point from a series of 2d projections. In *ECCV*, 2010. 8
- [30] G. Pavlakos, X. Zhou, K. G. Derpanis, and K. Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *CVPR*, 2017. 8
- [31] V. Ramakrishna, T. Kanade, and Y. Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *ECCV*, 2012. 8
- [32] E. Ribnick and N. Papanikolopoulos. 3d reconstruction of periodic motion from a single view. *IJCV*, 2010. 8
- [33] S. M. Seitz and C. R. Dyer. *Affine invariant detection of periodic motion*. University of Wisconsin-Madison, Computer Sciences Department, 1994. 8
- [34] S. M. Seitz and C. R. Dyer. Detecting irregularities in cyclic motion. In *Motion of Non-Rigid and Articulated Objects, IEEE Workshop on*, 1994. 8
- [35] J. Shi and J. Malik. Normalized cuts and image segmentation. *TPAMI*, 2000. 2, 5
- [36] P. H. Torr. An assessment of information criteria for motion model selection. In *CVPR*, 1997. 4
- [37] P. H. Torr and D. W. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *IJCV*, 1997. 5
- [38] R. White, K. Crane, and D. A. Forsyth. Capturing and animating occluded cloth. *TOG*, 2007. 7