

Efficient Non-Consecutive Feature Tracking for Robust Structure-From-Motion

Guofeng Zhang, *Member, IEEE*, Haomin Liu, *Student Member, IEEE*, Zilong Dong, Jiaya Jia, *Senior Member, IEEE*, Tien-Tsin Wong, *Member, IEEE*, and Hujun Bao, *Member, IEEE*

Abstract—Structure-from-motion (SfM) largely relies on feature tracking. In image sequences, if disjointed tracks caused by objects moving in and out of the field of view, occasional occlusion, or image noise are not handled well, corresponding SfM could be affected. This problem becomes severer for large-scale scenes, which typically requires to capture multiple sequences to cover the whole scene. In this paper, we propose an efficient non-consecutive feature tracking framework to match interrupted tracks distributed in different subsequences or even in different videos. Our framework consists of steps of solving the feature “dropout” problem when indistinctive structures, noise or large image distortion exists, and of rapidly recognizing and joining common features located in different subsequences. In addition, we contribute an effective segment-based coarse-to-fine SfM algorithm for robustly handling large data sets. Experimental results on challenging video data demonstrate the effectiveness of the proposed system.

Index Terms—Non-consecutive feature tracking, track matching, structure-from-motion, bundle adjustment.

I. INTRODUCTION

LARGE-SCALE 3D reconstruction [1]–[3] finds many practical applications. It primarily relies on SfM algorithms [4]–[8] to firstly estimate sparse 3D features and camera poses given the input of video or image collections.

Manuscript received August 24, 2015; revised January 17, 2016 and August 8, 2016; accepted August 9, 2016. Date of publication September 8, 2016; date of current version October 31, 2016. This work was supported in part by NSF of China under Grant 61272048 and Grant 61232011, in part by the Fundamental Research Funds for the Central Universities under Grant 2015XZZX005-05, in part by a Foundation for the Author of National Excellent Doctoral Dissertation of China under Grant 201245, and in part by the Research Grants Council of the Hong Kong SAR under Project 2150760, and Project CUHK417913. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Peter Tay. (Corresponding authors : Guofeng Zhang; Hujun Bao.)

G. Zhang and H. Bao are with the State Key Laboratory of CAD and CG, Zhejiang University, Hangzhou 310058, China, and also the Innovation Joint Research Center for Cyber-Physical-Society System, Zhejiang University, Hangzhou 310058, China (e-mail: zhangguofeng@cad.zju.edu.cn; bao@cad.zju.edu.cn).

H. Liu and Z. Dong are with the State Key Laboratory of CAD and CG, Zhejiang University, Hangzhou 310058, China (e-mail: 172753015@qq.com; zldong@cad.zju.edu.cn).

J. Jia is with The Chinese University of Hong Kong, Hong Kong. (e-mail: leojia@cse.cuhk.edu.hk).

T.-T. Wong is with The Chinese University of Hong Kong, Hong Kong, and also with the Shenzhen Research Institute, The Chinese University of Hong Kong, Hong Kong (e-mail: ttwong@cse.cuhk.edu.hk).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author. The material contains a video that shows the feature tracking and structure-from-motion results by our method for the authors’ four datasets. A few comparisons with other methods are also included. The total size of the video is 53.4 MB. Contact zhangguofeng@cad.zju.edu.cn for further questions about this work.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2016.2607425

Compared to images, videos contain denser geometrical and structural information, and are the main source of SfM in the movie, video and commercial industry. A common strategy for video SfM estimation is by employing feature point tracking [9], [10], which takes care of the temporal relationship among frames. It is also a basic tool for solving a variety of computer vision problems, such as camera tracking, video matching, and object recognition.

In this paper, we address two critical problems for feature point tracking, which could handicap SfM especially for large-scale scene modeling. The first problem is the vulnerability of feature tracking to object occlusions, illumination change, noise, and large motion, which easily causes occasional feature drop-out and distraction. This problem makes robust feature tracking from long sequences challenging.

The other problem is the inability of sequential feature tracking to cope with feature matching over non-consecutive subsequences. A typical scenario is that the tracked object moves out and then re-enters the field-of-view, which yields two discontinuous subsequences containing the same object. Although there are common features in the two subsequences, they are difficult to be matched/included in a single track using conventional tracking methods. Addressing this issue can alleviate the drift problem of SfM, which benefits high-quality 3D reconstruction. A naïve solution is to exhaustively search all features, which could consume much computation since many temporally far away frames simply share no content.

We propose an efficient non-consecutive feature tracking (ENFT) framework which can effectively address the above problems in two phases – that is, *consecutive point tracking* and *non-consecutive track matching*. We demonstrate their significance for SfM using challenging sequence data. *Consecutive point tracking* detects and matches invariant features in consecutive frames. A matching strategy is proposed to greatly increase the matching rate and extend lifetime of the tracks. Then in *non-consecutive track matching*, by rapidly computing a matching matrix, a set of disjoint subsequences with overlapping content can be detected. Common feature tracks scattered over these subsequences can also be reliably matched.

Our ENFT method reduces estimation errors for long loopback sequences. Given limited memory, it is generally intractable to use global bundle adjustment to refine camera poses and 3D points for very long sequences. Iteratively applying local bundle adjustment is difficult to effectively distribute estimation errors to all frames. We address this issue by adopting a segment-based coarse-to-fine SfM algorithm,

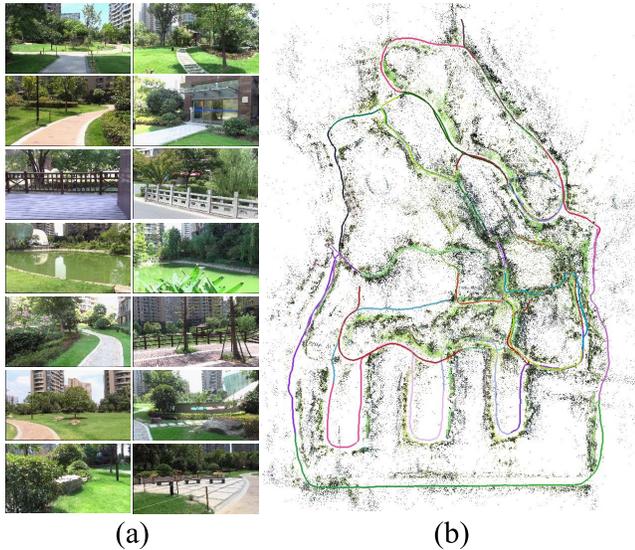


Fig. 1. A large-scale “Garden” example. (a) Snapshots of the input videos. (b) With the matched feature tracks, we register the recovered 3D points and camera trajectories in a large-scale 3D system. Camera trajectories are differently color-coded.

which globally optimizes structure and motion with limited memory.

Based on our ENFT algorithm and segment-based coarse-to-fine estimation scheme, we present the SfM system *ENFT-SfM*, which can effectively handle long loopback sequences and even multiple sequences. Fig. 1 shows an example containing 6 sequences with about 95,476 frames in total in a large-scale scene. Our system splits them to 37 shorter sequences, quickly computes many long and accurate feature tracks, efficiently estimates camera trajectories in different sequences, and finally registers them in a unified 3D system, as shown in Fig. 1(b). The whole process only takes about 90 minutes (excluding I/O) on a desktop PC, i.e., 17.7 FPS on average. Our supplementary video¹ contains the complete result.

Compared to our prior work [11], in this paper, we make a number of modifications to improve robustness and efficiency. Particularly, we improve the second-pass matching by formulating it as minimizing an energy function incorporating two geometric constraints, which not only boosts the matching accuracy but also reduces computation. The non-consecutive track matching algorithm is re-designed to perform feature matching and match-matrix update together. It is less sensitive to initialization and reduces the matching time. Finally, we propose a novel segment-based coarse-to-fine SfM method, which performs efficient global optimization for large data with only limited memory.

II. RELATED WORK

We review feature tracking and large-scale SfM methods in this section.

A. Feature Matching and Tracking

For video tracking, sequential matchers are used for establishing correspondences between consecutive frames.

Kanade-Lucas-Tomasi (KLT) tracker [9], [12] is widely used for small baseline matching. Other methods detect image features and match them considering local image patches [13], [14] or advanced descriptors [10], [15]–[17].

Both the KLT tracker and invariant feature algorithms depend on modeling feature appearance, and can be distracted by occlusion, similar structures, and noise. Generally, sequential matchers are difficult to match non-consecutive frames under large image transformation. Scale-invariant feature detection and matching algorithms [10], [17] are effective in matching images with large transformation. But they generally produce many short tracks in consecutive point tracking due primarily to the global indistinctiveness and feature dropout problems. In addition, invariant features are relatively sensitive to perspective distortion. Although variations, such as ASIFT [18], can improve matching performance under substantial viewpoint change, computation overhead increases owing to exhaustive viewpoint simulation. Cordes et al. [19] proposed a memory-based tracking method to extend feature trajectories by matching each frame to its neighbors. However, if an object re-enters the field-of-view after a long period of time, the size of neighborhood windows has to be very large. Besides, multiple-video setting was not discussed. In contrast, our method can not only extend track lifetime but also efficiently match common feature tracks in different subsequences by iteratively matching overlapping frame pairs and refining match matrix. The computation complexity is linear to the number of overlapping frame pairs.

There are methods using invariant features for object and location recognition in images/videos [20]–[24]. These methods typically use bag-of-words techniques to perform global localization and loop-closure detection in an image classification framework. Nistér and Stewénus [25] proposed using a hierarchical k -means algorithm to construct a vocabulary tree with feature descriptors, which can be used for large-scale image retrieval and location recognition. Cummins and Newman [26] proposed a probabilistic approach called FAB-MAP for location recognition and online loop closure detection, which models the world as a set of locations and computes the probability of belonging to previously visited locations for each input image. Later, they proposed using a sparse approximation for large scale location recognition [27]. However, FAB-MAP assumes the neighboring locations are not too close, so might perform less satisfyingly if we simply input a normal video sequence. In addition, existing methods generally divide the location recognition and non-consecutive feature matching into two separated phases [28]–[31]. Because the match matrix by bag-of-words only roughly reflects the match confidence, completely trusting it may lose many common features. In this paper, we introduce a novel strategy where the match matrix can be refined and updated along with non-consecutive feature matching. Our method can reliably and efficiently match the common features even with a coarse match matrix.

Engels et al. [32] proposed integrating wide-baseline local features with the tracked ones to improve SfM. The method creates small and independent submaps and links them via

¹<http://www.cad.zju.edu.cn/home/gfzhang/projects/tracking/featuretracking/ENFT-video.wmv>

feature recognition. This approach also cannot produce many long and accurate point tracks. Short tracks are not enough for drift-free SfM estimation. In comparison, our method is effective in high-quality point track estimation. We also address the ubiquitous nondistinctive feature matching problem in dense frames. Similar to the scheme of [33], we utilize track descriptors, instead of the feature descriptors, to reduce computation redundancy.

Wu et al. [34] proposed using dense 3D geometry information to extend SIFT features. In contrast, our method only uses sparse matches to estimate a set of homographies to represent scene motion, which also handles viewpoint change. It is general since geometry is not required.

B. Large-Scale Structure-From-Motion

State-of-the-art large-scale SfM methods can handle millions of images on a single PC in one day [3]. To this end, large image data are separated into a number of independent submaps, each is optimized independently. Steedly et al. [35] proposed a partitioning approach to decompose a large-scale optimization into multiple better-conditioned subproblems. Clemente et al. [36] proposed building local maps independently and stitching them with a hierarchical approach.

Ni et al. [37] proposed an out-of-core bundle adjustment (BA) for large-scale SfM. This method decomposes the data into multiple submaps, each of which has its own local coordinate system for optimization in parallel. For global optimization, an out-of-core implementation is adopted. Snavely et al. [38] proposed speeding up reconstruction by selecting a skeletal image set for SfM and then adding other images with pose estimation. Similarly, Konolige and Agrawal [39] selected a skeletal frame set and used reduced relative constraints for closing large loops. Each skeleton frame can actually be considered as a submap. A similar scheme is applied to iconic views [40], which are generated by clustering images with similar gist features [41]. In our work, a segment-based scheme is adopted, which first estimates SfM for each sequence independently, and then aligns the recovered submaps. Depending on estimation errors, we split each sequence to multiple segments, and perform segment-based refinement. This strategy can effectively handle large data and quickly reduce estimation errors during optimization.

Another line of research is to improve large-scale BA, which is a core component of SfM. Agarwal et al. [42] pointed out that connectivity graphs of Internet image collections are generally much less structured and accordingly presented an inexact Newton type BA algorithm. To speed up large-scale BA, Wu et al. [43] utilized multi-core CPUs or GPUs, and presented a parallel inexact Newton BA algorithm. Wu [44] also proposed preemptive feature matching that reduces matching image pairs, and an incremental SfM for full BA when the model is large enough. Pose graph optimization [45]–[47] was also widely used in realtime SfM and SLAM [48], [49], which uses the relative-pose constraints between cameras and is more efficient than full BA.

Most existing SfM approaches achieve reconstruction in an incremental way, which may risk drifting or local minima

when dealing with large-scale image sets. Crandall et al. [8] proposed combining discrete and continuous optimization to yield a better initialization for BA. By formulating SfM estimation as a labeling problem, belief propagation is employed to estimate camera parameters and 3D points. In the continuous step, Levenberg-Marquardt nonlinear optimization with additional constraints is used. This method is restricted to urban scenes, and assumes that the vertical vanishing point can be detected for rotation estimation, similar to the method proposed by Sinha et al. [50]. It also needs to leverage geotag contained in the collected images and takes complex discrete optimization. In contrast, our segment-based scheme can run on a common desktop PC with limited memory, even for large video data.

Real-time monocular SLAM methods [48], [49], [51], [52] typically perform tracking and mapping in parallel threads. The methods of [48] and [49] can close loops efficiently for large-scale scenes. However, they could still have difficulty in directly handling multiple sequences, as demonstrated in Figs. 8 and 9.

III. OUR APPROACH

Given a video sequence V with n frames, $V = \{I_t | t = 1, \dots, n\}$, our objective is to extract and match features in all frames in order to form a set of *feature tracks*. A feature track \mathcal{X} is defined as a series of feature points in images: $\mathcal{X} = \{\mathbf{x}_t | t \in f(\mathcal{X})\}$, where $f(\mathcal{X})$ denotes the frame set spanned by track \mathcal{X} . Each SIFT feature \mathbf{x}_t in frame t is associated with an appearance descriptor $\mathbf{p}(\mathbf{x}_t)$ [10] and we denote all descriptors in a feature track as $\mathcal{P}_{\mathcal{X}} = \{\mathbf{p}(\mathbf{x}_t) | t \in f(\mathcal{X})\}$.

With the detected m features in all frames, finding matchable ones generally requires a large amount of comparisons even using the k-d trees; meanwhile it inevitably induces errors due to the fact that a large number of features make descriptor space hardly distinctive, resulting in ambiguous matches. So it is neither reliable nor practical to only compare the feature descriptors to form tracks. Our ENFT method has two main steps to address this issue. The framework is outlined in Table I.

For reducing computation, we can extract one frame for every 3 ~ 5 frames to constitute a new sequence and then perform feature tracking on it. In the consecutive tracking stage, we employ a two-pass matching strategy to extend the track lifetime. Then in the non-consecutive tracking stage, we match common features in different subsequences. With the obtained feature tracks, a segment-based SfM scheme is employed to robustly recover the 3D structure and camera motion. Finally, if necessary, we propagate feature points from sampled frames to others. Since the 3D positions of these features have been computed, we can quickly estimate the camera poses of remaining frames with the obtained 3D-2D correspondences.

IV. CONSECUTIVE TRACKING

For video sequences, feature tracks are typically obtained by matching features between consecutive frames. However, due to illumination change, repeated texture, noise, and large

TABLE I
FRAMEWORK OVERVIEW OF ENFT-SFM

1.	Consecutive point tracking (Section IV):
1.1	Match the extracted SIFT features between consecutive frames with descriptor comparison.
1.2	Perform the second-pass matching to extend track lifetime.
2.	Non-consecutive track matching (Section V):
2.1	Use hierarchical k-means to cluster the constructed invariant tracks.
2.2	Estimate the match matrix with the grouped tracks.
2.3	Detect overlapping subsequences and join the matched tracks.
3.	Segment-based coarse-to-fine SfM (Section VI):
3.1	Estimate the submap for each sequence.
3.2	Match the common tracks among different sequences, and then use them to estimate the similarity transformations for each submap.
3.3	Use segment-based SfM to refine the aligned submaps.
4.	[Optional] Feature propagation with camera estimation:
4.1	Quickly propagate features from sampled frames to others.
4.2	Quickly estimate camera poses for remaining frames.

image distortion, features are easily dropped out or mismatched, resulting in breaking many tracks into shorter ones. In this section, we propose an *improved two-pass matching* strategy to alleviate this problem. The first-pass matching is the same as our prior method [11], which uses SIFT algorithm [10] with RANSAC [53] to obtain high-confidence matches and remove outliers. In the second pass matching, we firstly use the inlier matches to estimate a set of homographies $\{H_{t,t+1}^k | k = 1, \dots, N\}$ with multiple RANSAC procedures [11], [54]. To handle illumination change, we estimate global illumination variation $L_{t,t+1}$ between images I_t and I_{t+1} by computing the median intensity ratio between matched features. Here, I_t denotes the gray scale image of frame t .

We first linearly scale image I_t with $L_{t,t+1}$, and then transform it with homography $H_{t,t+1}^k$ to obtain the rectified image \hat{I}_t^k . Correspondingly, \mathbf{x}_t in image I_t is rectified to $\hat{\mathbf{x}}_t^k$ where $\hat{\mathbf{x}}_t^k \sim H_{t,t+1}^k \mathbf{x}_t$ in \hat{I}_t^k . The distance between a 2D point \mathbf{x}_{t+1}^k and the epipolar line $l_{t,t+1}(\mathbf{x}_t)$ is denoted by $d(\hat{\mathbf{x}}_t^k, l_{t,t+1}(\mathbf{x}_t))$. If $\hat{\mathbf{x}}_t^k$ largely deviates from the epipolar line (i.e., $d(\hat{\mathbf{x}}_t^k, l_{t,t+1}(\mathbf{x}_t)) > \tau_e$), we reject $H_{t,t+1}^k$ since it does not describe the motion of \mathbf{x}_t well. For each remaining $H_{t,t+1}^k$, we track \mathbf{x}_t to \mathbf{x}_{t+1}^k by minimizing the matching cost:

$$S_{t,t+1}^k(\mathbf{x}_{t+1}^k) = \sum_{\mathbf{y} \in W} |\hat{I}_t^k(\hat{\mathbf{x}}_t^k + \mathbf{y}) - I_{t+1}(\mathbf{x}_{t+1}^k + \mathbf{y})|^2 + \lambda_e d(\mathbf{x}_{t+1}^k, l_{t,t+1}(\mathbf{x}_t))^2 + \lambda_h \|\hat{\mathbf{x}}_t^k - \mathbf{x}_{t+1}^k\|^2 \quad (1)$$

where $\hat{\mathbf{x}}_t^k + \mathbf{y}$ are the points in the window W centered at $\hat{\mathbf{x}}_t^k$. Different from our prior method [11], the matching cost incorporates two geometric constraint terms, which encourage \mathbf{x}_{t+1}^k to be along the epipolar line and obey homography. $\|\cdot\|$ is the Euclidean distance and $|\cdot|$ is the absolute value. The corresponding weights are $\lambda_e = |W| \sigma_c^2 / \sigma_e^2$ and $\lambda_h = |W| \sigma_c^2 / \sigma_h^2$, where σ_c , σ_e , and σ_h account for the uncertainty of intensity, epipolar geometry and homography

transformation respectively. In our experiments, these values are by default $\sigma_c = 0.1$ (for intensity values normalized to $[0, 1]$), $\sigma_e = 2$ and $\sigma_h = 10$. Note that σ_h is relatively large because we do not require the points to strictly lie on the same plane. As long as the point is near the plane, H_k can alleviate the major distortion and provide a better matching condition.

Similar to KLT tracking, we solve for $S_{t,t+1}(\mathbf{x}_{t+1}^k)$ iteratively by taking the partial derivative w.r.t. \mathbf{x}_{t+1}^k and setting it to zero:

$$\frac{\partial S_{t,t+1}^k(\mathbf{x}_{t+1}^k)}{\partial \mathbf{x}_{t+1}^k} = 0. \quad (2)$$

$I_{t+1}(\mathbf{x} + \Delta \mathbf{x})$ is approximated by a Taylor expansion truncated up to its first order:

$$I_{t+1}(\mathbf{x} + \Delta \mathbf{x}) \approx I_{t+1}(\mathbf{x}) + g_{t+1}^\top(\mathbf{x}) \cdot \Delta \mathbf{x} \quad (3)$$

where g_{t+1}^\top is the image gradient in the $(t+1)^{th}$ frame. With the computed gradients, we propose an iterative solver to optimize (1) by first initializing \mathbf{x}_{t+1} as the midpoint between $\hat{\mathbf{x}}_t^k$ and its projection to $l_{t,t+1}(\mathbf{x}_t)$, as shown in Fig. 3. Then we iteratively update \mathbf{x}_{t+1} by solving (2). In iteration $i+1$, \mathbf{x}_{t+1} is updated as

$$\mathbf{x}_{t+1}^{k,(i+1)} = \mathbf{x}_{t+1}^{k,(i)} + \Delta \mathbf{x}$$

where $\mathbf{x}_{t+1}^{k,(i)}$ denotes the value of \mathbf{x}_{t+1}^k in iteration i . This procedure continues until $\Delta \mathbf{x}$ is sufficiently small.

The found match is denoted as \mathbf{x}_{t+1}^k . With the set of homographies $\{H_{t,t+1}^k | k = 1, \dots, N\}$, we can find several matches $\{\mathbf{x}_{t+1}^k | k = 1, \dots, N\}$. Only the best one $j = \min_k \sum_{\mathbf{y} \in W} |\hat{I}_t^k(\hat{\mathbf{x}}_t^k + \mathbf{y}) - I_{t+1}(\mathbf{x}_{t+1}^k + \mathbf{y})|$ is kept.

In case the feature motion cannot be described by any homographies or feature correspondence is indeed missing, the found match is actually an outlier. We detect it with the following conditions:

$$\begin{cases} \sum_{\mathbf{y} \in W} |\hat{I}_t^j(\hat{\mathbf{x}}_t^j + \mathbf{y}) - I_{t+1}(\mathbf{x}_{t+1}^j + \mathbf{y})| > \tau_c |W|; \\ d(\mathbf{x}_{t+1}^j, l_{t,t+1}(\mathbf{x}_t)) > \tau_e; \\ \|\hat{\mathbf{x}}_t^j - \mathbf{x}_{t+1}^j\| > \tau_h. \end{cases}$$

These conditions represent the constraints of color constancy, epipolar geometry and homography respectively. If any of them is satisfied, \mathbf{x}_{t+1}^j is treated as an outlier. τ_c is set to a small value (generally 0.02 in our experiments) since the image is rectified. The remaining two parameters are $\tau_e = 2$ and $\tau_h = 10$. Considering points may not strictly undergo planar transformation, τ_h is set to a relatively large value.

Fig. 2(c) shows the result after the second-pass matching. Compared to our prior method [11] (Fig. 2(d)), the improved two-pass matching method does not need to perform additional KLT matching. It thus runs faster. The computation time is only 18ms with GPU acceleration on a NVIDIA GTX780 display card. The number of credible matches also increases.

The two-pass matching can produce many long tracks. Each track has a group of descriptors. They are similar to each other in the same track due to the matching criteria. We compute average of the descriptors over the track, and denote it as *track descriptor* $\mathbf{p}(\mathcal{X})$. It is used in the following non-consecutive track matching.

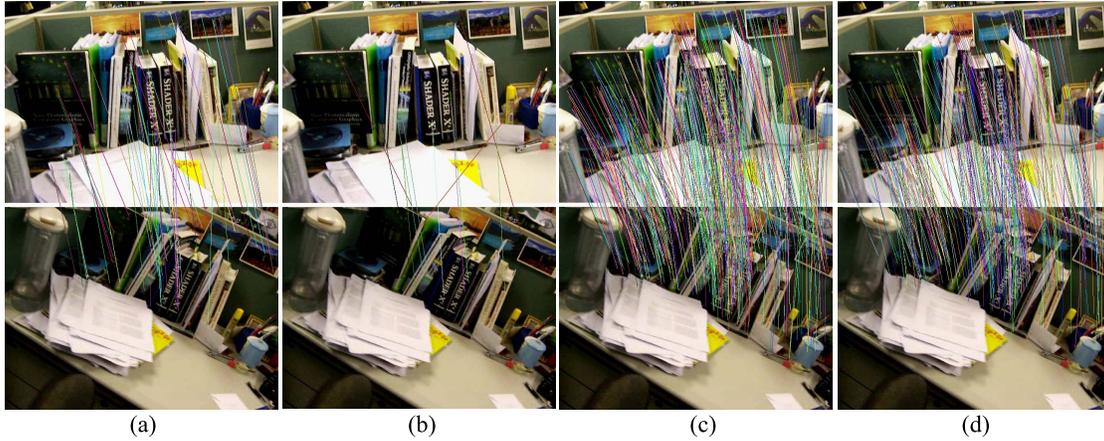


Fig. 2. Feature matching comparison. (a) First-pass matching by SIFT descriptor comparison. There are 958 features detected in the first image, but only 53 matches are found. (b) Additional match by directly searching the correspondences along the epipolar lines with SIFT descriptor comparison. Only 11 additional matches are found. (c) Second-pass matching with outlier rejection. 399 (i.e. 53 + 346) matches are obtained. (d) Matching result of [11]. 314 matches are obtained.

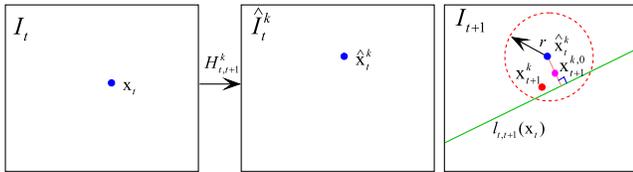


Fig. 3. Constrained spatial search with planar motion segmentation. Given homography $H_{t,t+1}^k$, we rectify I_t to \hat{I}_t^k such that $\hat{\mathbf{x}}_t^k \sim H_{t,t+1}^k \mathbf{x}_t$. Then we select the midpoint between $\hat{\mathbf{x}}_t^k$ and its projection to $l_{t,t+1}(\mathbf{x}_t)$ for initialization, and search the matched point by minimizing (1). The red dot \mathbf{x}_{t+1}^k is the result.

V. NON-CONSECUTIVE TRACK MATCHING

In this stage, we match features distributed in different subsequences, which is vital for drift-free SfM estimation. If we select all image pairs in a brute-force manner, the process can be intolerably costly for a long sequence. A better strategy is to estimate content similarity among different images first. We propose a non-consecutive track matching (NCTM) method to address this problem.

There are two steps. In the first step, similarity of different images is coarsely estimated by constructing a $n \times n$ symmetric match matrix M , where n is the number of frames. $M(i, j)$ stores overlapping confidence between images I_i and I_j . We use the same method of [11] to quickly estimate the initial matching matrix M , which first uses hierarchical K-means to cluster the track descriptors and then compute the similarity confidence of frame pairs by counting the number of potentially matched tracks that are clustered into the same leaf node.

For acceleration, we only select long tracks that span 5 or more keyframes to estimate overlap confidence. In our experiments, for the “Desktop” sequence, the initial match matrix estimation only takes 1.08 seconds, with a total of 5,935 selected feature tracks. Fig. 4(a) shows the initially estimated match matrix for the “Desktop” sequence. Bright pixels are with high overlapping confidence where many common features exist. Because we exclude track self-matching, the diagonal band of estimated match matrix has

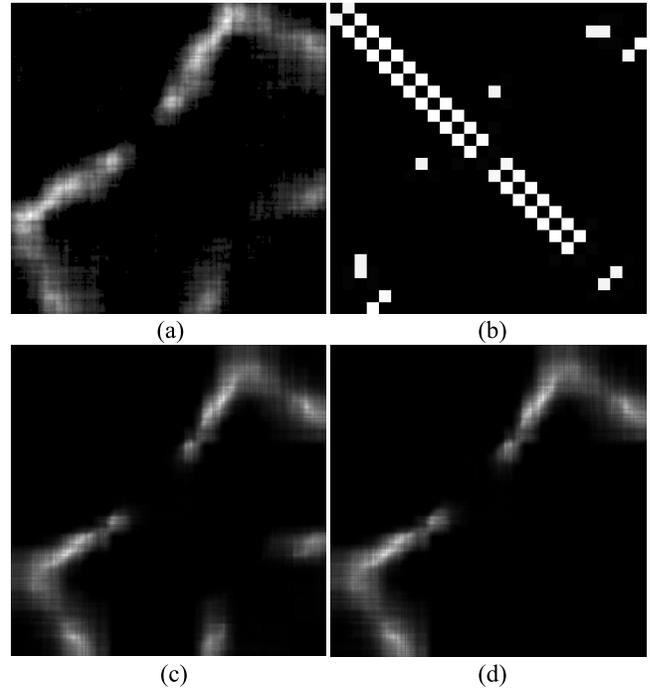


Fig. 4. Match matrix estimation for the “Desktop” sequence containing 941 frames. (a) Our initially estimated match matrix based on the keyframes. The matrix size is scaled for visualization. (b) Estimated match matrix by FAB-MAP [27] on the re-sampled sequence that contains 26 frames. (c) The final match matrix for all frames after our non-consecutive matching based on (a). (d) The final match matrix after our non-consecutive matching based on (b).

no value. Our method handles dense image sequences, unlike FAB-MAP [26], [27] that assumes sparsely sampled ones. When applying FAB-MAP to the original “Desktop” sequence, no loop is detected. So we manually sample the original sequence until common points between adjacent sampled frames are no more than 100. This generates 26 sampled frames. As shown in Fig. 4(b), in this case, a few overlapping image pairs are identified by FAB-MAP; but they are not enough to match many common features.

In the second step, with the initially estimated match matrix, we select the frame pairs with maximum overlapping confidence to perform feature matching, and update the match matrix iteratively. Matrix estimation and non-consecutive feature matching are benefitted from each other to simplify computation. Fig. 4(c) shows our finally estimated match matrix.

For speedup, we extract keyframes based on the result of consecutive feature tracking described in Section IV. Frame 1 is selected as the first keyframe. Then we select frame i as the second keyframe if it satisfies $N_1(1, i) \geq m_1$ and $N_1(1, i + 1) < m_1$, where $N_1(i, j)$ denotes the number of common features between frames i and j . Other keyframes are selected as follows. For the two recent keyframes with indices i_1 and i_2 in the original sequence, we select frame j ($j > i_2$) as the new keyframe if it is the farthest one from i_2 that satisfies $\{N_1(i_1, j) \geq m_1, N_2(i_1, i_2, j) \geq m_2\}$, where $N_2(i_1, i_2, j)$ denotes the number of common points among the three frames (i_1, i_2, j) . This step is repeated until all frames are processed. In our experiments, $m_1 = 100 \sim 500$ and $m_2 = 50 \sim 300$. Without special notice, the following procedures are only performed on keyframes.

A. Non-Consecutive Track Matching

Since the number of common features between two frames can be coarsely reflected by the initially estimated match matrix M , we select a frame pair (t_1^0, t_2^0) with the largest value in M to start matching. After matching (t_1^0, t_2^0) , the set of matched track pairs $C_{\mathcal{X}} = \{(\mathcal{X}_1, \mathcal{X}_2)\}$ approximately represent the number of common tracks for neighboring frame pairs. The matched track pairs in frame pair (t_1, t_2) can be expressed as

$$C_{\mathcal{X}}(t_1, t_2) = \{(\mathcal{X}_1, \mathcal{X}_2) | t_1 \in f(\mathcal{X}_1), t_2 \in f(\mathcal{X}_2), \\ \times (\mathcal{X}_1, \mathcal{X}_2) \in C_{\mathcal{X}}\}. \quad (4)$$

The number of common features in (t_1, t_2) can be approximated by $|C_{\mathcal{X}}(t_1, t_2)|$ as long as (t_1, t_2) shares sufficient common tracks with (t_1^0, t_2^0) . We maintain an updating match matrix M^* , computed as

$$M^*(t_1, t_2) = |C_{\mathcal{X}}(t_1, t_2)| \quad (5)$$

to propagate the overlapping confidence from (t_1^0, t_2^0) toward the neighboring frame pairs, and determine where the next matching should be performed. Details are given below.

1) *Main Procedure*: We first detect the largest element (t_1^0, t_2^0) in M . The value of $M(t_1^0, t_2^0)$ is also denoted as M_{\max} . If $M(t_1^0, t_2^0)$ is larger than a threshold, several common features may exist. After matching (t_1^0, t_2^0) , we collect and put the matched track pairs into $C_{\mathcal{X}}$ and update M^* according to Eq. (5). In particular, we set $M^*(t_1^0, t_2^0) = 0$, indicating (t_1^0, t_2^0) is matched. Next, we repeatedly select the largest element (t_1^k, t_2^k) in the updating matrix M^* , match (t_1^k, t_2^k) , and update $C_{\mathcal{X}}$ and M^* accordingly. This procedure continues until $M^*(t_1^k, t_2^k) < 50$. Then we go to another region by re-detecting the brightest point in M that has not been processed. The step ends if the brightest value is smaller than $0.1M_{\max}$.

TABLE II

NON-CONSECUTIVE TRACK MATCHING COMPARISON BETWEEN THE METHOD OF [11] AND OURS FOR THE ‘‘DESKTOP’’ AND ‘‘CIRCLE’’ SEQUENCES

Methods	Desktop Sequence		Circle Sequence	
	Merged Tracks	Time	Merged Tracks	Time
[11]	16, 279	81s	101, 948	132s
Our method	16, 827	35s	102, 583	55s

2) *Frame Pair Matching and Outlier Rejection*: When entering a new bright region, we perform the classical 2NN matching for (t_1^0, t_2^0) . Then each matching pair (t_1^k, t_2^k) is detected from the updating matrix M^* . Thus there are $M^*(t_1^k, t_2^k)$ common features found previously. We use these matches to estimate the fundamental matrix $F_{t_1^k, t_2^k}$ of frame pair (t_1^k, t_2^k) , and re-match those outlying features along the epipolar lines. We further search the correspondences for other unmatched features along epipolar lines.

Along with the fundamental matrix estimation between t_1^k and t_2^k , these $M^*(t_1^k, t_2^k)$ matches are classified into inliers and outliers. Since only part of matches are used to estimate $F_{t_1^k, t_2^k}$, the estimated $F_{t_1^k, t_2^k}$ could be biased. So we do not reject outliers immediately. Fortunately, each matched track pair $(\mathcal{X}_1, \mathcal{X}_2)$ undergoes multi-pass epipolar verification during processing the whole bright region. We record all the verification results for each $(\mathcal{X}_1, \mathcal{X}_2)$, and determine inliers/outliers after all bright regions are processed. Suppose $(\mathcal{X}_1, \mathcal{X}_2)$ is classified as an inlier match N_I times and as an outlier match N_O times. We reject $(\mathcal{X}_1, \mathcal{X}_2)$ if $N_I < s \cdot N_O$ ($s = 1 \sim 4$ in our experiments). In addition, we use the following strategy to remove the potential matching ambiguity. For example, a track \mathcal{X}_1 may find two corresponding tracks \mathcal{X}_2 and \mathcal{X}'_2 , where \mathcal{X}_2 and \mathcal{X}'_2 have overlapping frames. So the track matches $(\mathcal{X}_1, \mathcal{X}_2)$ and $(\mathcal{X}_1, \mathcal{X}'_2)$ conflict with each other. In this case, we simply select the best match with the largest N_I , and regard the other as an outlier.

3) *Benefits*: The proposed matching method outperforms previous ones in the following aspects. In our prior method [11], a rectangular region in the roughly estimated match matrix M is sought each time and local exhaustive track matching is performed for all frame pairs in it. It could involve a lot of unnecessary matching for non-overlapping frames and repeated feature comparison. Our current scheme only selects the frame pairs with sufficient overlapping, and matches each pair of frames and most tracks at most once. As shown in Table II, compared to the method of [11], our new non-consecutive track matching algorithm is more efficient. Both methods are implemented without GPU acceleration.

Standard image matching is to find a set of most similar images given the query one. This scheme has been extensively used in large-scale SfM [3], [7] and realtime SLAM systems for loop closure detection [28], [29], [55]. It, however, also may involve unnecessary matching for unrelated frame pairs and miss those with considerable common features. It is because image similarity based on appearance may not be sufficiently reliable. In contrast, we progressively expand frames with track matching. The expansion is not fully related to the

initial match matrix. Therefore a very rough matrix is enough to give a good starting point. Practically, as long as there is one good position, our system can extend it to the whole overlapping region accurately. To verify this, we provided two refined match matrices based on two different rough match matrices, as shown in Figs. 4(a) and (b). Although the two initially estimated match matrices are different and only based on keyframes, the finally estimated match matrices after our non-consecutive track matching are quite similar (except the bottom right area, where the initial match matrix by FAB-MAP does not provide any high confidence elements), which demonstrates the effectiveness of our method.

B. SfM for Multiple Sequences

Our method can be naturally extended to handle multiple sequences. Given one or multiple sequences, we first split long ones, making each new sequence generally contains only 1000 ~ 3001 frames. The splitted neighboring sequences can contain some overlapping frames for reliable matching. The sequence set is denoted as $\{V_i | i = 1, \dots, n\}$. Then we apply our feature tracking to each V_i , and estimate its 3D structure and camera motion using a keyframe-based incremental SfM scheme similar to that of [6]. The major modification is that we use known intrinsic camera parameters, and simply select an initial frame pair that has sufficient matches and a large baseline to start SfM. For each sequence pair, we use the fast matching matrix estimation algorithm [11] to estimate the rough match matrix such that related frames in any two different sequences can be found and common features can be matched by the algorithm introduced in Section V-A. Then we use the segment-based SfM method described in Section VI to efficiently recover and globally register 3D points and camera trajectories, as shown in Fig. 1(b).

VI. SEGMENT-BASED COARSE-TO-FINE SfM

With the independently reconstructed sequences and matched common tracks, we align them in a unified 3D coordinate system. For a long loopback sequence, error accumulation could be serious, making traditional bundle adjustment stuck in local optimum. It is because the first a few iterations of bundle adjustment aggregate accumulation errors at the joint loop points, which are hard to be propagated to the whole sequence. To address this problem, we split each sequence into multiple segments – each is with a similarity transformation. Only these transformations and overlapping points across different segments are optimized. We name it segment-based bundle adjustment and illustrate it in Fig. 6. Lim *et al.* [28] performed global adjustment by clustering keyframes into multiple disjoint sets (i.e. segments), which is conceptually similar to our idea. But the geodesic-distance-based segmentation to cluster frames could make inconsistent structure be put into a single body, complicating alignment-error reduction. This method also did not adaptively split the segments in a coarse-to-fine way to minimize the accumulation error. Local optimization within each body may not sufficiently minimize the error which is mainly caused by global misalignment.

In the beginning, we order all sequences and define the one that contains the maximum number of tracks merged with others as the reference. Without losing generality, we define it as sequence #1, denoted as V_1 . Its local 3D coordinate system is also set as the reference. Then with the common tracks among different sequences, we can estimate the coordinate transformation for each sequence j (i.e., V_j), denoted as $T_j = (s_j, R_j, t_j)$, where s_j is the scale factor, R_j is the rotational matrix, and t_j is the translation vector. For the reference sequence, s_1 have value 1, R_1 is an identity 3×3 matrix, and $t_1 = (0, 0, 0)^T$.

Each segment is assigned with a similarity transformation, and the relative camera motion between frames in each segment is fixed, so that the number of variables is small enough for efficient optimization. Different from [28], which clusters frames using geodesic distances, we propose clustering neighboring and geometrically consistent frames into segments. The position at which two consecutive frames are inconsistent is defined as a “split point”. We project the common points in each consecutive frame pair into the two images and check the re-projection error.

However, directly detecting the split points according to reprojection error is not optimal since it is generally large at loop closure points. Splitting such frame pairs does not help. We instead find split points that the re-projection error is most likely to be reduced. Assume each frame k is associated with a small similarity transformation T_k , which is parameterized as a 7-vector a_k (three Rodrigues components for rotation, 3D translation and scale). If we minimize the re-projection error w.r.t. a_k , the steepest descent direction is

$$g_k = \sum_{i=1 \dots N_k} A_i^T e_i \quad (6)$$

where N_k is the number of points visible in frame k , and A_i is the Jacobian matrix $A_i = \partial \pi(P_k X_i) / \partial a_k$. π is the projection function. e_i is the re-projection error $e_i = \mathbf{x}_i - \pi(P_k X_i)$, which is reduced along the direction of g_k . For two consecutive frames ($k, k+1$), if their g_k and g_{k+1} have similar directions, their re-projection errors both can be reduced with the same similarity transformation. Otherwise, these two frames are better to be assigned to different segments. The inconsistency between two consecutive frames is defined as the angle between the two steepest descent directions

$$C(k, k+1) = \arccos \frac{g_k^T \cdot g_{k+1}}{\|g_k\| \cdot \|g_{k+1}\|} \quad (7)$$

For verification, we group every 100 consecutive frames into one segment for the “Desktop” (Fig. 5(a)), and apply a certain transformation to each segment (Fig. 5(b)). As expected, the re-projection errors distribute in the whole overlapping regions. In contrast, the angle between the steepest descent directions reliably reflects the splitting result.

We progressively segment the sequences. At the t^{th} iteration, each sequence is divided into 2^t segments. We compute $C(k, k+1)$ for all k and detect the $2^t - 1$ split points with the largest $C(k, k+1)$. In order to evenly spread the split points across the whole sequence, we perform non-maximal suppression during selecting split points. While selecting the

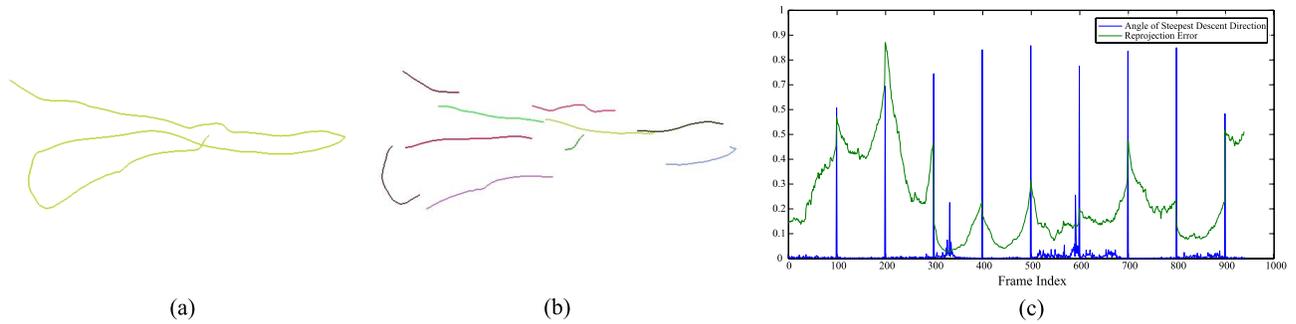


Fig. 5. Split point detection. (a) Original camera trajectory of the “Desktop” sequence. (b) Splitted camera trajectories. Each segment contains 100 frames. (c) Re-projection errors (green curve) and angles of steepest descending direction (blue curve). Values are all normalized to [0,1] for better comparison. The angle more accurately reflects the split result quality compared to the re-projection error.

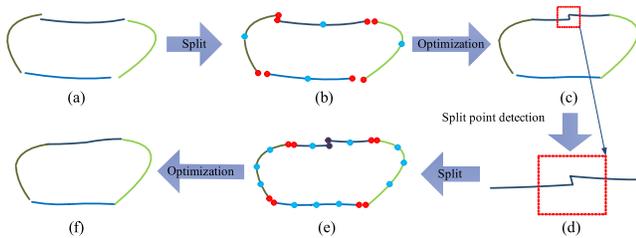


Fig. 6. Segment-based coarse-to-fine refinement. (a) Recovered camera trajectories marked with different colors. (b) Each sequence is split into 2 segments where endpoints and split points are highlighted. (c) Refined camera trajectories after the first iteration, where errors are redistributed. (d) “Split points”, which are joints of largely inconsistent camera motion for consecutive frames. (e) Sequence separation by split points. Two dark points denote the splitted two consecutive frames in a split point. (f) Refined camera trajectories after 2 iterations.

largest one, its neighboring $\frac{N_j}{2^t}$ candidates (N_j is the number of frames in sequence j) are suppressed and then select the next largest one from the remaining ones with non-maximal suppression. This procedure is repeated until $2^t - 1$ split points are selected. We put the consecutive frames in between two adjacent split points into a segment, and use the method described as follows to estimate the similarity transformations and submaps jointly for all segments. When the optimization is done, we detect split points for each sequence again, and re-separate the sequence into multiple segments. We can repeat this process until the average reprojection error is below a threshold or each segment contains only one frame. Errors are progressively propagated and reduced. The procedure of our segment-based coarse-to-fine refinement scheme is illustrated in Fig. 6.

Algorithm Details

Suppose the number of detected split points among all n sequences is m . We break the sequences into a total of $n' = n + m$ segments. Each of them is with a similarity transformation $T_j^w = (s_j^w, R_j^w, t_j^w)$, where $j = 1, \dots, n'$, w.r.t. the world coordinate. We use BA to refine the reconstructed 3D feature points with these similarity transformations. Different from traditional BA, the camera parameters inside each segment are fixed, we thus only update the similarity transformation. The procedure is to first transform one 3D point

in the world coordinate to a local one with parameters T^w . Then traditional perspective camera projection is employed to compute the re-projection error. Our BA function is written as

$$\min \sum_{i=1}^{N'} \sum_{j=1}^{n'} \sum_{k=1}^{n_j} w_{i,j,k} \|\pi(K_{j,k}(R_{j,k}(s_j^w R_j^w X_i + t_j^w) + t_{j,k}) - \mathbf{x}_{i,j,k})\|^2 \quad (8)$$

where n_j is the number of frames in the j -th segment, N' is the number of the 3D feature points, and n' is the number of the segments. π is the projection function. $\mathbf{x}_{i,j,k}$ is the image location of X_i in the k -th frame of the j -th subsequence. $K_{j,k}$, $R_{j,k}$, and $t_{j,k}$ are the intrinsic matrix, rotational matrix, and translation vector, respectively. $w_{i,j,k}$ is defined as

$$w_{i,j,k} = \begin{cases} 1, & \text{If point } i \text{ is visible in frame } k \text{ in sequence } j \\ 0, & \text{Otherwise} \end{cases}$$

We use Schur complement and forward substitution [56] to solve the normal equation, which separates the updating of rigid transformation and of 3D points in each iteration. It reduces the large linear system to a linear symmetric one with scale $7n' \times 7n'$ for updating transformation. It makes 3D point estimation much cheaper because each point can be updated independently by solving a 3×3 linear symmetric system. Moreover, since only a few segment pairs share common points, the Schur complement is rather sparse. In SBA [57], the system of Schur complement was explicitly constructed and solved by Cholesky decomposition. Wu et al. [43] implicitly built the Schur complement for parallel computing. They did not take full advantage of the sparsity property. For acceleration, sSBA [58] proposed to utilize the sparse structure of Schur complement and solve it with sparse Cholesky decomposition. We also utilize the sparsity and solve it with efficient preconditioned conjugate gradient similar to that of [43], which can significantly reduce the computation.

Because the size of the linear system is actually determined by n' , we can estimate n' based on the available memory. Once the size n' linear system is reached, SfM refinement is performed in the following two steps. In the first step, we only select the $m = n' - n$ split points to split the sequences, and solve (8) to refine the result. In the second step, we

TABLE III
RUNNING TIME OF ENFT-SFM

Datasets	Frames	Step	Sampled Frames	Resolution	Feature Tracking		SfM Estimation			Propagation		Reprojection Error
					CPT	KNCTM	Submap Estimation	Align	Refine	Feature Propagation	Camera Estimation	
Desktop	941	1	941	640 × 480	46.5s	5.8s	14.1s	-	-	-	-	1.26 pixels
Circle	2,129	3	710	960 × 540	63.1s	40.9s	13.8s	-	-	13.9s	4.2s	1.07 pixels
Street	22,799	5	5,537	960 × 540	7.4 min.	7.5 min.	176.0s	3.6s	32.0s	3.4 min.	65.1s	2.49 pixels
Garden	95,476	3 or 5	21,791	960 × 540	27.4 min.	31.2 min.	588.1s	2.5s	130.4s	15.6 min.	3.7 min.	2.28 pixels

perform a local BA for each sequence j iteratively by re-splitting sequence j to multiple segments with detected split points and refining them by solving (8) while fixing cameras and 3D points in other sequences. This process stops when all sequences are processed. This strategy makes it possible to efficiently and robustly handle large data with limited memory consumption.

Finally, we fix the 3D points and estimate the camera poses respectively for all frames. During the course of iterations, errors are quickly reduced.

VII. EXPERIMENTAL RESULTS

We evaluate our method with several challenging datasets. Running time is listed in Table III excluding I/O, which is obtained on a desktop PC with an Intel i7-4770K CPU, 8GB memory, and a NVIDIA GTX780 graphics card. The operating system is 64-bit Windows 7. Only the feature tracking component is accelerated by GPU. We use 64D descriptors for SIFT features. Our SIFT GPU implementation is inspired by [59] but runs faster. For SfM estimation, we optimize the code by applying SSE instructions, but only use a **single thread** without parallel computing. For the sequences captured by us, since the intrinsic matrix is known, we optimize the SfM code by incorporating this prior to improve the robustness and efficiency. Garden dataset contains 6 sequences, which are further splitted into 37 shorter sequences, from which we sample the frames by setting the step to 3 or 5. The source code and datasets can be found in our project website.²

As our consecutive point tracking can handle wide-baseline images, frame-by-frame tracking is generally unnecessary. For our datasets listed in Table III, we usually extract one frame for every 3 ~ 5 frames to apply feature tracking. We quickly propagate the feature points to other frames by KLT with GPU acceleration. This trick further saves computation. In addition, in order to reduce image noise and blur, for each input frame I_t , we perform matching with two past frames. One is the last frame I_{t-1} , and the other (denoted as $I_{t'}$) is the farthest frame that shares over 300 common features with I_{t-1} . Note that only a small number of features in $I_{t'}$ need to be matched with I_t , which does not increase computation much.

A. Quantitative Evaluation of Feature Tracking

We compare the feature tracking methods of consecutive SIFT matching (C-SIFT), our consecutive point tracking (CPT), brute-force SIFT matching (BF-SIFT), our

consecutive point tracking with non-consecutive track matching (CPT+NCTM), our consecutive point tracking with keyframe-based non-consecutive track matching (CPT+KNCTM).

C-SIFT extracts and matches SIFT features only in consecutive frames. It is a common strategy for feature tracking. The advantage is that the complexity is linear to the number of frames. However, feature dropout could occur due to global indistinctiveness or image noise, which causes producing many short tracks. The brute-force SIFT matching exhaustively compares extracted SIFT features, whose complexity is quadratic to the number of processed frames. In comparison, the complexity of our method (CPT+NCTM) is linear to the number of processed frames and the number of overlapping frame pairs while high quality results are guaranteed.

The “Circle” sequence contains 2,129 frames. To make computation feasible for a few prior methods, we select one frame for every 3 consecutive ones, which forms a new sequence containing 710 frames in total. Table IV lists the running time with GPU acceleration. Our consecutive point tracking (CPT) needs a bit more time than C-SIFT. But it significantly extends the lifetime of most tracks. With our non-consecutive track matching, common feature tracks scattered over disjoint subsequences are connected, further expanding track lifetime. Compared with the computationally most expensive BF-SIFT, our result (CPT+NCTM) obtains more long feature tracks and the computation is much faster. With keyframe-based acceleration, our non-consecutive track matching time is further significantly reduced (from 87.6s to 40.9s), without influencing much matching result. Table IV lists the average length of tracks for all tracks with length ≥ 1 . The computed average length is short because we also take into account unmatched features with track length 1. The quality of SfM computed by BF-SIFT, CPT+NCTM and CPT+KNCTM are quite comparable, as shown in our supplementary material.³

B. Comparison With Other SfM/SLAM Systems

We compare our ENFT-SfM system with state-of-the-art SfM/SLAM systems (i.e. ORB-SLAM [49], VisualSfM [60], [43], [44] and OpenMVG [61]) using our datasets and other public benchmark datasets (i.e. KITTI odometry dataset [62] and TUM RGB-D dataset [63]). Since VisualSfM and OpenMVG are mainly designed for unordered image datasets, we extract keyframes from the original sequences as input for

²<http://www.zjucv.net/ls-acts/ls-acts.html>

³<http://www.cad.zju.edu.cn/home/gfzhang/projects/tracking/featuretracking/ENFT-supple.pdf>

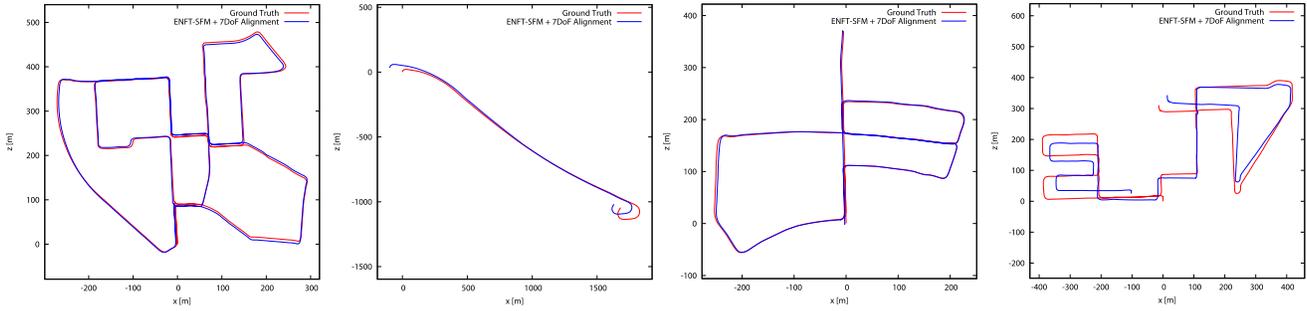


Fig. 7. The recovered camera trajectories by ENFT-SFM in KITTI 00, 01, 05 and 08 sequences.

TABLE IV
PERFORMANCE OF DIFFERENT ALGORITHMS

Algorithms	Running Time	Average Track Length
C-SIFT	38.4s	1.73
CPT	63.1s	2.28
CPT+NCTM	150.7s	3.10
CPT+KNCTM	104.0s	2.68
BF-SIFT	1086.4s	2.71

TABLE V
LOCALIZATION ERROR (RMSE (M)/COMPLETENESS) COMPARISON IN
KITTI ODOMETRY DATASET

Seq.	ENFT-SFM	ENFT-SFM (Keyframes)	ORB- SLAM	VisualSFM (Keyframes)	OpenMVG (Keyframes)
00	4.58 / 100%	4.76 / 100%	5.33	2.78 / 3.71%	5.83 / 0.7%
01	57.20 / 100%	53.96 / 100%	X	52.34 / 12.46%	8.79 / 2.08%
02	28.13 / 100%	28.26 / 100%	21.28	1.77 / 4.53%	50.36 / 3.74%
03	2.82 / 100%	2.94 / 100%	1.51	0.28 / 12.05%	3.53 / 8.43%
04	0.66 / 100%	0.66 / 100%	1.62	0.76 / 23.44%	5.14 / 14.06%
05	2.88 / 100%	3.48 / 100%	4.85	9.77 / 7.42%	22.42 / 9.07%
06	14.24 / 100%	14.43 / 100%	12.34	8.58 / 7.41%	3.16 / 3.37%
07	1.83 / 100%	2.03 / 100%	2.26	3.85 / 7.78%	7.75 / 5%
08	30.74 / 100%	28.32 / 100%	46.68	0.81 / 0.90%	17.82 / 2.58%
09	5.63 / 100%	5.88 / 100%	6.62	0.90 / 4.92%	14.26 / 3.36%
10	19.53 / 100%	18.49 / 100%	8.80	5.70 / 6.05%	27.06 / 7.01%

VisualSFM and OpenMVG. For fair comparison, our method processes both original sequences and extracted keyframes for KITTI odometry dataset and TUM RGB-D dataset.

For KITTI and TUM RGB-D datasets, we align recovered camera trajectories and ground truth by estimating a 7DoF similarity transformation. The RMSE and completeness of camera trajectories for all methods are listed in Tables V and VI. “X” denotes that the map cannot be accurately initialized or processed. The recovered camera trajectories of sequences 00, 01, 05 and 08 from KITTI odometry dataset by our method are shown in Fig. 7. The complete results are included in our supplementary material (see previous footnote link). Because sequences 01 and 08 do not contain loops, the drift cannot be corrected, leading to large RMSE.

For ORB-SLAM, we directly quote reported RMSE error of keyframe trajectory in their paper. Compared with ORB-SLAM, our method achieves comparable results in KITTI odometry dataset. We note only our method is able to process all sequences (the camera poses of some frames in TUM RGB-D sequences are not recovered due to extremely

TABLE VI
LOCALIZATION ERROR (RMSE (CM)/COMPLETENESS) COMPARISON IN
TUM RGB-D DATASET

Sequence	ENFT-SFM	ENFT-SFM (Keyframes)	ORB- SLAM	VisualSFM (Keyframes)	OpenMVG (Keyframes)
fr1_desk	2.71/99.84%	2.96/100%	1.69	2.74/100%	X
fr1_floor	4.08/96.70%	3.93/100%	2.99	53.11/69.23%	0.52/6.92%
fr1_xyz	1.25/100%	1.59/100%	0.90	1.43/100%	X
fr2_360_kidnap	13.57/91.47%	15.31/100%	3.81	10.08/50.91%	5.21/14.55%
fr2_desk	2.43/100%	2.27/100%	0.88	1.79/100%	1.38/13.95%
fr2_desk_person	2.46/100%	2.55/100%	0.63	1.92/100%	2.16/97.01%
fr2_xyz	0.81/100%	0.73/100%	0.30	0.71/100%	5.74/97.6%
fr3_long_office	1.21/100%	1.44/100%	3.45	1.15/100%	2.94/32.74%
fr3_nst_tex_far	3.60/86.58%	7.76/100%	X	7.29/100%	35.64/3.79%
fr3_nst_tex_near	1.87/100%	1.66/100%	1.39	1.13/100%	3.4/39.13%
fr3_sit_half	1.50/100%	1.55/100%	1.34	2.30/100%	0.68/9.3%
fr3_sit_xyz	0.84/100%	1.39/100%	0.79	1.28/100%	1.03/100%
fr3_str_tex_far	0.94/100%	0.95/100%	0.77	2.15/100%	1.12/100%
fr3_str_tex_near	1.86/100%	1.82/100%	1.58	0.95/100%	0.97/19.74%
fr3_walk_half	2.08/100%	2.21/100%	1.74	1.88/100%	X
fr3_walk_xyz	1.30/100%	1.74/100%	1.24	1.62/100%	X

serious motion blur, occlusion or there are not sufficient texture regions). We fix the parameters for both KITTI and TUM RGB-D datasets except for the maximum frame number for each sequence segment. It is set as 300 for KITTI odometry dataset and 1,500 for TUM RGB-D dataset respectively. Since the camera moves fast in KITTI odometry dataset, the maximum frame number for each segment should be smaller to reduce the accumulation error.

For our multi-sequence data, since ORB-SLAM cannot directly handle multiple sequences, we constitute multiple sequences into a single sequence by re-ordering the frame index. The input frame rate is set to 10fps for ORB-SLAM.⁴ The recovered camera trajectories by ORB-SLAM are shown in Figs. 8 and 9. The camera poses of many frames are not recovered due to unsuccessful relocalization. Although some loops are closed, the optimization is stuck in a local optimum. The reason is twofold. On the one hand, the matched common features among non-consecutive frames by a traditional bag-of-words place recognition method [64] are insufficient for robust SfM/SLAM. On the other hand, using pose graph optimization [46], [47] may not sufficiently minimize accumulation error, and traditional BA is easily stuck in a local optimum if a good starting point is not provided.

VisualSFM does not work that well in KITTI odometry dataset and our long sequences, as shown in Table V and

⁴We use ORB-SLAM2: https://github.com/raulmur/ORB_SLAM2.

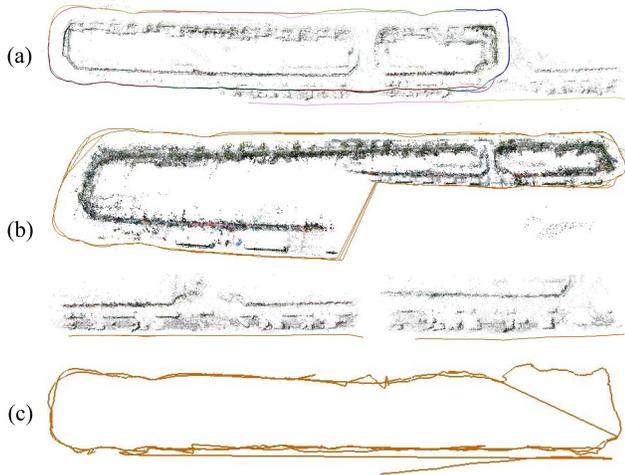


Fig. 8. Reconstruction comparison on the “Street” example. (a) SfM result of ENFT-SfM. (b) SfM result of VisualSfM, which is separated to 3 individual models. (c) The recovered camera trajectory by ORB-SLAM.

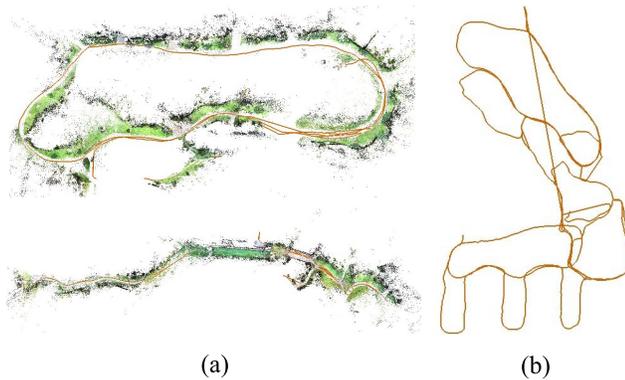


Fig. 9. Reconstruction comparison on the “Garden” example. (a) Two individual models reconstructed by VisualSfM. The reconstructed SfM result contains 60 individual models. (b) The recovered camera trajectory by ORB-SLAM.

Figs. 8 and 9. Note the matching time in our data is overly long for VisualSfM. We have to use our non-consecutive feature tracking algorithm to get the feature matching results. The produced SfM results still have the drifting problem and the whole camera trajectory is easily separated into multiple segments. We thus select the largest segment for computing RMSE and completeness. One reason for this drifting problem is the incremental SfM, which may not effectively eliminate accumulated errors. Another explanation is that sequence continuity/ordering is not completely utilized. Since the KITTI dataset is captured by an autonomous driving platform and each frame is only matched to its consecutive frames. Once camera tracking fails in one frame, the connection between two neighboring subsequences will be broken. In our experiments, OpenMVG usually performs worse than VisualSfM.

C. Results on General Image Collections

Although our segment-based SfM method is originally designed for handling sequences, it can be naturally extended to work with general image collections. The basic idea is

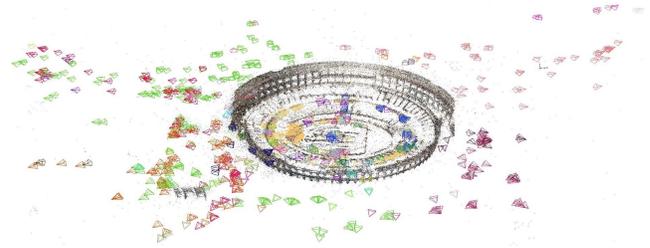


Fig. 10. The reconstruction result of “Colosseum” dataset by our method. Cameras in the same sequence are encoded with the same color.

to separate the unordered image data to a set of sequences according to their common matches.

We first select two images with the maximum number of common features to constitute an initial sequence. Then we select another image, which has the most common features with the head or tail frame, and add it into the sequence as the new head or tail. This process repeats until no image can be added. Then we begin to build another sequence based on remaining images. For some 3D points that have only one or two corresponding features in one sequence, we additionally select related images from other sequences to help estimate the 3D positions.

Fig. 10 shows our SfM result on Colosseum dataset [65], [66], which contains 1,164 images. We directly use the feature matching result obtained by VisualSfM. Because our current SfM implementation requires that the intrinsic camera parameters and radial distortion are known for each image, we calibrate the matched feature positions according to the calibrated parameters by VisualSfM. Then we use our extended segment-based SfM method to estimate camera poses and 3D points. The processing time of our SfM estimation in a single thread is 125 seconds, which is even shorter than that of VisualSfM enabling GPU (269 seconds).

D. Parameter Configuration and Limitation

The parameters can be easily set in our system because most of them are not sensitive and use default values. The most important parameter is τ_c , which controls the strength to mark outliers during feature tracking. A large τ_c could result in many matches, and introduce outliers. In our experiments, we conservatively set τ_c to a small value 0.02. By removing a small set of matches, the system becomes reliable for high-quality SfM. Fig. 11 shows the matching result with different τ_c . After the first-pass matching, 35 matches are obtained. The second-pass matching result with $\tau_c = 0.06$ is shown in Fig. 11(b). A few features that do not belong to the green book are included. These outliers are removed by using smaller τ_c values, as shown in (c) and (d). By setting $\tau_c = 0.02$, almost all outliers are removed and 95 reliable matches are obtained.

The proposed two-pass matching works best if the scene can be represented by multiple planes. For a video sequence with dense frames, this condition can be generally achieved because image transformation between two consecutive frames is small for viable approximation by one or multiple homographies.

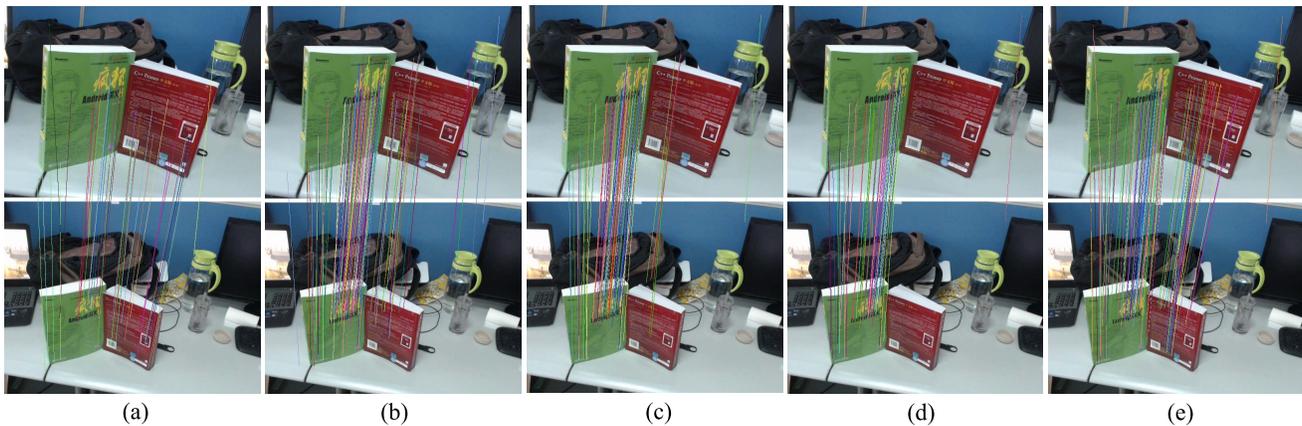


Fig. 11. Matching result with different τ_c . (a) First-pass matching result. (b-d) Results of the second-pass matching only using the homography corresponding to the left green book with $\tau_c = \{0.06, 0.04, 0.02\}$, respectively. The matches that do not belong to the green book are outliers. (e) Second-pass matching result using all homographies with $\tau_c = 0.02$. 95 matches are obtained.

We note even if the scene deviates from piecewise planarity, our second-pass matching still works as rectified images are close to the target ones. Our method may be not suitable for wide-baseline sparse images where the number of matches by first-pass matching is too small.

VIII. CONCLUSION AND DISCUSSION

We have presented a robust and efficient non-consecutive feature tracking (ENFT) method for robust SfM, which consists of two main steps, i.e., consecutive point tracking and non-consecutive track matching. Different from typical sequential matchers, e.g., KLT, we use invariant features and propose a two-pass matching strategy to significantly extend the track lifetime and reduce the feature sensitivity to noise and image distortion. The obtained tracks avail estimating a match matrix to detect disjointed subsequences with overlapping views. A new segment-based coarse-to-fine SfM estimation scheme is also introduced to effectively reduce accumulation error for long sequences. The presented ENFT-SfM system can handle tracking and registering large video datasets with limited memory consumption.

Our ENFT method greatly helps SfM, and considers feature tracking on non-deforming objects by tradition. Part of our future work is to handle dynamic objects. In addition, although the proposed method is based on SIFT features, there is no limitation to use other representations, e.g., SURF [17] and ORB [67], for further acceleration.

ACKNOWLEDGEMENTS

The authors thank Changchang Wu for his kind help in running VisualSfM in our datasets.

REFERENCES

- [1] M. Pollefeys *et al.*, "Detailed real-time urban 3D reconstruction from video," *Int. J. Comput. Vis.*, vol. 78, nos. 2–3, pp. 143–167, 2008.
- [2] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski, "Towards Internet-scale multi-view stereo," in *Proc. CVPR*, vol. 2010, pp. 1434–1441.
- [3] J.-M. Frahm *et al.*, "Building rome on a cloudless day," in *Proc. ECCV*, 2010, pp. 368–381.
- [4] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [5] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: Exploring photo collections in 3D," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 835–846, 2006.
- [6] G. Zhang, X. Qin, W. Hua, T.-T. Wong, P.-A. Heng, and H. Bao, "Robust metric reconstruction from challenging video sequences," in *Proc. CVPR*, Jun. 2007, pp. 1–8.
- [7] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, S. M. Seitz, and R. Szeliski, "Building rome in a day," in *Proc. ICCV*, 2009, pp. 72–79.
- [8] D. J. Crandall, A. Owens, N. Snavely, and D. Huttenlocher, "Discrete-continuous optimization for large-scale structure from motion," in *Proc. CVPR*, 2011, pp. 3001–3008.
- [9] C. Tomasi and T. Kanade, "Detection and tracking of point features," Carnegie Mellon Univ., Pittsburgh, PA, USA, Tech. Rep. CMU-CS-91-132, Apr. 1991.
- [10] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [11] G. Zhang, Z. Dong, J. Jia, T.-T. Wong, and H. Bao, "Efficient non-consecutive feature tracking for structure-from-motion," in *Proc. ECCV*, 2010, pp. 422–435.
- [12] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. IJCAI*, vol. 2, 1981, pp. 674–679.
- [13] D. Nistér, O. Naroditsky, and J. R. Bergen, "Visual odometry," in *Proc. CVPR*, 2004, pp. 652–659.
- [14] E. Royer, M. Lhuillier, M. Dhome, and J.-M. Lavest, "Monocular vision for mobile robot localization and autonomous navigation," *Int. J. Comput. Vis.*, vol. 74, no. 3, pp. 237–260, 2007.
- [15] K. Mikołajczyk and C. Schmid, "A performance evaluation of local descriptors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 10, pp. 1615–1630, Oct. 2005.
- [16] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide-baseline stereo from maximally stable extremal regions," *Image Vis. Comput.*, vol. 22, no. 10, pp. 761–767, 2004.
- [17] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "SURF: Speeded up robust features," *Comput. Vis. Image Understand.*, vol. 110, no. 3, pp. 346–359, 2008.
- [18] J.-M. Morel and G. Yu, "ASIFT: A new framework for fully affine invariant image comparison," *SIAM J. Imag. Sci.*, vol. 2, no. 2, pp. 438–469, 2009.
- [19] K. Cordes, O. Müller, B. Rosenhahn, and J. Ostermann, "Feature trajectory retrieval with application to accurate structure and motion recovery," in *Proc. Adv. Vis. Comput.*, Berlin, Germany, vol. 6938, 2011, pp. 156–167.
- [20] J. Sivic and A. Zisserman, "Video Google: A text retrieval approach to object matching in videos," in *Proc. ICCV*, 2003, pp. 1470–1477.
- [21] F. Schaffalitzky and A. Zisserman, "Automated location matching in movies," *Comput. Vis. Image Understand.*, vol. 92, nos. 2–3, pp. 236–264, 2003.

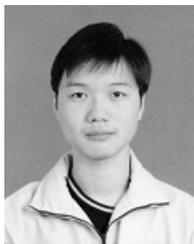
- [22] K. L. Ho and P. Newman, "Detecting loop closure with scene sequences," *Int. J. Comput. Vis.*, vol. 74, no. 3, pp. 261–286, 2007.
- [23] G. Schindler, M. Brown, and R. Szeliski, "City-scale location recognition," in *Proc. CVPR*, 2007, pp. 1–7.
- [24] A. Irshara, C. Zach, J.-M. Frahm, and H. Bischof, "From structure-from-motion point clouds to fast location recognition," in *Proc. CVPR*, 2009, pp. 2599–2606.
- [25] D. Nister and H. Stewenius, "Scalable recognition with a vocabulary tree," in *Proc. CVPR*, 2006, pp. 2161–2168.
- [26] M. Cummins and P. Newman, "FAB-MAP: Probabilistic localization and mapping in the space of appearance," *Int. J. Robot. Res.*, vol. 27, no. 6, pp. 647–665, 2008.
- [27] M. Cummins and P. Newman, "Appearance-only SLAM at large scale with FAB-MAP 2.0," *Int. J. Robot. Res.*, vol. 30, no. 9, pp. 1100–1123, 2011.
- [28] J. Lim, J.-M. Frahm, and M. Pollefeys, "Online environment mapping," in *Proc. CVPR*, 2011, pp. 3489–3496.
- [29] B. Clipp, J. Lim, J.-M. Frahm, and M. Pollefeys, "Parallel, real-time visual SLAM," in *Proc. IROS*, 2010, pp. 3961–3968.
- [30] E. Eade and T. Drummond, "Unified loop closing and recovery for real time monocular SLAM," in *Proc. BMVC*, 2008, pp. 6.1–6.10.
- [31] M. Havlena, A. Torii, and T. Pajdla, "Efficient structure from motion by graph optimization," in *Proc. ECCV*, vol. 6312, 2010, pp. 100–113.
- [32] C. Engels, F. Fraundorfer, and D. Nistér, "Integration of tracked and recognized features for locally and globally robust structure from motion," in *Proc. VISAPP Workshop Robot Perception*, 2008, pp. 13–22.
- [33] M. Grabner and H. Bischof, "Extracting object representations from local feature trajectories," in *Proc. Joint Hungarian-Austrian Conf. Image Process. Pattern Recognit.*, vol. 192, 2005, pp. 265–272.
- [34] C. Wu, B. Clipp, X. Li, J.-M. Frahm, and M. Pollefeys, "3D model matching with viewpoint-invariant patches (VIP)," in *Proc. CVPR*, 2008, pp. 1–8.
- [35] D. Steedly, I. A. Essa, and F. Dellaert, "Spectral partitioning for structure from motion," in *Proc. ICCV*, 2003, pp. 996–1003.
- [36] L. A. Clemente, A. J. Davison, I. D. Reid, J. Neira, and J. D. Tardós, "Mapping large loops with a single hand-held camera," in *Proc. Robotics: Sci. Syst.*, 2007.
- [37] K. Ni, D. Steedly, and F. Dellaert, "Out-of-core bundle adjustment for large-scale 3D reconstruction," in *Proc. ICCV*, 2007, pp. 1–8.
- [38] N. Snavely, S. M. Seitz, and R. Szeliski, "Skeletal sets for efficient structure from motion," in *Proc. CVPR*, 2008.
- [39] K. Konolige and M. Agrawal, "FrameSLAM: From Bundle Adjustment to Real-Time Visual Mapping," *IEEE Trans. Robot.*, vol. 24, no. 5, pp. 1066–1077, Oct. 2008.
- [40] X. Li, C. Wu, C. Zach, S. Lazebnik, and J.-M. Frahm, "Modeling and recognition of landmark image collections using iconic scene graphs," in *Proc. 10th Eur. Conf. Comput. Vis.*, 2008, pp. 427–440.
- [41] A. Oliva and A. Torralba, "Modeling the shape of the scene: A holistic representation of the spatial envelope," *Int. J. Comput. Vis.*, vol. 42, no. 3, pp. 145–175, 2001.
- [42] S. Agarwal, N. Snavely, S. M. Seitz, and R. Szeliski, "Bundle adjustment in the large," in *Proc. ECCV*, vol. 6312, 2010, pp. 29–42.
- [43] C. Wu, S. Agarwal, B. Curless, and S. M. Seitz, "Multicore bundle adjustment," in *Proc. CVPR*, 2011, pp. 3057–3064.
- [44] C. Wu, "Towards linear-time incremental structure from motion," in *Proc. 3DV*, 2013, pp. 127–134.
- [45] E. Olson, J. J. Leonard, and S. J. Teller, "Fast iterative alignment of pose graphs with poor initial estimates," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2006, pp. 2262–2269.
- [46] H. Strasdat, J. M. M. Montiel, and A. J. Davison, "Scale drift-aware large scale monocular SLAM," in *Proc. Robotics: Sci. Syst. VI*, 2010.
- [47] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "G²o: A general framework for graph optimization," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2011, pp. 3607–3613.
- [48] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *Proc. 13th Eur. Conf. Comput. Vis.*, vol. 8690, 2014, pp. 834–849.
- [49] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, Oct. 2015.
- [50] S. N. Sinha, D. Steedly, and R. Szeliski, "A multi-stage linear approach to structure from motion," in *Proc. Trends Topics Comput. Vis.*, vol. 6554, 2010, pp. 267–281.
- [51] G. Klein and D. W. Murray, "Parallel tracking and mapping for small AR workspaces," in *Proc. ISMAR*, 2007, pp. 225–234.
- [52] W. Tan, H. Liu, Z. Dong, G. Zhang, and H. Bao, "Robust monocular SLAM in dynamic environments," in *Proc. IEEE Int. Symp. Mixed Augmented Reality*, Oct. 2013, pp. 209–218.
- [53] M. A. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [54] Y. Jin, L. Tao, H. Di, N. I. Rao, and G. Xu, "Background modeling from a free-moving camera by multi-layer homography algorithm," in *Proc. 15th IEEE Int. Conf. Image Process.*, Oct. 2008, pp. 1572–1575.
- [55] R. O. Castle, G. Klein, and D. W. Murray, "Video-rate localization in multiple maps for wearable augmented reality," in *Proc. ISWC*, 2008, pp. 15–22.
- [56] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment—a modern synthesis," in *Proc. Workshop Vis. Algorithms*, 1999, pp. 298–372.
- [57] M. A. Lourakis and A. Argyros, "SBA: A Software Package for Generic Sparse Bundle Adjustment," *ACM Trans. Math. Softw.*, vol. 36, no. 1, pp. 1–30, Mar. 2009.
- [58] K. Konolige, "Sparse sparse bundle adjustment," in *Proc. Brit. Mach. Vis. Conf.*, 2010, pp. 1–11.
- [59] C. Wu. (2007). *SiftGPU: A GPU Implementation of Scale Invariant Feature Transform (SIFT)*. [Online]. Available: <http://cs.unc.edu/~ccwu/siftgpu>
- [60] C. Wu. (2013). *VisualSfM: A Visual Structure From Motion System*. [Online]. Available: <http://homes.cs.washington.edu/~ccwu/vsfm/>
- [61] P. Moulon, P. Monasse, R. Marlet, and Others. *OpenMVG: An Open Multiple View Geometry Library*, accessed on Jan. 12, 2016. [Online]. Available: <https://github.com/openMVG/openMVG>
- [62] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Proc. CVPR*, 2012, pp. 3354–3361.
- [63] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *Proc. Int. Conf. Intell. Robot. Syst. (IROS)*, Oct. 2012, pp. 573–580.
- [64] D. Galvez-López and J. D. Tardos, "Bags of binary words for fast place recognition in image sequences," *IEEE Trans. Robot.*, vol. 28, no. 5, pp. 1188–1197, Oct. 2012.
- [65] Y. Li, N. Snavely, and D. P. Huttenlocher, "Location recognition using prioritized feature matching," in *Proc. 11th Eur. Conf. Comput. Vis.*, vol. 6312, 2010, pp. 791–804.
- [66] Y. Lou, N. Snavely, and J. Gehrke, "Matchminer: Efficient spanning structure mining in large image collections," in *Proc. 12th Eur. Conf. Comput. Vis.*, vol. 7573, 2012, pp. 45–58.
- [67] E. Rublee, V. Rabaud, K. Konolige, and G. R. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *Proc. IEEE Int. Conf. Comput. Vis.*, Nov. 2011, pp. 2564–2571.



Guofeng Zhang (M'07) received the B.S. and Ph.D. degrees in computer science from Zhejiang University in 2003 and 2009, respectively. He received the National Excellent Doctoral Dissertation Award and the Excellent Doctoral Dissertation Award of China Computer Federation. He is currently an Associate Professor with the State Key Laboratory of CAD&CG, Zhejiang University. His research interests include structure-from-motion, SLAM, 3D reconstruction, augmented reality, video segmentation and editing.



Haomin Liu (S'13) received the B.S. degree in computer science from Xidian University in 2009, and the master's degree in computer science from Zhejiang University in 2012. He is currently pursuing the Ph.D. degree in computer science with the State Key Laboratory of CAD&CG, Zhejiang University. His research interests include Structure-from-Motion, SLAM, and augmented reality.



Zilong Dong received the B.S. and Ph.D. degrees in computer science from Zhejiang University, China, in 2004 and 2010, respectively. From 2010 to 2013, he held a post-doctoral position with Zhejiang University. Now he is a senior algorithm engineer in Hangzhou. His main research interests include SLAM, augmented reality, video segmentation and 3D reconstruction.



Jiaya Jia (SM'09) received the Ph.D. degree in computer science from the Hong Kong University of Science and Technology in 2004. He is currently a Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong (CUHK). He heads the research group focusing on computational imaging, machine learning, practical optimization, and low and high-level computer vision. He currently serves as an Associate Editor of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE and served as an

Area Chair for several-year ICCV and CVPR. He was also on the technical paper program committees of SIGGRAPH, SIGGRAPH Asia, ICCP, and 3DV, and Co-Chaired the Workshop on Interactive Computer Vision, in conjunction with ICCV 2007. He received the Young Researcher Award 2008 and the Research Excellence Award 2009 from CUHK.



Tien-Tsin Wong (M'94) received the B.Sc., M.Phil. and Ph.D. degrees in computer science from The Chinese University of Hong Kong CUHK in 1992, 1994, and 1998, respectively. He is currently a Professor with the Department of Computer Science and Engineering, CUHK, and the Director of Digital Visual Entertainment Laboratory at Shenzhen Research Institute, CUHK. His main research interests include computer graphics, computational manga, perception graphics, image-based rendering, GPU techniques, medical visualization, and computer vision. He received the IEEE TRANSACTIONS ON MULTIMEDIA Prize Paper Award 2005 and the Young Researcher Award 2004.



Hujun Bao (M'14) received the B.S. and Ph.D. degrees in applied mathematics from Zhejiang University in 1987 and 1993, respectively. He is currently a Cheung Kong Professor with State Key Laboratory of CAD&CG, Zhejiang University. His main research interest is computer graphics and computer vision, including geometry and vision computing, real-time rendering and mixed reality.