

Efficient Initial Pose-graph Generation for Global SfM

Daniel Barath^{1,2,3}, Dmytro Mishkin¹, Ivan Eichhardt², Ilia Shipachev¹, and Jiri Matas¹

¹Visual Recognition Group, Faculty of Electrical Engineering, Czech Technical University in Prague

²MPLab, SZTAKI, Budapest ³Department of Computer Science, ETH Zurich

Abstract

We propose ways to speed up the initial pose-graph generation for global Structure-from-Motion algorithms. To avoid forming tentative point correspondences by FLANN and geometric verification by RANSAC, which are the most time-consuming steps of the pose-graph creation, we propose two new methods – built on the fact that image pairs usually are matched consecutively. Thus, candidate relative poses can be recovered from paths in the partly-built pose-graph. We propose a heuristic for the A* traversal, considering global similarity of images and the quality of the pose-graph edges. Given a relative pose from a path, descriptor-based feature matching is made “light-weight” by exploiting the known epipolar geometry. To speed up PROSAC-based sampling when RANSAC is applied, we propose a third method to order the correspondences by their inlier probabilities from previous estimations. The algorithms are tested on 402 130 image pairs from the IDSfM dataset and they speed up the feature matching 17 times and pose estimation 5 times. Source code: <https://github.com/danini/pose-graph-initialization>

1. Introduction

Structure-from-Motion (SfM) has been intensively researched in computer vision for decades. Most of the early methods adopt an incremental strategy, where the reconstruction is built progressively and the images are carefully added one-by-one in the procedure [42, 39, 38, 1, 57, 44]. Recent studies [17, 18, 8, 28, 4, 10, 20, 36, 13, 9, 58] show that global approaches, considering all images simultaneously when reconstructing the scene, lead to comparable or better accuracy than incremental techniques while being significantly more efficient. Also, global methods are less dependent on local decisions or image ordering.

Typically, SfM pipelines consist of the following steps, see Fig. 2. First, features are extracted in all n images. Such step is easily parallelizable and has $\mathcal{O}(n)$ time complexity. These features are then often used to order the image pairs from the most probable to match to the most




Figure 1: Reconstruction by initializing Theia’s [52] global SfM with the pose-graph from the proposed algorithms.

difficult ones, e.g., via bag-of-visual-words [47]. Next, tentative correspondences are generated between all image pairs by matching the often high-dimensional (e.g., 128 for SIFT [27]) descriptors of the detected features. Then, the correspondences are filtered and relative poses are estimated between by applying RANSAC [15]. Usually, the feature matching and geometric estimation steps are by far the slowest parts, both having quadratic complexity in the number of images. Moreover, feature matching has a quadratic worst-case time complexity as it depends on the product of the number of features in the respective images. Finally, a global bundle adjustment obtains the accurate reconstruction from the pair-wise poses. Interestingly, this step has negligible time demand, i.e., a few minutes in our experiments, compared to the initial pose-graph generation.

This paper has three major contributions – three new algorithms which allow removing the need of RANSAC-based geometric estimation and, also, to make descriptor-based feature matching “light-weight”. *First*, a method is proposed exploiting the partly-built pose-graph to avoid the computationally demanding RANSAC-based robust estimation. To do so, we propose a heuristic for the A* [19] algorithm which guides the path-finding even without having a metric distance between the views. The lack of such a distance originates from the fact that the edges of a pose-graph represent relative poses and, thus, neither the global scale nor the length of any of the translations are known. *Second*, we propose a technique to make the expensive




Figure 2: The architecture of a global SfM pipeline.

descriptor-based feature matching “light-weight” by using the pose determined by A^* . This guided matching approach uses the fundamental matrix to efficiently select keypoints, which lead to correspondences consistent with the pose, via hashing. *Third*, an algorithm is proposed to adaptively re-rank the point-to-point correspondences based on their history – whether one or both of the points had been inliers in previous estimations. The method exploits the fact that these inlying feature points likely represent 3D points consistent with the rigid reconstruction of the scene. This adaptive ranking speeds up the robust estimation by guiding PROSAC [11] to find a good sample early. The proposed techniques were tested on the IDSfM dataset [56], see Fig. 1 for an example reconstruction. They consistently and significantly speed up the pose-graph generation.

1.1. Related Techniques

Robust estimation. To speed up robust estimation, there has been a number of algorithms proposed over the years. NAPSAC [34], PROSAC [11] and P-NAPSAC [7] modify the RANSAC sampling strategy to increase the probability of selecting an all-inlier sample early. PROSAC exploits an a priori predicted inlier probability rank of the points and starts the sampling with the most promising ones. NAPSAC is built on the fact the real-world data often are spatially coherent and selects samples from local neighborhoods, where the inlier ratio is likely high. P-NAPSAC combines the benefits of PROSAC and NAPSAC by first sampling locally, and then progressively blending into global sampling. The Sequential Probability Ratio Test [12] (SPRT), inspired by Wald’s theory, is applied for rejecting models early if the probability of being better than the previous best model falls below a threshold. All of the mentioned RANSAC improvements consider the case of a single, isolated two-view robust estimation. Here, we exploit information arising while performing estimation on some subset of the $\binom{N}{2}$ image pairs where some images are matched more than once.

Feature matching can be sped up in several ways, *e.g.*, by the use of binary descriptors [43, 2, 54] or by limiting the number of features detected, as often done in SLAM systems [33]. However, this often results in inaccurate camera poses for the general 3D reconstruction problem [23]. Often, approximate nearest neighbor algorithms are employed,

Symbols used in this paper	
$\mathcal{G} = (\mathcal{V}, \mathcal{E})$	- Directed graph,
$v \in \mathcal{V}$	- A vertex from the vertex set
$e = (v_i, v_j) \in \mathcal{E}$	- Edge between vertices v_i and v_j
$f \in \{e, e^{-1}\}$	- An edge or its inverse
$\phi(e) : \mathcal{E} \rightarrow \text{SE}(3)$	- Relative pose of edge e
$\rho(e) : \mathcal{E} \rightarrow \mathbb{R}$	- Quality of edge e
$\delta(v_i, v_j) : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}$	- Distance of vertices v_i and v_j
$\mathcal{W} \in \{(f_1, \dots, f_n) \mid 1 < n\}$	- A walk
$\rho(\mathcal{W}) : \{(f_1, \dots, f_n)\} \rightarrow \mathbb{R}$	- Quality of walk \mathcal{W}

such as kd-tree or product quantization [32, 25]. Hardware-based speed-ups include using a GPU [24]. None of these techniques consider that the matching is performed on a number of image pairs, where the relative pose might be known, at least approximately, prior to the matching.

Global image similarity. Matching an *unordered* image collection is usually a harder and more time consuming task than, matching, *e.g.*, a video sequence. There are two reasons for that. First, many image pairs might not have any commonly visible part of the scene and the time spent on matching attempts is wasted. Moreover, no-match is the worst case scenario for RANSAC, which will run the maximum number of iterations, often orders of magnitude more than in the matching-possible case. Second, the time spent on the estimation of epipolar geometry highly depends on the inlier ratio of the tentative correspondences [14]. The inlier ratio, in turn, depends on the difference between the two viewpoints: the bigger the difference, the fewer tentative correspondences are correct [30, 29]. A natural question would be – is it possible to order the image pairs from the most probable ones to the most difficult or impossible to match? Image retrieval techniques are commonly used for it, *e.g.*, one could re-use extracted local features to find the most promising candidates for matching via bag-of-visual-words [47] and then quickly reorder the preliminary list using geometric constraints [37, 45] as it is implemented in COLMAP SfM [44]. Such systems work well, but have significant memory footprint and are now overcome by CNN-based global descriptors [41, 40, 55], which are both faster to compute and provide more accurate results.

We use the following approach to generate a fully connected image similarity graph as a preliminary step. First, we extract GeM [41] descriptors with ResNet-50 [22] CNN, pre-trained on GLD-v1 dataset [35]. Then we calculate the inner-product similarity between all the descriptors, resulting in an $n \times n$ similarity matrix. The calculation of the similarity matrix is the only quadratic step of our pipeline. However, the scalar product operation is extremely fast. In practice, the creation and processing of the similarity matrix takes negligible time.




Figure 3: A schematic pose-graph, used for initializing global SfM algorithms. Vertices (images) are connected by edges representing relative pose $(t_{ij}, \mathbf{R}_{ij}) \in \text{SE}(3)$.

2. Relative Pose from Directed Walks

We propose an approach to speed up the pose-graph generation by avoiding running RANSAC when possible. The core idea exploits the fact that when estimating the relative pose of the $(t + 1)$ th image pair from an image collection, we are given a pose-graph consisting of t edges, *i.e.*, t view pairs. These t edges can be used to estimate the relative pose without running RANSAC.

In the rest of the paper, we assume that the view pairs are ordered by their similarity scores. Thus, we start the pose estimation from the most similar pair. Assume that we have matched t image pairs successfully and, thus, we are given pose-graph $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t)$, where $\mathcal{E}_t \subseteq \{(v_1, v_2) \mid v_1, v_2 \in \mathcal{V}\}$ are the $|\mathcal{E}_t| = t$ edges and \mathcal{V} is the set of images in the dataset, see Fig. 3 for an example. Function $\phi : \mathcal{E}_t \rightarrow \text{SE}(3)$ maps edge $e \in \mathcal{E}_t$ to its estimated relative pose.

When estimating the relative pose between the $(t + 1)$ th view pair, we are given two options. The traditional one is to run robust estimation on the corresponding points between the two images. The estimated pose $\mathbf{P} \in \text{SE}(3)$ is then added to the pose-graph as the pose of the new edge. Thus, $\mathcal{E}_{t+1} = \mathcal{E}_t \cup \{e = (v_s, v_d)\}^1$ and $\phi(e) = \mathbf{P}$. The problem with this step is that when having few inliers and, thus, low inlier ratio, the estimation can be often time-consuming. Due to this step being done approximately $\binom{|\mathcal{V}|}{2}$ times, the slow pair-wise pose estimation has a *severe impact* on the processing time of the entire pose-graph estimation.

Therefore, instead of estimating the pose blindly between a pair of views (v_s, v_d) , we propose to use the previously generated pose-graph \mathcal{G}_t . Let us assume that there exists a finite directed walk $\mathcal{W} = (f_{w_1}, f_{w_2}, \dots, f_{w_{n-1}})$, for which there is a sequence of vertices $(v_{w_1}, v_{w_2}, \dots, v_{w_n})$ such that $f_{w_i} \in \{e_{w_i}, e_{w_i}^{-1}\}$, $e_{w_i} = (v_{w_i}, v_{w_{i+1}})$ for $i = 1, 2, \dots, n - 1$, and $v_{w_1} = v_s, v_{w_n} = v_d$. See Fig. 4 for examples. The direction of edge e can be inverted as e^{-1} by inverting the relative pose as $\phi(e_i^{-1}) = \phi(e_i)^{-1}$ and swap-

¹ v_s – source view, v_d – destination view

ping its vertices as $e_i^{-1} = (v_{i+1}, v_i)$. We define the pose implied by walk \mathcal{W} recursively as

$$\begin{aligned} \phi(\mathcal{W}) &= \phi(f_{w_1}, f_{w_2}, \dots, f_{w_{n-1}}) \\ &= \phi(f_{w_1}, f_{w_2}, \dots, f_{w_{n-2}})\phi(f_{w_{n-1}}) \\ &= \phi(f_{w_1}, f_{w_2}, \dots, f_{w_{n-3}})\phi(f_{w_{n-2}})\phi(f_{w_{n-1}}) \quad (1) \\ &= \dots \\ &= \phi(f_{w_1})\phi(f_{w_2}) \dots \phi(f_{w_{n-1}}). \end{aligned}$$

Consequently, the relative pose between views v_s and v_d is calculated as $\phi(\mathcal{W})$ given a finite walk \mathcal{W} .

The problem with (1) is that a single incorrectly estimated pose $\phi(f)$, $f \in \mathcal{W}$, makes the entire $\phi(\mathcal{W})$ wrong. Therefore, we aim at finding *multiple walks* within a given distance, *i.e.*, the maximum depth is restricted to avoid infinitely long walks. The walks returned are evaluated sequentially and immediately, see Alg. 1. Whenever a new walk \mathcal{W} is found, its inlier ratio is calculated from pose $\phi(\mathcal{W})$ and the correspondences between the source and destination images, v_s and v_d , respectively.

Termination. There are two cases when the procedure of finding and testing walks terminates. They are as follows:

1. The process finishes when there are no more walks found within the maximum distance.
2. If there is a reasonably good pose \mathbf{P} found, the process terminates. We consider a relative pose reasonably good if it has at least I_{\min} inliers.²

Pose refinement. In case the pose is obtained successfully from one of the walks, it is calculated solely from the edges of pose-graph \mathcal{G}_t without considering the correspondences between images v_s and v_d . In order to improve the accuracy and obtain \mathbf{P}^* , we apply iteratively re-weighted least-squares fitting initialized by the newly estimated model \mathbf{P} . Finally, $\mathcal{E}_{t+1} = \mathcal{E}_t \cup \{e = (v_s, v_d)\}$ and $\phi(e) = \mathbf{P}^*$.

Failures. There are cases when at least a single walk exists between views v_s and v_d , but the implied pose is incorrect, *i.e.*, it does not lead to a reasonable number of inliers. In those cases, we apply the traditional approach, *i.e.*, RANSAC-based robust estimation [6].

Visibility. Deciding if there is at least a single walk in the pose-graph between views v_s and v_d can be done by the union-find algorithm [16] in $\mathcal{O}(1)$ time. On average, the time complexity of the update is $\mathcal{O}(\log(n))$.

Parallel graph building. Building graph \mathcal{G} , checking the visibility, finding and evaluating walks in parallel on multiple CPUs is efficiently doable by using readers-writer locking mechanisms where each thread matches the next best view pair. The readers are the processes trying to find walks between two views or the ones checking if view v_s is visible from v_d . A process becomes writer only when it adds a new edge to the pose-graph or updates the union-find method for visibility checking which both takes only a few operations.

²Typically, $I_{\min} \in [15, 30]$ in most of the recent SfMs [52, 44].




Figure 4: Example walks between vertices v_s and v_d with edge inversion allowed. The relative pose $\mathbf{P}_{sd} = (\mathbf{t}_{sd}, \mathbf{R}_{sd}) \in \text{SE}(3)$ from v_s to v_d is calculated as $\mathbf{P}_{sd} = \mathbf{P}_{nd} \dots \mathbf{P}_{12} \mathbf{P}_{s1}$, where \mathbf{P}_{s1} is the pose of the first edge in the walk and \mathbf{P}_{nd} is that of the last one, n is the length of the walk.

Algorithm 1 Pose from Pose-Graph.

Input: \mathcal{G}_t – current pose-graph; v_s, v_d – views to match
 V – visibility table; \mathcal{P} – point correspondences
 d_{\max} – maximum depth
Output: \mathbf{P} – pose; \mathcal{I} – inliers

```

1: if  $\neg \text{Visible}(V, v_s, v_d)$  then
2:   return
3: while  $\neg \text{Terminate}(\mathcal{I})$  do
4:    $\mathcal{W} \leftarrow \text{GetNextWalk}(\mathcal{G}_t, v_s, v_d, d_{\max})$ 
5:   if  $\text{EmptyWalk}(\mathcal{W})$  then
6:     break
7:    $\mathbf{P}_{\mathcal{W}} \leftarrow \phi(\mathcal{W})$ 
8:    $\mathcal{I}_{\mathcal{W}} \leftarrow \text{GetInliers}(\mathbf{P}_{\mathcal{W}}, \mathcal{P})$ 
9:   if  $|\mathcal{I}_{\mathcal{W}}| > |\mathcal{I}|$  then
10:     $\mathbf{P} \leftarrow \mathbf{P}_{\mathcal{W}}, \mathcal{I} \leftarrow \mathcal{I}_{\mathcal{W}}$ 

```

Algorithm 2 Epipolar Hashing.

Input: $\mathcal{K}_1, \mathcal{K}_2$ – sets of keypoints; \mathbf{R}, \mathbf{t} – relative pose;
 μ – inlier-outlier threshold; b – bin number
Output: \mathcal{P} – point correspondences

```

1:  $\mathbf{E} = [\mathbf{t}]_{\times} \mathbf{R}$  ▷ Get essential matrix
2:  $\mathcal{B} \leftarrow \text{Hashing}(\mathcal{K}_2, \mathbf{E}, b)$ 
3: for  $\mathbf{p}_1 \in \mathcal{K}_1$  do
4:    $\mathbf{d}_1 \leftarrow \text{Descriptor}(\mathbf{p}_1), \delta \leftarrow \infty, \mathbf{p}^* \leftarrow \mathbf{0}$ 
5:   for  $\mathbf{p}_2 \in \mathcal{B}(\mathbf{p}_1)$  do
6:     if  $\epsilon(\mathbf{p}_1, \mathbf{p}_2, \mathbf{E}) < \mu$  then ▷ Sampson dist.
7:        $\mathbf{d}_2 \leftarrow \text{Descriptor}(\mathbf{p}_2)$ 
8:       if  $|\mathbf{d}_2 - \mathbf{d}_1| < \delta$  then
9:          $\delta \leftarrow |\mathbf{d}_2 - \mathbf{d}_1|, \mathbf{p}^* \leftarrow \mathbf{p}_2$ 
10:   $\mathcal{P} \leftarrow \mathcal{P} \cup \{(\mathbf{p}_1, \mathbf{p}^*)\}$ 

```

2.1. Pose-graph Traversal

It is a rather important question how to find a walk between views v_s and v_d efficiently. There are a number of

graph traversals, however, most of them are not suitable for returning walks in a large graph in reasonable time. We choose the A^* [19] algorithm since it works well for such a task when a good heuristic exists. In this section, we propose a way of obtaining multiple walks in pose-graph \mathcal{G}_t by defining a heuristic for the A^* algorithm.

The objective is to define a heuristic which guides the A^* algorithm from node v_s to node v_d while visiting as few vertices as possible. Since we are given a graph of relative poses, we are unable to define a metric, measuring the Euclidean distance of a view pair. When having relative poses, both the global and local scales remain unknown and, thus, all translations have unit length. As a consequence, it is unclear whether two views are close to or far from each other. The proposed heuristic is composed of two functions.

First, the global similarity of views v_s and v_d is measured as $\delta(v_s, v_d)$, $\delta : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}$. It is determined via the inner-product of GeM [41] descriptors with ResNet-50 [22] CNN, pre-trained on GLD-v1 dataset [35] as described earlier. Second, reflecting the fact that a single incorrectly estimated edge severely affects the pose of the entire walk, we also consider the quality of edge e via function $\rho(e) : \mathcal{E} \rightarrow \mathbb{R}$. To our experiments, the inlier ratio is usually a good indicator of the pose quality. Function $\rho(e)$ returns the inlier ratio calculated given the pose $\phi(e)$ of the current edge and the point correspondences.

To measure the quality of the entire walk \mathcal{W} , we have to consider that a single incorrect pose makes $\phi(\mathcal{W})$ incorrect as well. Thus, the quality of \mathcal{W} is measured as $Q(\mathcal{W}) = \min_{f \in \mathcal{W}} \rho(f)$, i.e., the quality of the least accurate edge. To measure the similarity of walk \mathcal{W} between the destination view v_d , we define function $\Delta(\mathcal{W}, v_d) = \max_{f=(v_1, v_2) \in \mathcal{W}} \delta(v_2, v_d)$, i.e., the most similar vertex determines the similarity. The heuristic considering both the quality of the walk and similarity to the destination is as

$$h(\mathcal{W}) = \lambda \min_{f \in \mathcal{W}} \rho(f) + (1 - \lambda) \max_{f=(v_1, v_2) \in \mathcal{W}} \delta(v_2, v_d), \quad (2)$$

where $\lambda \in [0, 1]$ is a weighting parameter. Expression




Figure 5: *Epipolar Hashing*. Point p_2 in image C_2 is assigned to the bin in C_1 (gray area) which its corresponding epipolar line l selects. The bins are defined on the angles of the epipolar lines in C_1 . The epipoles are e_1 and e_2 .

$\min_{f \in \mathcal{W}} \rho(f)$ forces the A^* algorithm to find a walk maximizing the minimum inlier ratio along the walk. Expression $\max_{f=(v_1, v_2) \in \mathcal{W}} \delta(v_2, v_d)$ affects the graph traversal in a way such that it maximizes the maximum similarity to the destination view along the path.

3. Guided Matching with Pose

When matching an image collection, the most time-consuming process is often the local descriptor matching [23] *a.k.a.* establishing tentative correspondences. The reason is that it has $\mathcal{O}(n^2)$ complexity both, *w.r.t.* the number of local features, *i.e.* for a single image pair, and *w.r.t.* the number of images, *i.e.* for the whole collection. We have already addressed the second problem (see Section 1.1), but matching a single image pair still takes a significant amount of time. The common way to accelerate feature matching is by using approximate nearest neighbor search, instead of the exact one, *e.g.* using the kd-tree algorithm as implemented in FLANN [32]. Yet, even the approximate matching still takes a considerable amount of time and decreases the accuracy of the camera pose [23]. We propose an alternative solution instead – to exploit the poses coming from walks in the current pose-graph to establish tentative correspondences. These poses will be used to make the standard descriptor matching “light-weight” by checking only those correspondences which are consistent with the pose.

Guided feature matching with pose. Let us assume that we are given sets of keypoints $\mathcal{K}_i, \mathcal{K}_j$ in the i th and j th images, respectively, and a relative pose $\mathbf{P}_{ij} = (\mathbf{t}_{ij}, \mathbf{R}_{ij}) \in \text{SE}(3)$ from the i th image to the j th one. One can easily calculate essential matrix $\mathbf{E}_{ij} = [\mathbf{t}_{ij}]_{\times} \mathbf{R}_{ij}$ and use it to measure the distance ϵ of point pairs via the Sampson Distance or the Symmetric Epipolar Error [21]. Therefore, the objective is to find pairs of points $(\mathbf{p}_i, \mathbf{p}_j)$, where $\mathbf{p}_i \in \mathcal{K}_i$, $\mathbf{p}_j \in \mathcal{K}_j$ and $\epsilon(\mathbf{p}_i, \mathbf{p}_j, \mathbf{E}_{ij})$ is smaller than the inlier-outlier threshold. In contrast to the traditional approach, where the feature matching is defined over the high-dimensional descriptor vectors of all possible keypoints using the L_2 norm,

we propose to select a small subset of candidate matches using the essential matrix. Consequently, the descriptor matching becomes significantly faster.

Due to doing the matching in 2D, the procedure can be done by hashing instead of a brute-force or approximated pair-wise process. Using the essential matrix, finding possible pairs of a point in the source image degrades to finding points in the destination one where the corresponding epipolar lines project to the correct position, *i.e.*, onto the selected point in the source image. Therefore, the points in the destination image can be put into bins according to their epipolar lines in the source image. A straightforward choice is to define the bins on the angles of the epipolar lines as it is visualized in Fig. 5. We call this technique in the further sections *Epipolar Hashing* (EH). Note that EH is applicable even when the intrinsic camera parameters are unknown and, thus, we only have a fundamental matrix.

Let us denote the angle of the corresponding epipolar line l in the first image of point (x, y) in the second image as $\alpha_{(x,y)} \in [0, \pi)$. Due to the nature of epipolar geometry, certain $\alpha_{(x,y)}$ angles are impossible. Therefore, we define the interval consisting of the valid angles and, thus, which we will cover by a number of bins as $[a, b]$, where

$$\begin{aligned} a &= \min(\alpha_{(0,0)}, \alpha_{(w_2,0)}, \alpha_{(0,h_2)}, \alpha_{(w_2,h_2)}), \\ b &= \max(\alpha_{(0,0)}, \alpha_{(w_2,0)}, \alpha_{(0,h_2)}, \alpha_{(w_2,h_2)}). \end{aligned}$$

Point $(0, 0)$ is the top-left corner of the second image, w_2 is its width, and h_2 is its height. When hashing the points, the size of a bin will be $\frac{b-a}{\#\text{bins}}$. This is an important step in practice since sometimes the epipole is far outside the image and, thus, the range of angles is < 1 . Without the adaptive bin size calculation, the algorithm does not speed up the matching in such cases. Note that $[a, b]$ is $[0, \pi)$ when the epipole falls inside the image. When doing the traditional descriptor matching, we consider only those matches which are in the corresponding bin and have lower Sampson distance than the threshold used for determining the pose.

After the guiding is performed, descriptor matching is done on 2 to 30 possible candidates instead of all keypoints. To further clean it up, we apply standard SIFT ratio test [27, 23] with adaptive ratio threshold, depending on number of nearest neighbors – the smaller the pool, the stricter the ratio test is. Details are added to the supplementary material.

The matching process is applied after A^* if that finds a good pose. Since A^* requires a set of correspondences to determine if a pose is reasonably good, we use correspondences from those point tracks where the current images are visible. The multi-view tracklets are calculated and updated when a new image pair is matched successfully. Since both global and incremental SfM algorithms require point tracks, this step does not add to the final processing time.

Note that [46] proposed a guided matching where the points are put in a 2D grid and each neighbor selection re-

Table 1: Run-time of pose estimation on 402 130 view pairs from the 1DSfM dataset using GC-RANSAC, breadth-first traversal and A* with the proposed heuristic (unit: seconds).

method	avg	med	total
GC-RANSAC [6]	0.815	0.915	327 574
Breadth-first	2.916	0.703	1 429 423
A*	0.173	0.056	82 672

quires iterating through K grid cells which the corresponding epipolar lines select, thus, implying $\mathcal{O}(K)$ complexity. The proposed EH operates in 1D and all neighbors are from a single cell implying $\mathcal{O}(1)$ complexity.

4. Adaptive Correspondence Ranking

In this section, we propose a strategy to adaptively set the weight of the point correspondences for PROSAC sampling [11] when doing pair-wise relative pose estimation in large-scale problems. PROSAC exploits an a priori predicted inlier probability rank of the points and starts the sampling with the most promising ones. Progressively, samples which are less likely to lead to the sought model are drawn. The main idea of the proposed algorithm is based on the fact that features detected in one image and matched to the other ones often appear multiple times when matching the image collection. Therefore, correspondences containing points which were inliers earlier are to be used first in the PROSAC sampling. Conversely, points that were outliers in the previous images should be drawn later.

Assume that we are given the t -th image pair to match with sets of keypoints $\mathcal{K}_i, \mathcal{K}_j$. Each keypoint \mathbf{p} , from either set, has score $s_{\mathbf{p}}^{(t)} \in [0, 1]$ for determining its outlier rank among all keypoints. After successfully estimating the pose \mathbf{P}_{ij} of the image pair, we are given the probability $P((\mathbf{p}, \mathbf{q}) | \mathbf{P}_{ij})$ of (\mathbf{p}, \mathbf{q}) being outlier given pose \mathbf{P}_{ij} , where (\mathbf{p}, \mathbf{q}) is a tentative correspondence, $\mathbf{p} \in \mathcal{K}_i, \mathbf{q} \in \mathcal{K}_j$. Probability $P((\mathbf{p}, \mathbf{q}) | \mathbf{P}_{ij})$ can be calculated, e.g., as in MSAC [53], MLESAC [53] or MAGSAC++ [7] from the point-to-model residuals assuming normal or χ^2 distributions. Since we do not know how probabilities $P(\mathbf{p} | \mathbf{P}_{ij})$ and $P(\mathbf{q} | \mathbf{P}_{ij})$ relate, we assume that \mathbf{p} and \mathbf{q} being consistent with the rigid reconstruction are independent events and, thus, $P((\mathbf{p}, \mathbf{q}) | \mathbf{P}_{ij}) = P(\mathbf{p} | \mathbf{P}_{ij})P(\mathbf{q} | \mathbf{P}_{ij})$. To be able to decompose probability $P((\mathbf{p}, \mathbf{q}) | \mathbf{P}_{ij})$, we assume that $P(\mathbf{p} | \mathbf{P}_{ij}) = P(\mathbf{q} | \mathbf{P}_{ij}) = \sqrt{P((\mathbf{p}, \mathbf{q}) | \mathbf{P}_{ij})}$. This probability is then used to update score $s_{\mathbf{p}}$ and $s_{\mathbf{q}}$ after the t -th image pair matched as $s_{\mathbf{p}}^{(t+1)} = s_{\mathbf{p}}^{(t)}P(\mathbf{p} | \mathbf{P}_{ij})$ and $s_{\mathbf{q}}^{(t+1)} = s_{\mathbf{q}}^{(t)}P(\mathbf{q} | \mathbf{P}_{ij})$. Let us set $s_{\mathbf{p}}^{(0)} = 1$ since all keypoints are similarly likely to be outliers in the beginning.

When the $(t + 1)$ th image pair is matched by using PROSAC sampling, the correspondences are ordered ac-

Table 2: Run-time of matchers used for forming tentative correspondences with and without exploiting the relative pose in the proposed way (unit: seconds).

matcher	avg		med	
	w/o pose	pose	w/o pose	pose
Brute-force	7.609	1.078	1.047	1.139
FLANN [32]	0.992	0.318	0.728	0.137
Epipolar Hashing	–	0.057	–	0.046

cording to their outlier ranks $s_{\mathbf{p}}^{(0)}$ increasingly, such that the first one is the least likely to be an outlier.

5. Experiments

We tested the proposed algorithms on the 1DSfM dataset [56]. It consists of 13 scenes of landmarks with photos of varying sizes collected from the internet. 1DSfM provides 2-view matches with epipolar geometries and a reference reconstruction from incremental SfM (computed with Bundler [48, 49]) for measuring error. We used the SIFT features [27] as implemented in OpenCV with Root-SIFT [3] descriptors. In each image, 8000 keypoints are detected in order to have a reasonably dense point cloud reconstruction and precise pair-wise geometry camera poses [23]. We combined mutual nearest neighbor check with standard distance ratio test [26] to establish tentative point correspondences, as recommended in [23]. The bin number for Epipolar Hashing was set to 45. We matched all image pairs with global similarity higher than 0.4 with which we got accurate reconstruction in reasonable time. This leads to using only 1–2% (402 130 in total) of all pairs.

The methods are implemented in C++ using the Eigen and Sophus [51] libraries. The graph traversal algorithms are implemented by ourselves. For robust estimation, we always use the GC-RANSAC algorithm [6] with the five-point algorithm of Stewenius *et al.* [50]. Note that the used GC-RANSAC implementation contains PROSAC sampling [11], SPRT test [12] and a number of sample and model degeneracy tests to be as efficient as possible.

Alternatives for RANSAC. Pose-graph generation algorithms are compared in this section, including the proposed A*-based technique. The compared methods are:

1. The standard exhaustive matching (EM) where each tested image pair is matched by FLANN + GC-RANSAC.
2. A minimal spanning tree (MST) where the global similarity score is used as weights.
3. The proposed A*-based technique, where the pose comes from a path determined by A* if possible. Otherwise, the standard matching is applied.
4. Breadth-first (BF) traversal applied in the same way as the proposed A* algorithm.




Figure 6: The cumulative distribution functions of the errors (in degrees) and processing times (in seconds) of the pose-graphs generated by different initialization techniques on all scenes of the 1DSfM dataset. All algorithms returned 402 130 poses.

Table 3: The results of a global SfM [52] averaged over the scenes from the 1DSfM dataset [56]. The SfM is initialized with pose-graphs generated by the traditional exhaustive matching (EM), breadth-first-based (BF), A*-based (A*), and Spanning-tree-based (MST) pose-graph building. The reported properties are: pose-graph generator method (1st), number of views (2nd) and multi-view tracks (3rd) reconstructed by the SfM given an initial pose-graph; the total time of the pose-graph generation and bundle adjustment on a single CPU (4th); rotation error of the reconstructed poses in degrees (5th), position error in meters (6th) and focal length errors (7th).

			time	orientation err (°)			position err (m)			focal len. err ($\times 10^{-2}$)		
	# views	# tracks	(hours)	AVG	MED	STD	AVG	MED	STD	AVG	MED	STD
EM	437	76 837	202	9.8	6.6	6.2	10.9	8.4	13.9	2.0	1.2	3.7
BF	463	73 864	507	7.7	7.0	7.1	10.7	6.1	18.8	2.2	1.2	3.9
A* + FLANN	439	69 312	58	8.6	6.9	6.6	11.2	6.6	19.4	2.2	1.2	3.8
A* + EH	444	78 335	29	7.8	6.1	5.2	10.8	6.2	19.5	2.4	1.4	4.1
MST	84	10 334	4	25.4	10.6	7.6	12.5	23.0	7.9	2.4	1.5	3.5

The CDFs of the rotation and translation errors (angles) and run-times (seconds) are shown in Fig. 6. We do not include MST since it matches significantly fewer image pairs (9922) than the others (402 130). The A*-based technique leads to the most accurate rotation matrices while having similar translation errors as the breadth-first-based and exhaustive matching. The proposed A*-based technique leads to a significantly faster pose-graph generation than EM or BF. Note that the break-points in the run-time curves are caused by setting the maximum number of RANSAC iterations to 5000 to avoid extremely long runs.

The average, median and total processing times of the pair-wise pose estimation algorithms are shown in Table 1. The run-times contain those cases as well when no valid pose was found by the A* or breadth-first traversals and, thus, GC-RANSAC was applied to recover the pose. The A* algorithm leads to a speedup of almost an order of magnitude with its median time being approx. 20 times lower than that of GC-RANSAC. It validates the proposed heuristic that the breadth-first algorithm is significantly slower than A*. Consequently, the proposed heuristic guides the path-finding in the pose-graph successfully.

Matching with Pose. We compare the feature match-

ing speed with or without exploiting the pose determined by the A* algorithm. In FLANN and brute-force matching, this means that we find all candidate matches which lead to smaller epipolar error than the inlier-outlier threshold. The best candidate is then selected by descriptor-based matching. The run-times are reported in Table 2. Using the pose speeds up both the FLANN-based and brute-force algorithm significantly. The proposed Epipolar Hashing leads to a more than 17 times speedup compared to the traditional FLANN-based feature matching. Also, by the Epipolar Hashing, the neighbors are found precisely without approximation as done in FLANN.

On scene Madrid Metropolis, the number of reconstructed 3D points is 9486 without and 29 665 with using EH to find more correspondences.

Adaptive Ranking. The avg., median and total run-times (seconds) of the robust estimation using different correspondence ranking techniques for PROSAC are shown in Table 4. Three methods are compared: the uniform matching from RANSAC (unordered); PROSAC when the correspondences are ordered according to their SIFT ratios [11]; and the proposed adaptive re-ranking considering the prior information about the points from earlier es-

Table 4: Run-time of PROSAC using different ordering techniques: unordered (RANSAC-like uniform [15]), SIFT ratio [11] and the proposed adaptive one (unit: seconds).

ordering	avg	med	total
Unordered	0.736	0.768	295 997
SIFT ratio	0.664	0.698	266 849
Adaptive ordering	0.615	0.643	247 320

timations. While ordering the correspondences according to their SIFT ratios speeds up the estimation by 10 % compared to the uniform sampling, the proposed adaptive re-ranking leads to an additional 8 % speedup on average.

Applying Global SfM Algorithm. Once relative poses are estimated they are fed to the Theia library [52] that performs global SfM [10, 56]. That is, feature extraction, image matching and relative pose estimation were performed by our code either using the proposed algorithm or the traditional brute-force pair-wise matching using the 5PT [50] solver. The key steps of global SfM are robust orientation estimation, proposed by Chatterjee *et al.* [10], followed by robust nonlinear position optimization by the method of Wilson *et al.* [56]. The estimation of global rotations and positions enables triangulating 3D points, and the reconstruction is finalized by the bundle adjustment of camera parameters and point coordinates. Since the reconstruction always failed on scene Gendarmenmarkt, we did not consider that scene when calculating the errors.

Table 3 reports the results of Theia initialized by pose-graphs generated by the exhaustive matching (EM), breadth-first graph traversal (BF), the proposed A*-based graph-traversal, and by using a minimum spanning tree (MST). While the generation of the minimum spanning tree-based pose-graph is extremely fast, it can be seen that it is not good enough for a global SfM algorithm to provide a reconstruction of reasonable size. The avg. number of views reconstructed when initialized by MST is significantly lower than using other techniques. The proposed A*-based methods lead to similar number of views and similar error to the traditional approach. Note that the proposed method extremely cheaply creates cycles that are important for having stable results with a global SfM [31, 5].

Heuristic for A* Traversal. The A* traversal parameters are tuned on scene Alamo. For this purpose, the ground truth pose-graph is loaded and the pose is obtained by A* between all image pairs without direct connection. The parameters tuned are weight λ from (2) and the maximum depth allowed when obtaining the walks. They were then used for all other tests and scenes.

In Fig. 7, the avg. number of nodes visited by A* (left) and the ratio of accurate poses obtained (right) are plotted as the function of λ . Parameter $\lambda = 0$ means that there




Figure 7: Average # of nodes visited by A* (left) and ratio of accurate relative poses obtained (right) in scene Alamo plotted as the function of weight λ and the max. depth.

is no constraint on the edge quality, the only goal is to get to a node similar to the destination. Parameter $\lambda = 1$ is interpreted as walking on the highest quality edges without trying to get close to the destination. The number of nodes visited, *i.e.* proportional to the processing time, is nearly constant for $\lambda \in [0, 0.8]$. The highest success rate is achieved by $\lambda = 0.8$ and $\text{Depth}_{\max} = 5$. Since $\lambda = 0.8$ with $\text{Depth}_{\max} = 5$ also leads to a reasonably low number of nodes visited, we chose these values in all our experiments. Additional experiments are in the supplementary material.

6. Conclusions

The final bundle adjustment of global SfM algorithms has a negligible time demand compared to the initial pose-graph generation. To speed this step up by almost an order of magnitude, we proposed three new algorithms. The standard procedure (*i.e.*, feature matching by FLANN; pose estimation by RANSAC-like robust estimation) for estimating the pose-graph for all scenes from the 1DSfM dataset took a total of 726 487 seconds on a single CPU – approx. 202 hours. By using the proposed set of algorithms (*i.e.*, A*-based pose estimation; Epipolar Hashing for matching; adaptive re-ranking), the total run-time is reduced to 105 593 seconds (29 hours). In the experiments, A* found a valid pose in 93.8% of the image pairs. Thus, traditional FLANN-based feature matching and pose estimation by RANSAC was applied only to 6.2% of the image pairs.

Acknowledgements. J. Matas, D. Mishkin and D. Barath were supported by OP VVV funded project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics”. D. Barath was also supported by the Ministry of Innovation and Technology NRDI Office within the framework of the Autonomous Systems National Laboratory Program. D. Mishkin and I. Shipachev was supported by CTU student grant SGS20/171/OHK3/3T/13 and by the Austrian Ministry for Transport, Innovation and Technology, the Federal Ministry for Digital and Economic Affairs, and the Province of Upper Austria in the frame of the COMET center SCCH.

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. [1](#)
- [2] Alexandre Alahi, Raphaël Ortiz, and Pierre Vanderghenst. FREAK: Fast retina keypoint. *Conference on Computer Vision and Pattern Recognition*, 2012. [2](#)
- [3] Relja Arandjelovic and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *Conference on Computer Vision and Pattern Recognition*, pages 2911–2918, 2012. [6](#)
- [4] Mica Arie-Nachimson, Shahar Z Kovalsky, Ira Kemelmacher-Shlizerman, Amit Singer, and Ronen Basri. Global motion estimation from point matches. In *Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission*, pages 81–88. IEEE, 2012. [1](#)
- [5] Federica Arrigoni, Andrea Fusiello, and Beatrice Rossi. On computing the translations norm in the epipolar graph. In *International Conference on 3D Vision*, pages 300–308. IEEE, 2015. [8](#)
- [6] Daniel Barath and Jiří Matas. Graph-cut RANSAC. In *Conference on Computer Vision and Pattern Recognition*, pages 6733–6741, 2018. [3](#), [6](#)
- [7] Daniel Barath, Jana Noskova, Maksym Ivashechkin, and Jiri Matas. MAGSAC++, a fast, reliable and accurate robust estimator. In *Conference on Computer Vision and Pattern Recognition*, pages 1304–1312, 2020. [2](#), [6](#)
- [8] Matthew Brand, Matthew Antone, and Seth Teller. Spectral solution of large-scale extrinsic camera calibration as a graph embedding problem. In *European Conference on Computer Vision*, pages 262–273. Springer, 2004. [1](#)
- [9] Luca Carlone, Roberto Tron, Kostas Daniilidis, and Frank Dellaert. Initialization techniques for 3d slam: a survey on rotation estimation and its use in pose graph optimization. In *IEEE International Conference on Robotics and Automation*, pages 4597–4604. IEEE, 2015. [1](#)
- [10] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *International Conference on Computer Vision*, pages 521–528, 2013. [1](#), [8](#)
- [11] Ondrej Chum and Jiri Matas. Matching with PROSAC-progressive sample consensus. In *Conference on Computer Vision and Pattern Recognition*, volume 1, pages 220–226. IEEE, 2005. [2](#), [6](#), [7](#), [8](#)
- [12] Ondřej Chum and Jiří Matas. Optimal randomized RANSAC. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(8):1472–1482, 2008. [2](#), [6](#)
- [13] Zhaopeng Cui and Ping Tan. Global structure-from-motion by similarity averaging. In *International Conference on Computer Vision*, pages 864–872, 2015. [1](#)
- [14] Jana Noskova Daniel Barath, Jiri Matas. MAGSAC: marginalizing sample consensus. In *Conference on Computer Vision and Pattern Recognition*, 2019. [2](#)
- [15] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [1](#), [8](#)
- [16] Bernard A Galler and Michael J Fisher. An improved equivalence algorithm. *Communications of the ACM*, 7(5):301–303, 1964. [3](#)
- [17] Venu Madhav Govindu. Combining two-view constraints for motion estimation. In *Conference on Computer Vision and Pattern Recognition*, volume 2, pages II–II. IEEE, 2001. [1](#)
- [18] Venu Madhav Govindu. Lie-algebraic averaging for globally consistent motion estimation. In *Conference on Computer Vision and Pattern Recognition*, volume 1, pages I–I. IEEE, 2004. [1](#)
- [19] Peter Hart, Nils Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. [1](#), [4](#)
- [20] Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International Journal of Computer Vision*, 103(3):267–305, 2013. [1](#)
- [21] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. [5](#)
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, 2016. [2](#), [4](#)
- [23] Yuhe Jin, Dmytro Mishkin, Anastasiia Mishchuk, Jiri Matas, Pascal Fua, Kwang Moo Yi, and Eduard Trulls. Image matching across wide baselines: From paper to practice. *International Journal of Computer Vision*, 2020. [2](#), [5](#), [6](#)
- [24] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017. [2](#)
- [25] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128, 2011. [2](#)
- [26] David Lowe. Object recognition from local scale-invariant features. In *International Conference on Computer Vision*. IEEE, 1999. [6](#)
- [27] David Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. [1](#), [5](#), [6](#)
- [28] Daniel Martinec and Tomas Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [1](#)
- [29] Krystian Mikolajczyk and Cordelia Schmid. A Performance Evaluation of Local Descriptors. In *Conference on Computer Vision and Pattern Recognition*, pages 257–263, June 2003. [2](#)
- [30] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, Jiri Matas, Frederik Schaffalitzky, Timor Kadir, and Luc Van Gool. A comparison of affine region detectors. *International Journal of Computer Vision*, 65(1):43–72, 2005. [2](#)
- [31] Pierre Moulon, Pascal Monasse, and Renaud Marlet. Global fusion of relative motions for robust, accurate and scalable

- structure from motion. In *International Conference on Computer Vision*, pages 3248–3255, 2013. 8
- [32] Marius Muja and David Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Applications*, 2009. 2, 5, 6
- [33] Raúl Mur-Artal, José M. M. Montiel, and J. Tardós. ORB-SLAM: A Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics*, 31(5):1147–1163, 2015. 2
- [34] R. D. Myatt, Philip Torr, Slawomir Nasuto, John Bishop, and R. Craddock. NAPSAC: High noise, high dimensional robust estimation - it’s in the bag. In *British Machine Vision Conference*, 2002. 2
- [35] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-scale image retrieval with attentive deep local features. In *Conference on Computer Vision and Pattern Recognition*, pages 3456–3465, 2017. 2, 4
- [36] Onur Ozyesil and Amit Singer. Robust camera location estimation by convex programming. In *Conference on Computer Vision and Pattern Recognition*, pages 2674–2683, 2015. 1
- [37] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object Retrieval with Large Vocabularies and Fast Spatial Matching. In *Conference on Computer Vision and Pattern Recognition*, 2007. 2
- [38] Marc Pollefeys, David Nistér, Jan-Michael Frahm, Amir Akbarzadeh, Philippos Mordohai, Brian Clipp, Chris Engels, David Gallup, S-J Kim, Paul Merrell, et al. Detailed real-time urban 3d reconstruction from video. *International Journal of Computer Vision*, 78(2-3):143–167, 2008. 1
- [39] Marc Pollefeys, Luc Van Gool, Maarten Vergauwen, Frank Verbiest, Kurt Cornelis, Jan Tops, and Reinhard Koch. Visual modeling with a hand-held camera. *International Journal of Computer Vision*, 59(3):207–232, 2004. 1
- [40] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [41] Filip Radenović, Giorgos Tolias, and Ondrej Chum. Fine-tuning CNN image retrieval with no human annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 2, 4
- [42] Carsten Rother. *Multi-view reconstruction and camera recovery using a real or virtual reference plane*. PhD thesis, Numerisk analys och datalogi, 2003. 1
- [43] Ethan Rublee, Vincent Rabaud, Kurt Konolidge, and Gary Bradski. ORB: An Efficient Alternative to SIFT or SURF. In *International Conference on Computer Vision*, 2011. 2
- [44] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. 1, 2, 3
- [45] Johannes Lutz Schönberger, True Price, Torsten Sattler, Jan-Michael Frahm, and Marc Pollefeys. A vote-and-verify strategy for fast spatial verification in image retrieval. In *Asian Conference on Computer Vision*, 2016. 2
- [46] Rajvi Shah, Vanshika Srivastava, and PJ Narayanan. Geometry-aware feature matching for structure from motion applications. In *Winter Conference on Applications of Computer Vision*, pages 278–285. IEEE, 2015. 5
- [47] J. Sivic and A. Zisserman. Video Google: A Text Retrieval Approach to Object Matching in Videos. In *ICCV*, 2003. 1, 2
- [48] Noah Snavely, Steve Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. In *ACM Trans. Gr.*, volume 25, pages 835–846. ACM, 2006. 6
- [49] Noah Snavely, Steve Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *International Journal of Computer Vision*, 80(2):189–210, 2008. 6
- [50] H. Stewenius, C. Engels, and D. Nistér. Recent developments on direct relative orientation. *Journal of Photogrammetry and Remote Sensing*, 60(4):284–294, 2006. 6, 8
- [51] Hauke Strasdat. Sophus library. <https://github.com/strasdat/Sophus>. 6
- [52] Chris Sweeney. Theia multiview geometry library. <http://theia-sfm.org>. 1, 3, 7, 8
- [53] Philip H. S. Torr and Andrew Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Computer vision and image understanding*, 78(1):138–156, 2000. 6
- [54] Tomasz Trzcinski, Christos Marios Christoudias, and Vincent Lepetit. Learning image descriptors with boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3), 2015. 2
- [55] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval. In *Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [56] K. Wilson and N. Snavely. Robust Global Translations with IDSfM. In *European Conference on Computer Vision*, pages 61–75, 2014. 2, 6, 7, 8
- [57] Changchang Wu. Towards linear-time incremental structure from motion. In *International Conference on 3D Vision*, pages 127–134. IEEE, 2013. 1
- [58] Siyu Zhu, Runze Zhang, Lei Zhou, Tianwei Shen, Tian Fang, Ping Tan, and Long Quan. Very large-scale global sfm by distributed motion averaging. In *Conference on Computer Vision and Pattern Recognition*, pages 4568–4577, 2018. 1