

Finding Geometric Models by Clustering in the Consensus Space

Daniel Barath¹, Denys Rozumnyi^{2,1}, Ivan Eichhardt^{3,4}, Levente Hajder³, Jiri Matas²

¹Computer Vision and Geometry Group, ETH Zurich, Switzerland,

²VRG, Faculty of Electrical Engineering, CTU in Prague, Czech Republic,

³TMEIC Corporation Americas, Roanoke, VA, USA ⁴Eötvös Loránd University, Budapest, Hungary

Abstract

We propose a new algorithm for finding an unknown number of geometric models, e.g., homographies. The problem is formalized as finding dominant model instances progressively without forming crisp point-to-model assignments. Dominant instances are found via a RANSAC-like sampling and a consolidation process driven by a model quality function considering previously proposed instances. New ones are found by clustering in the consensus space. This new formulation leads to a simple iterative algorithm with state-of-the-art accuracy while running in real-time on a number of vision problems – at least two orders of magnitude faster than the competitors on two-view motion estimation. Also, we propose a deterministic sampler reflecting the fact that real-world data tend to form spatially coherent structures. The sampler returns connected components in a progressively densified neighborhood-graph. We present a number of applications where the use of multiple geometric models improves accuracy. These include pose estimation from multiple generalized homographies; trajectory estimation of fast-moving objects; and we also propose a way of using multiple homographies in global SfM algorithms. Source code: <https://github.com/danini/clustering-in-consensus-space>.

1. Introduction

Robust multi-instance model fitting is the problem of interpreting a set of data points as a mixture of noisy observations stemming from multiple instances of geometric models. Examples for such a problem are the estimation of plane-to-plane correspondences (i.e., homography matrices) in two images, and the retrieval of rigid motions in a dynamic scene captured by a moving camera. In the state-of-the-art algorithms, finding an unknown number of model instances is achieved by clustering the data points into disjoint sets, each representing a particular model instance. Robustness is achieved by considering an outlier model.

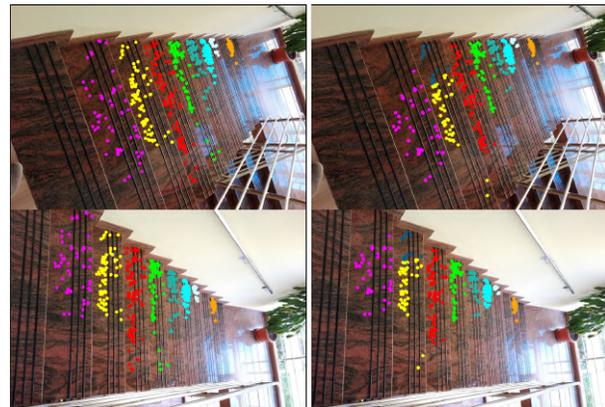


Figure 1. Multi-homography fitting with the proposed method in 0.04 secs (left), and with Prog-X [4] in 1.48 secs (right). Prog-X is one of the fastest SOTA algorithm. Outliers are not drawn.

Multi-instance model fitting has been studied since the early sixties. The Hough-transform [22, 23] is perhaps the first popular method for finding multiple instances of a single class [18, 37, 44, 67]. The RANSAC [16] algorithm was as well extended to deal with finding multiple instances. Sequential RANSAC [25, 60] detects instances in a sequential manner by repeatedly running RANSAC to recover a single instance and, then, removing its inliers from the point set. The greedy approach that makes RANSAC a powerful tool for recovering a single instance becomes its drawback when estimating multiple ones. Points are assigned not to the best but to the first instance, typically the one with the largest support, for which they cannot be deemed outliers. MultiRANSAC [71] forms compound hypotheses about n instances. In each iteration, MultiRANSAC draws samples of size n times m , where m is the number of points required for estimating a model instance, e.g., $m = 4$ for homographies. Besides requiring the number n of the instances to be known a priori, the increased sample size affects the problem complexity and, thus, the processing time severely.

Modern approaches for multi-model fitting [1, 3, 24, 32–34, 40, 62, 66] follow a two-step procedure. First, they gen-



Figure 2. **Left:** A case when assigning points to a single line (by color) prevents finding all 9 visible instances. Dashed black lines are not recovered. When fitting planes to 4 out of the 7 points, only a single plane can be found. **Middle, Right:** Examples where the point-to-model assignment fails at the intersection of planes.

erate many instances by repeatedly selecting minimal point sets and fitting model instances. Second, a subset of the hypotheses is selected interpreting the input data points the most. This selection is done in various ways. For instance, a popular group of methods [1, 3, 24, 40] optimizes point-to-model assignments by energy minimization using graph labeling techniques [8]. The energy originates from point-to-model residuals, label costs [14], and geometric priors [40] such as the spatial coherence of the data points. Another group of methods uses preference analysis based on the distribution of the residuals of data points [32–34, 68]. Also, there are techniques [62, 63, 69] approaching the problem as hyper-graph partitioning where the instances are represented by vertices, and the points by hyper-edges.

Prog-X [4] and CONSAC [26] discussed that the first, instance generation, step of the mentioned methods leads to a number of issues, *e.g.*, the instances are generated blindly, having no information about the data at hand. This approach severely restricts the out-of-the-box applicability of such techniques since the user either has to consider the worst-case scenario and, thus, generate an unnecessarily high number of instances; or requires some rule of thumb, *e.g.*, to generate twice the point number hypotheses that provides no guarantees of finding the sought instances. Prog-X approaches the problem via interleaving the model proposal and optimization steps. CONSAC further improves it by using a deep-learning-based guided sampling approach.

A common point of *all* state-of-the-art algorithms is formalizing the multi-model fitting problem as finding disjoint sets of data points each representing a model instance. There are two main practical issues with this assumption. First, in some cases, a point belongs to multiple instances and this assumption renders the problem *unsolvable*, see the left image of Fig. 2. Also, the point-to-model assignment is often unclear even if it is done by a human, especially, for points around the intersection of instances, see the right two plots of Fig. 2 for examples. The second issue stems from the recovery of disjoint point sets that usually requires a rather complex procedure, *e.g.* labeling via energy minimization, that affects the run-time severely.

The *main contribution* of this paper is a fundamentally

new problem formulation that does not require forming crisp point-to-model assignments, *i.e.*, a point can be assigned to multiple instances. This is different from the formulations used in the state-of-the-art algorithms for general multi-model fitting [1, 3, 4, 24, 26, 34, 40, 63]. This property allows the proposed method to be a simple iterative algorithm and, yet, to obtain results superior to the state-of-the-art both in terms of accuracy and run-time, being *real-time* on a number of problems, see Fig. 1, including ones where multi-model fitting algorithms generally are *not* real-time, *e.g.*, two-view motion detection. Also, this assumption relaxes the greedy nature of sequential algorithms as the ordering in which the instances are proposed becomes unimportant. As the *second* contribution, we discuss ways of exploiting multiple instances in popular applications – Structure-from-Motion, pose estimation for generalized and pin-hole cameras, and trajectory estimation of fast-moving objects. By considering multiple models, the accuracy is increased in almost all cases on several publicly available real-world datasets. As the *third* contribution, we propose a new sampler designed specifically for multi-instance model fitting. The sampler considers that real-world data tend to form spatially coherent structures. It returns the connected components in a gradually densified neighborhood-graph. While several samplers exist that exploit spatial properties of the data, *e.g.* [6, 38], the proposed one is *deterministic*.

2. Iterative Clustering in the Consensus Space

We propose a new algorithm for robust multi-instance model fitting that combines the advantages of state-of-the-art algorithms and, also, follows a new formulation that does not require crisp point-to-model assignments for finding the dominant model instances.

2.1. Idea and Schematic Algorithm

The proposed method is motivated by two observations about the nature of multi-model fitting problems. First, even though all of the state-of-the-art algorithms [1, 3, 4, 24, 26, 34, 40, 63] formalize the problem as a clustering where a set of data points (cluster) represents a model instance, this assumption is incorrect in a number of real-world scenes. Moreover, one of the primary reasons of multi-model fitting algorithms often being fairly slow stems from the optimization techniques, *e.g.* α -expansion in PEARL [24], needed to solve the point-to-model assignment problem.

Our second observation is that multi-model fitting can usually be rephrased as the problem of finding multiple *dominant* instances that are reasonably *different*. Ideally, a *dominant* instance is one that represents a real structure. Since this is not an algorithmically measurable property, we define being dominant as having a reasonably large support not shared with other dominant instances. We consider instances *different* if they are “far” on the model manifold

as proposed in Multi-X [3]. This simple formulation allows us to avoid applying complex procedures finessing to interpret point-point, model-model, and point-model interactions. Also, it further relaxes the greedy nature of the progressive model proposal strategy introduced in Prog-X [4] that enables to discover the data gradually. The pseudo-code of formalizing the multi-model problem as finding different dominant model instances is as follows:

Input: \mathcal{P} – data points
Output: \mathcal{I} – model instances

```

 $\mathcal{I} \leftarrow \emptyset$ 
while  $\neg$ Terminate() do
   $\mathcal{I} \leftarrow \mathcal{I} \cup \text{FindDominantInstances}(\mathcal{P})$ 
  while  $\neg$ Convergence() do
     $\mathcal{I} \leftarrow \text{SelectUniqueInstances}(\mathcal{I})$ 
     $\mathcal{I} \leftarrow \text{ImproveParameters}(\mathcal{I}, \mathcal{P})$ 

```

2.2. Finding Dominant Model Instances

Given a set of data points \mathcal{P} and a set of dominant model instances \mathcal{I} proposed in earlier iterations, the objective is to find a new dominant model instance $h \in \mathbb{R}^d$ which should be included in \mathcal{I} , where $d \in \mathbb{R}$ is the model dimension. In the first iteration, $\mathcal{I} = \emptyset$.

To do so, we start similarly to RANSAC by first drawing a random sample \mathcal{S} of data points. This is done by a state-of-the-art sampler, e.g., PROSAC [11] or P-NAPSAC [6]. Model instance h is estimated from sample \mathcal{S} . In order to decide about h being dominant or not, we define model quality function $Q : \mathbb{R}^d \times \mathcal{P}^* \times \mathbb{R} \rightarrow \mathbb{R}$ similarly as [4] to be calculated from the inliers of h not shared with other instances in \mathcal{I} , where \mathcal{P}^* is the power set of \mathcal{P} . Considering the RANSAC-like inlier counting, the implied quality is

$$Q_{\text{RSC}}(h, \mathcal{P}, \epsilon) = \sum_{\mathbf{p} \in \mathcal{P}} [\phi(h, \mathbf{p}) < \epsilon \wedge \phi(\mathcal{I}, \mathbf{p}) \geq \epsilon], \quad (1)$$

where $\epsilon \in \mathbb{R}$ is the inlier-outlier threshold and $\phi(\mathcal{I}, \mathbf{p}) = \min_{h \in \mathcal{I}} \phi(h, \mathbf{p})$ is the minimal point-to-model residual of point \mathbf{p} given the kept set of dominant instances \mathcal{I} . In order to use the recent advances of RANSAC, e.g. the loss function of MAGSAC++ [6] the currently most accurate method according to a recent survey [31], Q_{RSC} is reformulated considering a continuous loss function f . For practical reasons, we consider losses returning a value in-between 0 and 1. The implied quality function is

$$Q_f(h, \mathcal{P}, \epsilon) = |\mathcal{P}| - \sum_{\mathbf{p} \in \mathcal{P}} \max(f(h, \mathbf{p}), 1 - f(\mathcal{I}, \mathbf{p})), \quad (2)$$

where $f(\mathcal{I}, \mathbf{p}) = \min_{h \in \mathcal{I}} f(h, \mathbf{p})$ is the minimum loss of point \mathbf{p} given the set of kept instances \mathcal{I} . It can be easily seen that this quality function returns high score to those instances which do not share inliers with any of the instances

from \mathcal{I} . Otherwise, the quality is reduced according to the number and residuals of the inliers shared.

To determine whether instance h is dominant, we introduce parameter q_{\min} , and all model instances are considered dominant where $Q_f(h, \mathcal{P}, \epsilon) \geq q_{\min}$. This constraint can be interpreted as a lower bound for the number of perfectly fitting data points which are not shared with any of the instances from the maintained set in \mathcal{I} .¹

2.3. Clustering in the Consensus Space

The next step of the algorithm, after a set \mathcal{I} of dominant model instances have been found, is to select a subset of \mathcal{I} consisting of instances that represent different model instances and not noisy observations of the same one. We define a model-to-model residual function $\psi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ measuring the distance of two model instances.

Model-to-model residual. Defining a model-to-model residual function is a challenging problem. In the Multi-X algorithm [3], it was proposed to convert the model instances to point sets. The distance of two instances is the Hausdorff distance [43] of the point sets representing them. This solution is however challenging, since the conversion of geometric models to point sets in a robust manner is unclear in most of the cases. Even for homographies, there is a number of cases when this approach simply does not work, see Fig. 4 for examples. Instead, we follow the strategy proposed in the T-Linkage algorithm [32] to measure the model-to-model residual as the Tanimoto distance of the preference vectors [57] as follows:

$$f_{\text{T}}(v_a, v_b) = \frac{\langle v_a, v_b \rangle}{\|v_a\|^2 + \|v_b\|^2 - \langle v_a, v_b \rangle}. \quad (3)$$

The preference vector of a model instance h is $v \in [0, 1]^n$, where n is the number of input data points. Its i th coordinate is calculated as $v_i = 1 - f(h, \mathbf{p}_i)$, where f is the same loss function as what is used in the previous section and \mathbf{p}_i is the i th point from \mathcal{I} . Briefly, v_i is zero if the point-to-model residual is greater than the inlier-outlier threshold. Otherwise, it is from interval $(0, 1]$. In this case, the Tanimoto distance measures the overlap of the inlier sets of two model instances where the inlier assignment is done in a smooth manner. We will call the domain of preference vectors *consensus space* in the rest of the paper.

Note that, while the Tanimoto distance is not a proper metric over general vector spaces, it becomes one when the preference vector is $\in [0, +\infty)^n$ [30]. This property holds in our case. Also note that for $\mathbf{0}$ vectors, the distance is undefined. In our case, this never happens since each model is fit to a minimal sample of m ($=$ degrees of freedom) data points which consequently have 0 residuals. Thus, at least m elements of each preference vector are 1.

¹Such q_{\min} parameter is often used, e.g., in Structure-from-Motion algorithms (COLMAP uses $q_{\min} = 15$ [52]).

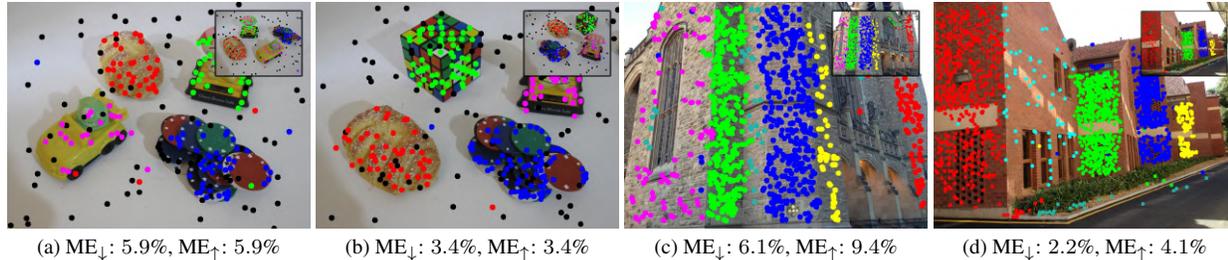


Figure 3. Image pairs used for multiple two-view motion and homography estimation, and point-to-model assignments (by color) determined by assigning each point to one of the instances returned by the proposed algorithm with the minimum point-to-model residual. Black points are outliers. For each image, the highest (ME_{\uparrow}) and lowest (ME_{\downarrow}) misclassification errors in five runs are reported. The least accurate results are shown. In (a–b), the worst and best results are identical. In (c–d), the difference is negligible. The proposed method finds all sought instances, the error originates from points assigned to the wrong instance. The selected scenes are the ones with the most ground truth instances to be found in the AdelaideRMF [66] dataset.



Figure 4. Examples where converting homographies to points and back is not robust. **Left:** top two corners are mapped to the same location. Thus, three matches remain for the homography recalculation. **Right:** the plane flips, thus the ordering of the points changes and the recovered homography will be incorrect.

Clustering. We formulate the problem of selecting different model instances as finding similar ones in \mathcal{I} which are then replaced by a single instance. A straightforward strategy for finding similar instances is to find clusters in the consensus space defined over the preference vectors.

In general, this clustering takes place in a large dimensional space, with as many dimensions as the number of input data points. In this particular setup however, we never have more than a few tens of instances to be clustered thanks to the iterative proposal strategy adapted from [4]. This means that the clustering is done on a few high-dimensional vectors that is very efficient with most of the clustering algorithms. Even if there are millions of points in the scene, a single model instance rarely has an extreme number of inliers and, thus, the indices of the non-zero elements in v can be stored, making the distance calculation efficient. In extreme cases, the min hash algorithm [9] can approximately find the inlier overlap in constant time.

After obtaining a set of instance clusters, the next step is to replace the instances in each cluster with a single one. Even though it would be straightforward to use the density modes, *e.g.* as in [13], it requires doing operations with the preference vectors, *e.g.*, averaging. However, such operations are undefined in the consensus space – the average of

two vectors is not necessarily the preference vector of the average instance. Thus, we replace each cluster with one of its elements that has the highest quality Q_f and, thus, is the most likely to represent the sought model parameters.

In the implementation, we use the DBSCAN [15, 53] density-based clustering that runs swiftly on our problem and returns accurate solutions. DBSCAN requires two parameters, *i.e.*, the minimum size c_{\min} of a cluster to be kept and a threshold ϵ_T to decide if two model instances are neighbors in the consensus space. The minimum size $c_{\min} = 1$ since single-element clusters also contain dominant model instances and, thus, should be kept. The setting of threshold ϵ_T is intuitive. Setting ϵ_T to 0 means that we consider models neighbors if and only if their preference vectors are exactly the same. Parameter $\epsilon_T = 1$ means that all methods are neighbors even if they do not share inliers.

2.4. Improving Instance Parameters

In order to improve the parameters of the instances kept by the clustering algorithm, we apply an iteratively re-weighted least-squares approach starting from the initial instance parameters. We use the robust MAGSAC++ weights.

The model optimization and clustering are applied repeatedly since during the optimization step two instances might become similar and, thus, should be put in the same cluster. This iteration stops when only one-element clusters are returned by the applied clustering algorithm.

2.5. Termination Criterion

To decide when the algorithm should terminate, we use the criterion proposed in [4] that is $n_i = (|\mathcal{P}| - |\mathcal{I}|) \sqrt[m]{1 - \sqrt[k]{1 - \mu}} \leq m + 1$, where μ is the required confidence in the results typically set to 0.95 or 0.99; k is the number of iterations; m is the size of the minimal sample; n_i and $|\mathcal{P}|$ are the number of inliers and points; and $|\mathcal{I}|$ is the cardinality of the united inlier sets of the kept model instances. This criterion is triggered if the probability of hav-

Algorithm 1 Connected Component Sampler: the next \mathcal{S} .

Input: $r, r_{\min}, r_{\max}, n_{\text{steps}}$ – current, min., max. neighborhood radius and partition number;

\mathcal{P} – data points; \mathcal{A} – neighborhood-graph; m – sample size

Output: \mathcal{S} – sample

if \neg Initialized(\mathcal{A}) **then** ▷ Run only once

$\mathcal{A} \leftarrow$ BuildNeighborhood(\mathcal{P}, r_{\max}) ▷ Radius is r_{\max}

$r \leftarrow r_{\min}$ ▷ The max. radius in \mathcal{A} for the next step

$\mathcal{C} \leftarrow$ GetConnectedComponents(\mathcal{A}, r)

while Empty(\mathcal{C}) $\wedge r \leq r_{\max}$ **do**

$r \leftarrow r + (r_{\max} - r_{\min}) / n_{\text{steps}}$

$\mathcal{C} \leftarrow$ GetConnectedComponents(\mathcal{A}, r)

$\mathcal{S} \leftarrow \emptyset$

if \neg Empty(\mathcal{C}) **then**

repeat ▷ Get the next largest *dominant* instance

$\mathcal{S} \leftarrow$ GetLargest(\mathcal{C}), $\mathcal{C} \leftarrow \mathcal{C} \setminus$ GetLargest(\mathcal{C})

until $|\mathcal{S}| \geq m \vee$ Empty(\mathcal{C})

if $|\mathcal{S}| < m$ **then**

$\mathcal{S} \leftarrow$ PROSAC(\mathcal{P}, m)

ing an unseen model with at least $m + 1$ inliers is smaller than $1 - \mu$. Since we have a criterion for an instance being dominant, the upper bound $m + 1$ for n_i can be replaced by q_{\min} to terminate when the probability of finding a dominant instance falls below $1 - \mu$.

3. Connected Component Sampling

There have been a number of algorithms proposed, *e.g.* PROSAC [11], P-NAPSAC [6], to find samples that consist of data points stemming from the same model instance early. When fitting multiple instances to real-world data, it usually is a reasonable assumption that the points form spatially coherent structures [2, 6, 24, 38]. We propose a deterministic sampling that returns the connected components in a progressively densified neighborhood-graph as samples. The algorithm is shown in Alg. 1.

The user-defined parameters are the minimum (r_{\min}) and maximum (r_{\max}) neighborhood radii and the number of steps when densifying the graph (n_{steps}). As initialization, the method first builds neighborhood-graph \mathcal{A} using the maximum radius. Then the connected components are selected from a sub-graph of \mathcal{A} where all edges are ignored that are larger than the current radius r . This is done to avoid building \mathcal{A} multiple times. The algorithm returns the largest connected component that has at least m points. If there is no such component, it increases the neighborhood size by changing r . Note that the returned sample \mathcal{S} is not necessarily a minimal sample. If r exceeds r_{\max} , there are no reasonable structures and, thus, it starts sampling from all data points in a global manner by the PROSAC sampler. Also

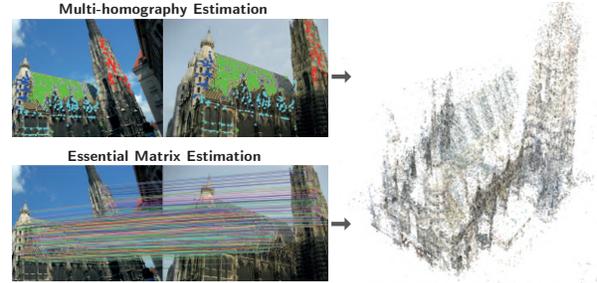


Figure 5. Multiple \mathbf{H} s contribute to the accurate reconstruction of Vienna Cathedral by [56]. The rot. and pos. errors decrease, respectively, by 9.8° and 5.0 m compared to using \mathbf{E} matrices only.

note that while PROSAC is a safe-guard for cases where the data is not spatially coherent, it was *never* executed in experiments of Sec. 4.

4. Experimental Results

Implementation Details. The proposed method is implemented in C++ using the Eigen library and the solver implementation from the GC-RANSAC [2] repository. We combine the algorithm with a number of components from USAC [42]. The included components are the following.

Sample degeneracy. The degeneracy tests of minimal samples are for rejecting clearly bad samples to avoid the sometimes expensive model estimation. For homographies, samples consisting of collinear points are rejected.

Sample cheirality. The test is for rejecting samples based on the assumption that both cameras observing a 3D surface must be on its same side. For homography fitting, we check if the ordering of the four point correspondences (along their convex hulls) in both images are the same.

Model degeneracy. The purpose of this test is to reject models early to avoid verifying them unnecessarily. For \mathbf{F} matrices, DEGENSAC [12] is applied to determine whether the epipolar geometry is affected by a dominant plane.

Parameters. Model-to-model threshold $\epsilon_T = 0.8$. This can roughly be interpreted as considering two model instances similar if more than 20% of their inliers are shared. The minimum quality $q_{\min} = 20$ and confidence $\mu = 0.99$. The parameters of the sampler are radii $r_{\min} = 20$, $r_{\max} = 200$, $r_{\text{steps}} = 5$. For point correspondences, the neighborhood is built on the joint 4D coordinate space. These parameters are used in *all* tested problems and on all datasets. Additional explanation of the hyper-parameters is in the supp. material.

4.1. Standard Benchmarks

To evaluate the proposed method on real-world problems, we use a number of publicly available datasets for homography, two-view motion, and motion fitting. The error is the misclassification error (ME), *i.e.*, the ratio of points assigned to the wrong cluster. The proposed method

	$ \mathcal{I} $ needed	Adelaide: Two-view motions			Adelaide: Homographies			Hopkins: Motions		
		19 scenes			19 scenes			155 scenes		
		avg.	std.	time	avg.	std.	time	avg.	std.	time
Proposed	no	5.3	4.4	0.05	3.1	3.5	0.11	4.4	6.3	0.04
Proposed (CC)	no	5.0	4.4	0.02	5.7	6.5	0.11	–	–	–
Prog-X [4]	no	10.7	8.7	14.38	6.6	5.9	1.03	8.4	10.3	0.02
Multi-X [3]	no	17.1	12.2	1.52	8.7	8.1	0.27	13.0	19.6	0.95
PEARL [24]	no	29.5	14.8	4.94	15.1	6.8	2.61	14.3	23.2	3.30
RPA [33]	yes	17.1	11.1	10.24	23.5	13.4	622.87	9.2	11.3	4.92
RansaCov [34]	yes	55.6	12.4	2.33	66.9	18.4	17.69	11.1	8.0	2.04
T-linkage [32]	no	46.7	15.6	2.69	54.8	22.2	57.84	27.2	15.6	0.95
MLink [35]	no	8.6	4.7	16.75	5.5	1.8	47.75	8.3	11.9	–
CONSAC [26]	no	–	–	–	5.2	6.5	8.1 / 21.0	–	–	–

Table 1. Avg. misclassification errors (in %; 5 runs on each scene), their std. and the run-times (secs) on two-view motion and homography fitting on the AdelaideRMF dataset [66], and motion fitting on the Hopkins dataset [59]. All methods use fixed parameters. For CONSAC, we report the times of running it on GPU and CPU. The second column ($|\mathcal{I}|$ needed) is “yes” for methods requiring the number of instances to fit. In the first row, the proposed method runs the P-NAPSAC [6] sampler. In the second one, the proposed CC sampler is used.

is designed to avoid assigning each data point to a single instance. Thus, we assigned each point to the model with the smallest residual. The results of the compared methods are copied from [4, 26, 35], where they were carefully tuned to achieve their best results with fixed parameters.

To test the proposed connected component-based sampler, we applied the proposed method with P-NAPSAC [6] and the proposed Connected Component Sampler, both of them exploiting the spatial nature of geometric data. We chose P-NAPSAC as a competitor, since it has a similar procedure, finding local structures by randomly sampling from gradually growing neighborhoods. The major difference between them is that P-NAPSAC is randomized and returns minimal samples, while the proposed Connected Component Sampler is deterministic and proposes larger-than-minimal samples as well.

Examples of multi-homography and two-view motion fitting are in Fig. 3. We chose the scenes from the AdelaideRMF [66] dataset with the most ground truth models to be found. Color denotes the point-to-model assignment done by assigning each point to the instance, outputted by the proposed method, with the smallest residual.

Two-view motion fitting is tested on the AdelaideRMF motion dataset consisting of 19 image pairs and correspondences manually assigned to two-view motion clusters. In this case, multiple \mathbf{F} matrices are to be found. For the proposal step, we used the 7PT algorithm [20]. In the IRLS fitting, we applied the norm. 8PT solver [19].

The avg. errors over five runs and their std. are shown in the left block of Table 1. The proposed method leads to state-of-the-art accuracy with both tested samplers. The proposed Connected Components Sampler (CC) improves both the accuracy and processing time. The proposed method with CC is twice as accurate as the second best competitor (MLink) while being two orders-of-magnitude

faster than the second fastest method (Multi-X). The proposed method runs in *real-time* on these scenes. On avg., out of the 45 motions in the dataset, the proposed method does not find 2 instances while returning 1 false positive.

Homography fitting is tested on the AdelaideRMF H dataset [66]. It consists of 19 image pairs with ground truth correspondences assigned manually to \mathbf{H} s. In these tests, we also included the errors of CONSAC [26]. Since the run-times are not reported in [26], we re-ran the algorithm both on GPU and CPU and calculated the avg. times.

We used the norm. 4PT algorithm both in the proposal and IRLS steps. The results are shown in the middle block of Table 1. The proposed method is almost twice as accurate as the second best one (CONSAC) while being significantly faster than all algorithms. It leads to the most accurate solutions while being the fastest. In this case, P-NAPSAC sampler leads to the best results. Out of the 52 \mathbf{H} s, the proposed method does not find 2 with 2 false positives.

Motion segmentation is tested on 155 videos of the Hopkins dataset [59]. It consists of 155 sequences divided into three categories: checkerboard, other, and traffic. The trajectories are inherently corrupted by noise, but no outliers are present. Motion segmentation in videos is the retrieval of sets of points undergoing rigid motions in a dynamic scene captured by a moving camera. It can be considered a subspace segmentation under the assumption of affine cameras. For such cameras, all feature trajectories associated with a single moving object lie in a 4D linear subspace in \mathbb{R}^{2F} , where F is the frame number [59].

The results are shown in the right part of Table 1. The proposed method leads to the lowest errors. It still runs in *real-time*. In this case, we used uniform sampling since building a neighborhood-graph (required both by the CC sampler and P-NAPSAC) on point trajectories is not trivial.

Relative Pose Estimation				
	avg. ϵ_R	med. ϵ_R	avg. ϵ_t	med. ϵ_t
E matrix	9.51	3.46	18.15	9.08
E from Hs	9.56	3.47	18.21	9.09
Pose averaging	8.71	3.69	34.34	25.27
Pose selection	8.33	3.34	17.84	8.92
Pose selection (CC)	8.24	3.31	17.81	8.89

Table 2. Relative rotation ϵ_R and translation ϵ_t errors ($^\circ$) on 435k image pairs from the 1DSfM dataset obtained by **E** matrix estimation; calculating **E** from the inliers of homographies (**E** from **Hs**); pose averaging on the poses decomposed from **E** and multiple **Hs**; and selecting the pose with the most inliers from the decomposed ones (Pose selection) with the proposed sampler (CC).

Global SfM Results				
	avg. ϵ_R	med. ϵ_R	avg. ϵ_p	med. ϵ_p
E matrix	11.15	6.58	10.25	8.93
E + mult. Hs	7.93	6.21	10.52	4.60
E + mult. Hs (CC)	5.56	5.61	9.57	3.99

Table 3. Rotation and position errors of the global SfM implemented in [56] when initialized with a poses estimated from **E** matrices, and via the proposed pose selection from **E** and multiple **Hs**.

4.2. Application: Relative Pose Estimation

In this section, we focus on improving relative pose estimation by exploiting multiple homographies. Pose estimation is a fundamental problem in a number of popular methods, *e.g.*, in Structure-from-Motion algorithms. While the usual procedure to estimate a relative pose uses epipolar geometry, it is well-known that the pose can also be obtained from a homography if the cameras are calibrated. However, in most pipelines, homographies are used only if the scene is degenerate for fundamental matrix estimation, *e.g.*, a single plane dominates the scene [12] or the camera undergoes purely rotational motion. In this section, we aim to propose a way of exploiting multiple homographies to improve the relative pose accuracy. See Fig. 5 for an example.

We downloaded the 1DSfM dataset [64] and applied COLMAP [52] to obtain a reconstruction that can be used as ground truth. Note that the 1DSfM dataset provides a ground truth, however, it was created by the Bundler algorithm [54] that is more than 10 years old. We use the following approach in order to find potentially matching image pairs. First, we extract GeM [41] descriptors with ResNet-50 [21] CNN, pre-trained on GLD-v1 dataset [39]. Then we calculate the inner-product similarity between the descriptors, resulting in an $n \times n$ similarity matrix. In the experiments, we use only the image pairs with similarity higher than 0.4 [5]. Finally, we estimated multiple homographies for all considered image pairs, 434 587 in total.

We tested the following approaches to recover the relative pose from multiple **Hs**:

1. Estimating the essential matrix [55] from the inliers of

	avg. ϵ_R	med. ϵ_R	avg. ϵ_p	med. ϵ_p
E4+2	1.19	0.50	0.033	0.025
H3+2	0.45	0.24	0.103	0.041
E + mult. Hs	0.32	0.26	0.026	0.022

Table 4. The avg. and med. rotation ϵ_R ($^\circ$) and position ϵ_p (m) errors on 23 190 image pairs from the KITTI dataset obtained by generalized **E** matrix (**E4+2**) [70] and **H** estimation [7] (**H3+2**); and by selecting the pose obtained from a generalized **E** matrix and a set of generalized homographies (**E** + mult. **Hs**).

the returned homographies.

2. Decomposing all found homographies and, also, the essential matrix relative poses and running pose averaging by [10, 65].
3. Decomposing each homography [36] and, also, the essential matrix to pose and selecting the one which has the most inliers determined by thresholding the re-projection error. The translation is then re-estimated by solving equation $\mathbf{p}_2^T [\mathbf{t}]_{\times} \mathbf{R} \mathbf{p}_1 = 0$ with known rotation **R**, where $[\mathbf{t}]_{\times}$ is the cross-product matrix of translation **t** and $[\mathbf{t}]_{\times} \mathbf{R}$ is the essential matrix. Details are in the supplementary material.

The results are reported in Table 2. Results when using the proposed connected component sampler (CC) are also shown. To measure the error in the rotation, we calculate the angular difference between the ground truth $\hat{\mathbf{R}}$ and estimated **R** ones as $\epsilon_R = \cos^{-1}((\text{tr}(\mathbf{R}\hat{\mathbf{R}}^T) - 1)/2)$. Since the translation is up to scale, the error is the angular difference ϵ_t of the ground truth and estimated translations. The avg. rotation and translation errors are improved by, respectively, 1.27 and 0.34 degrees compared to **E** estimation. CC sampler leads to the best results. Since it is extremely fast, the computational overhead is merely a few *ms*.

We applied the global SfM implemented in the Theia library [56] initialized with the poses estimated in the proposed way and, also, with the poses estimated using only essential matrices. The accuracy of the reconstruction is reported in Table 3. We report the average rotation (avg. ϵ_R , in degrees) and position errors (avg. ϵ_p , in meters) and, also, the median errors averaged over the scenes. The proposed algorithm with the CC sampler *significantly* reduces both the rotation and position errors of the reconstruction.

4.3. Application: Fast-moving Object Detection

In this section, we estimate the trajectories of objects that are significantly blurred by their motion. As defined in [47], an image *I* of such blurred object is formed as a composition of the blurred object appearance and the background

$$I = H * F + (1 - H * M) B, \quad (4)$$

where the sharp object appearance *F* with mask *M* encodes the object, blur kernel *H* encodes the trajectory, and *B* rep-

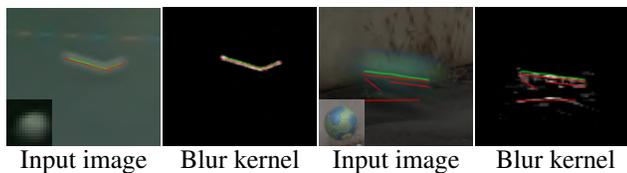


Figure 6. Multiple line segment fitting for trajectory estimation of fast-moving objects. Estimated line segments are in red, the ground truth is in green. Sharp object appearance is overlaid in the bottom left corner of the input image.

represents the background. Input image I and background B are assumed to be known. The unknowns in (4) are estimated either by alternating energy minimization with additional priors [27–29, 45, 46, 48] or more recently by learning from synthetic data [49, 50] and neural rendering [51].

The formation model in (4) encodes the trajectory by the blur kernel. However, there are no guarantees that the blur kernel corresponds to a physically plausible trajectory, which is assumed to be piece-wise linear due to bounces. Blur kernels also contain other responses due to other moving objects in the scene. In the extreme case, if two fast-moving objects intersect or fly close to each other, the blur kernel will contain multiple responses corresponding to each motion. In practice, the estimated blur kernels are noisy, with many outliers, and contain artifacts due to shadows, low contrast, and discretization. Motion blur priors [61] have been proposed to reduce these issues, but extracting the final continuous trajectory is still a challenging multi-instance model fitting task (see Fig. 6 for examples).

Recent methods [28, 48] address this task by employing Sequential RANSAC [25, 60] on the thresholded blur kernels. We extract blur kernels using the TbD method [28] from all sequences in the TbD [28] and TbD-3D [48] datasets. The TbD dataset is simpler since it contains mostly uniformly colored objects moving in the plane parallel to the camera plane. The TbD-3D dataset is more challenging with highly textured objects that are rotating and moving in 3D. The ground truth sub-frame object location is given from a high-speed camera. We estimate multiple line segments in each blur kernel and measure the average L_2 distance of each ground truth location to the closest fitted line segment. Table 5 shows the average error, its standard deviation, and average run-time for a wide range of state-of-the-art methods. We used the implementations provided by the authors. The proposed method outperforms all compared algorithms both in terms of accuracy and processing time, running in *real-time*. Additional results, *e.g.* demonstrating the effect of the proposed soft assignment, are in the supplementary material. Without considering soft assignment, continuous chains can not be found. This leads to losing short segments and affects the accuracy notably.

Dataset:	Easy (322) [28]			Challenging (470) [48]		
	avg.	std.	time	avg.	std.	time
Proposed	1.39	6.73	0.02	2.84	2.80	0.05
Prog-X [4]	1.87	6.80	0.24	3.74	3.22	0.09
PEARL [24]	1.39	6.74	0.05	4.83	6.17	0.08
J-Linkage [58]	1.73	6.72	4.02	4.85	6.51	4.52
T-Linkage [32]	1.71	6.71	7.07	4.46	5.21	33.65
RPA [33]	2.74	7.77	7.66	5.19	4.47	21.79
RansaCov [34]	1.48	6.74	2.09	3.90	4.83	7.62
Seq. RANSAC	1.66	6.72	0.68	6.08	7.50	0.98

Table 5. The avg. and std. accuracy (px) and run-time (secs) of multiple line segment detection for finding the trajectories of fast-moving objects. The number of images are in brackets.

4.4. Pose from Generalized Camera

To further test the pose selection technique from multiple homographies and an essential matrix, as proposed in Section 4.2, we downloaded the KITTI odometry dataset [17], where each frame consists of the images of two cameras. We considered the two cameras as a generalized one and estimated the pose between this camera and the left image of the next frame. We used the generalized essential matrix [70] (**E4+2**) and homography [7] (**H3+2**) solvers. For finding a single **E4+2** or **H3+2**, we used GC-RANSAC [2]. The methods were tested on a total of 23 190 frame pairs.

The results are in Table 4. The proposed technique (**E** + mult. **Hs**), selecting the best pose from the set decomposed from an essential matrix and multiple homographies, leads to the most accurate results in terms of average rotation and position errors. Its median rotation error is similar to **H3+2**. Its median position error is the lowest.

5. Conclusion

We propose a new multi-instance model fitting algorithm that is a simple iteration of instance proposal, clustering in the consensus space, and parameter re-estimation. Due to not forming crisp point-to-model assignments, the method runs in *real-time* on a number of vision problems. On two-view motion estimation, it is at least two orders-of-magnitude faster than the competitors. It leads to results superior to the state-of-the-art both in terms of accuracy and run-time on the standard benchmark datasets. Moreover, the proposed Connected Component sampler outperforms the recent P-NAPSAC on a number of real-world problems. In addition, we demonstrated on a total of 458 569 images or image pairs that using multiple model instances, *e.g.* homographies or line segments, is beneficial for various popular vision applications, *e.g.*, Structure-from-Motion.

Acknowledgements. This project was supported by the ETH Postdoc fellowship, by the OP VVV funded project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics” (RCI), and by the Robert Bosch GmbH. The conference presentation was supported by Czech Technical University in Prague (RCI project).

References

- [1] Paul Amayo, Pedro Piniés, Lina M Paz, and Paul Newman. Geometric multi-model fitting with a convex relaxation algorithm. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 8138–8146, 2018. 1, 2
- [2] Daniel Barath and Jiri Matas. Graph-cut RANSAC. *Proc. Conf. on Computer Vision and Pattern Recognition*, 2018. 5, 8
- [3] Daniel Barath and Jiri Matas. Multi-class model fitting by energy minimization and mode-seeking. In *Proc. European Conf. on Computer Vision*, 2018. 1, 2, 3, 6
- [4] Daniel Barath and Jiri Matas. Progressive-X: Efficient, anytime, multi-model fitting algorithm. In *Proc. Int. Conf. on Computer Vision*, pages 3780–3788, 2019. 1, 2, 3, 4, 6, 8
- [5] Daniel Barath, Dmytro Mishkin, Ivan Eichhardt, Ilya Shipachev, and Jiri Matas. Efficient initial pose-graph generation for Global SfM. In *Proc. Conf. on Computer Vision and Pattern Recognition*, 2021. 7
- [6] Daniel Barath, Jana Nuskova, Maksym Ivashechkin, and Jiri Matas. MAGSAC++, a fast, reliable and accurate robust estimator. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 1304–1312, 2020. 2, 3, 5, 6
- [7] Snehal Bhayani, Torsten Sattler, Daniel Barath, Patrik Beliansky, Janne Heikkila, and Zuzana Kukelova. Calibrated and partially calibrated semi-generalized homographies, 2021. 7, 8
- [8] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2004. 2
- [9] A. Z. Broder. On the resemblance and containment of documents. In *Compression and complexity of sequences 1997. proceedings*, pages 21–29. IEEE, 1997. 4
- [10] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *Proc. Int. Conf. on Computer Vision*, pages 521–528, 2013. 7
- [11] Ondrej Chum and Jiri Matas. Matching with PROSAC-progressive sample consensus. In *Proc. Conf. on Computer Vision and Pattern Recognition*. IEEE, 2005. 3, 5
- [12] Ondrej Chum, Tomas Werner, and Jiri Matas. Two-view geometry estimation unaffected by a dominant plane. In *Proc. Conf. on Computer Vision and Pattern Recognition*. IEEE, 2005. 5, 7
- [13] Dorin Comaniciu and Peter Meer. Mean shift analysis and applications. In *Proc. Int. Conf. on Computer Vision*, volume 2, pages 1197–1203. IEEE, 1999. 4
- [14] Andrew Delong, Lena Gorelick, Olga Veksler, and Yuri Boykov. Minimizing energies with hierarchical costs. *Int. Journal of Computer Vision*, 2012. 2
- [15] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996. 4
- [16] 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*, 1981. 1
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *Proc. Conf. on Computer Vision and Pattern Recognition*, 2012. 8
- [18] Nicolas Guil and Emilio L. Zapata. Lower order circle and ellipse Hough transform. *Pattern Recognition*, 1997. 1
- [19] Richard Hartley. In defense of the eight-point algorithm. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 1997. 6
- [20] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2003. 6
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. Conf. on Computer Vision and Pattern Recognition*, 2016. 7
- [22] P. V. C. Hough. Method and means for recognizing complex patterns, 1962. 1
- [23] John Illingworth and Josef Kittler. A survey of the Hough transform. *Computer Vision, Graphics, and Image Processing*, 1988. 1
- [24] Hossam Isack and Yuri Boykov. Energy-based geometric multi-model fitting. *Int. Journal of Computer Vision*, 2012. 1, 2, 5, 6, 8
- [25] Yasushi Kanazawa and Hiroshi Kawakami. Detection of planar regions with uncalibrated stereo using distributions of feature points. In *Proc. British Machine Vision Conf.*, 2004. 1, 8
- [26] Florian Kluger, Eric Brachmann, Hanno Ackermann, Carsten Rother, Michael Ying Yang, and Bodo Rosenhahn. CONSAC: Robust multi-model fitting by conditional sample consensus. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 4634–4643, 2020. 2, 6
- [27] Jan Kotera, Jiri Matas, and Filip Šroubek. Restoration of fast moving objects. *IEEE Trans. Image Processing*, 29:8577–8589, 2020. 8
- [28] Jan Kotera, Denys Rozumnyi, Filip Šroubek, and Jiri Matas. Intra-frame object tracking by deblatting. In *International Conference on Computer Vision Workshops*, Oct 2019. 8
- [29] Jan Kotera and Filip Šroubek. Motion estimation and deblurring of fast moving objects. In *Proc. Int. Conf. on Image Processing*, pages 2860–2864, Oct 2018. 8
- [30] Alan H Lipkus. A proof of the triangle inequality for the Tanimoto distance. *Journal of Mathematical Chemistry*, 26(1):263–265, 1999. 3
- [31] Jiayi Ma, Xingyu Jiang, Aoxiang Fan, Junjun Jiang, and Junchi Yan. Image matching from handcrafted to deep features: A survey. *Int. Journal of Computer Vision*, 129(1):23–79, 2021. 3
- [32] Luca Magri and Andrea Fusiello. T-Linkage: A continuous relaxation of J-Linkage for multi-model fitting. In *Proc. Conf. on Computer Vision and Pattern Recognition*, 2014. 1, 2, 3, 6, 8
- [33] Luca Magri and Andrea Fusiello. Robust multiple model fitting with preference analysis and low-rank approximation. In *Proc. British Machine Vision Conf.*, 2015. 1, 2, 6, 8
- [34] Luca Magri and Andrea Fusiello. Multiple model fitting as a set coverage problem. In *Proc. Conf. on Computer Vision and Pattern Recognition*, 2016. 1, 2, 6, 8

- [35] Luca Magri, Filippo Leveni, and Giacomo Boracchi. Multi-link: Multi-class structure recovery via agglomerative clustering and model selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1853–1862, 2021. 6
- [36] Ezio Malis and Manuel Vargas. *Deeper understanding of the homography decomposition for vision-based control*. PhD thesis, INRIA, 2007. 7
- [37] Jiri Matas, Csaba Galambos, and Josef Kittler. Robust detection of lines using the progressive probabilistic Hough transform. *Computer Vision and Image Understanding*, 2000. 1
- [38] D. R. Myatt, Philip Torr, Slawomir J. Nasuto, Mark J. Bishop, and R. Craddock. NAPSAC: High noise, high dimensional robust estimation - it's in the bag. In *Proc. British Machine Vision Conf.*, 2002. 2, 5
- [39] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-scale image retrieval with attentive deep local features. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 3456–3465, 2017. 7
- [40] Trung Thanh Pham, Tat-Jun Chin, Konrad Schindler, and David Suter. Interacting geometric priors for robust multi-model fitting. *IEEE Trans. Image Processing*, 2014. 1, 2
- [41] Filip Radenović, Giorgos Toliás, and Ondrej Chum. Fine-tuning CNN image retrieval with no human annotation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2018. 7
- [42] R. Raguram, O. Chum, M. Pollefeys, J. Matas, and J-M. Frahm. USAC: a universal framework for random sample consensus. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2013. 5
- [43] R Tyrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009. 3
- [44] Paul L. Rosin. Ellipse fitting by accumulating five-point fits. *Pattern Recognition Letters*, 1993. 1
- [45] Denys Rozumnyi, Jan Kotera, Filip Šroubek, and Jiri Matas. Non-causal tracking by deblatting. In Gernot A. Fink, Simone Frintrop, and Xiaoyi Jiang, editors, *German Conference on Pattern Recognition*, pages 122–135, Cham, 2019. Springer International Publishing. 8
- [46] D. Rozumnyi, J. Kotera, F. Šroubek, and J. Matas. Tracking by deblatting. *Int. Journal of Computer Vision*, 129(9):2583–2604, 2021. 8
- [47] Denys Rozumnyi, Jan Kotera, Filip Šroubek, Lukas Novotný, and Jiri Matas. The world of fast moving objects. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 4838–4846, July 2017. 7
- [48] Denys Rozumnyi, Jan Kotera, Filip Šroubek, and Jiri Matas. Sub-frame appearance and 6D pose estimation of fast moving objects. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 6777–6785, 2020. 8
- [49] Denys Rozumnyi, Jiří Matas, Filip Šroubek, Marc Pollefeys, and Martin R. Oswald. Fmodetect: Robust detection of fast moving objects. In *Proc. Int. Conf. on Computer Vision*, pages 3541–3549, October 2021. 8
- [50] Denys Rozumnyi, Martin R. Oswald, Vittorio Ferrari, Jiri Matas, and Marc Pollefeys. Defmo: Deblurring and shape recovery of fast moving objects. In *Proc. Conf. on Computer Vision and Pattern Recognition*, Nashville, Tennessee, USA, Jun 2021. 8
- [51] Denys Rozumnyi, Martin R. Oswald, Vittorio Ferrari, and Marc Pollefeys. Shape from blur: Recovering textured 3d shape and motion of fast moving objects. In *Proc. Conf. on Neural Information Processing Systems*, 2021. 8
- [52] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. 3, 7
- [53] Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems*, 42(3):1–21, 2017. 4
- [54] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3D. In *ACM siggraph*, pages 835–846. 2006. 7
- [55] H. Stewenius, C. Engels, and D. Nistér. Recent developments on direct relative orientation. *Journal of Photogrammetry and Remote Sensing*, 60(4):284–294, 2006. 7
- [56] Christopher Sweeney, Tobias Hollerer, and Matthew Turk. Theia: A fast and scalable structure-from-motion library. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 693–696, 2015. 5, 7
- [57] T. T. Tanimoto. Elementary mathematical theory of classification and prediction. 1958. 3
- [58] Roberto Toldo and Andrea Fusiello. Robust multiple structures estimation with J-Linkage. In *Proc. European Conf. on Computer Vision*, 2008. 8
- [59] Roberto Tron and Rene Vidal. A benchmark for the comparison of 3-d motion segmentation algorithms. In *Proc. Conf. on Computer Vision and Pattern Recognition*, 2007. 6
- [60] E. Vincent and Robert Laganière. Detecting planar homographies in an image pair. In *International Symposium on Image and Signal Processing and Analysis*, 2001. 1, 8
- [61] Filip Šroubek and Jan Kotera. Motion blur prior. In *Proc. Int. Conf. on Image Processing*, pages 928–932, 2020. 8
- [62] Hanzi Wang, Guobao Xiao, Yan Yan, and David Suter. Mode-seeking on hypergraphs for robust geometric model fitting. In *Proc. Int. Conf. on Computer Vision*, 2015. 1, 2
- [63] Hanzi Wang, Guobao Xiao, Yan Yan, and David Suter. Searching for representative modes on hypergraphs for robust geometric model fitting. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2018. 2
- [64] Kyle Wilson and Noah Snavely. Robust global translations with 1DSfM. In *Proc. European Conf. on Computer Vision*, 2014. 7
- [65] Kyle Wilson and Noah Snavely. Robust global translations with 1DSfM. In *Proc. European Conf. on Computer Vision*, pages 61–75. Springer, 2014. 7
- [66] Hoi Sim Wong, Tat-Jun Chin, Jin Yu, and David Suter. Dynamic and hierarchical multi-structure geometric model fitting. In *Proc. Int. Conf. on Computer Vision*, 2011. 1, 4, 6
- [67] Lei Xu, Erkki Oja, and Pekka Kultanen. A new curve detection method: randomized Hough transform (rht). *Pattern Recognition Letters*, 1990. 1

- [68] Wei Zhang and Jana Kosecká. Nonparametric estimation of multiple structures with outliers. In *Dynamical Vision*. Springer, 2007. [2](#)
- [69] Qing Zhao, Yun Zhang, Qianqing Qin, and Bin Luo. Quantized residual preference based linkage clustering for model selection and inlier segmentation in geometric multi-model fitting. *Sensors*, 20(13):3806, 2020. [2](#)
- [70] Enliang Zheng and Changchang Wu. Structure from motion using structure-less resection. In *Proc. Int. Conf. on Computer Vision*, pages 2075–2083, 2015. [7](#), [8](#)
- [71] Marco Zuliani, C. S. Kenney, and Bangalore Manjunath. The multiRANSAC algorithm and its application to detect planar homographies. In *Proc. Int. Conf. on Image Processing*. IEEE, 2005. [1](#)