# Using Spatial Order to Boost the Elimination of Incorrect Feature Matches 

Lior Talker<br>The University of Haifa<br>ltalke01@campus.haifa.ac.il

Yael Moses<br>The Interdisciplinary Center<br>Herzliya 46150, Israel<br>yael@idc.ac.il

Ilan Shimshoni<br>The University of Haifa<br>ishimshoni@mis.haifa.ac.il


#### Abstract

Correctly matching feature points in a pair of images is an important preprocessing step for many computer vision applications. In this paper we propose an efficient method for estimating the number of correct matches without explicitly computing them. In addition, our method estimates the region of overlap between the images. To this end, we propose to analyze the set of matches using the spatial order of the features, as projected to the $x$-axis of the image. The set of features in each image is thus represented by a sequence. This reduces the analysis of the matching problem to the analysis of the permutation between the sequences. Using the Kendall distance metric between permutations and natural assumptions on the distribution of the correct and incorrect matches, we show how to estimate the abovementioned values. We demonstrate the usefulness of our method in two applications: (i) a new halting condition for RANSAC based epipolar geometry estimation methods that considerably reduce the running time, and (ii) discarding spatially unrelated image pairs in the Structure-fromMotion pipeline. Furthermore, our analysis can be used to compute the probability that a given match is correct based on the estimated number of correct matches and the rank of the features within the sequences. Our experiments on a large number of synthetic and real data demonstrate the effectiveness of our method. For example, the running time of the image matching stage in the Structure-from-Motion pipeline may be reduced by about $99 \%$ while preserving about $80 \%$ of the correctly matched feature points.


## 1. Introduction

Matching feature points between a pair of images is a fundamental problem in computer vision. The estimation of epipolar geometry between images [13, 22], 3D structure reconstruction (SfM) [29, 30], and scene recognition [17] are typical examples of useful tasks that are based on feature matching. While many methods for feature matching exist, the critical stage of filtering incorrect matches is
costly when using algorithms such as Random Sample Consensus (RANSAC) [12, 13, 22, 3].

In this paper we propose to analyze the set of correct matches, without explicitly computing it, using the spatial order of the features in each image. We estimate the number of correct matches and the probability that a given match is correct. In addition, we estimate the overlap region of the pair of images and whether they overlap at all. Our method can be applied to sets of matching features irrespective of their descriptors (e.g., [21, 1, 4, 24]) or the matching method used to compute them. Our method can be used as a preprocessing step to improve the efficiency of existing methods such as RANSAC and SfM, as described below.

The basic idea is as follows. We represent the image features as sequences defined by their spatial order along the $x$-axis of the image. The matching between features in a pair of images induces a permutation that relates the spatial order in one image to that in the other image (see Fig. 1). The matching is analyzed using a measure of correlation between permutations, the Kendall distance metric [9, 10]. We use statistical assumptions on the distribution of correctly and incorrectly matched features; the spatial order of correctly matched features is usually preserved, while incorrectly matched features are expected to have random order. Note that the problem is not trivialized by these assumptions since simply computing the largest set of features that preserve their spatial order does not necessarily provide the correct set of matches (see discussion in Sec. 3). Hence, we also consider the interaction between correct and incorrect matches for obtaining our estimations.

Our estimation of the number of correct matches can be used to improve the running time of RANSAC methods [12, 13, 22, 3]. In adaptive RANSAC, a subset of matches in a pair of images is randomly sampled and used to compute the expected geometric transformation between the features (e.g., homography or epipolar geometry), which is then verified against all matches. The transformation with the largest set of inliers, features consistent with the computed geometric transformation, is chosen. Since the number of inliers is usually unknown, the number of required itera-
tions is high. This increases the running time of RANSAC, which is its major drawback. Our estimate of the number of correct matches can be used to improve the running time by halting when a consensus set of this size is obtained (see Sec. 6).

Computing the structure of the scene and the cameras' parameters from a set of images using SfM methods (e.g., [29, 30]) is a fundamental problem in computer vision with many applications. Such methods, which typically require tens or hundreds of images, strongly rely on correct matches. A major time consuming step in SfM methods is the robust matching of features between image pairs, which is typically obtained by RANSAC. Our method can significantly shorten the SfM pipeline [29, 30] by running RANSAC only on image pairs with a sufficiently large number of correct matches (see Sec. 6).

We show that the number of correct matches can also be used to compute the probability that a given match is correct. One application of such probability is in improving the performance of guided RANSAC, where matches that are more likely to be correct are sampled more often (e.g., [13, 3, 22, 5, 7, 6]). In existing methods the probability that a given match is correct is based only on feature descriptors, whereas in our method it is based on the features' order in the image. Thus, our method by itself or in combination with appearance based methods give a more accurate probability for the correctness of a match than when using the method based only on feature descriptors (see details in Sec. 5).

The main contribution of the paper is a novel analysis of the correct and incorrect matches with respect to their spatial order in the images. Given a set of putative matches, our method efficiently estimates the number of correct ones, the region of overlap between a pair of images, and for each match the probability that it is correct. These estimations can be used to improve the efficiency of RANSAC and SfM methods, as demonstrated empirically in Sec. 6.

## 2. Related Work

Feature point matching between a pair of images has been studied extensively in computer vision. Over the years many interest point detectors and patch descriptors were proposed to detect and match the images of the same 3D points in the scene (e.g., [21, 1, 4, 24]). Appearance based methods, such as [15], discard descriptors by learning their success rate in matching between their corresponding features, in order to increase the proportion of correct to incorrect matches. Other methods use local geometric structures between a number of matches in order to decrease the probability of mismatching, e.g., [23].

A major drawback of the RANSAC method for filtering incorrect matches is its running time; for example, running RANSAC on a pair of images with 1000 matches may


Figure 1. Correctly/incorrectly matched features are marked as green/yellow circles, respectively. The feature sequences are given by $[N]=\langle 1,2,3,4,5,6,7,8,9,10\rangle$ and $\sigma=$ $\langle 9,5,1,3,4,2,6,7,8\rangle$. For example, $\sigma(1)=9$ and $\sigma(5)=4$. Note that the green lines (correct matches) do not intersect, while most of the yellow lines (incorrect matches) do.
take a few seconds, which can be a major bottleneck in online applications involving large sets of images. To improve RANSAC's runtime, methods for guiding the sampling of matches according to their probability to be correct (instead of a uniform sampling) were proposed [13, 3, 22, 5, 7, 6]. For example, the similarity of feature appearance defined by Lowe's distance ratio between the closest and the second closest matches [21] is used as a match correctness likelihood by [13, 22]. This method ignores important geometric information. In our method, we use the spatial order of the matched features, which carries geometric information. A comparison between probabilities computed by our method and Lowe's distance ratio is provided in Sec. 6.

An inlier rate estimation was also used by [20] to compute a homography transformation between two images, which is guaranteed to find an approximation to the global optimum. The rate is computed by counting the number of homographies that agree with each inlier rate. The drawback of the method is the search process, which is time consuming and is applicable only to homographies.

## 3. Problem Formulation

Let $\mathcal{M}=\left\{\left(p_{i}, q_{j}\right)\right\}$ be a set of putative matches between two feature sequences, $p_{1}, \ldots, p_{N}$ and $q_{1}, \ldots, q_{N}$, in a pair of images, $I_{1}$ and $I_{2}$, respectively. The index $i$ of $p_{i}$ represents the position (the rank) of the feature in the sequence of $I_{1}$, when sorted according to the $x$ coordinate. Similarly,
the index $j$ of $q_{j}$ corresponds to the rank in the feature sequence of $I_{2}$. For the rest of the paper, let us represent $\mathcal{M}$ by two sequences of indexes, $[N]=\langle 1, \ldots, N\rangle$ and $\sigma$, where $\left(p_{i}, q_{\sigma(i)}\right) \in \mathcal{M}$ (see Fig. 1).

The set of matches, $\mathcal{M}$, can be partitioned into two disjoint sets, the correct ("Good") and the incorrect ("Bad") matches. Let $G$ and $B$ be the sets of indexes in $I_{1}$ for which the matching is correct and incorrect, respectively. In this case, the number of matches is given by $N=N_{G}+N_{B}$, where $N_{G}=|G|$ and $N_{B}=|B|$.

We analyze the spatial orders of matched features in the pair of images, which is represented by the permutation $\sigma$. We use statistical assumptions on the distribution of the correct and incorrect matches, which allows us to estimate the number of correct matches, $N_{G}$ without explicitly computing the set of correct matches. In addition, we propose a method to compute the probability that a given match is correct.

To do so, we use the Kendall [18] distance that corresponds to the sum of order inversions in $\sigma$. The two main statistical assumptions used are:
A1: The spatial order of the correctly matched features in $I_{1}$ is preserved in $I_{2}$. This assumption holds when (i) the cameras have no relative roll (rotation around the $z$-axis), and (ii) there are no "poles", i.e., features with significant difference in depth on different surfaces.
A2: The spatial order of incorrect matches is random. This can be justified since the spatial location of incorrectly matched features in different images is arbitrary.
Assumption A1 mostly holds since in natural scenes there are few inversions between correct matches. (This assumption is often used in stereo matching algorithms e.g., [31, 2].) Moreover, the roll value is often available in the camera so it can be compensated for.

We note here that, based on assumptions A1 and A2, it may seem that the set of correct matches can be obtained directly by finding the longest increasing (nonconsecutive) subsequence, $L(N)$. This approach yields poor results for two reasons. First, due to local order inversions between adjacent correct matches, caused by small changes in depth and camera roll, $|L(N)|$ is usually not a good estimation for $N_{G}$. Second, it is proven in [28] that for a random permutation of size $N$ the expectation of $|L(N)|$ is $2 \sqrt{N}$; thus, when $N_{G}$ is low, $L(N)$ may correspond to only incorrect matches or to a mixture of correct and incorrect matches. In contrast to this naïve approximation, our analysis also considers the distribution of the order of the incorrect matches and therefore better estimates $N_{G}$ even for low values.

## 4. The Number of Correct Matches

Two sequences are considered fully overlapped if the fields of view (FOV) of the corresponding images consist
of the same region of the scene. We first consider the case of fully overlapped sequences (Sec. 4.2) and then present a method to compute the regions of overlap between the pair of images (Sec. 4.3). The Kendall distance metric that is used to estimate $N_{G}$ is reviewed next.

### 4.1. The Kendall Distance

The Kendall distance is the number of pairwise order inversions in the two sequences $[N]$ and $\sigma$. Two pairs of matched features, $\left(p_{i}, q_{\sigma(i)}\right)$ and $\left(p_{j}, q_{\sigma(j)}\right)$, have order inversion if the orders $(i, j)$ and $(\sigma(i), \sigma(j))$ are inverted. Formally, let us define a binary function, $\eta_{\sigma}$ :

$$
\eta_{\sigma}(i, j)= \begin{cases}1 & (i<j) \wedge(\sigma(i)>\sigma(j))  \tag{1}\\ & \multicolumn{1}{c}{\text { or }} \\ & (i>j) \wedge(\sigma(i)<\sigma(j)) \\ 0 & \text { otherwise }\end{cases}
$$

The Kendall distance is thus given by:

$$
\begin{equation*}
K([N], \sigma)=\sum_{1 \leq i \leq N} \sum_{i<j \leq N} \eta_{\sigma}(i, j) \tag{2}
\end{equation*}
$$

To compute $N_{G}$ from the value $K=K([N], \sigma)$ we formulate $K$ as the sum of three terms:

$$
\begin{equation*}
K=K_{G}+K_{B}+K_{G B} \tag{3}
\end{equation*}
$$

where $K_{G}$ corresponds to the number of order inversions between correct matches, $K_{B}$ between incorrect matches, and $K_{G B}$ between pairs of correct and incorrect matches. That is,

$$
\begin{align*}
K_{G} & =\sum_{i \in G} \sum_{\substack{j \in G \\
i<j}} \eta_{\sigma}(i, j), \\
K_{B} & =\sum_{i \in B} \sum_{\substack{j \in B \\
i<j}} \eta_{\sigma}(i, j),  \tag{4}\\
K_{G B} & =\sum_{i \in G} \sum_{\substack{j \in B \\
i<j}} \eta_{\sigma}(i, j) .
\end{align*}
$$

For $N_{G}=N$ (i.e., $|B|=0$ ) it follows directly from $\mathbf{A 1}$ that $K=K_{G}=0$. On the other hand it follows directly from $\mathbf{A 2}$ that if $N_{G}<N$ (i.e., $|B|>0$ ), then $K>K_{B}>0$. Hence, a simple case to consider is when $K=0$ which implies $N_{G}=N$.

To obtain an explicit equation in $N_{G}$ for $K>0$, the terms in Eq. 3 are normalized by the maximal possible number of pairwise order inversions; that is, the number of pairs in each term. The number of pairs are given by $N(N-1) / 2$ in a sequence of length $N$, and by $N_{G} N_{B}$ between two disjoint sequences of lengths $N_{G}$ and $N_{B}$. That is,

$$
\begin{align*}
& \hat{K}_{B}=\frac{2 K_{B}}{N_{B}\left(N_{B}-1\right)}, \quad \hat{K}_{G}=\frac{2 K_{G}}{N_{G}\left(N_{G}-1\right)},  \tag{5}\\
& \hat{K}=\frac{2 K}{N(N-1)}, \quad \hat{K}_{G B}=\frac{K_{G B}}{N_{G} N_{B}}
\end{align*}
$$



Figure 2. A partial overlap between a pair of images from the ZuBuD dataset.

Using some algebraic manipulation after substituting the terms in Eq. 3 by Eq. 5 and replacing $N_{B}=N-N_{G}$, we obtain the following quadratic equation in $N_{G}$ :

$$
\begin{align*}
0= & N_{G}^{2}\left[\hat{K}_{G}+\hat{K}_{B}-2 \hat{K}_{G B}\right]+ \\
& N_{G}\left[2 N \hat{K}_{G B}-(2 N-1) \hat{K}_{B}-\hat{K}_{G}\right]+  \tag{6}\\
& N(N-1)\left(\hat{K}_{B}-\hat{K}\right) .
\end{align*}
$$

Note that $N$ is given and $\hat{K}$ can be directly computed from the set of matches, $\mathcal{M}$, using $E q$. 2. Using assumption A1 it directly follows that $\hat{K}_{G}=0$, from assumption $\mathbf{A 2}$ it follows that $E\left(\hat{K}_{B}\right)=1 / 2$, where $E(x)$ is the expected value of $x$ (see our proof in supplementary material; and alternative proof [11, p. 257]). We next show how to estimate $E\left(\hat{K}_{G B}\right)$ for fully overlapped and partially overlapped sequences.

### 4.2. Full Overlap

For fully overlapping sequences a third statistical assumption is used:

A3: The distribution of the features that correspond to correct matches within the features that correspond to incorrect matches is uniform in $[N]$ and $\sigma$; that is, the expected number of features in $G$ and $B$ in every interval of size $s$ is $s N_{G} / N$ and $s N_{B} / N$, respectively.

This assumption obviously does not hold for the nonoverlapping margins which contain only incorrect matches (see Sec. 4.3 which considers this case). In the full-overlap case, the following claim holds:

Claim 1. Under assumptions A1-A3, $E\left(\hat{K}_{G B}\right)=1 / 3$.
To prove this claim, let us first define $H_{\sigma}(i)$ to be the number of inversions in which a match, $(i, \sigma(i))$, participates. The desired value, $E\left(\hat{K}_{G B}\right)$, is obtained by averaging $E\left(H_{\sigma}(i)\right)$ over all $i \in G$ and normalizing it by $N_{G} N_{B}$. By definition, $H_{\sigma}(i)$ is the number of features from the left of $i$ that match features to the right of $\sigma(i)$ and those from the right of $i$ that match a features to the left of $\sigma(i)$. We can use assumption $\mathbf{A 2}$ to compute the expected value of $H_{\sigma}(i)$ for
$i \in G$ given the number of indexes with incorrect matches to the left and right of $i\left(\beta_{1}^{l}\right.$ and $\left.\beta_{1}^{r}\right)$ and to the left and right of $\sigma(i)\left(\beta_{2}^{l}\right.$ and $\left.\beta_{2}^{r}\right)$. The probability of inversion between $i \in G$ and $j<i$ is given by $\beta_{2}^{r} / N_{B}$. A similar argument is used for $j \in B$ and $j>i$. Hence we obtain:

$$
\begin{equation*}
E\left(H_{\sigma}(i) \mid i \in G, \beta_{1}^{l}, \beta_{1}^{r}, \beta_{2}^{l}, \beta_{2}^{r}\right)=\beta_{1}^{l} \frac{\beta_{2}^{r}}{N_{B}}+\beta_{2}^{l} \frac{\beta_{1}^{r}}{N_{B}} \tag{7}
\end{equation*}
$$

Using assumption A3 we can approximate $\sigma(i) \approx i$, and $E\left(\beta_{1}^{l}\right)=E\left(\beta_{2}^{l}\right)=i N_{B} / N$ and $E\left(\beta_{1}^{r}\right)=E\left(\beta_{2}^{r}\right)=(N-$ i) $N_{B} / N$; thus

$$
E\left(H_{\sigma}(i)\right)=2 \frac{i(N-i) N_{B}}{N^{2}}=2 N_{B} x(1-x)
$$

where $x=i / N$. The mean of $x(1-x)$ for $0 \leq x \leq 1$ is $1 / 6$, and thus
$E\left(\hat{K}_{G B}\right)=\frac{1}{N_{G} N_{B}} \sum_{i \in G} E\left(H_{\sigma}(i)\right)=\frac{2 N_{G} N_{B}}{N_{G} N_{B}} \frac{1}{6}=\frac{1}{3}$.
Substituting $E\left(\hat{K}_{G}\right), E\left(\hat{K}_{B}\right)$ and $E\left(\hat{K}_{G B}\right)$ with their obtained values in Eq. 6, we obtain the following quadratic equation in $N_{G}$ :

$$
\begin{equation*}
0=\frac{1}{6} N_{G}^{2}-\left(\frac{1}{2}-\frac{1}{3} N\right) N_{G}-N(N-1)\left(\frac{1}{2}-\hat{K}\right) \tag{8}
\end{equation*}
$$

There are two solutions to this equation; when $0 \leq \hat{K} \leq$ $1 / 2$, we take the only solution in the range $[0, N]$. When $\hat{K}>1 / 2$, we set $N_{G}=0$.

### 4.3. Partial Overlap

When the sequences are not fully overlapping (e.g., Fig. 2) the value of $E\left(K_{G B}\right)$ depends on the nonoverlapping margins of the sequences, defined below. In this case, the estimation of $N_{G}$ given by Eq. 8 is an underestimation (see Claim 2). We next describe a method to estimate $N_{G}$ by computing the regions of overlap, using this observation.

Formally, the sequence $[N]$ is partitioned into three intervals: $\mathcal{L}_{1}, \mathcal{O}_{1}$, and $\mathcal{R}_{1}$, defined by the lowest, $i_{L}^{1}$, and highest, $i_{H}^{1}$, indexes of the correct matches. Similarly, we define $\mathcal{L}_{2}, \mathcal{O}_{2}$, and $\mathcal{R}_{2}$ with the indexes $i_{L}^{2}$ and $i_{H}^{2}$. Note that an index of an incorrect match, $i \in \mathcal{O}_{1} \cap B$, is not necessarily matched to an index in $\mathcal{O}_{2}$. To use our results for the fully overlapped sequences (Sec. 4.2), we discard such indexes and define the fully overlapped subsequences, $\hat{\mathcal{O}}=\left(\hat{\mathcal{O}}_{1}, \hat{\mathcal{O}}_{2}\right)$ as follows:

$$
\begin{equation*}
\hat{\mathcal{O}}_{1}=\left\{i \mid\left(i \in \mathcal{O}_{1}\right) \wedge\left(\sigma(i) \in \mathcal{O}_{2}\right)\right\} \tag{9}
\end{equation*}
$$

and $\hat{\mathcal{O}}_{2}$ is the sequence of indexes of the matched features to $\hat{\mathcal{O}}_{1}$.

Note that $\hat{\mathcal{O}}_{1}$ and $\hat{\mathcal{O}}_{2}$ are defined by the 4 -tuple of indexes, $\omega^{*}=\left(i_{L}^{1}, i_{H}^{1}, i_{L}^{2}, i_{H}^{2}\right)$. Moreover, $N, N_{G}$ and $\hat{K}$,
on the subsequences defined by $\omega^{*}$ are given by $N\left(\omega^{*}\right)$, $N_{G}\left(\omega^{*}\right)$ and $\hat{K}\left(\omega^{*}\right)$. We can now use these new values in Eq. 8 to obtain:

$$
\begin{align*}
0= & \frac{1}{6} N_{G}^{2}\left(\omega^{*}\right)-\left(\frac{1}{2}-\frac{1}{3} N\left(\omega^{*}\right)\right) N_{G}\left(\omega^{*}\right) \\
& -N\left(\omega^{*}\right)\left(N\left(\omega^{*}\right)-1\right)\left(\frac{1}{2}-\hat{K}\left(\omega^{*}\right)\right) . \tag{10}
\end{align*}
$$

In this case we claim the following:
Claim 2. The expected maximal value of $N_{G}(\omega)$, for any $\omega$, is obtained for $\omega=\omega^{*}$, i.e.,

$$
\begin{equation*}
\max _{\omega} N_{G}(\omega)=N_{G}\left(\omega^{*}\right) \tag{11}
\end{equation*}
$$

We use assumptions A1-A3 to prove it (see the supplementary material). We use this claim to estimate $\omega^{*}$, which determines the region of overlap between the pair of images, as described below. Given $\omega^{*}$ we can estimate the desired $N_{G}$ using Eq. 10 .
Estimating $\omega^{*}$. To compute $N_{G}\left(\omega^{*}\right)$, we need to consider all 4-tuples $\omega \in[1, \ldots, N]^{4}$. For efficiency, we consider only a small subset of intervals. Formally, we consider all intervals in the sequence of the form $[i, j]=\left[t_{i} q+1, t_{j} q\right]$ where $0<t_{i}<t_{j} \leq N / q$ are integers. The algorithm that considers all these intervals in $I_{1}$ and $I_{2}$, we name $K 2 D$. To further improve efficiency, we propose the $K 1 D$ algorithm which considers sequentially the intervals in each image. That is, we first use the intervals, $\omega=\left(i_{L}^{1}, j_{H}^{1}, 1, N\right)$, to sample the values of $N_{G}(\omega)$. The resulting interval, $\left[i_{L}^{1}, j_{H}^{1}\right]$, is then used to fix the interval in $I_{1}$ and estimate the interval in $I_{2}$; that is, we use the intervals $\omega=\left(i_{L}^{1}, j_{H}^{1}, i_{L}^{2}, j_{H}^{2}\right)$ to arrive at our final estimate for $\omega^{*}$ and $N_{G}\left(\omega^{*}\right)$.
Kendall distance computation. The Kendall distance can be computed in $\mathcal{O}(N \log N)$ time using the merge sort algorithm [19], applied on $\sigma$. The basic idea is that the number of inversions can be computed at the merge stage (when merging two sorted arrays into one). The number of inversions that should be added to the count is the number of elements that remain in the left array when the next minimal element is taken from the right array.

For partial overlap, maximizing the estimated value of $N_{G}(\omega)$ in Eq. 11 requires multiple Kendall distance computations for various intervals. We compute these distances efficiently by first computing the Kendall distance for each of the non-overlapping intervals of size $q$ using the merge sort algorithm. The Kendall distance of an interval of size $d q$ is obtained by counting the inversions when merging two successive intervals of size $d_{1} q$ and $d_{2} q$, where $d_{1}+d_{2}=d$.

## 5. Matching Probabilities

We propose to compute a probability, $P_{K}(i)$, that a match, $m=\left(p_{i}, q_{\sigma(i)}\right)$, is correct, given the computed $H_{\sigma}(i)$ (defined in Sec. 4.2) and our estimation of


Figure 3. A graph of order inversions, $H_{\sigma}$, for an image from the Middlebury 2014 dataset. The $x$-axis corresponds to the rank of the feature in $I_{1}$, and the $y$-axis corresponds to the percentage of order inversions out of the maximal value, $N-1$. Computed by USAC, the blue and red dots correspond to the inliers and outliers, respectively. The green and cyan lines correspond to the range of inversions for the incorrect matches. The expected number of inversions for the correct matches is given by the green line.
$\theta=\left(\hat{\mathcal{O}}_{1}, \hat{\mathcal{O}}_{2}, N_{G}\right)$. This probability can then be used as a likelihood function for sampling matches, in particular in guided RANSAC methods.

When $i \notin \hat{\mathcal{O}}_{1}$ or $\sigma(i) \notin \hat{\mathcal{O}}_{2}$ we set $P_{K}(i)=0$ (or to a small value), since, by the definition of $\hat{\mathcal{O}}_{1}$ and $\hat{\mathcal{O}}_{2}, i \in B$. Hence we would like to compute $P_{K}(i)$ for the remaining indices. For a full overlap, a typical graph of the number of order inversions, $H_{\sigma}(i)$, as a function of the index $i$ is presented in Fig. 3. For $i \in B$ (blue dots), $H_{\sigma}(i)$ is approximately uniformly distributed in a range which depends on $i$. As expected from the analysis presented in Sec. 4.2, the distribution of $i \in G$ (red dots) is around the function $x(1-x)$, where $x=i / N$ and $N=\left|\hat{\mathcal{O}}_{1}\right|$.

To compute $P_{K}(i)$ for $i \in G$, we use the Bayes' theorem and the law of total probability:

$$
\begin{align*}
P_{K}(i)=P\left(i \in G \mid H_{\sigma}(i)\right) & =\frac{P\left(H_{\sigma}(i) \mid i \in G\right) P(i \in G)}{P\left(H_{\sigma}(i)\right)} \\
& =\frac{P_{H \mid G} P_{G}}{P_{H \mid G} P_{G}+P_{H \mid B} P_{B}}, \tag{12}
\end{align*}
$$

where $P_{H \mid G}=P\left(H_{\sigma}(i) \mid i \in G\right), P_{H \mid B}=P\left(H_{\sigma}(i) \mid i \notin\right.$ $G), P_{G}=P(i \in G)$ and $P_{B}=P(i \notin G)$. We assume w.l.g. that there is a full overlap; hence, the probabilities $P_{G}=N_{G} / N$ and $P_{B}=1-P_{G}$ are given by the ratio of the correct and the incorrect matches to $N$, respectively.

We next describe our estimation of $P_{H \mid G}$. Consider $\beta=$ $\left(\beta_{1}^{l}, \beta_{1}^{r}, \beta_{2}^{l}, \beta_{2}^{r}\right)$ (defined in Sec. 4.2). Note that $\beta_{1}^{r}$ and $\beta_{2}^{r}$ are given directly from $\beta_{1}^{l}, \beta_{2}^{l}, N$, and the estimated $N_{G}$. Since $\beta$ is unknown, we compute $P_{H \mid G}$ using the law of

| Dataset | Value | KGT | K | K2D | K1D |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Syn1 | $\mu\left(\mathbf{N}_{\mathbf{G}}\right)$ | 0.6 | 14.5 | 3.2 | 4 |
|  | $\mu(\mathcal{O})$ | n.a. | n.a. | 0.91 | 0.89 |
|  | Runtime | 2 | 2.2 | 83.2 | 11.2 |
| Syn2 | $\mu\left(\mathbf{N}_{\mathbf{G}}\right)$ | 0.5 | 10.4 | 3.2 | 3.7 |
|  | $\mu(\mathcal{O})$ | n.a. | n.a. | 0.9 | 0.8 |
|  | Runtime | 2.9 | 2 | 79.2 | 11.1 |

Table 1. The mean normalized absolute error (percentage) in the estimation of $N_{G}$ in the synthetic datasets.
total probability over the set of possible values, $\beta \in S_{\beta}$ :

$$
P\left(H_{\sigma}(i) \mid i \in G\right)=\sum_{\beta \in S_{\beta}} P\left(H_{\sigma}(i) \mid i \in G, \beta\right) P(\beta) .
$$

For efficiency, only $k$ values of $\beta \in S_{\beta}$ are considered by restricting the values of $\beta_{1}^{l}$ and $\beta_{2}^{l}$ to be around their expectations, $(i-1) P_{B}$ and $(\sigma(i)-1) P_{B}$, respectively (in our implementation we set $k=5$ ).

The probability $P(\beta)$ is modeled as the product of two hypergeometric distributions, one for $\beta_{1}^{l}$ and the other for $\beta_{2}^{l}$. Given a value for $\beta$, we model $P\left(H_{\sigma}(i) \mid i \in G, \beta\right)$ as a product of two hypergeometric distributions. In a similar way we estimate $P\left(H_{\sigma}(i) \mid i \in B, \beta\right)$. Substituting these terms into Eq. 12 yields $P_{K}(i)$. We approximate the hypergeometric PDFs by Gaussian PDFs with the same mean and standard deviation, and approximate $P_{H \mid B}$ as a uniform distribution whose boundaries we compute (see the supplementary material for the complete derivation).

## 6. Experiments

Our method was tested on both synthetic and real data. For each dataset we tested the following algorithms. The first, which we refer to as K , is the basic one where $N_{G}$ is estimated using the entire sequence and ignoring the margins; that is, without first computing the overlap (see Sec. 4.2). The other two algorithms compute $N_{G}, \mathcal{O}_{1}$ and $\mathcal{O}_{2}$. One uses the joint sampling of intervals from $I_{1}$ and $I_{2}$, while the other uses a separate sampling of intervals in $I_{1}$ and $I_{2}$ (see Sec. 4.3). These algorithms are referred to as K2D and K1D, respectively. For evaluation purpose we also consider KGT, where the ground truth of $\mathcal{O}_{1}$ and $\mathcal{O}_{2}$ are given, and compute only $N_{G}$.

To evaluate the accuracy of $N_{G}$ we measure the normalized absolute error, $\mu\left(N_{G}\right)=\left|N_{G}-N_{G}^{G T}\right| / N$, where $N_{G}^{G T}$ is the ground truth. Note that $\mu\left(N_{G}\right)$ is presented as a percentage in all tables. To evaluate the accuracy of $\mathcal{O}_{1}$ and $\mathcal{O}_{2}$, we measure the average of their normalized overlap with the ground truth; that is, the average between the two $\mu(\mathcal{O})=100\left|\mathcal{O} \cap \mathcal{O}_{G T}\right| /\left|\mathcal{O} \cup \mathcal{O}_{G T}\right|$, for $\mathcal{O} \in\left\{\mathcal{O}_{1}, \mathcal{O}_{2}\right\}$, where $\mathcal{O}_{G T}$ is the ground truth.


Figure 4. Scatter plot for synthetic experiment Syn2.

### 6.1. Synthetic Data

A matching between features in a pair of images is depicted by a permutation. We generated 500 different permutations of $N=1000$ indexes, using our assumptions A1A3, for various values of $\mathcal{O}_{1}, \mathcal{O}_{2}$ and $N_{G}$.

Test Syn1: We set $N_{G}=300$, the size $\left(>N_{G}\right)$ and location of $\mathcal{O}_{1}$ and $\mathcal{O}_{2}$ are randomly chosen. Tab. 1 presents $\mu\left(N_{G}\right), \mu(\mathcal{O})$ and the mean runtime (displayed in milliseconds). The error $\mu\left(N_{G}\right)$ is the largest for K, while K2D and K1D are similar and closer to KGT.

Test Syn2: We set $N_{G}$ between 0 to 1000, and randomly choose $\mathcal{O}_{1}$ and $\mathcal{O}_{2}$ as in "Test Syn1" (see results in Tab. 1). The scatter plot of the four algorithms with respect to the ground truth is presented in Fig. 4. The K algorithm, where $\left|\mathcal{O}_{1}\right|=\left|\mathcal{O}_{2}\right|=N$ is wrongly assumed, underestimates $N_{G}$, and $\mu\left(N_{G}\right)$ is relatively large. The estimation using K2D and K1D are similar and slightly underestimate of $N_{G}$. These results are probably due to the sampling used for estimating $\mathcal{O}_{1}$ and $\mathcal{O}_{2}$, since for KGT, the errors are negligible. For both tests, K1D is almost an order of magnitude faster than K2D.

### 6.2. Real Data

We use the Middlebury 2005\&2006 and 2014 datasets [26, 16, 25], for real data with ground truth (denoted by MFull05\&06 and MFull14, respectively). We also evaluate our method on the USAC dataset [22] and BLOGS dataset [3] (combined), and the ZuBuD dataset [27], for which no ground truth is available (denoted by U\&B and ZuBuD, respectively). We run the BEEM [13] and USAC [22] algorithms (both with their default settings), where the number of inliers returned is compared to $N_{G}$. The lowest and highest inlier indexes returned by these methods can be regarded as their estimations of $\mathcal{O}_{1}$ and $\mathcal{O}_{2}$.

Middlebury Full Overlap. The mean errors and runtimes (in milliseconds) are presented in Tab. 2. The error, $\mu\left(N_{G}\right)$, is similar for K, K2D and K1D, while K2D is much slower The K1D and BEEM are similar in $\mu\left(N_{G}\right)$, while USAC is more accurate. The runtime of K1D is between one and two orders of magnitude faster than USAC and BEEM's run-

| Dataset | \#im pairs | $N_{G}^{G T}$ | BEEM |  | USAC |  | K |  | K2D |  | K1D |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time |
| MFull14 | 23 | 1584 | 9.9 | 2065 | 3 | 1782 | 10 | 30 | 9.9 | 291 | 9.9 | 62 |
| MFull05\&06 | 27 | 2776 | 13 | 4198 | 3 | 1670 | 3.3 | 51 | 3.3 | 427 | 3.3 | 76 |
| MPar14 | 23 | 712 | 5 | 1825 | 3.6 | 1152 | 11.5 | 13 | 6.8 | 195 | 7.1 | 31 |
| MPar05\&06 | 27 | 1032 | 5.3 | 2344 | 4.1 | 1276 | 23.1 | 11 | 6.3 | 249 | 6.7 | 39 |

Table 2. The mean normalized absolute error (percentage), $\mu\left(N_{G}\right)$, in comparison to the ground truth (its mean, $N_{G}^{G T}$ ) and the mean runtime (in milliseconds) for each of the methods in the Middlebury datasets.

| Dataset | \#im pairs | $\left\|G_{B E E M}\right\|$ | BEEM |  | USAC |  | K |  | K2D |  | K1D |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time | $\mu\left(N_{G}\right)$ | Time |
| U\&B | 21 | 138 | n.a. | 1444 | 8.6 | 1440 | 11.7 | 11 | 8.8 | 142 | 9 | 18 |
| ZuBuD | 201 | 69 | n.a. | 692 | 7.3 | 1257 | 9.3 | 2 | 6.1 | 90 | 6.5 | 11 |

Table 3. The mean normalized absolute error (percentage), $\mu\left(N_{G}\right)$, in comparison to BEEM (its mean, $\left|G_{B E E M}\right|$ ) and the mean runtime (in milliseconds) for each of the methods in the USAC\&BLOGS and ZuBuD datasets.
time. Fig. 5 presents the complete comparison details.
Middlebury Partial Overlap. The overlap between image pairs in the Middlebury datasets is very large (about $90 \%$ of the image). Therefore, we randomly vertically cut them to obtain smaller overlaps. This enables us to test the partial-overlap method (Sec. 4.3) against the ground truth (Tab. 2). These datasets are denoted by MPartial05\&06 and MPartial14 for Middlebury 2005\&2006 and 2014, respectively. Both the errors of K2D and K1D are low and similar to the full overlap case, demonstrating the success of the partial overlap method. However, the running time is much lower for K1D. The error, $\mu\left(N_{G}\right)$, is large for method K since it ignores the margins. As in the full overlap case, BEEM is similar in accuracy to K2D and K1D, while USAC is more accurate. The similarity of the $\mathcal{O}$ evaluation, $\mu(\mathcal{O})$, is 0.89 and 0.89 for both the K2D and K1D methods. Fig. 5 presents the complete comparison details.
USAC \& BLOGS and ZuBuD. Since no ground truth is available for these datasets, the estimation results are compared to BEEM (see Tab. 3). The errors for K2D and K1D and USAC are similar ( K is larger as expected). K2D and K1D are one and two orders of magnitude faster than BEEM and USAC. Fig. 5 presents the complete comparison details.

### 6.3. Applications

We next evaluate three applications of $N_{G}$ estimation.
Halting condition. The classic halting condition [14, Chapter 11.6] for adaptive RANSAC methods is based on the probability that at iteration $t$ at least one consensus set was constructed from an uncontaminated minimal set of matches. This probability is computed as a function of the number of correct matches. As this number is usually unknown, the probability is updated when a larger consensus set is found. The method halts when the confidence in the solution is high based on the size of the consensus set. Since
our algorithm is able to estimate $N_{G}$, we are able to halt RANSAC instantly when a consensus set of at least $N_{G}$ matches has been found

To test this we run USAC with and without the halting condition on all the datasets and compare the inlier rate and runtime (see Fig. 7). The total runtime with the halting condition is $69 \%$ of the total runtime without the halting condition. This is achieved with nearly no loss of inliers (normalized absolute error of $0.78 \%$ ).

SfM image pairs. A major time-consuming stage in the pipeline of SfM methods [30] is running RANSAC on all image pairs to filter out incorrect matches. We propose to run RANSAC only on image pairs with sufficiently large $N_{G}$, avoiding spatially unrelated image pairs.

We use the dataset "Barcelona" [8], which consists of 191 images taken from all around a building. We consider the total number of inliers and the total runtime from running USAC on all image pairs as $100 \%$ of inliers and runtime. Our goal is to maintain the inliers, while reducing the runtime. We use a threshold, $t$, to discard image pairs for which our estimation of $N_{G}<t$, using the K2D and K1D algorithms. Fig. 6(a)\&(b) shows the reduction in inliers and runtime as a function of $t$.

For K2D, the total runtime decreases to $17 \%$ (or $4.5 \%$ ) while maintaining $91 \%$ (or $86 \%$ ) of the inliers. In both of these examples, all images remain with at least one other image for which the fundamental matrix could be calculated. For K1D, the total runtime decreases to $7 \%$ (or $1.2 \%$ !) while maintaining $87 \%$ (or $80 \%$ ) of the inliers. In the second example only one image out of the 191 images was lost. Note that the percentage of time spent discarding matches outliers in a typical SfM method (e.g., [30]) is $\sim 25 \%$ of the total runtime.
Matching probability. Computing matching probabilities, $P_{K}$, can be useful for several applications, including


Figure 5. Scatter plots for the real datasets. From left to right: K, K2D and K1D methods. Top against the ground truth (only in Middlebury datasets), and bottom is a comparison against BEEM. The axes correspond to $N_{G} / N$.


Figure 6. (a)-(b) Using our estimate for $N_{G}$ to discard pairs of images from a SfM pipeline in the "Barcelona" dataset; (a) percentage of the remaining correct matches as a function of the threshold; (b) percentage of the remaining runtime as a function of the threshold; (c) Comparison of sampling methods for Guided RANSAC using the pairs of images from all datasets (except for "Barcelona"). The $x$ axis corresponds to the rank of the matches when sorted according to the probability. The $y$ axis corresponds to the precision of correct matches; that is, a point on the graph, $(x, y)$, denotes the ratio of the number of correct matches from the $x$ matches to all $x$ matches.
guided RANSAC. For evaluating our estimate for $P_{K}$, we present the mean inlier precision of the $x \leq 200$ highest ranked matches, given by $P_{K}$ (Fig. 6(c)). It is compared with the Lowe ranking ([21]), denoted by $P_{L}$. The mean precision was calculated on all image pairs from all datasets. The curve for the mean $P_{K}$ (green) is below the curve for the mean $P_{L}$ (red) for the first 50 matches, and then vice versa. A combination of the two probabilities, given by $P_{C}=\frac{P_{K} P_{L}}{P_{K} P_{L}+\left(1-P_{K}\right)\left(1-P_{L}\right)}$. is the blue curve. Using low values of $x, P_{C}$ outperforms its components. For higher values of $x$, our estimation is better.

## 7. Conclusions and Future Work

In this paper we introduced a novel method to estimate the number of correct matches, and a method to estimate the region of overlap between a pair of images. We also derive a probability function for a match to be correct. All these were done using only the spatial order of a given set of matches, and some reasonable statistical assumptions. We demonstrated the effectiveness of this estimation on real datasets for an analysis of feature matches using only the
spatial order of features in each image. Our method successfully competes with methods that compute the set of inliers explicitly by recovering the epipolar geometry, but at a much lower cost. It is impressive to see what can be learned about matching using only spatial orders.

Acknowledgments: This work was partially supported by the Israel Science Foundation, grant no. 930/12, and by the Israeli Ministry of Science, grant no. 3-8744.


Figure 7. USAC for inlier rate estimation without the halting condition ( $y$-axis) vs with the halting condition ( $x$-axis).

## References

[1] H. Bay, T. Tuytelaars, and L. V. Gool. SURF: speeded up robust features. In ECCV. 2006.
[2] A. F. Bobick and S. S. Intille. Large occlusion stereo. International Journal of Computer Vision, 33(3):181-200, 1999.
[3] A. S. Brahmachari and S. Sarkar. Hop-diffusion monte carlo for epipolar geometry estimation between very widebaseline images. IEEE Trans. Patt. Anal. Mach. Intell., 35(3):755-762, 2013.
[4] M. Calonder, V. Lepetit, C. Strecha, and P. Fua. Brief: Binary robust independent elementary features. In ECCV. 2010.
[5] O. Chum and J. Matas. Matching with PROSAC-progressive sample consensus. In CVPR, 2005.
[6] O. Chum and J. Matas. Optimal randomized RANSAC. IEEE Trans. Patt. Anal. Mach. Intell., 30(8):1472-1482, 2008.
[7] O. Chum, J. Matas, and J. Kittler. Locally optimized RANSAC. In Pattern recognition, pages 236-243. Springer, 2003.
[8] A. Cohen, C. Zach, S. N. Sinha, and M. Pollefeys. Discovering and exploiting 3 d symmetries in structure from motion. In CVPR, 2012.
[9] M. Deza, L. N. Supérieure, and T. Huang. Metrics on permutations, a survey. In Journal of Combinatorics, Information and System Sciences. Citeseer, 1998.
[10] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In $W W W, 2001$.
[11] W. Feller. An introduction to probability theory and its applications, volume 2. John Wiley \& Sons, 2008.
[12] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Comm. of the ACM, 24(6):381-395, 1981.
[13] L. Goshen and I. Shimshoni. Balanced exploration and exploitation model search for efficient epipolar geometry estimation. IEEE Trans. Patt. Anal. Mach. Intell., 30(7):12301242, 2008.
[14] R. Hartley and A. Zisserman. Multiple View Geometry in Computer Vision. Cambridge University Press, 2003.
[15] W. Hartmann, M. Havlena, and K. Schindler. Predicting matchability. In CVPR, 2014.
[16] H. Hirschmüller and D. Scharstein. Evaluation of cost functions for stereo matching. In CVPR, 2007.
[17] E. Johns, D. Edward, and G. Z. Yang. Pairwise probabilistic voting: Fast place recognition without RANSAC. In $E C C V$. 2014.
[18] M. G. Kendall. A new measure of rank correlation. Biometrika, pages 81-93, 1938.
[19] J. Kleinberg and É. Tardos. Algorithm design. Pearson Education India, 2006.
[20] R. Litman, S. Korman, A. Bronstein, and S. Avidan. Inverting RANSAC: Global model detection via inlier rate estimation. In CVPR, 2015.
[21] D. G. Lowe. Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision, 60(2):91-110, 2004.
[22] R. Raguram, O. J. Chum, M. Pollefeys, J. Matas, and J. M. Frahm. USAC: a universal framework for random sample consensus. IEEE Trans. Patt. Anal. Mach. Intell., 35(8):2022-2038, 2013.
[23] S. Ramalingam, M. Antunes, D. Snow, L. G. Hee, and S. Pillai. Line-sweep: Cross-ratio for wide-baseline matching and 3D reconstruction. In CVPR, 2015.
[24] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. ORB: an efficient alternative to SIFT or SURF. In ICCV, 2011.
[25] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In Pattern Recognition, pages 31-42. Springer, 2014.
[26] D. Scharstein and C. Pal. Learning conditional random fields for stereo. In CVPR, 2007.
[27] H. Shao, T. Svoboda, and L. V. Gool. Zubud-zurich buildings database for image based recognition. Computer Vision Lab, Swiss Federal Institute of Technology, Switzerland, Tech. Rep, 260, 2003.
[28] J. M. Steele. Variations on the monotone subsequence theme of erdös and szekeres. In Discrete probability and algorithms, pages 111-131. Springer, 1995.
[29] R. Szeliski. Computer vision: algorithms and applications. Springer Science \& Business Media, 2010.
[30] C. Wu. Towards linear-time incremental structure from motion. http://ccwu.me/vsfm/, 2013.
[31] A. L. Yuille and T. Poggio. A generalized ordering constraint for stereo correspondence. Technical report, DTIC Document, 1984.

