Supplemental Material

A. Additional Experiments on the EPFL dataset

We test MatchFAME on the 6 EPFL datasets following the experimental setup of [13]. Each dataset includes 8 to 30 images, unlike the large number of images in the Photo Tourism datasets. Given each dataset, we generate and refine the initial keypoint matches with the same procedure introduced in [13]. We follow their convention and estimate the universe size with $\hat{m}=2\lceil M/n\rceil$. We implement MatchFAME with its default parameters, though with two changes described below. Indeed, the EPFL dataset contains a lot of noisy edges and thus the weights produced by PPM within the original MatchFAME algorithm are often small. Furthermore, note that $\mathbf{Proj}$ in (9) is not scale invariant and that the resulting small weights may lead to overly sparse refined matches. Therefore, we slightly changed the implementation of MatchFAME to overcome this issue. First, in order to obtain a dense initialization of partial permutations using MST, instead of assigning 1 to a random element for each zero column, we assign 1 to a random element for each zero row. Since the number of rows is larger than the number of columns, this modification results in a denser initialization of $[\boldsymbol{P}_j^{(0)}]_{j\in[n]}$ than that of the original MatchFAME. Second, to make sure that the final output is also sufficiently dense, we drop the step of the weights' normalization within the PPM iterations, which is described below (9) (this will increase the overall scale of the edge weights and thus the projected matrix is expected to be denser). We remark that these two changes help alleviate the over-sparseness of the final output and ends up with a higher ratio between the number of refined matches and the number of initial matches, which we denote by $\#M$.

In addition to this version of MatchFAME, we also test Spectral, MatchEIG and MatchALS with the same setting as [13]. Note that the 'ground truth' is obtained by estimating the projection distance of key points on the epipolar line instead of labeling by hand. Therefore the recall score is not a good benchmark on real data. We thus only report the resulting precision, number of remaining edges and runtime in Table 4.

MatchFAME achieves the highest precision of all methods in all datasets. Observing $\#M$, we note that MatchFAME has around 20% fewer matches remaining compared to all algorithms, but as long as there are enough matches for each edge, one can reliably compute relative rotations and translations for SfM tasks. We believe removing around 20% more matches is not an essential drawback. Furthermore, MatchFAME is faster than the other methods. In conclusion, MatchFAME can achieve a reasonable estimate of matches within a significant short amount of time.

B. Clarifications

We clarify some definitions and expand on various claims mentioned in the paper.

B.1. More on Cycle Consistency and Inconsistency

We referred to a cycle ijk as consistent whenever $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\leq \boldsymbol{X}_{ik}$, $\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}\leq \boldsymbol{X}_{ji}$ and $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}\leq \boldsymbol{X}_{kj}$. Note that $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}$ is a binary matrix with ones whenever there are paths of lengths 2 between keypoints of images i and k and $\boldsymbol{X}_{ik}$ is binary matrix with ones whenever there are paths of lengths 1 (single edges) between keypoints of image i and k. That is, $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\leq \boldsymbol{X}_{ik}$ means that if keypoints $t_i\in[m_i]$ and $t_k\in[m_k]$ (in images i and k, respectively) are both matched to a keypoint t_j in image j, then they are matched to each other. Therefore, any cycle ijk with corresponding partial permutations $\boldsymbol{X}_{ij}$, $\boldsymbol{X}_{jk}$, $\boldsymbol{X}_{ki}$ is consistent if and only if for any $t_i\in[m_i]$, $t_j\in[m_j]$ and $t_k\in[m_k]$: If two of the events $\boldsymbol{X}_{ij}(t_i,t_j)=1$, $\boldsymbol{X}_{jk}(t_j,t_k)=1$, $\boldsymbol{X}_{ki}(t_k,t_i)=1$ hold true, then the third one holds true as well.

This equivalent reformulation of cycle consistency further clarifies the definition of d_{ijk} in (4). For fixed $\boldsymbol{X}_{ij},\boldsymbol{X}_{jk},\boldsymbol{X}_{ki}\in\mathcal{P}^{l_1,l_2}$, the denominator of the fraction in (4) can be viewed as the number of combinations of three keypoints a, b, c, such that at least two of the three events

$$\boldsymbol{X}_{ij}(a,b)=1,\ \boldsymbol{X}_{jk}(b,c)=1, \text{ and } \boldsymbol{X}_{ki}(c,a)=1 \qquad (10)$$

hold. Furthermore, the numerator of the fraction in (4) can be viewed as the total number combinations of three keypoints a, b, c, such that all the three events in (10) hold. Thus, the fraction in (4) indeed measures the level of cycle consistency, and consequently d_{ijk} measures the cycle inconsistency.

We remark that an inequality of two full permutation matrices must be an equality. Therefore, for permutation synchronization the above definition of cycle consistency is equivalent with $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}=\boldsymbol{X}_{ik}$ (or equivalently, $\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}=\boldsymbol{X}_{ji}$ or $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}=\boldsymbol{X}_{kj}$ or $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}=\boldsymbol{I}$). That is, our definition of cycle consistency is a direct extension of the one in group synchronization.

B.2. Cycle-verifiability Helps in Verifying Matches in Cycles

We further interpret the cycle-verifiable condition and clarify its name. We claim that if ijk is a good cycle (w.r.t. ij) ensured by Definition 1 with $a\in I_i$ and $c\in I_j$, then one can verify whether a and c correctly match (i.e., $h(a)=h(c)$) using $b\in I_k$. Indeed, since b matches a and $k\in G_{ij}$, $h(a)=h(b)$. If b and c match then since $k\in G_{ij}$ $h(b)=h(c)$ and consequently $h(a)=h(c)$. Assume on the other hand that b and c do not match. If b matches another point c', then since $k\in G_{ij}$, $h(b)=h(c')\neq h(c)$. If b does not match any point in U_j, then since $k\in G_{ij}$, $h(b)\notin h(U_j)$ (otherwise there exists $c'\in U_j$ such that $h(c')=h(b)$ and since $k\in G_{ij}$ there has to be a match between b and c'.). Since $h(c)\in h(U_j)$ and $h(b)\notin h(U_j)$, $h(b)\neq h(c)$.

C. Proof of Theorem 1

The proof establishes two lemmas, Lemmas 1 and 2, and then uses them to conclude Theorem 1. It is rather technical

Algorithms / Dataset			Initial	MatchEig			Spectral			MatchALS			PPM			MatchFAME (ours)		
	n	$\hat{m}$	PR	PR	#M	T	PR	#M	T	PR	#M	T	PR	#M	T	PR	#M	T
Herz-Jesu-P25	25	517	89.6	94.2	73	72	92.2	81	125	93.3	83	9199	92.5	**88**	125	**95.0**	78	**15**
Herz-Jesu-P8	8	386	94.3	95.2	**97**	**1**	95.3	92	4	95.9	76	155	95.4	94	5	**95.9**	83	3
Castle-P30	30	445	71.8	84.7	55	64	80.6	72	99	80.4	76	13583	80.2	**77**	112	**87.9**	61	**15**
Castle-P19	19	314	70.1	79.7	57	23	76.3	**76**	21	77.0	74	1263	77.5	**76**	33	**83.0**	56	**4**
Entry-P10	10	432	75.4	79.9	78	11	82.1	78	30	77.3	77	322	80.7	**83**	34	**83.1**	69	**5**
Fountain-P11	11	374	94.2	95.4	81	8	95.4	93	14	95.7	82	333	95.6	**94**	18	**96.7**	81	**5**

Table 4. Performance on the EPFL datasets. n is the number of cameras; $\hat{m}$, the approximated m, is twice the averaged m_i over $i \in [n]$; PR refers to the precision $|\hat{E} \cap E_g|/|\hat{E}|$, which is expressed in percentage (the higher the better); #M is the ratio (expressed in percentage) between the number of refined matches and the number of initial matches; T is runtime in seconds.

and not so easy to motivate. In order to provide more intuition, we added some clarifying figures.

Convention for figures: In all of these figures, we designate by green lines good keypoint matches, by red lines bad keypoint matches and by dashed orange lines missing keypoint matches. All of these occur between keypoints of two different images. On the other hand, matches between keypoints in an image and universal 3D keypoints are designated by black dotted lines (these correspond to our formal h function). We further color the universal 3D keypoints (in U_{ij}), which represent elements of U_{ij}^{good}, by green. We also color the universal 3D keypoints, which represent elements of U_{ij}^{bad}, in red. In Figure 5, we slightly extend the latter convention and explain it in its caption.

Terminology Review: Recall that $n_\Delta = \text{Tr}(\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki})$, m is the number of all 3D keypoints, m_{ij} is the number of 3D keypoints that correspond to the 2D keypoints of images i or j and among these, m_{ij}^{bad} is the number of keypoints that match wrong keypoints in the other image, or match no keypoint if a ground-truth match exists. We also denote the number of the rest of points by m_{ij}^{good} (that is, $m_{ij}^{\text{good}} = m_{ij} - m_{ij}^{\text{bad}}$) and recall that these keypoints match the ground-truth keypoints or, do not match any keypoint, if no ground-truth matches exist.

C.1. Upper Bound for the Cycle Inconsistency of Good Edges

This section includes the proof of the following lemma:

Lemma 1. For any $ij \in E_g$, $d_{ijk} \le m(s_{ik}^* + s_{jk}^*)$.

We remark that in the case of group synchronization, in particular, PS, one can easily show that for any $ij \in E$, $|d_{ijk} - s_{ij}^*| \le s_{ik}^* + s_{jk}^*$ (see Lemma 1 of [11]). Consequently for $ij \in E_g$, $d_{ijk} \le s_{ik}^* + s_{jk}^*$. However, in PPS, without the full group structure with a bi-invariant metric, it is harder to prove the weaker bound of Lemma 1. The proof below involves various discrete combinatorial arguments.

Proof. Assume first that $n_\Delta = 0$ and note that (4) implies $d_{ijk} = 1$. Since $d_{ijk} \neq 0$ and $ij \in E_g$, s_{ik}^* and s_{jk}^* cannot be both zero (otherwise this and the fact that $ij \in E_g$ imply that ijk is cycle-consistent and thus $d_{ijk} = 0$). Without loss of generality, assume $s_{jk}^* > 0$. We note that

$$s_{jk}^* = \frac{m_{jk}^{\text{bad}}}{m_{jk}} \ge \frac{1}{m},$$

which implies the desired bound:

$$d_{ijk} = 1 = m \cdot \frac{1}{m} \le m s_{jk}^* \le m(s_{ik}^* + s_{jk}^*).$$

Assume next that $n_\Delta > 0$, or equivalently,

$$n_\Delta \ge 1. \tag{11}$$

The next arguments require additional definitions and observations. We recall that any element of I_i represents a 2D keypoint in image i. This keypoint is associated with the index vector (i,j), where $j = 1,...,m_i$ and we can thus view I_i as the set of m_i index vectors. For cycle ijk, $(a,b,c) \in I_i \times I_j \times I_k$ is an (i,jk) tuple if there is a match between a and both b and c (the match can be either good or bad). If (a,b,c) is an (i,jk) tuple and there is no match between b and c, then we refer to (a,b,c) as a bad (i,jk) tuple, otherwise, it is a good (i,jk) tuple. For example, in Figure 5, there are three (i,jk) tuples: $(p_{i,2},p_{j,1},p_{k,1})$, $(p_{i,3},p_{j,3},p_{k,2})$ and $(p_{i,4},p_{j,2},p_{k,3})$. We note that $(p_{i,2},p_{j,1},p_{k,1})$ and $(p_{i,4},p_{j,2},p_{k,3})$ are bad (i,jk) tuples and $(p_{i,3},p_{j,3},p_{k,2})$ is a good (i,jk) tuple. For cycle ijk, we denote by $A_{i,jk}$ the set of (i,jk) tuples in $I_i \times I_j \times I_k$, by $A_{i,jk}^{\text{bad}}$ the set of bad (i,jk) tuples and by $A_{i,jk}^{\text{good}}$ the set of good (i,jk) tuples.

Recall that for a cycle ijk, $n_i = \mathsf{nnz}(\boldsymbol{X}_{ki}\boldsymbol{X}_{ij})$, n_j and n_k are analogously defined, and $n_\Delta = \text{Tr}(\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki})$. Also recall the notation $p_{i,j}$ for elements of I_i. We note that (a,b,c) is an (i,jk) tuple if and only if (iff) $a = p_{i,u}$, $b = p_{j,v}$, $c = p_{k,w}$ and both $\boldsymbol{X}_{ki}(w,u) = 1$ and $\boldsymbol{X}_{ij}(u,v) = 1$ (so $p_{i,u}$ matches to both $p_{j,v}$ and $p_{k,w}$). We further note that the latter two requirements are equivalent with $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}(w,v) = 1$. Indeed, since $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}(w,v) = \sum_{u' \in [m_i]} \boldsymbol{X}_{ki}(w,u')\boldsymbol{X}_{ij}(u',v)$ and $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}$ is a partial permutation, then $\boldsymbol{X}_{ki}\boldsymbol{X}_{ij}(w,v) = 1$ iff $\boldsymbol{X}_{ki}(w,u) = \boldsymbol{X}_{ij}(u,v) = 1$ for some $u \in [m_i]$. Therefore

$$|A_{i,jk}| = n_i.$$

By the same way we note that

$$|A_{j,ki}| = n_j \text{ and } |A_{k,ij}| = n_k.$$

Similarly, note that (a,b,c) is a good (i,jk) tuple iff $a = p_{i,u}$, $b = p_{j,v}$, $c = p_{k,w}$, $\boldsymbol{X}_{ki}(w,u) = 1$, $\boldsymbol{X}_{ij}(u,v) = 1$ and $\boldsymbol{X}_{jk}(v,w) = 1$. The latter three requirements are equivalent with $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}(u,u) = 1$. Indeed, the following equation

$$\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}(u,u) = \sum_{v' \in [m_j], w' \in [m_k]} \boldsymbol{X}_{ij}(u,v')\boldsymbol{X}_{jk}(v',w')\boldsymbol{X}_{ki}(w',u)$$

and the fact that $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}$ is a partial permutation imply this equivalence. Therefore,

$$|A^{\text{good}}_{i,jk}| = n_\Delta \text{ and } |A^{\text{bad}}_{i,jk}| = n_i - n_\Delta.$$

Similarly, we conclude that

$$|A^{\text{bad}}_{j,ki}| = n_j - n_\Delta \text{ and } |A^{\text{bad}}_{k,ij}| = n_k - n_\Delta.$$

Using these observations and (4), one can rewrite d_{ijk} as follows

$$d_{ijk} = \frac{|A^{\text{bad}}_{i,jk}| + |A^{\text{bad}}_{j,ki}| + |A^{\text{bad}}_{k,ij}|}{|A^{\text{bad}}_{i,jk}| + |A^{\text{bad}}_{j,ki}| + |A^{\text{bad}}_{k,ij}| + 3n_\Delta}. \quad (12)$$

Let us assume that $(a,b,c) \in A^{\text{bad}}_{i,jk}$ and show that

$$h(c) \in U^{\text{bad}}_{ik} \cup U^{\text{bad}}_{jk}. \quad (13)$$

The assumption $ij \in E_g$ implies that there is a good match between a and b.

We claim that if there is also a good match between a and c, then $h(b) = h(a) = h(c) \in U_{ij} \cap U_{ik} \cap U_{jk}$. Indeed, assume $a = p_{i,u}$, $b = p_{j,v}$, $c = p_{k,w}$ and $h(a) = l$, i.e., $\boldsymbol{P}^*_i(u,l) = 1$. Because there exists a good match between a and b, $\boldsymbol{X}^*_{ij}(u,v) = \boldsymbol{X}_{ij}(u,v) = 1$. Since $\boldsymbol{X}^*_{ij}(u,v) = \boldsymbol{P}^*_i\boldsymbol{P}^{*T}_j(u,v) = \sum_{d \in [m]} \boldsymbol{P}^*_i(u,d)\boldsymbol{P}^*_j(v,d)$ and $\boldsymbol{P}^*_j$ is a partial permutation, $\boldsymbol{P}^*_j(v,l) = 1$ and thus $h(b) = l$. Similarly, since there exists a good match between a and c, $\boldsymbol{P}^*_k(w,l) = 1$ and $h(c) = l$. Therefore $h(a) = h(b) = h(c) = l \in U_{ij} \cap U_{ik} \cap U_{jk}$.

Since $(a,b,c) \in A^{\text{bad}}_{i,jk}$, there is no match between b and c and thus $h(c) \in U^{\text{bad}}_{jk}$, which implies (13).

If on the other hand, there is a bad match between a and c, then $h(c) \in U^{\text{bad}}_{ik}$, which also implies (13).

In view of (13), the function $f(a,b,c) = h(c)$ maps $A^{\text{bad}}_{i,jk}$ to $U^{\text{bad}}_{ik} \cup U^{\text{bad}}_{jk}$. We note that this function is injective. Indeed, since $\boldsymbol{X}_{ik}$ and $\boldsymbol{X}_{jk}$ are partial permutations, for any $h(c) \in U^{\text{bad}}_{ik} \cup U^{\text{bad}}_{jk}$, where $c \in I_k$, if there are a and b in I_i and I_j, respectively, such that there are matches between them and c and $f(a,b,c) = h(c)$, then a and b are unique. Figure 5 demonstrates f and its proven injectivity in a special case. In this case, $A^{\text{bad}}_{i,jk} = \{(p_{i,4}, p_{j,2}, p_{k,3}), (p_{i,2}, p_{j,1}, p_{k,1})\}$, $f(p_{i,4}, p_{j,2}, p_{k,3}) = h(p_{k,3}) = 5 \in U^{\text{bad}}_{ik}$ and $f(p_{i,2}, p_{j,1}, p_{k,1}) = h(p_{k,1}) = 2 \in U^{\text{bad}}_{jk}$.

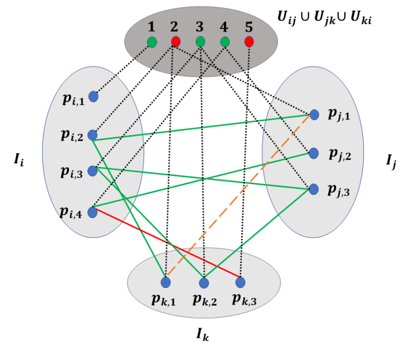

Figure 5. A demonstration for clarifying the definition of (i,jk) tuples, good (i,jk) tuples, bad (i,jk) tuples as well as the function f. The actual use of the figure is clarified when it is referred to. Unlike other figures (that focus on points in U_{ij} and not in $U_{ij} \cup U_{jk} \cup U_{ki}$), the red dots correspond to keypoints in $U^{\text{bad}}_{ij} \cup U^{\text{bad}}_{jk} \cup U^{\text{bad}}_{ik}$ and the green dots correspond to the rest of keypoints in $U_{ij} \cup U_{jk} \cup U_{ki}$. We note that $U^{\text{bad}}_{ik} = \{5\}$, $U^{\text{bad}}_{jk} = \{2\}$ and $U^{\text{bad}}_{ij} = \varnothing$.

By the cardinality property of an injective map,

$$|A^{\text{bad}}_{i,jk}| \le |U^{\text{bad}}_{ik} \cup U^{\text{bad}}_{jk}| \le |U^{\text{bad}}_{ik}| + |U^{\text{bad}}_{jk}| = m^{\text{bad}}_{ik} + m^{\text{bad}}_{jk}.$$

Similarly, the same bound holds for $|A^{\text{bad}}_{j,ki}|$ and $|A^{\text{bad}}_{k,ij}|$ and consequently,

$$|A^{\text{bad}}_{i,jk}| + |A^{\text{bad}}_{j,ki}| + |A^{\text{bad}}_{k,ij}| \le 3(m^{\text{bad}}_{ik} + m^{\text{bad}}_{jk}). \quad (14)$$

Therefore the combination of (11) and (14) with the definition of s^*_{ik} as well as s^*_{jk} (i.e., noting that $s^*_{ik} = m^{\text{bad}}_{ik}/m_{ik} = 1 - m^{\text{good}}_{ik}/m_{ik}$ and $s^*_{jk} = m^{\text{bad}}_{jk}/m_{jk} = 1 - m^{\text{good}}_{jk}/m_{jk}$) yields

$$\begin{aligned} d_{ijk} &= \frac{|A^{\text{bad}}_{i,jk}| + |A^{\text{bad}}_{j,ki}| + |A^{\text{bad}}_{k,ij}|}{|A^{\text{bad}}_{i,jk}| + |A^{\text{bad}}_{j,ki}| + |A^{\text{bad}}_{k,ij}| + 3n_\Delta} \\ &\le \frac{3(m^{\text{bad}}_{ik} + m^{\text{bad}}_{jk})}{3(m^{\text{bad}}_{ik} + m^{\text{bad}}_{jk}) + 3} \\ &\le \frac{3m^{\text{bad}}_{ik}}{3m^{\text{bad}}_{ik} + 3} + \frac{3m^{\text{bad}}_{jk}}{3m^{\text{bad}}_{jk} + 3} \\ &= \frac{3s^*_{ik}}{3s^*_{ik} + \frac{3}{m_{ik}}} + \frac{3s^*_{jk}}{3s^*_{jk} + \frac{3}{m_{jk}}} \\ &\le m_{ik}s^*_{ik} + m_{jk}s^*_{jk} \\ &\le m(s^*_{ik} + s^*_{jk}). \end{aligned}$$

□

C.2. Lower Bound for the Averaged Cycle Inconsistency Among Good Cycles

This section includes the proof of the following lemma:

Lemma 2. If $G=(V,E)$ is p_v-cycle verifiable, then

$$\frac{1}{3}p_v s_{ij}^* \leq \frac{1}{|G_{ij}|}\sum_{k\in G_{ij}} d_{ijk} \quad \forall ij\in E. \tag{15}$$

Proof. We assume several cases.

Case I: $ij\in E_g$. The left hand side of (15) is zero and its right hand side (RHS) is also zero since for any $ij\in E_g$ and $k\in G_{ij}$, $d_{ijk}=0$.

Case II: $ij\in E_b$ and $k\in G_{ij}$. Denote

$$N_{ijk}^{\text{bad}}=|A_{i,jk}^{\text{bad}}|+|A_{j,ki}^{\text{bad}}|+|A_{k,ij}^{\text{bad}}|,$$

and note that in view of (12)

$$\sum_{k\in G_{ij}} d_{ijk}=\sum_{k\in G_{ij}}\frac{N_{ijk}^{\text{bad}}}{N_{ijk}^{\text{bad}}+3n_\Delta}. \tag{16}$$

We thus need to lower bound the RHS of (16) in order to conclude (15).

We first derive the bound $n_\Delta \leq m_{ij}^{\text{good}}$. Figure 6 demonstrates the definitions below and the desired bound in a very special case. Let D_{ijk} denote the set of indices of diagonal entries of $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}$ that equal 1. Note that $|D_{ijk}|=n_\Delta$ due to the fact that $n_\Delta=\text{Tr}(\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki})$. Also, since $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}$ is of size $m_i\times m_i$, $D_{ijk}\subseteq[m_i]$. We can thus assign for $d\in D_{ijk}$, $p_{i,d}=a\in I_i$. Because $\boldsymbol{X}_{ij}\boldsymbol{X}_{jk}\boldsymbol{X}_{ki}(d,d)=1$, there exists $1\leq u\leq m_j$ and $1\leq v\leq m_k$ such that $\boldsymbol{X}_{ij}(d,u)=\boldsymbol{X}_{jk}(u,v)=\boldsymbol{X}_{ki}(v,d)=1$. Therefore, we note that for $b:=p_{j,u}\in I_j$ and $c:=p_{k,v}\in I_k$, there exist matches between a and b, b and c, as well as c and a. Since $k\in G_{ij}$, $jk\in E_g$ and $ki\in E_g$ and thus $h(b)=h(c)=h(a)$ (see the same argument in the paragraph below (13), where it is enough to just assume that either $jk\in E_g$ or $ki\in E_g$); we denote the latter common value by l. Therefore $\boldsymbol{P}_i^*(a,l)=\boldsymbol{P}_j^*(b,l)=\boldsymbol{X}_{ij}(a,b)=1$. By definition of U_{ij}^{good}, $l\in U_{ij}^{\text{good}}$. Let f_{ijk} be a function from D_{ijk} to U_{ij}^{good} such that $f_{ijk}(d)=l$. We note that it is injective since for any $d\neq d'\in D_{ijk}$, $p_{i,d}\neq p_{i,d'}$, therefore $h(p_{i,d})\neq h(p_{i,d'})$. By the cardinality property of an injective map, $n_\Delta\leq|U_{ij}^{\text{good}}|=m_{ij}^{\text{good}}$.

Next, We prove an upper bound of N_{ijk}^{bad}. We assume without loss of generality that $(a,b,c)\in A_{i,jk}^{\text{bad}}$. Since $ik,jk\in E_g$, the matches from a to c and from b to c are correct. Therefore, the match from a to b is wrong and $h(a)\in U_{ij}^{\text{bad}}$. Denote by $g_i: A_{i,jk}^{\text{bad}}\rightarrow U_{ij}^{\text{bad}}$ the function which maps $(a,b,c)\in A_{i,jk}^{\text{bad}}$ to $h(a)\in U_{ij}^{\text{bad}}$. Figure 7 illustrates g_i in a special case. This function is injective since for any $x\in U_{ij}^{\text{bad}}$, $g_i^{-1}(x)$ contains at most one element $(a,b,c)\in A_{i,jk}^{\text{bad}}$. Indeed, if $g_i(a,b,c)=x$, then there must exist $a\in I_i$ such that $h(a)=x$, c in I_k such that there is a match between a and c, and b in I_j such that there is

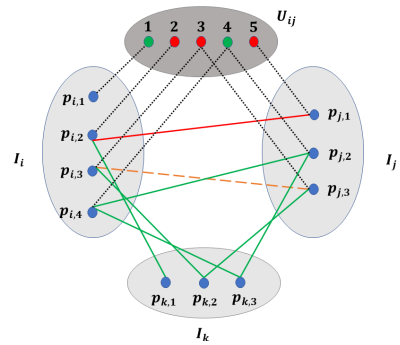

Figure 6. An illustration of n_Δ, D_{ijk}, f_{ijk} and m_{ij}^{good}. Note that n_Δ is equal to the number of green triangles with vertices in I_i,I_j,I_k. Since there is only one such triangle, $n_\Delta=1$. This triangle (with keypoints $p_{i,4},p_{j,2},p_{k,3}$) is associated with the keypoint in I_i with index 4 and thus $D_{ijk}=\{4\}$. Since it is also associated with the universal keypoint with index 4, the function f_{ijk} maps 4 to $4\in U_{ij}^{\text{good}}$. At last, note that $m_{ij}^{\text{good}}=|U_{ij}^{\text{good}}|=2$.

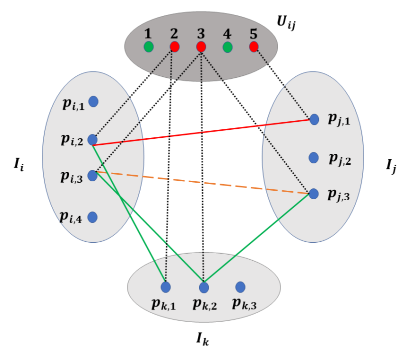

Figure 7. An illustration of g_i, $g_{ij,k}$ and the injectivity of both functions in a very special case. In this example, the only element of $A_{i,jk}^{\text{bad}}$ is $(p_{i,2},p_{j,1},p_{k,1})$. Since we defined $g_i(a,b,c)=h(a)$, we obtain that $g_i(p_{i,2},p_{j,1},p_{k,1})=h(p_{i,2})=2\in U_{ij}^{\text{bad}}$. Note that $U_{ij}^{\text{bad}}\cap h(I_k)=\{2,3\}$. Recall that the function $g_{ij,k}$ maps $x\in U_{ij}^{\text{bad}}\cap h(I_k)$ to the bad (j,ki) tuple or bad (k,ij) tuple that involves $c\in I_k$ with $h(c)=x$. In this example, the function $g_{ij,k}$ maps $2\in U_{ij}$ to the bad (i,jk) tuple $(p_{i,2},p_{j,1},p_{k,1})$. It maps $3\in U_{ij}$ to the bad (k,ij) tuple $(p_{k,2},p_{i,3},p_{j,3})$.

a match between a and b (and no match between b and c). Note that there is a match between at most one keypoint in I_k and a and thus there is at most one such c. Similarly, there is at most one such b. Since there is at most one keypoint in I_i which corresponds to the 3D keypoint x, there is at most one such a. The injectivity of g_i implies $|A_{i,jk}^{\text{bad}}|\leq|U_{ij}^{\text{bad}}|=m_{ij}^{\text{bad}}$. Similarly, $|A_{j,ki}^{\text{bad}}|\leq m_{ij}^{\text{bad}}$ and $|A_{k,ij}^{\text{bad}}|\leq m_{ij}^{\text{bad}}$. Thus, for any $k\in G_{ij}$

$$0\leq N_{ijk}^{\text{bad}}\leq 3m_{ij}^{\text{bad}}. \tag{17}$$

Next, we establish a lower bound of N_{ijk}^{bad}. For this purpose, we construct an injective map $g_{ij,k}$ from $U_{ij}^{\text{bad}}\cap h(I_k)$ to $A_{i,jk}^{\text{bad}}\cup A_{j,ki}^{\text{bad}}\cup A_{k,ij}^{\text{bad}}$. It will allow us to lower bound $N_{ijk}^{\text{bad}}=|A_{i,jk}^{\text{bad}}|+|A_{j,ki}^{\text{bad}}|+|A_{k,ij}^{\text{bad}}|$ by the cardinality of

$U_{ij}^{\text{bad}} \cap h(I_k)$. Note that $U_{ij}^{\text{bad}} \subseteq U_{ij} = h(I_i) \cup h(I_j)$. Therefore any element of $U_{ij}^{\text{bad}} \cap h(I_k)$ is either in $h(I_i)$ or $h(I_j) \backslash h(I_i)$.

In the case where $x \in U_{ij}^{\text{bad}} \cap h(I_k)$ and $x \in h(I_i)$, we will show that there exist either $(a,b,c) \in A_{i,jk}^{\text{bad}}$ or $(c,a,b) \in A_{k,ij}^{\text{bad}}$ such that $h(c) = x$. In the case where $x \in U_{ij}^{\text{bad}} \cap h(I_k)$ and $x \in h(I_j) \backslash h(I_i)$, then one can similarly show that there exists either $(b,c,a) \in A_{j,ki}^{\text{bad}}$ or $(c,a,b) \in A_{k,ij}^{\text{bad}}$ such that $h(c) = x$. These arguments induce a map $g_{ij,k}$ from $U_{ij}^{\text{bad}} \cap h(I_k)$ to $A_{i,jk}^{\text{bad}} \cup A_{j,ki}^{\text{bad}} \cup A_{k,ij}^{\text{bad}}$ which maps x to its corresponding bad tuple. Since $h(c) = x$, $g_{ij,k}$ is injective. Figure 7 illustrates $g_{ij,k}$ in a special case.

We thus assume that $x \in U_{ij}^{\text{bad}} \cap h(I_k)$ and $x \in h(I_i)$. Note that the latter requirement implies the existence of $a \in I_i$ such that $h(a) = x$. Since $x \in h(I_k)$, there exists $c \in I_k$ such that $h(c) = x$ and since $ik \in E_g$ there is a good match between a and c. Note that there cannot be a good match between a and any $b \in I_j$, otherwise $x \notin U_{ij}^{\text{bad}}$. Therefore, there are two cases to consider. In the first case there exists $b \in I_j$ such that there is a wrong match between a and b. This implies that $h(b) \neq h(a)$ and since we showed above that $h(a) = h(c)$, we conclude that $h(b) \neq h(c)$. The latter observation and the fact that $jk \in E_g$ imply that there is no match between b and c and thus (a,b,c) is a bad (i,jk) tuple, that is, $(a,b,c) \in A_{i,jk}^{\text{bad}}$. In the second case, there exists $b \in I_j$ such that $h(b) = h(a)$, but there is no match between b and a (the previous case considered the scenario where there exists $b \in I_j$ such that a and b match; furthermore, if $h(a) \neq h(b)$ for all $b \in I_j$, then $x \in U_{ij}^{\text{good}}$). Since $h(a) = h(c) = h(b)$ and $jk \in E_g$, there is a match between b and c. Therefore, (c,a,b) is a bad (k,ij) tuple, that is, $(c,a,b) \in A_{k,ij}^{\text{bad}}$. Following the above ideas, this concludes the injectivity of $g_{ij,k}$. This injectivity implies

$$\sum_{x \in U_{ij}^{\text{bad}}} 1_{\{x \in h(I_k)\}} = |U_{ij}^{\text{bad}} \cap h(I_k)| \leq |A_{i,jk}^{\text{bad}} \cup A_{j,ki}^{\text{bad}} \cup A_{k,ij}^{\text{bad}}| \leq |A_{i,jk}^{\text{bad}}| + |A_{j,ki}^{\text{bad}}| + |A_{k,ij}^{\text{bad}}| = N_{ijk}^{\text{bad}}. \quad (18)$$

In order to apply (18) we lower bound a certain sum of $1_{\{x \in h(I_k)\}}$. Our argument assumes that $x \in U_{ij}^{\text{bad}}$. Since $x \in U_{ij}$, we conclude WLOG that $x \in h(I_i)$. Therefore there exists $a \in I_i$ such that $h(a) = x$. By the p_v-cycle verifiability condition, a is verifiable w.r.t. ij in at least $p_v|G_{ij}|$ good cycles. For any such cycle ijk that a is verifiable in, let $b \in I_k$ match a (for convenience, we demonstrate x, a and b in Figure 8). Since $ik \in E_g$, the match between a and b is a good match and thus $x = h(a) = h(b)$. Since $b \in I_k$, $h(b) \in h(I_k)$ and thus $x \in h(I_k)$. That is, we have proved that if $k \in G_{ij}$ and a is verifiable in ijk, then $x \in h(I_k)$. We have at least $p_v|G_{ij}|$ such k's and thus

$$\sum_{k \in G_{ij}} 1_{\{x \in h(I_k)\}} \geq p_v|G_{ij}| \quad \text{for any } x \in U_{ij}. \quad (19)$$

We combine the above two inequalities as follows. Summing both sides of (18) over $k \in G_{ij}$, exchanging the order of summation and applying (19) result in

$$p_v|G_{ij}|m_{ij}^{\text{bad}} = p_v|G_{ij}||U_{ij}^{\text{bad}}| \leq \sum_{k \in G_{ij}} N_{ijk}^{\text{bad}}. \quad (20)$$

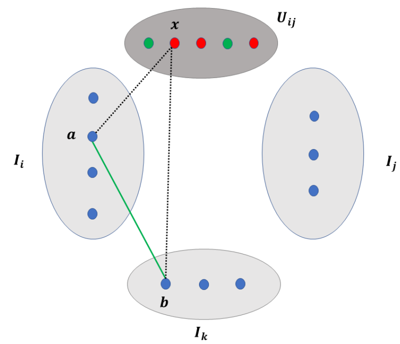

Figure 8. Visual demonstration of keypoints involved in the argument for bounding $\sum_{k \in G_{ij}} 1_{\{x \in h(I_k)\}}$.

Using the above bound we will bound from below

$$\sum_{k \in G_{ij}} \frac{N_{ijk}^{\text{bad}}}{N_{ijk}^{\text{bad}} + 3m_{ij}^{\text{good}}}$$

and we will then use (16) to conclude the desired inequality. We denote

$$F(x) = \frac{x}{x+\gamma} \quad \text{where } \gamma = 3m_{ij}^{\text{good}}. \quad (21)$$

Note that $F(0) = 0$,

$$F(3m_{ij}^{\text{bad}}) = \frac{3m_{ij}^{\text{bad}}}{3m_{ij}^{\text{bad}} + 3m_{ij}^{\text{good}}} = s_{ij}^*$$

and $F(x)$ is concave. Applying the definition of F, Jensen's inequality, (20) and (21) yield

$$\begin{aligned}
\sum_{k \in G_{ij}} \frac{N_{ijk}^{\text{bad}}}{N_{ijk}^{\text{bad}} + 3m_{ij}^{\text{good}}} &= \sum_{k \in G_{ij}} F(N_{ijk}^{\text{bad}}) \\
&\geq \sum_{k \in G_{ij}} \left(\left(1 - \frac{N_{ijk}^{\text{bad}}}{3m_{ij}^{\text{bad}}}\right) F(0) + \frac{N_{ijk}^{\text{bad}}}{3m_{ij}^{\text{bad}}} F(3m_{ij}^{\text{bad}}) \right) \\
&= \sum_{k \in G_{ij}} \frac{N_{ijk}^{\text{bad}}}{3m_{ij}^{\text{bad}}} F(3m_{ij}^{\text{bad}}) \\
&\geq \frac{1}{3} p_v |G_{ij}| F(3m_{ij}^{\text{bad}}) = \frac{1}{3} p_v |G_{ij}| s_{ij}^*.
\end{aligned}$$

The combination of this inequality with (16) concludes the proof of the lemma. □

C.3. Conclusion of Theorem 1

We prove the main theorem by induction, using Lemmas 1 and 2. For $t=0$, the definition of $s_{ij}^{(0)}$, Lemma 2 and the definition of λ imply that for all $ij\in E$:

$$s_{ij}^{(0)}=\frac{\sum_{k\in N_{ij}}d_{ijk}}{|N_{ij}|}\geq\frac{\sum_{k\in G_{ij}}d_{ijk}}{|N_{ij}|}\geq\frac{p_v}{3}\frac{|G_{ij}|}{|N_{ij}|}s_{ij}^*\geq\frac{p_v}{3}(1-\lambda)s_{ij}^*.$$

We further note by using again the above definitions and the fact that for all $ij\in E$ $0\leq d_{ijk}\leq 1$ that for $ij\in E_g$,

$$s_{ij}^{(0)}=\frac{\sum_{k\in N_{ij}}d_{ijk}}{|N_{ij}|}=\frac{\sum_{k\in B_{ij}}d_{ijk}}{|N_{ij}|}\leq\frac{\sum_{k\in B_{ij}}1}{|N_{ij}|}\leq\lambda\leq\frac{1}{2\beta_0}.$$

Therefore, the theorem is proved when $t=0$.

Next, we assume that the theorem holds for iterations $0,1,\cdots,t$ and show that it also holds for iteration $t+1$. Applying the definition of $s_{ij}^{(t+1)}$, the positivity of the terms in the sum, the induction assumption $\frac{1}{2\beta_t}\geq\max_{ij\in E_g}s_{ij}^{(t)}$, Lemma 2 and the definition of λ, we obtain for any $ij\in E_b$

$$\begin{aligned}s_{ij}^{(t+1)}&=\frac{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}d_{ijk}}{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}}\\&\geq\frac{\sum_{k\in G_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}d_{ijk}}{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}}\\&\geq\frac{\sum_{k\in G_{ij}}e^{-1}d_{ijk}}{|N_{ij}|}\\&\geq\frac{p_v}{3e}\frac{|G_{ij}|}{|N_{ij}|}s_{ij}^*\\&\geq\frac{(1-\lambda)p_v}{3e}s_{ij}^*.\end{aligned}\tag{22}$$

Note that $xe^{-\alpha x}\leq\frac{1}{\alpha e}$ for any $\alpha>0$ and $x\geq 0$. In particular, for $\alpha=\beta_t(1-\lambda)p_v/3e$ and $x=s_{ik}^*+s_{jk}^*$,

$$e^{-\beta_t(s_{ik}^*+s_{jk}^*)(1-\lambda)p_v/3e}(s_{ik}^*+s_{jk}^*)\leq\frac{3}{\beta_t(1-\lambda)p_v}.\tag{23}$$

Applying the definition of $s_{ij}^{(t+1)}$, the fact that $d_{ijk}=0$ for any $ij\in E_g$ and $k\in G_{ij}$, Lemma 1, the induction assumption $s_{ij}^{(t)}\geq\frac{(1-\lambda)p_v}{3e}s_{ij}^*$ for $ij\in E_b$ (for the numerator) and the positivity of the relevant terms (for the denominator), the induction assumption $s_{ij}^{(t)}\leq 1/(2\beta_t)$ for all $ij\in E_g$, (23) and the definition of λ, we obtain for all $ij\in E_g$

$$\begin{aligned}s_{ij}^{(t+1)}&=\frac{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}d_{ijk}}{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}}\\&=\frac{\sum_{k\in B_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}d_{ijk}}{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}}\\&\leq\frac{\sum_{k\in B_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}m(s_{ik}^*+s_{jk}^*)}{\sum_{k\in N_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}}\\&\leq\frac{\sum_{k\in B_{ij}}e^{-\beta_t(s_{ik}^*+s_{jk}^*)\frac{(1-\lambda)p_v}{3e}}m(s_{ik}^*+s_{jk}^*)}{\sum_{k\in G_{ij}}e^{-\beta_t(s_{ik}^{(t)}+s_{jk}^{(t)})}}\\&\leq\frac{m\sum_{k\in B_{ij}}e^{-\beta_t(s_{ik}^*+s_{jk}^*)\frac{(1-\lambda)p_v}{3e}}(s_{ik}^*+s_{jk}^*)}{|G_{ij}|e^{-1}}\\&\leq\frac{m\sum_{k\in B_{ij}}\frac{3}{\beta_t(1-\lambda)p_v}}{|G_{ij}|e^{-1}}\\&=\frac{3m|B_{ij}|}{|G_{ij}|e^{-1}}\cdot\frac{1}{\beta_t(1-\lambda)p_v}\\&\leq\frac{6em\lambda}{(1-\lambda)^2p_v}\cdot\frac{1}{2\beta_t}.\end{aligned}$$

We note that the assumption $\lambda<1+\frac{3em}{p_v}-\sqrt{\frac{3em}{p_v}(2+\frac{3em}{p_v})}$ is equivalent with $\frac{6em\lambda}{(1-\lambda)^2p_v}<1$. Therefore by taking $\beta_{t+1}=r\beta_t$ with $1<r<\frac{(1-\lambda)^2p_v}{6em\lambda}$, we guarantee that for any $ij\in E_g$, $s_{ij}^{(t+1)}\leq\frac{1}{2\beta_{t+1}}$, that is, $\max_{ij\in E_g}s_{ij}^{(t+1)}\leq\frac{1}{2\beta_{t+1}}=\frac{1}{2\beta_0r^t}$. This implication and (22) conclude the proof of the theorem.

D. Discussion of a Possible Theoretical Extension

Although our current analysis assumes no noise on the set of good edges, one can relax this assumption. Indeed, one can assume sufficiently small noise on good edges so that for all cycles ijk and a sufficiently small positive constant δ: $|d'_{ijk}-d_{ijk}|<\delta$, where d_{ijk} and d'_{ijk} are respectively the cycle inconsistencies with and without noise on good edges. Using a basic perturbation analysis, similarly as in the proof of Theorem 1, with a carefully chosen set of the reweighting parameters β_t, one can prove approximate separation of good and bad edges. In particular, the maximum value of the estimated s_{ij} on good edges is proportional to δ. Removing the bad edges (with estimated s_{ij} larger than this threshold), one can then approximately solve the PPS problem with a subsequent spectral solver. An approximate recovery theorem for the absolute partial permutations using the filtered edges can be established using spectral graph theory.