A Robust Translation Synchronization Algorithm

Zihang He*
UT Austin
Austin, TX 78712
zhhe@utexas.edu

Hang Ruan*
Shanghai Jiaotong University
Shanghai, 200240
ruanhang1123@gmail.com

Qixing Huang
UT Austin
Austin, TX 78712
huangqx@utexas.edu

Abstract

This paper introduces a robust translation synchronization approach which takes relative directions between pairs of images as inputs and outputs absolute image locations. Our approach is based on a generalized eigenvalue formulation, which contains edge weights in relative directions and vertex weights in absolute image translations. We present a rigorous stability analysis to determine how to set these weights optimally. Specifically, optimal vertex weights are always identical, whereas optimal edge weights depend on the magnitudes of relative translations and variances of relative directions. These results lead to an iterative synchronization procedure, which progressively removes outliers in the inputs by adaptively adjusting the edge weights. Experimental results justify our theoretical results and show that our approach outperforms state-of-the-art baseline approaches on both synthetic and real datasets.

1. Introduction

This paper studies the classical pose synchronization problem of computing the absolute 6D poses of a collection of images from pairwise relative 6D poses along a graph of these images. A 6D pose has a rotation component and a translation component. Pose synchronization typically proceeds in two phases. The first phase performs rotation synchronization, which determines the absolute rotation of each image. The resulting rotations transform the relative translations expressed in local coordinate frames into relative translations in a global coordinate system. The second phase then determines the absolute translations of the input images.

As we cannot predict the absolute scale of the underlying objects, translation synchronization is much more challenging than rotation synchronization. Specifically, the relative translation associated with each pair of images is a direction that offers only two constraints. Therefore, even though a connected graph of image pairs is sufficient for rotation synchronization, translation synchronization needs more image pairs. Computationally, the fundamental challenge lies in jointly estimating the scales of relative translations and absolute translations. This challenge also makes it difficult to obtain closed-form solutions for translation synchronization.

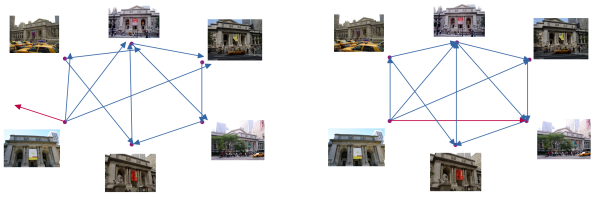

(a) Input graph of relative directions. Purple points are the camera locations. The blue arrows are the input relative translations. The red arrow is an outlier relative prediction.

(b) Recovered global absolute translations. Purple points are the camera locations. Blue arrows are the predicted relative translations

Figure 1. (Left) Input graph of relative directions, in which the color of each edge indicates the directional error. (Right) Recovered global absolute translations, overlaid with the ground-truth absolute translations.

In this paper, we introduce a novel translation synchronization that admits a closed-form solution. Our approach solves a generalized eigenvalue problem. The formulation has a set of edge weights associated with the relative directions and a set of vertex weights associated with the input images. We present a rigorous stability analysis of this generalized eigenvalue formulation. We show that the optimal vertex weights are identical, whereas the optimal edge weights depend on the variances and magnitudes of relative translations. Based on these results, we then develop an iterative translation synchronization formulation that adaptively adjusts the edge weights, in which estimations of variances and magnitudes of relative translations are derived from current predicted absolute translations. We present theoretical results on the robust recovery conditions of our approach,

showing the robustness of our approach against corrupted measurements.

We have evaluated our approach on both synthetic and real datasets. Experimental results on synthetic datasets justify our theoretical results. Experimental results on both synthetic and real datasets show that our approach outperforms state-of-the-art approaches. We also present an advanced analysis to show the effectiveness of our iterative scheme.

2. Related Work

We discuss related work in three categories, namely, map synchronization, translation synchronization, and spectral techniques.

2.1. Pose synchronization

Map synchronization, which generalizes pose synchronization, estimates consistent maps among a collection of objects from maps estimated between pairs of objects in isolation. Maps can take different forms, including transformations (the case of pose synchronization) [5, 28, 31], point-based maps [12, 17], and functional maps [13]. Map synchronization typically utilizes the generic constraint that composite maps along cycles should be equal to identify maps. In the case of point-based maps, Huang and Guibas [11] show the equivalence between the cycle-consistency constraint and the positive semidefiniteness of the matrix that encodes pairwise maps in blocks. In fact, many state-of-the-art map synchronization approaches solve low-rank matrix recovery problems to recover consistent maps. Optimization strategies include convex optimization [8, 11, 13, 24, 28, 31, 33], non-convex optimization [5, 7, 14, 18, 19, 34], and spectral techniques [1, 2, 4, 12, 16, 17, 22, 23, 26, 29, 30]. Our approach introduces a novel matrix encoding of the directional constraints for translation synchronization and falls into the categories of non-convex optimization and spectral techniques.

2.2. Translation synchronization

Translation synchronization requires optimizing both the scales of the directional constraints and the absolute translations. Because of this, few existing approaches admit closed-form solutions. Govindu [9] utilized the cross product and constructed a linear system of to estimate translations. Özyesil et al. [21] proposed iteratively reweighted least squares algorithm for the least squared deviations (LUD) solver in L^1 the norm of the relaxed displacement cost function. Hand et al. [10] proposed the ShapeFit algorithm by optimizing a convex program in the translation direction. Wilson proposed 1DSfm [32] to remove outliers by projecting from 3d to 1d and optimize the direction-based cost function on the remaining accurate matches. Zhuang et al. [35] analyzed LUD and ShapeFit and proposed the bilinear angle-based translation averaging algorithm (BATA) that optimizes a relaxed direction cost function. In particular, BATA shows that the loss based on the angle between the input direction and the optimized direction is superior to standard distance norms between directions. Our approach also uses angles to define the loss term but introduces a novel spectral formulation that admits a closed-form solution.

Combining displacement and direction cost functions, Lalit et al. [20] alternatively update translations by direction and displacement-based methods. In contrast, our approach is based on a spectral formulation that admits a closed-form solution. We analyze its stability to derive an iterative approach to robust translation synchronization.

2.3. Spectral techniques

Our approach is motivated by the stability of eigenvalues and eigenvectors of a symmetric matrix under perturbation. We refer to a survey article on this topic [6]. Compared to convex and non-convex optimization formulations, spectral techniques are simple and efficient. Unlike robust low-rank matrix recovery, in which outliers may be unbounded, outliers in pairwise maps have bounded norms. This fact improves the robustness of applying spectral techniques in map synchronization [3, 15, 22, 23, 25, 27]. Unlike most prior work that studies robust recovery under random and independent noise, this paper presents a robustness result under adversarial noise.

3. Problem Statement and Approach Overview

3.1. Problem Statement

Consider a connected graph $\mathcal{G} = (\mathcal{I}, \mathcal{E})$ of $n = |\mathcal{I}|$ image objects. Each edge $(i, j) \in \mathcal{E}$ is associated with a measurement $\boldsymbol{v}_{ij}^{\text{inp}}$ ($\|\boldsymbol{v}_{ij}^{\text{inp}}\| = 1$) of the normalized translation from I_i to I_j. Specifically, denote $\boldsymbol{t}_i^{gt}$ as the underlying ground-truth translation of I_i, we have

$$\boldsymbol{v}_{ij}^{\text{inp}} \approx \frac{\boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}}{\|\boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}\|}. \tag{1}$$

In particular, we expect that for a subset of the edges $\{(i, j)\}$, $\boldsymbol{v}_{ij}^{\text{inp}}$ have large errors.

Translation synchronization aims at recovering $\boldsymbol{t}_i$ from $\boldsymbol{v}_{ij}^{\text{inp}}$ up to a global scale and a global translation.

3.2. Approach Overview

A fundamental challenge of translation synchronization is that (1) is non-linear. Our approach addresses this challenge using a generalized spectral formulation by introducing weights associated with edges $\mathcal{E}$ and vertices $\mathcal{I}$ (Section 4.1). When there is no input noise, this generalized formulation allows us to derive simple conditions on which the translation synchronization problem has a unique solution (Section 4.2). When the input has noise, the accuracy of the output depends on these weights. We present a rigorous stability analysis on the relations between output errors and these weights (Section 4.3). The results lead to an iterative approach that adaptively adjusts these weights to progressively eliminate outliers in the inputs and average inliers in the inputs (Section 4.4). Finally, we present a robust recovery condition for our approach under adversarial input noise (Section 4.5).

3.3. Notations

For each $(i,j) \in \mathcal{E}$, we define $\boldsymbol{e}_{ij} \in \mathbb{R}^n$ as the edge indicator function whose element $e_{ij}(k)$ satisfies $e_{ij}(i) = 1$, $e_{ij}(j) = -1$, and $e_{ij}(k) = 0$ otherwise.

We will also use several matrix norms. Given a block matrix $A \in \mathbb{R}^{3m\times 3n}$ and a block vector $\boldsymbol{v} \in \mathbb{R}^{3m}$, we define the L^1 norm of A and the L^∞ norm of $\boldsymbol{v}$ as

$$\|A\|_1 = \max_{1\le i\le m} \sum_{j=1}^{n} \|A_{ij}\|, \quad \|\boldsymbol{v}\|_\infty = \max_{1\le i\le m} \|\boldsymbol{v}_i\|. \tag{2}$$

Moreover, $\|A\|$ denotes the spectral norm of matrix A. $A^\dagger$ denotes the pseudo-inverse of A.

4. Approach

This section presents our technical approach.

4.1. Spectral Translation Synchronization

Let s_{ij} be the latent absolute scale parameter of edge $(i,j) \in \mathcal{E}$. We can write out the relation between $\boldsymbol{v}_{ij}$ and $\boldsymbol{t}_i$ and $\boldsymbol{t}_j$ as

$$s_{ij}\boldsymbol{v}_{ij}^{\text{inp}} \approx (\boldsymbol{t}_i - \boldsymbol{t}_j). \tag{3}$$

When $\boldsymbol{t}_i$ and $\boldsymbol{t}_j$ are fixed, the optimal s_{ij} is given by

$$\begin{aligned} s_{ij} &= \operatorname*{argmin}_s \quad \|s\boldsymbol{v}_{ij}^{\text{inp}} - (\boldsymbol{t}_i - \boldsymbol{t}_j)\|^2 \\ &= \frac{\boldsymbol{v}_{ij}^{\text{inp}^T}(\boldsymbol{t}_i - \boldsymbol{t}_j)}{\boldsymbol{v}_{ij}^{\text{inp}^T}\boldsymbol{v}_{ij}^{\text{inp}}} = \boldsymbol{v}_{ij}^{\text{inp}^T}(\boldsymbol{t}_i - \boldsymbol{t}_j). \end{aligned} \tag{4}$$

Denote $\boldsymbol{t} = (\boldsymbol{t}_i) \in \mathbb{R}^{3n}$. Substituting Eq. (4) into Eq. (3) and introducing a weight $w_{ij} > 0$ for each edge $(i,j) \in \mathcal{E}$, we arrive at the following quadratic loss term among $\boldsymbol{t}$:

$$\begin{aligned} f(\boldsymbol{t}) &= \sum_{(i,j)\in\mathcal{E}} w_{ij}\|(I_3 - \boldsymbol{v}_{ij}^{\text{inp}}\boldsymbol{v}_{ij}^{\text{inp}^T})(\boldsymbol{t}_i - \boldsymbol{t}_j)\|^2 \\ &= \sum_{(i,j)\in\mathcal{E}} w_{ij}(\boldsymbol{t}_i - \boldsymbol{t}_j)^T(I_3 - \boldsymbol{v}_{ij}^{\text{inp}}\boldsymbol{v}_{ij}^{\text{inp}^T})(\boldsymbol{t}_i - \boldsymbol{t}_j). \end{aligned} \tag{5}$$

Remark 1. *Let θ_{ij} be the angle between $\boldsymbol{v}_{ij}^{\text{inp}}$ and $\frac{\boldsymbol{t}_i - \boldsymbol{t}_j}{\|\boldsymbol{t}_i - \boldsymbol{t}_j\|}$. It is easy to see that*

$$\begin{aligned} &(\boldsymbol{t}_i - \boldsymbol{t}_j)^T(I_3 - \boldsymbol{v}_{ij}^{\text{inp}}\boldsymbol{v}_{ij}^{\text{inp}^T})(\boldsymbol{t}_i - \boldsymbol{t}_j) \\ =& \sin(\theta_{ij})^2\|\boldsymbol{t}_i - \boldsymbol{t}_j\|^2. \end{aligned}$$

Therefore, our formulation can be viewed as a variant of the angle-based loss proposed in [35].

Minimizing Eq. (5) directly has a trivial solution, in which $\boldsymbol{t}_i = 0$. We apply the standard approach of enforcing a quadratic normalization constraint on $\boldsymbol{t}_i$:

$$\sum_{i=1}^{n} s_i\|\boldsymbol{t}_i\|^2 = 1 \tag{6}$$

where $s_i > 0$ is the weight associated with image I_i.

Introduce a generalized connection Laplacian matrix $L \in \mathbb{R}^{3n\times 3n}$ whose blocks are given by

$$L_{ij} = \begin{cases} \sum_{j\in\mathcal{N}_i} w_{ij}(I_3 - \boldsymbol{v}_{ij}^{\text{inp}}\boldsymbol{v}_{ij}^{\text{inp}^T}) & i = j \\ -w_{ij}(I_3 - \boldsymbol{v}_{ij}^{\text{inp}}\boldsymbol{v}_{ij}^{\text{inp}^T}) & (i,j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}$$

It is clear that we can rewrite

$$L = \sum_{(i,j)\in\mathcal{E}} w_{ij}E_{ij}(I_3 - \boldsymbol{v}_{ij}^{\text{inp}}\boldsymbol{v}_{ij}^{\text{inp}^T})E_{ij}^T, \qquad E_{ij} = \boldsymbol{e}_{ij}\otimes I_3.$$

Denote $S = \text{diag}(\{s_i\}) \otimes I_3$. We can then reformulate the optimization problem of minimizing Eq. (5) subject to the constraint Eq. (6) as

$$\min_{\boldsymbol{t}} \ \boldsymbol{t}^T L\boldsymbol{t} \qquad s.t. \quad \boldsymbol{t}^T S\boldsymbol{t} = 1 \tag{7}$$

It is easy to see that the optimal solution $\boldsymbol{t}^\star$ of Eq. (7) is given by the following generalized eigen-value problem:

$$L\boldsymbol{t}^\star = \lambda S\boldsymbol{t}^\star. \tag{8}$$

The following proposition describes some basic properties of L.

Proposition 1. *L is positive semidefinite. It has three zero eigenvalues. The corresponding eigenvectors are given by $\mathbf{1} \otimes I_3$.*

Proof: See Section B.1. □

Therefore, we look for $\boldsymbol{t}^\star$ as the eigenvector that corresponds to the fourth-smallest eigenvalue of Eq. (8) and satisfies $(\mathbf{1} \otimes I_3)^T \boldsymbol{t}^\star = 0$.

4.2. Uniqueness of Translation Synchronization

This section studies the uniqueness problem of translation synchronization, that is, under what conditions of $\boldsymbol{t}_i^{gt}$ and $\mathcal{E}$, one can recover $\boldsymbol{t}_i^{gt}$ from $\boldsymbol{v}_{ij}^{gt} = \frac{\boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}}{\|\boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}\|}, \forall (i,j) \in \mathcal{E}$. Although there are a lot of translation synchronization approaches, this uniqueness problem is underexplored. We show a simple necessary and sufficient uniqueness condition using Eq. (8).

Without losing generality, we assume that $\boldsymbol{t}^{gt}$ is normalized so that

$$\sum_{i=1}^{n} \boldsymbol{t}_i^{gt} = 0, \qquad \sum_{i=1}^{n} s_i \|\boldsymbol{t}_i^{gt}\|^2 = 1.$$

Proposition 2. *Suppose that the input directions are exact, that is, $\boldsymbol{v}_{ij} = \frac{\boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}}{\|\boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}\|}$, then the fourth eigenvalue of Eq. (8) $\lambda_4 = 0$. If and only if the fifth eigenvalue $\lambda_5 > 0$, $\boldsymbol{t}^{gt}$ is the unique fourth eigenvector of Eq. (8) and translation synchronization has a unique solution.*

Proof: See Section B.2. □

It is difficult to develop explicit conditions of $\boldsymbol{t}_i^{gt}$ on which $\lambda_5^{gt} > 0$. However, in many practical settings, $\boldsymbol{t}_i^{gt}$ usually lie close to a plane, e.g., images captured by a moving vehicle. Therefore, we study the properties of λ_5^{gt} when projecting $\boldsymbol{v}_{ij}^{gt}$ onto a plane.

First, we have the relation between the uniqueness of the 2D projection problem and the uniqueness of the original problem.

Proposition 3. *Let $\boldsymbol{n}$ be the normal to the plane of consideration. Denote $\boldsymbol{v}_{ij}^{gt,2D} = (I_3 - \boldsymbol{n}\boldsymbol{n}^T)\boldsymbol{v}_{ij}^{gt} / \|(I_3 - \boldsymbol{n}\boldsymbol{n}^T)\boldsymbol{v}_{ij}^{gt}\|$ as the projected direction along (i,j). Introduce L^{2D} as the connection Laplacian derived from $\boldsymbol{v}_{ij}^{gt,2D}$. With $\lambda_5^{gt,2D}$ we denote the fifth smallest eigenvalue of L^{2D}. We have, if $\lambda_5^{gt,2D} > 0$, then $\lambda_5^{gt} > 0$.*

Proof: See Section B.3. □

We can also understand why 2D uniqueness is more difficult than 3D uniqueness as follows. We have $3n-4$ variables in 3D and each edge offers two constraints. Therefore, we need $|\mathcal{E}| \geq \frac{3n-4}{2}$ to have a unique solution. We have $2n-3$ variables in 2D and each edge offers one constraint. Therefore, we need more edges, i.e., $|\mathcal{E}| \geq 2n-3$, to have a unique solution. This also means that if the $\mathcal{G}$ is a 2D planar graph, then we do not have unique solutions.

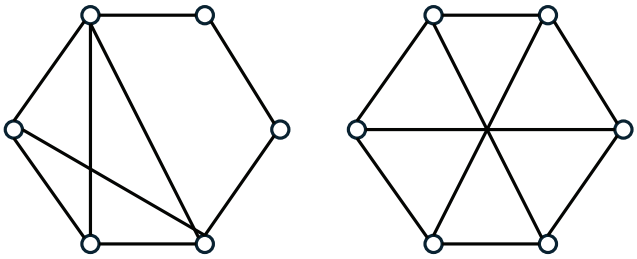

Figure 2. Connected graphs with six vertices. (Left) An example graph that does not have unique solutions. (Right) An example graph that has unique solutions.

In Section C of the supp. material, we present a necessary uniqueness condition on the topological structure of $\mathcal{E}$. It is also a sufficient condition in a probabilistic sense. Figure 2 shows examples of topological graphs that admit unique solutions and do not admit unique solutions.

4.3. Local Stability Analysis

When $\boldsymbol{v}_{ij}^{\text{inp}}$ are not exact, the fourth eigenvector $\boldsymbol{u}_4$ of generalized eigen-decomposition probelm in (8) will deviate from $\boldsymbol{t}^{gt}$. The deviations depend on w_{ij}, s_i, and the edge set $\mathcal{E}$. This section presents local stability results that provide insights into how to set w_{ij} and s_i to minimize the deviations of $\boldsymbol{u}_4$.

We begin by describing a noise model for $\boldsymbol{v}_{ij}^{\text{inp}}$. Denote $\boldsymbol{v}_{ij}^{gt}$ as the ground truth of $\boldsymbol{v}_{ij}^{\text{inp}}$. Let θ_{ij} be the angle between $\boldsymbol{v}_{ij}^{\text{inp}}$ and $\boldsymbol{v}_{ij}^{gt}$. We assume $\boldsymbol{v}_{ij}^{\text{inp}}$ is a random perturbation of $\boldsymbol{v}_{ij}^{gt}$ such that the perturbation is circular symmetric and $\mathbb{E}(\sin^2(\theta_{ij})) = \sigma_{ij}^2$.

Given $\{w_{ij}, s_i, \sigma_{ij}\}$, we define a stability score of $\boldsymbol{u}_4$ as

$$f(\{w_{ij}\},\{s_i\},\{\sigma_{ij}\}) = \mathop{\mathbb{E}}_{\{\boldsymbol{v}_{ij}^{\text{inp}}\}} \frac{\|\boldsymbol{u}_4 - \boldsymbol{t}^{gt}\|^2}{\|\boldsymbol{t}^{gt}\|^2}. \tag{9}$$

Our analysis focuses on the regime in which σ_{ij} is small, i.e., $\boldsymbol{v}_{ij}^{gt\,T}(\boldsymbol{v}_{ij}^{\text{inp}} - \boldsymbol{v}_{ij}^{gt}) \approx 0$. As we shall discuss next, we have simple and analytical results on the optimal w_{ij} and s_i in this regime. We will use these results to derive our algorithm, as we find them to be effective on arbitrary inputs.

The following theorem shows that when $\{\sigma_{ij}, w_{ij}\}$ are fixed, the optimal values of s_i to minimize $f(\{w_{ij}\},\{s_i\},\{\sigma_{ij}\})$ are identical.

Theorem 1. *Suppose* $\sum_{i=1}^{n} s_i = n$. *Then*

$$\{1\} = \underset{\{s_i\}}{\operatorname{argmin}} \quad f(\{w_{ij}\}, \{s_i\}, \{\sigma_{ij}\}).$$

Proof: See Section D.1 and Section D.2. □

The following theorem characterizes the optimal values of $\{w_{ij}\}$ given σ_{ij}.

Theorem 2. *Factoring out the global scale of* w_{ij}, *we have*

$$\{\frac{1}{\sigma_{ij}^2 \|\boldsymbol{t}_{ij}^{gt}\|^2}\} = \underset{\{w_{ij}\}}{\operatorname{argmin}} \quad f(\{w_{ij}\}, \{1\}, \{\sigma_{ij}\}).$$

where $\boldsymbol{t}_{ij}^{gt} = \boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}$.

Proof: See Section D.1 and Section D.3. □

From the analysis above, we can always set $s_i = 1$. The values of w_{ij} depend on the estimations of σ_{ij} and $\|\boldsymbol{t}_{ij}^{gt}\|$, which can be determined in an iterative manner. This leads to our iterative spectral translation synchronization algorithm, which is introduced next.

4.4. Iterative Spectral Translation Synchronization

Our algorithm consists of an initialization phase and an alternating phase of location optimization and weight update.

Initialization phase. We set $w_{ij} = 1, \forall (i,j) \in \mathcal{E}$ at iteration $k = 1$. On real datasets where graphs may contain weakly connected components, we find that taking the whole graph as the input can lead to degenerate solutions, i.e., $\|\boldsymbol{t}_i\| \to 1$ for some i. To address this issue, we iteratively prune vertices until $\min_i \|\boldsymbol{t}_i\| < \delta$ where δ is a user specified threshold. Translation synchronization is performed on the remaining sub-graph. Please refer to Section A in the supp. material for more details.

Alternating step II: location optimization. Given the current edge weights w_{ij} and vertex weights $s_i = 1$ at the current iteration k, we solve (8) to obtain an estimation of the image locations $\boldsymbol{t}_i$.

Alternating step III: weight update. Given the current image locations $\boldsymbol{t}_i, 1 \le i \le n$ at iteration k, we estimate

$$\sigma_{ij}^2 = \|\boldsymbol{v}_{ij}^{\text{inp}} - \frac{\boldsymbol{t}_{ij}}{\|\boldsymbol{t}_{ij}\|}\|^2$$

where $\boldsymbol{t}_{ij} = \boldsymbol{t}_i - \boldsymbol{t}_j$. Similarly, we predict $\boldsymbol{t}_{ij}^{gt}$ using $\boldsymbol{t}_{ij}$

Based on these two estimations, we set

$$w_{ij} = \frac{\sigma_k^2}{\sigma_k^2 + \|\sigma_{ij}^2\| \|\boldsymbol{t}_{ij}\|^2}$$

where σ_k is introduced to avoid overconfident estimations of σ_{ij} and $\boldsymbol{t}_{ij}^{gt}$. We set σ_k using the geometric decaying scheme, i.e.,

$$\sigma_k = \sigma_{\max}(\frac{\sigma_{\min}}{\sigma_{\max}})^{s_k}, \quad s_k = \frac{k-1}{k_{\max}-1}.$$

where $k_{\max}$ is the maximum number of iterations. In all of our experiments, we set $\sigma_{\max} = 1$, $\sigma_{\min} = 10^{-3}$, $k_{\max} = 30$. During optimization, we also truncate $w_{ij} = 0$ when $w_{ij} \le 0.01$. This operation removes outlier contributions and improves performance.

4.5. Robust Recovery Conditions

In this section, we present a robust recovery condition on our approach. We will use the ground-truth connection Laplacian matrix in which all input directions are exact.

$$L^{gt}(\boldsymbol{w}) = \sum_{(i,j)\in\mathcal{E}} w_{ij} E_{ij} (I_3 - \boldsymbol{v}_{ij}^{gt} \boldsymbol{v}_{ij}^{gt\,T}) E_{ij}^T. \tag{10}$$

where $\boldsymbol{w} = (w_{ij}) \in \mathbb{R}^{|\mathcal{E}|}$ collects all edge vectors. In the following, we first present a noise model. We then present the robust recovery result.

Noise model. Most previous work performs analysis by assuming that the noisy observations are independent. However, this assumption does not hold as the pairwise matches are computed from the input images, which exhibit correlations. In this paper, we present a worst-case analysis where the noisy observations can be adversarial.

Our noise model considers a mixture of inliners and outliers. Specifically, given a fixed observation graph $\mathcal{G} = (\mathcal{I}, \mathcal{E})$. The edge set $\mathcal{E} = \mathcal{E}^{\text{in}} \cup \mathcal{E}^{\text{out}}$ decomposes into an inlier set $\mathcal{E}^{\text{in}}$ and an outlier set $\mathcal{E}^{\text{out}}$. With $\mathcal{N}_i^{\text{in}} = \{j | (i,j) \in \mathcal{E}^{\text{in}}\}$ and $\mathcal{N}_i^{\text{out}} = \{j | (i.j) \in \mathcal{E}^{\text{out}}\}$ we denote the neighbors of i that form inlier and outlier edges.

For each $(i,j) \in \mathcal{E}_{\text{in}}$, we assume $\boldsymbol{v}_{ij}$ is an arbitrary unit length vector in the neighborhood of $\boldsymbol{v}_{ij}^{gt}$ that satisfies

$$\|(I_3 - \boldsymbol{v}_{ij}^{gt} \boldsymbol{v}_{ij}^{gt\,T}) \boldsymbol{v}_{ij}\| \le \epsilon \tag{11}$$

where ϵ is a sufficiently small constant. For each edge $(i,j) \in \mathcal{E}_{\text{out}}$, $\boldsymbol{v}_{ij}$ can be an arbitrary unit vector. Based on this noise model, we proceed to present the following robust recovery result.

Denote $\boldsymbol{t}_{ij}^{\star} = \boldsymbol{t}_i^{\star} - \boldsymbol{t}_j^{\star}$ and $\boldsymbol{t}_{ij}^{gt} = \boldsymbol{t}_i^{gt} - \boldsymbol{t}_j^{gt}$, where $\boldsymbol{t}_i^{\star}$ and $\boldsymbol{t}_i^{gt}$ as the output translation of our approach and the ground-truth translation of image i.

Theorem 3. *Consider three universal small constants*

$c_1, c_2, c_3 \geq 1$. *Denote*

$$\delta_1(\sigma) = \max_{1\leq i\leq n} \Big(\epsilon|\mathcal{N}_i^{\text{in}}| + c_1\sigma^2|\mathcal{N}_i^{\text{out}}|\Big),$$
$$\delta_2(\sigma) = \max_{1\leq i\leq n} \Big(\epsilon \sum_{j\in\mathcal{N}_i^{\text{in}}} \|\boldsymbol{t}_{ij}^{gt}\| + c_1\sigma^2 \sum_{j\in\mathcal{N}_i^{\text{out}}} \|\boldsymbol{t}_{ij}^{gt}\|\Big),$$
$$\delta_3(\sigma) = \epsilon^2 \sum_{(i,j)\in\mathcal{E}^{\text{in}}} \|\boldsymbol{t}_{ij}^{gt}\|^2 + c_1\sigma^2 \sum_{(i,j)\in\mathcal{E}^{\text{out}}} \|\boldsymbol{t}_{ij}^{gt}\|^2.$$

Suppose $\delta_1(1) \leq \frac{1}{3}\lambda_5^{gt}$ *and*

$$c_2(\delta_1(1) + \delta_3(1))\|L^{gt+}(\mathbf{1})\|_1 < 1$$

Then under mild conditions in $\sigma_{\min}$, $\sigma_{\max}$, *and* $k_{\max}$, *our interactive spectral translation synchronization approach converges to a solution* $\boldsymbol{t}_i^\star$ *that satisfies*

$$\|\boldsymbol{t}_{ij}^\star - \boldsymbol{t}_{ij}^{gt}\| \leq (1-\beta)\frac{c_3\delta_2^{\min}\|E_{ij}^T L^{gt}(\mathbf{1})^\dagger\|_1}{1 - c_3(\delta_1^{\min} + \delta_3^{\min})\|L^{gt}(\mathbf{1})^\dagger\|_1} + \beta\|\boldsymbol{t}_{ij}^{gt}\|, \qquad \forall (i,j)\in\mathcal{E}, \tag{12}$$

and

$$\|\boldsymbol{t}^\star - \boldsymbol{t}^{gt}\|_\infty \leq (1-\beta)\frac{c_3\delta_2^{\min}\|L^{gt}(\mathbf{1})^\dagger\|_1}{1 - c_3(\delta_1^{\min} + \delta_3^{\min})\|L^{gt}(\mathbf{1})^\dagger\|_1} + \beta\|\boldsymbol{t}_{ij}^{gt}\|, \tag{13}$$

where

$$\beta \leq \frac{{\delta_1^{\min}}^2}{2(\lambda_5^{gt} - \delta_1^{\min} - \delta_3^{\min})^2} \tag{14}$$

and $\delta_i^{\min} = \delta_i(\sigma_{\min})$

Proof: See Section E. □

Theorem 3 shows that our approach can effectively remove input outliers whenever their noise level is upper-bounded. Note that we do not place any assumptions on their correlations.

5. Experimental Results

We begin with the experimental setup in Section 5.1. We then present results on synthetic and real datasets in Section 5.2 and Section 5.3, respectively. Finally, we present an analysis of our approach in Section 5.4.

5.1. Experimental Setup

Synthetic datasets. We describe a synthetic dataset as $\mathcal{D}(n, p_{\text{edge}}, t, p_{\text{noise}}, \sigma)$. Here n denotes the number of images that are randomly sampled from the unit sphere. $p_{\text{edge}} \in [0,1]$ denotes the percentage of edges. When $t = r$, these edges follow the Erdős–Rényi model. When $t = g$, these edges connect nearest neighbors, meaning $\mathcal{G}$ is a geometric graph. p_{noise} denotes the percentage of edges, in which the associated directions are random outliers. σ is the variance of the inliers. We fix $n = 100$ and use two values for each other hyper-parameter, i.e., $p_{\text{edge}} \in \{0.3, 0.7\}$, $t \in \{'r', 'g'\}$, $p_{\text{noise}} \in \{0.1, 0.4\}$, $\sigma \in \{0.01, 0.03\}$. For each configuration of hyper-parameters, we sample 20 times. In total, we have 16 synthetic datasets. As $n = 100$, we simplify the notation as $\mathcal{D}(p_{\text{edge}}, t, p_{\text{noise}}, \sigma)$.

Real datasets. We perform evaluations on the dataset provided by [32]. To ensure uniqueness, we remove the images whose degree is less than 3.

Baseline techniques. The baselines include five state-of-the-art methods: LUD [21], 1DSfm [32], BATA [35], ShapeFit [10], Fused- TA [20].

Evaluation protocol. We compare our method with the baseline methods using the protocol in [32], which is based on estimating the optimal scaling and global translation to align the output of an algorithm with the ground-truth. On synthetic datasets, we report the mean translation error of each algorithm. On real datasets, we report the median translation error of each algorithm. This is because our approach may remove a small number of images. We include these images when calculating the median. Note that we find these images have large errors among baseline approaches as well and do not affect the median value.

	BATA	Fused_TA	1dSFM	LUD	ShapeFit	TranSync
$\mathcal{D}(0.7, r, 0.1, 0.01)$	3.87	<u>3.16</u>	30.6	4.15	2.86	**1.53**
$\mathcal{D}(0.7, g, 0.1, 0.01)$	3.73	<u>2.98</u>	29.1	3.38	2.43	**1.32**
$\mathcal{D}(0.7, r, 0.1, 0.03)$	5.68	<u>5.49</u>	34.3	11.91	8.30	**5.31**
$\mathcal{D}(0.7, g, 0.1, 0.03)$	4.94	<u>4.83</u>	30.5	10.03	7.14	**4.49**
$\mathcal{D}(0.7, r, 0.4, 0.01)$	<u>13.15</u>	14.96	106	26.26	102	**1.93**
$\mathcal{D}(0.7, g, 0.4, 0.01)$	<u>10.82</u>	13.89	100	15.28	17.4	**1.70**
$\mathcal{D}(0.7, r, 0.4, 0.03)$	<u>13.93</u>	17.45	108	44.19	228	**6.75**
$\mathcal{D}(0.7, g, 0.4, 0.03)$	<u>11.59</u>	16.00	96.3	28.90	78.6	**5.79**
$\mathcal{D}(0.3, r, 0.1, 0.01)$	6.93	5.44	70.8	6.57	<u>4.55</u>	**2.58**
$\mathcal{D}(0.3, g, 0.1, 0.01)$	4.37	4.34	55.5	<u>3.97</u>	21.4	**1.61**
$\mathcal{D}(0.3, r, 0.1, 0.03)$	10.12	<u>9.46</u>	69.6	18.93	14.0	**8.97**
$\mathcal{D}(0.3, g, 0.1, 0.03)$	6.21	<u>6.18</u>	50.9	11.11	8.64	**5.54**
$\mathcal{D}(0.3, r, 0.4, 0.01)$	<u>32.66</u>	57.87	226	86.14	592	**9.19**
$\mathcal{D}(0.3, g, 0.4, 0.01)$	<u>16.02</u>	20.83	219	38.36	267	**2.22**
$\mathcal{D}(0.3, r, 0.4, 0.03)$	<u>29.54</u>	40.80	232	96.23	508	**18.29**
$\mathcal{D}(0.3, g, 0.4, 0.03)$	<u>16.70</u>	23.46	234	46.64	297	**7.28**

Table 1. Translation errors($\times 10^{-3}$) on synthetic datasets. The top-performing approach is bold-faced. The second best is underlined.

5.2. Results on Synthetic Datasets

Table 1 presents the quantitative results of our approach and the baseline approaches in synthetic data sets. Overall, our approach outperforms all baseline approaches by salient margins. The top baseline approaches are BATA and FusedTA, which show different behaviors in different data sets.

BATA shows a strong performance among baseline approaches in data sets with high noise levels, that is, $p_{\text{noise}} = 0.4$. The error reductions of our approach on these datasets range 38.1% and 86.1%. The greatest improvement is in $\mathcal{D}(0.3, g, 0.4, 0.01)$. The improvements are great when $\sigma = 0.01$. The smallest improvement is in $\mathcal{D}(0.3, r, 0.4, 0.03)$. We observe similar behaviors when $\sigma = 0.03$. These salient improvements show the strength of our approach in handling outliers in the inputs. In particular, our approach excels when there are clear separations between inliers and outliers, i.e., there are large variations in magnitudes of σ_{ij}. This shows the strength of our local analysis.

Fused_TA (LUD and ShapeFit) show good performance in the small outlier regime or large σ. The error reductions of our approach range from 3.3% to 59.4%. The greatest improvement is in $\mathcal{D}(0.3, g, 0.1, 0.01)$. We obtain similar improvements when $\sigma = 0.01$. The smallest improvement is in $\mathcal{D}(0.7, g, 0.1, 0.03)$. We observe a similar relative performance when $\sigma = 0.03$. Moreover, the relative improvements when $p_{\text{edge}} = 0.3$ are greater than when $p_{\text{edge}} = 0.7$. Those results are consistent with our discussion that our formulation is very effective in pruning outliers. When p_{noise} and the variations in σ_{ij} are small, our approach is restricted to the linear approximation of $\boldsymbol{u}_4$ used to determine w_{ij} and s_i. This shows certain limitations when $\sigma = 0.03$ (See Figure 4).

Our approach and baseline approaches perform better on geometric graphs than random graphs. This is counter-intuitive because the spectral gap, which is closely related to stability of eigenvectors under noisy observations, is smaller on geometric graphs than on random graphs. However, in Appendix D.3 we show that for our approach the difference between $\boldsymbol{u}_4$ and $\boldsymbol{u}_4^{gt}$ is related to $\text{Tr}(L^{gt\dagger})$. In geometric graphs, this quantity is smaller than their random counterparts. This explains why geometric graphs have smaller translation errors than random graphs.

5.3. Results on Real Datasets

Table 2 shows the quantitative results in the 1DSFM benchmark datasets. In general, our approach still outperforms all baseline approaches. To better understand the relative performance of our approach and baseline approaches, we divide the fifteen datasets into four groups: (ds) that consists of Alamo, Montreal-Notre Dame, Notre Dame, Piazza del Popolo, which are dense graphs with small noise ratios, (dl) that consists of Ellis Island, Madrid Metropolis, and NYC Library, which are dense graphs with large noise ratios, (ss) that consists of Roman Forum, Tower of London, Vienna Cathedral, Yorkminster, which are sparse graphs with small noise ratios, and (sl) that consists of Gandarmenmarkt, Piccadilly, Trafalgar, Union Square, which are sparse graphs with large noise ratios.

	BATA	Fused_TA	1dSFM	LUD	ShapeFit	TranSync
Alamo	0.87	1.03	0.86	2.77	2.56	**0.82**
Montreal Notre Dame	0.73	2.06	1.23	1.20	1.96	**0.68**
Notre Dame	1.28	1.49	**0.92**	1.67	1.32	**0.92**
Piazza del Popolo	1.76	1.79	2.80	1.77	1.66	**1.65**
Roman Forum	6.04	42.94	**5.40**	24.77	46.01	5.65
Tower of London	2.86	4.52	9.40	8.93	47.49	**2.46**
Vienna Cathedral	2.95	5.38	3.90	5.91	11.23	**2.54**
Yorkminster	1.73	2.66	6.36	5.34	8.09	**1.56**
Ellis Island	4.09	7.70	3.46	8.02	16.24	**3.56**
Madrid Metropolis	3.78	7.07	8.12	8.84	28.80	**2.57**
NYC Library	1.06	1.90	2.00	2.22	9.83	**0.92**
Gendarmenmarkt	58.30	31.48	46.72	30.19	33.18	**27.81**
Piccadilly	1.74	7.25	2.45	4.00	14.59	**1.56**
Trafalgar	5.90	15.35	6.66	13.26	53.38	**5.12**
Union Square	5.48	7.10	**4.82**	7.32	10.60	6.12

Table 2. Translation errors for real dataset from 1dsfm. Units are specified by the ground-truth.

In (ds), which are easier datasets, our approach is comparable to the top performing baselines. In these data sets, all approaches can successfully remove small fractions of outliers in the inputs and perform a prediction from the remaining inputs.

In (ss), which are harder than (ds) because of the sparsity of the graph, our approach outperforms the baseline approaches in average. This is attributed to our interactive scheme and the local analysis that shows how to properly reweight the measurements.

In (ds) and (dl), which have large noise ratios, our approach outperforms baseline approaches except Union Square. This again shows the advantage of our approach, which is based on principled analysis to reweight inputs for translation synchronization.

Note that the relative improvements of our approach in real data are smaller than the relative improvements in synthetic data. One reason is that there is no clear separation of inliers and outliers in real data. However, our approach can still remove outliers and properly reweight the inputs to balance their errors to obtain considerably improved results.

5.4. Advanced Analysis

We proceed to analyze our iterative spectral translation synchronization approach. For this analysis, we show average results on the synthetic datasets, where we know the distributions of input errors.

Effects of the alternating minimization. As shown in Figure 3(a), the iterative procedure can effectively reduce the prediction error. The relative improvements depend on the type of observations. On dense graphs with small noise ratio, where the initialization

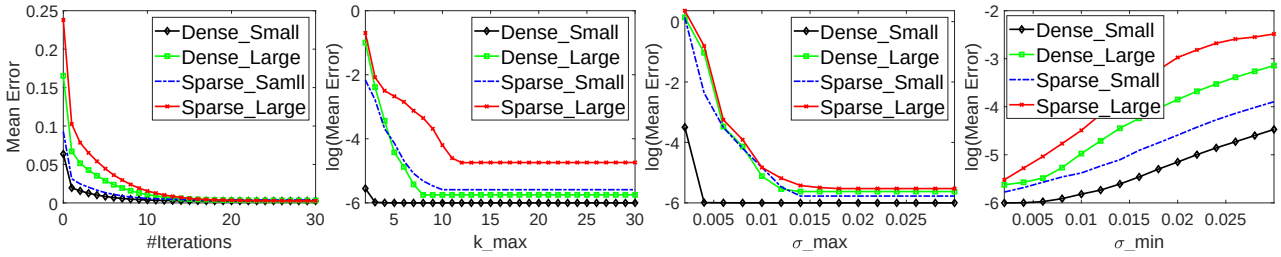

(a) (b) (c) (d)

Figure 3. Advanced analysis of our approach. We show average translation error among all synthetic datasets. Above are on dense graphs with small noise ratio and bottom are on sparse graphs with large noise ratio (a) Effects of alternating minimization. (b) Effects of varying $k_{\max}$. (c) Effects of varying $\sigma_{\max}$. (d) Effects of varying $\sigma_{\min}$.

stage already returns a good solution, the improvement obtained from the iterative procedure is not significant. However, on sparse graphs with large noise ratio, the initial solution has large error as outliers participate in the estimation. Our reweighted scheme can substantially reduce the prediction error as the outliers have small weights in the end.

Effects of varying k_{max}. Figure 3(b) shows the results when varying $k_{\max}$. We can see that when $k_{\max}$ is above 10 across all datasets, the output of our algorithm becomes steady. In addition, on more challenging datasets, i.e., sparse graphs with large noise ratios, $k_{\max}$ that is required to reach a steady state is larger. This is expected, as when using a small $k_{\max}$, the weights of inliers may not differentiate from those of outliers during the alternating procedure due to inaccurate initial solution.

Effects of varying $\sigma_{\max}$. Figure 3(c) shows the effects when varying $\sigma_{\max}$. We can see that for easy datasets, there is not much difference when choosing a small $\sigma_{\max}$. This is because the initial solution is close to the underlying ground truth. The weighting scheme with a small $\sigma_{\max}$ can differentiate inliers and outliers. However, for hard data sets, that is, sparse graphs with large noise ratios, it is important to start from a large $\sigma_{\max}$. This allows us to gradually remove outliers at different noise levels.

Effects when varying $\sigma_{\min}$. Figure 3(d) shows the effects when varying $\sigma_{\min}$. We can see that decreasing the value of $\sigma_{\min}$ improves the prediction accuracy. However, $\sigma_{\min}$ can not be too small, i.e., smaller than variances of inliers. In this case, the weighting scheme cannot differentiate inliers and outliers, and we see that the prediction error drastically goes up. This effect is particularly salient on hard datasets, i.e., sparse graphs with large noise ratios.

6. Conclusions and Limitations

In this paper, we have introduced a robust translation synchronization approach that iteratively solves gen-

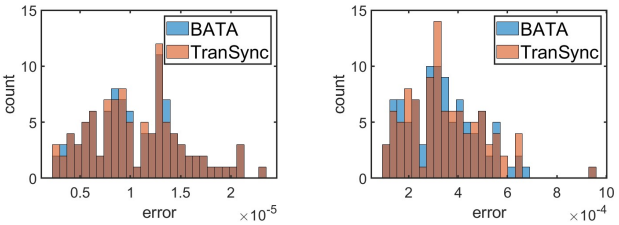

(a) (b)

Figure 4. Comparison with our approach and BATA under $\mathcal{N}(0.7, g, 0, \sigma)$. (Left) $\sigma = 10^{-3}$. (Right) $\sigma = 0.03$.

eralized eigenvalue problems. We present a rigorous stability analysis of this spectral formulation. The results are used to determine how to adaptively set the edge weights to progressively remove input outliers and average inliers of the input. Experimental results justify the effectiveness of our approach.

One limitation of our approach is that we derive the optimal values of s_i and σ_{ij} in the infinitesimal regime, in which we have simple expressions of $\boldsymbol{v}_{ij}^{\text{inp}}$ and $\boldsymbol{u}_4$. The optimal w_{ij} and s_i are effective in removing outliers in $\boldsymbol{v}_{ij}^{\text{inp}}$. On the other hand, when all $\boldsymbol{v}_{ij}^{\text{inp}}$ are inliers with nonnegligible variance, we find that the linear approximation we used is sub-optimal when σ is large. As shown in Figure 4, our approach slightly outperforms BATA when $\sigma = 10^{-3}$. However, BATA slightly outperforms our approach when $\sigma = 0.03$. In the future, we plan to address this issue using second-order approximations of $\boldsymbol{u}_4$.

In the future, we also plan to integrate translation synchronization and rotation synchronization into a single optimization formulation, since the constraints for translation synchronization depend on rotations. Another direction is to explore learning-based approaches to set the edge weights, which have been shown to be effective for pose synchronization of RGB-D scans.

Acknowledgment. Q. H. acknowledges the support from NSF IIS-2047677, IIS-2413161, and GIFTs from Adobe and Google.

References

[1] Federica Arrigoni, Andrea Fusiello, and Beatrice Rossi. Camera motion from group synchronization. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 546–555. IEEE, 2016. 2

[2] Federica Arrigoni, Beatrice Rossi, and Andrea Fusiello. Spectral synchronization of multiple views in se (3). *SIAM Journal on Imaging Sciences*, 9(4):1963–1990, 2016. 2

[3] Chandrajit Bajaj, Tingran Gao, Zihang He, Qixing Huang, and Zhenxiao Liang. SMAC: simultaneous mapping and clustering using spectral decompositions. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 334–343. PMLR, 2018. 2

[4] Florian Bernard, Johan Thunberg, Peter Gemmar, Frank Hertel, Andreas Husch, and Jorge Goncalves. A solution for multi-alignment by transformation synchronisation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2161–2169, 2015. 2

[5] Avishek Chatterjee and Venu Madhav Govindu. Efficient and robust large-scale rotation averaging. In *ICCV*, pages 521–528. IEEE Computer Society, 2013. 2

[6] Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Spectral methods for data science: A statistical perspective. *Found. Trends Mach. Learn.*, 14(5):566–806, 2021. 2

[7] Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun. Robust reconstruction of indoor scenes. In *CVPR*, pages 5556–5565. IEEE Computer Society, 2015. 2

[8] Frank Dellaert, David M Rosen, Jing Wu, Robert Mahony, and Luca Carlone. Shonan rotation averaging: Global optimality by surfing $so(n)$. In *European Conference on Computer Vision*, pages 292–308. Springer, 2020. 2

[9] Venu Madhav Govindu. Combining two-view constraints for motion estimation. In *Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001*, pages II–II. IEEE, 2001. 2

[10] Paul Hand, Choongbum Lee, and Vladislav Voroninski. Shapefit: Exact location recovery from corrupted pairwise directions. *Communications on Pure and Applied Mathematics*, 71(1):3–50, 2018. 2, 6

[11] Qixing Huang and Leonidas Guibas. Consistent shape maps via semidefinite programming. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing*, pages 177–186, 2013. 2

[12] Qi-Xing Huang, Guo-Xin Zhang, Lin Gao, Shi-Min Hu, Adrian Butscher, and Leonidas J. Guibas. An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Trans. Graph.*, 31(6): 167:1–167:11, 2012. 2

[13] Qixing Huang, Fan Wang, and Leonidas J. Guibas. Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.*, 33(4):36:1–36:11, 2014. 2

[14] Xiangru Huang, Zhenxiao Liang, Chandrajit Bajaj, and Qixing Huang. Translation synchronization via truncated least squares. In *NIPS*, 2017. 2

[15] Xiangru Huang, Zhenxiao Liang, Xiaowei Zhou, Yao Xie, Leonidas J. Guibas, and Qixing Huang. Learning transformation synchronization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 8082–8091. Computer Vision Foundation / IEEE, 2019. 2

[16] Xiangru Huang, Zhenxiao Liang, Xiaowei Zhou, Yao Xie, Leonidas J Guibas, and Qixing Huang. Learning transformation synchronization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8082–8091, 2019. 2

[17] Vladimir G. Kim, Wilmot Li, Niloy J. Mitra, Stephen DiVerdi, and Thomas Funkhouser. Exploring collections of 3d models using fuzzy correspondences. *ACM Trans. Graph.*, 31(4):54:1–54:11, 2012. 2

[18] Spyridon Leonardos, Xiaowei Zhou, and Kostas Daniilidis. Distributed consistent data association via permutation synchronization. In *ICRA*, pages 2645–2652. IEEE, 2017. 2

[19] Shaohan Li, Yunpeng Shi, and Gilad Lerman. Fast, accurate and memory-efficient partial permutation synchronization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 15714–15722. IEEE, 2022. 2

[20] Lalit Manam and Venu Madhav Govindu. Fusing directions and displacements in translation averaging. In *International Conference on 3D Vision (3DV)*, 2024. 2, 6

[21] Onur Özyesil and Amit Singer. Robust camera location estimation by convex programming. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 2674–2683. IEEE Computer Society, 2015. 2, 6

[22] Deepti Pachauri, Risi Kondor, and Vikas Singh. Solving the multi-way matching problem by permutation synchronization. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 1860–1868, 2013. 2

[23] Deepti Pachauri, Risi Kondor, Gautam Sargur, and Vikas Singh. Permutation diffusion maps (PDM) with application to the image association problem in computer vision. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 541–549, 2014. 2

[24] D.M. Rosen, L. Carlone, A.S. Bandeira, and J.J. Leonard. SE-Sync: A certifiably correct algorithm for synchronization over the special Euclidean group. *Intl. J. of Robotics Research*, 38(2–3):95–125, 2019. 2

[25] Yanyao Shen, Qixing Huang, Nati Srebro, and Sujay Sanghavi. Normalized spectral map synchronization.

In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4925–4933, 2016. 2

[26] Yanyao Shen, Qixing Huang, Nati Srebro, and Sujay Sanghavi. Normalized spectral map synchronization. In *NIPS*, pages 4925–4933, 2016. 2

[27] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30:20–36, 2011. 2

[28] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and computational harmonic analysis*, 30:20–36, 2011. 2

[29] Amit Singer and Hau tieng Wu. Vector diffusion maps and the connection laplacian. *Communications on Pure and Applied Mathematics*, 65:1067–1144, 2012. 2

[30] Yifan Sun, Zhenxiao Liang, Xiangru Huang, and Qixing Huang. Joint map and symmetry synchronization. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part V*, pages 257–275, 2018. 2

[31] Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2:145–193, 2013. 2

[32] Kyle Wilson and Noah Snavely. Robust global translations with 1dsfm. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III 13*, pages 61–75. Springer, 2014. 2, 6

[33] Yuxin Chen, Leonidas J. Guibas, and Qi-Xing Huang. Near-optimal joint object matching via convex relaxation. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 100–108, 2014. 2

[34] Xiaowei Zhou, Menglong Zhu, and Kostas Daniilidis. Multi-image matching via fast alternating minimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4032–4040, 2015. 2

[35] Bingbing Zhuang, Loong-Fah Cheong, and Gim Hee Lee. Baseline desensitizing in translation averaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4539–4547, 2018. 2, 3, 6