



Depth Descent Synchronization in $SO(D)$

Tyler Maunu¹ · Gilad Lerman²

Received: 17 November 2020 / Accepted: 27 August 2022 / Published online: 3 January 2023
 © The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

We give robust recovery results for synchronization on the rotation group, $SO(D)$. In particular, we consider an adversarial corruption setting, where a limited percentage of the observations are arbitrarily corrupted. We develop a novel algorithm that exploits Tukey depth in the tangent space of $SO(D)$. This algorithm, called Depth Descent Synchronization, exactly recovers the underlying rotations up to an outlier percentage of $1/(D(D-1)+2)$, which corresponds to $1/4$ for $SO(2)$ and $1/8$ for $SO(3)$. In the case of $SO(2)$, we demonstrate that a variant of this algorithm converges linearly to the ground truth rotations. We implement this algorithm for the case of $SO(3)$ and demonstrate that it performs competitively on baseline synthetic data.

Keywords Robust synchronization · Structure from motion · Nonconvex optimization · Multiple rotation averaging

1 Introduction

The typical synchronization problem involves recovery of n group elements from pairwise measurements between them. It arises, for example, when solving the Structure from Motion (SfM) problem. One subproblem of SfM is to recover the three-dimensional orientations and positions of cameras from pairwise orientations and positions in relation to a scene (Özyeşil et al., 2017). Here, we specifically focus on robust synchronization over $SO(D)$, the rotation group for $\mathbb{R}^D$. That is, given pairwise rotations in $SO(D)$, some of which are corrupted, we aim to recover the original set of n rotations.

We assume n unknown, ground truth elements of $SO(D)$, which we denote by $R_1^*, \dots, R_n^*$. We form a graph $G([n], E)$, where $[n] := \{1, \dots, n\}$ indexes the n unknown elements and E designates the edges for which measurements of relative rotations are taken. For each $jk \in E$, we are provided with

the measurement

$$R_{jk}^* = R_j^* R_k^{*\top}. \quad (1)$$

We can think of R_{jk}^* in the following way: If we are oriented in the coordinate system with respect to node k , then R_{jk}^* rotates our coordinate system into the coordinate system we would see if we were sitting at node j . This synchronization formulation extends to any given group, where one wishes to recover $(g_1, \dots, g_n)$, an n -tuple of elements in the group, given measurements of the group ratios $g_i g_j^{-1}$, $i, j = 1, \dots, n$.

In reality, we cannot hope to exactly measure all the pairwise rotations in (1). In many real systems, both noisy and corrupted measurements occur: our focus here is on *adversarially corrupted measurements*. That is, within the measurement graph G , the corruption model is assumed to be fully adversarial. Our model is specified by partitioning the measured data into two parts:

1. We observe corrupted (or “bad”) edges $E_b \subset E$, where all edges in E_b have a corresponding arbitrary corruption. The adversary is allowed to choose E_b (and thus may to some degree influence the connectivity of $E \setminus E_b$) as well as the corrupted values R_{jk}^* for $jk \in E_b$. For each node, the adversary is only allowed to corrupt a limited fraction of edges.
2. The rest of the observed edges are uncorrupted (or “good”) edges $E_g = E \setminus E_b$, where each edge in E_g has an associated measurement given by (1).

Communicated by Laurent Najman.

✉ Tyler Maunu
 maunu@brandeis.edu
 Gilad Lerman
 lerman@umn.edu

¹ Department of Mathematics, Brandeis University, Waltham, USA

² School of Mathematics, University of Minnesota, Minneapolis, USA

Theoretically guaranteed methods for robust synchronization are still lacking, especially in adversarial and nonconvex settings. The development of these methods is important because in practice measurements are usually quite corrupted, especially in applied problems like Structure from Motion (Özyeşil et al., 2017). The results we establish here are concerned with *exact recovery*. That is, given a set of corrupted measurements, we wish to exactly recover $\mathbf{R}_1^*, \dots, \mathbf{R}_n^*$. We will show that this is possible for a nonconvex method even in the presence of a significant amount of arbitrary corruption.

Our method falls into the class of *multiple rotation averaging algorithms* (Govindu, 2004; Martinec & Pajdla, 2007; Hartley et al., 2013). These methods are effectively coordinate descent algorithms, which are highly efficient algorithms for nonconvex programs that are notoriously hard to analyze. While their analysis is challenging, it is imperative to develop a theoretical understanding of these methods and their robust counterparts (Hartley et al., 2011; Chatterjee & Govindu, 2017). Moreover, as we discuss later, there are few robustness guarantees for group synchronization with adversarial corruption. Among the limited guarantees, none cover our model, and we thus make a significant contribution to this area. This work is also of general appeal to the nonconvex optimization community since we are able to prove convergence results in the complex nonconvex landscape of robust multiple rotation averaging. Furthermore, some energy landscapes associated with this problem exhibit many local minima and spurious fixed points, which we are able to avoid with our new method.

1.1 Contributions of This Work

The main contributions of this work follow.

1. As a warm-up, we develop an adversarially robust algorithm for synchronization in $\text{SO}(2)$, which we call Trimmed Averaging Synchronization (TAS). In Theorem 3, under a generic condition on the measurement graph G , which we call the “well-connectedness”, and proper initialization, we show that it can tolerate a fraction of outliers per node that is bounded above by $1/4$. We further prove that it converges linearly for fully connected observation graphs.
2. To extend this result to $\text{SO}(D)$, we develop a new algorithm that we call *Depth Descent Synchronization (DDS)* based on Tukey depth in the tangent space of $\text{SO}(D)$. To our knowledge, this is the first application of a manifold version of Tukey depth in an applied setting.
3. Assuming well-connectedness and good initialization, the DDS algorithm exactly recovers an underlying signal in the presence of a significant amount of adversarial outliers for D sufficiently small. This result is given in

Theorem 6 and is the first guarantee of robustness to adversarial corruption for a multiple rotation averaging algorithm. This result extends elegantly to sparse random graphs, where we show that it achieves the information theoretic rate with respect to graph sparsity for Erdős-Rényi observation graphs in Sect. 4.5. These bounds are state-of-the-art for rotation synchronization algorithms in the cases of $D = 2$ and $D = 3$, which are the primary cases of interest in computer vision applications like SfM.

4. We show that this algorithm can be efficiently implemented for $\text{SO}(3)$. We run baseline experiments that show it performs competitively on some baseline synthetic data for $\text{SO}(3)$ synchronization, which arises in the important application of Structure from Motion.

While we carefully review related work later in Sect. 2, we emphasize here our contributions in terms of the most relevant works. Again, we note that we study an efficient, nonconvex algorithm for rotation synchronization that has guarantees for adversarial outliers.

A robustness result for $\text{SO}(D)$ synchronization based on semidefinite programming is given in Wang and Singer (2013). However, the probabilistic model in this work is very restrictive (see details in Sect. 2), and the proposed method is slow for large n .

Huang et al. (2019) use a truncated least squares framework to do robust rotation synchronization. The truncated least squares framework was originally proposed by Huang et al. (2017) in the context of translation synchronization, and sequentially filters those pairwise measurements that are furthest from the current estimated pairwise measurements. Following this, Huang et al. (2019) show that this can be extended to rotation synchronization. Under an appropriate choice of a thresholding parameter, they demonstrate that it is possible to exactly recover the ground truth in the presence of outliers if a certain generic condition is satisfied. This method has two downsides. First, one must repeatedly compute the lowest eigenvectors of the graph connection Laplacian, which has a higher memory cost for dense graphs than DDS and may also have issues of numerical stability. Second, the bound on the fraction of outliers that they present is not clear in general settings since it depends on the $(1, \infty)$ -norm of the pseudoinverse of the graph connection Laplacian. This quantity is hard to control, and so it is unclear how their bound scales with various parameters as well as what it would state for arbitrary outliers. On the other hand, Huang et al. (2017) guarantees success of the truncated least squares method for adversarial outliers when considering translation synchronization in 1-dimension, but we do not see how to extend these to the problem of rotation synchronization. In review, it was pointed out to us that the proof of Huang et al. (2019) can be extended to show that it can tolerate a con-

stant fraction of corruption in the rotation synchronization problem. However, this constant may be very small, and in particular is much smaller than our bounds for the cases of $SO(2)$ and $SO(3)$. We discuss this more in Sect. 5.

The only existing result for adversarial robust synchronization was recently given by Lerman and Shi (2020). They propose a general method, called Cycle-Edge Message Passing (CEMP), for group synchronization that is guaranteed to be robust to adversarial corruption. However, their method uses information from 3-cycles, that is, triangles in the graph, and so it is less efficient than typical multiple rotation averaging schemes by an order of n (the ratio between the number of triangles and the number of edges in the graph). Beyond this, multiple rotation averaging algorithms are also attractive because they are more memory efficient. A caveat to our current work is that our new method is not as efficient as previous multiple rotation averaging algorithms. In particular, we require the computation of a depth-based estimator, and so each rotation update has complexity $O(n_j^3 \log(n_j))$ for $SO(3)$, where n_j is the number of neighbors of the node to be updated. Therefore, we do not claim that DDS is uniformly most efficient for adversarially robust synchronization in terms of time complexity, although it is more computationally efficient than CEMP for very sparse graphs. However, our depth descent method does still have the benefit of more efficient memory complexity than Lerman and Shi (2020).

Beyond computational efficiency, the theoretical guarantees are also different: we bound the ratio of corrupted edges, whereas Lerman and Shi (2020) bound the ratio of corrupted triangles. The method of Lerman and Shi (2020) degrades with extremely sparse graphs, due to the fact that they need to ensure that it contains sufficiently many triangles, whereas we require a well-connectedness condition of the graph G that extends to extremely sparse cases. Other well-connectedness conditions appear, sometimes implicitly, in works minimizing energy functions (Wang & Singer, 2013; Hand et al., 2018; Lerman et al., 2018; Huang et al., 2017). Finally, CEMP is tailored to finding the corruption level in the graph, but it does not have complete guarantees for the recovery of the underlying rotations themselves.

1.2 Notation

Bold uppercase letters will be used to denote matrices, while bold lowercase letters will be used to denote vectors. For a set $\mathcal{X}$ in a Hilbert space, the convex hull is denoted by $\text{conv}(\mathcal{X})$. The sphere in $\mathbb{R}^D$ is written as S^{D-1} . For an ordered tuple of n rotations, $\mathbf{R}_1, \dots, \mathbf{R}_n \in SO(D)$, we write $(\mathbf{R}) = (\mathbf{R}_1, \dots, \mathbf{R}_n)$.

1.3 Structure of the Rest of the Paper

We now outline the structure of this paper. First, we review related work in Sect. 2. We then discuss the specific case of synchronization over $SO(2)$ in Sect. 3 and give a simple, adversarially robust algorithm, called Trimmed Averaging Synchronization. Following this, in Sect. 4 we develop our novel Depth Descent Synchronization algorithm, which utilizes Tukey depth to yield robust updates. Coupled with this, we develop its theoretical guarantees of robustness and convergence. In Sect. 5, we compare our results with existing theory. Section 6 presents some baseline experiments demonstrating the practical implementation of our proposed method.

2 Related Work

Interest in the synchronization problem has grown in recent years due to applications in computer vision and image processing, such as SfM (Govindu, 2004; Martinec & Pajdla, 2007; Arie-Nachimson et al., 2012; Hartley et al., 2013; Tron & Vidal, 2009; Ozyesil et al., 2015; Boumal, 2016), cryo-electron microscopy (Wang & Singer, 2013) and Simultaneous Localization And Mapping (SLAM) (Rosen et al., 2019).

The most common formulation for solving rotation and other group synchronization problems involve a non-convex least squares formulation that can be addressed by spectral methods (Singer, 2011) or semidefinite relaxation (Bandeira et al., 2017). On the other hand, the work of Wang and Singer (2013) uses a semidefinite relaxation of a least absolute deviations formulation to obtain a robust estimate for $SO(d)$ synchronization. They prove recovery for the pure optimizer of this convex problem in a restricted setting. In this setting the full graph is complete, every edge is corrupted with a certain probability p (in the case of $SO(2)$, they require that $p \leq 0.543$ and for $SO(3)$ they require $p \leq 0.5088$) and the corrupted group ratios are distributed uniformly on $SO(D)$. In practice, they advocate using an alternating direction augmented Lagrangian to solve their optimization problem. One may also use methods like the Burer-Monteiro formulation (Boumal et al., 2020), although current guarantees require the rank of the semidefinite program to be at least $O(\sqrt{n})$, which results in storing iterates much larger than the underlying signal that is a matrix of size $n \times D$ (Waldspurger & Waters, 2020). Another recent work tries to leverage a low-rank plus sparse decomposition for robust synchronization (Arrigoni et al., 2018). However, this work does not contain robustness guarantees.

2.1 Robust Synchronization Methods

For a survey of robust rotation synchronization, see Tron et al. (2016). Some early works on rotation synchronization include Govindu (2001, 2006); Martinec and Pajdla (2007), with later follow-up works by Hartley et al. (2013); Chatterjee and Govindu (2013, 2017). The later works discuss some least absolute deviations based approaches to multiple rotation averaging which we will discuss later. For theoretical foundations on averaging rotations, one can consult Moakher (2002). For foundational work on optimization on the manifold $SO(d)$, see (Taylor & Kriegman, 1994; Arora, 2009).

Robust multiple rotation averaging algorithms were studied in Hartley et al. (2011) and Hartley et al. (2013). There, the authors used a least absolute deviations formulation over $SO(3)$ using successive averaging with a Weiszfeld algorithm and a gradient-based algorithm. The authors also give a counterexample that shows that local minima exist and thus the global minimum of their problem may be hard to find in general. However, the authors give no guarantee of convergence or recovery in any setting. Also, we have found that this method may suffer from suboptimal fixed points in general, which we analyze in more detail in a forthcoming work [see also Section 5 of Maunu and Lerman (2020)].

2.2 Adversarially Robust Synchronization

A work that does contain guarantees is that of Lerman and Shi (2020), which considers a message-passing procedure that incorporates consistent information from cycles. This algorithm was guaranteed to be robust for the adversarial setting and applies to any compact group. Although its adversarial setting is very general, it requires a bound on the ratio of corrupted cycles per edge and not on the ratio of corrupted edges. Furthermore, the use of cycles results in a potentially more computationally intensive algorithm than the one in this work that only uses pairwise information.

Guarantees for exact recovery with adversarial, or partially adversarial, corruption appear in few other synchronization problems. The adversarial corruption in $\mathbb{Z}_2$ synchronization is very special since there is a single choice to corrupt a group ratio. Under a special probabilistic model, Bandeira (2018) established asymptotic and probabilistic exact recovery for the SDP relaxation of the least squares energy function of $\mathbb{Z}_2$ synchronization. The model assumes that $G([n], E)$ is an Erdős-Rényi graph with probability p of connection, edges are randomly corrupted with probability q and $p(1 - 2q)^2 \leq 0.5$. Hand et al. (2018) and Lerman et al. (2018) established asymptotic exact recovery under a probabilistic model for solutions of the different problem of location recovery from pairwise orientations. In this problem ratios of the Euclidean group are normalized to the sphere. They assume an i.i.d. Gaussian generative probabilis-

tic model for the ground truth locations and an Erdős-Rényi model for the graph $G([n], E)$ and further bounded the ratio of maximal degree of $G([n], E_b)$ over n . In both works, these bounds approach zero as n approaches infinity, unlike the constant bound of this work.

Robust permutation synchronization was studied by Huang and Guibas (2013), where they give a maximum corruption percentage of 1/4 in the case of fully connected observation graphs. Huang et al. (2017) analyzed a robust algorithm for one-dimensional translation synchronization that uses a truncated least squares formulation. They show that their method achieves a maximum corruption percentage of 1/6 for fully connected graphs. However, their generic condition is rather complicated and in order to interpret it they must make the fully connected assumption and they also restrict the maximal degree of $G([n], E_b)$. In both (Huang & Guibas, 2013; Huang et al., 2017), the bounds degrade for sparser observation graphs.

The results of Huang et al. (2017) were extended to the problem of rotation synchronization in Huang et al. (2019). Here, the authors show that a truncated least squares formulation for rotation synchronization can recover the underlying signal assuming a generic bound that includes the $(1, \infty)$ -norm of the pseudoinverse of the graph connection Laplacian.

2.3 Synchronization in Other Settings

In contrast to corrupted settings, some works have considered estimation in a noisy setting. Bandeira et al. (2017) study maximum likelihood estimation of the angular synchronization problem and show that the associated semidefinite relaxation is tight. More recently, message-passing algorithms have been used for maximum likelihood estimation in the Gaussian setting (Perry et al., 2018). Other recent results leverage multiple phases to obtain better results in noisy settings (Gao & Zhao, 2019). Minimax estimation under the squared loss over $SO(2)$ is considered in Gao and Zhang (2021), and optimization methods for the squared error over subgroups over the orthogonal group are considered in Liu et al. (2020).

Another related problem over $SO(2)$ is the synchronization of Kuramoto oscillators. In particular, a primary question is the minimal graph connectivity requirement ensuring that the energy landscape is nice. The weakest known requirement is that every vertex is connected to at least $0.7889n$ other vertices (Lu & Steinerberger, 2020). The conjectured bound is $0.75n$, which is reminiscent of the bound we require for local recovery with adversarial corruption over $SO(2)$.

2.4 Nonconvex Optimization

Optimization problems cast over $SO(D)$ are usually nonconvex. We can think of our method as attempting to solve a

nonconvex problem over $SO(D)$ as well. Therefore, our work also fits in with the growing body of work analyzing nonconvex energy landscapes and procedures (Dauphin et al., 2014; Hardt, 2014; Jain et al., 2014; Netrapalli et al., 2014; Yi et al., 2013; Zhang & Yang, 2018; Ge et al., 2015; Lee et al., 2016; Arora et al., 2015; Mei et al., 2018; Ge et al., 2016; Boumal, 2016; Sun et al., 2015b, a; Lerman & Maunu, 2017; Cherapanamjeri et al., 2017; Ma et al., 2018; Maunu et al., 2019).

2.5 Tukey Depth

In our method, we use a notion of tangent space depth, that is, using Tukey or halfspace depth (Tukey, 1974) in the tangent space of the manifold $SO(D)$, to create a provable robust method. Tangent space depth, for a general manifold, first appears to have been discussed in Mizera & Volau (2002), where the author proves existence and depth bounds for maximum tangent depth estimators. Earlier work on Tukey depth includes Rado (1946), which proves a depth lower bound for general measures, and Danzer et al. (1963), which discusses the relation to Helly's Theorem. More recently, the classical reference of Donoho and Gasko (1992) proves bounds on the maximum depth achieved in a dataset under ellipticity conditions. Computation of depth contours was considered in Liu (2017); Hammer et al. (2022). Recently, an interesting connection between depth estimators and generative adversarial networks has been exhibited (Gao et al., 2019), which may perhaps lead to more computationally efficient estimators.

2.6 Notions of Robustness

We finish by clarifying our setting in the context of robustness. In order to quantify our notion of robustness, we introduce the following terminology. Recall that we have an underlying graph $G([n], E)$ corresponding to the pairwise measurements, where E is partitioned into an inlier set, E_g , and outlier set E_b . For any $j \in [n]$, we define its neighborhood as well as its inlier and outlier neighborhoods as

$$\begin{aligned} E^j &= E_g^j \cup E_b^j, & E_g^j &:= \{k \in [n] : jk \in E_g\}, \\ E_b^j &:= \{k \in [n] : jk \in E_b\}. \end{aligned} \quad (2)$$

We will denote by α_0 the maximum fraction of outliers per node. That is, α_0 is the maximum of $\#(E_b^j)/n_j$ over all $j \in [n]$, where throughout the rest of the paper $\#(\cdot)$ denotes the number of points in a set and n_j is the degree of node j , $n_j = \#(E^j)$.

The following notion of recovery threshold is related to the notion of a breakdown point in robust statistics. However, our goal is somewhat different, since we desire an exact estimator rather than an approximation, as is typically considered for

classical breakdown points. This is similar to the notion of RSR breakdown point given in Section 1.1 of Maunu and Lerman (2019).

Definition 1 (Recovery Threshold) The recovery threshold of a robust rotation synchronization algorithm is the largest value of α_0 such that the algorithm outputs an estimator ($\hat{\mathbf{R}}$) that satisfies (20).

The simplest information-theoretic bound for the recovery threshold is $\alpha_0 \leq 1/2$. Indeed, if $\alpha_0 > 1/2$, then an adversary could easily choose E_b to have a subgraph that dominates E_g with a consistent set of measurements for an alternative signal ($\mathbf{R}^b = (\mathbf{R}_1^b, \dots, \mathbf{R}_n^b)$). That is, the observations would be

$$\mathbf{R}_{jk} = \begin{cases} \mathbf{R}_j^{*\top} \mathbf{R}_k^*, & jk \in E_g \\ \mathbf{R}_j^{b\top} \mathbf{R}_k^b, & jk \in E_b. \end{cases} \quad (3)$$

If the adversary chooses the partition E_g and E_b properly, then one could easily think that $\mathbf{R}^b$ is the true underlying signal. For our method, we obtain recovery thresholds for α_0 that are smaller than $1/2$.

On the other hand, the information theoretic bound may be much higher in special models. For example, suppose that $G([n], E)$ is an Erdős-Rényi graph with parameter p . Suppose further that each edge in E is corrupted independently with probability q , and the corrupted measurements are i.i.d. uniform on $SO(D)$. Wang and Singer (2013) call this the uniform corruption model. Then, Singer (2011) and Chen et al. (2016) established the following information theoretic threshold for the rotation synchronization problem:

$$q = 1 - \Omega\left(\sqrt{\frac{1}{p} \frac{\log(n)}{n}}\right). \quad (4)$$

Notice that one needs $p \gtrsim \log(n)/n$ to ensure that the underlying Erdős-Rényi graph is connected. For fixed p , notice that one can take q arbitrarily close to 1 (and so α_0 is then very close to 1) as long as n is sufficiently large. We later discuss how our main results extend to this model in Sect. 4.5, where we show that the DDS algorithm achieves optimal recovery rates with respect to p (i.e., it can tolerate extremely sparse observation graphs).

3 An Adversarially Robust Algorithm for $SO(2)$ Synchronization

To begin to build motivation for our method, we consider the case of synchronization over $SO(2)$, where the method becomes considerably simpler due to its 1-dimensional manifold structure. First, Sect. 3.1 gives definitions of some geometrical objects on $SO(2)$, which we identify with $\mathbb{C}_1$




Fig. 1 The angular distance function $d_L(e^{i\theta}, 1)$

for mathematical convenience. Then, in Sect. 3.2, we define our SO(2) synchronization method, which we call Trimmed Averaging Synchronization (TAS) and is a special case of our later DDS algorithm. Finally, Sect. 3.3 discusses the initialization and well-connectedness assumptions and uses these to give an adversarial recovery guarantee for the TAS algorithm.

3.1 The Geometry of $\mathbb{C}_1$

We define a few structures related to the manifold $\mathbb{C}_1$. The tangent space can be identified with $\mathbb{R}$. Let $v \in T_z\mathbb{C}_1$ be a unit direction in the tangent space at z_j (i.e., $v = \pm 1$). The geodesic originating at z_j in the direction v is given by $\gamma(t) = e^{ivt}z_j$, $t \in [0, \pi/|v|]$. The exponential map and inverse exponential map (logarithm map) on this 1-dimensional manifold are given by

$$\text{Exp}_z(\theta) = e^{i\theta}z, \theta \in (-\pi, \pi], \quad \text{Log}_z(y) = \arg(y\bar{z}). \quad (5)$$

Finally, the *cut-locus* of a point $z \in \mathbb{C}_1$ is defined as the set of points for which there is not a unique geodesic from z . It is not hard to see that this is given by $\text{cut}(z) = \{-z\}$.

Recall that we seek an underlying signal $z^* \in \mathbb{C}_1^n$. Notice that its elements, $z_j^* \in \mathbb{C}_1$ for $j \in [n]$, can be parameterized by angles, $z_j^* = e^{i\theta_j^*}$. This angle is also known as the argument of the complex number, and so we write $\arg(e^{i\theta}) = \theta$, where $\theta \in (-\pi, \pi]$. The angular, or geodesic, distance between z_1 and $z_2 \in \mathbb{C}_1$ is

$$d_L(z_1, z_2) = |\arg(z_1\bar{z}_2)|. \quad (6)$$

For later reference, we plot the extended angular distance function in Fig. 1.

Recall that if $jk \in E_g$, then the edge measurement is correct, that is, $z_{jk} = z_{jk}^*$, where $z_{jk}^* := z_j^* \bar{z}_k^*$ is defined analogously to (1). For $jk \in E_b$, the measurement z_{jk} is assumed to be an arbitrary element of $\mathbb{C}_1$. From the mea-

surements z_{jk}^* , $jk \in E_g$, z^* is only identified up to a global rotation, due to the ambiguity that $z_j^* \bar{z}_k^* = z_j^* y y \bar{z}_k^*$, $y \in \mathbb{C}_1$, and so $z^* y$ generates the same pairwise measurements as z^* .

To deal with this ambiguity, the following function will be used to demonstrate convergence of a sequence to z^* :

$$\delta(z) = \max_{jk \in E} d_L(\bar{z}_j^* z_j, \bar{z}_k^* z_k). \quad (7)$$

We call the products $\bar{z}_j^* z_j$ the *unrotated estimates*, and (7) measures the maximum distance between unrotated estimates. Notice that $\delta(z) = 0 \iff z = z^* y$ for some rotation $y \in \mathbb{C}_1$. Therefore, convergence of $\delta(z)$ to zero indicates convergence of z to z^* , and an algorithm exactly recovers z^* iff $\delta(z) \rightarrow 0$.

3.2 Trimmed Averaging Synchronization

A natural way to solve the rotation synchronization problem involves energy minimization. The simplest strategy (Govindu, 2001; Martinec & Pajdla, 2007) attempts to minimize

$$\min_{z \in \mathbb{C}_1^n} F_L(z) := \sum_{jk \in E} d_L^2(z_j, z_{jk} z_k). \quad (8)$$

A coordinate descent strategy to solve (8) involves updating z_j by solving the averaging problem

$$\min_{z \in \mathbb{C}_1} \sum_{k \in E^j} d_L^2(z, z_{jk} z_k). \quad (9)$$

Applying this sequentially over the indices $j = 1, \dots, n$ results in the multiple rotation averaging (MRA) discussed in more detail in Hartley et al. (2013). While such coordinate descent strategies generally lack coordination across all objects like global synchronization methods, they lead to algorithms that are more memory efficient and that can be decentralized easily.

A simple way to robustify (9) is to select the average of all points that fall within a trimmed set, which results in a trimmed averaging procedure. To account for the 1-dimensional manifold structure of SO(2), we propose to do this trimming in the tangent space, which yields the TAS algorithm. An illustration of one trimmed averaging step is given in Fig. 2. Here, at the point $z_j = z = i$, the observations $z_{jk} z_k$, $k \in E^j$ are given in blue and projected to the tangent space. The trimmed average update would take the average of the green points and ignore the red points.

More concretely, for a discrete $\mathcal{X} \subset \mathbb{R}$ and a fraction $0 < p < 1$, we write the p th quantile of $\mathcal{X}$ by $\mathcal{X}_p$. It is convenient to define the trimming operator

$$\mathcal{T}_\tau \mathcal{X} = \{x \in \mathcal{X} : \mathcal{X}_\tau \leq x \leq \mathcal{X}_{1-\tau}\}. \quad (10)$$




Fig. 2 Illustration of the TAS algorithm at a fixed step and a fixed node j . The measurement is $z_j = z = i$. After projecting into the tangent space, the outermost points in red are filtered, and the green points are averaged. This trimmed average is then projected back to the manifold (Color figure online)

We also denote the average of a dataset $\mathcal{X} \subset \mathbb{R}$ by $\text{ave}(\mathcal{X})$. That is, $\text{ave}(\mathcal{X}) = (\sum_{x \in \mathcal{X}} x) / \#(\mathcal{X})$.

For clarity, we give the TAS algorithm in Algorithm 1. To allow for damping of the updates, we include the step-size parameter $\eta \in (0, 1]$. When $\eta < 1$, we refer to the algorithm as Damped TAS or DTAS for short.

Algorithm 1 Damped Trimmed Averaging Synchronization

Require: $z(0)$, number of iterations T , damping parameter $\eta \in (0, 1]$, trimming parameter τ
for $t = 1, \dots, T$ **do**
 $j = t \bmod n$
 $z_j(t+1) \leftarrow \text{Exp}_{z_j(t+1)} \left[\eta \cdot \text{ave} \left(\mathcal{T}_\tau \left\{ \text{Log}_{z_j(t+1)} (z_{jk} z_k(t)) : k \in E^j \right\} \right) \right]$
 $z_k(t+1) \leftarrow z_k(t), k \neq j$
end for
return $z(T)$

Due to the simplified geometry of $\text{SO}(2)$, we will show in the following that using this trimmed rotation averaging scheme converges to the underlying solution linearly when the percentage of outliers is at most $\alpha_0 < 1/4$ when G is fully connected. In the case where $(G, [n])$ is not fully connected, the result is a corollary of our later Theorem 6 under a connectedness assumption on $(G, [n])$. We note that this fraction is similar to the one given in Lerman and Shi (2020), although there the bound is formulated for corrupted triangles in the graph.

3.3 Recovery Guarantees for DTAS

We begin by discussing the assumptions that will make a synchronization problem tractable for TAS. The first assumption we require is a good initialization, which is common in the analysis of such nonconvex methods.

Assumption 1 The initial set of rotations $z(0) \in \mathbb{C}_1^n$ lies within a $\pi/2$ -neighborhood of z^* : that is, there exists a $w \in \mathbb{C}_1$ such that

$$d_{\angle}(\overline{z_j^*} z_j, w) < \pi/2, \quad j = 1, \dots, n. \quad (11)$$

Note that this is equivalent to the assumption that $\delta(z) < \pi$.

While the corrupted edges and measurements are arbitrary, we require an assumption on the underlying graph $(G, [n])$. It essentially requires that the graph is sufficiently well connected.

Assumption 2 (ζ -Well-connectedness condition) For a fixed $\zeta \in (0, 1]$, for any $J \subset [n]$ such that $\#(J) \leq n/2$, there exists an index $j \in J$ such that

$$\left(\frac{2}{\zeta} - 1 \right) \# \left[E^j \cap ([n] \setminus J) \right] > \# \left[E^j \cap J \right]. \quad (12)$$

In words, this assumption requires that inside any set of at most $n/2$ nodes, there is a node that is connected to a significant number of nodes outside this set. The condition in this assumption is equivalent to requiring that

$$\# \left[E^j \cap J \right] < \left(1 - \frac{\zeta}{2} \right) n_j. \quad (13)$$

We include a discussion of this condition and its connection with random graphs, conductance, and expanders later in Sect. 4.5.

For the case of Assumption 2 with $\zeta = 1$, we just call the graph well-connected. While fully connected graphs satisfy this condition, there exist many more examples of graphs satisfying it with $\zeta = 1$ as well, and we give some examples of some simple graphs that meet this assumption in Fig. 3.




The following theorem gives the main recovery result for the DTAS algorithm. While the algorithm converges linearly, the rate we derive depends on n and is worst-case. In the few simulations we have run, the algorithm seems to converge at a faster rate that merits more study. Also, a more complicated proof may yield linear convergence in the general case of well-connected G , but for sake of brevity, we only prove it for the fully connected case.

Theorem 3 Suppose that $\alpha_0 < \zeta/4$, Assumption 1 holds, G satisfies Assumption 2 with parameter ζ , and $[z(t)]_{t \in \mathbb{N}}$ is the sequence generated by DTAS, for $\eta \in (0, 1)$ and $\tau = \zeta/4$. Then, $\delta(z(t)) \rightarrow 0$, and the algorithm exactly recovers z^* . Furthermore, in the case where G is fully connected and $\eta \in ((n-1)/(4n), (n-1)/n)$, the DTAS algorithm linearly converges to z^* .

Proof (Theorem 3) The proof of convergence under well-connectedness follows from the fact that, when updating index j at iteration t , the selection rule defined by choosing the trimmed average yields a point in the interior of $\mathcal{T}_\tau(\{\text{Log}_{z_j(t)}(z_{jk} z_k(t)) : k \in E^j\})$. The proof then follows from the proof of Theorem 6.

To see linear convergence in the fully connected case, we prove that all unrotated estimates $\overline{z_j^*} z_j(t)$ contract during

Fig. 3 Examples of graphs that satisfy the well-connectedness condition for $n = 4, 5$ and 6 . In each of these graphs, (12) is satisfied with $\zeta = 1$: that is, all subsets J of size at most $n/2$, there exists a node $j \in J$ such that $\#(E^j \cap ([n] \setminus J)) > \#(E^j \cap J)$



each pass over the dataset. Denote

$$\delta_j = \theta_j(t) - \theta_j^* \in [-\delta(z(t))/2, \delta(z(t))/2]. \quad (14)$$

These are the translation of the unrotated estimates to the angular coordinates of the points $z_1(t), \dots, z_n(t)$. Also, define the sets

$$\begin{aligned} I_+(t) &= \{k : \arg(\overline{z_j^*} z_k(t)) > 0\}, \\ I_-(t) &= \{k : \arg(\overline{z_j^*} z_k(t)) \leq 0\}. \end{aligned} \quad (15)$$

In this proof, we will write $\delta = \delta(z(t))$ as a shorthand. Notice that we must have

$$\min(\#I_+(t), \#I_-(t)) \leq n/2,$$

unless $I_+(t) = I_-(t) = [n]$, in which case $z(t) = z^*w$ for some $w \in \mathbb{C}_1$ and $z(t)$ recovers z^* .

For the update with respect to index j , all good pairwise measurements must lie in $-\delta_j + [-\delta/2, \delta/2]$. Since there are at least $3n_j/4$ good measurements, all trimmed points must lie in this interval as well. Therefore, for all $j \in I_-(t)$, after updating we have

$$\overline{z_j^*} z_j(t) \in \exp[i[-\delta/2, \eta\delta/2]]. \quad (16)$$

Using this fact, we will now show that the indices in $I_+(t)$ must move inwards. Indeed, since $\#I_+(t) \leq n/2$, we must have $\#(E_g^j \cap I_-(t)) \geq 1$ for all $j \in I_+(t)$. Therefore, for each trimmed mean for $j \in I_+(t)$, we have

$$\begin{aligned} \text{ave} \left(\mathcal{T}_{0.25} \left(\left\{ \text{Log}_{z_j(t)}(z_{jk} z_k(t)) : k \neq j \right\} \right) \right) \\ \leq \frac{2}{n-1} \left[\eta \frac{\delta}{2} - \delta_j + \left(\frac{n-1}{2} - 1 \right) \left(\frac{\delta}{2} - \delta_j \right) \right] \\ \leq \left(\frac{n-3}{n-1} - \frac{2}{n-1} \eta \right) \cdot \frac{\delta}{2} - \delta_j. \end{aligned} \quad (17)$$

Thus,

$$\eta \text{ave} \left(\mathcal{T}_{0.25} \left(\left\{ \text{Log}_{z_j(t)}(z_{jk} z_k(t)) : k \in E^j \right\} \right) \right) + \delta_j$$

$$\begin{aligned} &\leq \eta \left(\frac{n-3-2\eta}{n-1} \right) \cdot \frac{\delta}{2} + (1-\eta)\delta_j \\ &\leq \eta \left(\frac{n-3}{n-1} \right) \cdot \frac{\delta}{2} + (1-\eta)\frac{\delta}{2} \\ &= \frac{\delta}{2} \left(1 - \eta \left[\frac{2}{n-1} \right] \right). \end{aligned} \quad (18)$$

After the coordinate update, we have that for all $j \in I_+(t)$,

$$\overline{z_j^*} z_j(t) \in \exp \left(i \left[-\frac{\delta}{2}, \left(1 - \eta \left(\frac{2}{n-1} \right) \frac{\delta}{2} \right) \right] \right). \quad (19)$$

After repeating this argument for all j over the course of an epoch, or pass over all indices $j = 1, \dots, n$, this yields that

$$\overline{z_j^*} z_j(t+1) \in \exp \left(i \left[-\frac{\delta}{2}, \left(1 - \eta \left(\frac{2}{n-1} \right) \frac{\delta}{2} \right) \right] \right),$$

as long as $\eta < (n-1)/(n+1)$. The width of this interval is $(n-1-\eta)\delta(z(t))/(n-1)$, which yields the desired result. $\square$

4 Robust Synchronization over $\text{SO}(D)$

We now move on to our more general synchronization algorithm and result over $\text{SO}(D)$. This section presents a novel algorithm for robust synchronization over the rotation group, $\text{SO}(D)$. We assume a fixed observation graph G that encodes which pairwise rotations we observe. The pairwise rotations are written as $\mathbf{R}_{jk} \in \text{SO}(D)$, where the good edges $jk \in E_g$ have the associated observation $\mathbf{R}_j^* \mathbf{R}_k^{*\top}$, and the bad edges are arbitrarily chosen from G and have arbitrarily corrupted measurements.

To proceed, we must make clear our goal for the synchronization problem, since there is a well known ambiguity, similar to the one encountered for $\text{SO}(2)$ synchronization in Sect. 3.1—we can only recover $\mathbf{R}^*$ up to right multiplication by an element of $\text{SO}(D)$. This is because, after this multiplication, one arrives at the same pairwise measurements in (1).

This is a form of rotational symmetry in the nonconvex problem, which may be leveraged to develop tractable nonconvex programs (Zhang et al., 2020). Exactly recovering the ground truth measurements $(\mathbf{R}^*) = (\mathbf{R}_1^*, \dots, \mathbf{R}_n^*) \in \text{SO}(D)^n$ up to right multiplication by $\mathbf{S} \in \text{SO}(D)$ is equivalent to finding a set of rotations $(\mathbf{R}) = (\mathbf{R}_1, \dots, \mathbf{R}_n)$ such that

$$\mathbf{R}_1^{*\top} \mathbf{R}_1 = \dots = \mathbf{R}_n^{*\top} \mathbf{R}_n = \mathbf{S}, \quad (20)$$

for some $\mathbf{S} \in \text{SO}(D)$. We refer to the set of rotations $\mathbf{R}_j^{*\top} \mathbf{R}_j$, $j = 1, \dots, n$, as *unrotated estimates* since, when $(\mathbf{R}) = (\mathbf{R}^* \mathbf{S})$, they reveal the normalization factor that multiplies each element of $(\mathbf{R}^*)$ from the right. One could extend this discussion to the case of approximate recovery by requiring that the unrotated estimates are approximately equal.

With our goal now in mind, we begin in Sect. 4.1 by discussing the geometry of the manifold $\text{SO}(D)$ and presenting some basic geometric results that will be used in our main theorem. Following this, Sect. 4.2 reviews the concept of half-space depth from robust statistics, which will be the core tool that we use to construct our algorithm. In Sect. 4.3, we give outline the DDS algorithm. Then, in Sect. 4.4, we give the theoretical guarantees that constitute the main innovations of this work. We finish in Sect. 4.5 with a discussion of the assumptions we make in our theorem.

4.1 The Manifold Structure of $\text{SO}(D)$

The rotation synchronization problem is obviously a robust recovery problem on the product Riemannian manifold $\text{SO}(D)^n$. Therefore, in the following, we freely use concepts from Riemannian geometry, and specifically those concepts related to the geometry of $\text{SO}(D)$.

4.1.1 Riemannian Geometry of $\text{SO}(D)$

The set of rotations $\text{SO}(D)$ is a $D(D-1)/2$ -dimensional Lie group that has a natural Riemannian structure. The bi-invariant distance metric $d : \text{SO}(D) \times \text{SO}(D) \rightarrow [0, \lfloor \frac{D}{2} \rfloor \pi]$ is given by

$$d(\mathbf{R}_1, \mathbf{R}_2) = \|\log(\mathbf{R}_1 \mathbf{R}_2^\top)\|_F, \quad (21)$$

where $\log$ is the matrix logarithm. The corresponding Lie algebra is $\mathfrak{so}(D)$, the set of $D \times D$ skew-symmetric matrices. The tangent space of $\text{SO}(D)$ at $\mathbf{R} \in \text{SO}(D)$ is

$$T_{\mathbf{R}} \text{SO}(D) = \{\mathbf{A}_s \in \mathbb{R}^{D \times D} : \mathbf{R}^\top \mathbf{A}_s \in \mathfrak{so}(D)\}.$$

Notice that every tangent vector $v \in T_{\mathbf{R}} \text{SO}(D)$ has a corresponding element of $\mathfrak{so}(D)$, which we denote by $v_{\mathfrak{so}(D)}$. The corresponding Riemannian metric (which is an inner product and thus should not be confused with a distance metric)

for $v, w \in T_{\mathbf{R}} \text{SO}(D)$ is given by $\langle v, w \rangle_{\mathbf{R}} = \text{Tr}(v^\top w)/2 = \text{Tr}(v_{\mathfrak{so}(D)}^\top w_{\mathfrak{so}(D)})/2$. Equipped with this metric, $\text{SO}(D)$ is a Riemannian manifold with nonnegative sectional curvature. An open ball with respect to the metric d is written as $B(\mathbf{R}, r)$, where the radius is r and the center is $\mathbf{R}$. Its closure and boundary are $\overline{B}(\mathbf{R}, r)$ and $\partial B(\mathbf{R}, r)$, respectively.

The exponential map is given by

$$\text{Exp}_{\mathbf{R}} : T_{\mathbf{R}} \text{SO}(D) \rightarrow \text{SO}(D), \quad \text{Exp}_{\mathbf{R}}(\mathbf{U}) = \mathbf{R} \exp(\mathbf{R}^\top \mathbf{U}), \quad (22)$$

where $\exp$ is the matrix exponential. The logarithmic map is the inverse of this:

$$\text{Log}_{\mathbf{R}} : \text{SO}(D) \rightarrow T_{\mathbf{R}} \text{SO}(D), \quad \text{Log}_{\mathbf{R}}(\mathbf{S}) = \mathbf{R} \log(\mathbf{R}^\top \mathbf{S}). \quad (23)$$

The geodesic between $\mathbf{R}, \mathbf{S} \in \text{SO}(D)$ is written as $\overrightarrow{\mathbf{R}\mathbf{S}}(t) = \text{Exp}_{\mathbf{R}}(t \text{Log}_{\mathbf{R}}(\mathbf{S}))$, for $t \in [0, 1]$. In the following, we use the notation for a halfspace of $T_{\mathbf{R}} \text{SO}(D)$,

$$\mathcal{H}(\mathbf{R}, v) = \{u \in T_{\mathbf{R}} \text{SO}(D) : \langle u, v \rangle_{\mathbf{R}} > 0, \text{Exp}_{\mathbf{R}}(tu) \text{ is a geodesic for } t \in [0, 1]\}. \quad (24)$$

4.1.2 Local Convexity Properties of $\text{SO}(D)$

We continue by recalling some local convexity properties of manifolds like $\text{SO}(D)$. The following result on the convexity of sufficiently small balls is standard in the literature (Karcher, 1977; Afsari, 2009; Petersen, 2016).

Theorem 4 (Convexity of small balls) *In a closed ball $\overline{B}(\mathbf{C}, r) \subset \text{SO}(D)$ with $r < \pi/2$, the squared distance metric d^2 is strictly convex. This implies, in particular, that for all $\mathbf{R}_0, \mathbf{R}_1 \in \overline{B}(\mathbf{C}, r)$, $\overrightarrow{\mathbf{R}_0 \mathbf{R}_1}(t) \in B(\mathbf{C}, r)$ for all $t \in (0, 1)$.*

Note that the closed ball in the previous theorem has the property that the interior of any nonconstant geodesic lies strictly in the interior of the ball. The following result is also readily apparent. It states that, for boundary points on a sufficiently small ball in $\text{SO}(D)$, all interior directions are contained in a halfspace.

Corollary 1 *Let $\overline{B}$ be a ball on $\text{SO}(D)$ with radius $r < \pi/2$ and $\mathbf{R} \in \partial \overline{B}$. Then, there is a halfspace $\mathcal{H} \subset T_{\mathbf{R}} \text{SO}(D)$ such that $\text{Log}_{\mathbf{R}} \overline{B} \subset \mathcal{H}$.*

An important component of our later theoretical results relies on showing that the radius of the smallest ball containing a set of rotations shrinks. The final lemma of this section states that if a discrete set of rotations $\mathcal{R}$ is contained in a ball, and if half of the ball contains no boundary measurements, then the set $\mathcal{R}$ is actually contained in a ball of

smaller radius. Thus, for such a set of rotations, this gives us a sufficient condition for decreasing the radius of the smallest containing ball.

Lemma 1 Let $B = B(\mathbf{C}, r) \subset SO(D)$ and $\mathcal{R} \subset \overline{B}$ a finite set of rotations. Suppose that there exists $v \in T_{\mathbf{C}} SO(D)$ such that

$$\text{Exp}_{\mathbf{C}}[\overline{\mathcal{H}(\mathbf{C}, -v)}] \cap \partial B \cap \mathcal{R} = \emptyset.$$

Then, $\mathcal{R}$ is contained in a ball with radius less than r .

Proof (Lemma 1) Let $c(t)$ be the geodesic $\text{Exp}_{\mathbf{C}}(tv)$. We claim that, for t sufficiently small, $\mathcal{R} \subset B(c(t), r)$, which is an open ball. Heuristically, one should expect this to be true, since the closed halfspace $\overline{\mathcal{H}(\mathbf{C}, -v)}$ contains no boundary points, and so moving the center a small amount in the v direction keeps all points within the ball.

By the first order approximation to $d^2(c(t), \mathbf{R})$ and since $\angle(v, \text{Log}_{\mathbf{C}} \mathbf{R}) < \pi/2$, we have

$$d(c(t), \mathbf{R})^2 < d(c(0), \mathbf{R})^2, \quad \forall \mathbf{R} \in \mathcal{R} \cap \text{Exp}_{\mathbf{C}}(\overline{\mathcal{H}(\mathbf{C}, v)}), \quad (25)$$

for t sufficiently small. On the other hand, since

$$\text{Exp}_{\mathbf{C}}[\overline{\mathcal{H}(\mathbf{C}, -v)}] \cap \partial B \cap \mathcal{R} \subset B(\mathbf{C}, r),$$

the distance to $\mathbf{C}$ over all $\mathbf{R} \in \text{Exp}_{\mathbf{C}}[\overline{\mathcal{H}(\mathbf{C}, -v)}] \cap \partial B \cap \mathcal{R}$ is bounded away from r . By continuity of $d(c(t), \mathbf{S})$ for all $\mathbf{R} \in \mathcal{R}$, this implies that there is an ϵ such that

$$d(c(t), \mathbf{R}) < r, \quad \forall t \in (0, \epsilon). \quad (26)$$

Putting (25) and (26) together implies that a small shift of the ball results in a new center such that $\max_{\mathbf{R} \in \mathcal{R}} d(c(t), \mathbf{R}) < r$. In turn, this means that all points of $\mathcal{R}$ lie in a ball $B(c(t), r')$ with $r' < r$. $\square$

4.2 Tukey Depth and its Properties

We will use the concept of *Tukey depth* to determine descent directions on the manifold $SO(D)^n$, although other notions of depth could potentially be used as well (see Ch. 58 of Toth et al. (2017) for a discussion of different notions of depth). In Euclidean space, the Tukey depth of a point $\mathbf{x} \in \mathbb{R}^D$ in a dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^D$ is given by

$$\text{depth}(\mathbf{x}, \mathcal{X}) = \min_{\mathbf{u} \in S^{D-1}} \#\{\mathbf{x}_i \in \mathcal{X} : \mathbf{u}^\top (\mathbf{x}_i - \mathbf{x}) \geq 0\} \quad (27)$$

The depth is therefore the minimum number of points contained in any halfspace that has $\mathbf{x}$ in its separating hyperplane. A natural robust estimator is then the point of maximum

depth, which is also called the Tukey median. The β -depth level set for $\beta \in [0, 1]$ is defined by

$$\mathcal{D}_\beta(\mathcal{X}) = \{\mathbf{y} \in \mathbb{R}^D : \text{depth}(\mathbf{y}, \mathcal{X}) \geq \beta \#(\mathcal{X})\}. \quad (28)$$

This level set is convex and compact, and its boundary is made up of hyperplanes defined by sets of D points (Liu et al., 2019). This function will be used in the construction of our algorithm.

As an example, consider the 1-dimensional dataset $\mathcal{X} = \{x_i\}$. Here, the formulation of depth is quite simple:

$$\text{depth}(x, \mathcal{X}) = \min(\#\{x_i \leq x\}, \#\{x_i \geq x\}). \quad (29)$$

With this in mind, the β -depth level set is $\mathcal{D}_\beta(\mathcal{X}) = [x_{(\lceil \beta n \rceil)}, x_{(\lfloor (1-\beta)n \rfloor)}]$, where $x_{(i)}$ denotes the i th order statistic. The Tukey median in this case is just the median.

We recall the following theorem, which bounds the maximum possible depth within a general dataset. Notice that, in particular, this guarantees that the depth level set $\mathcal{D}_\beta(\mathcal{X})$ is nonempty for all $\beta \leq 1/(D+1)$.

Proposition 1 (Rado, 1946) Suppose that $\mathcal{X}$ is a set of n points. Then, the maximum depth in $\mathcal{X}$ is bounded below by $\lceil n/(D+1) \rceil$.

A particularly useful property of Tukey depth is that it is *affine equivariant*, that is, it is stable under affine transformations.

Lemma 2 (Donoho & Gasko, 1992)

$$\text{depth}(\mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{A}\mathcal{X} + \mathbf{b}) = \text{depth}(\mathbf{x}, \mathcal{X}).$$

This implies, in particular, that things behave nicely if we change the inner product on $\mathbb{R}^D$.

$$\begin{aligned} \min_{\mathbf{u} \in S^{D-1}} \#\{\mathbf{x}_i \in \mathcal{X} : \mathbf{u}^\top \mathbf{A}(\mathbf{x}_i - \mathbf{x}) \geq 0\} &= \text{depth}(\mathbf{A}\mathbf{x}, \mathbf{A}\mathcal{X}) \\ &= \text{depth}(\mathbf{x}, \mathcal{X}). \end{aligned} \quad (30)$$

With this property, we can map between the depth regions in $(\mathbb{R}^D, \langle \cdot, \cdot \rangle)$ and $(\mathbb{R}^D, \langle \cdot, \mathbf{A} \cdot \rangle)$ by

$$\mathbf{A}\mathcal{D}_\beta(\mathcal{X}) = \mathcal{D}_\beta(\mathbf{A}\mathcal{X}). \quad (31)$$

This affine equivariance implies, in particular, that Proposition 1 extends to datasets in tangent spaces of manifolds.

We finish with a simple lemma on depth level sets that will be the key to the robustness guarantee for our later algorithm. In simple words, this lemma guarantees sufficient conditions for a depth level set to contain a nonzero value in a dataset containing many zeros. In the following, datasets are represented by multisets and may contain duplicate points. For a

halfspace $H(\mathbf{0}, \mathbf{v}) := \{\mathbf{x} \in \mathbb{R}^D : \mathbf{v}^\top \mathbf{x} \geq 0\} \subset \mathbb{R}^D$, we write its separating hyperplane as $L(\mathbf{0}, \mathbf{v})$.

Lemma 3 Suppose that we have a dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathbb{R}^D$ and a subset $\mathcal{Y} \subset \mathcal{X}$ that satisfies the properties i) $\#(\mathcal{Y}) > n - n(\zeta/(2D+2))$, ii) There exists closed halfspace $\overline{H(\mathbf{0}, \mathbf{v})} \subset \mathbb{R}^D$ such that $\overline{H(\mathbf{0}, \mathbf{v})} \supset \mathcal{Y}$, and iii) the only points of $\mathcal{Y}$ in $L(\mathbf{0}, \mathbf{v})$ are $\mathbf{0}$. Then, $\mathcal{D}_{\zeta/(2D+2)}(\mathcal{X}) \subset \text{conv}(\mathcal{Y}) \subset (H(\mathbf{0}, \mathbf{v}) \cup \{\mathbf{0}\})$. Beyond this, if $\#(\mathcal{Y} \cap L(\mathbf{0}, \mathbf{v})) < (1 - \zeta/2)n$, then $(\mathcal{D}_{\zeta/(2D+2)}(\mathcal{X})) \cap H(\mathbf{0}, \mathbf{v}) \neq \emptyset$.

Proof (Lemma 3) It is obvious by the properties of depth that $\mathcal{D}_{\zeta/(2D+2)}(\mathcal{X}) \subset \text{conv}(\mathcal{Y}) \subset (H(\mathbf{0}, \mathbf{v}) \cup \{\mathbf{0}\})$, since any point on the boundary of $\text{conv}(\mathcal{Y})$ has depth less than $\zeta/(2D+2)$. By Proposition 1, there is a point of depth at least $\zeta/(2D+2)$, and so $\mathcal{D}_{\zeta/(2D+2)}(\mathcal{X})$ is nonempty.

Suppose that less than $(1 - \zeta/2)n$ points in $\mathcal{Y}$ are zero, that is, $\#(\mathcal{Y} \cap L(\mathbf{0}, \mathbf{v})) < (1 - \zeta/2)n$. We claim that $\mathcal{D}_{\zeta/(2D+2)}(\mathcal{X}) \cap H(\mathbf{0}, \mathbf{v}) \neq \emptyset$. Define the auxiliary set $\mathcal{Z} = \mathcal{X} \setminus (\mathcal{Y} \cap L(\mathbf{0}, \mathbf{v}))$, that is, $\mathcal{Z}$ removes the $\mathbf{0}$ values in $\mathcal{Y}$ from $\mathcal{X}$. Since $\#(\mathcal{Y} \cap L(\mathbf{0}, \mathbf{v})) < (1 - \zeta/2)n$, we have that $m = \#(\mathcal{Z}) \geq (\zeta/2)n$. Within $\mathcal{Z}$, there is a point $\hat{\mathbf{z}}$ of depth at least $m/(D+1)$. Further, since $n \leq 2m/\zeta$ and $\#(\mathcal{X} \cap H(\mathbf{0}, -\mathbf{v})) < n\zeta/(2D+2)$, we have

$$\begin{aligned} \#(H(\mathbf{0}, \mathbf{v}) \cap \mathcal{Z}) &> m - \frac{n\zeta}{2D+2} \\ &\geq m - \frac{m}{D+1} \\ &= m \frac{D}{D+1}, \end{aligned}$$

which implies that $\hat{\mathbf{z}}$ must lie in $H(\mathbf{0}, \mathbf{v})$. On the other hand, since $m \geq \zeta n/2$, we have that

$$\text{depth}(\hat{\mathbf{z}}, \mathcal{X}) \geq \frac{m}{D+1} \geq \frac{\zeta n}{2(D+1)},$$

which means that $\hat{\mathbf{z}}$ is a point of depth at least $\zeta/(2D+2)$ in $\mathcal{X}$. $\square$

4.3 Depth Descent Synchronization

We now use the results of the previous sections to derive the DDS algorithm.

4.3.1 The General DDS Algorithm

We assume a selection rule $\mathcal{S}_R$ on convex, compact subsets of $T_R \text{SO}(D)$ for all $R \in \text{SO}(D)$. In particular, for our theorem, we assume that this selection rule chooses a nonzero point from the convex set if possible, and otherwise outputs zero. To select this point, one selection rule could be to take a point uniformly at random, or another could be to take the

center of mass. As our theorem makes clear, the choice of this selection rule does not affect the exact recovery result, but it may change convergence rates. Since we do not give quantitative convergence rates in this work, we leave this choice as arbitrary.

In any case, given a selection rule $\mathcal{S}_R$ over such subsets of $T_R \text{SO}(d)$, suppose our estimated rotations at time t are $(R(t))$. Our update direction at this time for index $j = t \bmod n$ is defined by

$$\mathbf{v}_j(t) = \mathcal{S}_R(\mathcal{D}_\beta(\{\text{Log}_{R_j(t)} R_{jk} R_k(t) : k \in E^j\})). \quad (32)$$

In words, the direction $\mathbf{v}_j(t)$ is a direction in the tangent space at $R_j(t)$ that is sufficiently deep with respect to the neighbor measurements $\{\text{Log}_{R_j(t)} R_{jk} R_k(t) : k \in E^j\}$. Given $\mathbf{v}_j(t)$, the algorithm updates

$$\begin{aligned} j &= t \bmod n, \quad R_j(t+1) = \text{Exp}_{R_j(t)}(\eta(t)\mathbf{v}_j(t)), \\ R_k(t+1) &= R_k(t) \text{ for } k \in [n] \setminus \{j\}. \end{aligned} \quad (33)$$

for a chosen step size $\eta(t) \in (0, \eta^*(D)]$. Our theory below restricts this step size according to Theorem 4.2 of Afsari et al. (2013): in the case of $D = 2$ or 3 , one can take $\eta^*(D) = 1$, while for $D > 3$ the upper bound is more restrictive (the reader can consult the discussion in Afsari et al. (2013) for a more thorough discussion of the bound). For sake of clarity, we write the full DDS algorithm in Algorithm 2.

Algorithm 2 Depth Descent Synchronization for $\text{SO}(D)$

Require: $R(0)$, number of iterations T , selection rule $\mathcal{S}$, $\eta \in (0, \eta^*(D)]$ step size, β depth parameter
for $t = 1, \dots, T$ **do**
 $j = t \bmod n$
 $\mathbf{v}_j(t) = \mathcal{S}_R(\mathcal{D}_\beta(\{\text{Log}_{R_j(t)} R_{jk} R_k(t) : k \in E^j\}))$
 $R_j(t+1) \leftarrow \text{Exp}_{R_j(t)}(\eta \mathbf{v}_j(t))$
 $R_k(t+1) \leftarrow R_k(t), k \neq j$
end for
return $R(T)$

One benefit of DDS is that there is no need to tune the step size, which is due to the local convexity properties of the manifold. As we outline in our main theorem, the step size can be directly selected from the guidance of Afsari et al. (2013).

As for computational complexity, at least in low-dimensions, depth regions can be calculated efficiently for small datasets (Liu et al., 2019). In particular, the most straightforward algorithm involves an exhaustive search over hyperplanes spanned by D -subsets of $\mathcal{X}$ that cut off βn points in $\mathcal{X}$. The time complexity of this method for $\mathbb{R}^3$ is $O(n^3 \log(n))$, and so could be used for moderately sized datasets. Translating this to our problem in $\text{SO}(3)$, the time complexity for updating $j \in [n]$ is $O(n_j^3 \log(n_j))$, and so we see that this method is efficient

for sparser graphs. In an Erdős-Rényi model, the complexity to update all n rotations is $O(n^4 p^3)$ in expectation, where p is the Erdős-Rényi parameter.

As discussed in the introduction, the time complexity of the DDS algorithm is not necessarily more efficient than that of Lerman and Shi (2020). Indeed, the complexity of their message-passing algorithm is $O(n^3)$ for a single update to all rotations, while for our method it is $O(\sum_j n_j^3 \log(n_j))$ for $\text{SO}(3)$ (and much larger for higher dimensions). Therefore, DDS has better complexity for sparse graphs, while Lerman and Shi (2020) has better complexity for dense graphs. On the other hand, our complexity is uniformly better in $\text{SO}(2)$, where depth contours can be easily found in $O(n \log(n))$ by sorting, and it thus takes $O(n^2 \log(n))$ time to update all nodes. Finally, we also note that our method has better scaling in terms of memory usage: the multiple rotation averaging scheme takes $O(n_j)$ memory while the message-passing scheme takes $O(n^3)$.

4.3.2 The Approximate DDS Algorithm

To make the DDS algorithm more computationally efficient, we employ a few strategies to develop an approximate DDS algorithm.

First, so that we do not need to resort to computing full depth contours, we instead take the average of the deepest points in the set $\{\text{Log}_{\mathbf{R}_j(t)} \mathbf{R}_{jk} \mathbf{R}_k(t) : k \in E^j\}$ as our update direction at each iteration.

Second, to avoid computation of the full depth of every point in this set, we instead use an approximation of depth based on sampling. Suppose that we wish to approximate the depth of the vectors in $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_n\} \subset \mathbb{R}^D$ with respect to $\mathcal{Y}$. We can sample a set of vectors $\mathbf{u}_1, \dots, \mathbf{u}_m \in S^{D-1}$. Then, for each $\mathbf{y}_i \in \mathcal{Y}$, the approximation of depth, $\widetilde{\text{depth}}$, is

$$\begin{aligned} \widetilde{\text{depth}}(\mathbf{y}_i, \mathcal{Y}) = \min_{j \in [m]} \min \Big(\#(\{\mathbf{y}_k \in \mathcal{Y} : \mathbf{y}_k^\top \mathbf{u}_j \geq \mathbf{y}_i^\top \mathbf{u}_j\}), \\ \#(\{\mathbf{y}_k \in \mathcal{Y} : \mathbf{y}_k^\top \mathbf{u}_j \leq \mathbf{y}_i^\top \mathbf{u}_j\}) \Big). \end{aligned} \quad (34)$$

Notice that $\widetilde{\text{depth}}$ replaces the minimum over $\mathbf{u} \in S^{D-1}$ in (27) by the minimum over the discrete set of vectors $\{\mathbf{u}_1, \dots, \mathbf{u}_m\}$. For $\mathbb{R}^3$ (which corresponds to the tangent space for $\text{SO}(3)$), computation of the full depth for all points would take $O(n_j^3)$ time, where for each $\mathbf{x}_i$ one would need to search over all planes defined by triplets $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k$, for distinct i, j, k . On the other hand, the computation of the approximate depth using (34) takes $O(n_j^2 m)$ and can be more efficiently implemented due to the fact that the $\mathbf{u}_1, \dots, \mathbf{u}_m$ are shared between all $\mathbf{y}_i$.

The approximate DDS algorithm is given in Algorithm 3. Here, we use the notation $S_{\mathbf{R}_j(t)} \text{SO}(D)$ for the set of all unit vectors in $T_{\mathbf{R}_j} \text{SO}(D)$.

Algorithm 3 Approximate Depth Descent Synchronization for $\text{SO}(D)$

Require: $\mathbf{R}(0)$, number of iterations T , selection rule $\mathcal{S}$, $\eta \in (0, 1)$ step size, m : number of depth vectors
for $t = 1, \dots, T$ **do**
 $j = t \bmod n$
 $\mathcal{Y}(t) = \{\text{Log}_{\mathbf{R}_j(t)} \mathbf{R}_{jk} \mathbf{R}_k(t) : k \in E^j\}$
 $\mathbf{u}_1, \dots, \mathbf{u}_m \stackrel{i.i.d.}{\sim} \text{Unif}(S_{\mathbf{R}_j(t)} \text{SO}(D))$
 $\mathbf{v}_j(t) = \arg\max_{\mathbf{y} \in \mathcal{Y}(t)} \left[\min_{\mathbf{u}_i} \left(\min(\#\{\mathbf{y}' \in \mathcal{Y}(t) : \mathbf{u}_i^\top (\mathbf{y}' - \mathbf{y}) \geq 0\}, \#\{\mathbf{y}' \in \mathcal{Y}(t) : \mathbf{u}_i^\top (\mathbf{y}' - \mathbf{y}) \leq 0\}) \right) \right]$
 $\mathbf{R}_j(t+1) \leftarrow \text{Exp}_{\mathbf{R}_j(t)}(\eta \mathbf{v}_j(t))$
 $\mathbf{R}_k(t+1) \leftarrow \mathbf{R}_k(t), k \neq j$
end for
return $\mathbf{R}(T)$

4.4 Exact Recovery for DDS

For many nonconvex methods, good initialization is quite important (Li et al., 2019; Maunu et al., 2019; Qu et al., 2019; Chi et al., 2019; Qu et al., 2020). We only obtain a recovery result for DDS if we can initialize in a suitable neighborhood of $\mathbf{R}^*$.

Definition 2 A set of n rotations $(\mathbf{R})$ lies within a ρ -neighborhood of $(\mathbf{R}^*)$ if the unrotated estimates $\mathbf{R}_j^{*\top} \mathbf{R}_j$ all lie in a ball of radius ρ . That is, there exists a $\mathbf{C} \in \text{SO}(D)$ such that

$$d(\mathbf{R}_j^{*\top} \mathbf{R}_j, \mathbf{C}) < \rho, \quad \forall j = 1, \dots, n.$$

With this terminology, our main assumption is that we can initialize in a $\pi/2$ -neighborhood of $(\mathbf{R}^*)$. This is the analog of Assumption 1 for $\text{SO}(D)$.

Assumption 5 The initial set of rotations $(\mathbf{R}(0))$ for our algorithm lies within a $\pi/2$ -neighborhood of $(\mathbf{R}^*)$.

As we discuss in Sect. 4.5, we believe that this assumption is not so restrictive, and we later give some intuition for how this might occur in a real scenario.

The following theorem constitutes the main theoretical result of this work. It states that, with proper initialization and well-connectedness, the recovery threshold of Algorithm 2 is $1/(D(D-1)+2)$. As examples when $\zeta = 1$, in the case of $\text{SO}(2)$, Theorem 6 yields a corruption threshold of $\alpha_0 < 1/4$. In the case of $\text{SO}(3)$, Theorem 6 yields a corruption level of $\alpha_0 < 1/8$.

Theorem 6 Suppose that $\alpha_0 < \zeta/(D(D-1)+2)$, Assumptions 5 and 2 hold, and $[(\mathbf{R}(t))]_{t \in \mathbb{N}}$ is generated by (33) with $\beta = \zeta/(D(D-1)+2)$. Further assume in the case of $D = 2, 3$ that $\eta \in (0, 1]$, and in the case of $D > 3$ that η is chosen according to Theorem 4.2 of Afsari et al. (2013). Then, $d(\mathbf{R}_j^* \mathbf{R}_j(t), \mathbf{R}_k^* \mathbf{R}_k(t)) \rightarrow 0$ for all j, k , and the DDS algorithm exactly recovers $(\mathbf{R}^*)$.

Proof (Theorem 6) To aid in the proof, we denote the smallest ball enclosing our unrotated estimates as

$$B(t) := \operatorname{argmin}_{B(C, \rho)} \rho, \text{ s.t. } \mathbf{R}_1^* \mathbf{R}_1(t), \dots, \mathbf{R}_n^* \mathbf{R}_n(t) \in \overline{B(C, \rho)}. \quad (35)$$

The center of $B(t)$ is $C(t)$ and its radius is $r(B(t))$. Our goal will be to show that $r(B(t)) \rightarrow 0$, $t \rightarrow \infty$.

The proof of the theorem is broken into three parts. In the first part, we prove that the sequence $[(\mathbf{R}(t))]_{t \in \mathbb{N}}$ remains in a nested sequence of balls. In the second part, we show that, after sufficiently many iterations, the radius of the smallest enclosing ball must shrink. We finish in the third part by appealing to a general convergence theorem for monotonic algorithms.

Part I: $B(t+1) \subseteq B(t)$: First, we show that at time t , no matter which index is updated, the unrotated estimates remain in $\overline{B(t)}$. This is true at $t = 0$, so assume that it is true at a time t . Let $j = t \bmod n$ and consider the pairwise measurements in the tangent space at $\mathbf{R}_j(t)$: for each $k \in E_g^j$, the corresponding point in the tangent space is given by $\operatorname{Log}_{\mathbf{R}_j(t)}(\mathbf{R}_j^* \mathbf{R}_k^* \mathbf{R}_k(t))$. By assumption, we have that $\mathbf{R}_j^* \mathbf{R}_k^* \mathbf{R}_k(t) \in \mathbf{R}_j^* \overline{B(t)}$ for all k . Since $\alpha_0 < \zeta/(D(D-1)+2)$ and $\beta = \zeta/(D(D-1)+2)$,

$$\begin{aligned} & \mathcal{D}_{\zeta/(2D+2)}(\{\operatorname{Log}_{\mathbf{R}_j(t)}(\mathbf{R}_{jk} \mathbf{R}_k(t)) : k \in E_g^j\}) \\ & \subset \operatorname{conv}(\{\operatorname{Log}_{\mathbf{R}_j(t)}(\mathbf{R}_{jk} \mathbf{R}_k(t)) : k \in E_g^j\}), \end{aligned}$$

since the set in the right-hand side of the display contains more than a $1 - \beta$ fraction of points. We can now apply Theorem 3.7 of Afsari et al. (2013). This follows from the fact that the update direction $\mathbf{v}_j(t)$ is the gradient of the Frechét mean function for a weighted combination of $\{\operatorname{Log}_{\mathbf{R}_j(t)}(\mathbf{R}_{jk} \mathbf{R}_k(t)) : k \in E_g^j\}$ (since it lies in the convex hull of these points). More formally, letting $m = \#(E_g^j)$, since $\mathbf{v}_j(t) \in \operatorname{conv}(\{\operatorname{Log}_{\mathbf{R}_j(t)}(\mathbf{R}_{jk} \mathbf{R}_k(t)) : k \in E_g^j\})$, there exist weights $a_1, \dots, a_m$ such that

$$\begin{aligned} \mathbf{v}_j(t) &= \sum_{i=1}^m a_i \operatorname{Log}_{\mathbf{R}_j(t)} \mathbf{R}_{jk_i} \mathbf{R}_{k_i}(t) \\ &= \sum_{i=1}^m a_i \operatorname{Log}_{\mathbf{R}_j(t)} \mathbf{R}_j^* \mathbf{R}_{k_i}^* \mathbf{R}_{k_i}(t) \end{aligned}$$

$$= -\operatorname{grad} \sum_{i=1}^m a_i d^2(\mathbf{R}_{k_i}^* \mathbf{R}_{k_i}(t), \cdot) \Big|_{\mathbf{R}_j(t)}.$$

Therefore, choosing the step size as in Afsari et al. (2013) implies that $\mathbf{R}_j^* \mathbf{R}_j(t+1) \in \overline{B(t)}$, and further that $\mathbf{R}_j^* \mathbf{R}_j(t+1) \in B(t)$ when $\mathbf{v}_j(t) \neq 0$. In turn, this implies that $\overline{B(t+1)} \subseteq \overline{B(t)}$ and, if $\mathbf{R}_j(t) \in B(t)$, then $\mathbf{R}_j(t+1) \in B(t)$ as well (i.e., interior points cannot move to the boundary).

In the case of $\operatorname{SO}(3)$, Theorem 3.7 of Afsari et al. (2013) tells us that choosing $\eta \in (0, 1]$ suffices. The case of $\operatorname{SO}(D)$ for $D > 3$ is dealt with in a similar way using Theorem 4.2 of Afsari et al. (2013).

Part II: $r(B(s+\Delta_s)) < r(B(s))$: We now must show that after sufficiently many iterations, the radius of $B(t)$ strictly decreases. To this end, fix a time s . At this time, at least one unrotated estimates $\mathbf{R}_j^* \mathbf{R}_j(s)$ must lie on the boundary $\partial B(s)$. For convenience, define the index set $J(s)$ of boundary rotations at time s by

$$J(s) := \{j : \mathbf{R}_j^* \mathbf{R}_j(s) \in \partial B(s)\}.$$

We will show that there exists a $\Delta_s > 0$ such that at some future time $s + \Delta_s$,

$$r(B(s+\Delta_s)) < r(B(s)). \quad (36)$$

To this end, pick a direction w uniformly at random from $T_{C(s)} \operatorname{SO}(D)$ such that $\|w\|_{C(s)} = 1$. This vector separates $T_{C(s)}$ into two halfspaces, and thus partitions $B(s)$ into two halves. One of these halves contains at most $n/2$ points $\operatorname{Log}_{C(s)}(\mathbf{R}_j^* \mathbf{R}_j(s))$, and we will denote the corresponding halfspace of $T_{C(s)} \operatorname{SO}(D)$ by $\mathcal{H}$. Since the direction w is chosen uniformly at random, there are no points on the boundary of this halfspace with probability 1.

Let $K(s)$ denote the set

$$K(s) := \{k : \mathbf{R}_k^* \mathbf{R}_k(s) \in \operatorname{Exp}_{C(s)}(\mathcal{H}) \cap \partial B(s)\},$$

that is, the set of indices corresponding to boundary unrotated estimates at time s . At each time $t = s + m$ for $m > 0$, if none of the unrotated estimates $\mathbf{R}_k^* \mathbf{R}_k(t)$, $k \in K(s)$, lie in $\partial B(s)$, then set $\Delta_s = m$ and we can apply Lemma 1 to yield (36).

Otherwise, by Assumption 2, there is at least one index $k \in K(s)$ such that $\mathbf{R}_k^* \mathbf{R}_k(s)$ is in $\partial B(s)$ and

$$\#(E^k \setminus K(s)) > \#(E^k \cap K(s)).$$

Suppose that we update this index k at time $t = s + m$. We are in a situation where we can apply Lemma 3, with

$$\mathcal{X} = \{\operatorname{Log}_{\mathbf{R}_j(t)} \mathbf{R}_{jk} \mathbf{R}_k(t) : k \in E^j\},$$

$$\mathcal{V} = \{\text{Log}_{\mathbf{R}_j(t)} \mathbf{R}_{jk} \mathbf{R}_k(t) : k \in E_g^j\},$$

which yields that $\mathbf{v}_j(s+m) \in \text{conv}(\mathcal{V})$ and $\mathbf{v}_j(s+m) \neq 0$. Thus, for sufficiently small η , $\mathbf{R}_j(s+m+1) \in B(s)$. Repeating this sequentially for all elements of $K(s)$, there must exist a Δ_s such that

$$\mathbf{R}_j^{\star\top} \mathbf{R}_j(s + \Delta_s) \in B(s), \quad \forall j \in K(s).$$

We are left in a situation where $\mathbf{R}_1^{\star\top} \mathbf{R}_1(s + \Delta_s), \dots, \mathbf{R}_n^{\star\top} \mathbf{R}_n(s + \Delta_s) \in \overline{B(s)}$, and there exists a w such that $\mathcal{H}(C(s), -w)$ contains no boundary points. By appealing to Lemma 1, we know that $(\mathbf{R}(s + \Delta_s))$ lies in a ball of radius smaller than $r(B(s))$, and therefore $r(B(s))$ has a strictly monotonic subsequence.

Part III: Strict monotonicity implies convergence: By Mizera & Volauf (2002), as long as $\beta \leq 1/(D(D-1)/2+1)$, the point to set mapping

$$\mathbf{R}_j(t) \mapsto D_\beta(\{\log_{\mathbf{R}_j(t)}(\mathbf{R}_{jk} \mathbf{R}_k(t)) : k \in E^j\})$$

is non-empty and outer semicontinuous with respect to the empirical measure on

$$\{\text{Log}_{\mathbf{R}_j(t)} \mathbf{R}_{jk} \mathbf{R}_k(t) : k \in E^j\}.$$

Therefore, the associated algorithm (33) is upper semicontinuous in the sense of Theorem 3.1 of Meyer (1976), and we obtain convergence of $\mathbf{R}(t)$ to a fixed point of (33).

We finish by examining fixed points of the algorithm (33). Suppose that $\mathbf{R}$ is a fixed point of this sequence, such that $\mathbf{v}_j$ as defined by (32) is zero.

The fixed point is characterized by $d(\mathbf{R}_j, \mathbf{R}_j^{\star\top} \mathbf{R}_k^{\star\top} \mathbf{R}_k) = 0$ for at least $(1 - \zeta/2)n_j$ measurements k . By ζ -well-connectedness, for all subsets J of size at most $n/2$, there is an index j such that $\# [E^j \cap J] < (1 - \zeta/2)n_j$, because otherwise Lemma 3 would yield a nonzero update direction. Thus, there is an index $k \in E^j \setminus [J]$ such that $d(\mathbf{R}_j, \mathbf{R}_j^{\star\top} \mathbf{R}_k^{\star\top} \mathbf{R}_k) = 0$. This implies that $d(\mathbf{R}_j, \mathbf{R}_j^{\star\top} \mathbf{R}_k^{\star\top} \mathbf{R}_k) = 0$ for all j, k . $\square$

4.5 Discussion of Assumptions

The ζ -well-connectedness condition in Assumption 2 bears some similarity to the notions of conductance and graph expansion. From the perspective of graph theoretical results, we note that a sufficient condition for Assumption 2 with $\zeta = 1$ is for the conductance of the graph to be greater than or equal to $1/2$. This follows from a simple pigeonhole argument. It is unclear if this condition holds for Erdős-Rényi graphs. Indeed, if one uses Cheeger's inequality, one would need the spectral gap to be greater than or equal to 1, but for Erdős-Rényi graphs one only expects this gap to concentrate around 1 in practice (Hoffman et al., 2021).

A sufficient condition for (13) is for the conductance to be bounded below by $\zeta/2$, which can be achieved with high probability for any fixed $\zeta < 1$ by Erdős-Rényi graphs when $p \gtrsim \log(n)/n$ (see, for example, Hoffman et al. (2021)), as well as expander graphs (see, for example, Friedman (2003)).

Our result in Theorem 6 holds with high probability for the uniform corruption model discussed in Sect. 2.6 with $q < \alpha_0 = \zeta/(D(D-1)+2)$ and $p \gtrsim \log(n)/n$. This means that DDS achieves the information theoretically optimal rate with respect to p in this model. Indeed, we can read (4) as $p = \Omega\left(\frac{1}{(1-q)^2} \frac{\log(n)}{n}\right)$.

Assumption 5 requires that we initialize the DDS algorithm so that the unrotated estimates lie in sufficiently small ball. This can be achieved in practice for cameras whose orientations lie close enough together. That is, suppose that all of the rotations $\mathbf{R}_1^{\star\top}, \dots, \mathbf{R}_n^{\star\top}$ lie in $B(S, \rho)$ for all j , for some S and $\rho < \pi/2$. Then, if the initial point for the DDS algorithm is chosen to be $(\mathbf{I}, \dots, \mathbf{I})$, then it is not hard to see that

$$\mathbf{R}_j^{\star\top} \mathbf{I} = \mathbf{R}_j^{\star\top} \in B(S, \rho), \quad \forall j \in [n]. \quad (37)$$

which directly shows that Assumption 5 holds. Notice that $B(S, \pi/2)$ is a large ball that essentially makes up half of the manifold $\text{SO}(3)$, since the distance from any point to its cut-locus is π . For example, if one considers reconstruction of an object from many images taken from points on a sphere that surrounds this object, then our requirement would essentially boil down to needing all of the images being taken from a single hemisphere.

We conjecture that one can weaken this initialization condition to only require that neighboring unrotated estimates are close to each other, but we leave weakening of this assumption to future work.

5 Comparison with Existing Theory

The main works we must compare our theoretical guarantees with are Lerman and Shi (2020) and Huang et al. (2019). For context, we primarily consider the practical cases of $D = 2$ or $D = 3$, which are the cases of interest for us and of high relevance for the computer vision community. We note that our bounds degrade for larger D . While our method does extend to larger D , we do not have the motivation for studying such a problem in the context of computer vision.

In these regimes, our breakdown points for adversarial outliers are fractions of $1/4$ for $D = 2$ and $1/8$ for $D = 3$ corrupted edges per node, assuming the generic well-connectedness condition.

As mentioned in the review of related work, Lerman and Shi (2020) achieve a bound on the ratio of corrupted triangles

per node up to $1/4$. This bound is a bit different from ours, as we instead bound the number of corrupted edges per node. Furthermore, as we mentioned earlier, the method of Lerman and Shi (2020) degrades for sparse graphs, whereas ours elegantly extends to sparse graphs with our generic well-connectedness condition.

As for the second work, during review it was pointed out to us that one can extend the results of Huang et al. (2019) to a bound for arbitrary outliers. This method truncates the graph connection Laplacian at each iteration to filter outliers. A modification of the proof in Huang et al. (2019, Section C.7) demonstrates a best possible bound on the fraction of corrupted edges per node of $1/30$. Furthermore, the upper bound of $1/30$ is under the assumption of a fully connected graph, and it degrades for sparser graphs. The complicated dependencies of the upper bound on the graph connectivity make it hard to estimate in real scenarios. Compared to our results, this bound is much worse for $D = 2$ and $D = 3$, and so more work should be done to see if such a spectral method can achieve better bounds in low-dimensional cases or not. Thus, in the adversarial case for $D = 2$ and $D = 3$, our bounds are state-of-the-art.

6 Empirical Evaluation

We run some synthetic experiments on the DDS and related algorithms

The algorithms we compare with are MRA and L1-MRA (Hartley et al., 2013), IRLS after L1-MRA initialization (Chatterjee & Govindu, 2013), LTS (Huang et al., 2019), CEMP (Lerman & Shi, 2020), and MPLS (Shi & Lerman, 2020). Default parameters of all methods are used. For CEMP, the method computes the corruptions levels first to determine which edges are most corrupted. Then, using these corruption levels, it finds a minimum spanning tree, from which one can fix $R_1 = I$ and then propagate from this to find the other rotations along this tree. For LTS, we implement the truncation step with parameter $\gamma = 0.96$ and run for 40 iterations. The approximate DDS algorithm is run for 40 epochs (or passes over the data, which means we take $T = 40n$) with a step size of $\eta = 0.7$ and number of depth vectors $m = 20$.

We compute the distance between the estimated rotations ($\hat{R}$) and (R^*) by first aligning them by solving

$$S = \operatorname{argmin}_{S' \in \operatorname{SO}(3)} \sum_{i=1}^n \|R_i^* - \hat{R}_i S'\|^2. \quad (38)$$

The error is then computed as

$$\operatorname{err}(\hat{R}, R^*) = \max_{i=1, \dots, n} d_{\angle}(\hat{R}_i S, R_i^*). \quad (39)$$

All algorithms take less than a minute to run on each individual dataset on a Macbook Air with a 1.6 GHz Dual-Core Intel Core i5 and 8 GB of RAM.

Figure 4 presents a first comparison of these algorithms on synthetic data. The model is the uniform corruption model, which is discussed in Sect. 2.6. The graph is Erdős-Rényi on $n = 50$ nodes with varying parameter p , and each edge on this graph is corrupted with probability q . The underlying rotations, $(R_1^*, \dots, R_n^*)$, are distributed uniformly on $\operatorname{SO}(D)$. The bad measurements, R_{jk}^b , are also uniformly distributed on $\operatorname{SO}(D)$. For each set of parameters ($p = 0.1, 0.2, \dots, 0.5$ and $q = 0.05, 0.1, \dots, 0.3$), 5 datasets are generated and the color represents the mean of the $\log_{10}$ -errors over these experiments. As we can see, the approximate DDS algorithm performs on par with the other most competitive methods (CEMP and MPLS).

As a second experiment, Fig. 5 presents a more challenging adversarial example on synthetic data. Here, the outliers form a consistent set of measurements themselves, and a similar corruption model is discussed in Section 7.3 of Lerman and Shi (2020), although here we extend this to $\operatorname{SO}(3)$ and use a different model for the underlying rotations.

The graph is an Erdős-Rényi graph on $n = 50$ nodes with parameter p , and each edge on this graph is corrupted with probability q . The ground truth rotations approximately come from a geodesic on $\operatorname{SO}(3)$, and the outliers are self-consistent measurements that come from (approximately) another geodesic on $\operatorname{SO}(3)$. The ground truth rotations are

$$R_i^* = \operatorname{Exp}_I \left(-s_i(v + \xi_i) \right), \quad (40)$$

where v is a fixed vector drawn uniformly from the sphere, $\xi_i \sim N(0, 10^{-4}I)$, and $s_i = -1 + 2(i-1)/50$. The outliers generated by pairwise measurements between another set of rotations

$$R_i^b = \operatorname{Exp}_I \left(-s_i(v' + \xi'_i) \right), \quad (41)$$

where again v' is a fixed vector drawn uniformly from the sphere, $\xi'_i \sim N(0, 0.5I)$, and $s_i = -1 + 2(i-1)/50$. As before, for each set of parameters p and q , 5 datasets are generated and the color represents the mean of the $\log_{10}$ -errors over these experiments. As we can see again, the approximate DDS algorithm performs well in this experiment, and in fact it performs on par with the most competitive rotation synchronization algorithms (CEMP and MPLS).

In both experiments, we note that the other competitive algorithms are CEMP and MPLS (Lerman & Shi, 2020; Shi & Lerman, 2020). As mentioned earlier, CEMP, and thus also MPLS that uses ideas of CEMP, have higher memory cost than DDS.




Fig. 4 Rotation synchronization experiment with uniform outliers. Here, p is the parameter of the Erdős-Rényi graph, and q is the percentage of corrupted edges. The underlying rotations are distributed

uniformly in $SO(D)$, and the corrupted measurements of group ratios are also uniform on $SO(D)$. The color represents the mean of the $\log_{10}$ -errors over the 10 generated datasets (Color figure online)




Fig. 5 Rotation synchronization experiment with adversarial outliers. Here, p is the parameter of the Erdős-Rényi graph, and q is the percentage of corrupted edges which are uniformly distributed across this graph. The underlying rotations follow the model in (40), and the cor-

rupted measurements are pairwise measurements between rotations generated by the separate set (41). The color represents the mean of the $\log_{10}$ -errors over the 10 generated datasets (Color figure online)

7 Conclusion

In this work, we developed the first adversarial robustness guarantees for a multiple rotation averaging algorithm. Our novel algorithm relies on finding descent directions using Tukey depth in the tangent space of $SO(D)$. To our knowledge, this represents the first application of manifold Tukey depth in an applied setting. In the case of $D = 2$ and $D = 3$, which most frequently arise in practice, our recovery thresholds are $1/4$ and $1/8$, and the algorithm can be implemented efficiently. These are the most important settings for many computer vision applications, and the results are state-of-the-

art for arbitrary outliers. We also show how to speed up the algorithm with some approximations, and this approximate algorithm performs competitively on simple synthetic experiments. Future work should also examine if it is possible to extend the analysis to the more practical approximate DDS algorithm.

In practice, we have not found the DDS algorithm to be the uniformly best choice of method. Yet, we must comment that this has not been the main aim of this work. Instead of offering a new state-of-the-art practical method for rotation synchronization, the purpose of this work has been to offer a first justification for the robustness and benign energy landscape

of a multiple rotation averaging algorithm. This also points to new avenues for research in the development of multiple rotation averaging algorithms. In some sense, using a depth based method is similar to trimming procedures, where one throws away a portion of the measurements at each iteration. We believe that a fruitful direction for the development of new practical methods will be in exploring trimmed rotation averaging algorithms.

In terms of other strong contenders for multiple rotation averaging algorithms, in the adversarial setting, it is possible to show that the L_1 minimization method of Hartley et al. (2011) may suffer from sub-optimal stationary points. Such a proof was given in an earlier version of this manuscript which can be accessed at Maunu and Lerman (2020). On the other hand, we observe that such methods perform well on common SfM datasets and giving theoretical justification for this fact should be the subject of future work.

Another direction for future work is to examine the limits of our analysis. In particular, it would be interesting to know if tighter analyses can yield larger recovery thresholds. At least for the cases of $SO(2)$ and $SO(3)$, which arise in applications, the depth descent estimator discussed in this paper has significant recovery thresholds, while also being computationally tractable. It is not clear what the optimal bounds for recovery with adversarial corruption are in general. Furthermore, if one moves away from adversarial corruption and instead considers special models of data like the uniform corruption model, the bounds could be much better.

Two more concrete directions for future work would be to carry out further examination of the ζ -well-connectedness condition in Assumption 2. In particular, it would be interesting to see if it can be relaxed at all, what its implications are, and when it actually holds.

Finally, perhaps the most important direction for future work is to give theoretically justified algorithms for a larger range of algorithms employed for SfM (Özyeşil et al., 2017; Bianco et al., 2018). Indeed, such theoretical work can lead to new and improved algorithms and also to the development of novel state-of-the-art pipelines.

Acknowledgements Gilad Lerman was supported by NSF awards DMS-1821266 and DMS-2152766.

References

- Afsari, B., Tron, R., & Vidal, R. (2013). On the convergence of gradient descent for finding the Riemannian center of mass. *SIAM Journal on Control and Optimization*, 51(3), 2230–2260.
- Afsari, Bijan. (2009). *Means and averaging on Riemannian manifolds*. PhD thesis, University of Maryland, College Park.
- Arie-Nachimson, M., Kovalsky, S. Z., Kemelmacher-Shlizerman, I., Singer, A., & Basri, R. (2012). Global motion estimation from point matches. In *2012 Second international conference on 3D imaging, modeling, processing, visualization & transmission*, (pp. 81–88). IEEE.
- Arora, R. (2009). On learning rotations. In *Advances in neural information processing systems*, (pp. 55–63).
- Arora, S., Ge, R., Ma, T., & Moitra, A. (2015). Simple, efficient, and neural algorithms for sparse coding. In *Conference on learning theory* (pp. 113–149). PMLR.
- Arrigoni, F., Rossi, B., Fragneto, P., & Fusiello, A. (2018). Robust synchronization in $SO(3)$ and $SE(3)$ via low-rank and sparse matrix decomposition. *Computer Vision and Image Understanding*, 174, 95–113.
- Bandeira, A. S., Boumal, N., & Singer, A. (2017). Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *Mathematical Programming*, 163(1–2), 145–167.
- Bandeira, A. S. (2018). Random Laplacian matrices and convex relaxations. *Foundations of Computational Mathematics*, 18(2), 345–379. <https://doi.org/10.1007/s10208-016-9341-9>.
- Bianco, S., Ciocca, G., & Marelli, D. (2018). Evaluating the performance of structure from motion pipelines. *Journal of Imaging*, 4(8), 98.
- Boumal, N. (2016). Nonconvex phase synchronization. *SIAM Journal on Optimization*, 26(4), 2355–2377.
- Boumal, N., Voroninski, V., & Bandeira, A. S. (2020). Deterministic guarantees for Burer–Monteiro factorizations of smooth semidefinite programs. *Communications on Pure and Applied Mathematics*, 73(3), 581–608.
- Chatterjee, A., & Govindu, V. M. (2017). Robust relative rotation averaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 958–972.
- Chatterjee, A., & Govindu, V. M. (2013). Efficient and robust large-scale rotation averaging. In *Proceedings of the IEEE international conference on computer vision*, (pp. 521–528).
- Chen, Y., Suh, C., & Goldsmith, A. J. (2016). Information recovery from pairwise measurements. *IEEE Transactions on Information Theory*, 62(10), 5881–5905.
- Cherapanamjeri, Y., Jain, P., & Netrapalli, P. (2017). Thresholding based outlier robust PCA. In *COLT*, (pp. 593–628).
- Chi, Y., Lu, Y., & Chen, Y. (2019). Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20), 5239–5269.
- Danzer, L., Grünbaum, B., & Klee, V. (1963) Helly’s theorem and its relatives. In *Proc. Symp. Pure Math.*, (vol. 7, pp. 101–180). Amer. Math. Soc.,.
- Dauphin, Y. N., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., & Bengio, Y. (2014). Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in neural information processing systems*, (pp. 2933–2941).
- Donoho, D. L., & Gasko, M. (1992). Breakdown properties of location estimates based on halfspace depth and projected outlyingness. *The Annals of Statistics*, 20(4), 1803–1827.
- Friedman, J., et al. (2003). Relative expanders or weakly relatively Ramanujan graphs. *Duke Mathematical Journal*, 118(1), 19–35.
- Gao, C., & Zhang, A. Y. (2021). Exact minimax estimation for phase synchronization. *IEEE Transactions on Information Theory*, 67(12), 8236–8247.
- Gao, C., Liu, J., Yao, Y., & Zhu, W. (2019). Robust estimation and generative adversarial networks. In *7th international conference on learning representations*. ICLR.
- Gao, T., & Zhao, Z. (2019). Multi-frequency phase synchronization. In *International conference on machine learning* (pp. 2132–2141). PMLR.
- Ge, R., Huang, F., Jin, C., & Yuan, Y. (2015). Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Proceedings of The 28th conference on learning theory*, (pp. 797–842).

- Ge, R., Lee, J. D., & Ma, T. (2016). Matrix completion has no spurious local minimum. *Advances in Neural Information Processing Systems*, 29.
- Govindu, V. M. (2001). Combining two-view constraints for motion estimation. In *Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001*, (vol. 2, pages II–II). IEEE.
- Govindu, V. M. (2004). Lie-algebraic averaging for globally consistent motion estimation. In *Proceedings of the 2004 IEEE Computer society conference on computer vision and pattern recognition, 2004. CVPR 2004.*, (vol. 1, pages I–I). IEEE.
- Govindu, V. M. (2006). Robustness in motion averaging. In *Asian conference on computer vision*, (pp. 457–466). Springer.
- Hammer, H. L., Yazidi, A., & Rue, H. (2022). Estimating Tukey depth using incremental quantile estimators. *Pattern Recognition*, 122, 108339.
- Hand, P., Lee, C., & Voroninski, V. (2018). Exact simultaneous recovery of locations and structure from known orientations and corrupted point correspondences. *Discrete & Computational Geometry*, 59(2), 413–450.
- Hardt, M. (2014). Understanding alternating minimization for matrix completion. In *FOCS*, (pp. 651–660). IEEE.
- Hartley, R., Trumppf, J., Dai, Y., & Li, H. (2013). Rotation averaging. *International Journal of Computer Vision*, 103(3), 267–305.
- Hartley, Richard, Aftab, Khurram, & Trumppf, Jochen. (2011). L1 rotation averaging using the weiszfeld algorithm. In *Computer vision and pattern recognition (CVPR), 2011 IEEE Conference on*, (pp. 3041–3048). IEEE.
- Hoffman, C., Kahle, M., & Paquette, E. (2021). Spectral gaps of random graphs and applications. *International Mathematics Research Notices*, 2021(11), 8353–8404.
- Huang, Qi-Xing, & Guibas, Leonidas. (2013). Consistent shape maps via semidefinite programming. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH symposium on geometry processing, SGP '13*, (pp. 177–186). Eurographics Association.
- Huang, Xiangru, Liang, Zhenxiao, Bajaj, Chandrajit, & Huang, Qixing. (2017). Translation synchronization via truncated least squares. In *Advances in neural information processing systems 30: Annual conference on neural information processing systems 2017, 4–9 December 2017, Long Beach, CA, USA*, (pp. 1459–1468).
- Huang, Xiangru, Liang, Zhenxiao, Zhou, Xiaowei, Xie, Yao, Guibas, Leonidas J., & Huang, Qixing. (2019). Learning transformation synchronization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, (pp. 8082–8091).
- Jain, Prateek, Tewari, Ambuj, & Kar, Purushottam. (2014). On iterative hard thresholding methods for high-dimensional m-estimation. In *Advances in neural information processing systems*, (pp. 685–693).
- Karcher, H. (1977). Riemannian center of mass and mollifier smoothing. *Communications on Pure and Applied Mathematics*, 30(5), 509–541.
- Lee, J. D., Simchowitz, M., Jordan, M. I., & Recht, B. (2016). Gradient descent converges to minimizers. University of California, Berkeley. 1050, 16.
- Lerman, G., & Maunu, T. (2017). Fast, robust and non-convex subspace recovery. *Information and Inference: A Journal of the IMA*, 7(2), 277–336.
- Lerman, G., & Shi, Y. (2022). Robust group synchronization via cycle-edge message passing. *Foundations of Computational Mathematics*, 22(6), 1665–1741.
- Lerman, G., Shi, Y., & Zhang, T. (2018). Exact camera location recovery by least unsquared deviations. *SIAM Journal on Imaging Sciences*, 11(4), 2692–2721.
- Li, X., Ling, S., Strohmmer, T., & Wei, K. (2019). Rapid, robust, and reliable blind deconvolution via nonconvex optimization. *Applied and computational harmonic analysis*, 47(3), 893–934.
- Liu, H., Yue, M.-C., & So, A. (2020). A unified approach to synchronization problems over subgroups of the orthogonal group. arXiv preprint [arXiv:2009.07514](https://arxiv.org/abs/2009.07514).
- Liu, X. (2017). Fast implementation of the tukey depth. *Computational Statistics*, 32(4), 1395–1410.
- Liu, X., Mosler, K., & Mozharovskiy, P. (2019). Fast computation of tukey trimmed regions and median in dimension $p > 2$. *Journal of Computational and Graphical Statistics*, 28(3), 682–697.
- Lu, J., & Steinerberger, S. (2020). Synchronization of Kuramoto oscillators in dense networks. *Nonlinearity*, 33(11), 5905.
- Ma, C., Wang, K., Chi, Y., & Chen, Y. (Jul 2018). Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval and matrix completion. In *PMLR*, (vol. 80, pp. 3345–3354, 10–15).
- Martinec, D., & Pajdla, T. (2007). Robust rotation and translation estimation in multiview reconstruction. In *2007 IEEE conference on computer vision and pattern recognition*, (pp. 1–8). IEEE.
- Maunu, T., & Lerman, G. (2019). Robust subspace recovery with adversarial outliers. arXiv preprint [arXiv:1904.03275](https://arxiv.org/abs/1904.03275).
- Maunu, T., Zhang, T., & Lerman, G. (2019). A well-tempered landscape for non-convex robust subspace recovery. *Journal of Machine Learning Research*, 20(37), 1–59.
- Maunu, Tyler, & Lerman, Gilad. (2020). Depth descent synchronization in $SO(D)$. arXiv preprint [arXiv:2002.05299v2](https://arxiv.org/abs/2002.05299v2).
- Mei, S., Bai, Y., & Montanari, A. (2018). The landscape of empirical risk for nonconvex losses. *The Annals of Statistics*, 46(6A), 2747–2774.
- Meyer, R. R. (1976). Sufficient conditions for the convergence of monotonic mathematical programming algorithms. *Journal of Computer and System Sciences*, 12, 108–121.
- Mizera, I. (2002). On depth and deep points: A calculus. *The Annals of Statistics*, 30(6), 1681–1736.
- Mizera, I., & Volauf, M. (2002). Continuity of halfspace depth contours and maximum depth estimators: Diagnostics of depth-related methods. *Journal of Multivariate Analysis*, 83(2), 365–388.
- Moakher, M. (2002). Means and averaging in the group of rotations. *SIAM Journal on Matrix Analysis and Applications*, 24(1), 1–16.
- Netrapalli, Praneeth, Niranjan, U. N., Sanghavi, Sujay, Anandkumar, Animashree, & Jain, Prateek. (2014). Non-convex robust pca. In *Advances in neural information processing systems*, (pp. 1107–1115).
- Ozyesil, O., Singer, A., & Basri, R. (2015). Stable camera motion estimation using convex programming. *SIAM Journal on Imaging Sciences*, 8(2), 1220–1262.
- Özyeşil, O., Voroninski, V., Basri, R., & Singer, A. (2017). A survey of structure from motion*. *Acta Numerica*, 26, 305–364.
- Perry, A., Wein, A. S., Bandeira, A. S., & Moitra, A. (2018). Message-passing algorithms for synchronization problems over compact groups. *Communications on Pure and Applied Mathematics*, 71(11), 2275–2322.
- Petersen, P. (2016). *Riemannian geometry* (3rd ed., Vol. 171). Springer.
- Qu, Q., Zhang, Y., Eldar, Y., & Wright, J. (2019). Convolutional phase retrieval via gradient descent. *IEEE Transactions on Information Theory*, 66(3), 1785–1821.
- Qu, Q., Zhu, Z., Li, X., Tsakiris, M., Wright, J., & Vidal, R. (2020). Finding the sparsest vectors in a subspace: Theory, algorithms, and applications. arXiv preprint [arXiv:2001.06970](https://arxiv.org/abs/2001.06970).
- Rado, R. (1946). A theorem on general measure. *Journal of the London Mathematical Society*, 1(4), 291–300.
- Rosen, D. M., Carlone, L., Bandeira, A. S., & Leonard, J. J. (2019). Se-sync: A certifiably correct algorithm for synchronization over the special euclidean group. *The International Journal of Robotics Research*, 38(2–3), 95–125.
- Shi, Yunpeng, & Lerman, Gilad. (2020). Message passing least squares framework and its application to rotation synchronization. In

- International conference on machine learning*, (pp. 8796–8806). PMLR.
- Singer, A. (2011). Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1), 20–36.
- Sun, J., Qu, Q., & Wright, J. (May 2015a). Complete dictionary recovery over the sphere. In *Sampling Theory and Applications (SampTA), 2015 international conference on*, (pp. 407–410). <https://doi.org/10.1109/SAMPAT.2015.7148922>.
- Sun, J., Qu, Q., & Wright, J. (2015b). When are nonconvex problems not scary? arXiv preprint [arXiv:1510.06096](https://arxiv.org/abs/1510.06096).
- Taylor, C. J., & Kriegman, D. J. (1994). Minimization on the Lie group $SO(3)$ and related manifolds. *Yale University*, 16, 155.
- Toth, C. D., O'Rourke, J., & Goodman, J. E. (2017). *Handbook of discrete and computational geometry*. Chapman and Hall/CRC.
- Tron, R., & Vidal, R. (2009). Distributed image-based 3-d localization of camera sensor networks. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, (pp. 901–908). IEEE.
- Tron, R., Zhou, X., & Daniilidis, K. (2016). A survey on rotation optimization in structure from motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, (pp. 77–85).
- Tukey, John W. (1974). T6: Order statistics, in mimeographed notes for statistics 411. Department of Statistics, Princeton University.
- Waldspurger, I., & Waters, A. (2020). Rank optimality for the Burer-Monteiro factorization. *SIAM Journal on Optimization*, 30(3), 2577–2602.
- Wang, L., & Singer, A. (2013). Exact and stable recovery of rotations for robust synchronization. *Information and Inference*.
- Wang, L., & Singer, A. (2013). Exact and stable recovery of rotations for robust synchronization. *Information and Inference*. <https://doi.org/10.1093/imaiai/iat005>.
- Yi, X., Park, D., Chen, Y., & Caramanis, C. (2016). Fast algorithms for robust PCA via gradient descent. In *NIPS*, (pp. 4152–4160).
- Zhang, T., & Yang, Y. (2018). Robust PCA by manifold optimization. *Journal of Machine Learning Research*, 19(80), 1–39.
- Zhang, Y., Qu, Q., & Wright, J. (2020). From symmetry to geometry: Tractable nonconvex problems. arXiv preprint [arXiv:2007.06753](https://arxiv.org/abs/2007.06753).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.