

A Polynomial-time Solution for Robust Registration with Extreme Outlier Rates

Heng Yang and Luca Carlone

Laboratory for Information & Decision Systems (LIDS)
Massachusetts Institute of Technology, Cambridge, MA 02139
Email: {hankyang, lcarlone}@mit.edu

Abstract—We propose a robust approach for the registration of two sets of 3D points in the presence of a large amount of outliers. Our first contribution is to reformulate the registration problem using a *Truncated Least Squares* (TLS) cost that makes the estimation insensitive to a large fraction of spurious point-to-point correspondences. The second contribution is a general framework to decouple rotation, translation, and scale estimation, which allows solving in cascade for the three transformations. Since each subproblem (scale, rotation, and translation estimation) is still non-convex and combinatorial in nature, our third contribution is to show that (i) TLS scale and (component-wise) translation estimation can be solved exactly and in polynomial time via an *adaptive voting* scheme, (ii) TLS rotation estimation can be relaxed to a semidefinite program and the relaxation is tight in practice, even in the presence of an extreme amount of outliers. We validate the proposed algorithm, named TEASER (*Truncated least squares Estimation And SEMidefinite Relaxation*), in standard registration benchmarks showing that the algorithm outperforms RANSAC and robust local optimization techniques, and favorably compares with Branch-and-Bound methods, while being a polynomial-time algorithm. TEASER can tolerate up to 99% outliers and returns highly-accurate solutions.

I. INTRODUCTION

Point cloud registration is a fundamental problem in robotics and computer vision and consists in finding the best transformation (rotation, translation, and potentially scale) that aligns two point clouds. It finds applications in motion estimation and 3D reconstruction [27, 6, 17, 58], object recognition and localization [19, 53, 56, 40], panorama stitching [3], and medical imaging [2, 50], to name a few.

When the ground-truth correspondences between the point clouds are known and the noise follows a zero-mean Gaussian distribution, the registration problem can be readily solved, since elegant closed-form solutions [28, 1] exist for the case of isotropic noise, and recently proposed convex relaxations [9] are empirically tight even in the presence of large anisotropic noise. In practice, however, the correspondences are either unknown, or contain a high ratio of outliers. Large outlier rates are typical of 3D keypoint detection and matching techniques [52, 46]. Therefore, it is common to use the aforementioned methods within a RANSAC scheme [21].

While RANSAC is a popular approach for several robust vision and robotics problems, its runtime grows exponentially with the outlier ratio [11] and it can perform poorly with extreme outlier rates. The capability of tolerating a large amount of outliers is of paramount importance in applications where




Fig. 1. We address 3D point cloud registration with extreme outlier rates. (a) Bunny dataset spoiled with 80% outlier correspondences: finding/localizing the bunny is challenging even for a human; (b) The proposed *Truncated least squares Estimation And SEMidefinite Relaxation* (TEASER) is able to find the correct inlier correspondences (green) and compute the correct registration result in polynomial time. TEASER is robust to 99% outlier correspondences.

the correspondences are unknown and when operating in the clutter (e.g., object pose estimation in the wild). Moreover, even when the correspondences are known but uncertain, it is desirable to develop registration techniques that can afford stronger performance guarantees compared to RANSAC.

This paper is motivated by the goal of designing an approach that (i) can solve registration globally (without relying on an initial guess), (ii) can tolerate extreme amounts of outliers (e.g., when 99% of the measurements are outliers), (iii) runs in polynomial time, and (iv) provides formal performance guarantees. The related literature, reviewed in Section II, fails to simultaneously address all these aspects, and only includes techniques that are robust to moderate amounts (e.g., 60%) of outliers and lack optimality guarantees (e.g., FGR [59]), or are globally optimal but run in exponential time in the worst case, such as branch-and-bound (BnB) methods (e.g., Go-ICP [55]).

Contribution. Our first contribution (presented in Section III) is to reformulate the registration problem using a *Truncated Least Squares* (TLS) cost that is insensitive to a large fraction of spurious data. We name the resulting problem the *Truncated Least Squared Registration* (TR) problem.

The second contribution (Section IV) is a general framework to decouple scale, rotation, and translation estimation. The idea of decoupling rotation and translation has appeared in related work, e.g., [39, 37, 11]. The novelty of our proposal is threefold: (i) we develop invariant measurements to estimate the scale ([37, 11] assume the scale is given), (ii) we make the decoupling formal within the framework of *unknown-but-*

bounded noise [42], and (iii) we provide a general graph-theoretic framework to derive these invariant measurements.

The decoupling allows solving in cascade for scale, rotation, and translation. However, each subproblem is still combinatorial in nature. Our third contribution is to show that (i) in the scalar case TLS estimation can be solved exactly in polynomial time using an *adaptive voting* scheme, and this enables efficient estimation of the scale and the (component-wise) translation; (ii) we can prune a large amount of outliers by finding a *maximal clique* of the graph defined by the invariant measurements; (iii) we can formulate a tight *semidefinite programming (SDP) relaxation* to estimate the rotation, (iv) we can provide per-instance bounds on the performance of the SDP relaxation. To the best of our knowledge, this is the first polynomial-time algorithm for outlier-robust registration with computable performance guarantees.

We validate the proposed algorithm, named *Truncated least squares Estimation And Semidefinite Relaxation (TEASER)*, in standard registration benchmarks as well as robotics datasets, showing that the algorithm outperforms RANSAC and robust local optimization techniques, and favorably compares with Branch-and-Bound methods, while being a polynomial-time algorithm. TEASER can tolerate up to 99% outliers (Fig. 1) and returns highly-accurate solutions.

II. RELATED WORK

There are two established paradigms for the registration of 3D point clouds: *Correspondence-based* and *Simultaneous Pose and Correspondence* methods.

A. Correspondence-based Methods

Correspondence-based methods first detect and match 3D keypoints between point clouds using local [52, 24, 46, 47] or global [19, 33] descriptors to establish putative correspondences, and then either use closed-form solutions [28, 1] in a RANSAC [21] scheme, or apply robust optimization methods [59, 11] to gain robustness against outliers. 3D keypoint matching is known to be less accurate compared to 2D counterparts like SIFT and ORB, thus causing much higher outlier rates, e.g., having 95% spurious correspondences is considered common [11]. Therefore, a robust backend that can deal with extreme outlier rates is highly desirable.

Registration without outliers. Horn [28] and Arun [1] show that optimal solutions (in the maximum likelihood sense) for scale, rotation, and translation can be computed in closed form when the points are affected by isotropic zero-mean Gaussian noise. Olsson *et al.* [44] propose a method based on BnB that is globally optimal and allows point-to-point, point-to-line, and point-to-plane correspondences. Recently, Briales and Gonzalez-Jimenez [9] propose a semidefinite relaxation that can deal with anisotropic Gaussian noise, and has per-instance optimality guarantees.

Robust registration. Probably the most widely used robust registration approach is based on RANSAC [21, 15], which has enabled several early applications in vision and robotics [25, 41]. Despite its efficiency in the low-noise and

low-outlier regime, RANSAC exhibits slow convergence and low accuracy with large outlier rates [11], where it becomes harder to sample a “good” consensus set. Other approaches resort to *M-estimation*, which replaces the least squares objective function with robust costs that are less sensitive to outliers [51, 5, 35]. Zhou *et al.* [59] propose *Fast Global Registration (FGR)* that uses the Geman-McClure cost function and leverages graduated non-convexity to solve the resulting non-convex optimization. Since graduated non-convexity has to be solved in discrete steps, FGR does not guarantee global optimality [11]. Indeed, FGR tends to fail when the outlier ratio is high (>80%), as we show in Section VI. Bustos and Chin [11] propose a *Guaranteed Outlier REmoval (GORE)* technique, that uses geometric operations to significantly reduce the amount of outlier correspondences before passing them to the optimization backend. GORE has been shown to be robust to 95% spurious correspondences [11]. However, GORE does not estimate the *scale* of the registration and has exponential worst-case time complexity due to the possible usage of BnB (see Algorithm 2 in [11]).

B. Simultaneous Pose and Correspondence Methods

Simultaneous Pose and Correspondence (SPC) methods alternate between finding the correspondences and computing the best transformation given the correspondences.

Local methods. The *Iterative Closest Point (ICP)* algorithm [4] is considered a milestone in point cloud registration. However, ICP is prone to converge to *local minima* and it only performs well given a good initial guess. Multiple variants of ICP [22, 48, 57, 38, 16, 32, 7] have proposed to use robust cost functions to improve convergence. Probabilistic interpretations have also been proposed to improve ICP convergence, for instance interpreting the registration problem as a minimization of the Kullback-Leibler divergence between two Gaussian Mixture Models [30, 43, 31, 13]. All these methods rely on iterative local search, do not provide global optimality guarantees, and typically fail without a good initial guess.

Global methods. Global SPC approaches compute a globally optimal solution without initial guesses, and are usually based on BnB, which at each iteration divides the parameter space into multiple sub-domains (branch) and computes the bounds of the objective function for each sub-domain (bound). A series of geometric techniques have been proposed to improve the bounding tightness [26, 8, 55, 14, 12] and increase the search speed [55, 36]. However, the runtime of BnB increases exponentially with the size of the point cloud and it can be made worse by the explosion of the number of local minima resulting from high outlier ratios [11]. Global SPC registration can be also formulated as a mixed-integer program [29], though the runtime remains exponential.

III. ROBUST REGISTRATION WITH TRUNCATED LEAST SQUARES COST

In the robust registration problem, we are given two 3D point sets $\mathcal{A} = \{\mathbf{a}_i\}_{i=1}^N$ and $\mathcal{B} = \{\mathbf{b}_i\}_{i=1}^N$, with $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^3$,

such that:

$$\mathbf{b}_i = s\mathbf{R}\mathbf{a}_i + \mathbf{t} + \mathbf{o}_i + \boldsymbol{\epsilon}_i \quad (1)$$

where $s > 0$, $\mathbf{R} \in \text{SO}(3)$, and $\mathbf{t} \in \mathbb{R}^3$ are an unknown scale, rotation, and translation, $\boldsymbol{\epsilon}_i$ models measurement noise, and \mathbf{o}_i is a vector of zeros for *inliers*, or a vector of arbitrary numbers for *outliers*. In words, if the i -th correspondence $(\mathbf{a}_i, \mathbf{b}_i)$ is an inlier, \mathbf{b}_i corresponds to a 3D transformation of \mathbf{a}_i (plus noise), while if $(\mathbf{a}_i, \mathbf{b}_i)$ is an outlier correspondence, \mathbf{b}_i is just an arbitrary vector. $\text{SO}(3) \doteq \{\mathbf{R} \in \mathbb{R}^{3 \times 3} : \mathbf{R}^\top \mathbf{R} = \mathbf{I}_3, \det(\mathbf{R}) = +1\}$ is the set of proper rotation matrices (where \mathbf{I}_d is the identity matrix of size d). We consider a correspondence-based setup, where we need to compute $(s, \mathbf{R}, \mathbf{t})$ given putative correspondences $(\mathbf{a}_i, \mathbf{b}_i)$, $i = 1, \dots, N$.

Registration without outliers. When $\boldsymbol{\epsilon}_i$ is a zero-mean Gaussian noise with isotropic covariance $\sigma_i^2 \mathbf{I}_3$, and all the correspondences are correct (i.e., $\mathbf{o}_i = \mathbf{0}, \forall i$), the Maximum Likelihood estimator of $(s, \mathbf{R}, \mathbf{t})$ can be computed in closed form by decoupling the estimation of the scale, translation, and rotation, using Horn’s [28] or Arun’s method [1].

Robust registration. In practice, a large fraction of the correspondences are *outliers*, due to incorrect keypoint matching. Despite the elegance of the closed-form solutions [28, 1], they are not robust to outliers, and a single “bad” outlier can compromise the correctness of the resulting estimate. Hence, we propose a *truncated least squares registration* formulation that can tolerate extreme amounts of spurious data.

Truncated Least Squares Registration. We depart from the Gaussian noise model and assume the noise is *unknown but bounded* [42]. Formally, we assume the noise $\boldsymbol{\epsilon}_i$ in (1) is such that $\|\boldsymbol{\epsilon}_i\| \leq \beta_i$, where β_i is a given bound.

Then we adopt the following *Truncated Least Squares Registration (TR)* formulation :

$$\min_{s>0, \mathbf{t} \in \mathbb{R}^3, \mathbf{R} \in \text{SO}(3)} \sum_{i=1}^N \min \left(\frac{1}{\beta_i^2} \|\mathbf{b}_i - s\mathbf{R}\mathbf{a}_i - \mathbf{t}\|^2, \bar{c}^2 \right) \quad (2)$$

which computes a least squares solution of measurements with small residual ($\frac{1}{\beta_i^2} \|\mathbf{b}_i - s\mathbf{R}\mathbf{a}_i - \mathbf{t}\|^2 \leq \bar{c}^2$), while discarding measurements with large residuals (when $\frac{1}{\beta_i^2} \|\mathbf{b}_i - s\mathbf{R}\mathbf{a}_i - \mathbf{t}\|^2 > \bar{c}^2$ the i -th summand becomes constant and does not influence the optimization). The constant \bar{c}^2 is typically chosen to be 1, while one may use a different \bar{c}^2 to be stricter or more lenient towards potential outliers.

IV. DECOUPLING SCALE, ROTATION, AND TRANSLATION ESTIMATION

We propose a polynomial-time algorithm that decouples the estimation of scale, translation, and rotation in problem (2). The key insight is that we can reformulate the measurements (1) to obtain quantities that are invariant to a subset of the transformations (scaling, rotation, translation).

A. Translation Invariant Measurements (TIM)

While the absolute positions of the points in \mathcal{B} depend on the translation \mathbf{t} , the relative positions are invariant to




Fig. 2. Graph topologies for generating TIMs in the Bunny dataset.

\mathbf{t} . Mathematically, given two points \mathbf{b}_i and \mathbf{b}_j from (1), the relative position of these two points is:

$$\mathbf{b}_j - \mathbf{b}_i = s\mathbf{R}(\mathbf{a}_j - \mathbf{a}_i) + (\mathbf{o}_j - \mathbf{o}_i) + (\boldsymbol{\epsilon}_j - \boldsymbol{\epsilon}_i) \quad (3)$$

where the translation \mathbf{t} cancels out in the subtraction. Therefore, we can obtain a *Translation Invariant Measurement (TIM)* by computing $\bar{\mathbf{a}}_{ij} \doteq \mathbf{a}_j - \mathbf{a}_i$ and $\bar{\mathbf{b}}_{ij} \doteq \mathbf{b}_j - \mathbf{b}_i$, and the TIM satisfies the following generative model:

$$\bar{\mathbf{b}}_{ij} = s\mathbf{R}\bar{\mathbf{a}}_{ij} + \mathbf{o}_{ij} + \boldsymbol{\epsilon}_{ij} \quad (\text{TIM})$$

where $\mathbf{o}_{ij} \doteq \mathbf{o}_j - \mathbf{o}_i$ is zero if both the i -th and the j -th measurements are inliers (or it is an arbitrary vector otherwise), while $\boldsymbol{\epsilon}_{ij} \doteq \boldsymbol{\epsilon}_j - \boldsymbol{\epsilon}_i$ is the measurement noise. It is easy to see that if $\|\boldsymbol{\epsilon}_i\| \leq \beta_i$ and $\|\boldsymbol{\epsilon}_j\| \leq \beta_j$ then $\|\boldsymbol{\epsilon}_{ij}\| \leq \beta_i + \beta_j \doteq \beta_{ij}$.

The advantage of the TIMs in eq. (TIM) is that their generative model only depends on two unknowns, s and \mathbf{R} . The number of TIMs is upper-bounded by $\binom{N}{2} = N(N-1)/2$, where pairwise relative measurements between all pairs of points are computed. For computational reasons, one might want to downsample the TIMs. Theorem 1 below provides a graph-theoretic way to create the TIMs.

Theorem 1 (Translation Invariant Measurements). *Define the vectors $\mathbf{a} \in \mathbb{R}^{3N}$ (resp. $\mathbf{b} \in \mathbb{R}^{3N}$), obtained by concatenating all vectors \mathbf{a}_i (resp. \mathbf{b}_i) in a single column vector. Moreover, define an arbitrary graph \mathcal{G} with nodes $\{1, \dots, N\}$ and an arbitrary set of edges \mathcal{E} . Then, the vectors $\bar{\mathbf{a}} = (\mathbf{A} \otimes \mathbf{I}_3)\mathbf{a}$ and $\bar{\mathbf{b}} = (\mathbf{A} \otimes \mathbf{I}_3)\mathbf{b}$ are TIMs, where $\mathbf{A} \in \mathbb{R}^{|\mathcal{E}| \times N}$ is the incidence matrix of \mathcal{G} , and \otimes is the Kronecker product.*

A proof of the theorem is given in the Supplementary Material. Three potential graph topologies for generating TIMs are illustrated in Fig. 2.

B. Translation and Rotation Invariant Measurements (TRIM)

While the relative locations of pairs of points (TIMs) still depends on the rotation \mathbf{R} , their distances are invariant to both \mathbf{R} and \mathbf{t} . Therefore, to build rotation invariant measurements, we compute the norm of each TIM vector:

$$\|\bar{\mathbf{b}}_{ij}\| = \|s\mathbf{R}\bar{\mathbf{a}}_{ij} + \mathbf{o}_{ij} + \boldsymbol{\epsilon}_{ij}\| \quad (4)$$

We now note that for the inliers ($\mathbf{o}_{ij} = \mathbf{0}$) it holds (using $\|\epsilon_{ij}\| \leq \beta_{ij}$ and the triangle inequality):

$$\|s\mathbf{R}\bar{\mathbf{a}}_{ij}\| - \beta_{ij} \leq \|s\mathbf{R}\bar{\mathbf{a}}_{ij} + \epsilon_{ij}\| \leq \|s\mathbf{R}\bar{\mathbf{a}}_{ij}\| + \beta_{ij} \quad (5)$$

hence we can write (4) equivalently as:

$$\|\bar{\mathbf{b}}_{ij}\| = \|s\mathbf{R}\bar{\mathbf{a}}_{ij}\| + \tilde{o}_{ij} + \tilde{\epsilon}_{ij} \quad (6)$$

with $|\tilde{\epsilon}_{ij}| \leq \beta_{ij}$, and $\tilde{o}_{ij} = 0$ if both i and j are inliers or is an arbitrary scalar otherwise. Recalling that the norm is rotation invariant and that $s > 0$, and dividing both members of (6) by $\|\bar{\mathbf{a}}_{ij}\|$, we obtain new measurements $s_{ij} \doteq \frac{\|\bar{\mathbf{b}}_{ij}\|}{\|\bar{\mathbf{a}}_{ij}\|}$:

$$s_{ij} = s + o_{ij}^s + \epsilon_{ij}^s \quad (\text{TRIM})$$

where $\epsilon_{ij}^s \doteq \frac{\tilde{\epsilon}_{ij}}{\|\bar{\mathbf{a}}_{ij}\|}$, and $o_{ij}^s \doteq \frac{\tilde{o}_{ij}}{\|\bar{\mathbf{a}}_{ij}\|}$. It is easy to see that $|\epsilon_{ij}^s| \leq \beta_{ij}/\|\bar{\mathbf{a}}_{ij}\|$ since $|\tilde{\epsilon}_{ij}| \leq \beta_{ij}$. We define $\alpha_{ij} \doteq \beta_{ij}/\|\bar{\mathbf{a}}_{ij}\|$.

Eq. (TRIM) describes a *translation and rotation invariant measurement* (TRIM) whose generative model is only function of the unknown scale s . A summary table of the invariant measurements and a remark on the novelty of creating TIMs and TRIMs is presented in the Supplementary Material.

C. Our Registration Algorithm: Truncated least squares Estimation And SEMidefinite Relaxation (TEASER)

We propose a decoupled approach to solve in cascade for the scale, the rotation, and the translation in (2). The approach, named *Truncated least squares Estimation And SEMidefinite Relaxation* (TEASER), works as follows:

- 1) we use the TRIMs to estimate the scale \hat{s}
- 2) we use \hat{s} and the TIMs to estimate the rotation $\hat{\mathbf{R}}$
- 3) we use \hat{s} and $\hat{\mathbf{R}}$ to estimate the translation $\hat{\mathbf{t}}$ from the original TLS problem (2).

The pseudocode is also summarized in Algorithm 1.

Algorithm 1: Truncated least squares Estimation And SEMidefinite Relaxation (TEASER).

- 1 **Input:** points $(\mathbf{a}_i, \mathbf{b}_i)$ and bounds β_i ($i = 1, \dots, N$), threshold \bar{c}^2 (default: $\bar{c}^2 = 1$), graph edges \mathcal{E} (default: \mathcal{E} describes the complete graph);
 - 2 **Output:** $s, \mathbf{R}, \mathbf{t}$;
 - 3 % Compute TIM and TRIM
 - 4 $\bar{\mathbf{b}}_{ij} = \mathbf{b}_j - \mathbf{b}_i$, $\bar{\mathbf{a}}_{ij} = \mathbf{a}_j - \mathbf{a}_i$, $\beta_{ij} = \beta_i + \beta_j \quad \forall (i, j) \in \mathcal{E}$
 - 5 $s_{ij} = \frac{\|\bar{\mathbf{b}}_{ij}\|}{\|\bar{\mathbf{a}}_{ij}\|}$, $\alpha_{ij} = \frac{\beta_{ij}}{\|\bar{\mathbf{a}}_{ij}\|} \quad \forall (i, j) \in \mathcal{E}$
 - 6 % Decoupled estimation of $s, \mathbf{R}, \mathbf{t}$
 - 7 $\hat{s} = \text{estimate_s}(\{s_{ij}, \alpha_{ij} : \forall (i, j) \in \mathcal{E}\}, \bar{c}^2)$
 - 8 $\hat{\mathbf{R}} = \text{estimate_R}(\{\bar{\mathbf{a}}_{ij}, \bar{\mathbf{b}}_{ij}, \beta_{ij} : \forall (i, j) \in \mathcal{E}\}, \bar{c}^2, \hat{s})$
 - 9 $\hat{\mathbf{t}} = \text{estimate_t}(\{\mathbf{a}_i, \mathbf{b}_i, \beta_i : i = 1 \dots, N\}, \bar{c}^2, \hat{s}, \hat{\mathbf{R}})$
 - 10 **return:** $\hat{s}, \hat{\mathbf{R}}, \hat{\mathbf{t}}$
-

The following section describes how to implement the functions in lines 7-9 of Algorithm 1. In particular, we show how to obtain global and robust estimates of scale (`estimate_s`) in Section V-A, rotation (`estimate_R`) in Section V-B, and translation (`estimate_t`) in Section V-C.




Fig. 3. (a) confidence interval for each measurement s_k (every s in the k -th interval satisfies $\frac{\|s - s_k\|^2}{\alpha_k^2} \leq \bar{c}^2$); (b) cardinality of the consensus set for every s and middle-points m_j for each interval with constant consensus set.

V. SOLVING THE REGISTRATION SUBPROBLEMS

A. Robust Scale Estimation

The generative model (TRIM) describes linear measurements s_{ij} of the unknown scale s , affected by bounded noise $|\epsilon_{ij}^s| \leq \alpha_{ij}$ including potential outliers (when $o_{ij}^s \neq 0$). Again, we estimate the scale given the measurements s_{ij} and the bounds α_{ij} using a TLS estimator:

$$\hat{s} = \arg \min_s \sum_{k=1}^K \min \left(\frac{(s - s_k)^2}{\alpha_k^2}, \bar{c}^2 \right) \quad (7)$$

where for simplicity we numbered the measurements from 1 to $K = |\mathcal{E}|$ and adopted the notation s_k instead of s_{ij} .

The following theorem shows that one can solve (7) in polynomial time by a simple enumeration.

Theorem 2 (Optimal TLS Scale Estimation). *For a given $s \in \mathbb{R}$, define the consensus set of s as $\mathcal{C}(s) = \{k : \frac{(s - s_k)^2}{\alpha_k^2} \leq \bar{c}^2\}$. Then, for any $s \in \mathbb{R}$, there are at most $2K - 1$ different non-empty consensus sets. If we name these sets $\mathcal{C}_1, \dots, \mathcal{C}_{2K-1}$, then the solution of (7) can be computed by enumeration as:*

$$\hat{s} = \arg \min \left\{ f_s(\hat{s}_i) : \hat{s}_i = \left(\sum_{k \in \mathcal{C}_i} \frac{1}{\alpha_k^2} \right)^{-1} \sum_{k \in \mathcal{C}_i} \frac{s_k}{\alpha_k^2}, \forall i \right\} \quad (8)$$

where $f_s(\cdot)$ is the objective function of (7).

Theorem 2, whose proof is given in the Supplementary Material, is based on the insight that the consensus set can only change at the boundaries of the intervals $[s_k - \alpha_k \bar{c}, s_k + \alpha_k \bar{c}]$ (Fig. 3(a)) and there are at most $2K$ such boundaries. The theorem also suggests a straightforward *adaptive voting* algorithm to solve (7), with pseudocode given in Algorithm 2. The algorithm first builds the boundaries of the intervals shown in Fig. 3(a) (line 4). Then, for each interval, it evaluates the consensus set (line 12, see also Fig. 3(b)). Since the consensus set does not change within an interval, we compute it at the interval centers (line 6, see also Fig. 3(b)). Finally, the cost of each consensus set is computed and the smallest cost is returned as optimal solution (line 17).

The interested reader can find a discussion on the relation between TLS and *consensus maximization* (a popular approach for outlier detection [49, 37]) in the Supplementary Material.

Maximal clique inlier selection (MCIS). The graph theoretic interpretation of Theorem 1 offers further opportunities to prune outliers. Considering the TRIMs as edges in the graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ (where the vertices \mathcal{V} are the 3D points and the edge

Algorithm 2: Adaptive Voting.

1 Input: s_k, α_k, \bar{c} ;
2 Output: \hat{s} , scale estimate solving (7);
3 % Define boundaries and sort
4 $\mathbf{v} = \text{sort}([s_1 - \alpha_1 \bar{c}, s_1 + \alpha_1 \bar{c}, \dots, s_K - \alpha_K \bar{c}, s_K + \alpha_K \bar{c}])$

5 % Compute middle points
6 $m_i = \frac{v_{2i-1} + v_{2i}}{2}$ for $i = 1, \dots, 2K - 1$
7 % Voting
8 for $i = 1, \dots, 2K - 1$ **do**
9 $\mathcal{S}_i = \emptyset$
10 **for** $k = 1, \dots, K$ **do**
11 **if** $m_i \in [s_k - \alpha_k \bar{c}, s_k + \alpha_k \bar{c}]$ **then**
12 $\mathcal{S}_i = \mathcal{S}_i \cup \{k\}$ % add to consensus set
13 **end**
14 **end**
15 end
16 % Enumerate consensus sets and return best
17 return:

$$\arg \min \left\{ f_s(\hat{s}_i) : \hat{s}_i = \left(\sum_{k \in \mathcal{C}_i} \frac{1}{\alpha_k^2} \right)^{-1} \sum_{k \in \mathcal{C}_i} \frac{s_k}{\alpha_k^2}, \forall i \right\}$$

set \mathcal{E} induces the TIMs and TRIMs per Theorem 1), we can use the scale estimate \hat{s} from Algorithm 2 to prune edges (i, j) in the graph whose associated TRIM s_{ij} is such that $\frac{(s_{ij} - \hat{s})^2}{\alpha_{ij}^2} > \bar{c}^2$. This allows us to obtain a pruned graph $\mathcal{G}'(\mathcal{V}, \mathcal{E}')$, with $\mathcal{E}' \subseteq \mathcal{E}$, where gross outliers are discarded. The following result ensures that inliers form a clique in the graph $\mathcal{G}'(\mathcal{V}, \mathcal{E}')$ enabling an even more substantial rejection of outliers.

Theorem 3 (Maximal Clique Inlier Selection). *Edges corresponding to inlier TIMs form a clique in \mathcal{E}' , and there is at least one maximal clique in \mathcal{E}' that contains all the inliers.*

A proof of Theorem 3 is presented in the Supplementary Material. Theorem 3 allows us to prune outliers by finding the maximal cliques of $\mathcal{G}'(\mathcal{V}, \mathcal{E}')$. Although finding the maximal cliques of a graph takes exponential time in general, there exist efficient approximation algorithms based on heuristics [10, 45, 54]. In addition, under high outlier rates, the graph $\mathcal{G}'(\mathcal{V}, \mathcal{E}')$ is sparse and the maximal clique problem can be solved quickly in practice [20]. In this paper, we choose the maximal clique with largest cardinality, i.e., the *maximum clique*, as the inlier set to pass to rotation estimation. Section VI-A shows that this method drastically reduces the number of outliers.

In summary, the function `estimate_s` in Algorithm 1 first calls Algorithm 2, then computes the maximum clique in the resulting graph to reject all measurements outside the clique.

What if the scale is known? In some registration problems, the scale is known, e.g., the scale of the two point clouds is the same. In such a case, we can skip Algorithm 2 and set \hat{s} to be the known scale. Moreover, we can still use the MCIS method to largely reduce the number of outliers.

B. Robust Rotation Estimation

The generative model (TIM) describes measurements $\bar{\mathbf{b}}_{ij}$ affected by bounded noise $\|\epsilon_{ij}\| \leq \beta_{ij}$ including potential outliers (when $\mathbf{o}_{ij} \neq \mathbf{0}$). Again, we estimate \mathbf{R} from the estimated scale \hat{s} , the measurements $\bar{\mathbf{b}}_{ij}$ and the bounds β_{ij} using a TLS estimator:

$$\hat{\mathbf{R}} = \arg \min_{\mathbf{R} \in \text{O}(3)} \sum_{k=1}^K \min \left(\frac{\|\bar{\mathbf{b}}_k - \hat{s} \mathbf{R} \bar{\mathbf{a}}_k\|^2}{\beta_k^2}, \bar{c}^2 \right) \quad (9)$$

where for simplicity we numbered the measurements from 1 to $K = |\mathcal{E}|$ and adopted the notation $\bar{\mathbf{a}}_k, \bar{\mathbf{b}}_k$ instead of $\bar{\mathbf{a}}_{ij}, \bar{\mathbf{b}}_{ij}$. We have also relaxed $\mathbf{R} \in \text{SO}(3)$ to $\mathbf{R} \in \text{O}(3)$, where $\text{O}(3) \doteq \{\mathbf{R} \in \mathbb{R}^{3 \times 3} : \mathbf{R}^T \mathbf{R} = \mathbf{I}_3\}$ is the *Orthogonal Group*, which includes proper rotations and reflections. For simplicity of notation, in the following we drop \hat{s} and assume that $\bar{\mathbf{a}}_1, \dots, \bar{\mathbf{a}}_K$ have been corrected by the scale ($\bar{\mathbf{a}}_k \leftarrow \hat{s} \bar{\mathbf{a}}_k$).

A fundamental contribution of this paper is to develop a tight convex relaxation for (9). The relaxation is tight even in presence of a large number (90%) of outliers and provides per-instance suboptimality guarantees. Before presenting the relaxation, we introduce a binary formulation that is instrumental to develop the proposed relaxation.

Binary formulation and Binary cloning. The first insight behind our convex relaxation is the fact that we can write the TLS cost (9) in additive form using auxiliary binary variables (a property recently leveraged in a different context by [35]):

$$\min_{\substack{\mathbf{R} \in \text{O}(3), \\ \theta_k \in \{-1, +1\}, \forall k}} \sum_{k=1}^K \frac{(1 + \theta_k)}{2} \frac{\|\bar{\mathbf{b}}_k - \mathbf{R} \bar{\mathbf{a}}_k\|^2}{\beta_k^2} + \frac{(1 - \theta_k)}{2} \bar{c}^2 \quad (10)$$

The equivalence can be easily understood from the fact that $\min(x, y) = \min_{\theta \in \{-1, +1\}} \frac{(1+\theta)}{2} x + \frac{(1-\theta)}{2} y$.

We conveniently rewrite (10) by replacing the binary variables with suitable (orthogonal) matrices.

Proposition 4 (Binary cloning). *Problem (10) is equivalent to the following optimization problem*

$$\begin{aligned} \min_{\mathbf{R}, \mathbf{R}_k, \forall k} \sum_{k=1}^K \frac{\|\bar{\mathbf{b}}_k - \mathbf{R} \bar{\mathbf{a}}_k + \mathbf{R}_k^T \bar{\mathbf{b}}_k - \mathbf{R}_k \bar{\mathbf{a}}_k\|^2}{4\beta_k^2} + \frac{(1 - \mathbf{e}_1^T \mathbf{R}_k^T \mathbf{R}_k \mathbf{e}_1)}{2} \bar{c}^2 \\ \text{subject to} \quad \mathbf{R}^T \mathbf{R} = \mathbf{I}_3, \quad \mathbf{R}_k^T \mathbf{R}_k = \mathbf{I}_3, \\ \mathbf{R}^T \mathbf{R}_k \in \{-\mathbf{I}, +\mathbf{I}\}, \quad k = 1, \dots, K \end{aligned} \quad (11)$$

where we introduced a matrix $\mathbf{R}_k \in \mathbb{R}^{3 \times 3}$ for each $k = 1, \dots, K$, and defined the vector $\mathbf{e}_1 \doteq [1 \ 0 \ 0]^T$.

A formal proof of Proposition 4 is given in the Supplementary Material. We name the re-parametrization in Proposition 4 *binary cloning*, since we now have K clones of \mathbf{R} (namely $\mathbf{R}_k = \theta_k \mathbf{R} \in \{\mathbf{R}, -\mathbf{R}\}$, $k = 1, \dots, K$) that are in charge of rejecting outliers: when $\mathbf{R}_k = \mathbf{R}$ the k -th term in the objective becomes $\frac{\|\bar{\mathbf{b}}_k - \mathbf{R} \bar{\mathbf{a}}_k\|^2}{\beta_k^2}$ (i.e., k is treated as an inlier, similarly to choosing $\theta_k = +1$), while when $\mathbf{R}_k = -\mathbf{R}$ the k -th term is equal to \bar{c}^2 (i.e., k is treated as an outlier, similarly to choosing $\theta_k = -1$). This reparametrization enables our relaxation.

Convex relaxation. The proposed relaxation is presented in Proposition 5. The main goal of this paragraph is to provide

the intuition behind our relaxation, while the interested reader can find a formal derivation in the Supplementary Material.

Let us define a $3 \times 3(K+2)$ matrix $\mathbf{X} = [\mathbf{I}_3 \ \mathbf{R} \ \mathbf{R}_1 \ \dots \ \mathbf{R}_K]$, stacking all unknown variables in (11). We observe that the matrix $\mathbf{Z} \doteq \mathbf{X}^\top \mathbf{X}$ contains all linear and quadratic terms in \mathbf{R} and \mathbf{R}_k :

$$\mathbf{Z} \doteq \mathbf{X}^\top \mathbf{X} = \begin{bmatrix} \overbrace{\mathbf{I}_3}^I & \overbrace{\mathbf{R}}^R & \overbrace{\mathbf{R}_1}^{R_1} & \dots & \overbrace{\mathbf{R}_K}^{R_K} \\ * & \mathbf{I}_3 & \mathbf{R}^\top \mathbf{R}_1 & \dots & \mathbf{R}^\top \mathbf{R}_K \\ * & * & \mathbf{I}_3 & \dots & \mathbf{R}_1^\top \mathbf{R}_K \\ * & \vdots & \vdots & \ddots & \vdots \\ * & * & * & \dots & \mathbf{I}_3 \end{bmatrix} \quad (12)$$

Now it is easy to see that the cost function in (11) can be rewritten as a function of \mathbf{Z} , by noting that $\mathbf{R}, \mathbf{R}^\top \mathbf{R}_k, \mathbf{R}_k$ (appearing in the cost) are entries of \mathbf{Z} , see (12). The constraints in (11) can be similarly written as a function of \mathbf{Z} . For instance, the constraints $\mathbf{R}^\top \mathbf{R} = \mathbf{I}_3$ and $\mathbf{R}_k^\top \mathbf{R}_k = \mathbf{I}_3$ simply enforce that the block diagonal entries of \mathbf{Z} are identity matrices. Similarly, $\mathbf{R}^\top \mathbf{R}_k \in \{-\mathbf{I}, +\mathbf{I}\}$ can be rewritten as a (non-convex) constraint involving off-diagonal entries of \mathbf{Z} . Finally, the fact that $\mathbf{Z} \doteq \mathbf{X}^\top \mathbf{X}$ implies \mathbf{Z} is positive semidefinite and has rank 3 (number of rows in \mathbf{X} , see (12)).

According to the discussion so far, we can reparametrize problem (11) using \mathbf{Z} , and we can then develop a convex relaxation by relaxing all the resulting non-convex constraints. This is formalized in the following proposition.

Proposition 5 (TLS Rotation Estimation: Convex Relaxation). *The following convex program is a relaxation of (11):*

$$\begin{aligned} & \min_{\mathbf{Z} \succeq 0} \quad \text{tr}(\bar{\mathbf{Q}}\mathbf{Z}) & (13) \\ \text{subject to} \quad & [\mathbf{Z}]_{RR} = \mathbf{I}_3, \quad [\mathbf{Z}]_{R_k R_k} = \mathbf{I}_3, \quad [\mathbf{Z}]_{II} = \mathbf{I}_3, \\ & [\mathbf{Z}]_{RR_k} = (\mathbf{e}_1^\top [\mathbf{Z}]_{RR_k} \mathbf{e}_1) \mathbf{I}_3, \quad \forall k \\ & \|[\mathbf{Z}]_{IR} \pm [\mathbf{Z}]_{IR_k}\| \leq 1 \pm (\mathbf{e}_1^\top [\mathbf{Z}]_{RR_k} \mathbf{e}_1), \quad \forall k \\ & \|[\mathbf{Z}]_{IR_k} \pm [\mathbf{Z}]_{IR_{k'}}\| \leq 1 \pm (\mathbf{e}_1^\top [\mathbf{Z}]_{R_k R_{k'}} \mathbf{e}_1), \quad \forall k, k' \end{aligned}$$

where $\bar{\mathbf{Q}}$ is a known $3(K+2) \times 3(K+2)$ symmetric matrix (expression given in the supplementary), and $[\mathbf{Z}]_{RR_k}$ denotes a 3×3 block of \mathbf{Z} whose row indices correspond to the location of \mathbf{R} in \mathbf{X} (cf. with indices at the top of the matrix in eq. (12)) and column indices correspond to the location of \mathbf{R}_k in \mathbf{X} (similarly, for $[\mathbf{Z}]_{RR}, [\mathbf{Z}]_{R_k R_k}, [\mathbf{Z}]_{II}$, etc.).

The convex program (13) can be solved in polynomial time using off-the-shelf convex solvers, such as `cvx` [23]. It is a relaxation, in the sense that the set of feasible solutions of (13) includes the set of feasible solutions of (11). Moreover, it enjoys the typical per-instance guarantees of convex relaxations.

Proposition 6 (Guarantees for TLS Rotation Estimation). *Let \mathbf{Z}^* be the optimal solution of the relaxation (13). If \mathbf{Z}^* has rank 3, then it can be factored as $\mathbf{Z}^* = (\mathbf{X}^*)^\top (\mathbf{X}^*)$, where $\mathbf{X}^* \doteq [\mathbf{I}_3, \mathbf{R}^*, \mathbf{R}_1^*, \dots, \mathbf{R}_K^*]$ is the first block row of \mathbf{Z}^* . Moreover, $\mathbf{R}^*, \mathbf{R}_1^*, \dots, \mathbf{R}_K^*$ is an optimal solution for (11).*

Empirically, we found that our relaxation is tight (i.e., numerically produces a rank-3 solution) even when 90% of the TIMs are outliers. Even when the relaxation is not tight, one can still project \mathbf{Z}^* to a feasible solution of (11) and obtain an upper-bound on how suboptimal the resulting solution is.

In summary, the function `estimate_R` in Algorithm 1 solves the convex program (13) (e.g., using `cvx`) to obtain a matrix \mathbf{Z}^* and extracts the rotation estimate $\hat{\mathbf{R}}$ from \mathbf{Z}^* . In particular, $\hat{\mathbf{R}} = [\mathbf{Z}^*]_{IR}$ if \mathbf{Z}^* has rank 3 (Proposition 6), or $\hat{\mathbf{R}}$ is computed as the projection of $[\mathbf{Z}^*]_{IR}$ to $O(3)$ otherwise.

C. Robust Translation Estimation

Since we already presented a polynomial-time solution for scalar TLS in Section V-A, we propose to solve for the translation component-wise, i.e., we compute the entries t_1, t_2, t_3 of \mathbf{t} independently (see the Supplementary Material for details):

$$\min_{t_j} \sum_{i=1}^N \min \left(\frac{1}{\beta_i^2} \left| [\mathbf{b}_i - \hat{s} \mathbf{R} \mathbf{a}_i]_j - t_j \right|^2, \bar{c}^2 \right), \quad j = 1, 2, 3 \quad (14)$$

where $[\cdot]_j$ denotes the j -th entry of a vector.

In summary, the function `estimate_t` in Algorithm 1 calls Algorithm 2 three times (one for each entry of \mathbf{t}) and returns the translation estimate $\hat{\mathbf{t}} = [t_1 \ t_2 \ t_3]$.

VI. EXPERIMENTS AND APPLICATIONS

The goal of this section is to (i) test the performance of our scale, rotation, translation solvers and the MCIS pruning (Section VI-A), (ii) evaluate TEASER against related techniques in benchmarking datasets (Section VI-B), (iii) evaluate TEASER with *extreme* outliers rates (Section VI-C), and (iv) show an application of TEASER for object localization in an RGB-D robotics dataset (Section VI-D). In all tests we set $\bar{c}^2 = 1$.

Implementation details. We implemented TEASER in matlab and used `cvx` to solve the convex relaxation (13). Moreover, we used the algorithm in [20] to find all the maximal cliques in the pruned TIM graph (see Theorem 3).

A. Testing TEASER's Subproblems

Testing setup. We use the Bunny point cloud from the Stanford 3D Scanning Repository [18] and resize it to be within the $[0, 1]^3$ cube. The Bunny is first downsampled to $N = 50$ points, and then a random transformation $(s, \mathbf{R}, \mathbf{t})$ (with $1 \leq s \leq 5$ and $\|\mathbf{t}\| \leq 1$) is applied according to eq. (1). To generate the bounded noise ϵ_i , we sample $\epsilon_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, until the resulting vector satisfies $\|\epsilon_i\| \leq \beta_i = \beta$. We set $\sigma = 0.01$ and $\beta = 0.0554$ such that $\mathbb{P}(\|\epsilon_i\|^2 / \sigma^2 > \beta^2) \leq 10^{-6}$ (this bound stems from the fact that for Gaussian ϵ_i , $\|\epsilon_i\|^2$ follows a Chi-square distribution with 3 degrees of freedom). To generate outliers, we replace a fraction of \mathbf{b}_i with vectors uniformly sampled inside the sphere of radius 5. We test increasing outlier ratios $\{0, 0.2, 0.4, 0.6, 0.7, 0.8, 0.9\}$. All statistics are computed over 40 Monte Carlo runs.

Scale solver. Given two point clouds \mathcal{A} and \mathcal{B} , we first create $N(N-1)/2$ TIMs corresponding to a complete graph and then use Algorithm 2 to solve for the scale. We compute both *maximum consensus* [49] and TLS estimates of the scale.

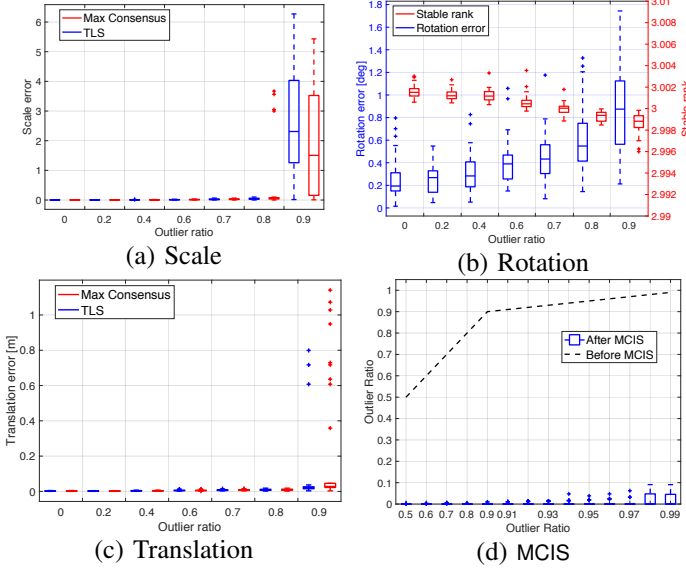


Fig. 4. Results for scale, rotation, translation estimation, and impact of maximal clique pruning for increasing outlier ratios.

Fig. 4(a) shows box plots of the scale error with increasing outlier ratios. The scale error is computed as $|s^* - s_{gt}|$, where s^* is the scale estimate and s_{gt} is the ground-truth. We observe the TLS solver is robust against 80% outliers, while maximum consensus failed three times in that regime.

Rotation Solver. We apply a random rotation \mathbf{R} to the Bunny, and fix $s = 1$ and $\mathbf{t} = \mathbf{0}$. Two metrics are boxplotted in Fig. 4(b) to show the performance of the rotation solver: (i) the *stable rank* of \mathbf{Z}^* , the optimal solution of SDP relaxation (13), where the *stable rank* is defined by the squared ratio between the Frobenius norm and the spectral norm; (ii) the rotation estimation error, defined as $|\arccos((\text{tr}(\mathbf{R}_{gt}^T \mathbf{R}^*) - 1)/2)|$, *i.e.* the geodesic distance between the rotation estimate \mathbf{R}^* and the ground-truth \mathbf{R}_{gt} . We observe the stable rank is numerically close to 3 (Proposition 5) even with 90% of outliers, and the rotation error remains below 2 degrees.

Translation Solver. We apply a random translation \mathbf{t} to the Bunny, and fix $s = 1$ and $\mathbf{R} = \mathbf{I}_3$. Fig. 4(c) shows component-wise translation estimation using both *maximum consensus* and TLS are robust against 80% outliers. The translation error is defined as $\|\mathbf{t}^* - \mathbf{t}_{gt}\|$, the 2-norm of the difference between the estimate \mathbf{t}^* and the ground-truth \mathbf{t}_{gt} .

Maximal Clique Inlier Selection. We downsample Bunny to $N = 1000$ and fix the scale to $s = 1$ when applying the random transformation. We first prune the outlier TIMs/TRIMs (edges) that are not consistent with the scale $s = 1$, while keeping all the points (nodes), to obtain the graph \mathcal{G}' . Then we compute the *maximum clique* in \mathcal{G}' using the algorithm in [20], and remove all edges and nodes outside the clique, obtaining a pruned graph \mathcal{G}'' . Fig. 4(d) shows the outlier ratio in \mathcal{G}' (label: “Before MCIS”) and \mathcal{G}'' (label: “After MCIS”). The MCIS procedure effectively reduces the amount of outliers to below 10%, facilitating rotation and translation estimation, which, in isolation, can already tolerate more than 90% outliers.

B. Benchmarking on Standard Datasets

Testing setup. We benchmark TEASER against two state-of-the-art robust registration techniques: *Fast Global Registration* (FGR) [59] and *Guaranteed Outlier REmoval* (GORE) [11]. In addition, we test two RANSAC variants: a fast version where we terminate RANSAC after a maximum of 1,000 iterations (RANSAC (1K)) and a slow version where we terminate RANSAC after 60s (RANSAC (60s)). Four datasets, Bunny, Armadillo, Dragon and Buddha, from the Stanford 3D Scanning Repository are selected and downsampled to $N = 100$ points. The tests below follow the same protocol of Section VI-A. In the Supplementary Material, we provide an example of the performance of TEASER on registration problems with high noise ($\sigma = 0.1$).

Known Scale. We first evaluate the compared techniques with known scale $s = 1$. Fig. 5(a) shows the rotation and translation error at increasing outlier ratios for the Bunny dataset. TEASER, GORE and RANSAC are robust against up to 90% outliers, although TEASER tends to produce more accurate estimates than GORE, and RANSAC typically requires over 10^5 iterations for convergence at 90% outlier rate. FGR can only resist 70% outliers and RANSAC (1K) starts breaking at 60% outlier rate. These conclusions are confirmed by the results on the other three datasets (Armadillo, Dragon, Buddha), which are given in the Supplementary Material due to space constraints.

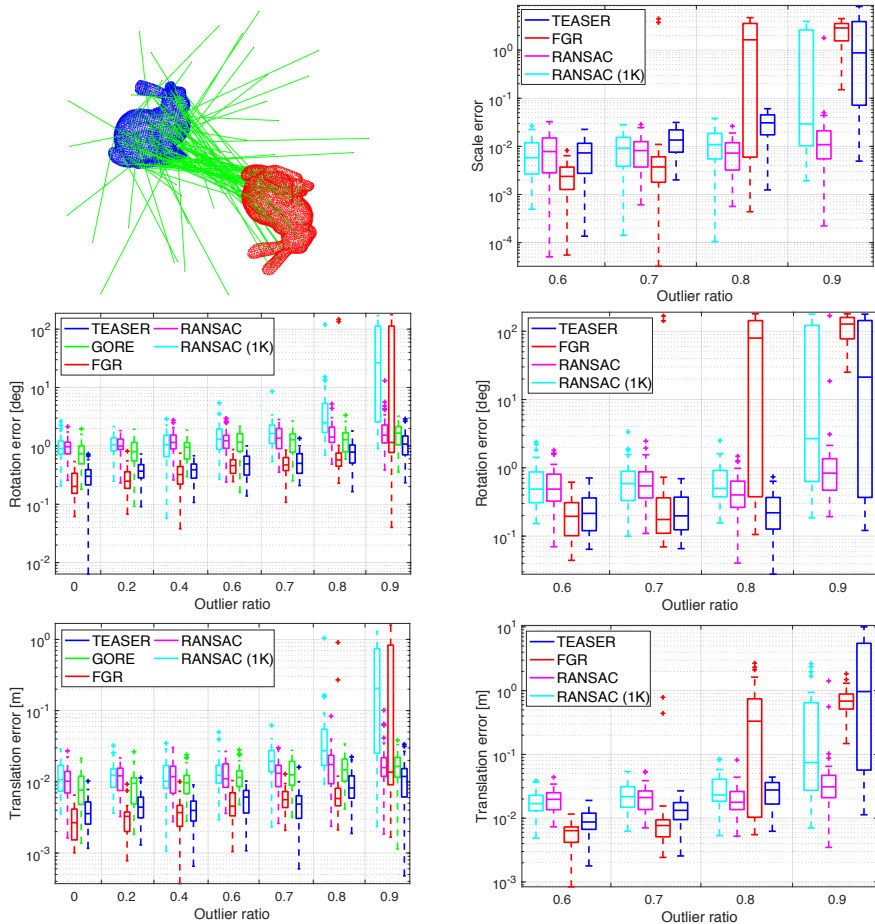
Unknown Scale. GORE is unable to solve for the scale, hence we only benchmark TEASER against FGR (although the original algorithm in [59] did not solve for the scale, we extend it by using Horn’s method to compute the scale at each iteration), RANSAC (1K) and RANSAC (60s). Fig. 5(b) plots the scale, rotation and translation error for increasing outlier ratios on the Bunny dataset. All the compared techniques perform well when the outlier ratio is below 60%. FGR has the lowest breakdown point and fails at 80%. RANSAC (1K) and TEASER only fail at 90% outlier ratio when the scale is unknown. Although RANSAC with 60s timeout outperforms other methods at 90% outlier rate, it typically requires more than 10^5 iterations to converge, which is not practical for real-time applications.

C. Testing under Extreme Outlier Rates

We further benchmark the performance of TEASER under extreme outlier rates from 95% to 99% with known scale and $N = 1000$ correspondences on the Bunny. We replace RANSAC (1K) with RANSAC (10K), since RANSAC (1K) already performs poorly at 90% outlier ratio. Fig. 6 shows the boxplots of the rotation and translation errors. Both TEASER and GORE are robust against up to 99% outliers, while RANSAC with 60s timeout can resist 98% outliers with about 10^6 iterations. RANSAC (10K) and FGR perform poorly under extreme outlier ratios. While GORE and TEASER are both robust against 99% outliers, TEASER produces slightly lower estimation errors.

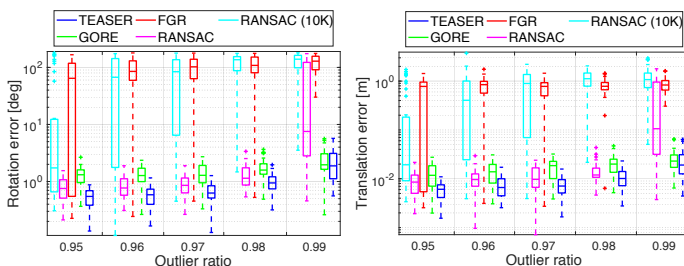
D. Application: Object Pose Estimation and Localization

We use the large-scale point cloud datasets from [34] to test TEASER in *object pose estimation and localization* applications. We first use the ground-truth object labels to extract the *cereal box/cap* out of the scene and treat it as the object, then



(a) Registration with known scale (b) Registration with unknown scale (c) Registration on robotics datasets

Fig. 5. Benchmark result. (a) Boxplots of rotation and translation errors for the five compared methods on the Bunny dataset with known scale (the top figure shows an example with 50% outliers). (b) Boxplots of scale, rotation and translation errors for four registration methods on the Bunny dataset with unknown scale. (c) Successful object pose estimation by TEASER on a real RGB-D dataset. Blue lines are the original FPFH [47] correspondences with outliers, green lines are the inlier correspondences computed by TEASER, and the final registered object is highlighted in red.



(a) Rotation Error (b) Translation Error

Fig. 6. Estimation errors under extreme outlier rates (known scale).

apply a random transformation to the scene, to get an object-scene pair. To register the object-scene pair, we first use FPFH feature descriptors [47] to establish putative correspondences. Then, TEASER is used to find the relative pose. Fig. 5(c) shows the noisy FPFH correspondences, the inlier correspondences obtained by TEASER, and successful localization and pose estimation of the *cereal box*. The Supplementary Material provides results on more object-scene pairs.

VII. CONCLUSION

We propose a *Truncated Least Squares* approach to compute the relative transformation (scale, rotation, translation) that

aligns two point clouds in the presence of extreme outlier rates. We present a general graph-theoretic framework to decouple rotation, translation, and scale estimation. We provide a polynomial-time solution for each subproblem: scale and (component-wise) translation estimation can be solved exactly via an *adaptive voting* scheme, while rotation estimation can be relaxed to a semidefinite program. The resulting polynomial-time approach, named TEASER (*Truncated least squares Estimation And SEMidefinite Relaxation*), outperforms RANSAC and robust local optimization techniques, and favorably compares with BnB methods. TEASER can tolerate up to 99% outliers and returns highly-accurate solutions.

While running in polynomial time, the general-purpose SDP solver used in our current implementation scales poorly in the problem size. Current research effort is devoted to developing specialized SDP solvers to allow using TEASER to solve large-scale registration problems in real-time.

ACKNOWLEDGMENTS

This work was partially funded by ARL DCIST CRA W911NF-17-2-0181, ONR RAIDER N00014-18-1-2828, and the Google Daydream Research Program.

REFERENCES

- [1] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Trans. Pattern Anal. Machine Intell.*, 9(5):698–700, sept. 1987.
- [2] M. A. Audette, F. P. Ferrie, and T. M. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Med. Image Anal.*, 4(3):201–217, 2000.
- [3] J. C. Bazin, Y. Seo, R. I. Hartley, and M. Pollefeys. Globally optimal inlier set maximization with unknown rotation and focal length. In *European Conf. on Computer Vision (ECCV)*, pages 803–817, 2014.
- [4] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(2), 1992.
- [5] M. J. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *Intl. J. of Computer Vision*, 19(1):57–91, 1996.
- [6] G. Blais and M. D. Levine. Registering multiview range data to create 3d computer objects. *IEEE Trans. Pattern Anal. Machine Intell.*, 17(8):820–824, 1995.
- [7] S. Bouaziz, A. Tagliasacchi, and M. Pauly. Sparse iterative closest point. In *ACM Symp. Geom. Process.*, pages 113–123. Eurographics Association, 2013.
- [8] T. M. Breuel. Implementation techniques for geometric branch-and-bound matching methods. *Comput. Vis. Image Underst.*, 90(3):258–294, 2003.
- [9] J. Briales and J. Gonzalez-Jimenez. Convex Global 3D Registration with Lagrangian Duality. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [10] C. Bron and J. Kerbosch. Algorithm 457: finding all cliques of an undirected graph. *Communications of the ACM*, 16(9):575–577, 1973.
- [11] Á. P. Bustos and T. J. Chin. Guaranteed outlier removal for point cloud registration with correspondences. *IEEE Trans. Pattern Anal. Machine Intell.*, 40(12):2868–2882, 2018.
- [12] A. Parra Bustos, T. J. Chin, and D. Suter. Fast rotation search with stereographic projections for 3d registration. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3930–3937, 2014.
- [13] D. Campbell and L. Petersson. An adaptive data representation for robust point-set registration and merging. In *Intl. Conf. on Computer Vision (ICCV)*, pages 4292–4300, 2015.
- [14] D. Campbell and L. Petersson. Gogma: Globally-optimal gaussian mixture alignment. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5685–5694, 2016.
- [15] C. S. Chen, Y. P. Hung, and J. B. Cheng. RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. Pattern Anal. Machine Intell.*, 21(11):1229–1234, 1999.
- [16] D. Chetverikov, D. Stepanov, and P. Krsek. Robust euclidean alignment of 3d point sets: the trimmed iterative closest point algorithm. *Image and Vision Computing*, 23(3):299–309, 2005.
- [17] S. Choi, Q. Y. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5556–5565, 2015.
- [18] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH*, pages 303–312, 1996.
- [19] B. Drost, M. Ulrich, N. Navab, and S. Ilic. Model globally, match locally: Efficient and robust 3D object recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 998–1005, 2010.
- [20] D. Eppstein, M. Löffler, and D. Strash. Listing all maximal cliques in sparse graphs in near-optimal time. In *International Symposium on Algorithms and Computation*, pages 403–414. Springer, 2010.
- [21] M. Fischler and R. Bolles. Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography. *Commun. ACM*, 24: 381–395, 1981.
- [22] S. Granger and X. Pennec. Multi-scale EM-ICP: A fast and robust approach for surface registration. In *European Conf. on Computer Vision (ECCV)*, 2002.
- [23] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming. URL <http://cvxr.com/cvx>.
- [24] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, and J. Wan. 3D object recognition in cluttered scenes with local surface features: a survey. *IEEE Trans. Pattern Anal. Machine Intell.*, 36(11):2270–2287, 2014.
- [25] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
- [26] R. I. Hartley and F. Kahl. Global optimization through rotation space search. *Intl. J. of Computer Vision*, 82(1): 64–79, 2009.
- [27] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. Rgb-d mapping: Using kinect-style depth cameras for dense 3d modeling of indoor environments. *Intl. J. of Robotics Research*, 31(5):647–663, 2012.
- [28] B. K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. Amer.*, 4(4): 629–642, Apr 1987.
- [29] G. Izatt, H. Dai, and R. Tedrake. Globally optimal object pose estimation in point clouds with mixed-integer programming. In *Proc. of the Intl. Symp. of Robotics Research (ISRR)*, 2017.
- [30] B. Jian and B. C. Vemuri. A robust algorithm for point set registration using mixture of gaussians. In *Intl. Conf. on Computer Vision (ICCV)*, volume 2, pages 1246–1251. IEEE, 2005.
- [31] B. Jian and B. C. Vemuri. Robust point set registration using gaussian mixture models. *IEEE Trans. Pattern Anal. Machine Intell.*, 33(8):1633–1645, 2011.
- [32] S. Kaneko, T. Kondo, and A. Miyamoto. Robust matching of 3d contours using iterative closest point algorithm

- improved by m-estimation. *Pattern Recognition*, 36(9): 2041–2047, 2003.
- [33] H.S. Koppula, A. Anand, T. Joachims, and A. Saxena. Semantic labeling of 3d point clouds for indoor scenes. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [34] K. Lai, L. Bo, X. Ren, and D. Fox. A large-scale hierarchical multi-view RGB-D object dataset. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1817–1824. IEEE, 2011.
- [35] P. Lajoie, S. Hu, G. Beltrame, and L. Carlone. Modeling perceptual aliasing in SLAM via discrete-continuous graphical models. *IEEE Robotics and Automation Letters*, 2019.
- [36] X. Li, Y. Liu, Y. Wang, C. Wang, M. Wang, and Z. Song. Fast and globally optimal rigid registration of 3d point sets by transformation decomposition. *arXiv preprint arXiv:1812.11307*, 2019.
- [37] Y. Liu, C. Wang, Z. Song, and M. Wang. Efficient global point cloud registration by matching rotation invariant features through translation search. In *European Conf. on Computer Vision (ECCV)*, September 2018.
- [38] L. Maier-Hein, A. M. Franz, T. R. dos Santos, M. Schmidt, M. Fangerau, H. P. Meinzer, and J. M. Fitzpatrick. Convergent iterative closest-point algorithm to accommodate anisotropic and inhomogeneous localization error. *IEEE Trans. Pattern Anal. Machine Intell.*, 34(8): 1520–1532, 2012.
- [39] A. Makadia, A. Patterson, and K. Daniilidis. Fully automatic registration of 3d point clouds. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 1297–1304, 2006.
- [40] P. Marion, P. R. Florence, L. Manuelli, and R. Tedrake. Label fusion: A pipeline for generating ground truth labels for real rgbd data of cluttered scenes. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018.
- [41] P. Meer, D. Mintz, A. Rosenfeld, and D. Y. Kim. Robust regression methods for computer vision: A review. *Intl. J. of Computer Vision*, 6(1):59–70, Apr 1991.
- [42] M. Milanese. *Estimation and Prediction in the Presence of Unknown but Bounded Uncertainty: A Survey*, pages 3–24. Springer US, Boston, MA, 1989.
- [43] A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE Trans. Pattern Anal. Machine Intell.*, 32(12):2262–2275, 2010.
- [44] C. Olsson, F. Kahl, and M. Oskarsson. Branch-and-bound methods for euclidean registration problems. *IEEE Trans. Pattern Anal. Machine Intell.*, 31(5):783–794, 2009.
- [45] B. Pattabiraman, M. M. A. Patwary, A. H. Gebremedhin, W. K. Liao, and A. Choudhary. Fast algorithms for the maximum clique problem on massive graphs with applications to overlapping community detection. *Internet Mathematics*, 11(4-5):421–448, 2015.
- [46] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz. Aligning point cloud views using persistent feature histograms. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 3384–3391. IEEE, 2008.
- [47] R.B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3212–3217. Citeseer, 2009.
- [48] R. Sandhu, S. Dambreville, and A. Tannenbaum. Point set registration via particle filtering and stochastic dynamics. *IEEE Trans. Pattern Anal. Machine Intell.*, 32(8):1459–1473, 2010.
- [49] P. Speciale, D. P. Paudel, M. R. Oswald, T. Kroeger, L. V. Gool, and M. Pollefeys. Consensus maximization with linear matrix inequality constraints. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5048–5056, July 2017. doi: 10.1109/CVPR.2017.536.
- [50] G. K. L. Tam, Z. Q. Cheng, Y. K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X. F. Sun, and P. L. Rosin. Registration of 3d point clouds and meshes: a survey from rigid to nonrigid. *IEEE Trans. Vis. Comput. Graph.*, 19(7):1199–1217, 2013.
- [51] K. Mac Tavish and T. D. Barfoot. At all costs: A comparison of robust cost functions for camera correspondence outliers. In *Computer and Robot Vision (CRV), 2015 12th Conference on*, pages 62–69. IEEE, 2015.
- [52] F. Tombari, S. Salti, and L. Di Stefano. Performance evaluation of 3d keypoint detectors. *Intl. J. of Computer Vision*, 102(1-3):198–220, 2013.
- [53] J. M. Wong, V. Kee, T. Le, S. Wagner, G. L. Mariottini, A. Schneider, L. Hamilton, R. Chipalkatty, M. Hebert, D. M. S. Johnson, et al. Segicp: Integrated deep semantic segmentation and pose estimation. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5784–5789. IEEE, 2017.
- [54] Q. Wu and J.K. Hao. A review on algorithms for maximum clique problems. *European Journal of Operational Research*, 242(3):693–709, 2015.
- [55] J. Yang, H. Li, D. Campbell, and Y. Jia. Go-ICP: A globally optimal solution to 3D ICP point-set registration. *IEEE Trans. Pattern Anal. Machine Intell.*, 38(11):2241–2254, November 2016. ISSN 0162-8828.
- [56] A. Zeng, K. T. Yu, S. Song, D. Suo, E. Walker, A. Rodriguez, and J. Xiao. Multi-view self-supervised deep learning for 6d pose estimation in the amazon picking challenge. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1386–1383. IEEE, 2017.
- [57] C. Zhang, S. Du, J. Liu, and J. Xue. Robust 3d point set registration using iterative closest point algorithm with bounded rotation angle. *Signal Processing*, 120:777–788, 2016.
- [58] J. Zhang and S. Singh. Visual-lidar odometry and mapping: Low-drift, robust, and fast. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 2174–2181. IEEE, 2015.
- [59] Q. Y. Zhou, J. Park, and V. Koltun. Fast global registration. In *European Conf. on Computer Vision (ECCV)*, pages 766–782. Springer, 2016.