# Outlier-Robust Spatial Perception: Hardness, General-Purpose Algorithms, and Guarantees

Vasileios Tzoumas,* Pasquale Antonante,* Luca Carlone

***Abstract*— Spatial perception is the backbone of many robotics applications, and spans a broad range of research problems, including localization and mapping, point cloud alignment, and relative pose estimation from camera images. Robust spatial perception is jeopardized by the presence of incorrect data association, and in general, outliers. Although techniques to handle outliers do exist, they can fail in unpredictable manners (e.g., RANSAC, robust estimators), or can have exponential runtime (e.g., branch-and-bound). In this paper, we advance the state of the art in outlier rejection by making three contributions. First, we show that even a simple linear instance of outlier rejection is *inapproximable*: in the worst-case one cannot design even a quasi-polynomial time algorithm that computes an approximate solution efficiently. Our second contribution is to provide the first per-instance sub-optimality bounds to assess the approximation quality of a given outlier rejection outcome. Our third contribution is to propose a simple general-purpose algorithm, named *adaptive trimming*, to remove outliers. Our algorithm leverages recently-proposed global solvers that are able to solve outlier-free problems, and iteratively removes measurements with large errors. We demonstrate the algorithm on three spatial perception problems: 3D registration, two-view geometry, and SLAM. The results show that our algorithm outperforms several state-of-the-art methods across applications while being a general-purpose method.**

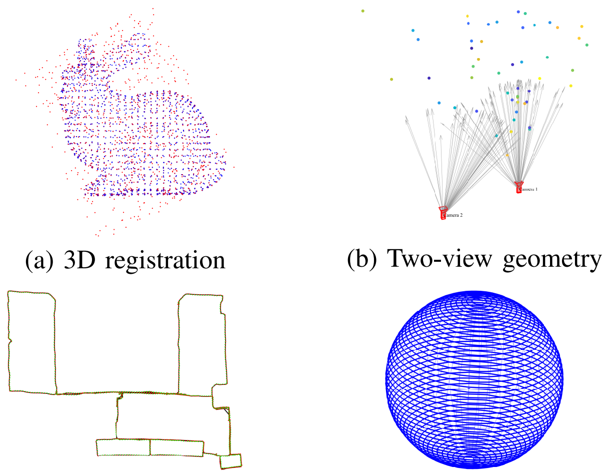

(a) 3D registration (b) Two-view geometry

(c) 2D SLAM (d) 3D SLAM

Fig. 1. We investigate outlier rejection across multiple spatial perception problems, including (a) 3D registration, (b) two-view geometry, and (c-d) SLAM. We provide inapproximability results and performance bounds. We also propose an algorithm, ADAPT, that outperforms RANSAC and other specialized methods. ADAPT tolerates up to 90% outliers in 3D registration, and up to 50% outliers in two-view geometry and most SLAM datasets.

## I. Introduction

*Spatial perception* is concerned with the estimation of a geometric model that describes the state of the robot, and/or the environment the robot is deployed in. As such, spatial perception includes a broad set of robotics problems, including motion estimation [1], object detection, localization and tracking [2], multi-robot localization [3], dense reconstruction [4], and Simultaneous Localization and Mapping (SLAM) [5]. Spatial perception algorithms find applications beyond robotics, including virtual and augmented reality, and medical imaging [2], to mention a few.

Safety-critical applications, including self-driving cars, demand robust spatial perception algorithms that can estimate correct models (and assess their performance) in the presence of measurement noise and outliers. While we currently have several approaches that can tolerate large measurement noise (e.g., [6], [7], [8]), these algorithm tend to catastrophically fail in the presence of outliers resulting from incorrect data association, sensor malfunction, or even adversarial attacks.

In this paper, we focus on the analysis and design of *outlier-robust general-purpose* algorithms for robust estimation applied to spatial perception. Our proposal is motivated by three observations. First, recent years have seen a convergence of the robotics community towards optimization-based approaches for spatial perception. Therefore, despite the apparent heterogeneity of the perception landscape, it is possible to develop *general-purpose* methods to reject outliers (e.g., M-estimators [9] and *consensus maximization* [10] can be thought as general estimation tools). Second, the research community has developed global solutions to many perception problems *without* outliers, from well-established techniques for point cloud registration [8], to very recent solvers for SLAM [6] and two-view geometry [7]. These global solvers offer unprecedented opportunities to tackle robust estimation *with* outliers. Third, the literature still lacks a satisfactory answer to provably-robust spatial perception.

The literature on outlier-robust spatial perception is currently divided between *fast* approaches (that mainly work in the low-outlier regime, without performance guarantees) and *provably-robust* approaches (that can tolerate many outliers, but have exponential runtime). While we postpone a comprehensive literature review to [11], it is instructive to briefly review this dichotomy. *Fast* approaches include RANSAC [12], M-estimators [9], and measurement-consistency checking [13], [14]. These methods fall short of providing performance guarantees. In particular, RANSAC is known to become slow and brittle with high outlier rates ($>50\%$) [10], and does not scale to high-dimensional problems, while M-estimators have a breakdown point of zero, meaning that a single "bad" outlier can compromise the results. On

*The authors contributed equally to this work.

V. Tzoumas, P. Antonante, L. Carlone are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.
{vtzoumas,antonap,lcarlone}@mit.edu

the other hand, provably-robust methods, typically based on *branch-and-bound* [15], [16], [17], [18], [10], can tolerate more than $50\%$ of outliers [19], but do not scale to large problems and are relatively slow for robotics applications. Overall, the first goal of this paper is to understand whether we can resolve this divide, and design algorithms that are both efficient and provably robust.

**Contributions.** We propose a *Minimal Trimmed Squares (MTS)* formulation for outlier-robust estimation. MTS encapsulate a wide spectrum of commonly-used outlier-robust formulations in the literature, such as the popular *maximum consensus* [20], *Linear Trimmed Squares* [21], and *truncated least-squares* [22]. In particular, MTS aims to compute a "good" estimate by rejecting a minimal set of measurements.

Our first contribution (Section III) is a negative result: we show that outlier rejection is *inapproximable*. In the worst-case, there exist no quasi-polynomial algorithm that can compute (even an approximate) solution to the outlier rejection problem. We prove that this remains true, surprisingly, even if the algorithm knows the true number of outliers and even if we allow the algorithm to reject more measurements than necessary. Our conclusions largely extend previously-known negative results [20], which already ruled-out the possibility of designing polynomial-time approximation methods.

Our second contribution (Section IV) is to derive the first per-instance sub-optimality bounds to assess the quality of a given outlier rejection solution. While in the worst case we expect efficient algorithms to perform poorly, we can still hope that in typical problem instances a polynomial-time algorithm can compute good solutions, and we can use the proposed sub-optimality bounds to assess the performance of such an algorithm. Our bounds are algorithm-agnostic (e.g., they also apply to RANSAC) and can be computed efficiently.

Our third contribution (Section V) is a *general-purpose* algorithm for outlier rejection, named *Adaptive Trimming (*ADAPT*)*. ADAPT leverages recently-proposed global solvers that solve outlier-free problems and adaptively removes measurements with large residual errors. Despite its simplicity, our experiments show that it outperforms RANSAC and even specialized state-of-the-art methods for robust estimation.

We conclude the paper by providing an experimental evaluation across multiple spatial perception problems (Section VI). The experiments show that ADAPT can tolerate up to 90% outliers in 3D registration (with a runtime similar to existing methods), and up to 50% outliers in two-view geometry and most SLAM datasets. The experiments also show that the proposed sub-optimality bounds are effective in assessing the outlier rejection outcomes.

We report extra results and proofs in [11].

## II. Outlier Rejection: a Minimally Trimmed Squares Formulation

Many estimation problems in robotics and computer vision can be formulated as non-linear least squares problems:

$$\min_{x\in\mathbb{X}} \sum_{i\in\mathcal{M}} \|h_i(y_i,x)\|^2, \tag{1}$$

where we are given measurements $y_i$ of an unknown variable $x$, with $i \in \mathcal{M}$ ($\mathcal{M}$ is the measurement set), and we want to estimate $x$, potentially restricted to a given domain $\mathbb{X}$ (e.g., $x$ is a pose, and $\mathbb{X}$ is the set of 3D poses). The least squares problem in eq. (1) looks for the $x$ that minimizes the (squares of) the *residual errors* $h_i(y_i,x)$, where the $i$-th residual error captures how well $x$ explains the measurement $y_i$. The problem in eq. (1) typically results from maximum likelihood and maximum a posteriori estimation [5], [23], under the assumption that the measurement noise is Gaussian.

Both researchers and practitioners are well-aware that least squares formulations are sensitive to outliers, and that the estimator in eq. (1) fails to produce a meaningful estimate of $x$ in the presence of gross outliers $y_i$. Therefore, in this paper we address the following question:

*Can we compute an accurate estimate of $x$ that is insensitive to the presence of outlying measurements?*

We formulate the resulting robust estimation problem as one where select a few outliers, such that the remaining measurements (the inliers) can be explained with small error.

*Problem 1 (Minimally Trimmed Squares (MTS)):* Let $\mathcal{M}$ denote a set of measurements of an unknown variable $x$, and let $y_i$ denote the $i$-th measurement. Also denote with $h_i(y_i,x)$ the residual error that quantifies how well $x$ fits the measurement $y_i$. Then, the *minimally trimmed squares* problem consists in estimating the unknown variable $x$ by solving the following optimization problem:

$$\min_{\mathcal{O}\subseteq\mathcal{M}} \min_{x\in\mathbb{X}} |\mathcal{O}|, \text{ s.t. } \sum_{i\in\mathcal{M}\setminus\mathcal{O}} \|h_i(y_i,x)\|^2 \leq \epsilon_{\mathcal{M}\setminus\mathcal{O}}, \tag{2}$$

where one searches for the smallest set of outliers $\mathcal{O}$ ($|\cdot|$ is the cardinality of a set) among the given measurements $\mathcal{M}$, such that the remaining measurements $\mathcal{M}\setminus\mathcal{O}$ (i.e., the inliers) can be explained with small error, i.e., $\sum_{i\in\mathcal{M}\setminus\mathcal{O}}\|h_i(y_i,x)\|^2 \leq \epsilon_{\mathcal{M}\setminus\mathcal{O}}$ for some $x \in \mathbb{X}$, and where $\epsilon_{\mathcal{M}\setminus\mathcal{O}}$ is a given *outlier-free* bound. ⌟

*Remark 1 (Generality and applicability):* In this paper we address robustness in non-linear and non-convex estimation problems as the ones arising in robotics and computer vision. Therefore, the algorithms and bounds presented in this paper hold for any function $h_i(y_i,x)$ and any domain $\mathbb{X}$. In contrast with related work [24], [25], we do not assume the number of outliers to be known in advance (an unrealistic assumption in perception problems). Indeed, MTS looks for the smallest set of outliers. Finally, while the formulation (2) requires to set an outlier-free threshold, we will propose an algorithm (Section V) that automatically computes a suitable threshold without any prior knowledge about the measurement noise. ⌟

In summary, MTS is a general non-linear and non-convex outlier rejection framework. We exemplify its generality by presenting in [11] its explicit form for 3 robust perception problems: 3D registration, two-view geometry, and SLAM.

## III. Outlier Rejection is Inapproximable

We show that MTS is *inapproximable* even by quasi-polynomial-time algorithms. We start with some definitions and present our key result in Theorem 4.

*Definition 2 (Approximability):* Consider the MTS Problem 1. Let $\mathcal{O}^\star$ be an optimal solution, let $k^\star \doteq |\mathcal{O}^\star|$, and $\epsilon \doteq \epsilon_{\mathcal{M}\setminus\mathcal{O}^\star}$, that is, $\epsilon$ is the outlier-free bound when the measurements $\mathcal{O}^\star$ are the rejected outliers. Also, consider a number $\lambda > 1$. We say that an algorithm makes MTS $(\lambda, \epsilon)$-*approximable* if it returns a set $\mathcal{O}$, and a parameter $x$, such that: the cardinality $|\mathcal{O}|$ is at most $\lambda_1 k^\star$; and the residual error $\sum_{i\in\mathcal{M}\setminus\mathcal{O}} \|h_i(y_i, x)\|^2$ is at most $\epsilon$. ⌟

Definition 2 allows some slack in the quality of the MTS's solution: it allows the desired error to be achieved with more measurement rejections than necessary.

*Definition 3 (Quasi-polynomial algorithm):* An algorithm is said to be *quasi-polynomial* if it runs in $2^{O[(\log m)^c]}$ time, where $m$ is the size of the input and $c$ is constant. ⌟

Any polynomial algorithm is also quasi-polynomial, since $m^k = 2^{k \log m}$. Yet, a quasi-polynomial algorithm is asymptotically faster than an exponential-time algorithm, since exponential algorithms run in $O(2^{m^c})$ time, for some $c > 0$.

*Theorem 4 (Inapproximability):* Consider the linear and convex MTS problem (3) below:

$$\min_{x\in\mathbb{R}^n} \min_{\mathcal{O}\subseteq\mathcal{M}} |\mathcal{O}|, \text{ s.t. } \sum_{i\in\mathcal{M}\setminus\mathcal{O}} \|y_i - a_i^\top x\|^2 \leq \epsilon, \quad (3)$$

Let $x^\star$ be the optimal value of the variable to be estimated, $m$ be the number of measurements ($m \doteq |\mathcal{M}|$), $\mathcal{O}^\star$ be the optimal solution, and set $k^\star \doteq |\mathcal{O}^\star|$. Then, for any $\delta \in (0, 1)$, there exist a polynomial $\lambda_1(m)$ and a function $\lambda_2(m) = 2^{\Omega(\log^{1-\delta} m)}$ and instances of MTS (i.e., measurements $y_i$, vectors $a_i$, and outlier-free bound $\epsilon$) where $\epsilon = \lambda_2(m)$, such that unless $\mathrm{NP}{\in}\mathrm{BPTIME}(m^{\mathrm{poly} \log m})$,[1] there is no quasi-polynomial algorithm making MTS $(\lambda_1(m), \lambda_2(m))$-approximable. This holds true even if the algorithm knows $k^\star$, and that $x^\star$ exist. ⌟

Theorem 4 stresses the extreme hardness of MTS. Even if we inform the algorithms with the true number of outliers, it is impossible in the worst-case for even quasi-polynomial algorithms to find a good set of inliers. Surprisingly, this remains true even if we allow the algorithms to cheat by rejecting more measurements than $k^\star$ (i.e., $\lambda_1 k^\star$).

Thinking beyond the worst-case, it becomes important to derive *per-instance* bounds that, for a given MTS instance, can evaluate how far an algorithm is from the optimal MTS solution. In order words, since we cannot guarantee that any efficient algorithm will do well in the worst-case, we are happy with evaluating (a posteriori) if an algorithm computed a good solution for a given problem instance. For this reason, in the next section we develop the first per-instance sub-optimality bound for Problem 1.

## IV. Performance Guarantees

We present the first per-instance (i.e., a posteriori) sub-optimality bound for the MTS Problem 1. The bound is algorithm-agnostic (does not take assumption on the way $\mathcal{O}$ is generated), and is computable in $O(1)$ time. Also, we demonstrate its informativeness via simulations.

*Theorem 5 (A posteriori sub-optimality bound):* Consider the MTS problem (2) and let $\mathcal{O}^\star$ be an optimal solution to (2). Also, for any candidate solution $\mathcal{O}$, let:

- $r(\mathcal{O}) \doteq \min_{x\in\mathbb{X}} \sum_{i\in\mathcal{M}\setminus\mathcal{O}} \|h_i(y_i, x)\|^2$; i.e., $r(\mathcal{O})$ is the minimum residual error given the rejection $\mathcal{O}$;
- $r_k^\star \doteq \min_{\mathcal{O}\subseteq\mathcal{M}, |\mathcal{O}|\leq k} r(\mathcal{O})$; i.e., $r_k^\star$ is the optimal residual error when at most $k$ measurements are rejected;
- $r^\star \doteq r(\mathcal{O}^\star)$; i.e., $r^\star$ is the residual error for the optimal outlier rejection $\mathcal{O}^\star$.

Then, given a candidate solution $\mathcal{O}$, the following bound relates the residual error $r(\mathcal{O})$ of the candidate solution with the residual error of an optimal solution rejecting the same number of outliers:

$$\frac{r(\mathcal{O}) - r^\star_{|\mathcal{O}|}}{r(\emptyset) - r^\star_{|\mathcal{O}|}} \leq \chi_{\mathcal{O}}, \quad (4)$$

where

$$\chi_{\mathcal{O}} \doteq \frac{r(\mathcal{O})}{r(\emptyset) - r(\mathcal{O})}. \quad (5)$$

⌟

Eq. (4) quantifies the distance between the residual of the candidate solution and the residual of an optimal solution rejecting the same number of outliers $|\mathcal{O}|$. The smaller $\chi_{\mathcal{O}}$, the closer the candidate selection is to the optimal selection. For example, when $\chi_{\mathcal{O}} = 0$, then $r(\mathcal{O}) = r^\star_{|\mathcal{O}|}$, i.e., we conclude that the algorithm returned a globally optimal solution (restricted to the ones rejecting $|\mathcal{O}|$ measurements).

*Remark 6 (Quality of the bound):* In [11], we showcase the quality of the bound (5) in terms of tightness by considering instances of the linear estimation problem in eq. (3). In particular, our results demonstrate that the bound predicts well the actual sub-optimality ratio, and its quality improves for increasing number of outliers. ⌟

We remark that the bound (5) can be also used to quantify the performance of existing algorithms, including RANSAC. Next, we step forward to a novel general-purpose algorithm for outlier rejection that empirically returns accurate solutions (and for which our bound $\chi_{\mathcal{O}}$ is typically close to zero).

## V. A General-Purpose Algorithm: ADAPT

We introduce a novel algorithm for outlier rejection that we name *Adaptive Trimming (*ADAPT*)*. The algorithm starts by processing all measurements and at each iteration it trims measurements with residuals larger than a threshold. It is *adaptive* in that it dynamically decides the threshold at each iteration (hence relaxing the need for a threshold $\epsilon_{\mathcal{M}\setminus\mathcal{O}}$). Moreover, it is not greedy in that it can reject multiple measurements at each iteration while it keeps revisiting the quality of previously rejected outliers.[2]

*Assumption 7 (Global solver):* ADAPT assumes the availability of a black-box solver that can (even approximately) solve the outlier-free problem (1) to optimality.

Luckily, for all problems in the experimental Section VI (3D registration; two-view geometry; SLAM), there exist (outlier-free) global solvers, including [6], [7], [8].

**Description of** ADAPT**.** The preudo-code of ADAPT is given in Algorithm 1. Here, we use the additional notation:

- Let $x^\star(\mathcal{O}) \in \arg\min_{x\in\mathbb{X}} \sum_{i\in\mathcal{M}\setminus\mathcal{O}} \|h_i(y_i, x)\|^2$; i.e., $x^\star(\mathcal{O})$ is an estimator of $x$ given an outlier selection $\mathcal{O}$.

[1] The complexity hypothesis $\mathrm{NP}{\notin}\mathrm{BPTIME}(m^{\mathrm{poly} \log m})$ means there is no randomized algorithm which outputs solutions to problems in NP with probability $2/3$, after running for $O(m^{(\log m)^c})$ time, for a constant $c$ [26].

[2] In our tests we found that a greedy algorithm similar to [27] tends to converge to poor outlier rejection decisions and is typically slow for practical applications, since it has quadratic runtime in the number of measurements.

**Algorithm 1:** Adaptive Trimming (ADAPT)

**Input:**

- $v$: minimum nr. of measurements required by global solver;
- $\gamma$: discount factor for outlier threshold (default $\gamma = 0.99$);
- $\delta$: convergence threshold;
- $T$: nr. of iterations to decide convergence (default $T = 2$);
- $\bar{g}$: maximum nr. of extra rejections per iteration.

**Output:** outlier set $\mathcal{O}$.

1: $t \leftarrow 0$; $\mathcal{O}_t \leftarrow \emptyset$; $g \leftarrow \bar{g}$; $c \leftarrow 0$; $\tau \leftarrow \max_{i \in \mathcal{M}} r_i(\emptyset)$;
2: **while** true **do**
3: $\quad t \leftarrow t + 1$; $\mathcal{O}_t \leftarrow \mathcal{O}_{t-1}$;
4: $\quad$ **while** $\mathcal{O}_t = \mathcal{O}_{t-1}$ **do** {discount threshold & update}
5: $\quad\quad \mathcal{I} \leftarrow$ indices of $g$ largest $r_i(\mathcal{O}_{t-1})$ across $i \in \mathcal{M}$;
6: $\quad\quad \mathcal{O}_t \leftarrow \{i \in \mathcal{I}$ and $r_i(\mathcal{O}_{t-1}) \geq \tau\}$;
7: $\quad\quad$ **if** $\mathcal{O}_t = \mathcal{O}_{t-1}$ or $\mathcal{O}_t = \emptyset$ **then** {discount}
8: $\quad\quad\quad \tau \leftarrow \gamma \min\left\{\tau, \max_{i \in \mathcal{M} \setminus \mathcal{O}_t} r_i(\mathcal{O}_t)\right\}$;
9: $\quad g \leftarrow g + \bar{g}$;
10: $\quad$ **if** $|\mathcal{O}_t| = |\mathcal{M}| - v$ **then** {terminate}
11: $\quad\quad$ **return** $\mathcal{O}_t$.
12: $\quad$ **if** $|r(\mathcal{O}_t) - r(\mathcal{O}_{t-1})| \leq \delta$ **then** {check convergence}
13: $\quad\quad c \leftarrow c + 1$;
14: $\quad\quad$ **if** $c = T$ **then** {terminate}
15: $\quad\quad\quad$ **return** $\mathcal{O}_t$.
16: $\quad$ **else** {reset convergence counter}
17: $\quad\quad c \leftarrow 0$.

- Let $r_i(\mathcal{O}) \doteq \|h_i(y_i, x^\star(\mathcal{O}))\|^2$; i.e., $r_i(\mathcal{O})$ is the residual of the measurement $i$, given an outlier selection $\mathcal{O}$.

A description for each of ADAPT's steps is found in [11].

*Remark 8 (Complexity and practicality):* The termination condition in line 10 guarantees the termination of the algorithm with at most $|\mathcal{M}| - v$ calls of the global solver. ADAPT terminates faster as one increases the outlier group size $\bar{g}$, the convergence thresholds $\delta$, and/or as one decreases the discount factor $\gamma$ and the number $T$ of iterations to decide convergence. Overall, the linear runtime (in the number of measurements) of ADAPT makes the algorithm practical in real-time applications where fast global solvers are available.

*Remark 9 (vs. RANSAC):* While RANSAC builds an inlier set by sampling small (minimal) sets of measurements, ADAPT iteratively prunes the overall set of measurements. Arguably, this gives ADAPT a "global vision" of the measurement set as we showcase in the experimental section. RANSAC assumes the availability of fast minimal solvers, while ADAPT assumes the availability of fast global (non-minimal) solvers. Finally, RANSAC is not suitable for high-dimensional problems where is becomes exponentially more difficult to sample an outlier-free set [19]. On the other hand, ADAPT is deterministic and guaranteed to terminate in a number of iterations bounded by the number of measurements.

## VI. Experiments and Applications

We evaluate ADAPT against the state of the art in 3D registration, two-view geometry, and SLAM. ADAPT outperforms RANSAC in terms of accuracy and scalability, and often outperforms specialized outlier rejection methods (in particular for SLAM) while being a general-purpose algorithm. Finally, the tests show that the performance bounds of Section IV are informative and can be used to assess the outlier rejection outcomes. All results are averaged over 10 Monte Carlo runs. Details are outlined in [11].

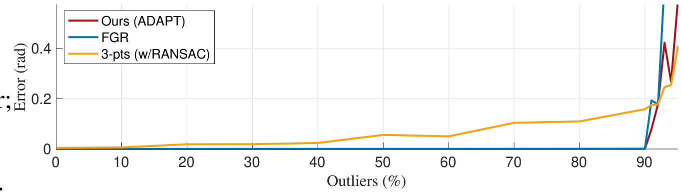


(a) Rotation Error

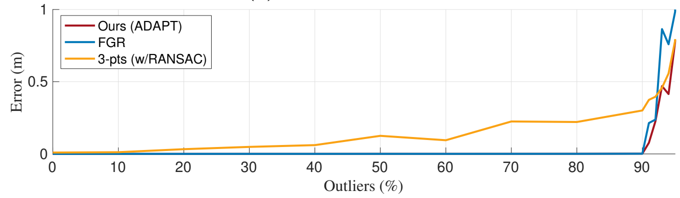


(b) Translation Error

Fig. 2. 3D registration: rotation and translation errors for ADAPT, FGR [28], and RANSAC on the Bunny dataset for increasing outlier percentages.

### A. Robust Registration

**Experimental setup.** We test ADAPT on two standard datasets for 3D registration: the Stanford Bunny and the ETH Hauptgebaude [29]. In each iteration, ADAPT uses Horn's method [8] as global solver. We benchmark ADAPT against *Fast Global Registration* (FGR) [28] and the three-point RANSAC. We set the maximum number of iterations in RANSAC to 1000 and use default parameters for FGR.

**Results.** Fig. 2 shows the (average) translation and rotation errors for the estimates computed by ADAPT, FGR, and RANSAC on the Bunny dataset for increasing outlier percentages. ADAPT performs on-pair with FGR which is a specialized robust solver for 3D registration and they both achieve practically zero error for up to 90% of outliers, after which they both break. RANSAC starts performing distinctively worse early on. We obtain similar results on the ETH dataset hence for space reasons we report them in [11].

For both the Bunny and ETH datasets, we compute the sub-optimality bound for the result of ADAPT, using Theorem 5. The plot of the bound is given in [11]; the bound remains around $10^{-5}$, confirming that ADAPT remains close to the optimal outlier selection. The runtime of ADAPT is comparable to FGR and is reported in [11].

### B. Robust Two-view Geometry

**Experimental setup.** We tested ADAPT on synthetic data and the MH_01 sequence of the *EuRoC* dataset [30]. In each iteration, ADAPT uses QCQP relaxation [7] as global solver. We benchmarked ADAPT against Nister's five-point [31] and the eight-point algorithm [32] within RANSAC.

**Results.** Per Fig. 3, ADAPT and five-point perform on-pair till 40% of outliers. Beyond that point, the five-point method attains considerably higher errors than ADAPT (50% to 100% more in rotation; and more than 300% more in translation).

For the synthetic dataset, the typical value for the sub-optimality bound achieved by ADAPT is 0.2: ADAPT makes

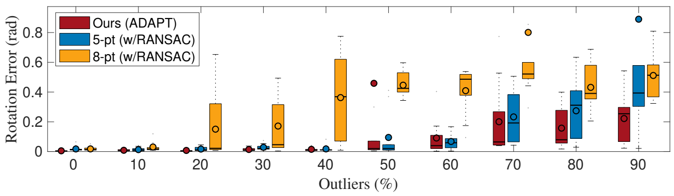


(a) Rotation Error

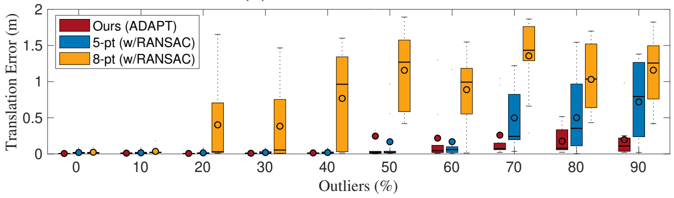


(b) Translation Error

Fig. 3. Two-view geometry: rotation and translation errors for ADAPT, five- and eight-point RANSAC, on a synthetic dataset for increasing outliers.

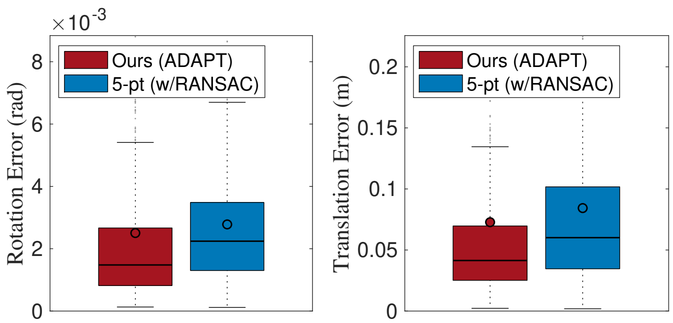


Fig. 4. Two-view geometry: rotation and translation errors for ADAPT, and five-point RANSAC, on the sequence MH_01 of the EuRoC dataset.

a rejection that achieves an error that is at most 20% away from the optimal, even in the presence of 90% of outliers.

The runtime of ADAPT is reported in [11]: our approach is one order of magnitude slower than the five-point method, mainly due to the relatively high runtime of the global solver [7], which is called in each iteration.

### C. Robust SLAM

**Experimental setup.** We test ADAPT on standard 2D and 3D SLAM benchmarking datasets and report missing results here in [11]. We use *SE-Sync* as the global solver for ADAPT. We test the following datasets, described in [33], [6]: *MIT* (2D), *Intel* (2D), *CSAIL* (2D), and *Sphere2500* (3D). We also test a simulated $5 \times 5 \times 5$ 3D grid dataset (results in the supplemental [34]). We benchmark ADAPT against DCS [35]; we report DCS results for three choices of the robust kernel size: $\{1,10,100\}$ (the default value is 1, see [35]).

**Results.** ADAPT outperforms DCS (for any chosen kernel size) across all datasets and outlier percentages.

For example, in *2D SLAM* (similar observation holds for the *3D SLAM case*; see [11]): in the *MIT* dataset (Fig. 5a), a particularly challenging dataset, ADAPT is insensitive to up to 20% of outliers. All variants of DCS fail to produce an error smaller than 10 meters even in the absence of outliers. ADAPT leverages SE-sync, which is a global solver, hence is able to converge to the correct solution. And in the *CSAIL* dataset (Fig. 5b), ADAPT also dominates DCS while DCS performance is acceptable when the kernel size is equal to 100.

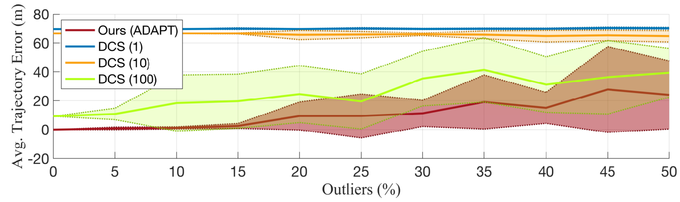


(a) MIT

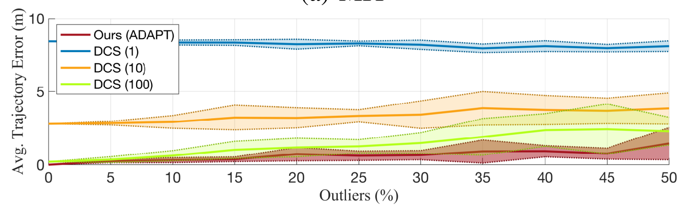


(b) CSAIL

Fig. 5. 2D SLAM: Average Trajectory Error of ADAPT and DCS for increasing outliers in the MIT and CSAIL datasets.

The typical value for the sub-optimality bound achieved by ADAPT, per Theorem 5, is 0.01 to 0.01.

ADAPT is one to two orders slower than DCS. This is due to the repeated calls to SE-sync and is further aggravated in the 3D case by the fact that SE-sync tends to be slow in the presence of outliers (the Riemannian staircase method [6] requires multiple staircase iterations since the rank of the relaxation increases in the presence of outliers).

## VII. Conclusion

We proposed a *minimally trimmed squares (MTS)* formulation for outlier-robust estimation. We proved that the resulting outlier rejection problem is inapproximable. We derived the first a posteriori theoretical performance bounds. Finally, we proposed a linear-time, general-purpose algorithm for outlier rejection, and showed that it outperforms several specialized methods across three spatial perception problems (3D registration, two-view geometry, SLAM).

## References


[1] D. Scaramuzza and F. Fraundorfer, "Visual odometry: Part I the first 30 years and fundamentals," 2011.

[2] 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.*, vol. 19, no. 7, pp. 1199–1217, 2013.

[3] S. Roumeliotis and G. Bekey, "Distributed multi-robot localization," *IEEE Trans. Robot. Automat.*, vol. 18, no. 5, pp. 781–795, October 2002.

[4] T. Whelan, S. Leutenegger, R. Salas-Moreno, B. Glocker, and A. Davison, "ElasticFusion: Dense SLAM without a pose graph," in *Robotics: Science and Systems (RSS)*, 2015.

[5] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.

[6] D. Rosen, L. Carlone, A. Bandeira, and J. Leonard, "SE-Sync: a certifiably correct algorithm for synchronization over the Special Euclidean group," *Intl. J. of Robotics Research*, 2018.

[7] J. Briales, L. Kneip, and J. Gonzalez-Jimenez, "A certifiably globally optimal solution to the non-minimal relative pose problem," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[8] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Amer.*, vol. 4, no. 4, pp. 629–642, Apr 1987.

[9] M. Bosse, G. Agamennoni, and I. Gilitschenski, "Robust estimation and applications in robotics," *Foundations and Trends in Robotics*, vol. 4, no. 4, pp. 225–269, 2016.

[10] 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)*, July 2017, pp. 5048–5056.

[11] V. Tzoumas, P. Antonante, and L. Carlone, "Outlier-robust spatial perception: Hardness, general-purpose algorithms, and guarantees," *arXiv preprint:1903.11683*, 2019.

[12] 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.*, vol. 21, no. 11, pp. 1229–1234, 1999.

[13] Y. Latif, C. D. C. Lerma, and J. Neira, "Robust loop closing over time." in *Robotics: Science and Systems (RSS)*, 2012.

[14] J. G. Mangelson, D. Dominic, R. M. Eustice, and R. Vasudevan, "Pairwise consistent measurement set maximization for robust multi-robot map merging," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2018.

[15] 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)*, 2014, pp. 803–817.

[16] R. I. Hartley and F. Kahl, "Global optimization through rotation space search," *Intl. J. of Computer Vision*, vol. 82, no. 1, pp. 64–79, 2009.

[17] Y. Zheng, S. Sugimoto, and M. Okutomi, "Deterministically maximizing feasible subsystem for robust model fitting with unit norm constraint," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 1825–1832.

[18] H. Li, "Consensus set maximization with guaranteed global optimality for robust geometry estimation," in *Intl. Conf. on Computer Vision (ICCV)*, 2009, pp. 1074–1080.

[19] Á. P. Bustos and T. J. Chin, "Guaranteed outlier removal for point cloud registration with correspondences," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 40, no. 12, pp. 2868–2882, 2018.

[20] T.-J. Chin, Z. Cai, and F. Neumann, "Robust fitting in computer vision: Easy or hard?" in *European Conf. on Computer Vision (ECCV)*, 2018.

[21] P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Detection*. John Wiley & Sons, New York, NY, 1987.

[22] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[23] F. Dellaert and M. Kaess, "Factor graphs for robot perception," *Foundations and Trends in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.

[24] P. J. Rousseeuw and M. Hubert, "Robust statistics for outlier detection," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 1, no. 1, pp. 73–79, 2011.

[25] 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.*, vol. 38, no. 11, pp. 2241–2254, Nov. 2016.

[26] S. Arora and B. Barak, *Computational complexity: A modern approach*. Cambridge University Press, 2009.

[27] G. Nemhauser, L. Wolsey, and M. Fisher, "An analysis of approximations for maximizing submodular set functions – I," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.

[28] Q. Y. Zhou, J. Park, and V. Koltun, "Fast global registration," in *European Conf. on Computer Vision (ECCV)*. Springer, 2016, pp. 766–782.

[29] F. Pomerleau, M. Liu, F. Colas, and R. Siegwart, "Challenging data sets for point cloud registration algorithms," *The International Journal of Robotics Research*, vol. 31, no. 14, pp. 1705–1711, Dec. 2012.

[30] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The euroc micro aerial vehicle datasets," *The International Journal of Robotics Research*, 2016. [Online]. Available: http://ijr.sagepub.com/content/early/2016/01/21/0278364915620033.abstract

[31] D. Nistér, "An efficient solution to the five-point relative pose problem," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 26, no. 6, pp. 756–770, 2004.

[32] R. I. Hartley, "In defense of the eight-point algorithm," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 6, pp. 580–593, June 1997.

[33] L. Carlone and A. Censi, "From angular manifolds to the integer lattice: Guaranteed orientation estimation with application to pose graph optimization," *IEEE Trans. Robotics*, vol. 30, no. 2, pp. 475–492, 2014.

[34] V. Tzoumas, P. Antonante, and L. Carlone, "Outlier-robust spatial perception: Hardness, general-purpose algorithms, and guarantees," Tech. Rep., 2019, https://www.dropbox.com/s/wtrk2a3zynurc5u/2019c-IROS-outliers-supplemental.pdf?dl=0.

[35] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard, "Robust map optimization using dynamic covariance scaling," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013.