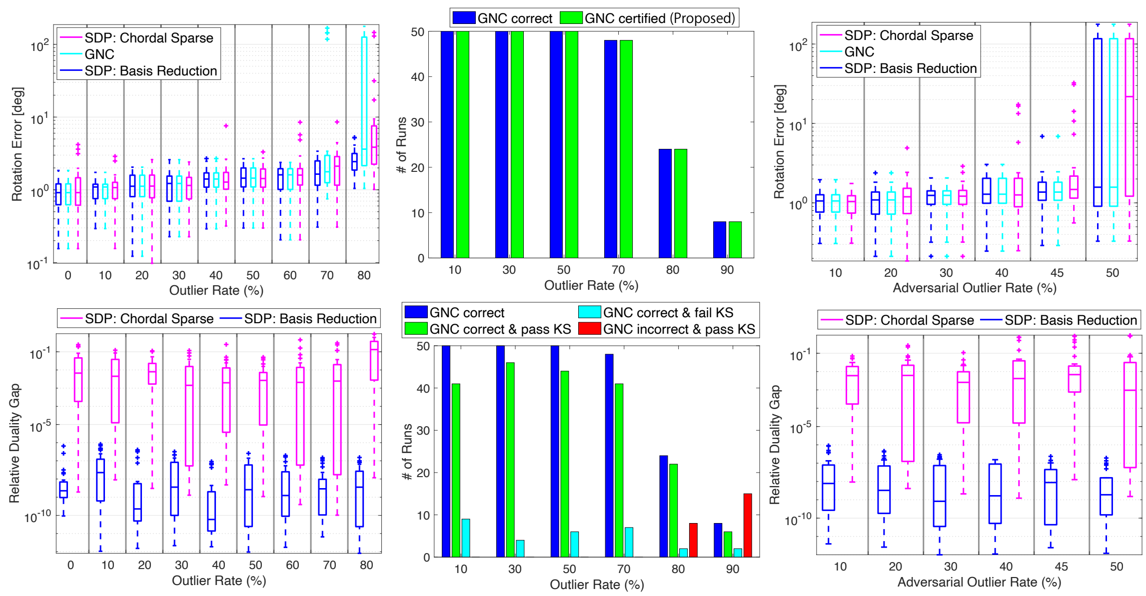

(a) Comparison to baselines for solving primal TLS problem. Top: rotation error, bottom: duality gap.

(b) Comparison to baselines for verifying solution correctness. Top: DRS (ours), bottom: KS test.

(c) Performance w.r.t. increasing adversarial outliers. Top: rotation error, bottom: duality gap.

Novelty and Significance (R1, R4). (i) **Motivation:** *Outliers* are ubiquitous in computer vision and robotics [101]. (ii) **Hardness:** *Globally optimal* outlier-robust geometric perception is NP-hard and intractable [28]. (iii) **Limitation of SOTA:** SOTA algorithms are divided into *fast heuristics* (real-time but no guarantees) and *global solvers* (optimal but exponential time). (iv) **Significance of this paper:** (1) ***Theoretical:*** The *first polynomial-time* method for solving generic outlier-robust geometric perception with *a posteriori* global optimality guarantees. The tightness of the relaxation discovers a class of non-convex problems that admits *hidden convexity*, which has significant potential for theoretical study.[1] (2) ***Algorithmic:*** The *first* tractable method for designing *dual optimality certifiers* that run orders of magnitude faster than SOTA SDP solvers (*cf.* Table 1).[2] (3) ***Broader Impact:*** Optimality guarantees (*e.g.,* rejecting wrong solutions as in Fig. 1(b)) are crucial for safety-critical applications. (v) **Novelties:** (1) First to reformulate TLS as POP with structured sparsity (*cf.* Prop. 5(i)(ii)); (2) First to empirically show Lasserre's hierarchy is (surprisingly) tight for *outlier-robust* problems (Kahl'07 IJCV, Yang'20 CVPR assume outlier-free) with *binary variables* (vs. MAX-CUT relaxation is not tight); (3) First to use *basis reduction* to improve efficiency but keep tightness (vs. chordal relaxation [91] is not tight (Fig. (a))). (4) First to propose scalable certifiers using DRS and chordal initialization.

Comparison with Baselines (R1-4). (i) Fig. (a) compares the performance of our primal solver (SDP: Basis Reduction) versus two (SOTA) baselines, GNC (best heuristics) [97] and SDP: Chordal Sparse (efficient SDP relaxation) [91]. Our primal relaxation is significantly tighter than chordal sparse relaxation [91], and the accuracy and robustness of our estimates dominates both baselines. (ii) Our DRS approach is the *first* mathematically rigorous approach for verifying solution correctness. We compare it with a heuristic method that performs Kolmogorov–Smirnov (KS) test on the squared residuals with a χ^2 distribution (*i.e.,* tests normality of the residuals classified as inliers). Fig. (b) shows that KS test has many false positives/negatives, while ours has zero. *All results are for Single Rotation Averaging, similar results hold for the other applications considered in the paper.*

Adversarial Outliers (R2). We performed tests with an adversarial outlier model (where outliers follow a different model and are consistent with each other) and test our algorithm (SDP: Basis Reduction) against two SOTA baselines. Fig. (c) shows our method dominates both baselines, is insensitive to adversarial outliers until the maximum breakdown point 50% (the tightness of the relaxation implies the optimal solution may not be the ground-truth solution at 50%).

LTS (R2). LTS minimizes sum of the K smallest squared residuals: $\min_{\boldsymbol{x}\in\mathcal{X}} \sum_{i=1}^{K} r_i^2(\boldsymbol{x})$, with $r_1^2 \leq \cdots \leq r_N^2$. This is equivalent to: $\min_{\boldsymbol{x}\in\mathcal{X},\theta_i^2=1} \sum_{i=1}^{N} \frac{1+\theta_i}{2} r_i^2(\boldsymbol{x})$, subject to: $\sum_{i=1}^{N} \theta_i = 2K - N$, which ensures the number of θ_i's with value $+1$ is exactly K (smallest K). Therefore, LTS can be written as a POP, and our framework can be applied. However, tightness of performing relaxation for LTS is not guaranteed and basis reduction may need extra care.

Others. (i) Theoretical breakdown of TLS is 50%. Empirical robustness can be higher if outliers are not adversarial ([99], over 95%). [101] establishes breakdown for a specific problem. (ii) $N = 100$ is common for real problems. But surely there is still room for scalability improvements (*e.g.,* BM factorization [18]). (iii) Main paper is rigorous thus hard to follow. Supplementary provides details for non-expert readers, and we will open source our implementation.

[1] Moreover, the tightness of the relaxations further motivates developing fast SDP solvers, which is a major line of research on its own (many related work in NeurIPS). As SDP solvers become more efficient, these problems eventually can be solved in real-time.

[2] The goal of this paper is NOT to *replace* existing fast heuristics, but to *enhance* them with a fast *certification* that allows asserting the quality of their estimates and rejecting failure cases (*cf.* Fig. 1(b)), for safety-critical applications.