Outlier-Robust Spatial Perception:
Hardness, Algorithms, Guarantees

Vasileios Tzoumas
Research Scientist
Laboratory of Decision and Information Systems
Massachusetts Institute of Technology

Recap…

Perception as least squares optimization

When Gaussian measurement noise, maximum likelihood estimation (MLE) gives:

$$\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{N} r^2(\boldsymbol{y}_i, \boldsymbol{x})$$

Estimate ← $\min_{\boldsymbol{x} \in \mathcal{X}}$

Measurements/data

Residual

Examples:

Point cloud registration:

$$\min_{\substack{R \in SO(3) \\ t \in \mathbb{R}^3}} \sum_{(i,j) \in \mathcal{M}} \| R p_i + t - p'_j \|^2$$

Rotation + translation

Point clouds (data)

Correspondences between $p_i, {p'}_j$

SLAM:

$$\min_{\substack{T_i \in \mathrm{SE}(3) \\ i=1,\ldots,n}} \sum_{(i,j) \in \mathcal{M}} \| T_j - T_i \bar{T}_{ij} \|_F^2$$

Pose ← $\min_{T_i \in \mathrm{SE}(3)}$

Relative pose measurement

Loop closures between T_i, T_j

Outliers compromise least squares solutions

But if some y_i are outliers, solution of $\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{N} r^2(\boldsymbol{y}_i, \boldsymbol{x})$ can be wrong:

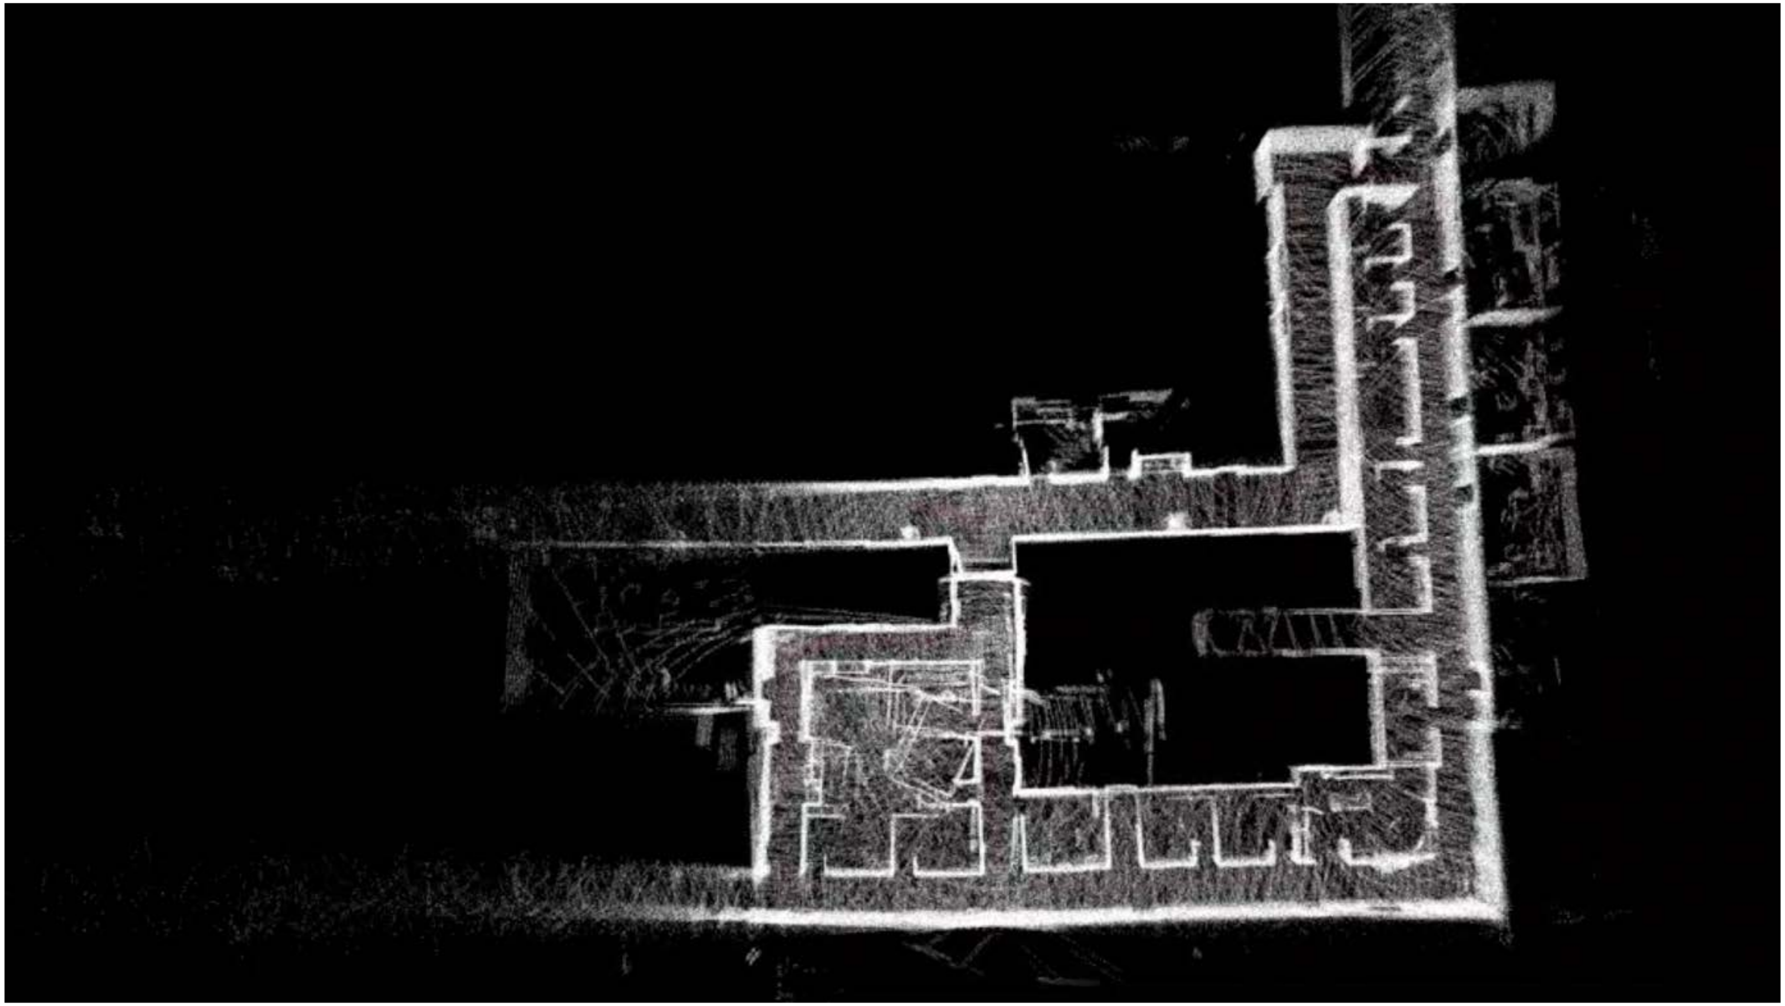

Outlier-robust least squares reformulations

L : Robust-cost "least squares"

$$\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho(r(x, y_i), \bar{c})$$

R : Outlier rejection "least squares"

$$\min_{x \in \mathcal{X}, \mathcal{O} \subseteq \mathcal{M}} |\mathcal{O}| \quad s.t. \quad \left\| r\left(x, y_{\mathcal{M} \setminus \mathcal{O}}\right) \right\|^2 \leq \epsilon$$

Possible choices for ρ:

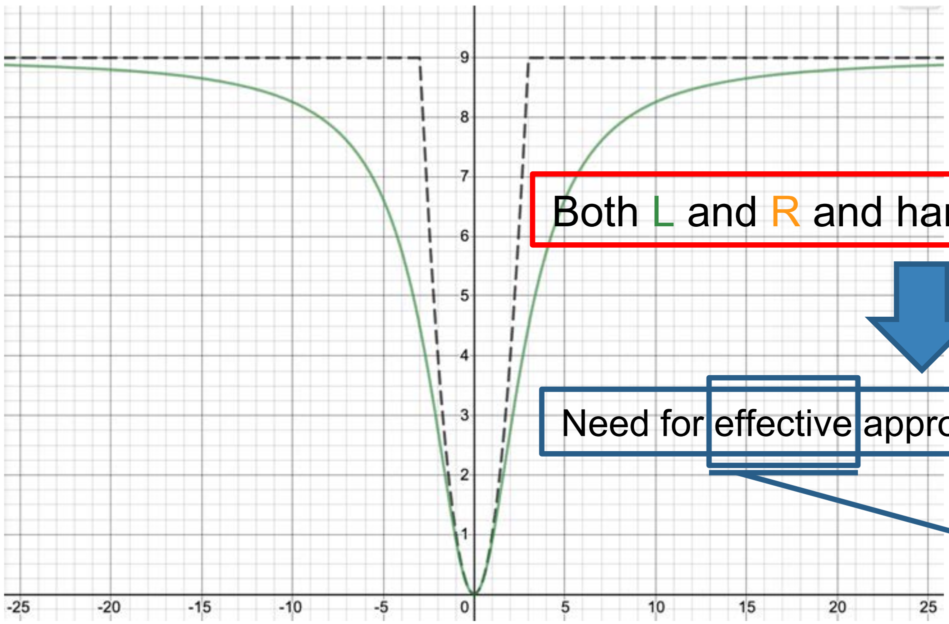

Both L and R and harder than NP-hard

Need for effective approximation algorithms

- Fast
- Finds correct x despite many outliers

Today's focus

Methods to solve $\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho\big(r(x, y_i)\big)$

Why solving $\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho\big(r(x, y_i)\big)$ can be hard?

Recall:

Possible choices for ρ:

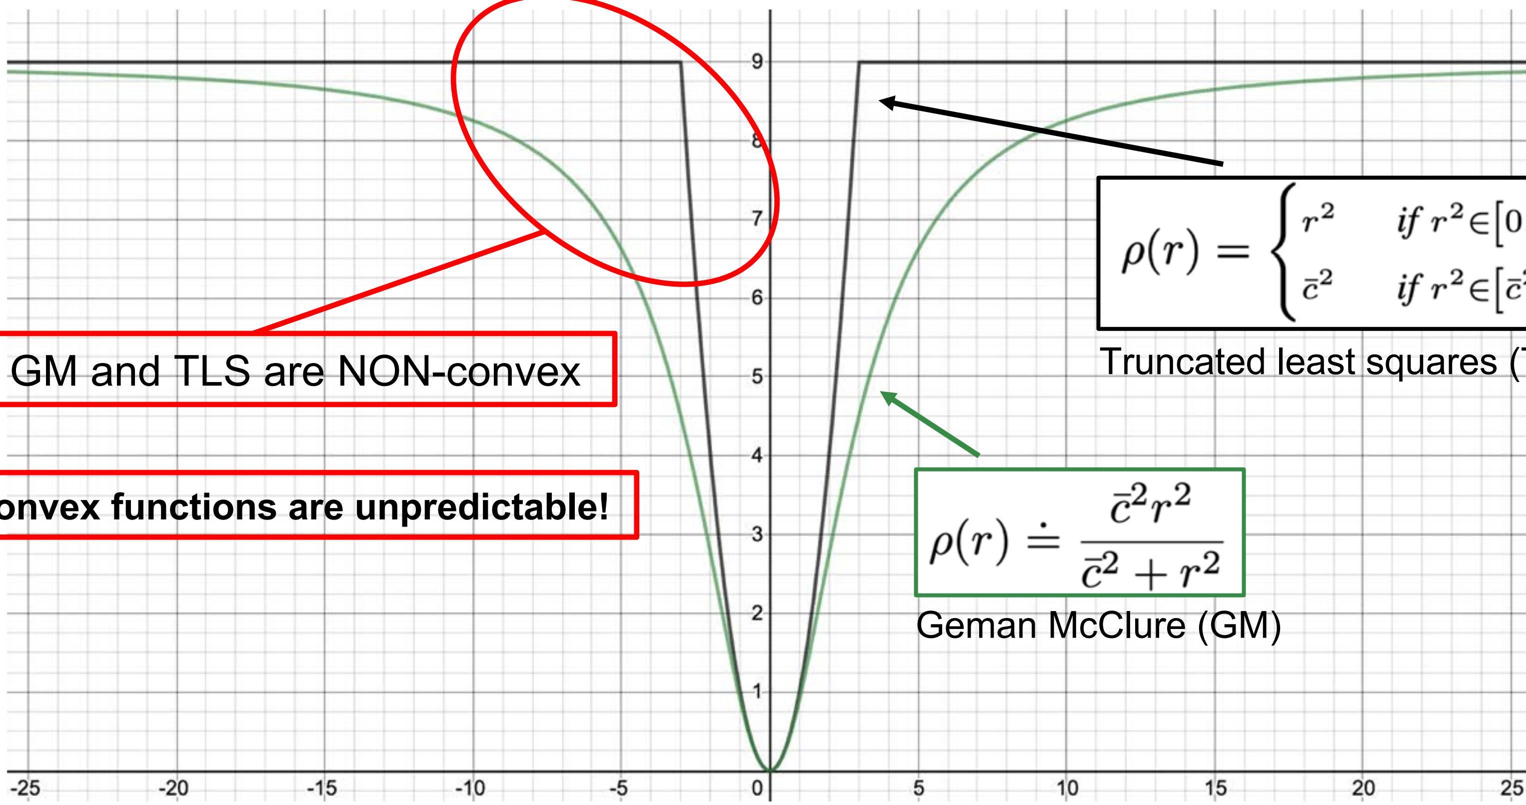

Both GM and TLS are NON-convex

Non-convex functions are unpredictable!

$$\rho(r) = \begin{cases} r^2 & \text{if } r^2 \in [0, \bar{c}^2] \\ \bar{c}^2 & \text{if } r^2 \in [\bar{c}^2, +\infty) \end{cases}$$

Truncated least squares (TLS)

$$\rho(r) \doteq \frac{\bar{c}^2 r^2}{\bar{c}^2 + r^2}$$

Geman McClure (GM)

Why solving $\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho(r(x, y_i))$ can be hard?

Example (revisited):

- $x_{true} = 0$
- Measurements' model: $y_1 = x + gaussian\ noise\ of\ \mu = 0, \sigma = 1$
  $y_2 = x + gaussian\ noise\ of\ \mu = 0, \sigma = 1$
  $y_3 = 2x + gaussian\ noise\ of\ \mu = 0, \sigma = 1$
- Observed measurements: $y_1 = y_2 = 0$, $y_3 = 10$
- $\rho(r) \doteq \frac{\bar{c}^2 r^2}{\bar{c}^2 + r^2}$

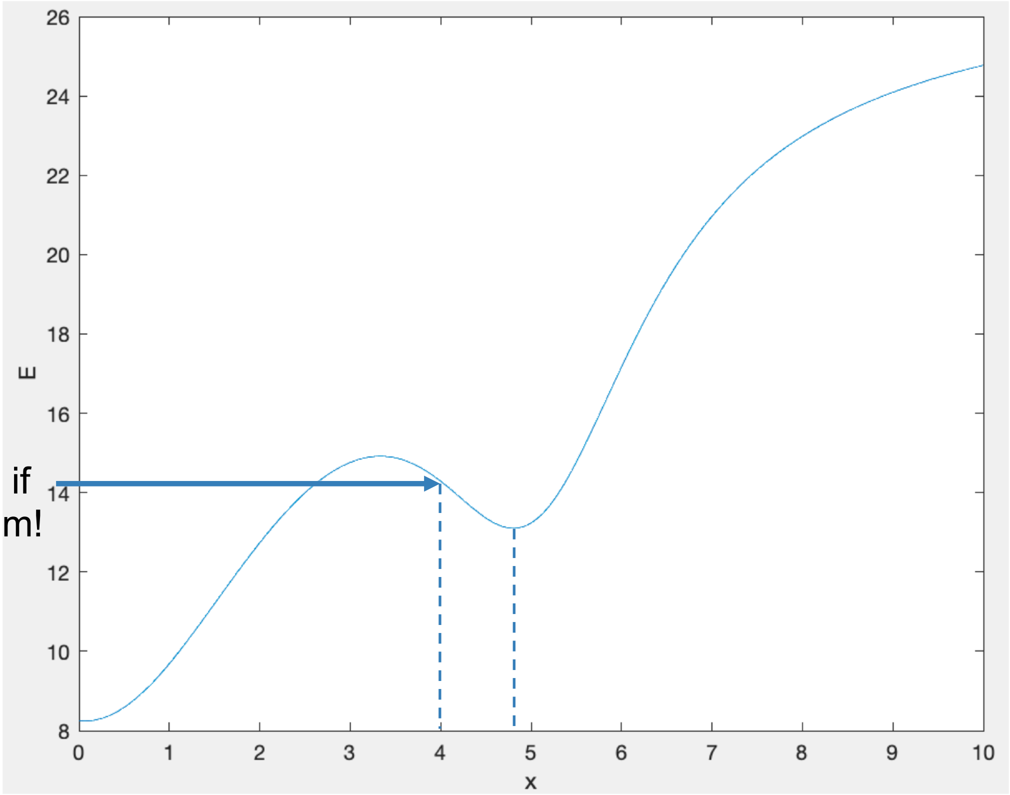

If we stand here (x = 4), we cannot be sure if visible minimum at x = 4.8 is global minimum!

Solving (optimally) non-convex problems is hard

No methods exist that guarantee convergence in general:

- Typically rely on good **initial guess** on x to converge to optimality (e.g., Gauss-Newton (GN))

- Alternate direction of optimization:

  (i) Decompose x into k subvectors: $x = (x_1, x_2, \ldots, x_k)$
  (ii) Consecutively optimize x_i given current $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots x_k$

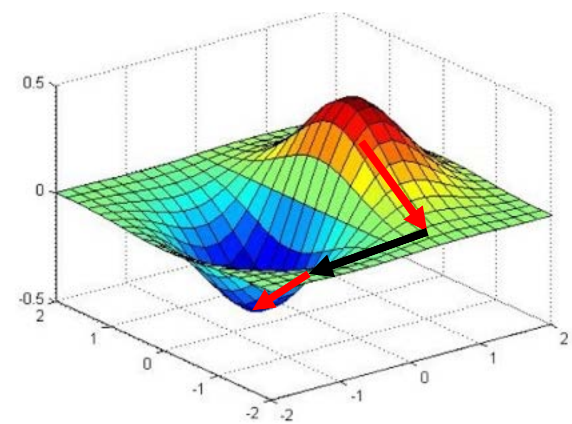

Overcoming ρ's non-convexity

Intermediate point:

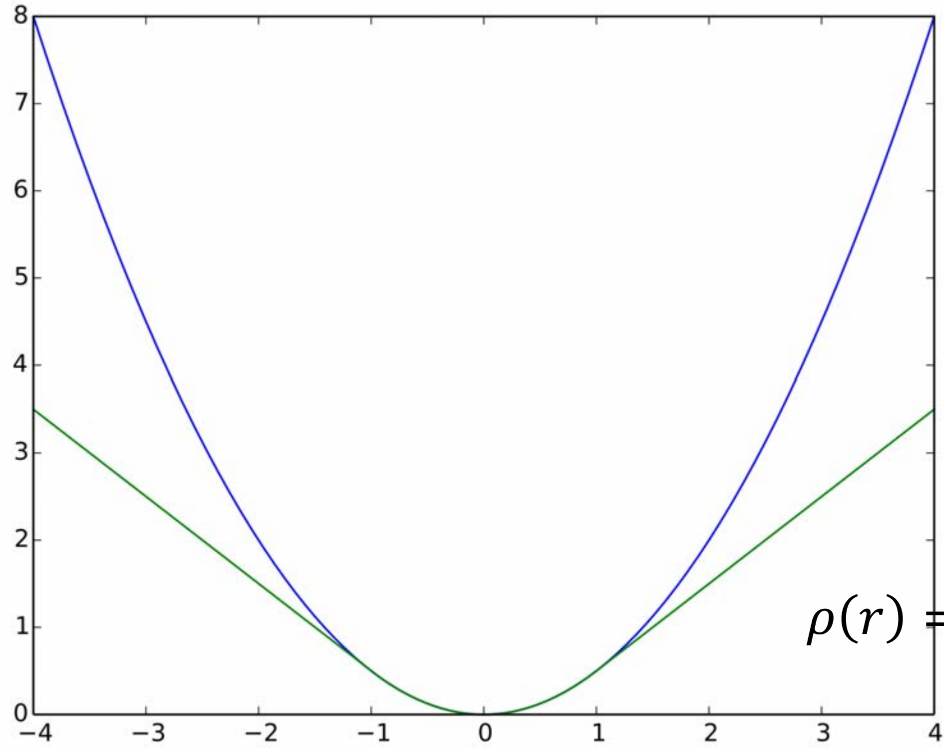

Huber loss:
Still convex!

$$\rho(r) = \begin{cases} \frac{1}{2} r^2, |r| \leq \bar{c} \\ \bar{c}\left(|r| - \frac{1}{2}\bar{c}\right), o.w. \end{cases}$$

Overcoming ρ's non-convexity

$$\rho_\mu(r) = \frac{\mu \bar{c}^2 r^2}{\mu \bar{c}^2 + r^2}$$

- For $\mu = 1$ becomes **GM**
- For $\mu = +\infty$ becomes r^2

Idea: Start from least squares and **gradually** go to ρ-"least squares":

1. Start with least squares, solve it, take initial estimate for x; call it x_0;

2. Use x_0 as initial guess to solve via GN a slightly non-least squares problem (by starting, e.g., with $\mu = 100$), solve it, take x_1;

3. Set $x_0 = x_1$, and go to Step 2 (using now $\mu = \mu/2$), and so forth, until $\mu = 1$.

Similarly for TLS:

For similar approaches (on GN and graduated non-convexity), see paper *Iterated Lifting for Robust Cost Optimization* by C. Zach '14

$$\rho_\mu(r) = \begin{cases} r^2 & \text{if } r^2 \in \left[0, \frac{\mu}{\mu+1}\bar{c}^2\right] \\ 2\bar{c}|r|\sqrt{\mu(\mu+1)} - \mu(\bar{c}^2 + r^2) & \text{if } r^2 \in \left[\frac{\mu}{\mu+1}\bar{c}^2, \frac{\mu+1}{\mu}\bar{c}^2\right] \\ \bar{c}^2 & \text{if } r^2 \in \left[\frac{\mu+1}{\mu}\bar{c}^2, +\infty\right) \end{cases}$$

- For $\mu = +\infty$ becomes x^2
- For $\mu = 0$ becomes **TLS**

GN based approaches are fast but not necessarily optimal (can converge to local minima)!
Maybe we can do better at Step 2?

Optimal solvers and graduated non-convexity

Idea: Utilize the known optimal solvers for the least squares problem

$$\min_{\boldsymbol{x} \in \mathcal{X}} \sum_{i=1}^{N} r^2(\boldsymbol{y}_i, \boldsymbol{x})$$

to solve the outlier-robust cost “least squares” problem of Step 2 in previous slide:

$$\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho_\mu\big(r(y_i, x)\big)$$

Optimal solvers and graduated non-convexity

Idea: "Transform" $\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho_\mu(r(y_i, x))$ to a least squares problem:

Find appropriate set $\mathcal{W}$ and function Φ_{ρ_μ} such that:

$$\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho_\mu(r(y_i, x)) = \min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \min_{w_i \in \mathcal{W}} [\, w_i r^2(y_i, x) + \Phi_{\rho_\mu}(w_i) \,] \triangleq E$$

scalar

Finding Φ_{ρ_μ} in the general case

- For **GM**: $\Phi_{\rho_\mu}(w_i) = \mu\bar{c}^2\left(\sqrt{w_i} - 1\right)^2$, where $w_i = \left(\frac{\mu\bar{c}^2}{r_i^2 + \mu\bar{c}^2}\right)^2$
- For **TLS**: $\Phi_{\rho_\mu}(w_i) = \dfrac{\mu(1-w_i)}{\mu + w_i}\bar{c}^2, \quad w_i = \begin{cases} 0 & \text{if } r_i^2 \in \left[\frac{\mu+1}{\mu}\bar{c}^2, +\infty\right] \\ \frac{\bar{c}}{r_i}\sqrt{\mu(\mu+1)} - \mu & \text{if } r_i^2 \in \left[\frac{\mu}{\mu+1}\bar{c}^2, \frac{\mu+1}{\mu}\bar{c}^2\right] \\ 1 & \text{if } r_i^2 \in \left[0, \frac{\mu}{\mu+1}\bar{c}^2\right]. \end{cases}$

In the general case, a method to find Φ_{ρ_μ} described in

On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision
by Michael J. Black and Anand Rangarajan IJCV '96

given that the following conditions hold, where $\phi(z) \triangleq \rho(\sqrt{z})$:

- $\lim_{z\to 0} \phi'(z) = 1$
- $\lim_{z\to\infty} \phi'(z) = 0$
- $\phi''(z) < 0$

Optimal solvers and graduated non-convexity

The steps in previous slide become (substituting GN step 2 with steps 3-5 below):[1]

1. Initialize $\mu \gg (e.g., 100)$ and $t = 0$.
2. Start by solving the least squares $\min_{\boldsymbol{x}\in\mathcal{X}} \sum_{i=1}^{N} r^2(\boldsymbol{y}_i, \boldsymbol{x})$ and let $x^{(t)}$ be the solution.
3. **Weight update**: Update $w^{(t)}$, given the fixed $x^{(t)}$:

$$\boldsymbol{w}^{(t)} = \underset{w_i\in[0,1]}{\arg\min} \sum_{i=1}^{N} \left[w_i r^2(\boldsymbol{y}_i, \boldsymbol{x}^{(t)}) + \Phi_{\rho_\mu}(w_i) \right]$$

4. $t = t + 1$.
5. **Variable update**: Update $x^{(t)}$, given the $w^{(t-1)}$ found at Step 3:

$$\boldsymbol{x}^{(t)} = \underset{\boldsymbol{x}\in\mathcal{X}}{\arg\min} \sum_{i=1}^{N} w_i^{(t-1)} r^2(\boldsymbol{y}_i, \boldsymbol{x})$$

6. $\mu = \mu/2$, and go to Step 3 until $\mu = 1$.

[1] Yang, Antonante, Tzoumas, Carlone, *Graduated Non-Convexity for Robust Spatial Perception: From Non-Minimal Solvers to Global Outlier Rejection,* IEEE Robotics and Automation Letters (RA-L), 2020, and IEEE ICRA 2020.

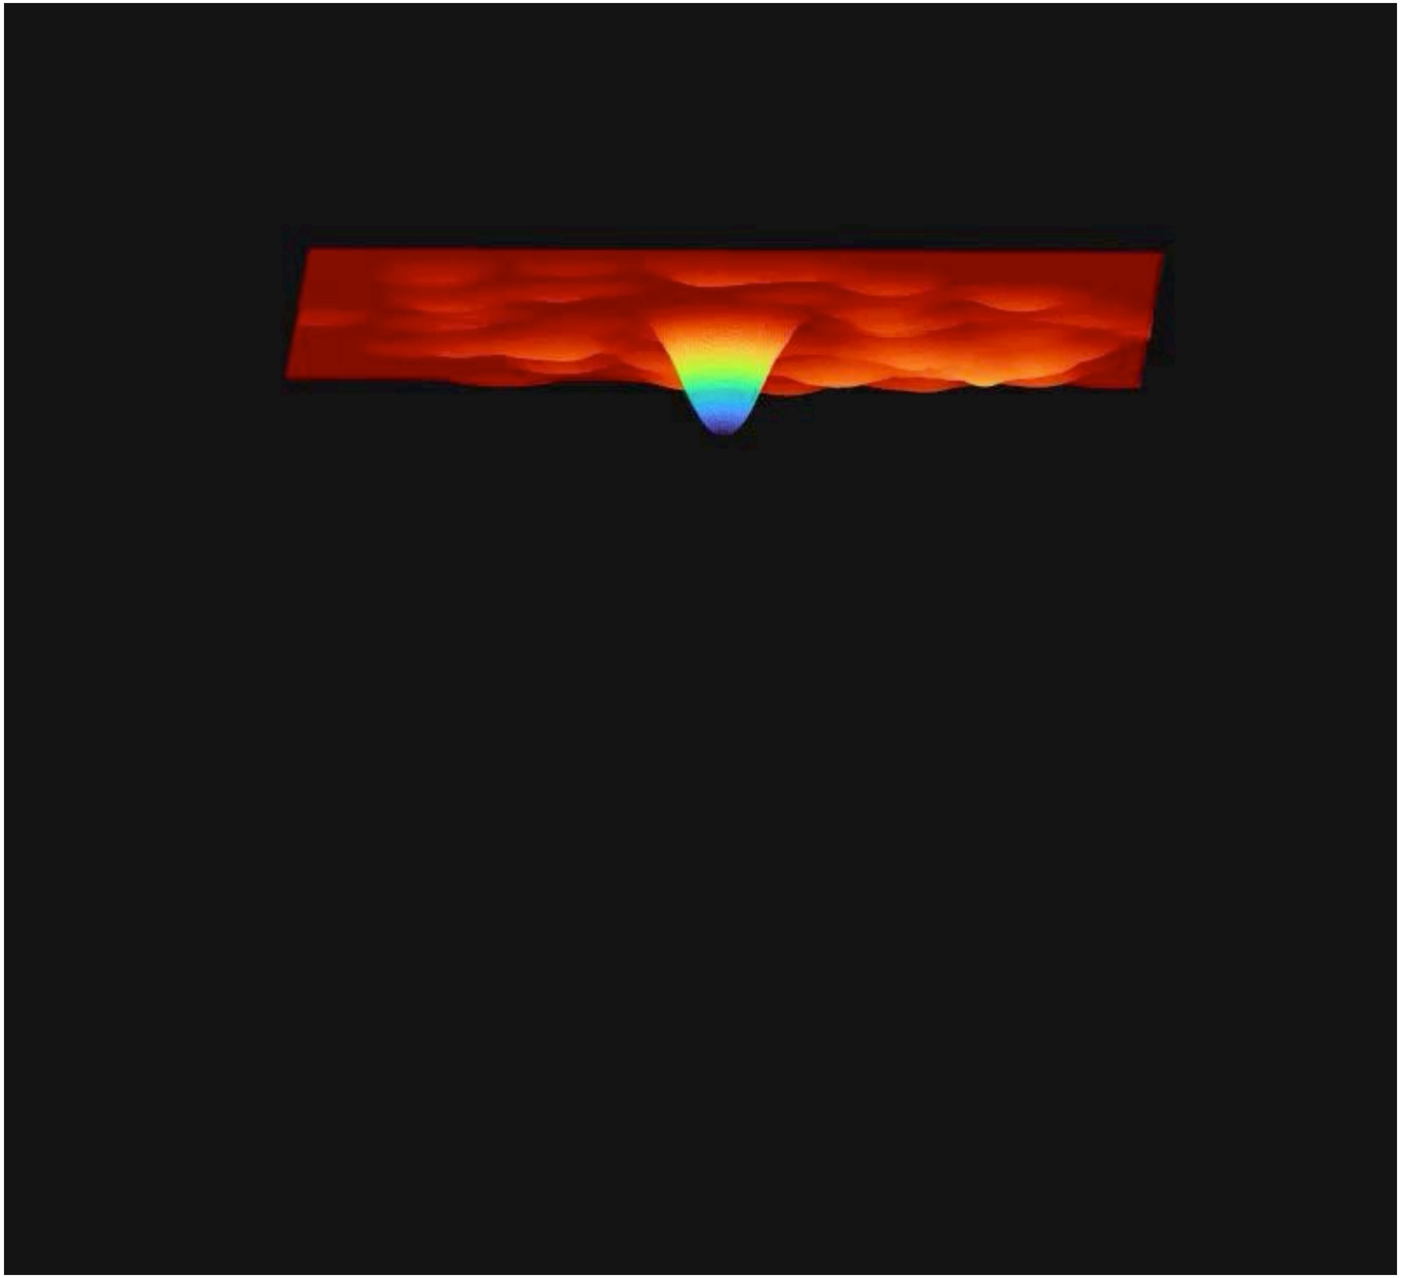

Experimental results

Experimental results[1]

3D registration

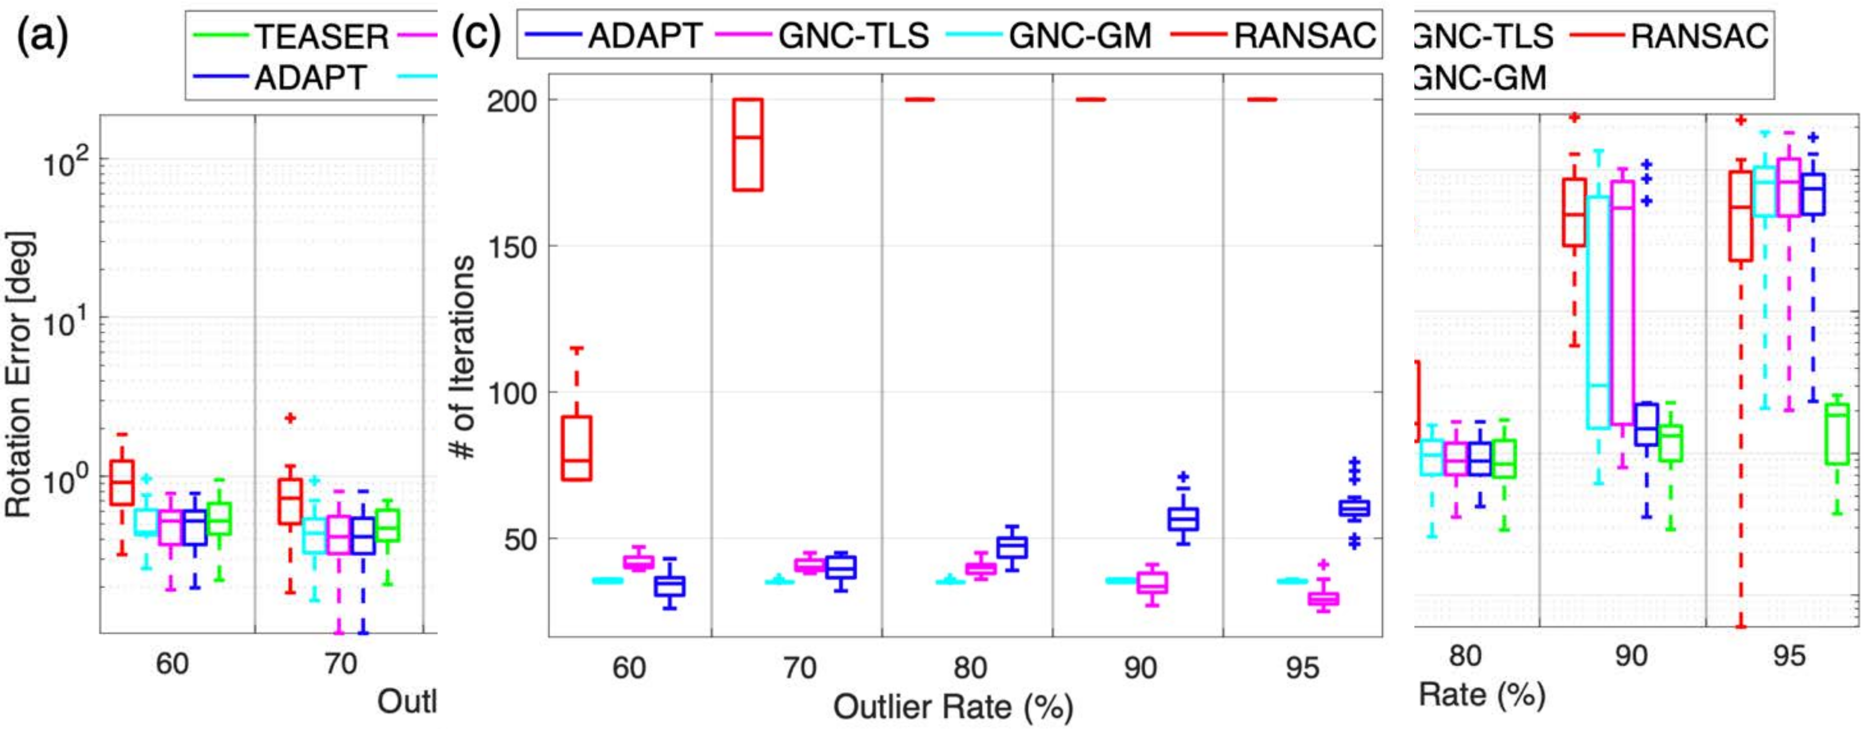

[1] Yang, Antonante, Tzoumas, Carlone, *Graduated Non-Convexity for Robust Spatial Perception: From Non-Minimal Solvers to Global Outlier Rejection,* IEEE Robotics and Automation Letters (RA-L), 2020, and IEEE ICRA 2020.

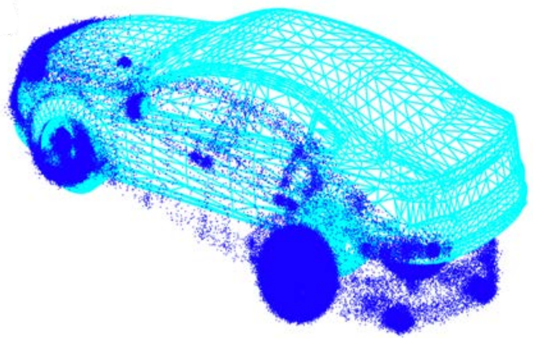

Experimental results

Mesh registration

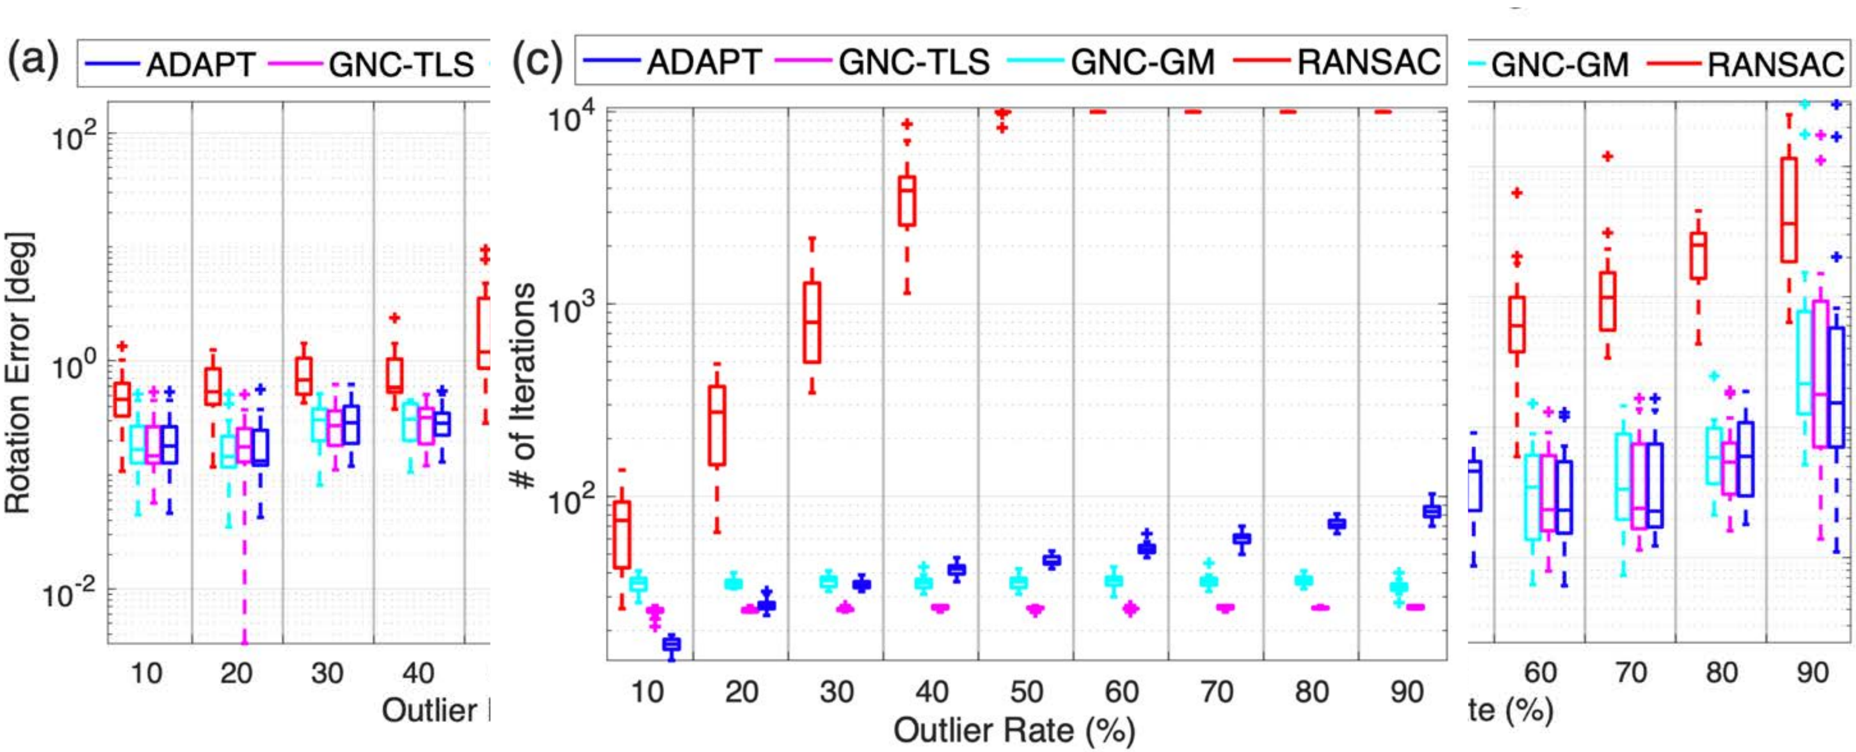

Experimental results

Pose graph optimization

CSAIL

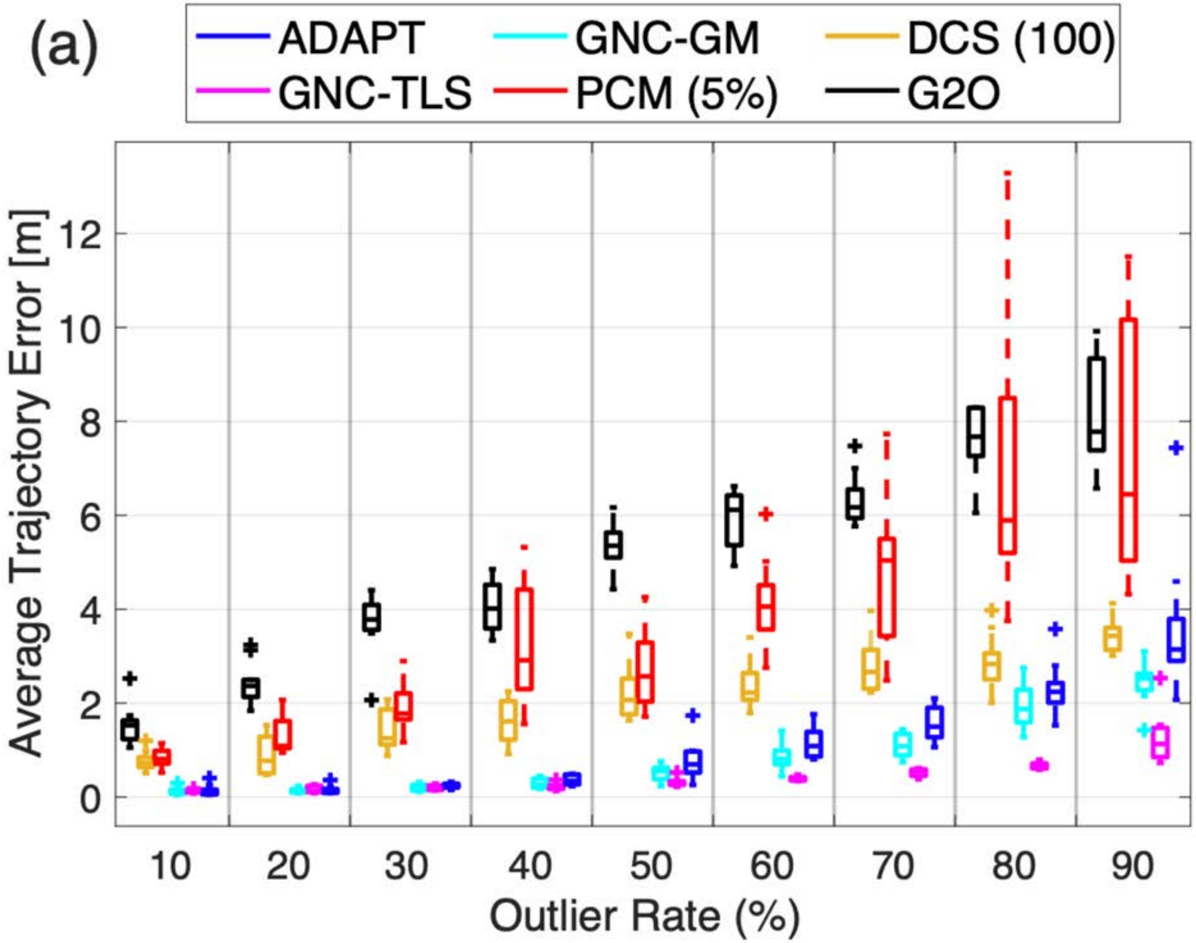

Experimental results

Pose graph optimization

CSAIL

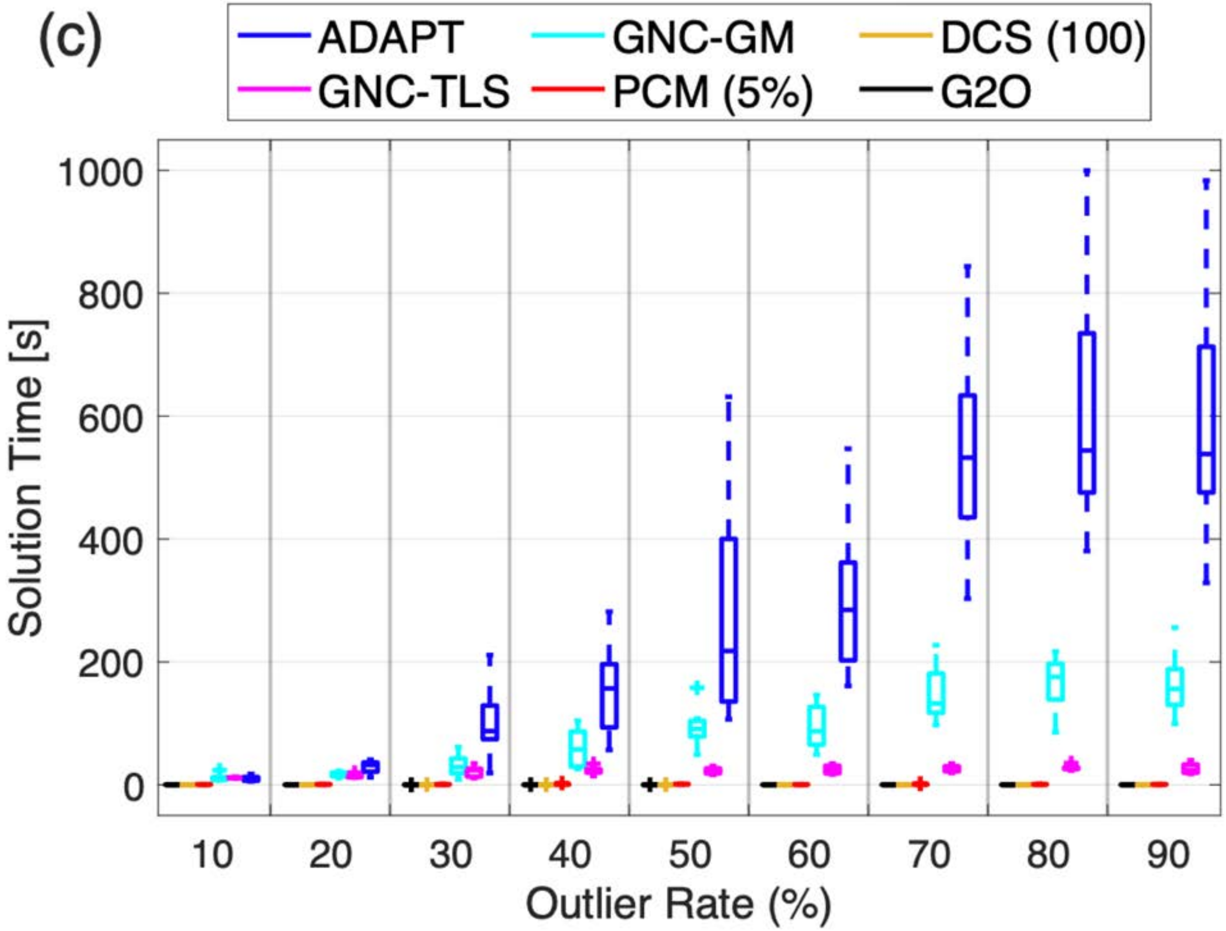

Experimental results

Pose graph optimization

INTEL

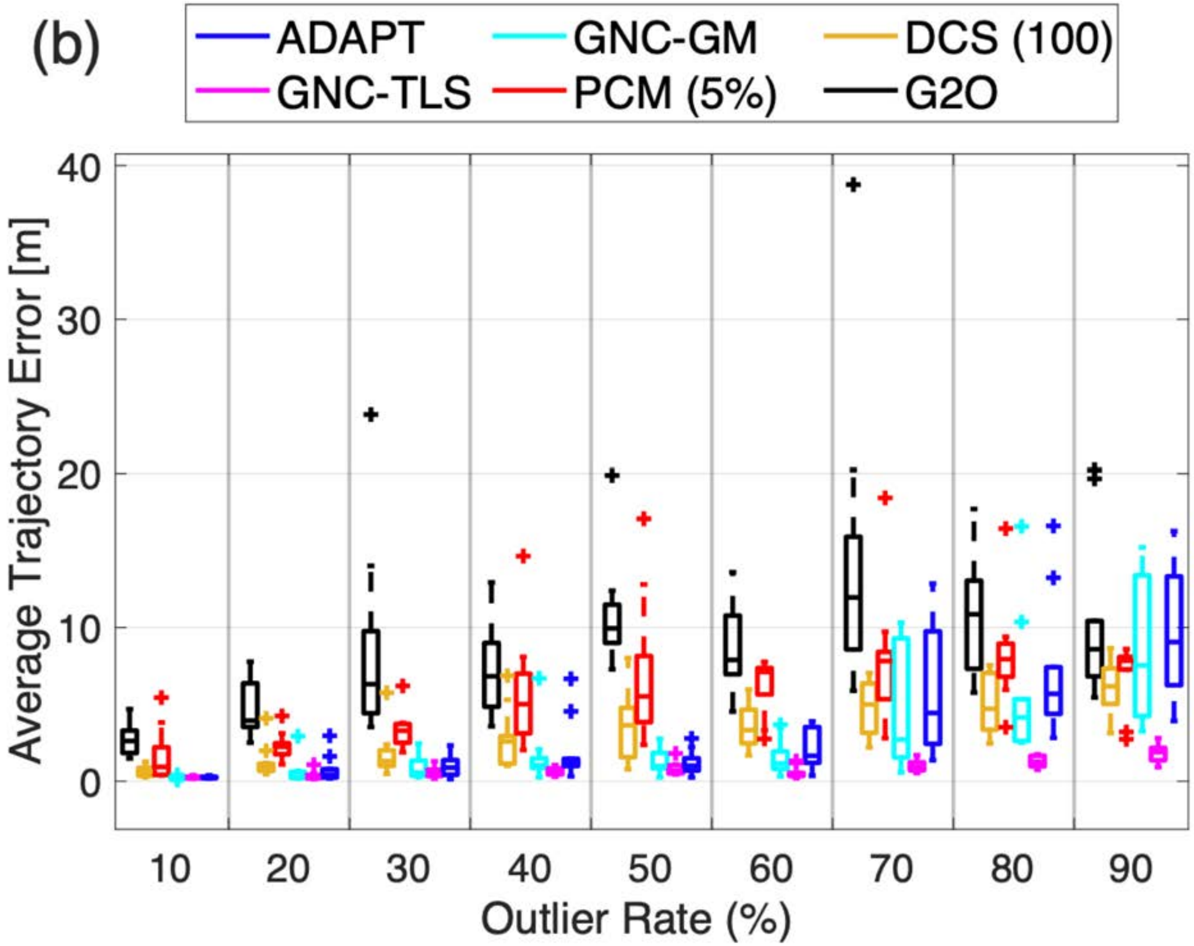

Experimental results

Shape alignment