Outlier-Robust Spatial Perception:
Hardness, Algorithms, Guarantees

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

Recap…

This slide has been intentionally left blank.

Perception as least squares optimization

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

Measurements/data

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

Estimate

Residual

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:

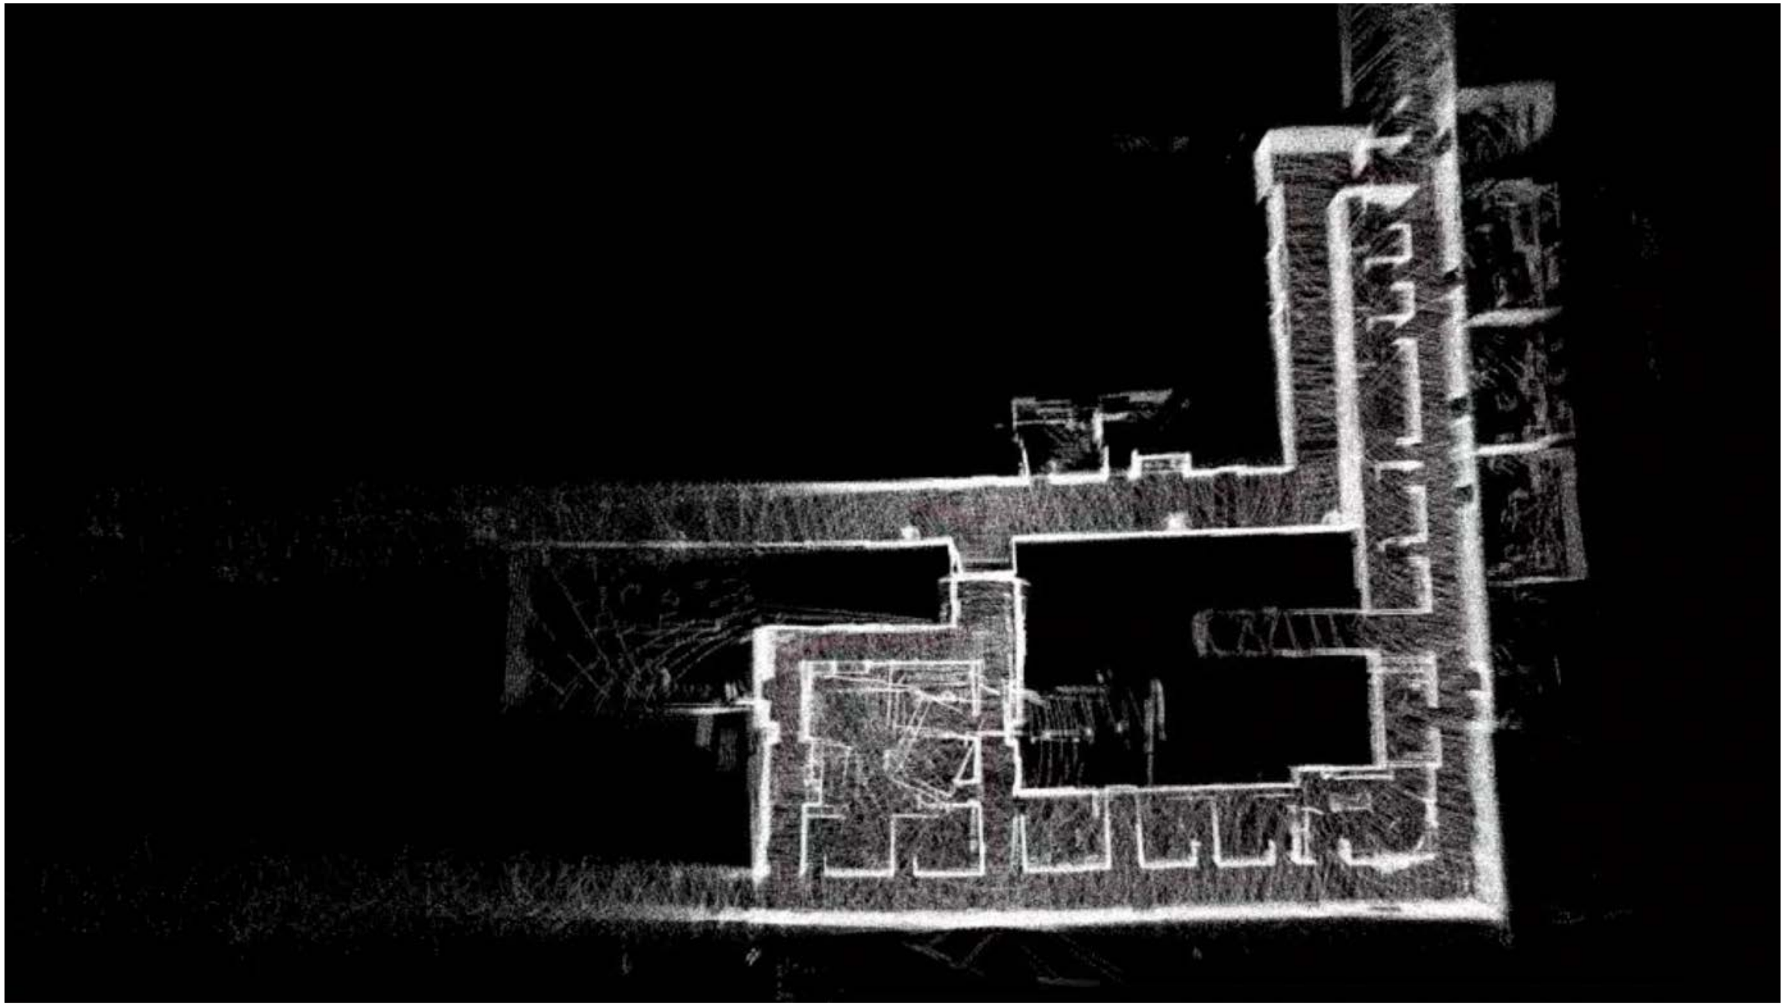

Outlier-robust least squares reformulations

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"

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

Both L and R and harder than NP-hard

Need for effective approximation algorithms

- Fast
- Finds correct x despite many outliers

Last lecture's focus

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

Optimal solvers and graduated non-convexity

Final algorithm (GM case):[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 $x^{(t)}$:

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

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

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

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.

Today's focus

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

Why $\min_{x,\mathcal{O}} |\mathcal{O}| \ s.t. \ \|r(x, y_{\mathcal{M}\setminus\mathcal{O}})\|^2 \leq \epsilon$ can be hard?

Recall: Possible instances of the problem:

Maximum consensus:

$$\min_{x \in \mathcal{X}, \mathcal{O} \subseteq \mathcal{M}} |\mathcal{O}| \quad s.t. \quad r(x, y_i) \leq \bar{c}, \qquad \forall\, i \in \mathcal{M}\setminus\mathcal{O}$$

($\|\cdot\|_\infty$ norm above)

(Outlier rejection) “least squares:”

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

($\|\cdot\|_2$ norm above)

Both are combinatorial problems

↓

Guaranteed outlier removal requires exponential time, e.g., via branch and bound (BnB)

Guaranteed outlier removal via BnB[1]

[1]*Guaranteed Outlier Removal with Mixed Integer Linear Programs*, Chin et al. CVPR 16

We'll develop a method to verify whether a measurement is an outlier

Let's re-write $\min_{x \in \mathcal{X}, \mathcal{O} \subseteq \mathcal{M}} |\mathcal{O}| \;\; s.t. \;\; r(x, y_i) \leq \bar{c}, \; \forall \, i \in \mathcal{M} \backslash \mathcal{O}$ as:

$$\begin{aligned} \underset{,\mathbf{z}}{\text{minimize}} \quad & \sum_i z_i \\ \text{subject to} \quad & \left| \mathbf{x}^T \boldsymbol{\theta}_i - y_i \right| \leq \; + z_i M \\ & z_i \in \{0, 1\}, \end{aligned}$$

large

(P)

Otherwise, y_k is an outlier!

where, for simplicity, $r(x, (\theta_i, y_i)) = x^T \theta_i - y_i$.

Assume y_k is an inlier; then optimal value of AUX-P equals the value of P:

$$\begin{aligned} \underset{,\mathbf{z}}{\text{minimize}} \quad & \sum_{i \neq k} z_i \\ \text{subject to} \quad & \left| \mathbf{x}^T \boldsymbol{\theta}_i - y_i \right| \leq \; + z_i M, \\ & z_i \in \{0, 1\}, \\ & \left| \mathbf{x}^T \boldsymbol{\theta}_k - y_k \right| \leq \; . \end{aligned}$$

(AUX-P)

Guaranteed outlier removal via BnB[1]

Goal: Show that P and AUX-P have different values to prove y_k is outlier

$$\underset{,\mathbf{z}}{\text{minimize}} \quad \sum_{i} z_i$$
$$\text{subject to} \quad \left|\mathbf{x}^T \boldsymbol{\theta}_i - y_i\right| \leq \; + z_i M$$
$$z_i \in \{0, 1\},$$

(P)

$$\underset{,\mathbf{z}}{\text{minimize}} \quad \sum_{i \neq k} z_i$$
$$\text{subject to} \quad \left|\mathbf{x}^T \boldsymbol{\theta}_i - y_i\right| \leq \; + z_i M,$$
$$z_i \in \{0, 1\},$$
$$\left|\mathbf{x}^T \boldsymbol{\theta}_k - y_k\right| \leq \; .$$

(AUX-P)

Recall:

Finding values of AUX-P and P is hard → Approximation method to reach **Goal**

- Find **upper** bound $\hat{u}$ to P's value
- Find **lower** bound α^k to AUX-P

Lemma 1 *If $\alpha^k > \hat{u}$, then $\{\boldsymbol{\theta}_k, y_k\}$ is a true outlier.*

Guaranteed outlier removal via BnB[1]

Lemma 1 *If $\alpha^k > \hat{u}$, then $\{、\,_k, y_k\}$ is a true outlier.*

How to efficiently find $\hat{u}$ and α^k?

- **Upper** bound $\hat{u}$ to P's value:
  - a fast way to find $\hat{u}$ is by using RANSAC

- **Lower** bound α^k to AUX-P:
  - **Use BnB instead:**[2] BnB is an iterative method, where at each iteration t finds lower bound α_t^k, and an upper bound γ_t^k to the value of AUX-P (tighter after each iteration; terminates when $\alpha_t^k = \gamma_t^k$, in the worst-case after exponential time).

Run BnB until $\alpha_t^k > \hat{u}$ ($\Rightarrow y_k$ outlier) or $\gamma_t^k \leq \hat{u}$ ($\Rightarrow \alpha_t^k \leq \hat{u}$)

[2]https://web.stanford.edu/class/ee364b/lectures/bb_slides.pdf

Faster methods for $\min_{x,\mathcal{O}} |\mathcal{O}| \;\; s.t. \;\; \|r(x, y_{\mathcal{M}\setminus\mathcal{O}})\|^2 \leq \epsilon$

Previous BnB method can be effective for even > 95% of outliers, but slow…

↓

Approximation algorithms

- **RANSAC**: ineffective > 50% of outliers;
  impractical for SLAM

- **Greedy algorithms:**[1] Can fail for > 50% of outliers (can quickly hit local minima);
  Quadratic running time so impractical for SLAM

- **Adaptive trimming (ADAPT):**[2,3] Has been observed to withstand: < 90% registration
  < 70-80% two-view
  < 70% SLAM

  Linear running time (slower than GNC in SLAM)

[1]Nemhauser, Wolsey, Fisher 78; Rousseeuw 87

[2]Tzoumas, Antonante, Carlone, IROS 19

[3]Antonante, Tzoumas, Yang, Carlone, arXiv:2007.15109, 2020.

ADAPT: ADAPtive Trimming

ADAPT adaptively rejects measurements with large residuals:

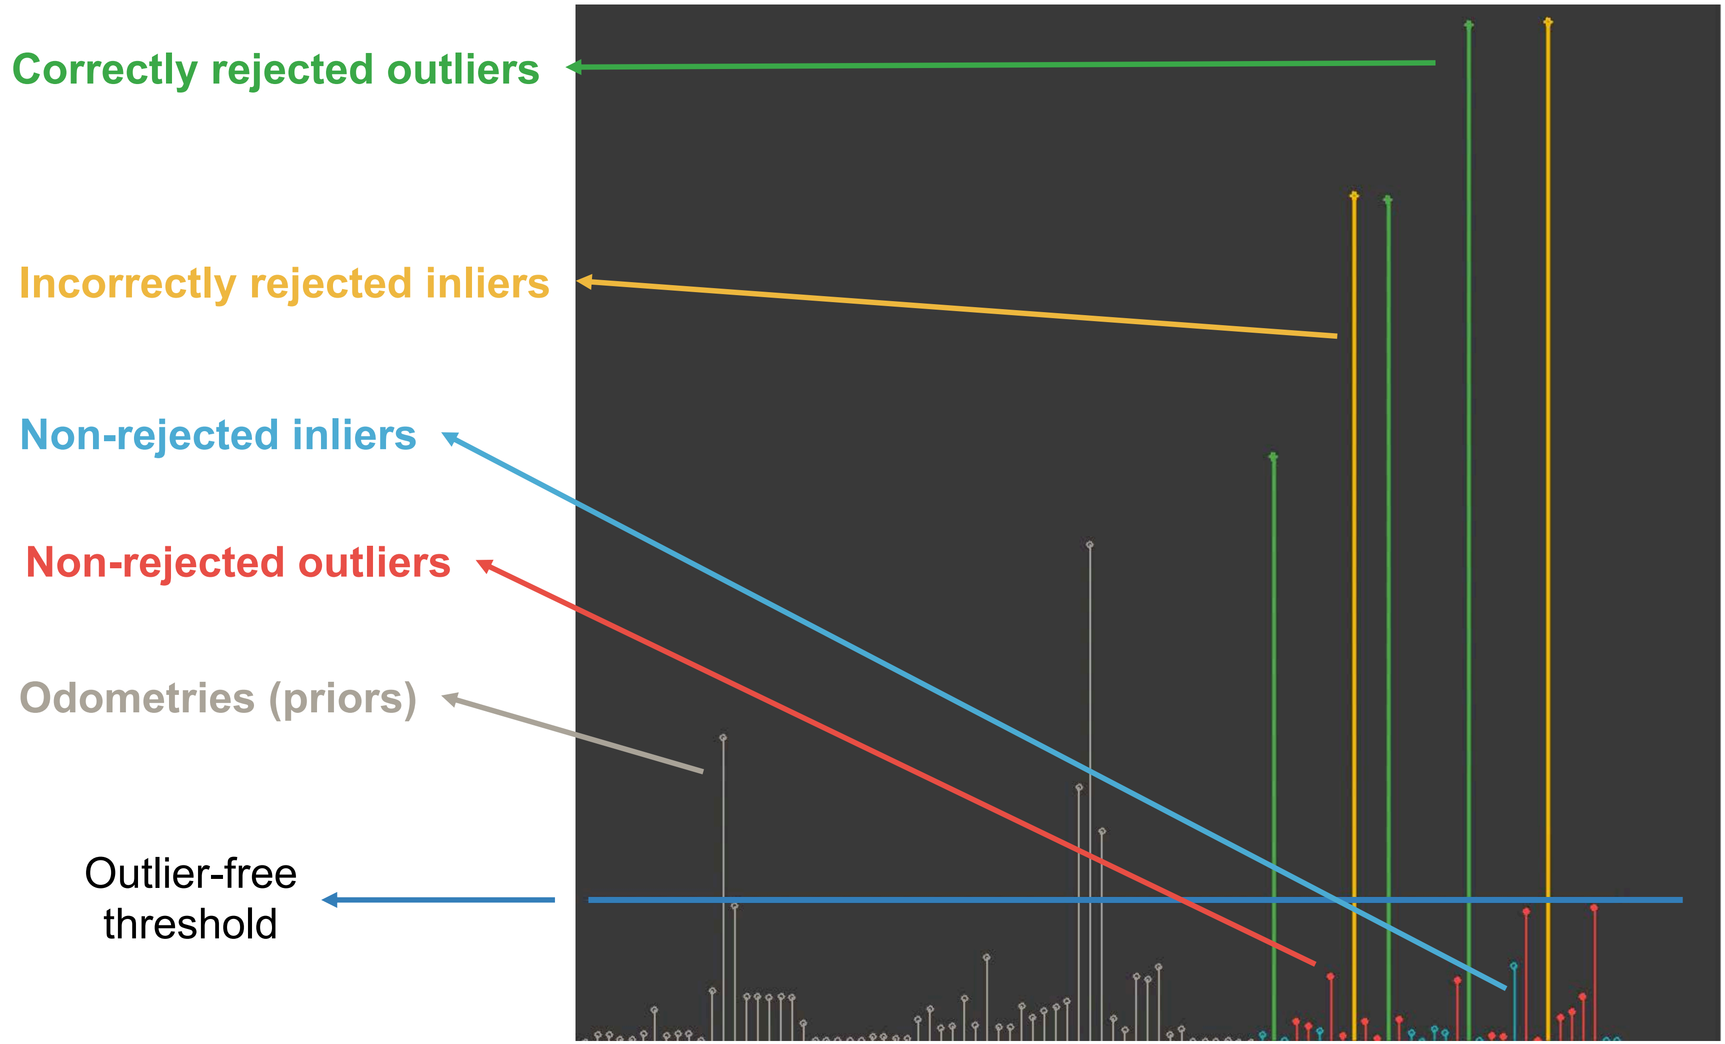

ADAPT on SLAM 2D grid

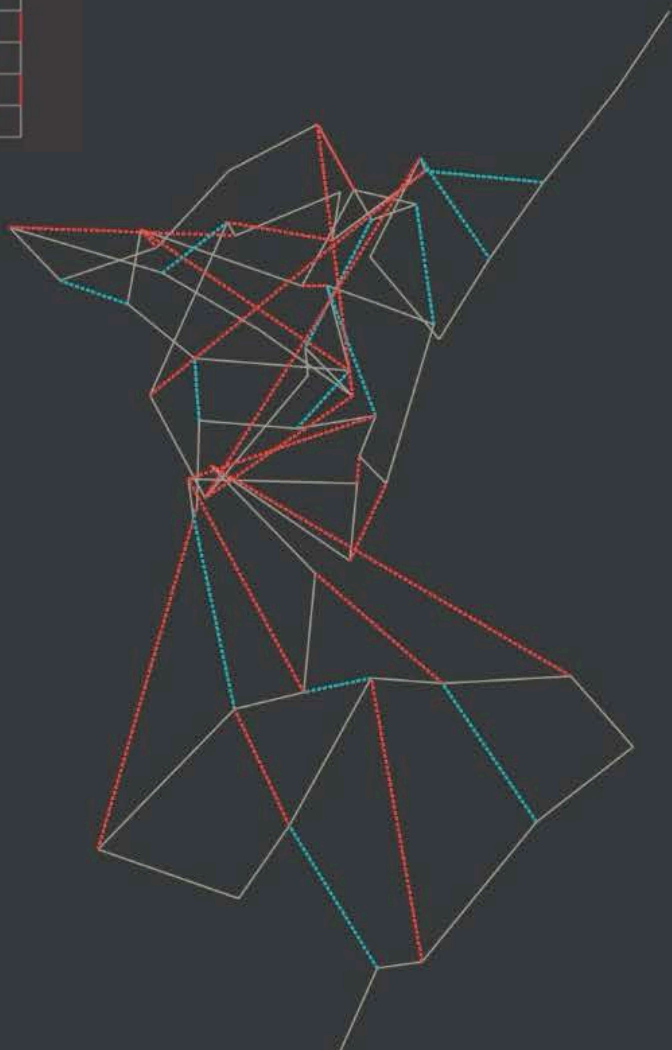

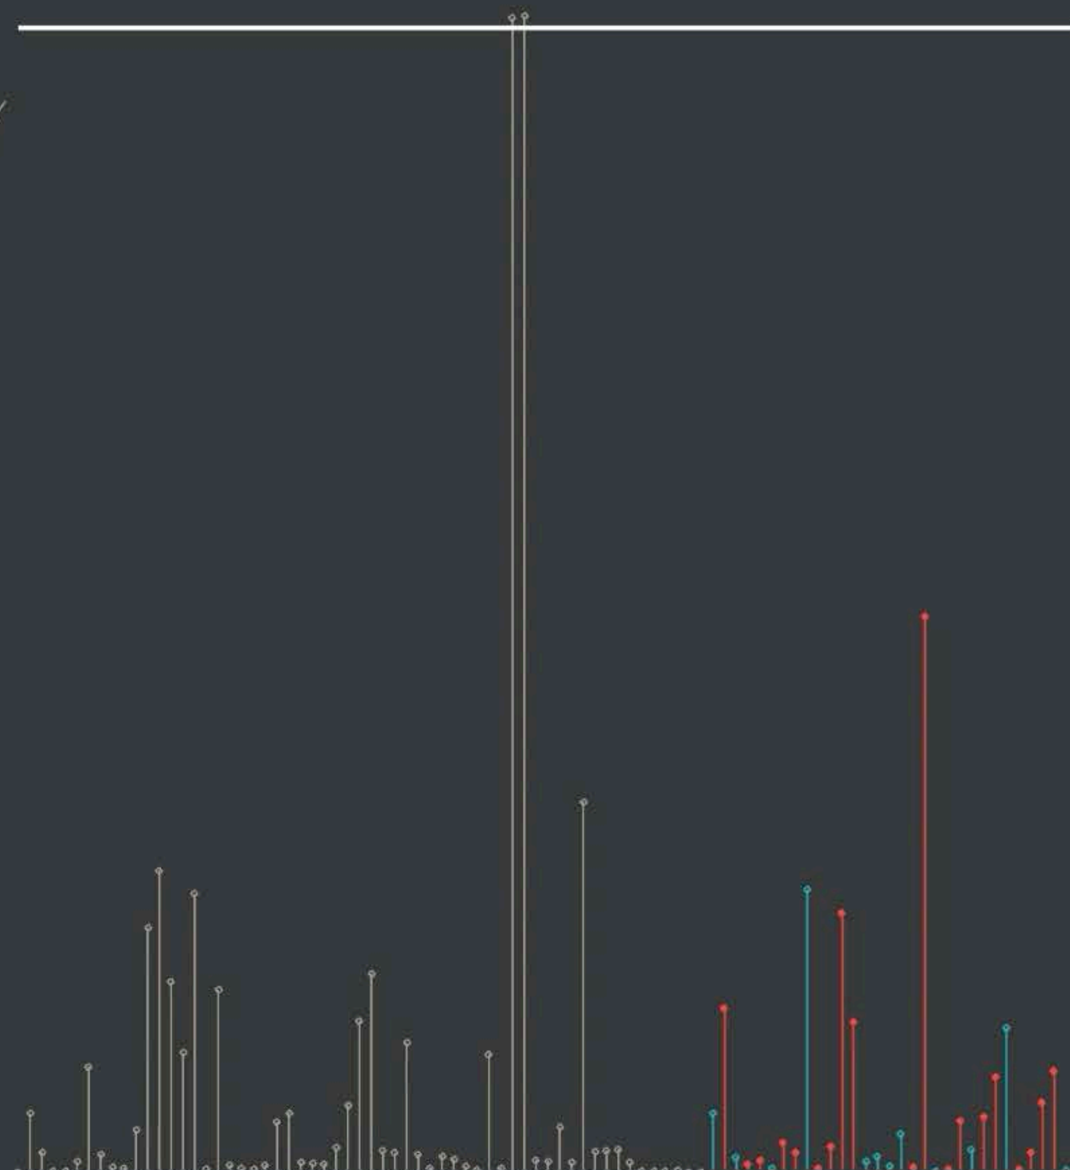

ADAPT

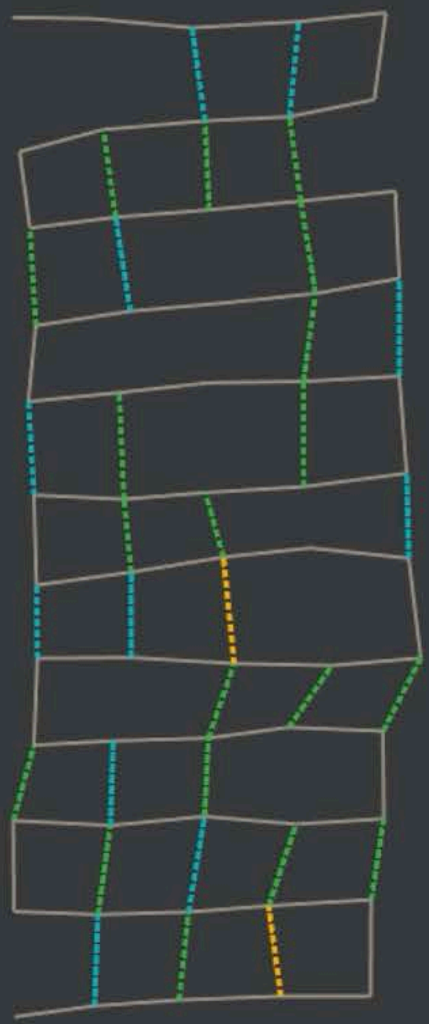

Ground truth

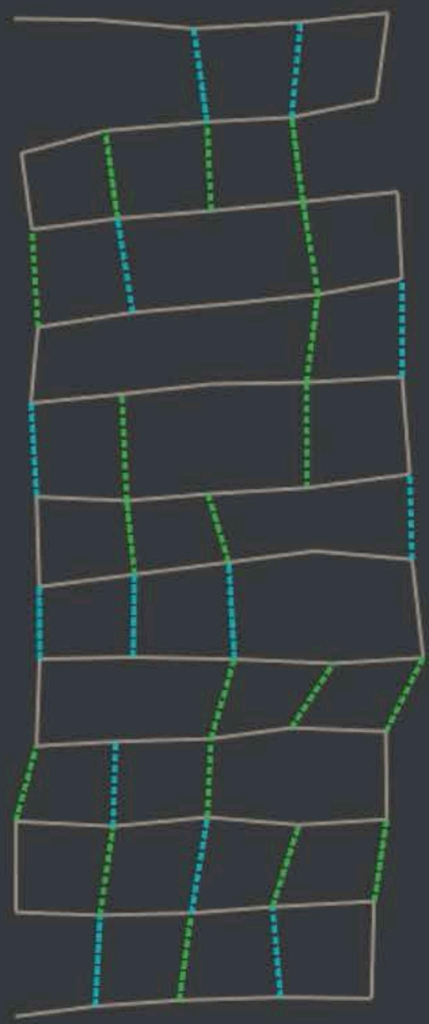

Experimental results

This slide has been intentionally left blank.

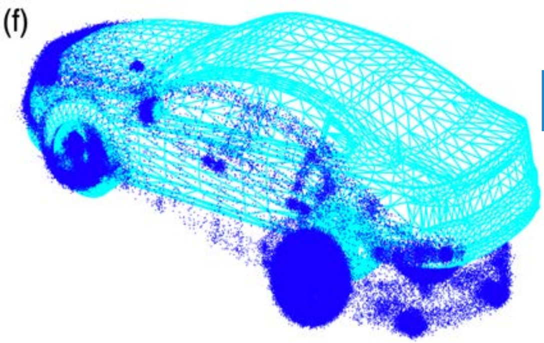

Experimental results[1,2]

Mesh registration

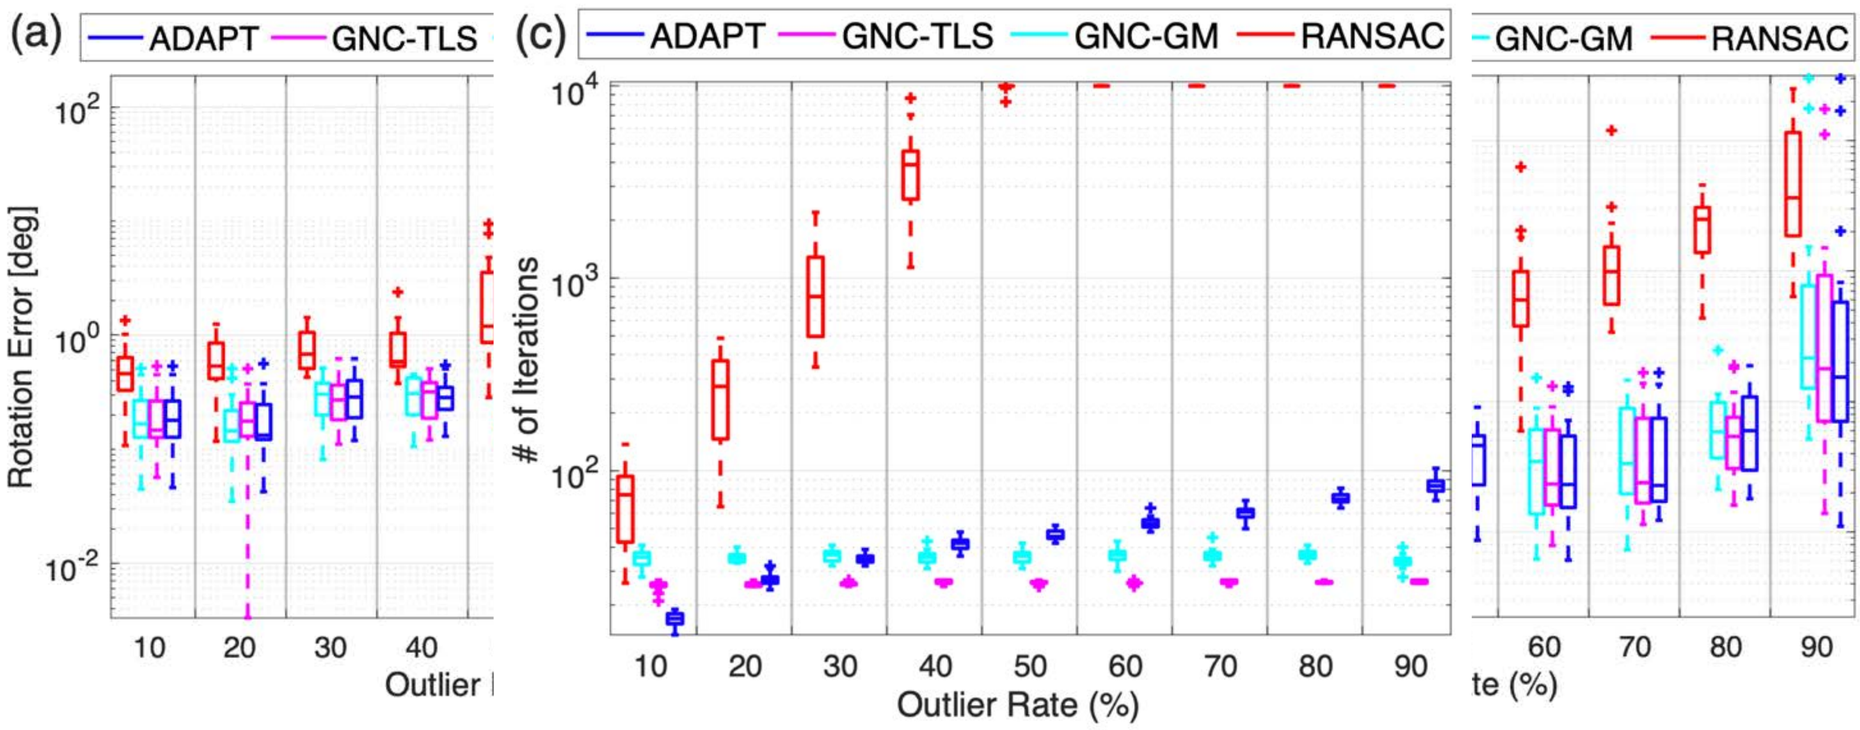

[1] Yang, Antonante, Tzoumas, Carlone, *Graduated Non-Convexity for Robust Spatial Perception: From Non-Minimal Solvers to Global Outlier Rejection,* IEEE Robotics and Automation Letters (RA-L), 2020, and IEEE ICRA 2020.
[2] Antonante, Tzoumas, Yang, Carlone, *Outlier-robust estimation: Hardness, Minimally-Tuned Algorithms, and Applications,* arXiv:2007.15109, 2020.

Experimental results

Pose graph optimization

CSAIL

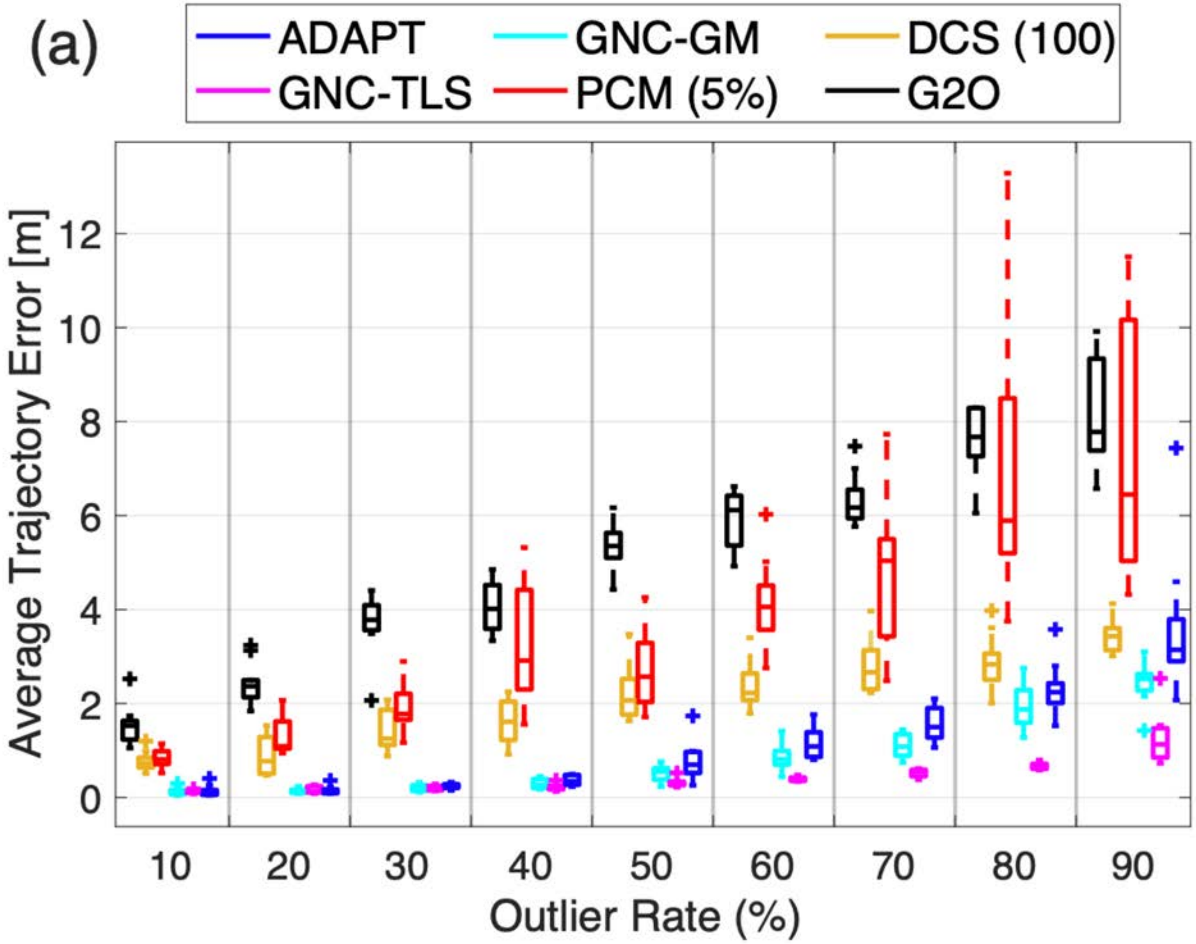

Experimental results

Pose graph optimization

CSAIL

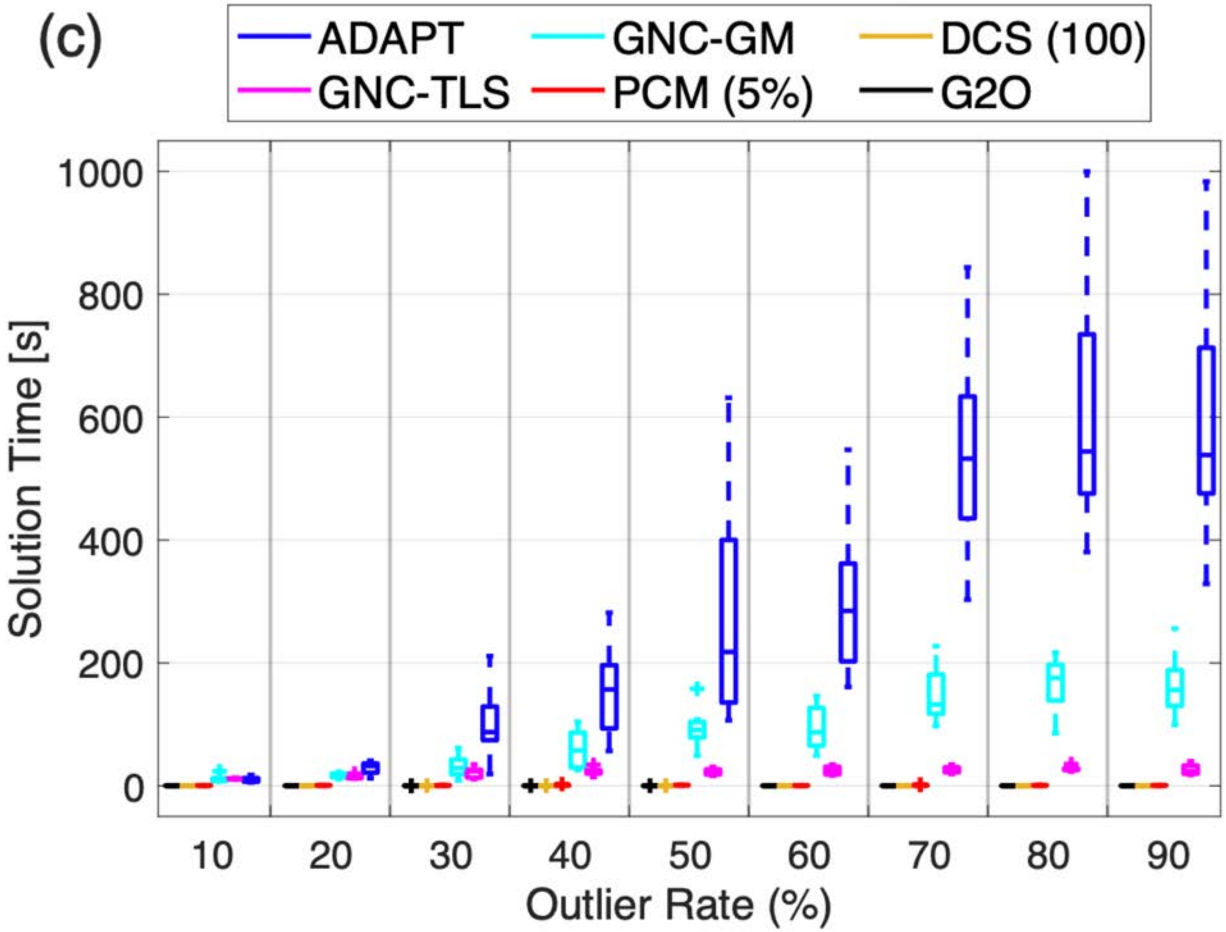

Experimental results

Shape alignment

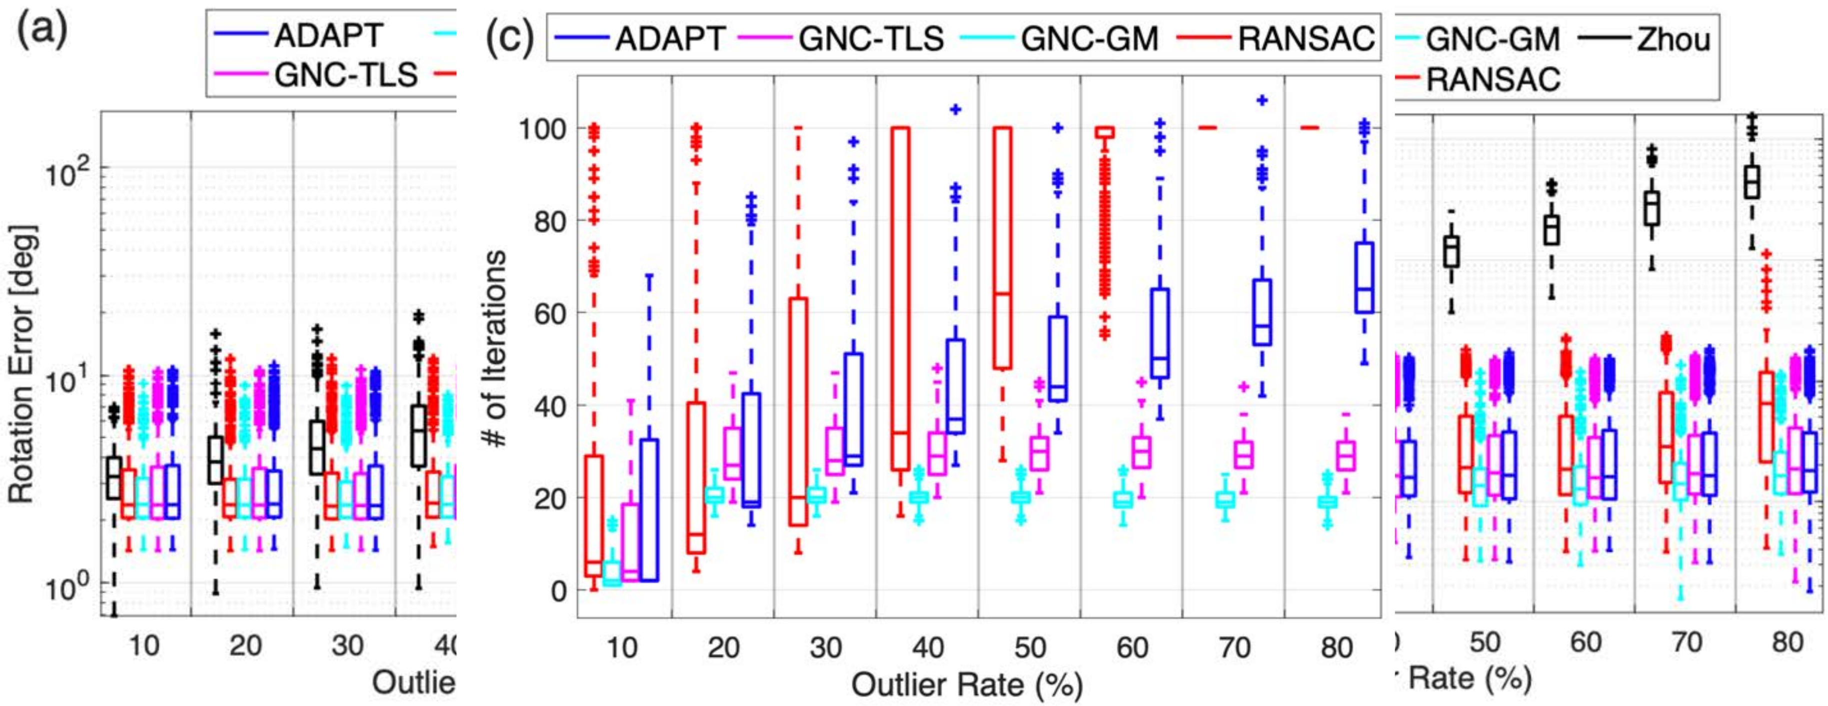

What if $\bar{c}$ is unknown?

Extension of Graduated Non-Convexity (GNC) and ADAPT to unknown $\bar{c}$:

Antonante, Tzoumas, Yang, Carlone, *Outlier-robust estimation: Hardness, Minimally-Tuned Algorithms, and Applications,* arXiv:2007.15109, 2020.

Certifiable Outlier-Robust Optimization?

Extension of Graduated Non-Convexity (GNC) and ADAPT to unknown $\bar{c}$:

Yang, Carlone, *One Ring to Rule Them All: Certifiably Robust Geometric Perception with Outliers,* NeurIPS, 2020.