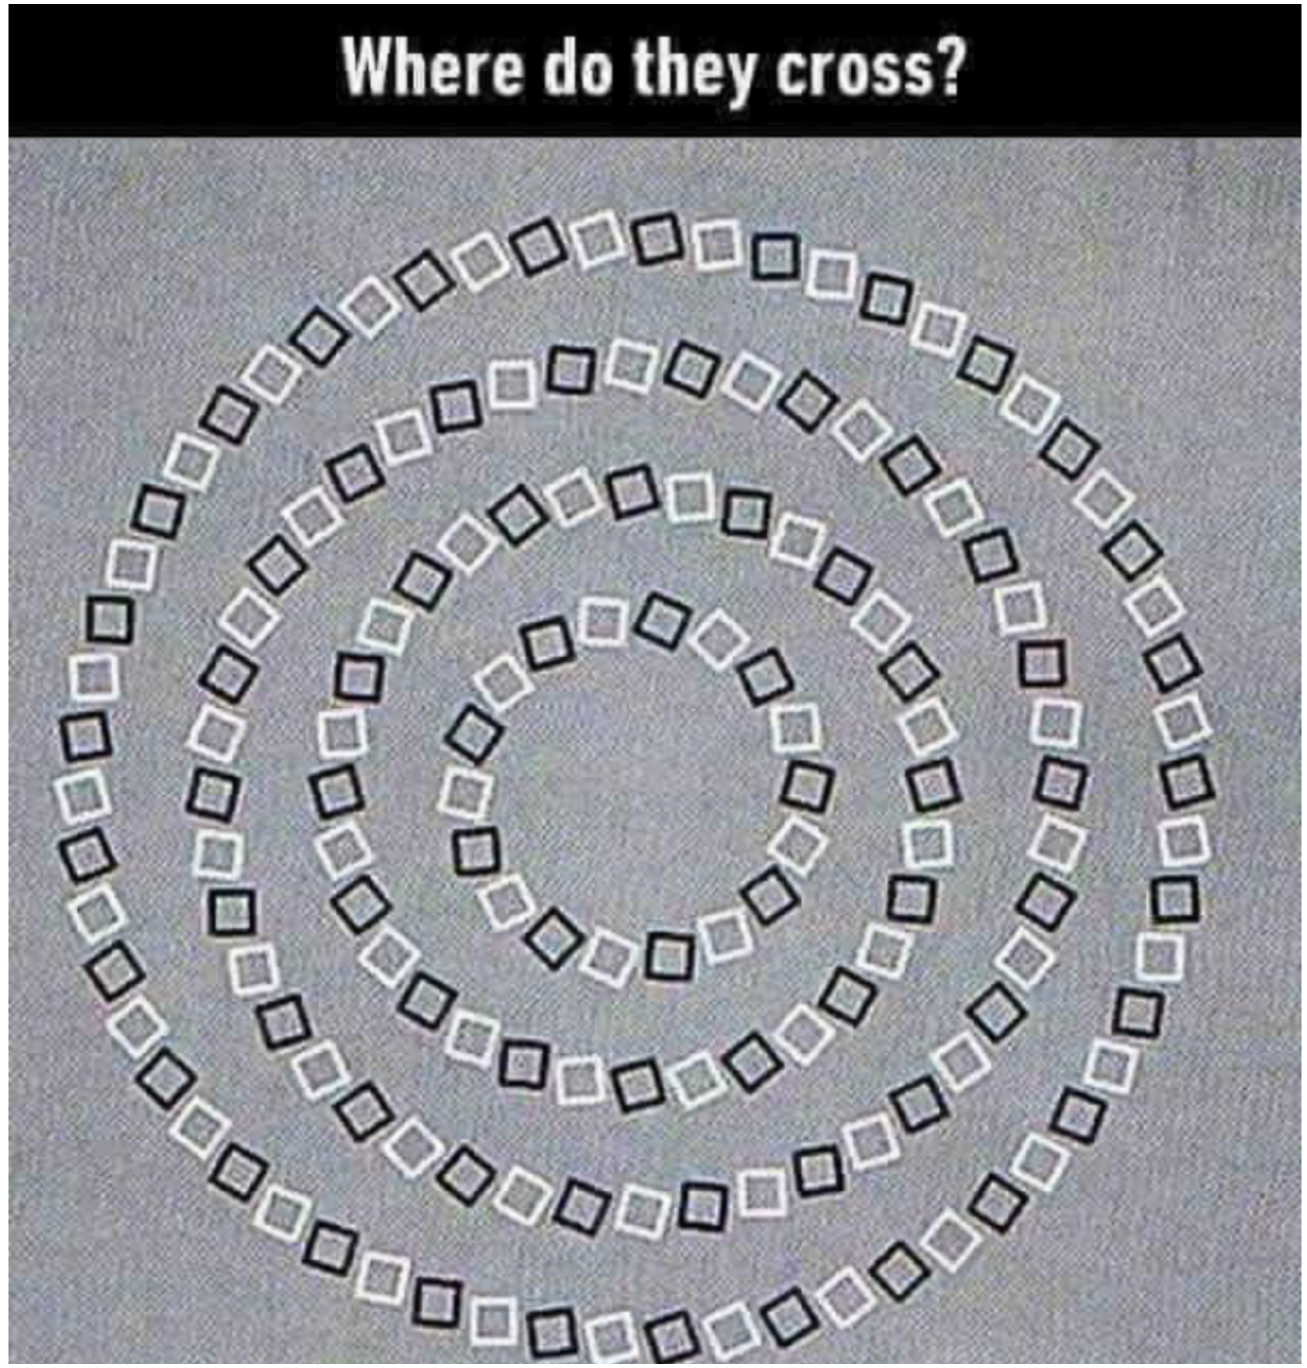

16.485: VNAV - Visual Navigation for Autonomous Vehicles

Luca Carlone

Lecture 29-30: Robust Estimation

based on slides by Vasileios Tzoumas and Cyrill Stachniss

Today and Next Lecture

- **Robust estimation**:
  - Motivations: outliers, data association
  - Formulations: M-estimation & Maximum Consensus
- **Solvers for robust estimation**:
  - (RANSAC)
  - Iteratively Reweighted Least Squares (IRLS)
  - Max-mixture
  - Switchable constraints
  - Graduated non-convexity
  - Others: BnB, SDP relaxations, graph-theoretic pruning

Some problems in VNAV

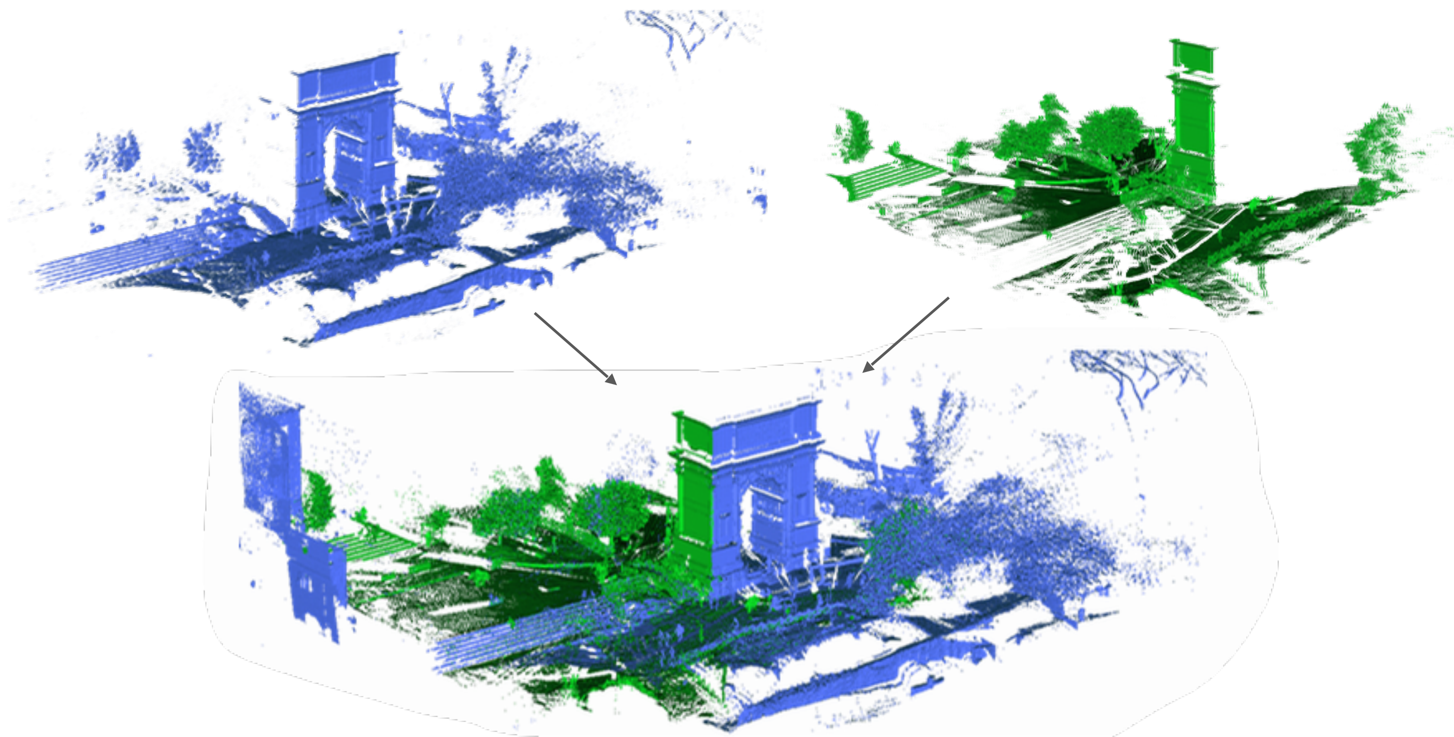

Point cloud registration

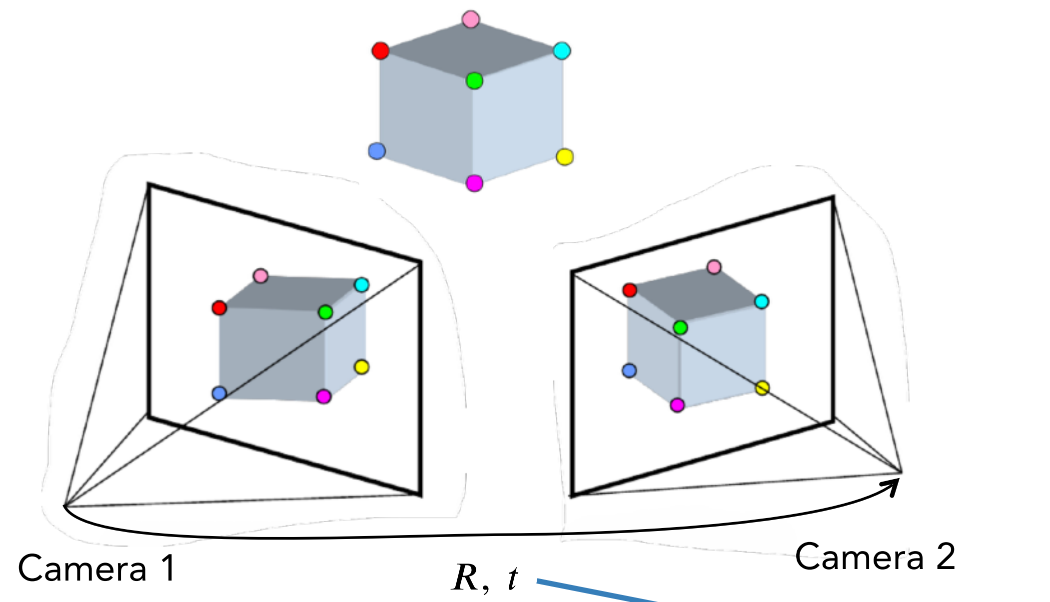

Two-view geometry

R: rotation
t: translation

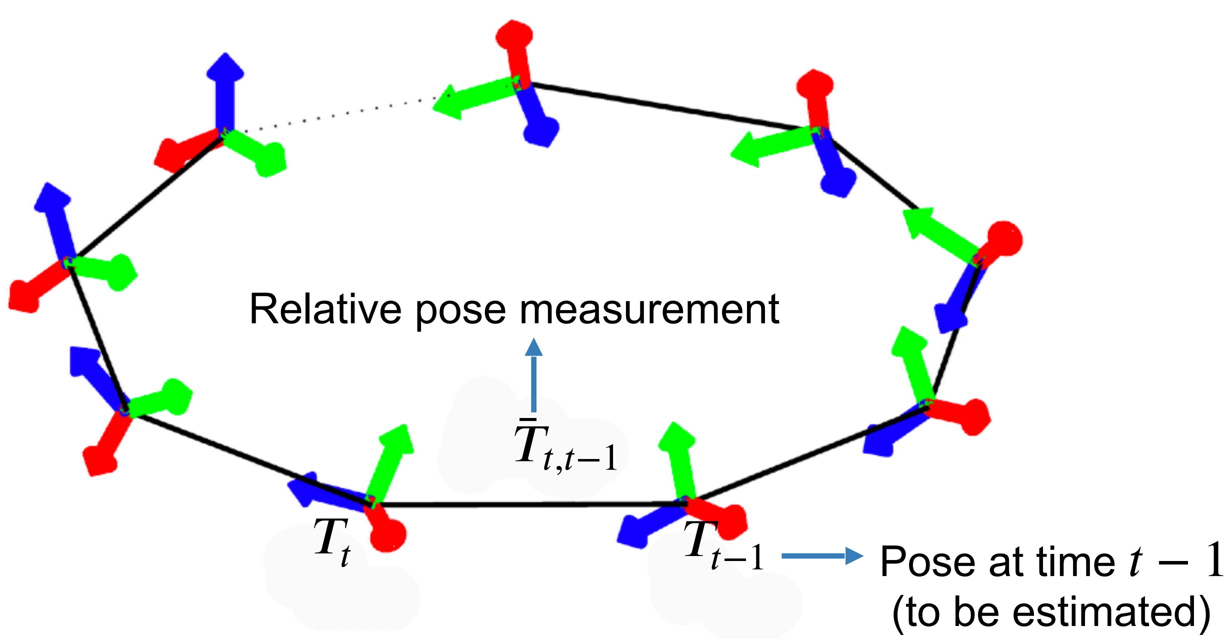

Pose graph optimization

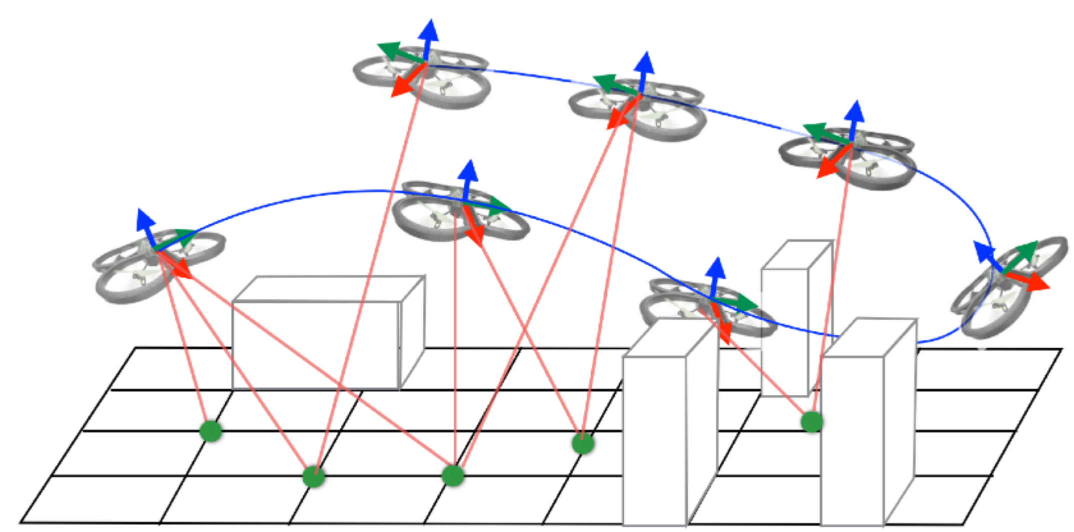

Landmark-based SLAM

Outliers in 2-view geometry

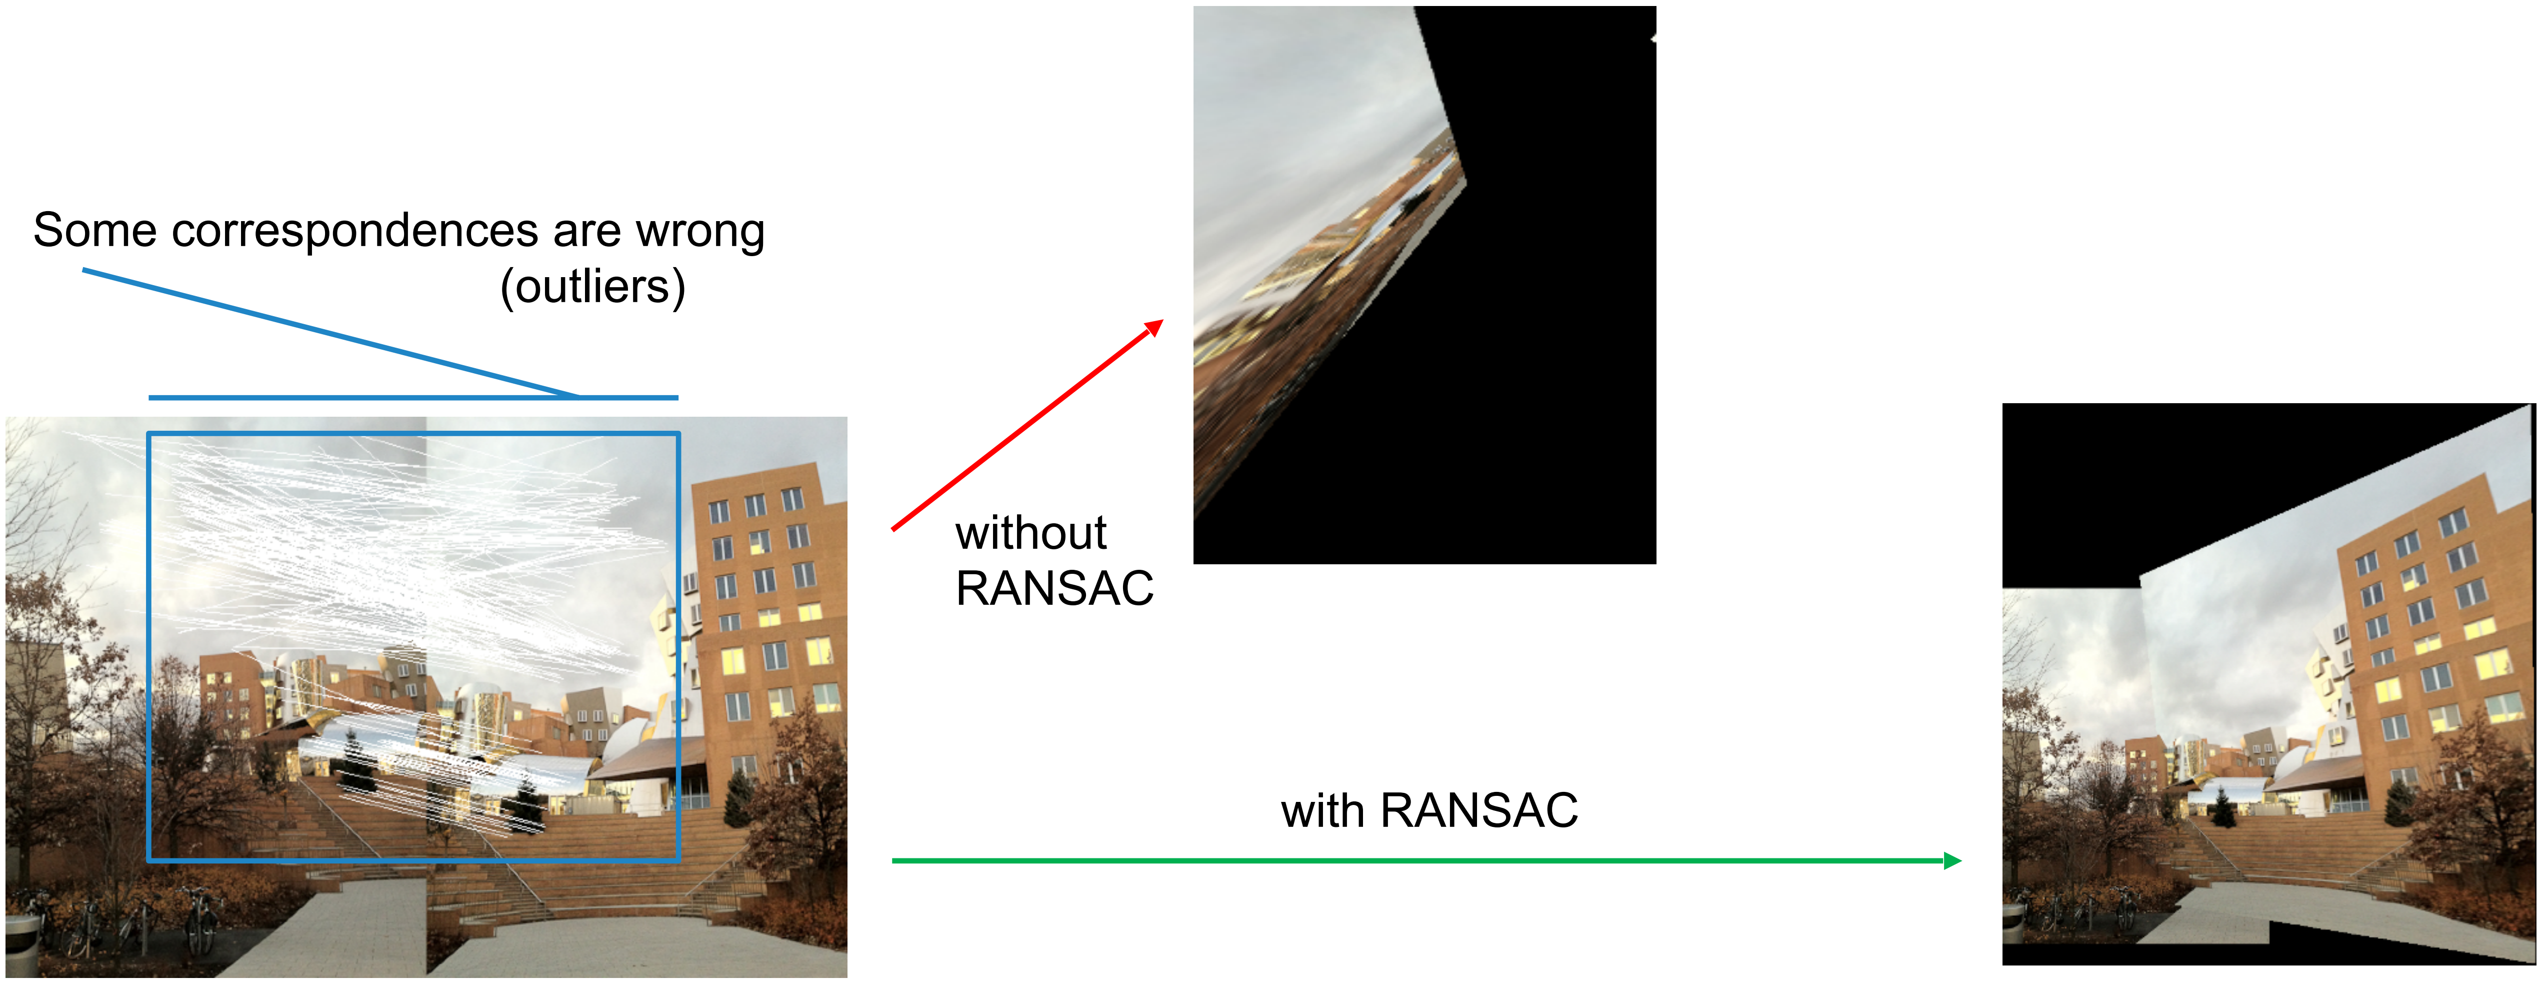

Outliers: uninformative/incorrect measurements

RANSAC to the rescue but only applies to problems
(i) where estimation can be performed from a small set of measurements
(ii) for which a fast minimal solver is available
(iii) for which there are not many outliers

Outliers in pose graph optimization

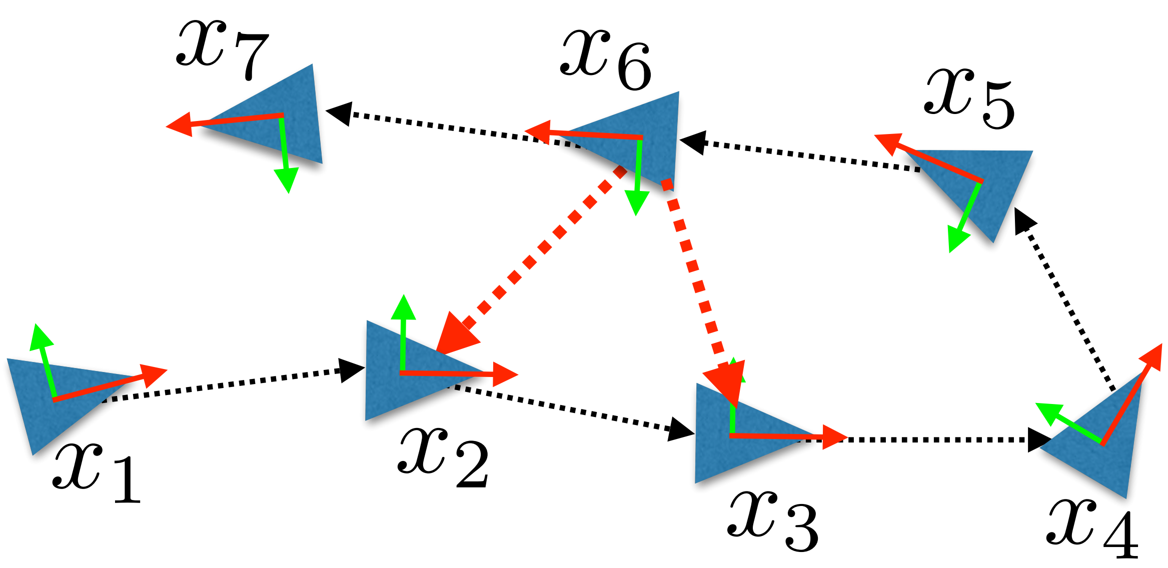

outliers: completely incorrect measurements (Perceptual Aliasing)

0 outliers

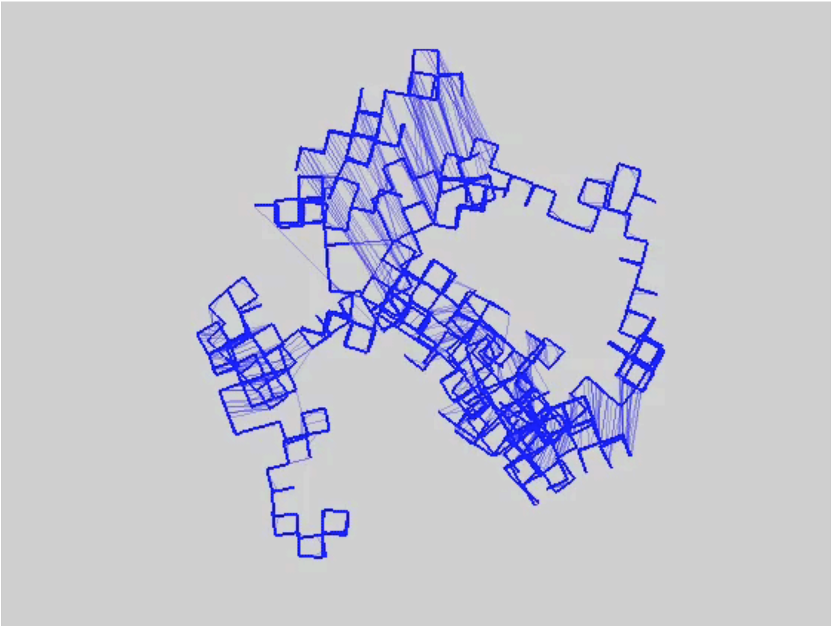

5 outliers

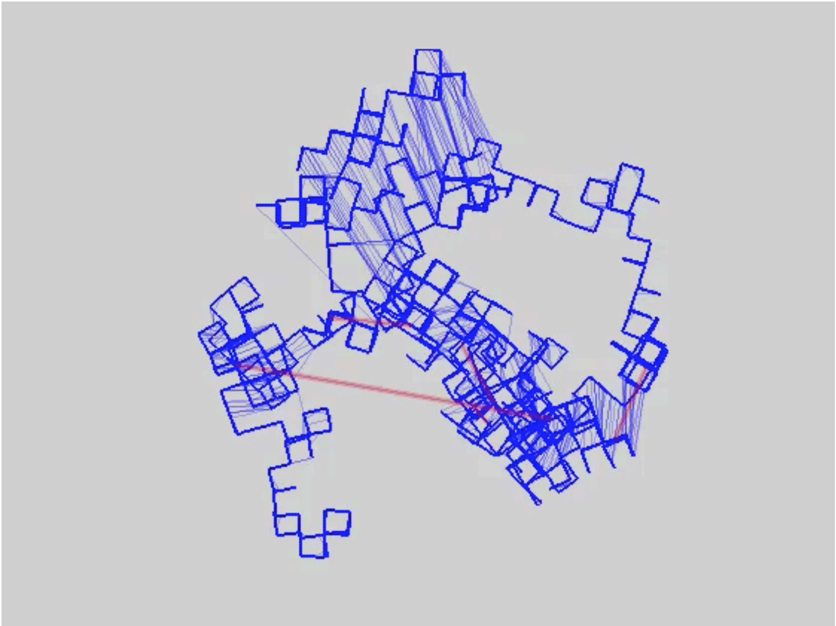

20 outliers

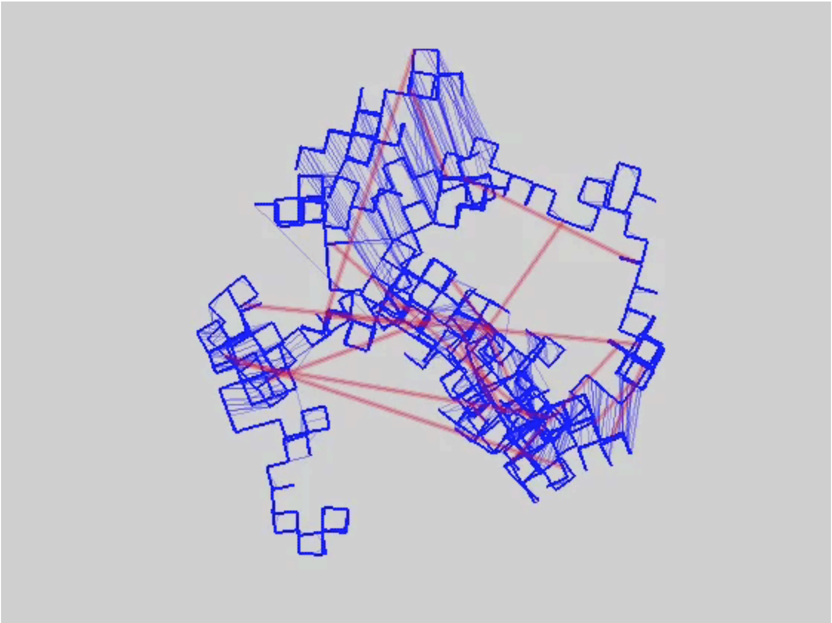

Outliers in pose graph optimization

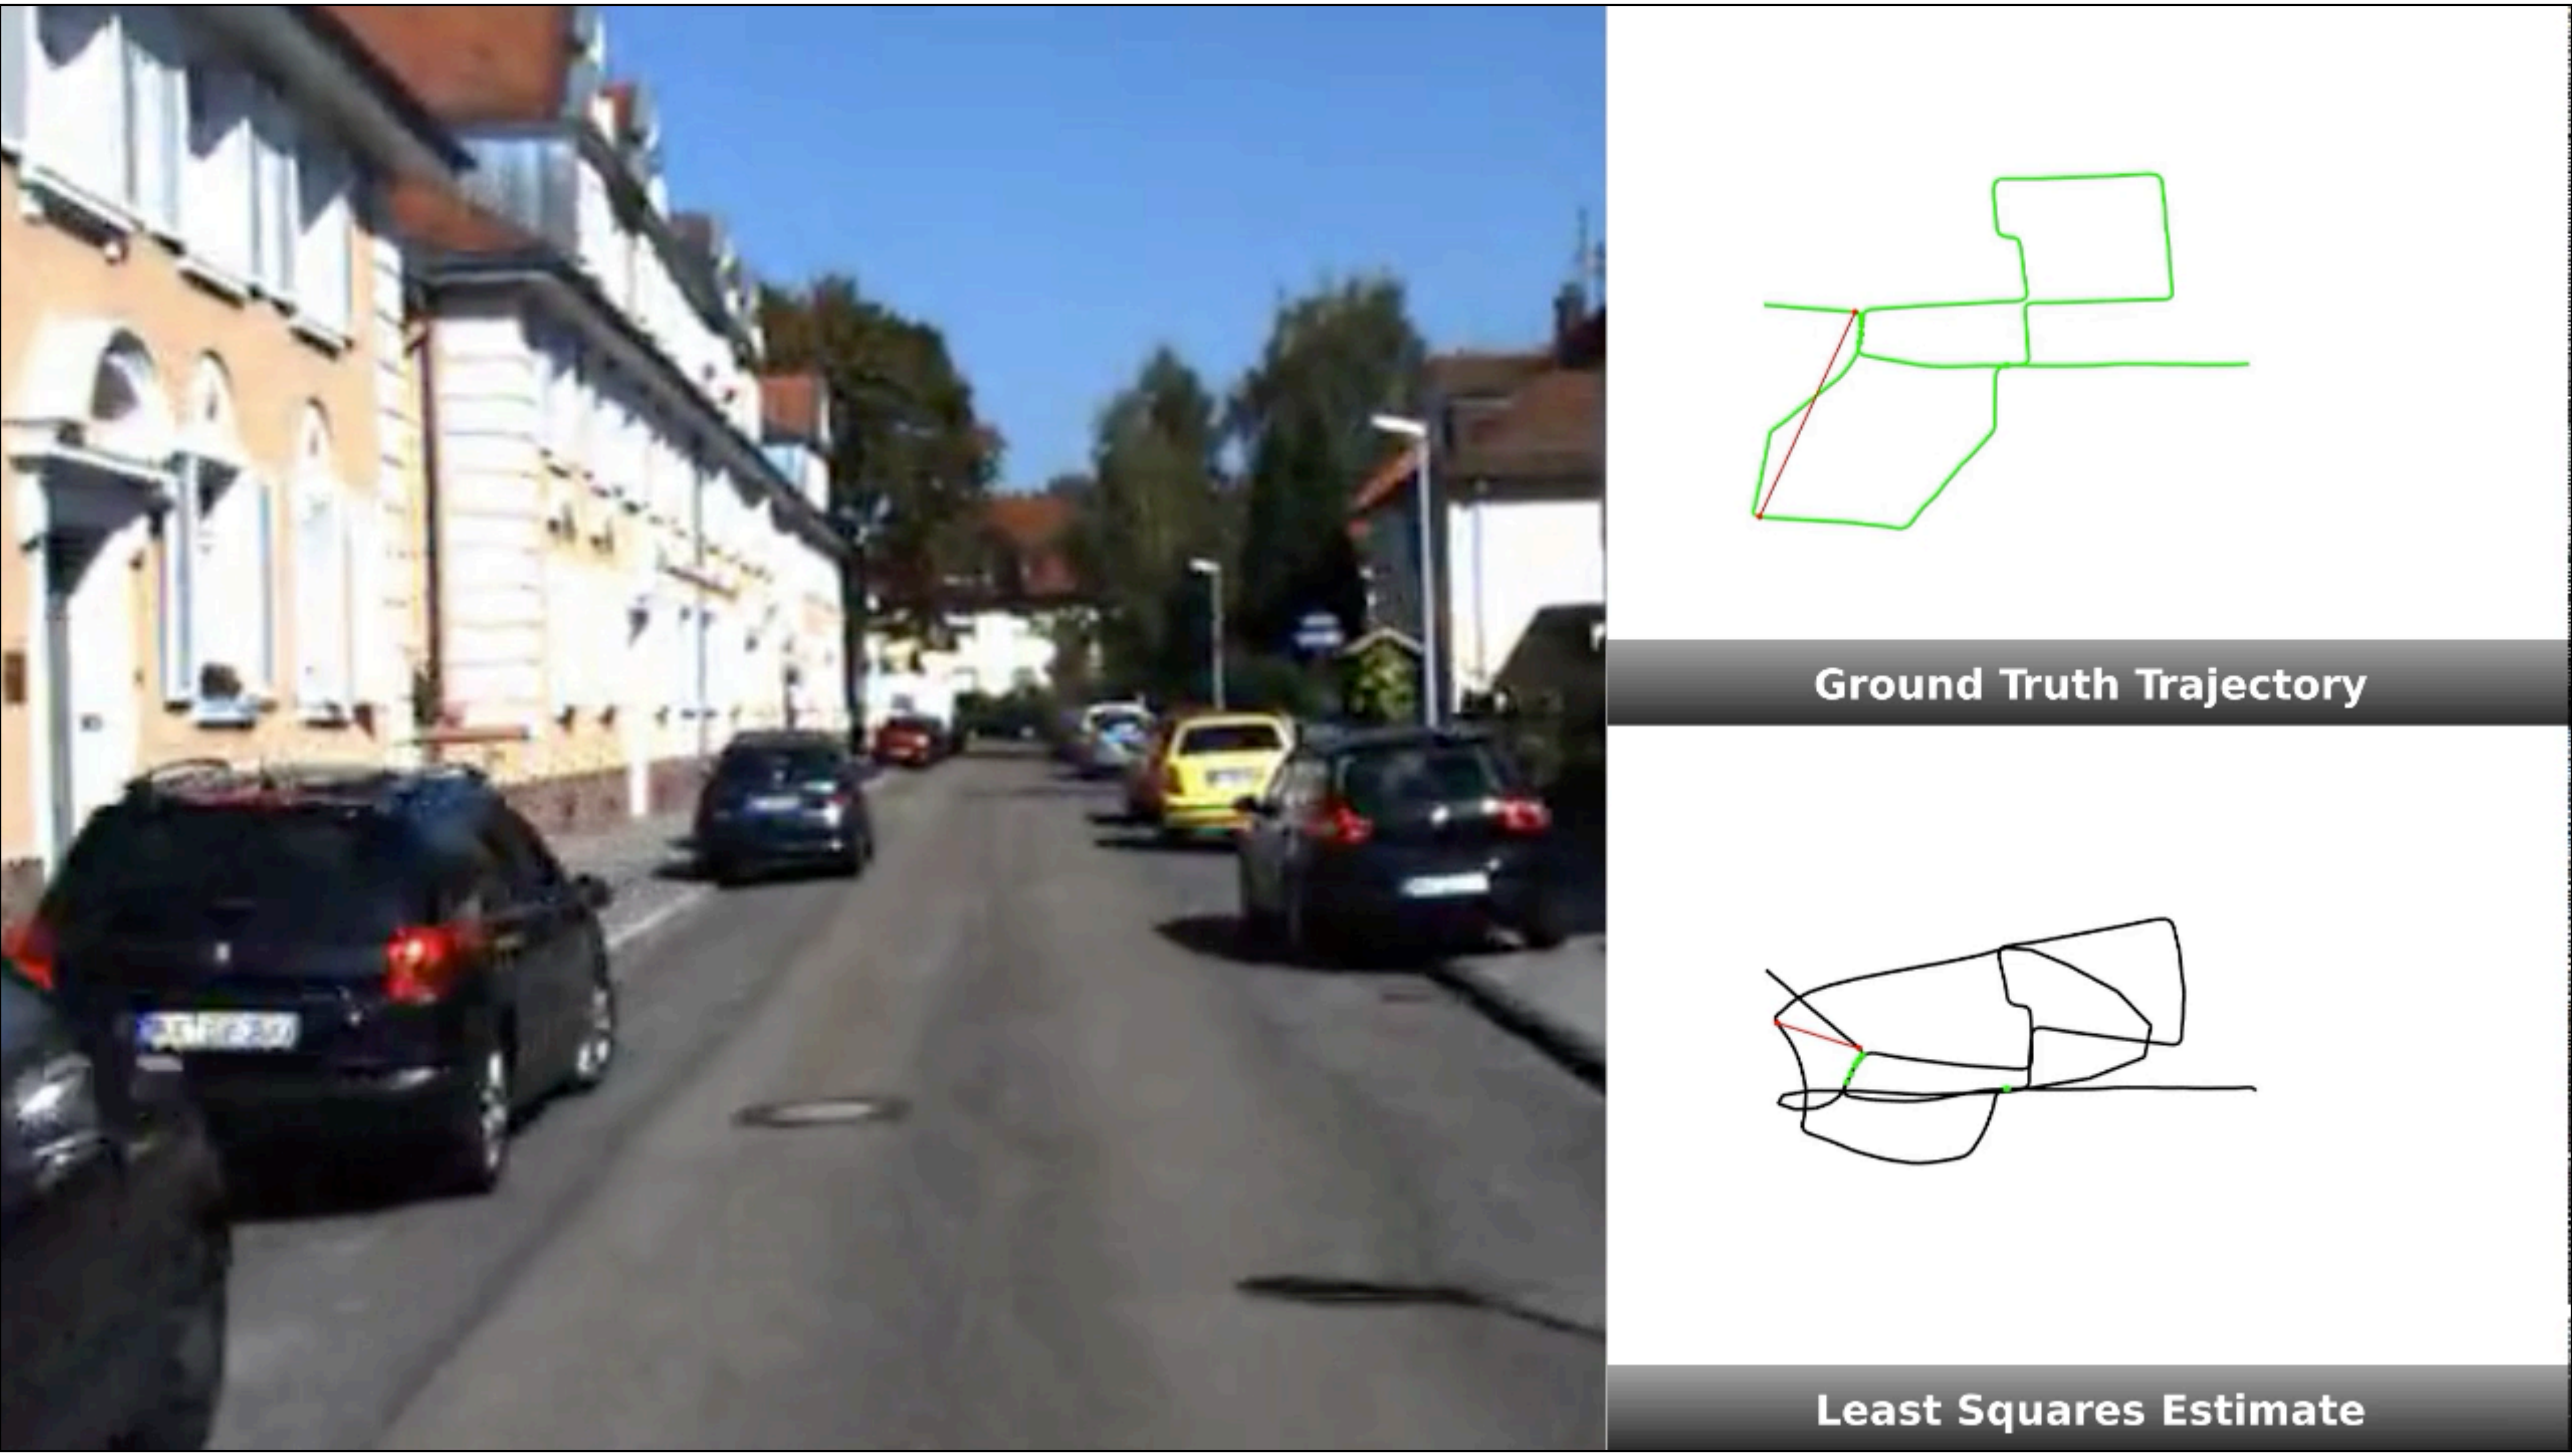

Outliers in visual SLAM

Outliers in pose graph optimization

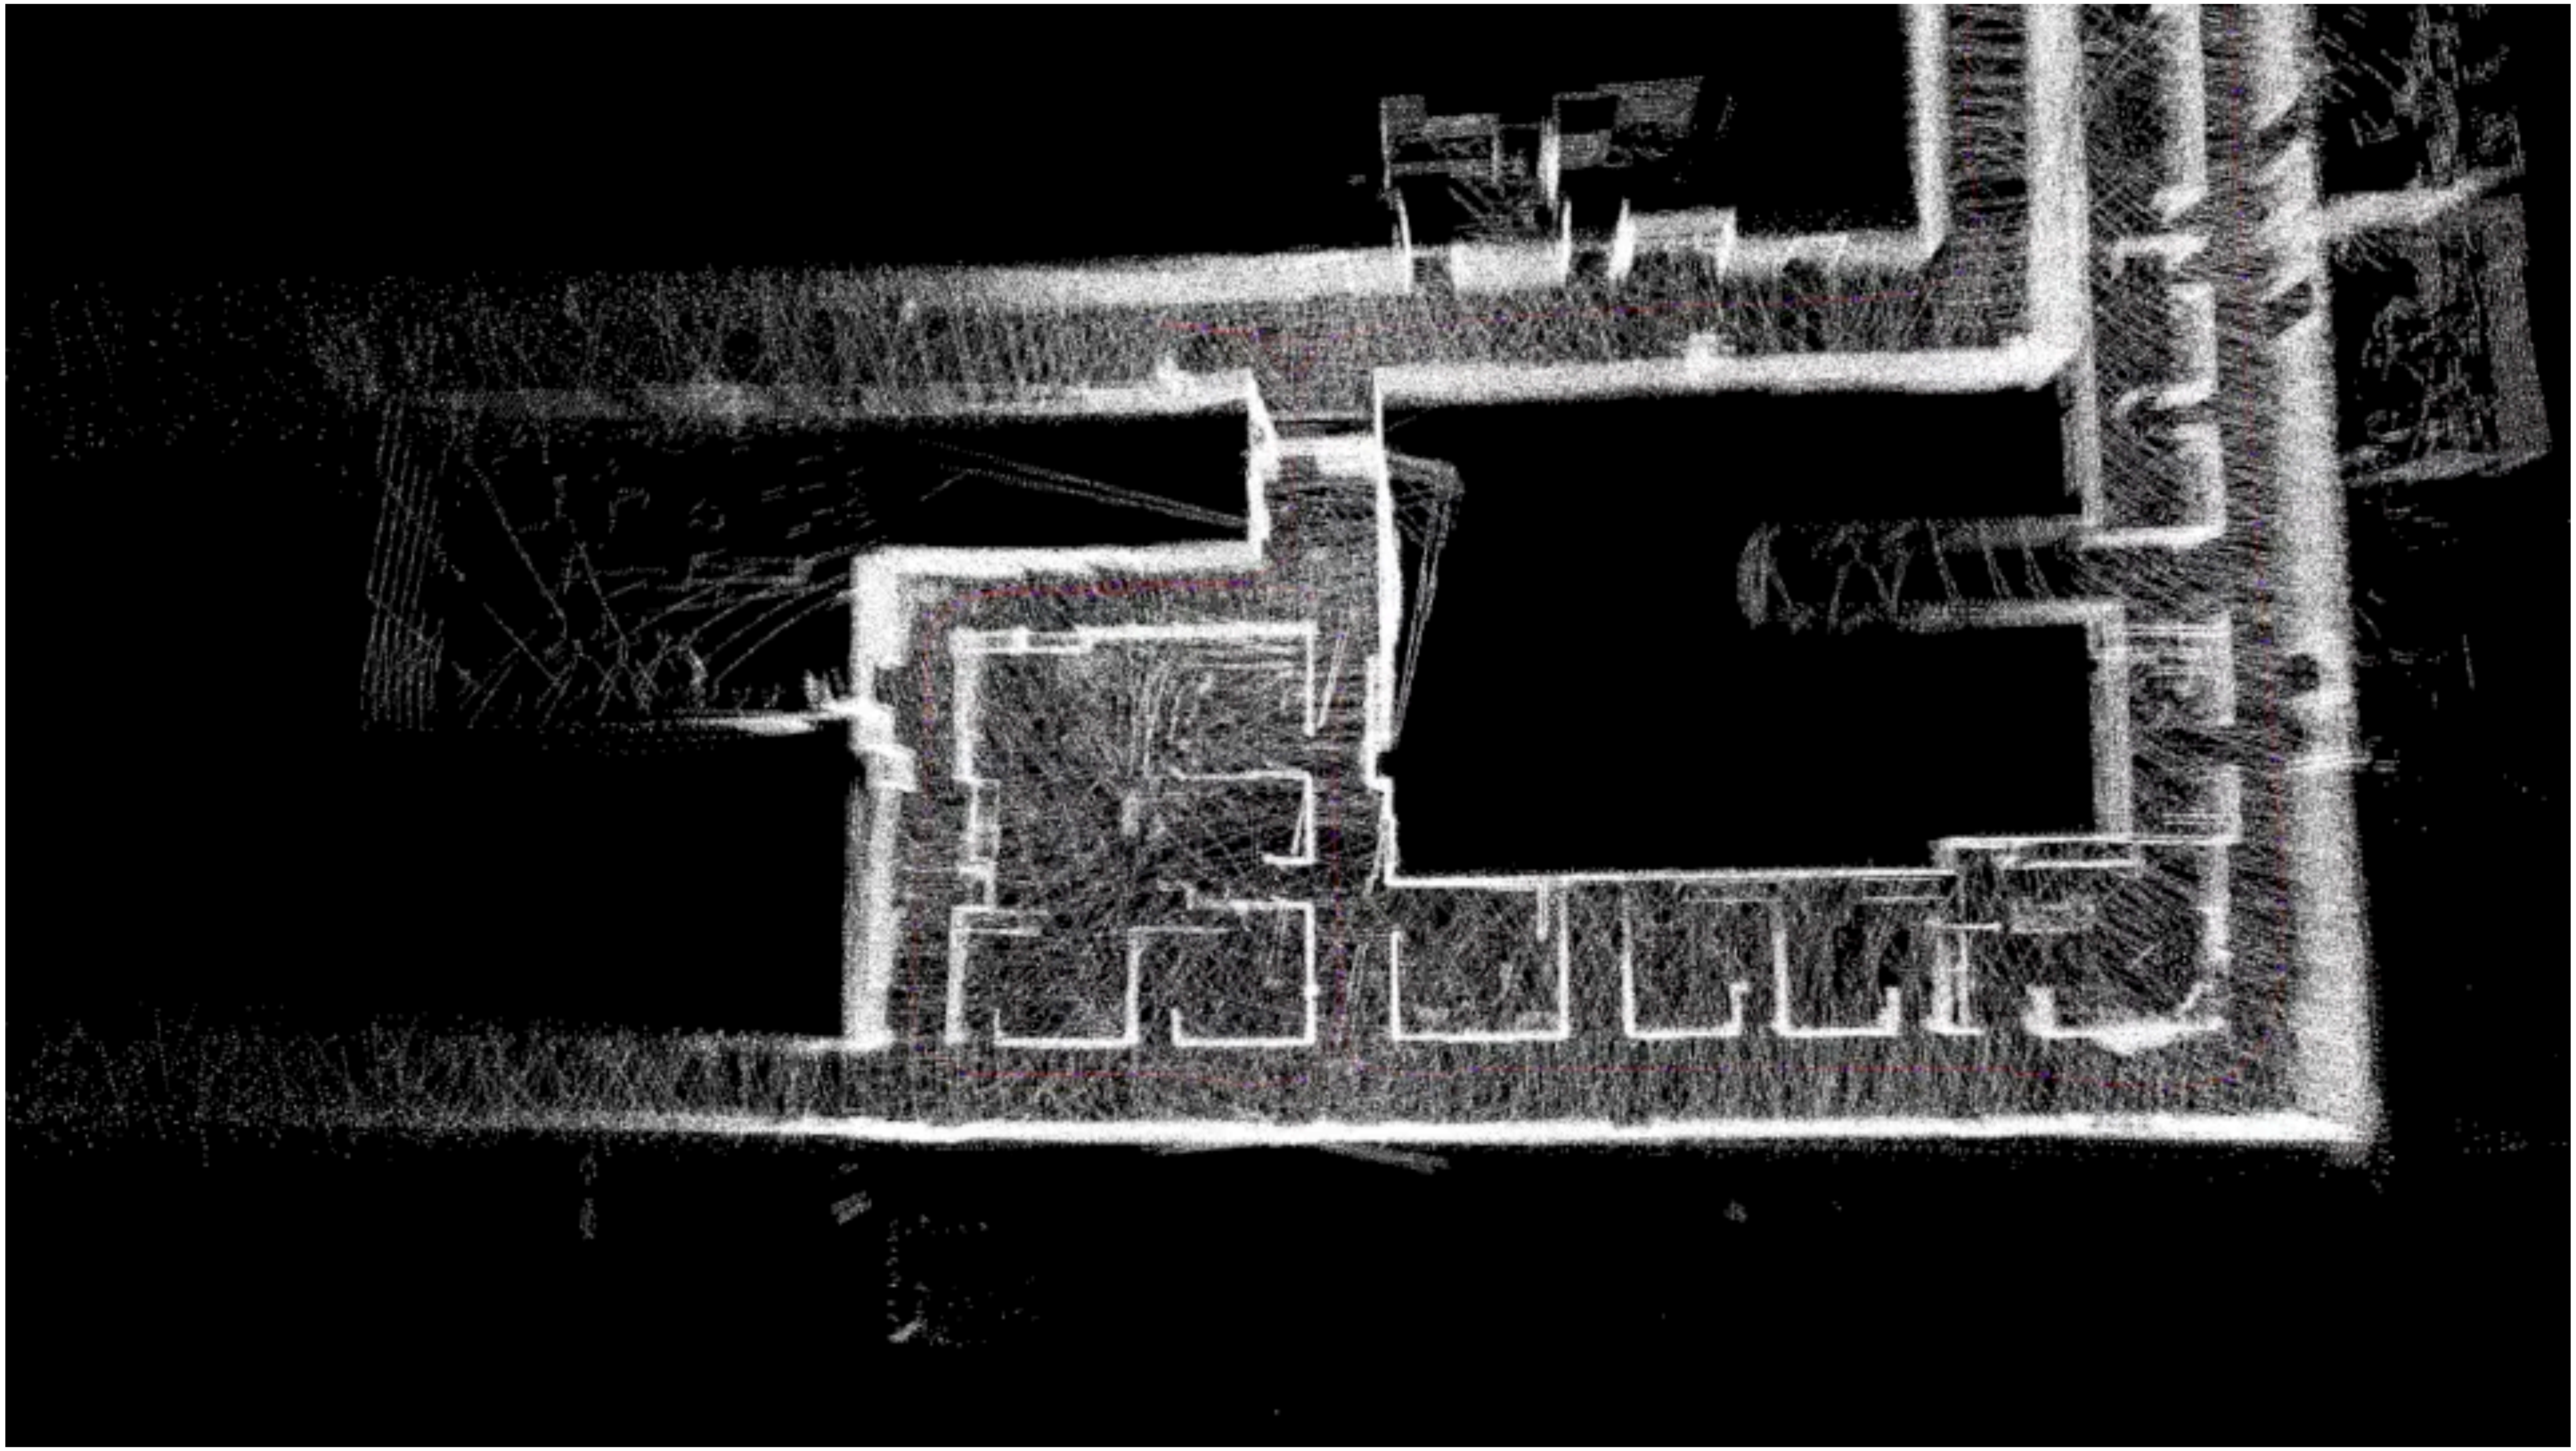

Outliers in lidar-based SLAM

Outliers in landmark-based SLAM

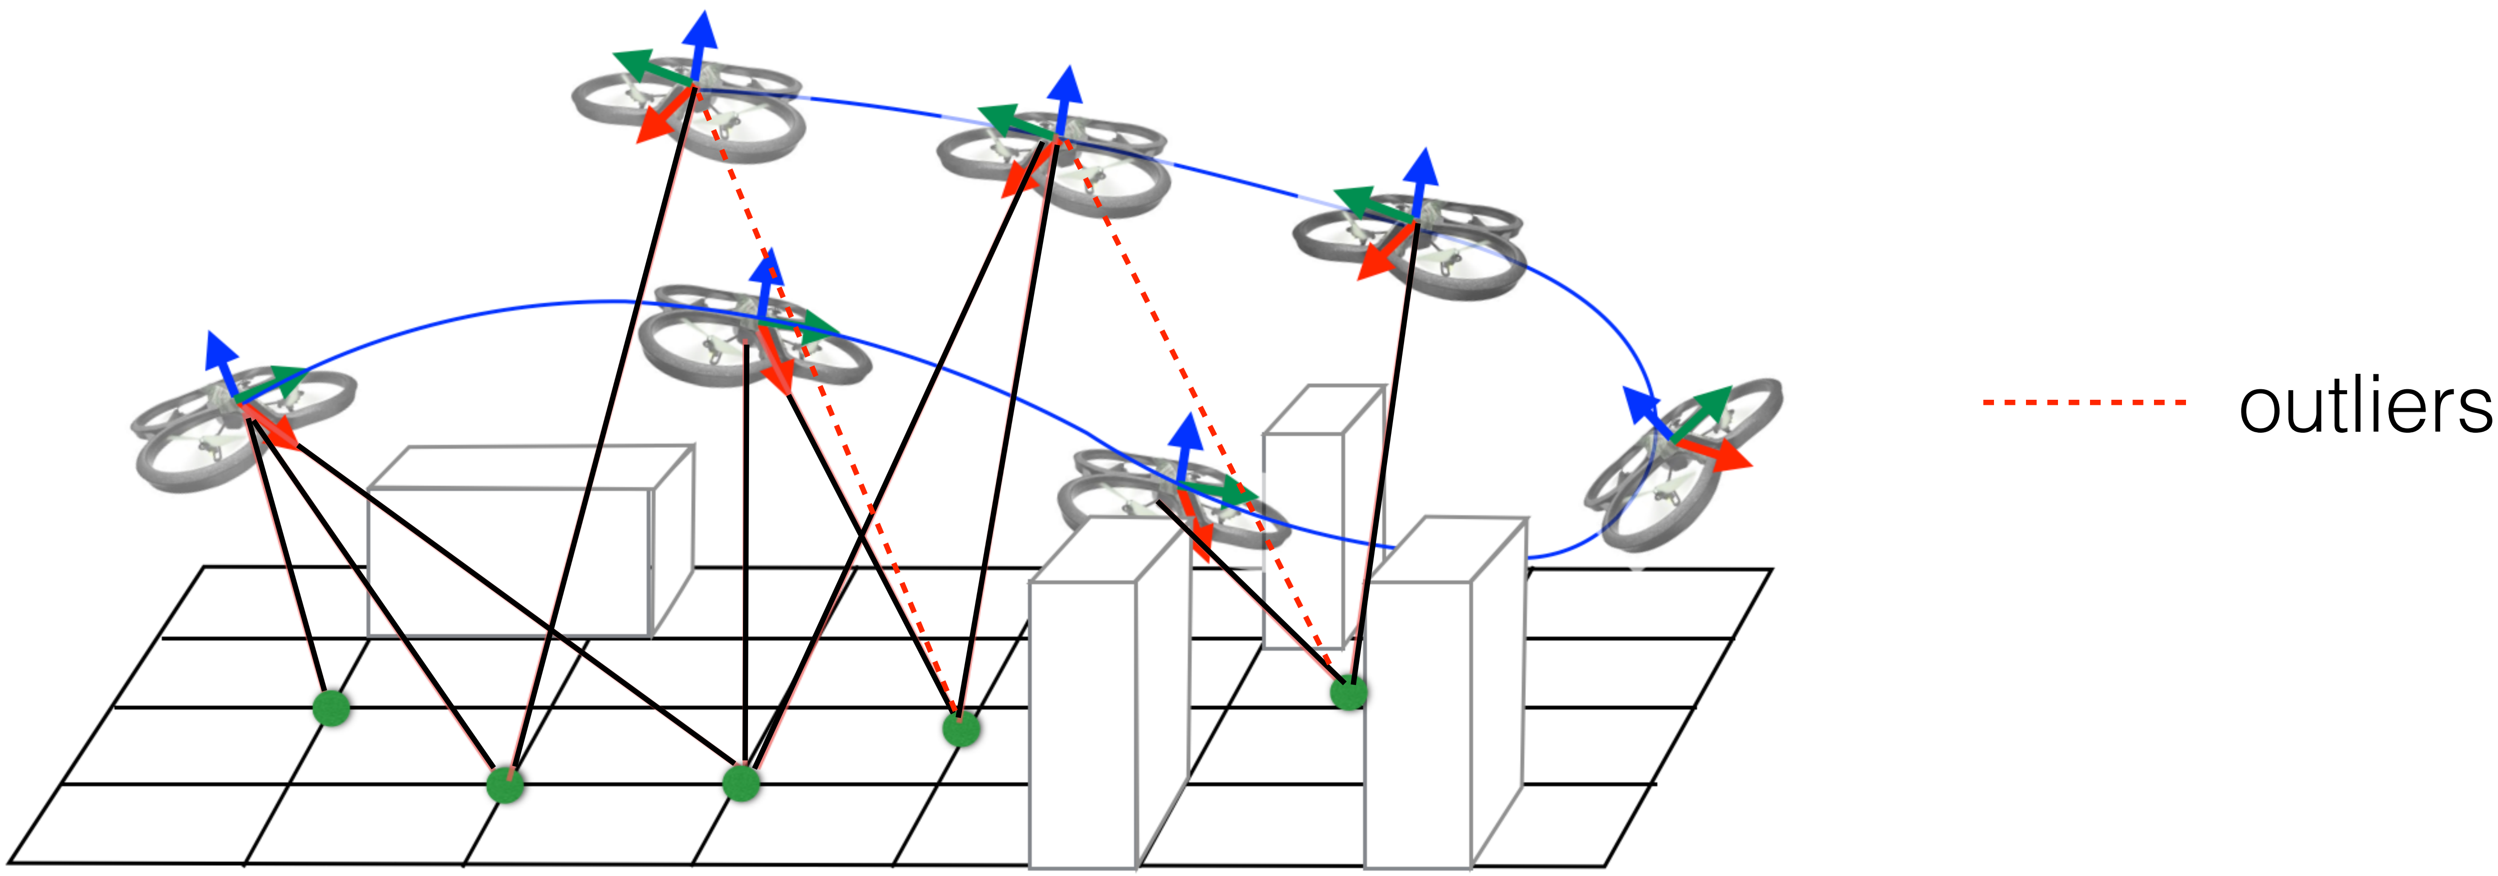

Data association: association of a measurement with the variables being measured:

$$\bar{\boldsymbol{y}}_{k,t} = h_i(\boldsymbol{T}_t^{\mathrm{w}}, \boldsymbol{l}_k^{\mathrm{w}}) + \boldsymbol{\epsilon}_l$$

Measurement robot pose landmark

Outliers are typically the result of incorrect data association

So far in VNAV

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

Example: Point Cloud Registration

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

Pose

Point clouds (data)

Correspondences between p_i, p'_j

Example: Pose Graph Optimization

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

Poses

Relative pose measurement

Odometry and Loop closures between T_i, T_j

Why do least squares fail with outliers?

Least squares problems penalize large residuals a **LOT** (due to square)

Least squares find an estimate x to minimize large residuals

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

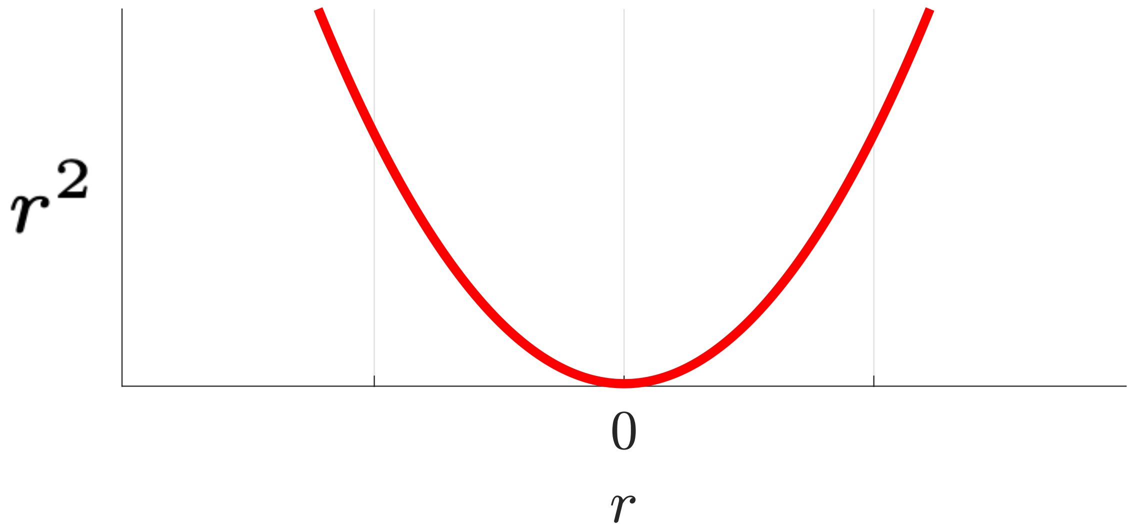

Example:

- $x_{true} = 0$
- Measurement 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: $y_1 = y_2 = 0,\ y_3 = 10$

Least squares opt. solution is $x = 3.33 \neq x_{true} = 0!$

Robust Estimation: M-Estimation

Use robust loss function that down-weighs the influence of outliers

$$\min_{\boldsymbol{x} \in \mathbb{X}} \sum_{i \in \mathcal{M}} \rho(r(\boldsymbol{x}, \boldsymbol{y}_i))$$

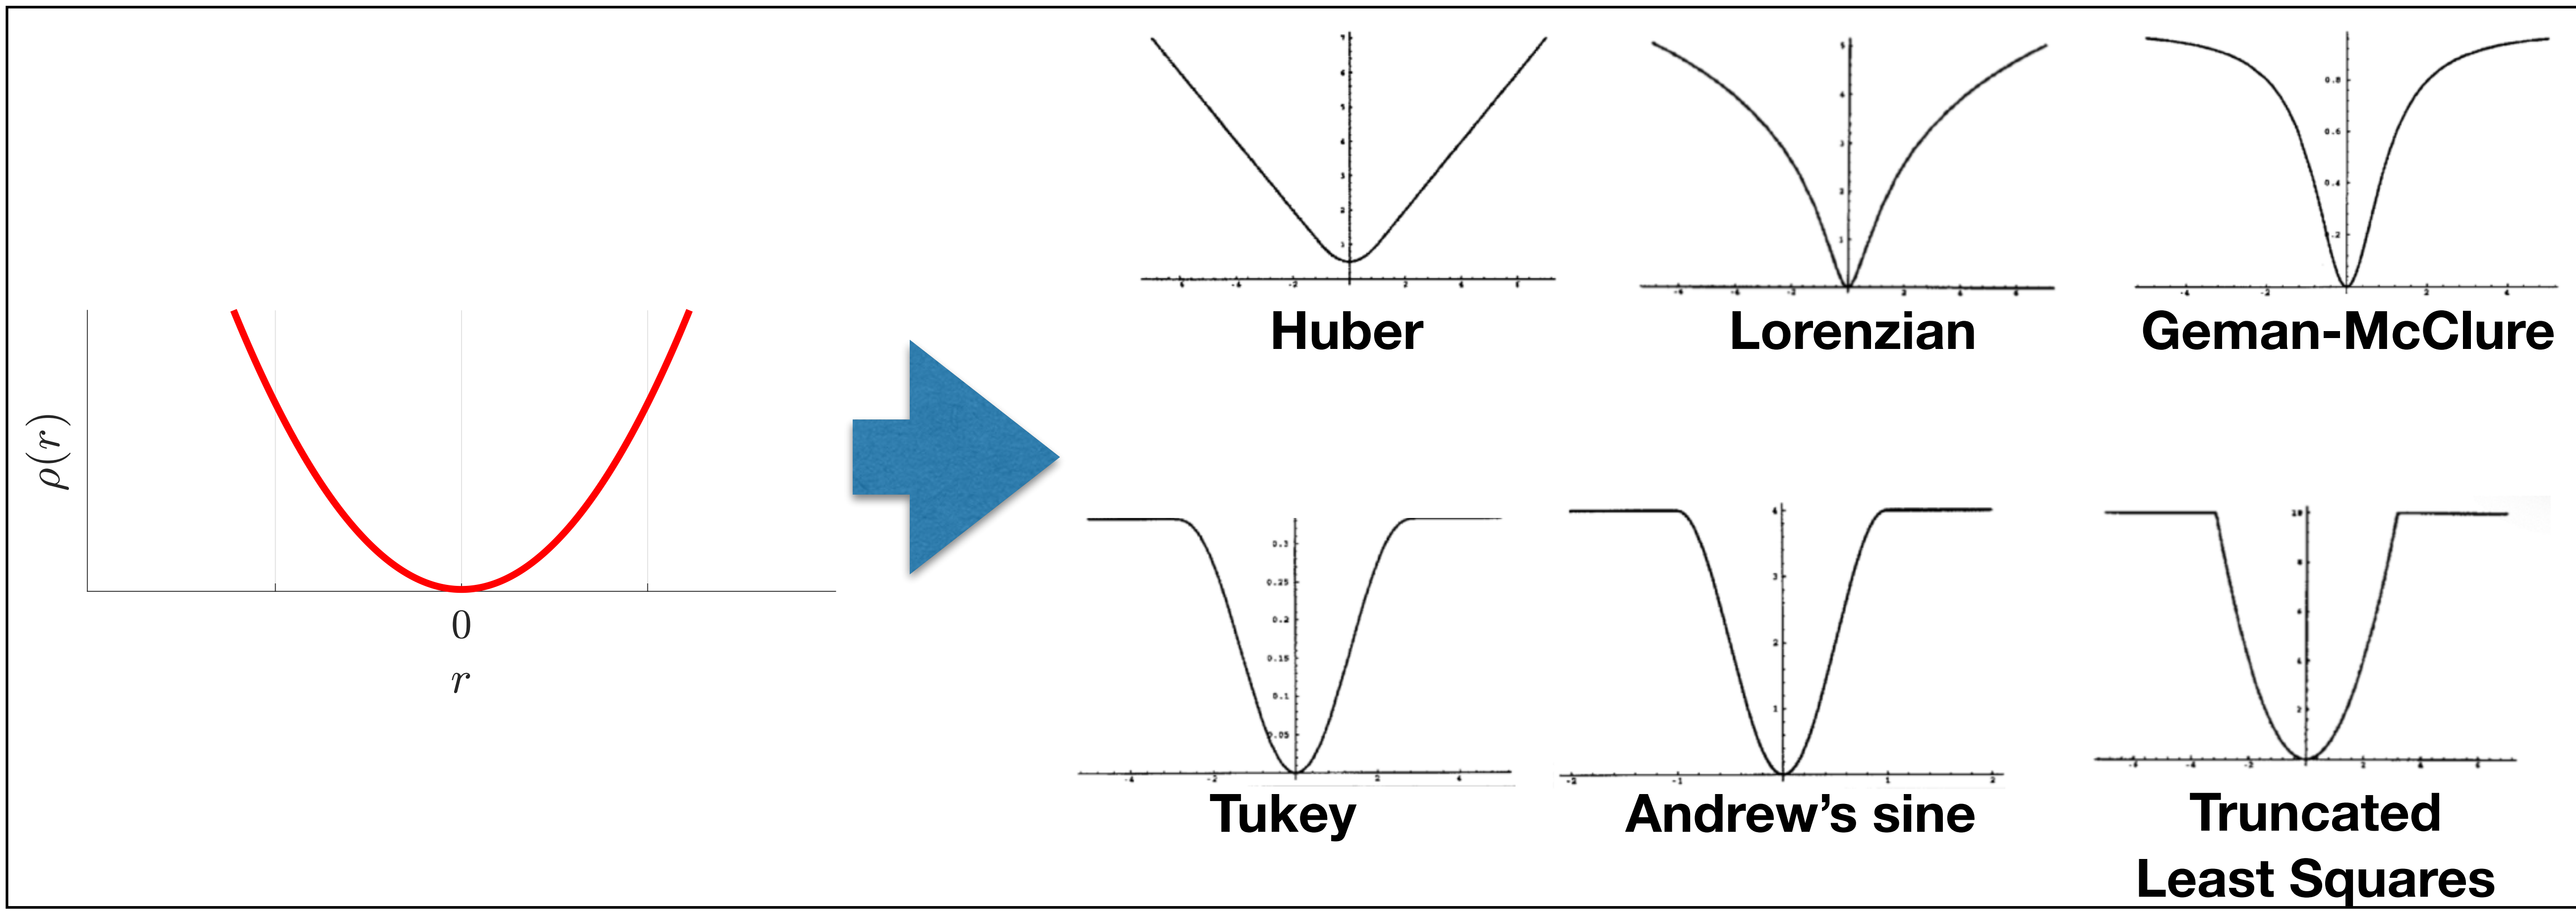

… among many others

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

"Maximum likelihood-type" estimators

Robust Estimation: Maximum Consensus

$$\min_{\boldsymbol{x} \in \mathbb{X}} \sum_{i \in \mathcal{M}} \rho(r(\boldsymbol{x}, \boldsymbol{y}_i))$$

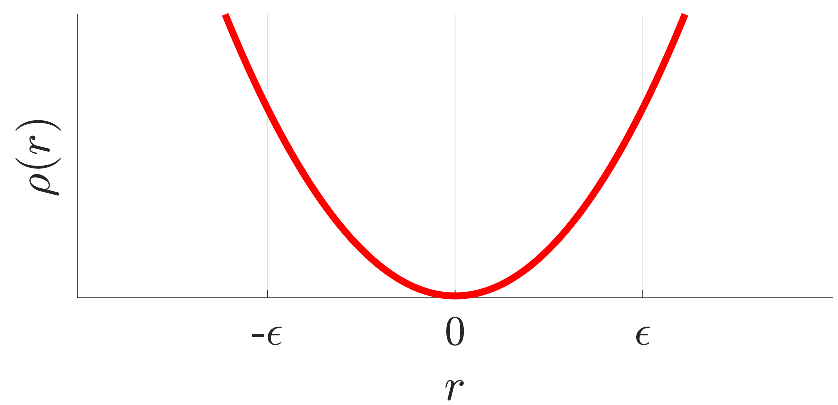

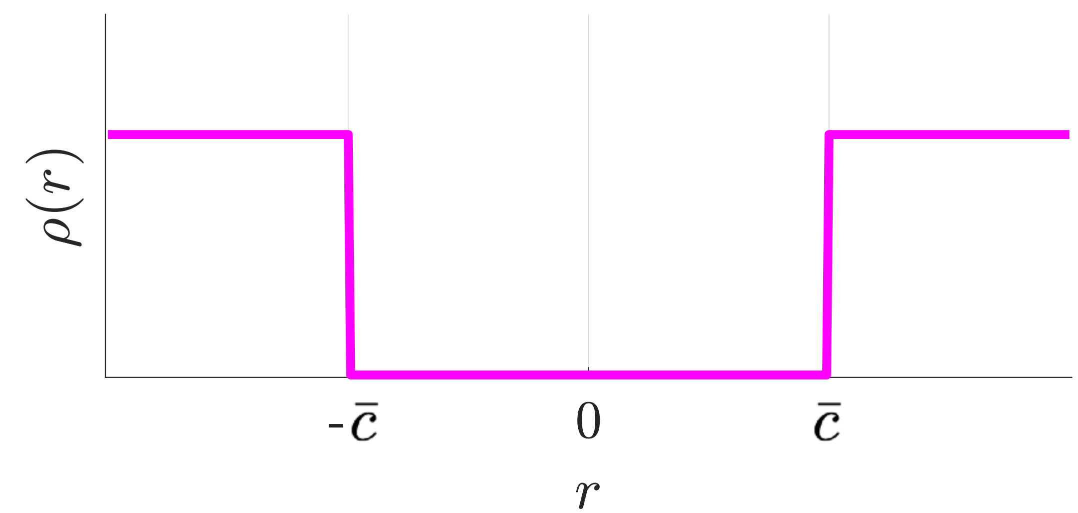

Function counts number of outliers

Minimize nr. of outliers = maximize nr. of inliers

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

Robust Estimation: Hardness

Inapproximability result: solving robust estimation problems to optimality is intractable for common choices of loss functions

Theorem (Chin et al. ‘18, Antonante et al. ‘19)

- Let $\mathcal{O}^\star$ be the true number of outliers.
- Let $\epsilon = p_1\left(\left|\mathcal{M}\right|\right)$, where p_1 is a polynomial in number of measurements.
- Let $p_2(\left|\mathcal{M}\right|)$ another polynomial.

Then:

No quasi-polynomial algorithm can reject less than $p_2(\left|\mathcal{M}\right|)\left|\mathcal{O}^\star\right|$ measurements.

slower than polynomial
faster than exponential

no constant approximation factor

Theorem applies to both:

Truncated least squares: $\min_{x \in \mathcal{X}} \sum_{i \in \mathcal{M}} \rho\left(r\left(x, y_i\right)\right)$ with $\rho(r) = \begin{cases} r^2 & if \;\; r \in [0,\bar{c}] \\ \bar{c}^2 & if \;\; r \in [\bar{c}, +\infty] \end{cases}$

Maximum consensus: $\min_{\mathcal{O} \subseteq \mathcal{M}} |\mathcal{O}| \quad s.t. \quad r\left(x, y_i\right) \leq \bar{c}, \;\; \forall\, i \in \mathcal{M} \backslash \mathcal{O}$

Today and Next Lecture

- **Robust estimation**:
  - Motivations: outliers, data association
  - Formulations: M-estimation & Maximum Consensus
- **Solvers for robust estimation**:
  - (RANSAC)
  - Iteratively Reweighted Least Squares (IRLS)
  - Max-mixture
  - Switchable constraints
  - Graduated non-convexity
  - Others: BnB, SDP relaxations, graph-theoretic pruning

RANSAC

$$\min_{\boldsymbol{x}\in\mathbb{X}} \sum_{i\in\mathcal{M}} \rho(r(\boldsymbol{x},\boldsymbol{y}_i))$$

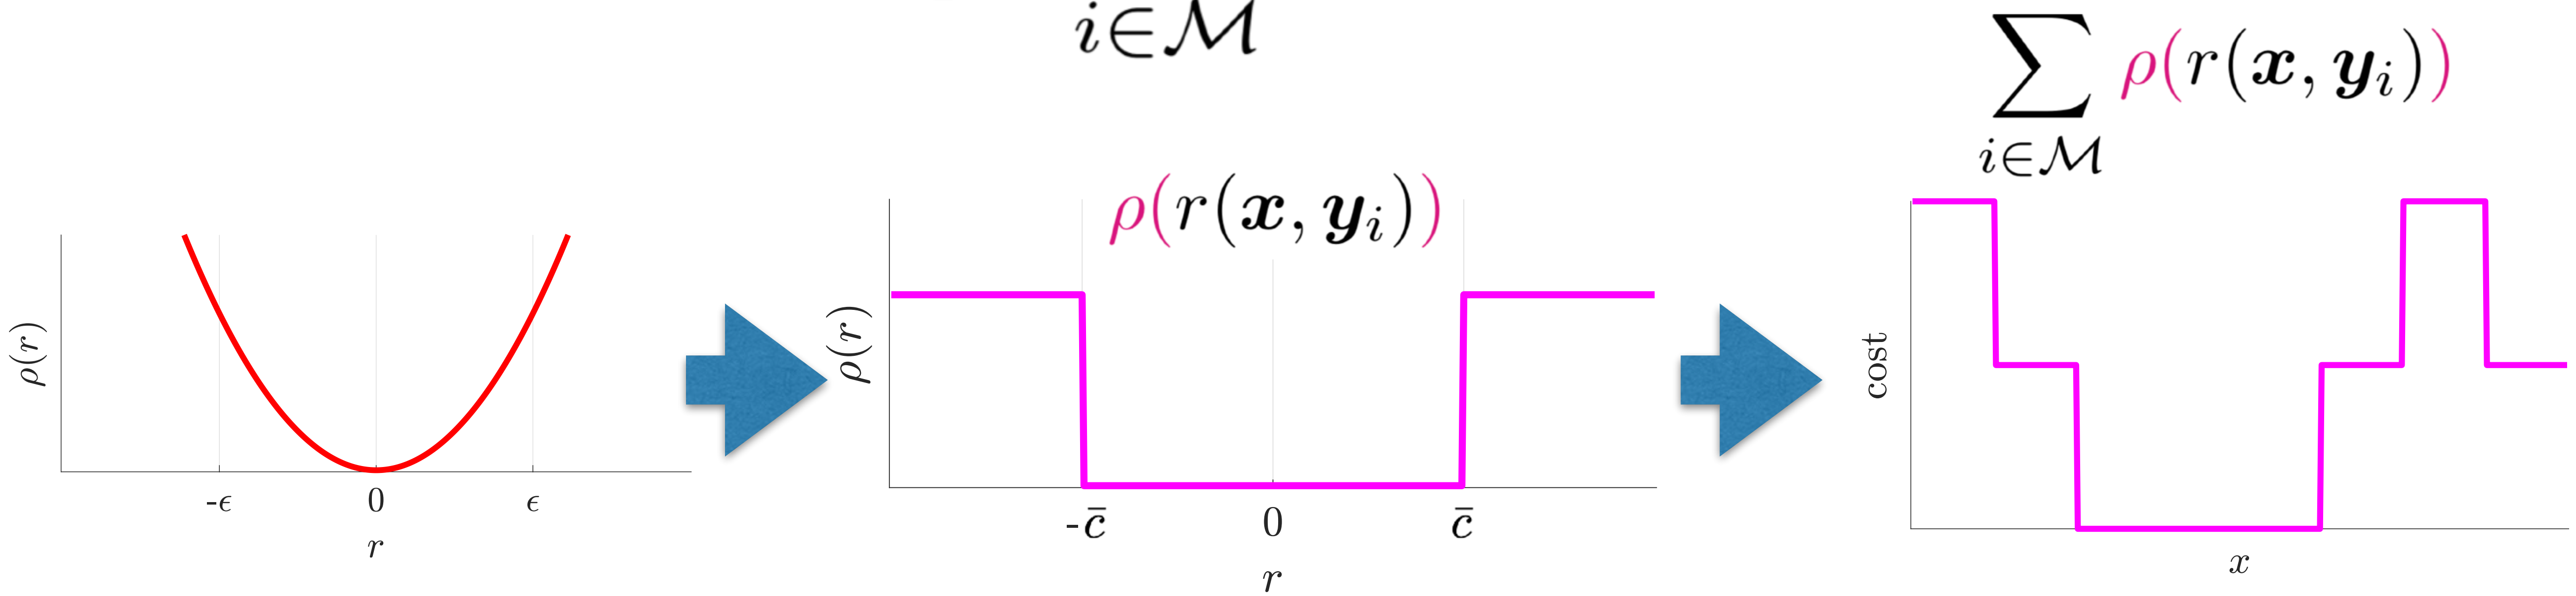

$$\min_{\mathcal{O}\subseteq\mathcal{M}} |\mathcal{O}| \quad s.t. \quad r(x, y_i) \leq \bar{c}, \quad \forall\, i \in \mathcal{M}\backslash\mathcal{O}$$

RANSAC samples small set of measurements to build an estimate of x and hope that it minimizes the cost

Nr. Iterations in RANSAC increases exponentially in the percentage of outliers and the number of points used by the minimal solver

Iteratively Reweighted Least Squares (IRLS)

GTSAM Robust Noise Model

Fan Jiang[†], Yetong Zhang[†]

February 2020

1 Introduction

In gtsam, we solve the problem of reducing the error of a factor graph. For each factor i, we have observation function h_i, and the measurement value z_i. Then the measurement error vector e_i is defined as

$$e_i = h_i(x_i) - z_i$$

Then, our objective of reducing the error of the factor graph becomes

$$\min_x err_{graph}(x) = \min_x \sum_i err_i(e_i)$$

Normally, we are concerned with the least square problem, where the error function for each factor is defined as

$$err(e) = \frac{1}{2}\|e\|^2_\Sigma$$

where Σ si the covaraince matrix associated with the measurement. Then, our objective becomes:

$$\min_x \sum_i \frac{1}{2}\|h_i(x_i) - z_i\|^2_{\Sigma_i}$$

Iteratively Reweighted Least Squares (IRLS)

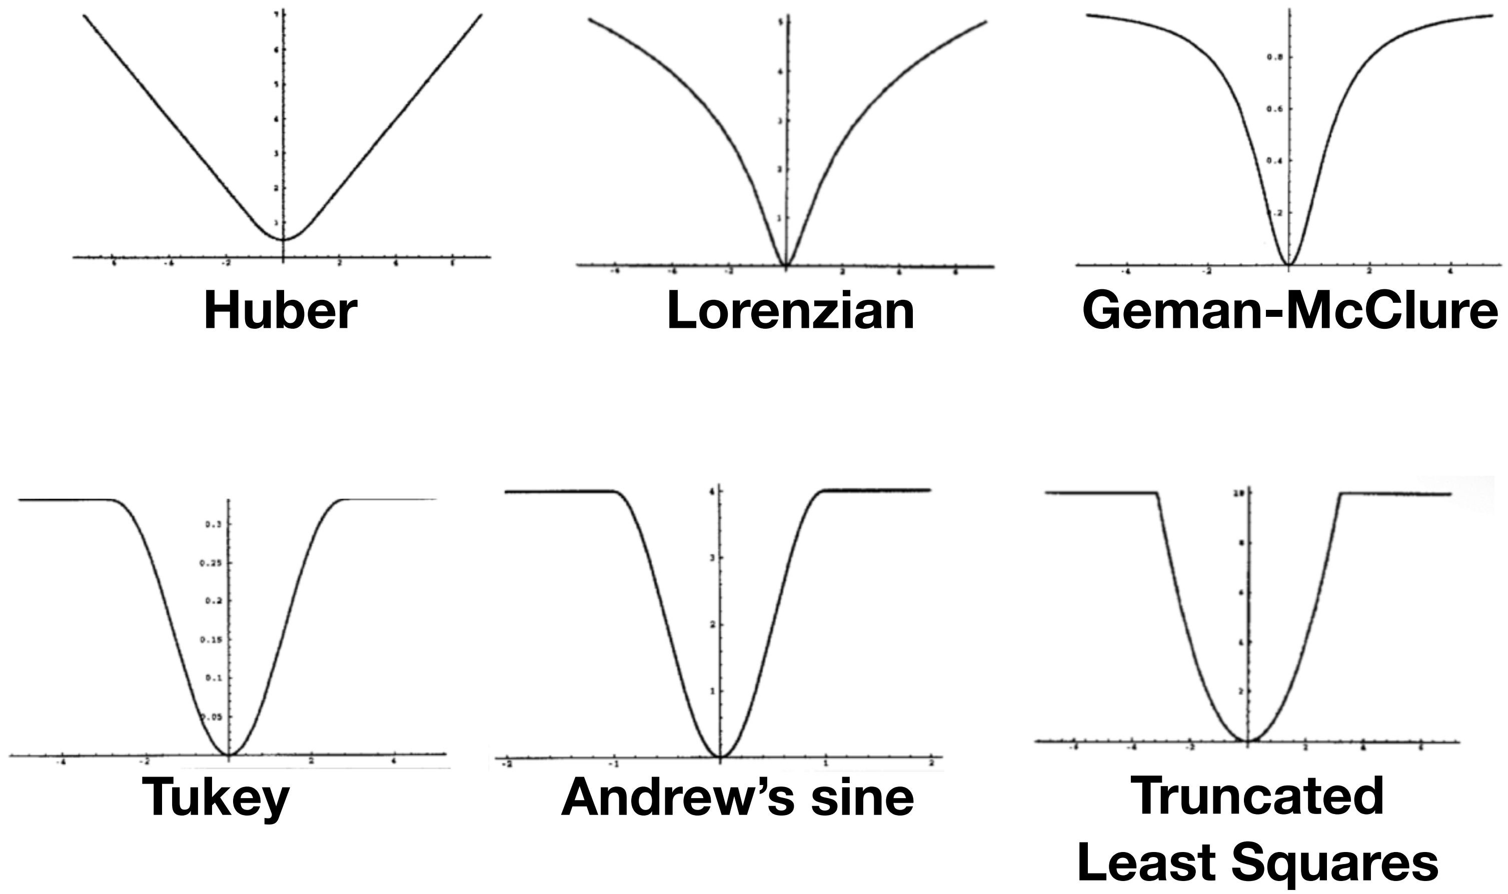

Start from initial guess and at each iteration convert the problem into a weighted nonlinear least squares:

$$\min_{\boldsymbol{x} \in \mathbb{X}} \sum_{i \in \mathcal{M}} \rho(r(\boldsymbol{x}, \boldsymbol{y}_i)) \quad \Rightarrow \quad \min_{\boldsymbol{x} \in \mathbb{X}} \sum_{i \in \mathcal{M}} w_i r_i^2(\boldsymbol{x}, \boldsymbol{y}_i)$$

Iterative Reweighted Least Squares (IRLS)

Loss functions

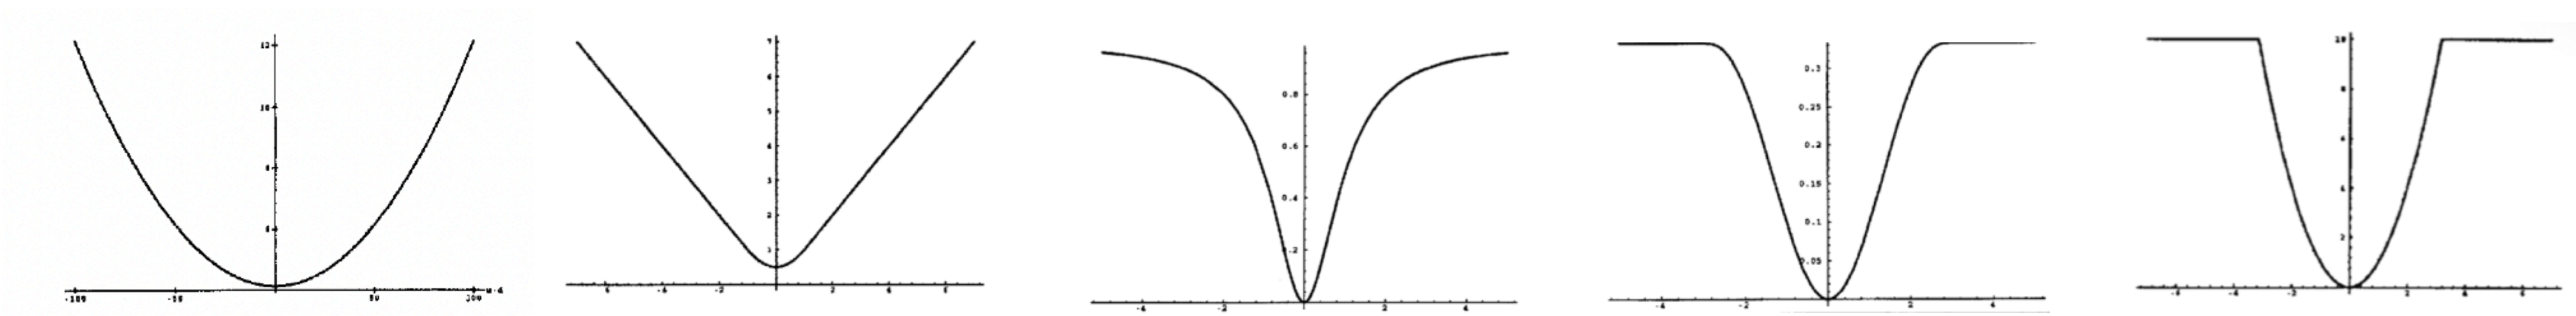

Quadratic | Huber | Geman-McClure | Tukey | Truncated Least Squares

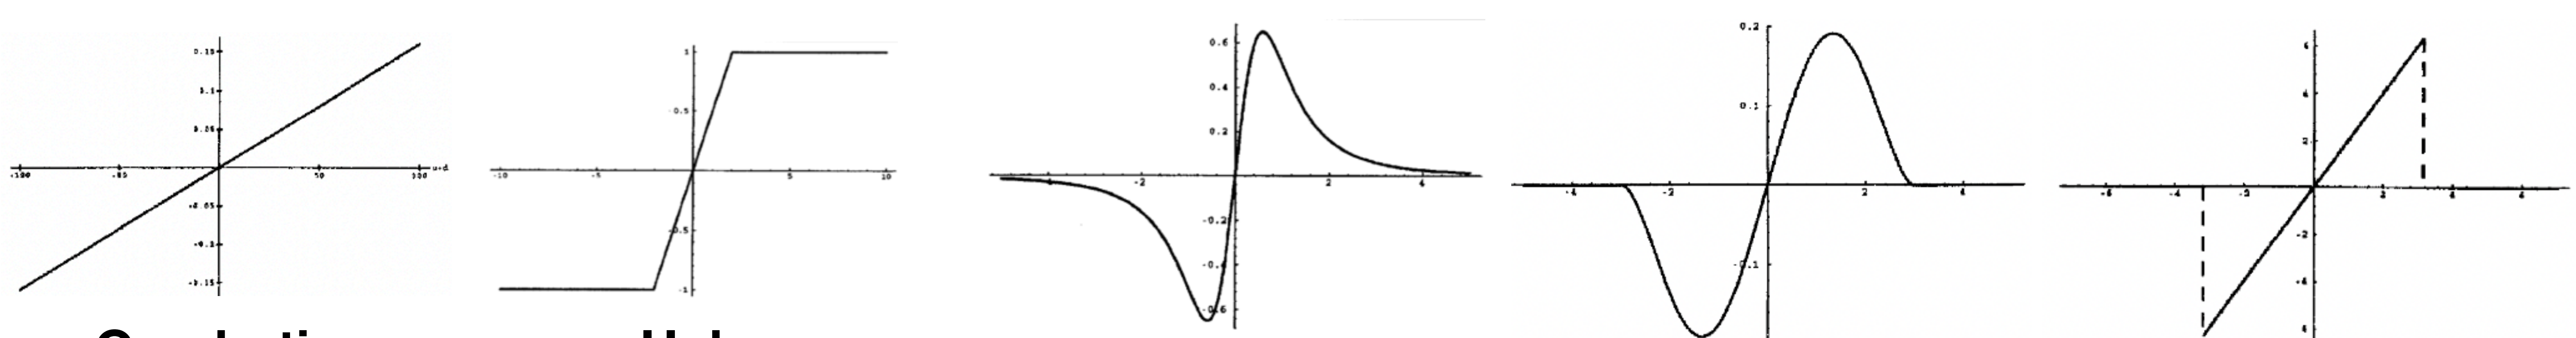

Quadratic | Huber | Geman-McClure | Tukey | Truncated Least Squares

~Influence functions

[Black, Rangarajan, "On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision, IJCV'96.]

Robust Estimation with Max-Mixture

Inference on networks of mixtures for robust robot mapping

Edwin Olson
Computer Science and Engineering,
University of Michigan,
2260 Hayward Street,
Ann Arbor, Michigan
Email: ebolson@umich.edu

Pratik Agarwal
Computer Science and Engineering,
University of Michigan,
2260 Hayward Street,
Ann Arbor, Michigan
Email: pratikag@umich.edu

***Abstract*— The central challenge in robotic mapping is obtaining reliable data associations (or "loop closures"): state-of-the-art inference algorithms can fail catastrophically if even one erroneous loop closure is incorporated into the map. Consequently, much work has been done to push error rates closer to zero. However, a long-lived or multi-robot system will still encounter errors, leading to system failure.**

We propose a fundamentally different approach: allow richer error models that allow the probability of a failure to be explicitly modeled. In other words, we optimize the map while simultaneously determining which loop closures are correct from within a single integrated Bayesian framework. Unlike earlier multiple-hypothesis approaches, our approach avoids exponential memory complexity and is fast enough for real-time performance.

We show that the proposed method not only allows loop closing errors to be automatically identified, but also that in extreme cases, the "front-end" loop-validation systems can be unnecessary. We demonstrate our system both on standard benchmarks and on the real-world datasets that motivated this work.

I. Introduction

Robot mapping problems are often formulated as an inference problem on a factor graph: variable nodes (representing the location of robots or other landmarks in the environment) are related through factor nodes, which encode geometric relationships between those nodes. Recent Simultaneous Localization and Mapping (SLAM) algorithms can rapidly find maximum likelihood solutions for maps, exploiting both fundamental improvements in the understanding of

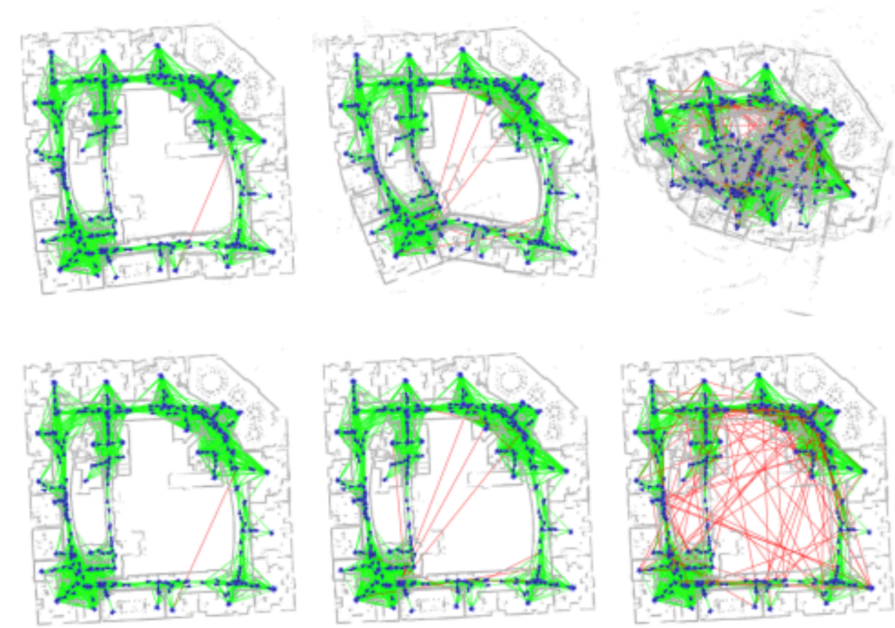

Fig. 1. Recovering a map in the presence of erroneous loop closures. We evaluated the robustness of our method by adding erroneous loop closures to the Intel data set. The top row reflects the posterior map as computed by a state-of-the-art sparse Cholesky factorization method with 1, 10, and 100 bad loop closures. The bottom row shows the posterior map for the same data set using our proposed max mixture method. While earlier methods produce maps with increasing global map deformation, our proposed method is essentially unaffected by the presence of the incorrect loop closures.

tifying and validating loop closures and constructing a factor graph; the back-end then performs inference (often maximum likelihood) on this factor graph. In most of the literature, it is assumed that the loop closures found by the front-end have

Robust Estimation with Max-Mixture

Mathematical Model

- We can express a multi-modal belief by a sum of Gaussians

$$p(\mathbf{z} \mid \mathbf{x}) = \eta \exp(-\frac{1}{2}\mathbf{e}_{ij}^T \mathbf{\Omega}_{ij} \mathbf{e}_{ij})$$

$\Downarrow$

$$p(\mathbf{z} \mid \mathbf{x}) = \sum_k w_k \eta_k \exp(-\frac{1}{2}\mathbf{e}_{ij_k}^T \mathbf{\Omega}_{ij_k} \mathbf{e}_{ij_k})$$

Sum of Gaussians with k modes

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Robust Estimation with Max-Mixture

Problem

- During error minimization, we consider the negative log likelihood

$$-\log p(\mathbf{z} \mid \mathbf{x}) = \frac{1}{2}\mathbf{e}_{ij}^T \mathbf{\Omega}_{ij} \mathbf{e}_{ij} - \log \eta$$

$\Downarrow$

$$-\log p(\mathbf{z} \mid \mathbf{x}) = -\log \sum_k w_k \eta_k \exp(-\frac{1}{2}\mathbf{e}_{ij_k}^T \mathbf{\Omega}_{ij_k} \mathbf{e}_{ij_k})$$

The log cannot be moved inside the sum!

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Robust Estimation with Max-Mixture

Max-Mixture Approximation

- Instead of computing the sum of Gaussians at x, compute the maximum of the Gaussians

$$p(\mathbf{z} \mid \mathbf{x}) = \sum_k w_k \eta_k \exp(-\frac{1}{2}\mathbf{e}_{ij_k}^T \mathbf{\Omega}_{ij_k} \mathbf{e}_{ij_k})$$

$$\simeq \max_k w_k \eta_k \exp(-\frac{1}{2}\mathbf{e}_{ij_k}^T \mathbf{\Omega}_{ij_k} \mathbf{e}_{ij_k})$$

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Robust Estimation with Max-Mixture

Max-Mixture Approximation

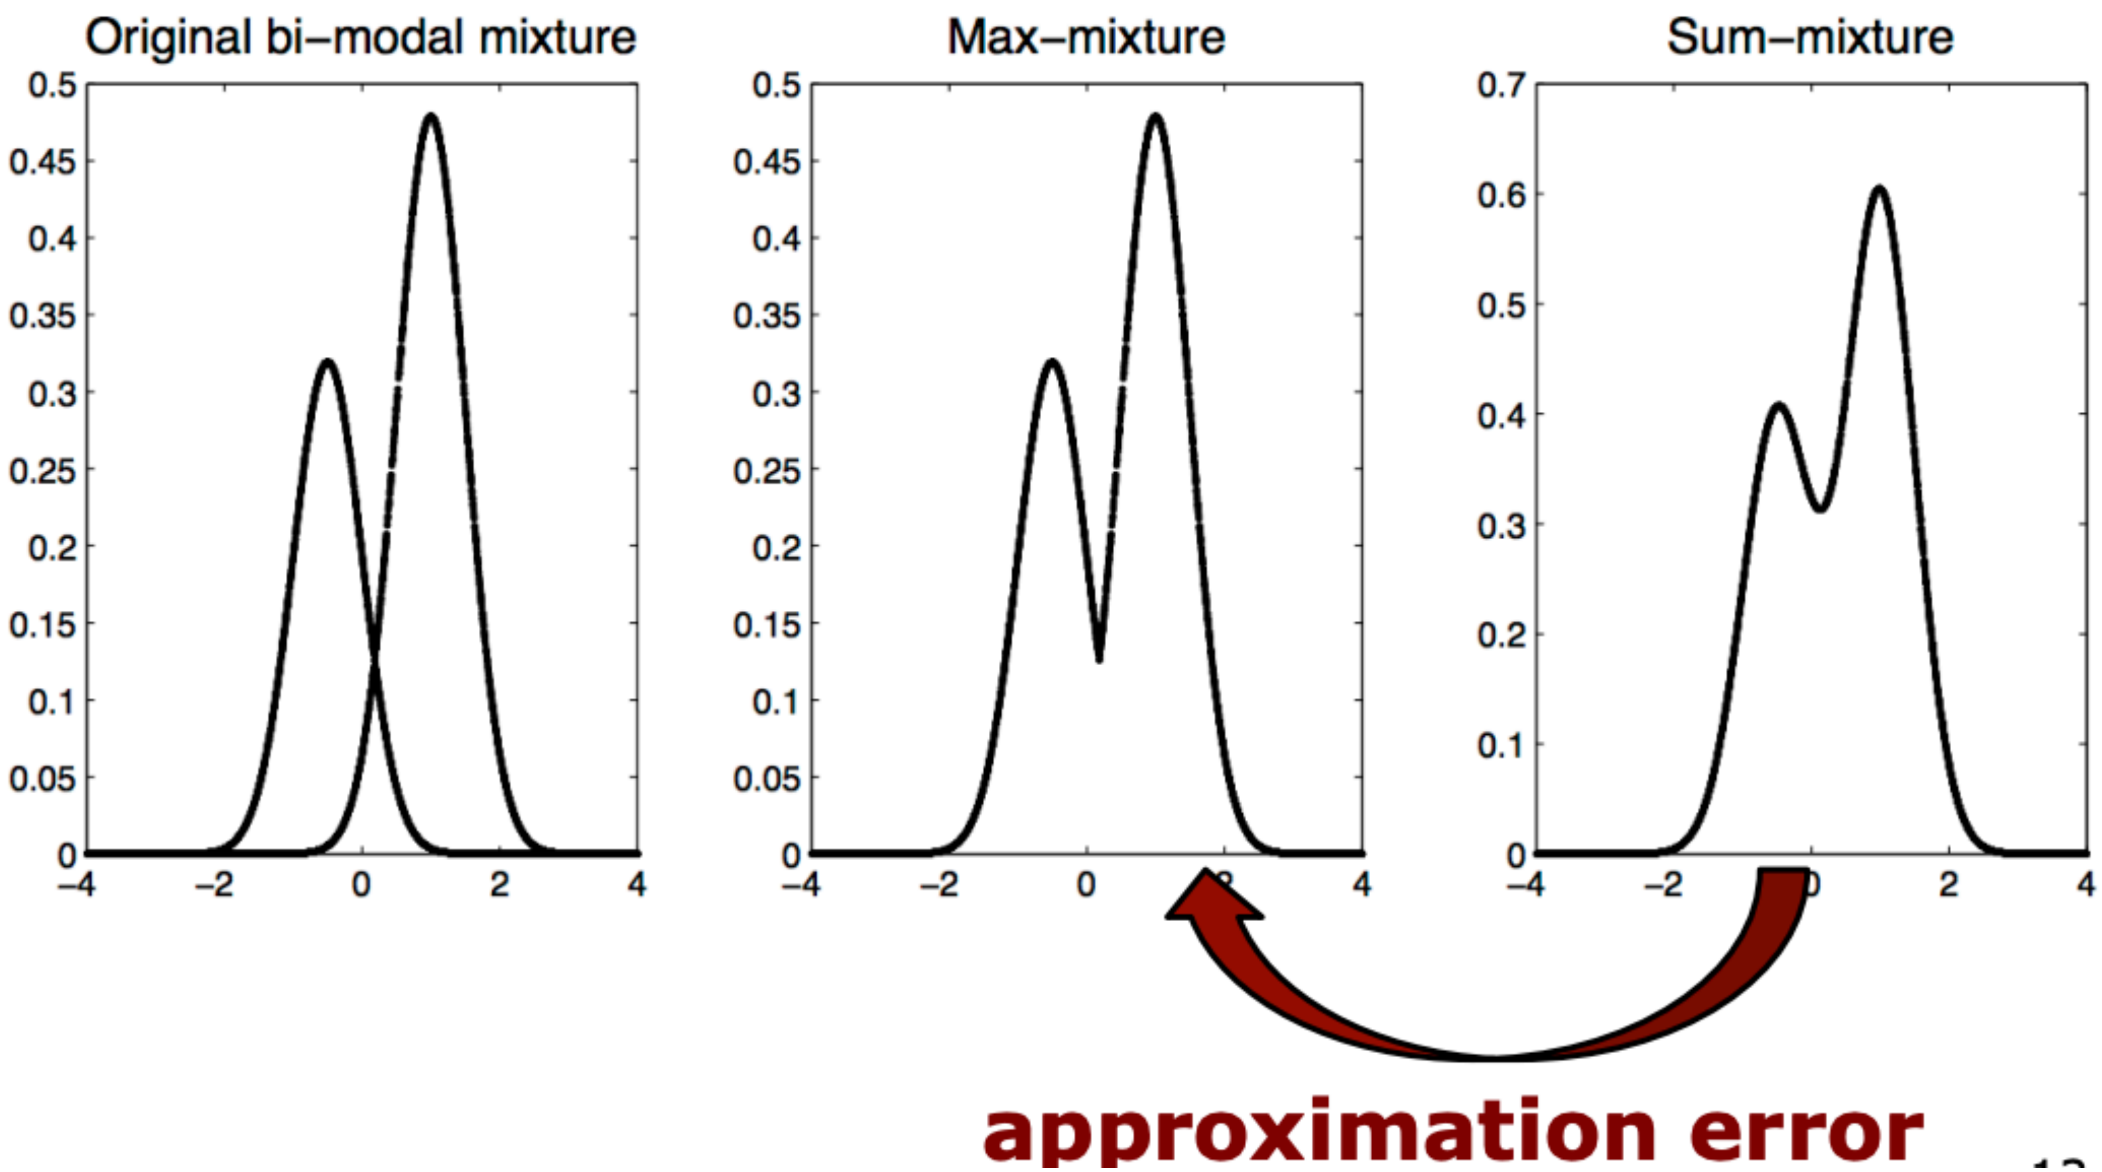

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Robust Estimation with Max-Mixture

Log Likelihood Of The Max-Mixture Formulation

- The log can be moved inside the max operator

$$p(\mathbf{z} \mid \mathbf{x}) \simeq \max_k w_k \eta_k \exp(-\frac{1}{2}\mathbf{e}_{ij_k}^T \boldsymbol{\Omega}_{ij_k} \mathbf{e}_{ij_k})$$

$$\log p(\mathbf{z} \mid \mathbf{x}) \simeq \max_k -\frac{1}{2}\mathbf{e}_{ij_k}^T \boldsymbol{\Omega}_{ij_k} \mathbf{e}_{ij_k} + \log(w_k \eta_k)$$

or: $$-\log p(\mathbf{z} \mid \mathbf{x}) \simeq \min_k \frac{1}{2}\mathbf{e}_{ij_k}^T \boldsymbol{\Omega}_{ij_k} \mathbf{e}_{ij_k} - \log(w_k \eta_k)$$

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Robust Estimation with Max-Mixture

- Easy to integrate in the optimizer:

1. Evaluate all k components
2. Select the component with the maximum log likelihood
3. Perform the optimization as before using only the max components (as a single Gaussian)

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Robust Estimation with Max-Mixture

Performance (Gauss vs. MM)

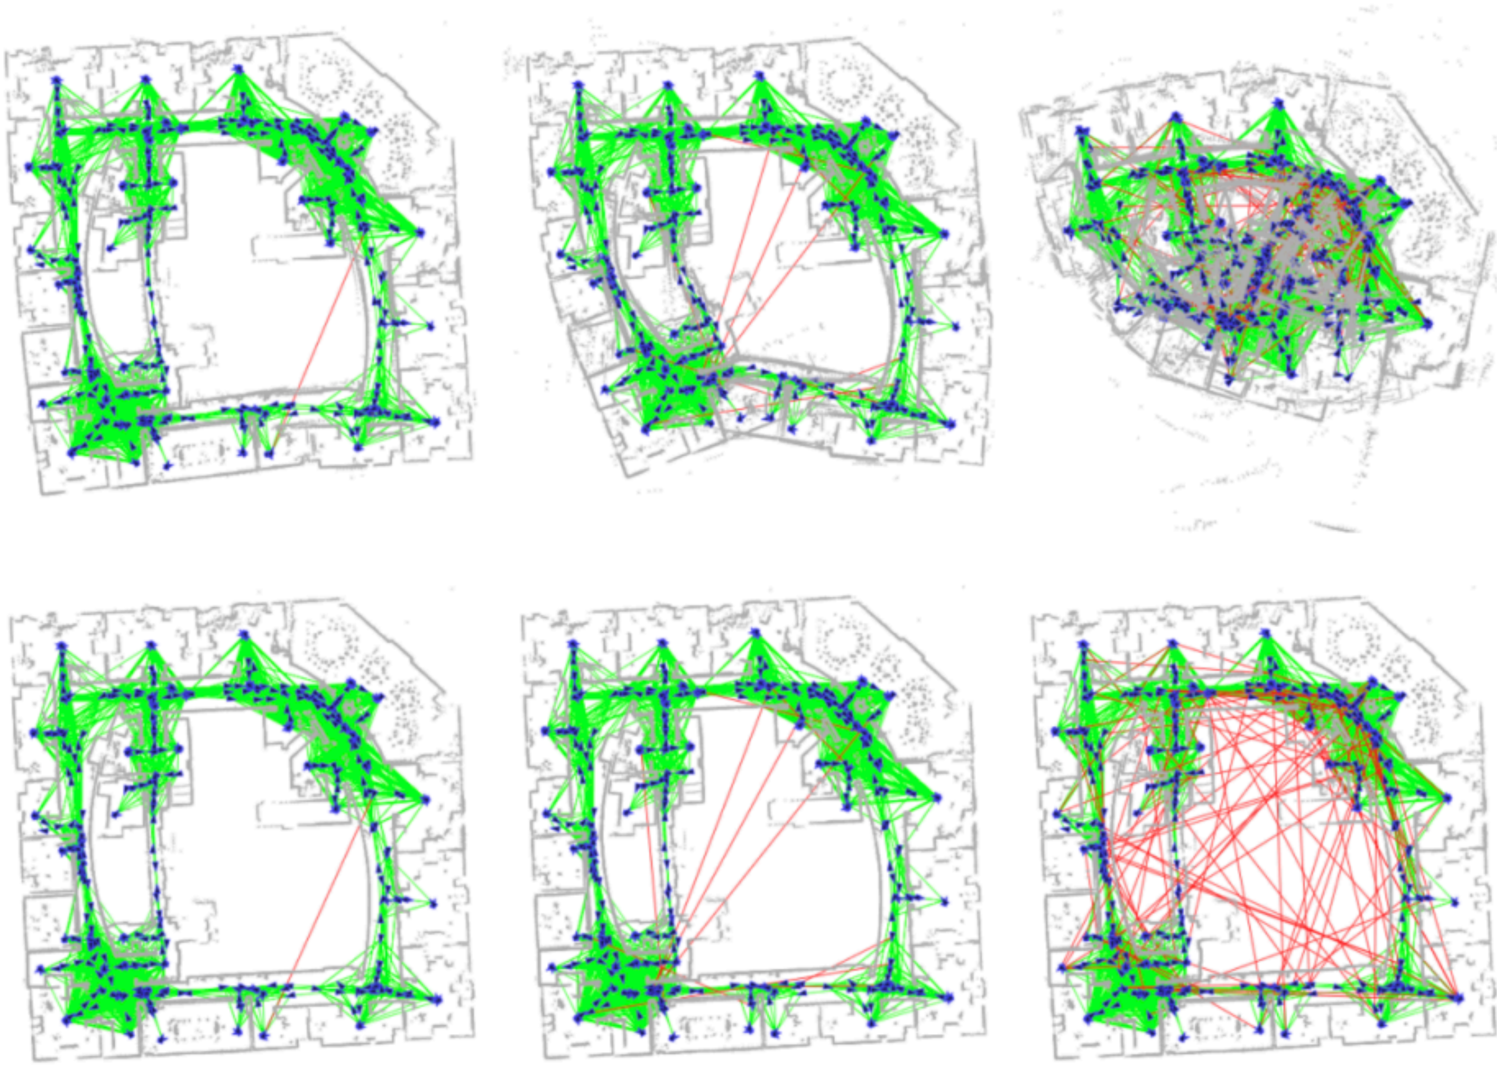

Performance (1 outlier)

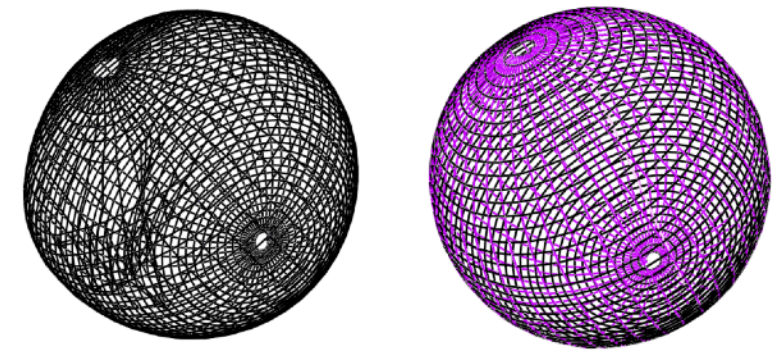

Gauss-Newton MM Gauss-Newton

Performance (10 outliers)

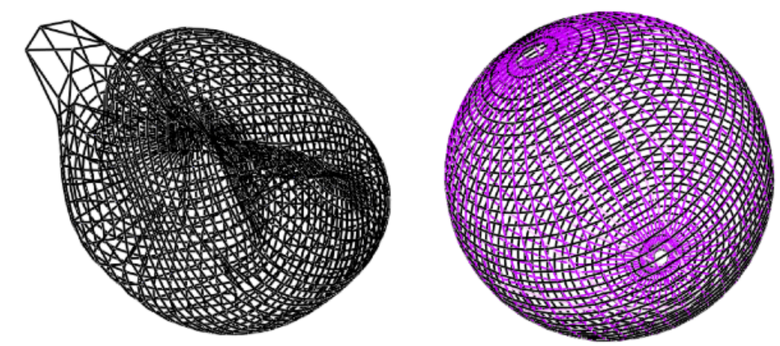

Gauss-Newton MM Gauss-Newton

Performance (100 outliers)

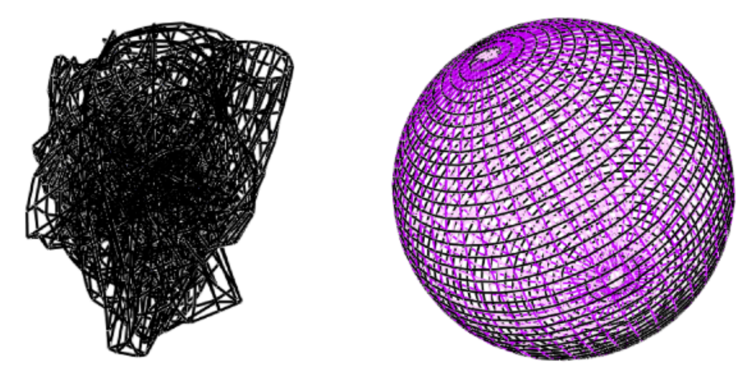

Gauss-Newton MM Gauss-Newton

[slides courtesy of Cyrill Stachniss]

[E. Olson and P. Agarwal, Inference on networks of mixtures for robust robot mapping, RSS 2012]

Today and Next Lecture

- **Robust estimation**:
  - Motivations: outliers, data association
  - Formulations: M-estimation & Maximum Consensus
- **Solvers for robust estimation**:
  - (RANSAC)
  - Iteratively Reweighted Least Squares (IRLS)
  - Max-mixture
  - Switchable constraints
  - Graduated non-convexity
  - Others: BnB, SDP relaxations, graph-theoretic pruning

Robust Estimation with Switchable Constraints

To appear in Proc. of IEEE Conf. on Intelligent Robots and Systems (IROS), 2012. DOI: not yet available
©2012 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Switchable Constraints for Robust Pose Graph SLAM

Niko Sünderhauf and Peter Protzel

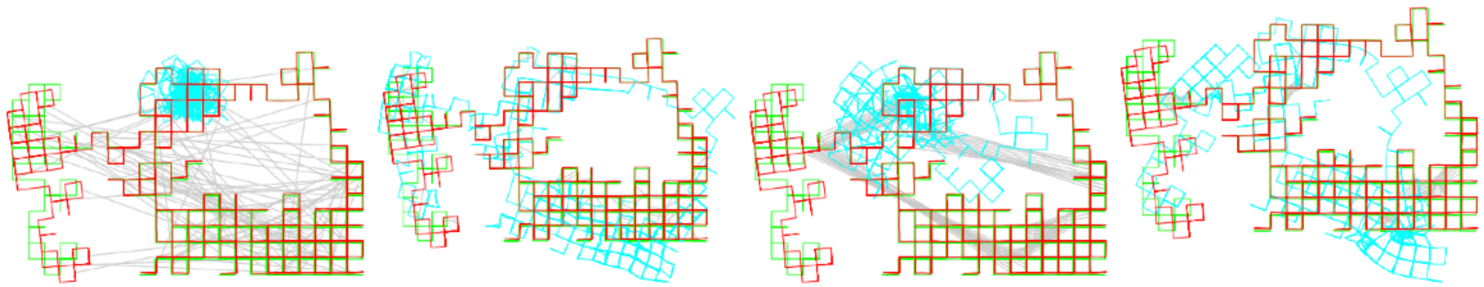

Fig. 1: Exemplary results of the proposed robust SLAM back-end on the synthetic Manhattan world dataset [10] that contains 3500 poses and 2099 loop closures. In these examples, we corrupted the dataset by introducing 100 additional wrong loop closures that might have been produced due to data association errors (e.g. failed place recognition) in the SLAM front-end. Current back-ends like g^2o [6] are not able to converge to a correct solution (shown in blue) despite being supported by so called robust cost functions like the Huber function [1]. Our robust solution (red) that uses switchable constraints correctly discards the wrong loop closure candidates (visible as grey links) *during* the optimization and converges to a correct solution. For comparison, the ground truth is plotted in green. Our robust back-end was able to cope with 1000 outliers on a number of 2D and 3D datasets. Notice that the outlier loop closure constraints have been added following different policies (from left to right: random, local, random group, local group) which are explained later on.

Robust Map Optimization using Dynamic Covariance Scaling

Pratik Agarwal, Gian Diego Tipaldi, Luciano Spinello, Cyrill Stachniss, and Wolfram Burgard

***Abstract*—Developing the perfect SLAM front-end that produces graphs which are free of outliers is generally impossible due to perceptual aliasing. Therefore, optimization back-ends need to be able to deal with outliers resulting from an imperfect front-end. In this paper, we introduce dynamic covariance scaling, a novel approach for effective optimization of constraint networks under the presence of outliers. The key idea is to use a robust function that generalizes classical gating and dynamically rejects outliers without compromising convergence speed. We implemented and thoroughly evaluated our method on publicly available datasets. Compared to recently published state-of-the-art methods, we obtain a substantial speed up without increasing the number of variables in the optimization process. Our method can be easily integrated in almost any SLAM back-end.**

I. Introduction

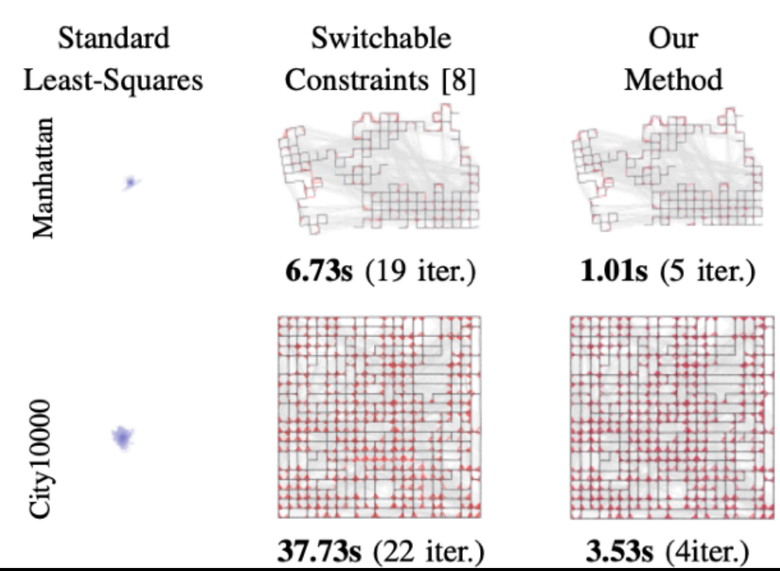

Robust Estimation with Switchable Constraints

$$X^*, S^* = \underset{X,S}{\operatorname{argmin}} \underbrace{\sum_i \| f(\mathbf{x}_i, \mathbf{u}_i) - \mathbf{x}_{i+1} \|^2_{\boldsymbol{\Sigma}_i}}_{\text{Odometry Constraints}} + \underbrace{\sum_{ij} \| s_{ij} \cdot (f(\mathbf{x}_i, \mathbf{u}_{ij}) - \mathbf{x}_j) \|^2_{\boldsymbol{\Lambda}_{ij}}}_{\text{Switchable Loop Closure Constraints}} \quad (1)$$

Also see: **Dynamic Covariance Scaling (DCS)**, which eliminates the switch variables, making the optimization more efficient

[N. Sunderhauf and P. Protzel, Switchable Constraints for Robust Pose Graph SLAM, IROS 2012]

[P. Agarwal, G. Tipaldi, L. Spinello, C. Stachniss, W. Burgard: "Robust Map Optimization Using Dynamic Covariance Scaling", ICRA 2013.

Switchable Constraints vs. IRLS

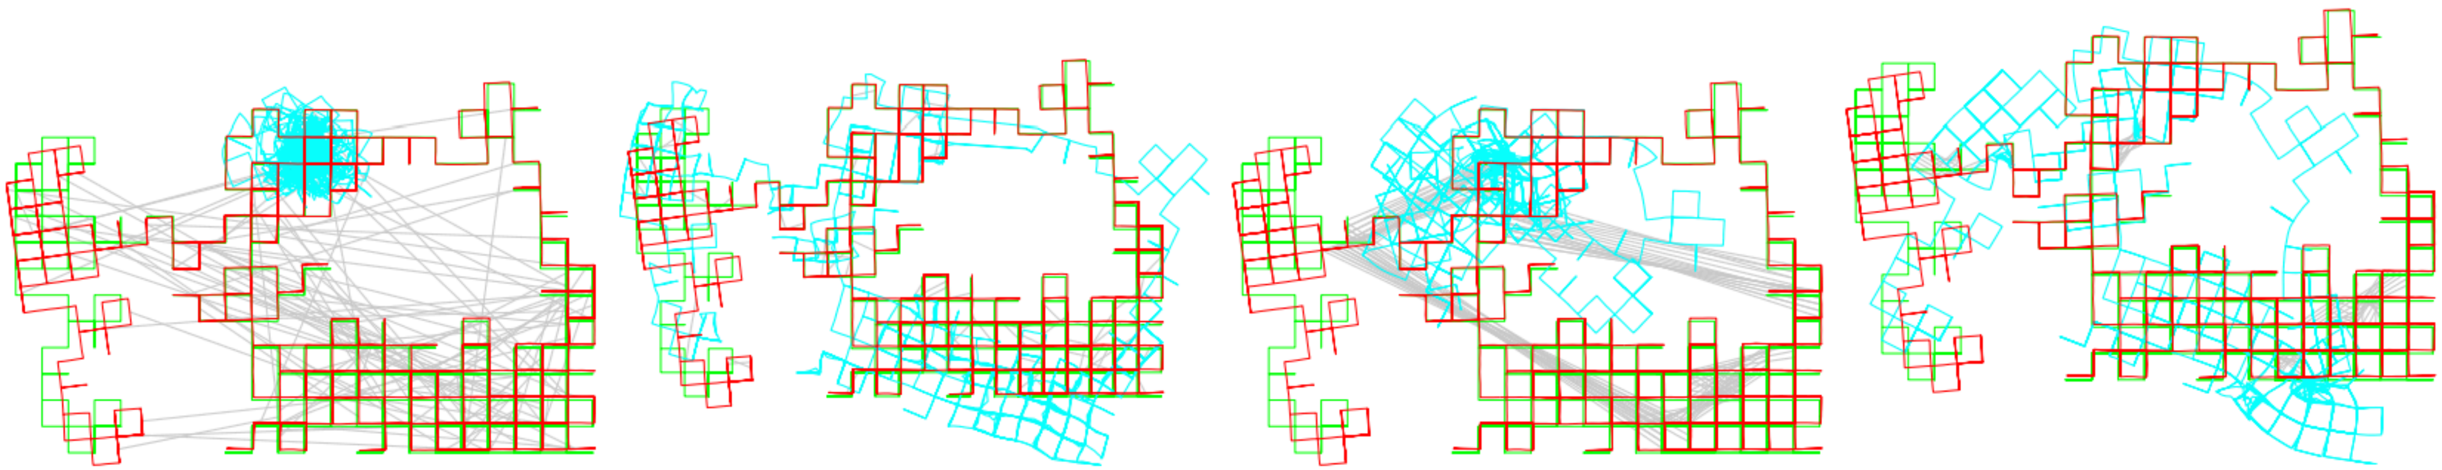

Fig. 1: Exemplary results of the proposed robust SLAM back-end on the synthetic Manhattan world dataset [10] that contains 3500 poses and 2099 loop closures. In these examples, we corrupted the dataset by introducing 100 additional wrong loop closures that might have been produced due to data association errors (e.g. failed place recognition) in the SLAM front-end. Current back-ends like g^2o [6] are not able to converge to a correct solution (shown in blue) despite being supported by so called robust cost functions like the Huber function [1]. Our robust solution (red) that uses switchable constraints correctly discards the wrong loop closure candidates (visible as grey links) *during* the optimization and converges to a correct solution. For comparison, the ground truth is plotted in green. Our robust back-end was able to cope with 1000 outliers on a number of 2D and 3D datasets. Notice that the outlier loop closure constraints have been added following different policies (from left to right: random, local, random group, local group) which are explained later on.

[N. Sunderhauf and P. Protzel, Switchable Constraints for Robust Pose Graph SLAM, IROS 2012]

Robust Estimation with Switchable Constraints

	Switchable Constraints				Max-Mixtures				RRR				Best
	RMSE [m]			Time [s]	RMSE [m]			Time [s]	RMSE [m]			Time [s]	
Dataset	median	mean	max	mean	median	mean	max	mean	median	mean	max	mean	
Manhattan	1.16	1.36	26.42	9.7	1.18	1.49	38.28	13.9	7.38	11.64	37.40	9.8	SC
City	0.063	0.063	0.063	38.8	0.058	0.251	64.18	47.7	0.94	1.60	5.11	523.3	SC
Ring	-	4.39	-	0.07	-	15.06	-	0.12	-	5.21	-	0.19	SC
RingCity	-	1.82	-	0.41	-	41.13	-	2.0	-	4.18	-	54.0	SC
Bovisa-04	-	2.39	-	1.1	-	11.81	-	1.4	-	3.01	-	5.9	SC
Bovisa-06	-	9.38	-	1.1	-	7.67	-	1.4	-	3.95	-	2.9	RRR
Bicocca	2.73	2.67	2.98	0.8	3.93	3.92	5.59	1.1	1.10	1.64	2.96	2.29	RRR

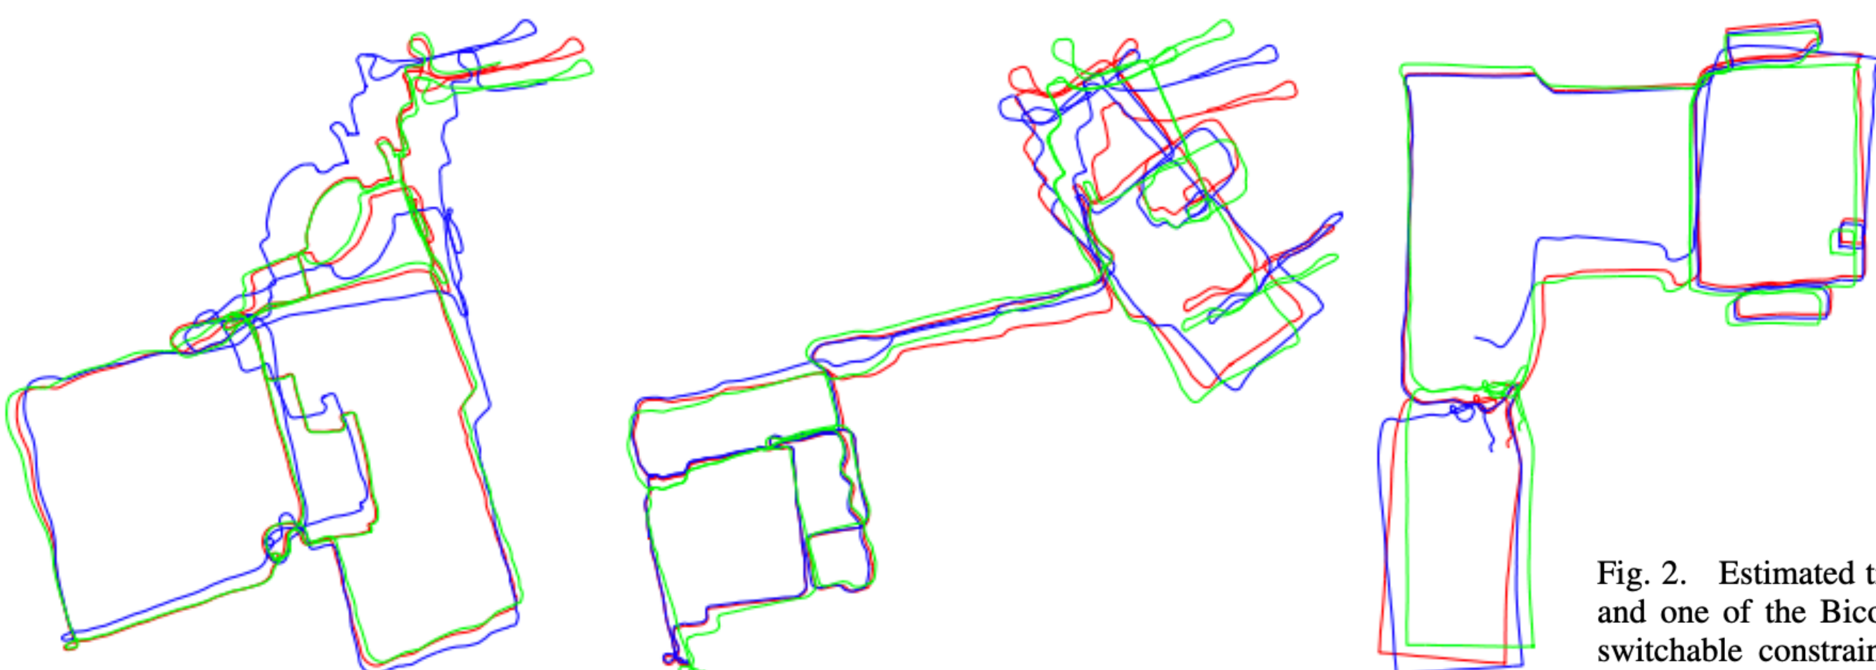

Fig. 2. Estimated trajectories for the Bovisa-04 (left), Bovisa-06 (middle), and one of the Bicocca (right) datasets. Colors indicate the used method: switchable constraints (red), max-mixtures (blue), RRR (green). The trajectories are plotted in their optimal alignment with the ground truth (not shown) according to the `get_ATE()` error function from the Rawseeds toolkit.

[N. Sunderhauf and P. Protzel, Switchable constraints vs. max-mixture models vs. RRR - A comparison of three approaches to robust pose graph SLAM, ICRA 2013]

Graduated non-convexity

First insight: equivalence between M-estimation and formulations with switchable constraints:

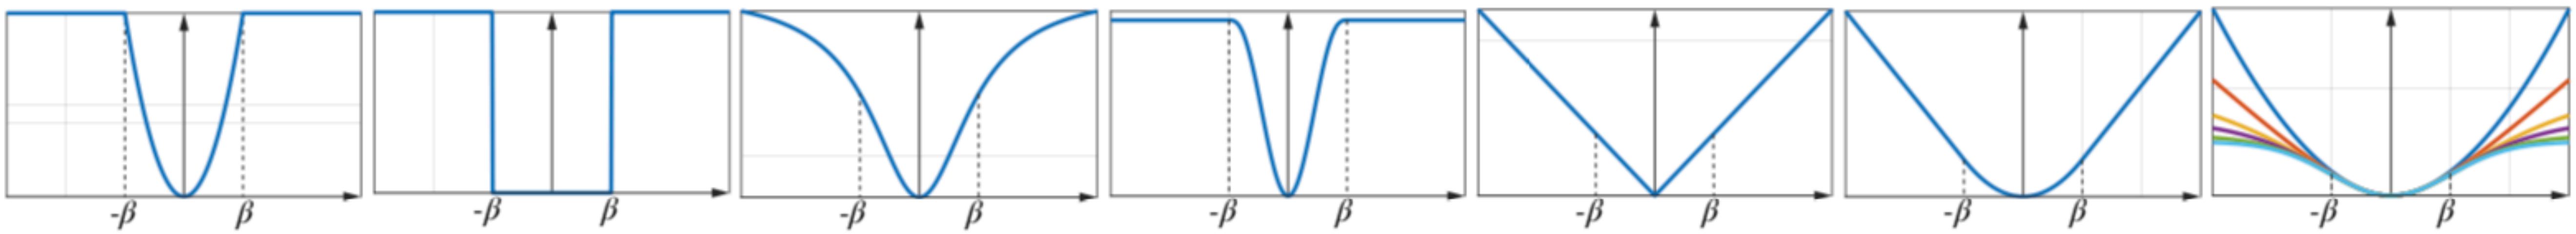

(a) TLS (b) MC (c) GM (d) TB (e) L1 (f) Huber (g) Adaptive

Black-Rangarajan duality:

Theorem 1 [Informal - Black, Rangarajan, 1996] We can rewrite common robust loss functions by adding auxiliary variables θ_i (one for each measurement)

$$\arg\min_{x \in \mathbb{X}} \sum_{i \in \mathcal{M}} \rho\left(r_i(\boldsymbol{x}, \boldsymbol{y}_i)\right)$$

Robust loss function

$$\arg\min_{x \in \mathbb{X}, \theta_i \in [0,1]} \sum_{i \in \mathcal{M}} \theta_i r_i^2(\boldsymbol{x}, \boldsymbol{y}_i) + \Phi_\rho(\theta_i)$$

"Outlier process"

[Black, Rangarajan, "On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision, IJCV'96.]
[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Graduated non-convexity

Second insight: alternation-based solver

$$\arg\min_{\substack{x \in \mathbb{X}, \\ \theta_i \in \{0,1\}, \forall i}} \sum_{i \in \mathcal{M}} \theta_i \|r_i(x, y_i)\|^2 + (1 - \theta_i)\bar{c}^2$$

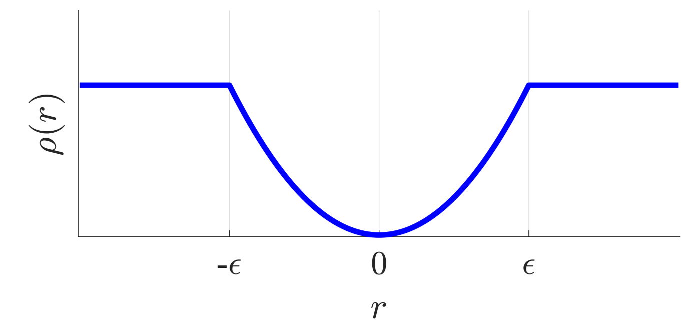

Potential approach: Alternating Minimization (Block Coordinate

a **Variable Update:** fix weights θ_i, optimize variable x

✓ becomes a weighted least squares problem

b **Weight Update:** fix variable x, optimize θ_i

✓ splits into scalar optimization problems

✓ can be solved in closed form

ISSUE: approach easily gets stuck in local minima

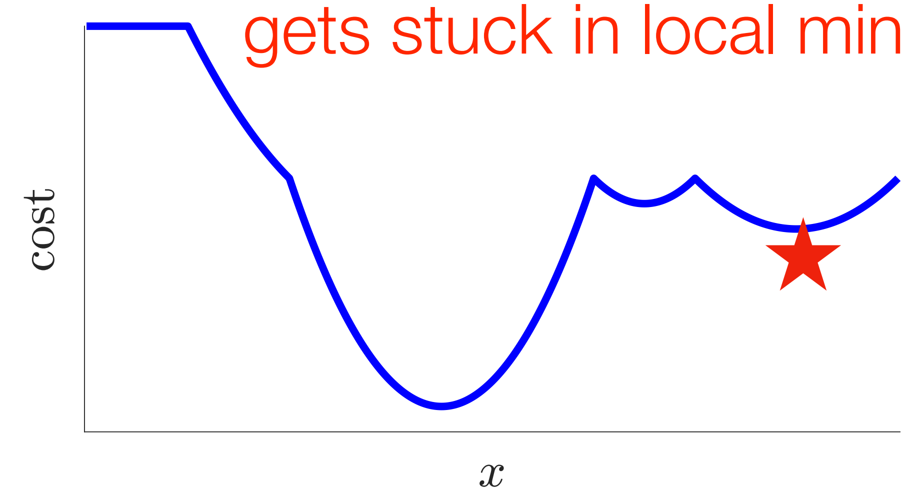

[https://mit-spark.github.io/GlobalOptimization-ICCV2019/]

Graduated non-convexity

Second insight: alternation-based solver

$$\arg\min_{\substack{x \in \mathbb{X}, \\ \theta_i \in \{0,1\}, \forall i}} \sum_{i \in \mathcal{M}} \theta_i \|r_i(x, y_i)\|^2 + (1 - \theta_i)\bar{c}^2$$

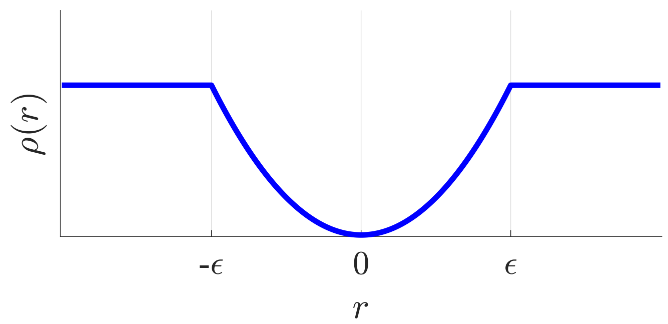

Key idea to avoid getting stuck in local minima:

- start from a convex approximation of the cost function
- gradually increase non-convexity until you recover

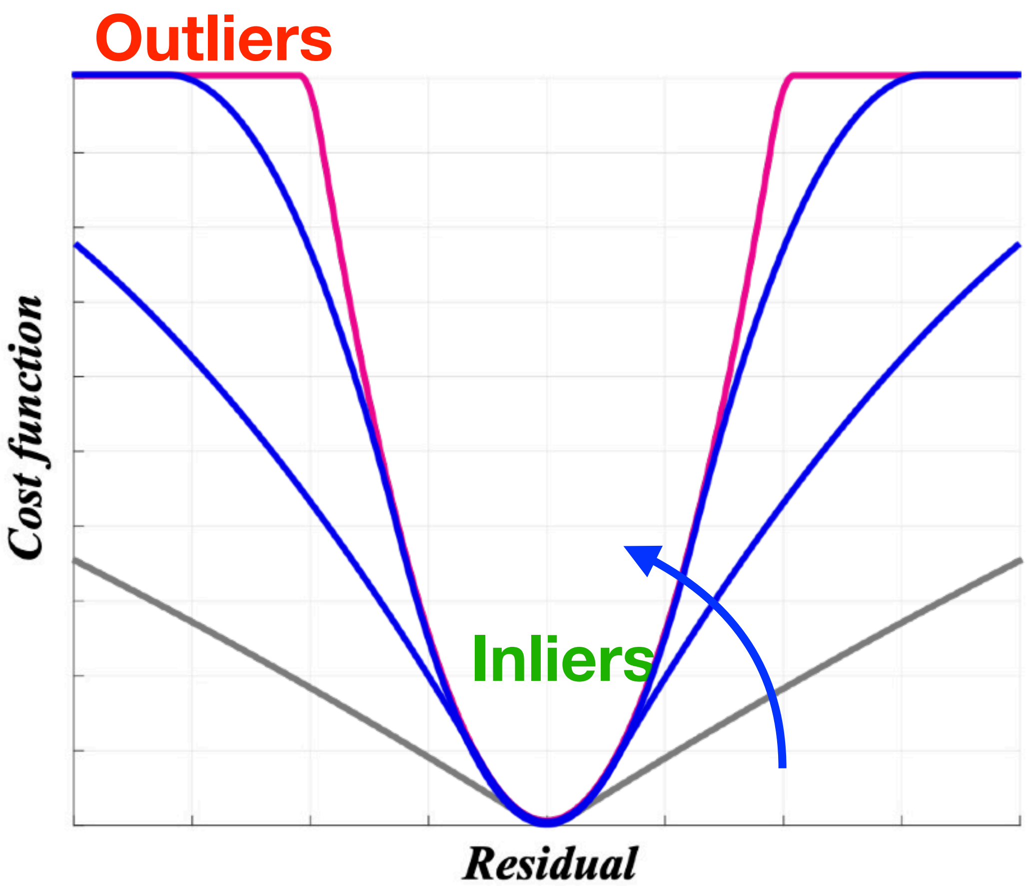

[Black, Rangarajan, "On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision, IJCV'96.]
[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Graduated non-convexity

Truncated Least Squares Loss

$$\rho(r(\boldsymbol{x}, \boldsymbol{y}_i))$$

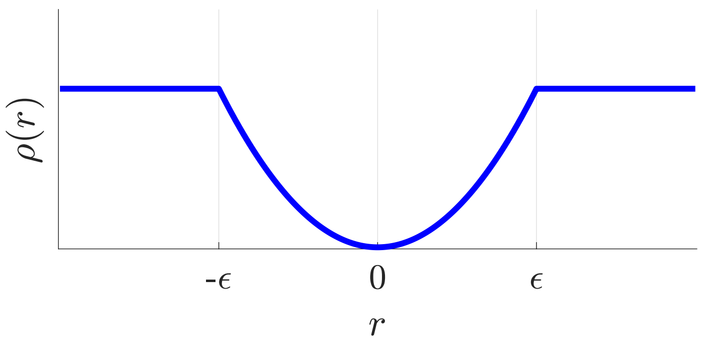

Non-convex cost (hard to optimize)

$$\sum_{i \in \mathcal{M}} \rho(r(\boldsymbol{x}, \boldsymbol{y}_i))$$

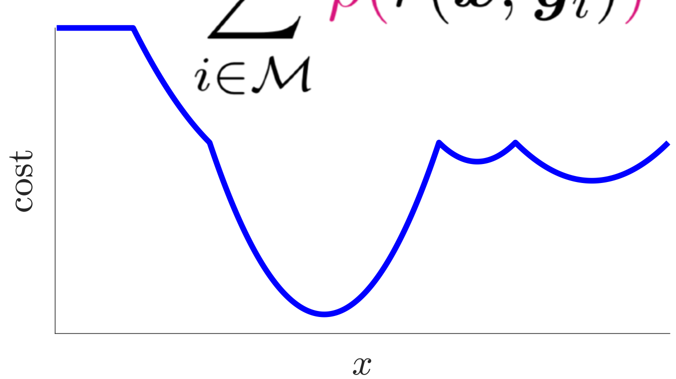

Graduated Non-Convexity

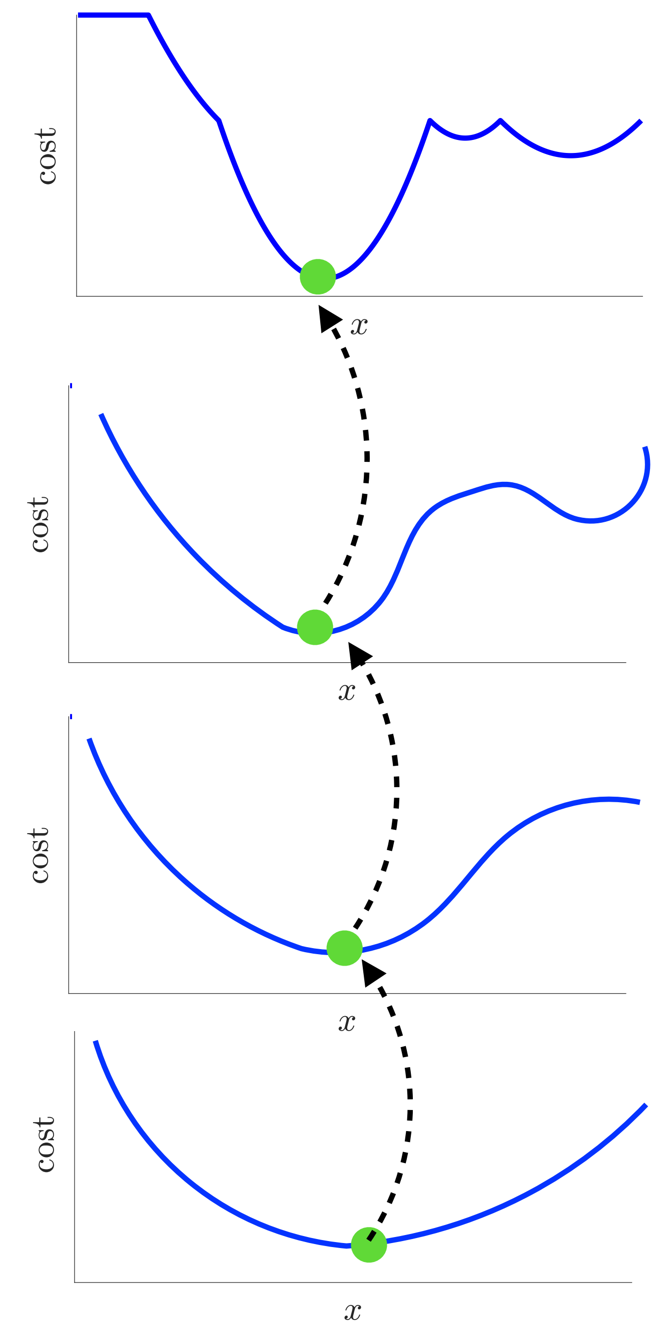

[Black, Rangarajan, "On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision, IJCV'96.]
[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Graduated non-convexity

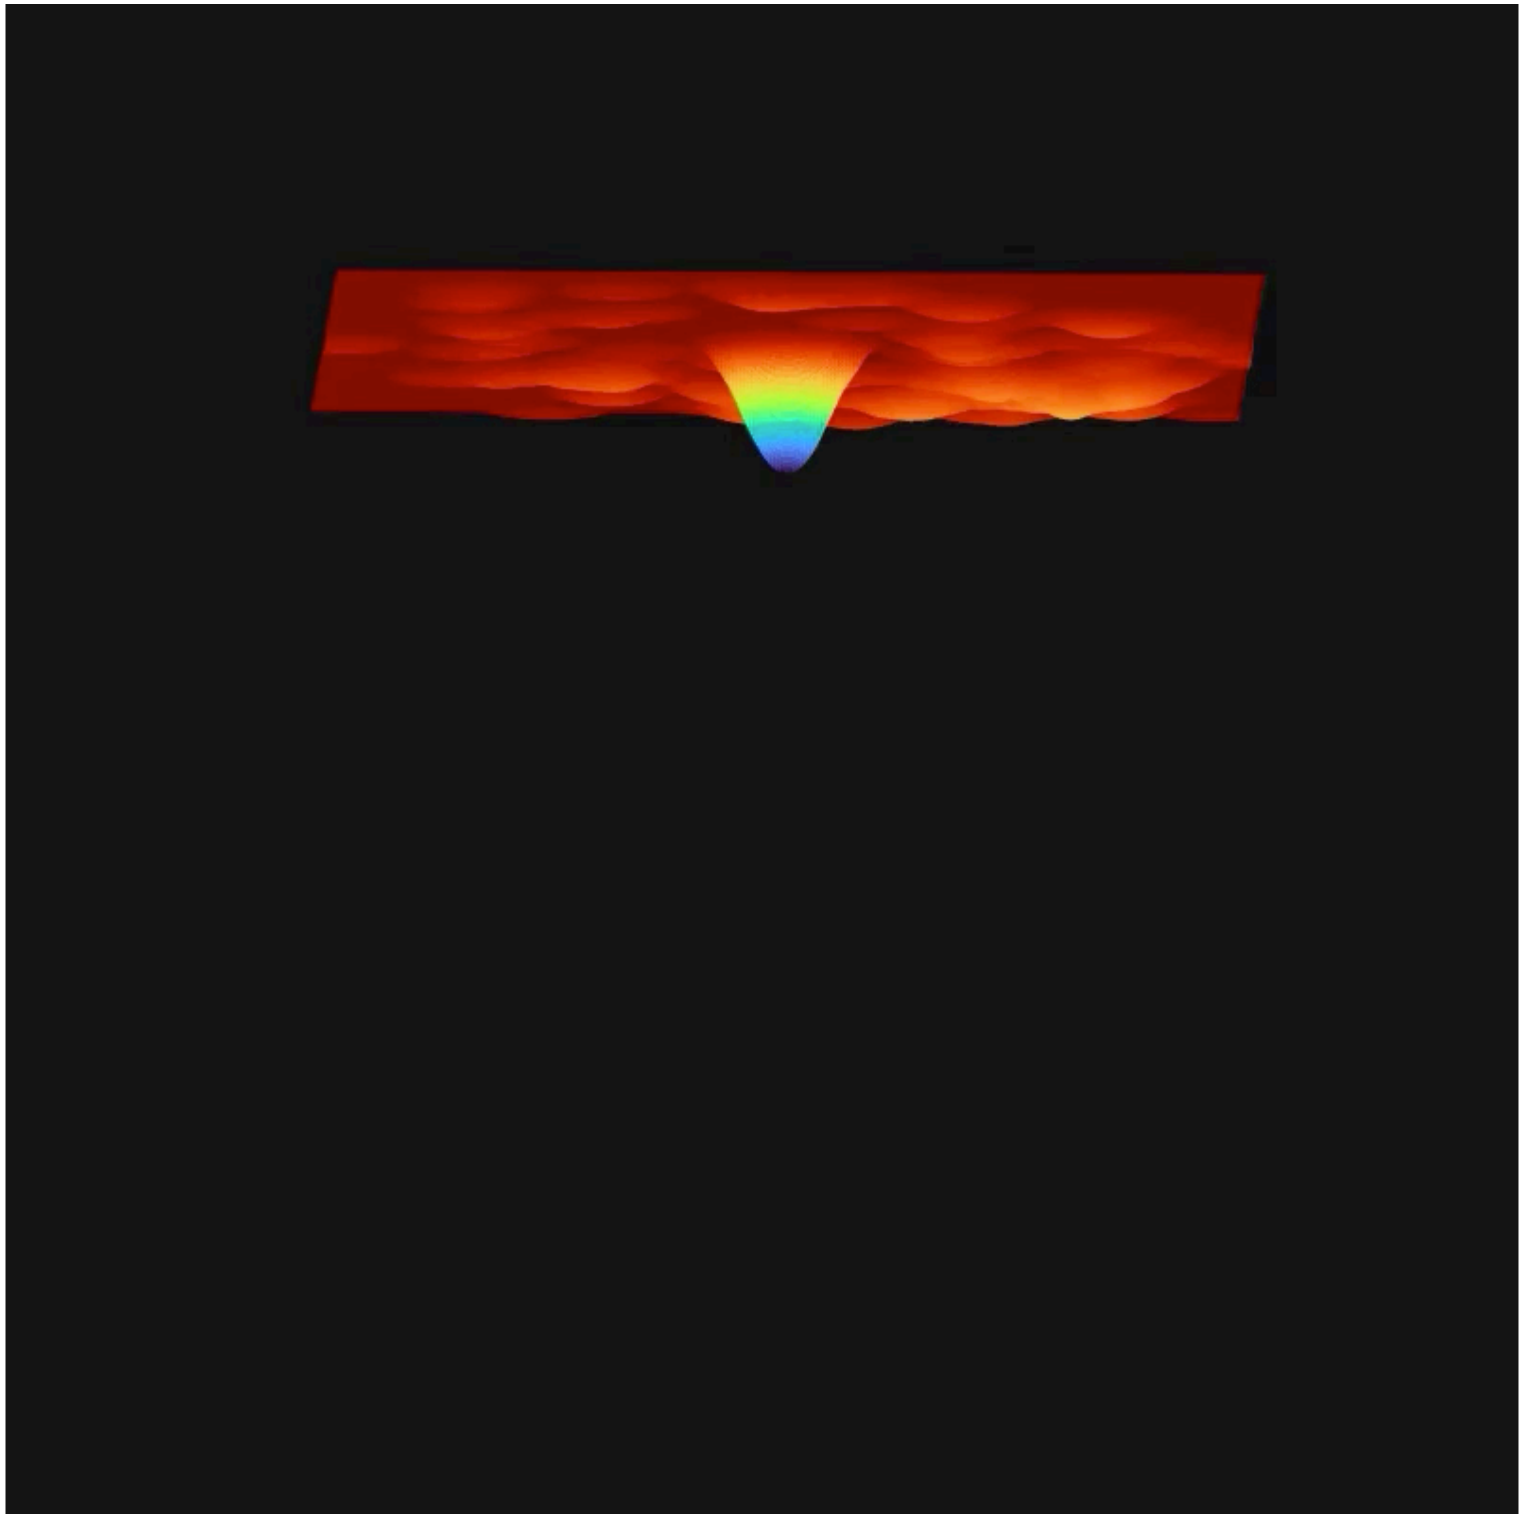

[Black, Rangarajan, “On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision, IJCV’96.]
[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Graduated non-convexity algorithm

$$\arg\min_{\substack{x\in\mathbb{X},\\ \theta_i\in\{0,1\},\forall i}} \sum_{i\in\mathcal{M}} \theta_i \|r_i(x,y_i)\|^2 + (1-\theta_i)\bar{c}^2$$

Surrogate function with parameter μ

$$\arg\min_{\substack{x\in\mathbb{X},\\ \theta_i\in[0,1],\forall i}} \sum_{i\in\mathcal{M}} \theta_i \|r_i(x,y_i)\|^2 + \frac{\mu(1-\theta_i)}{\mu+\theta_i}\bar{c}^2$$

Graduated Non-Convexity (GNC)

1. **Initialization:** set $\mu \to 0$
   - a. Set all weights $\theta_i = 1$
   - b. Variable Update (weighted least square)
2. While cost function decrease
   - a. Weight Update (closed-form)
   - b. Variable Update (weighted least square)
   - c. Increase Non-Convexity: $\mu_t = \delta \cdot \mu_{t-1}, \delta > 1$

Intuition

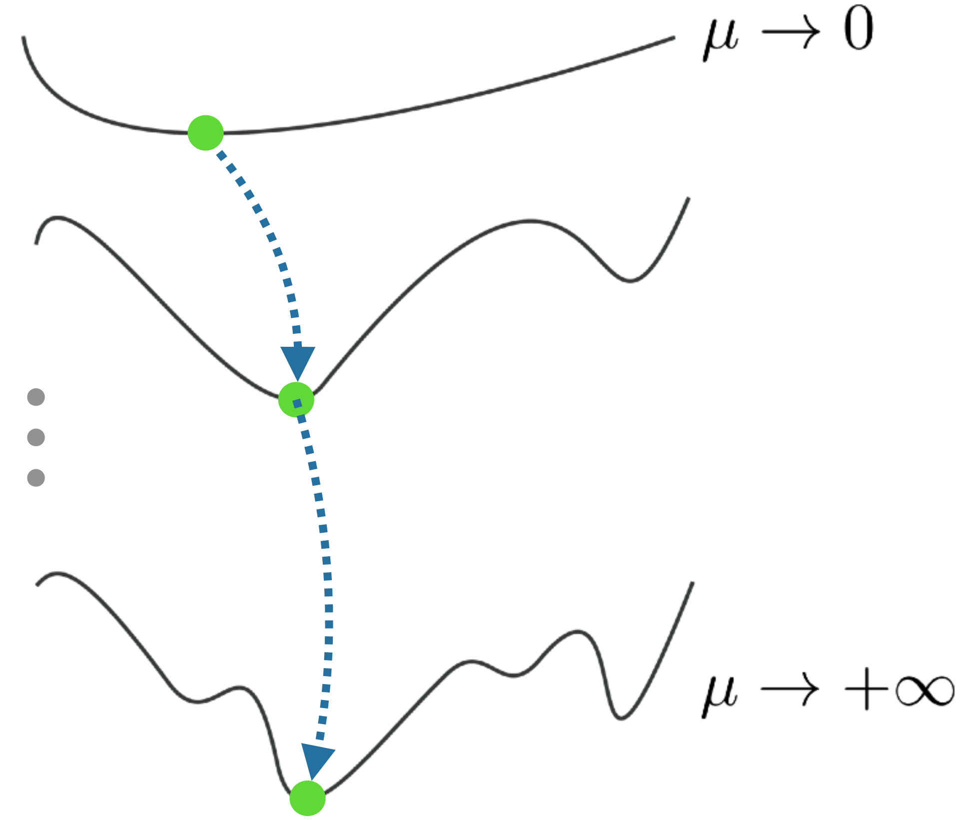

[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Graduated non-convexity for SLAM

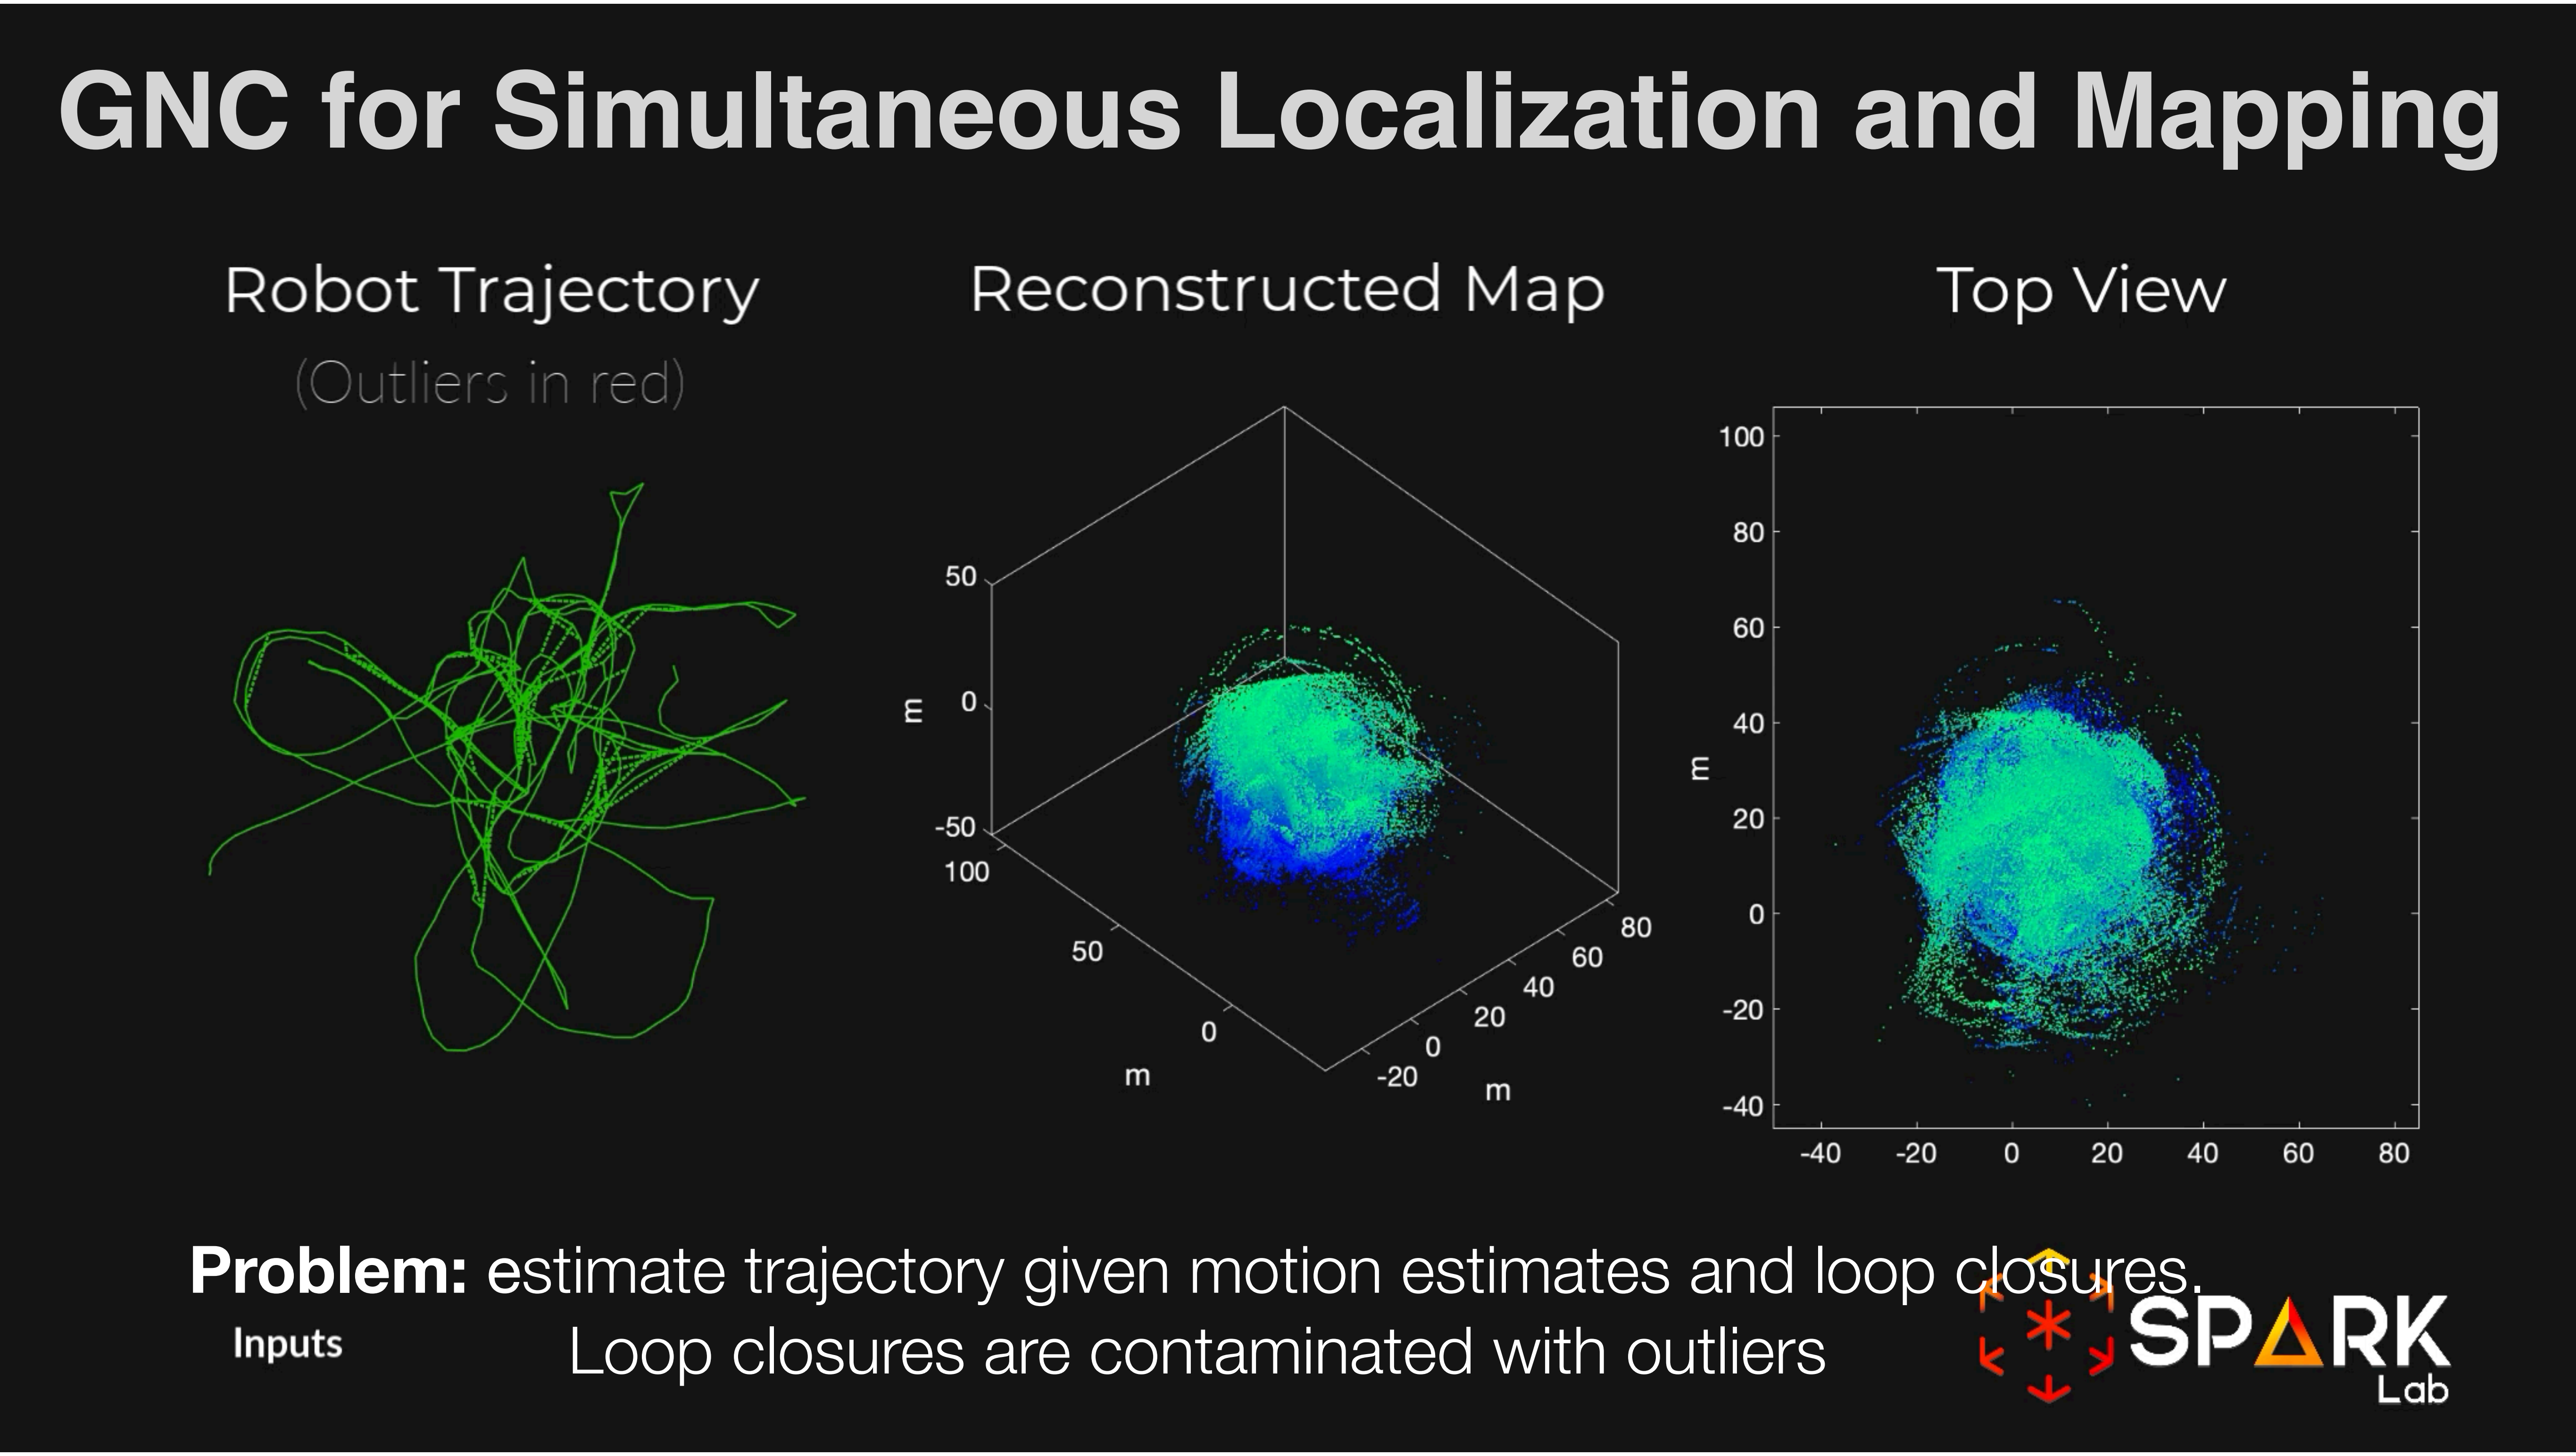

[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Pose Graph Optimization Results

CSAIL

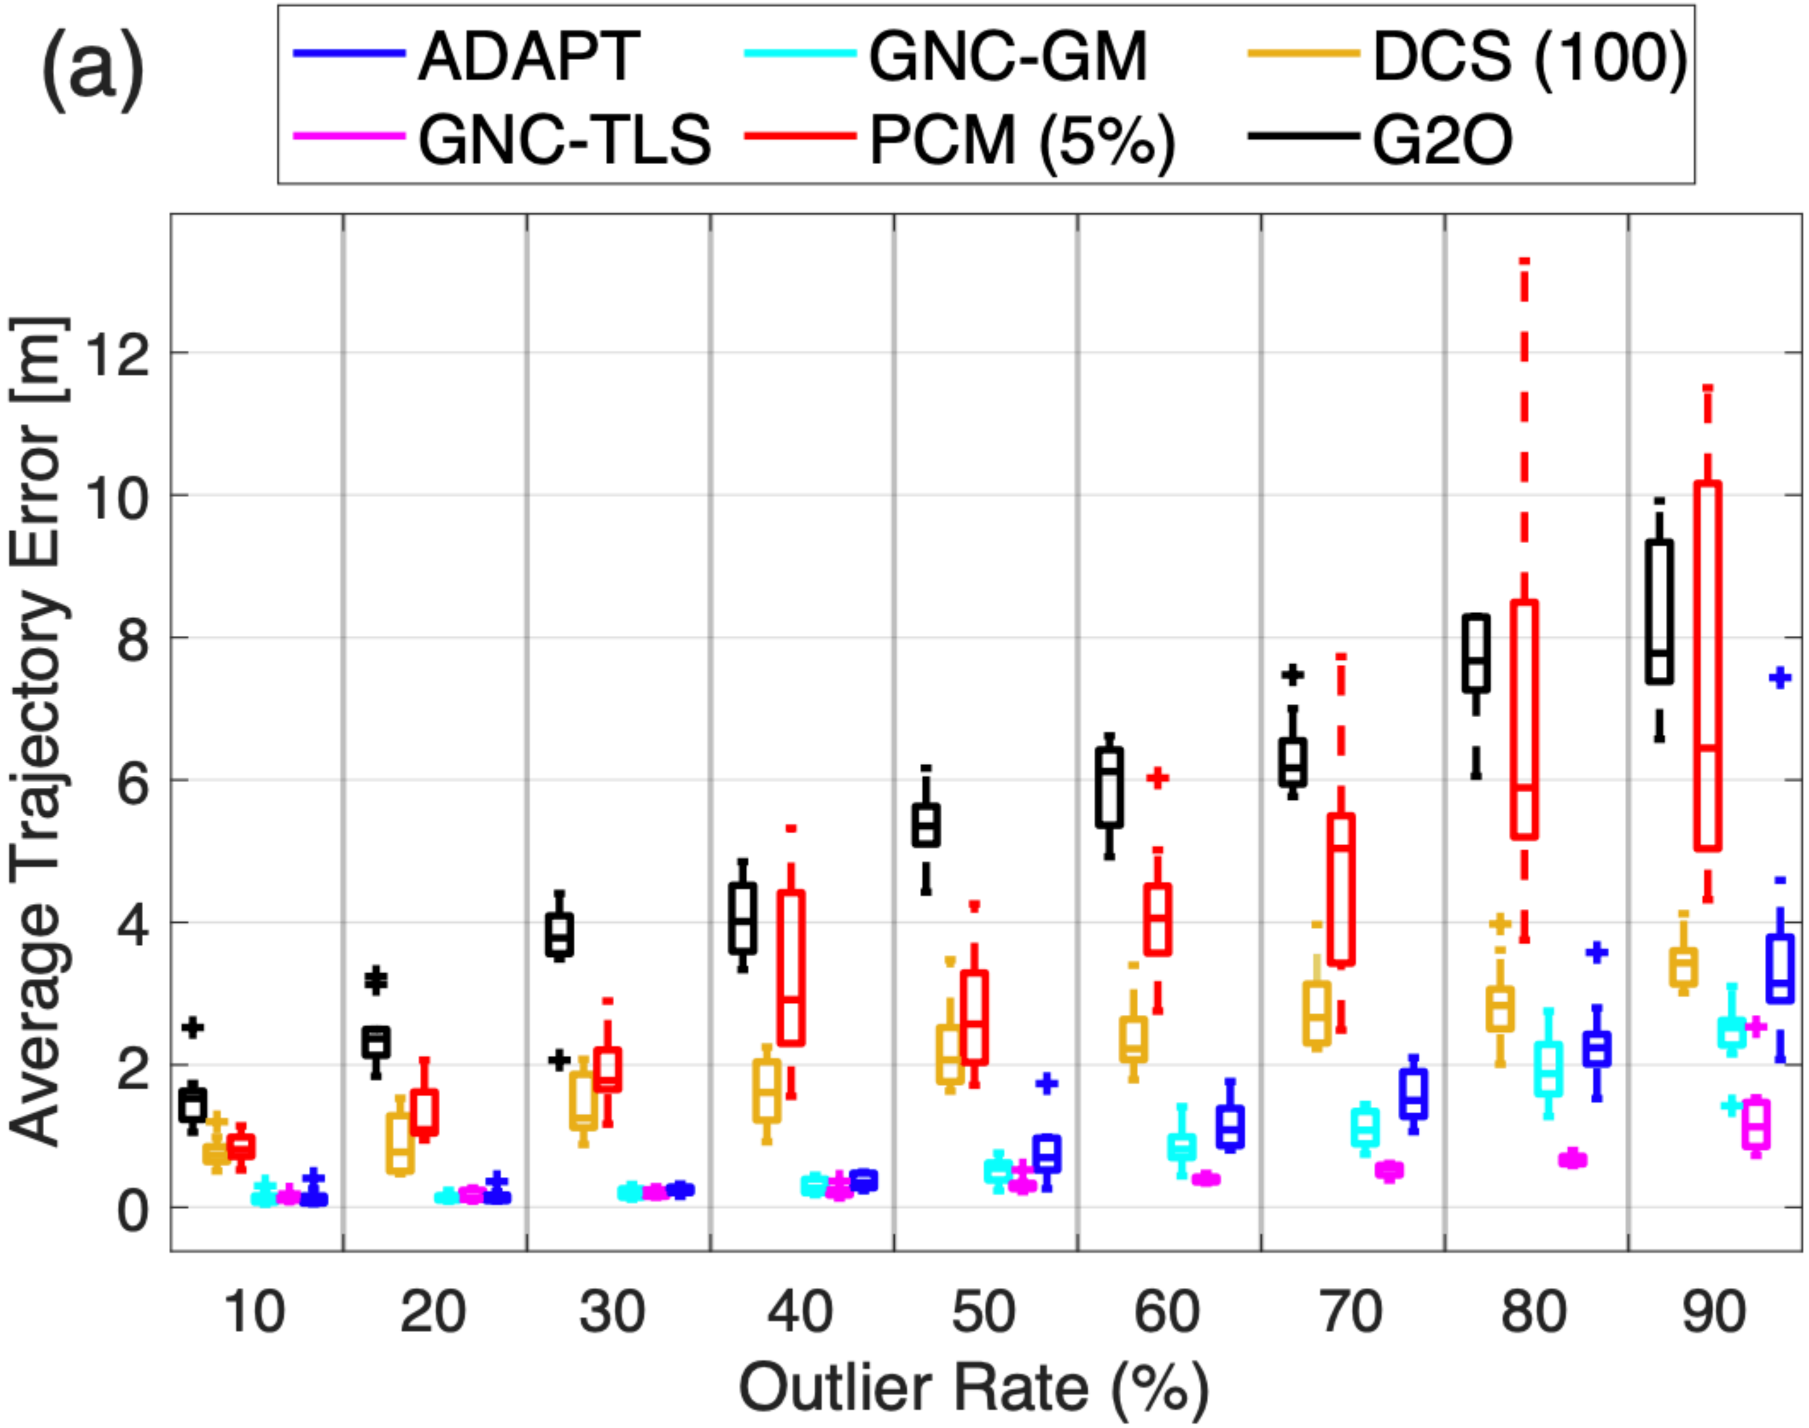

[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Pose Graph Optimization Results

INTEL

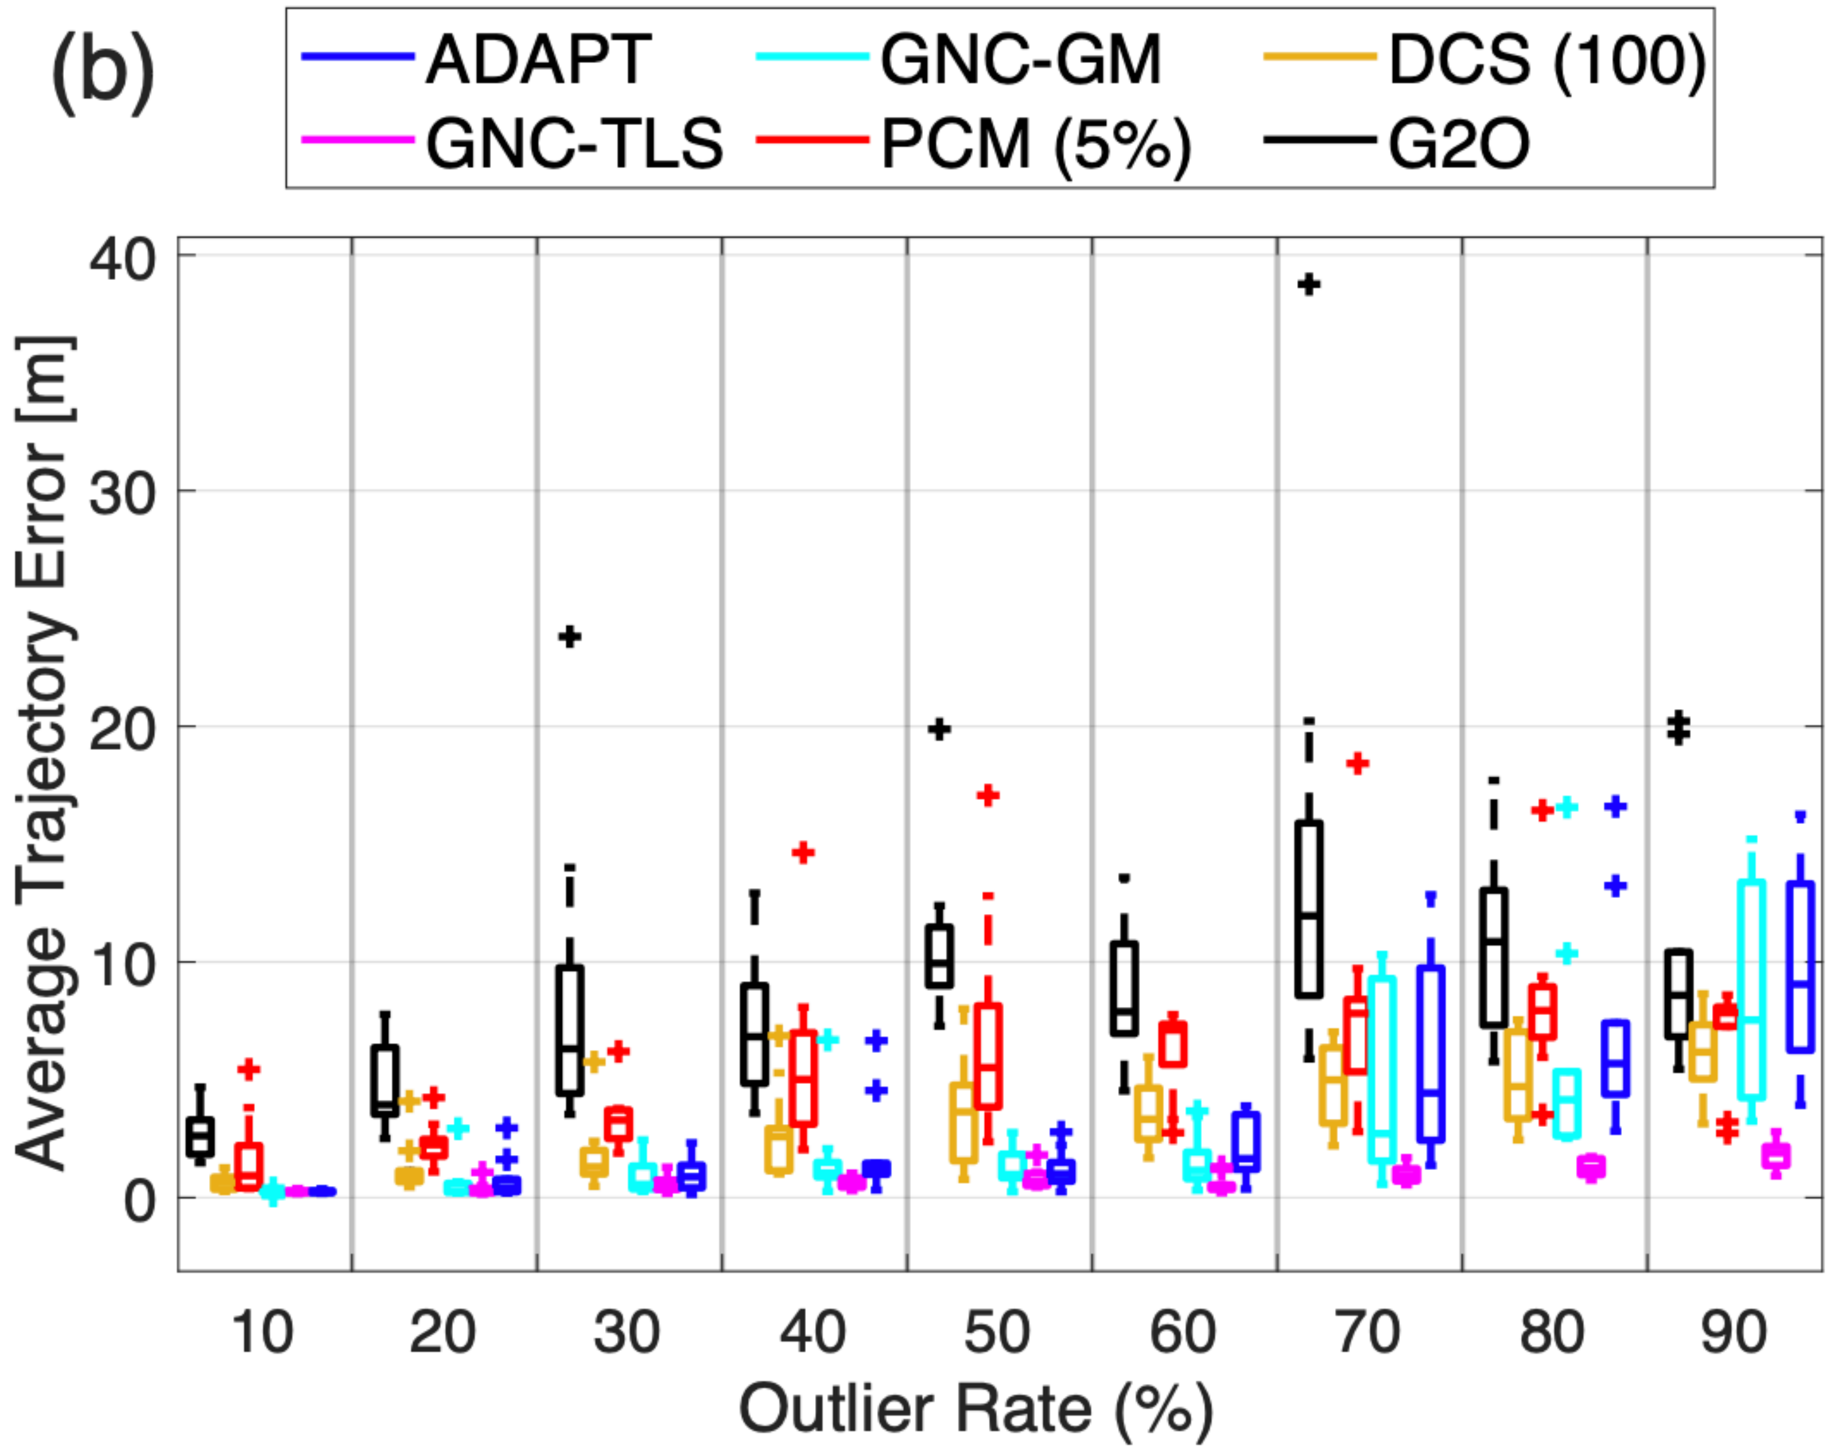

[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Other applications of GNC

Mesh Registration

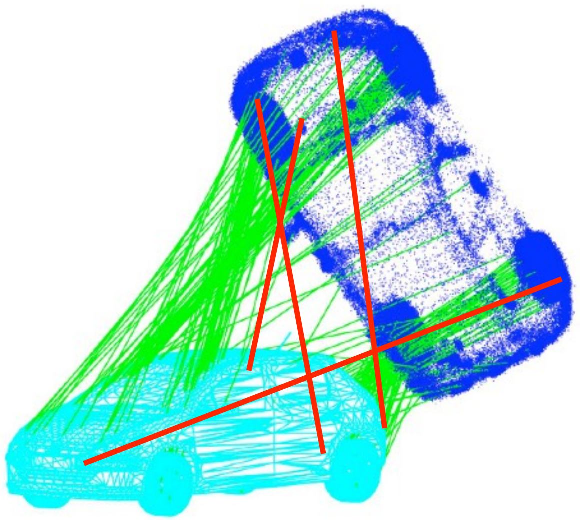

up to

80

outliers

Shape Alignment

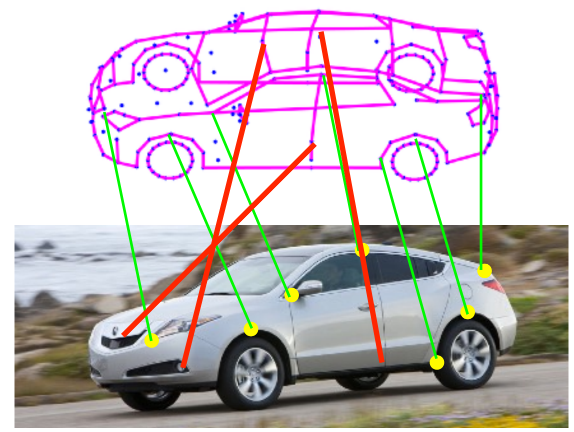

up to

70

outliers

SLAM

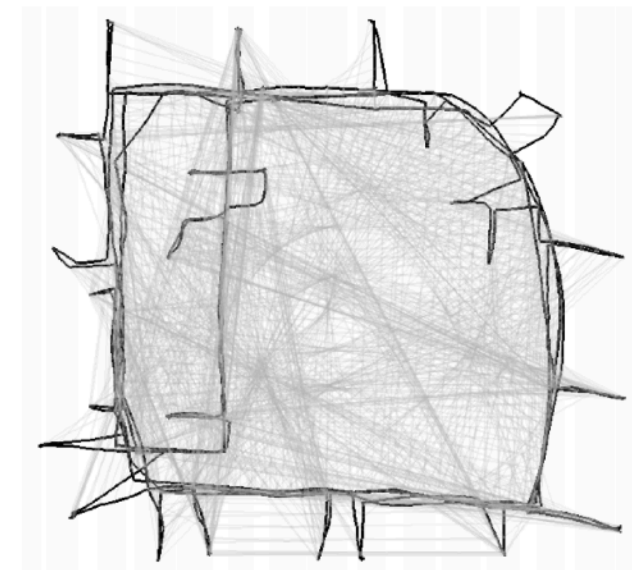

up to

90

outliers

No need for initial guess (as opposed to local solvers)
No need for minimal solver (as opposed to RANSAC)
GNC implementation available in Matlab and GTSAM

[Yang, Antonante, Tzoumas, Carlone. Graduated non-convexity for robust spatial perception: from non-minimal solvers to global outlier rejection. RAL 2020. (best paper in robot vision at ICRA 2020)]

Today and Next Lecture

- **Robust estimation**:
  - Motivations: outliers, data association
  - Formulations: M-estimation & Maximum Consensus
- **Solvers for robust estimation**:
  - (RANSAC)
  - Iteratively Reweighted Least Squares (IRLS)
  - Max-mixture
  - Switchable constraints
  - Graduated non-convexity
  - Others: BnB, SDP relaxations, graph-theoretic pruning

Convex relaxations, graph theory

Pairwise Consistent Measurement Set Maximization for Robust Multi-robot Map Merging

Joshua G. Mangelson, Derrick Dominic, Ryan M. Eustice, and Ram Vasudevan

***Abstract*— This paper reports on a method for robust selection of inter-map loop closures in multi-robot simultaneous localization and mapping (SLAM). Existing robust SLAM methods assume a good initialization or an "odometry backbone" to classify inlier and outlier loop closures. In the multi-robot case, these assumptions do not always hold. This paper presents an algorithm called Pairwise Consistency Maximization (PCM) that estimates the largest *pairwise internally consistent set* of measurements. Finding the largest pairwise internally consistent set can be transformed into an instance of the maximum clique problem from graph theory, and by leveraging the associated literature it can be solved in real-time. This paper evaluates how well PCM approximates the combinatorial gold standard using simulated data. It also evaluates the performance of PCM on synthetic and real-world data sets in comparison with DCS, SCGP, and RANSAC, and shows that PCM significantly outperforms these methods.**

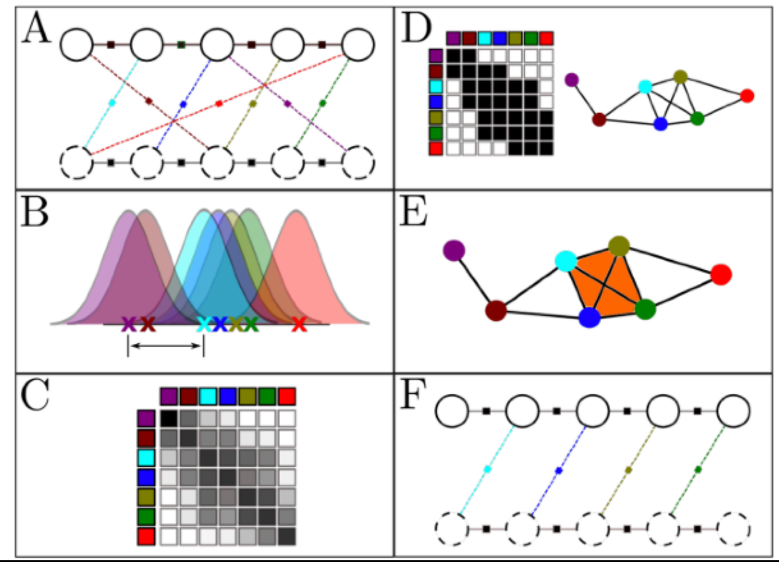

Certifiable Outlier-Robust Geometric Perception: Exact Semidefinite Relaxations and Scalable Global Optimization

Heng Yang, *Student Member, IEEE,* and Luca Carlone, *Senior Member, IEEE*

Abstract—We propose the first general and scalable framework to design *certifiable* algorithms for robust geometric perception in the presence of outliers. Our first contribution is to show that estimation using common robust costs, such as truncated least squares (TLS), maximum consensus, Geman-McClure, Tukey's biweight, among others, can be reformulated as *polynomial optimization problems* (POPs). By focusing on the TLS cost, our second contribution is to exploit *sparsity* in the POP and propose a sparse *semidefinite programming* (SDP) relaxation that is much smaller than the standard Lasserre's hierarchy while preserving *exactness*, *i.e.*, the SDP recovers the optimizer of the nonconvex POP with an *optimality certificate*. Our third contribution is to solve the SDP relaxations at an unprecedented scale and accuracy by presenting STRIDE, a solver that blends *global descent* on the convex SDP with fast *local search* on the nonconvex POP. Our fourth contribution is an evaluation of the proposed framework on six geometric perception problems including single and multiple rotation averaging, point cloud and mesh registration, absolute pose estimation, and category-level object pose and shape estimation. Our experiments demonstrate that (i) our sparse SDP relaxation is exact with up to 60%–90% outliers across applications; (ii) while still being far from real-time, STRIDE is up to 100 times faster than existing SDP solvers on medium-scale problems, and is the only solver that can solve large-scale SDPs with hundreds of thousands of constraints to high accuracy; (iii) STRIDE provides a safeguard to existing fast heuristics for robust estimation (*e.g.*, RANSAC or Graduated Non-Convexity), *i.e.*, it certifies global optimality if the heuristic estimates are optimal, or detects and allows escaping local optima when the heuristic estimates are suboptimal.

Index Terms—certifiable algorithms, outlier-robust estimation, robust fitting, robust estimation, polynomial optimization, semidefinite programming, global optimization, moment/sums-of-squares relaxation, large-scale convex optimization