ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Computer Vision
and Geometry Lab



Visual Mapping and Pose Estimation for Self-Driving Cars

Marc Pollefeys

8 July 2014

slides courtesy of Gim Hee Lee

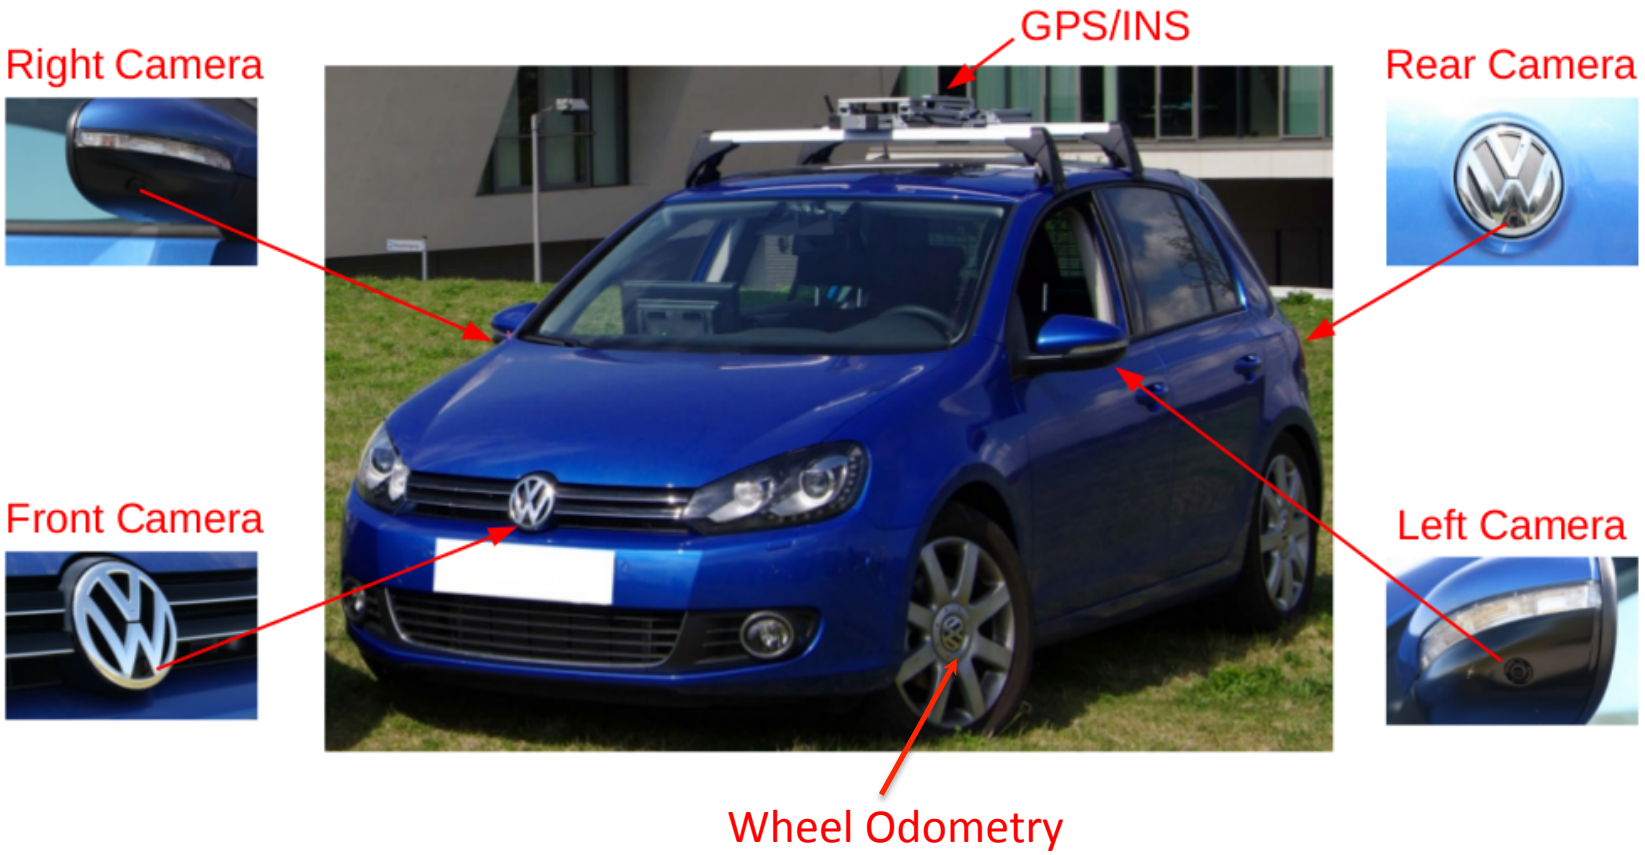
Turning Mobile Phones into 3D-Scanners





Our Self-Driving Car - **Grobi**





Toward Automated Driving in Cities using Close-to-Market Sensors, an Overview of the V-Charge Project
Paul Furgale, ... Gim Hee Lee, ... Marc Pollefeys and Roland Siegwart
IEEE Intelligent Vehicle Symposium (IV) 2013

Why Multi-Camera System?

- Cameras are low-cost and easy to maintain.
- Aesthetic: Car appearance not changed.
- Configured be chosen to maximize field-of-view.
- Absolute scale directly from the epipolar geometry.

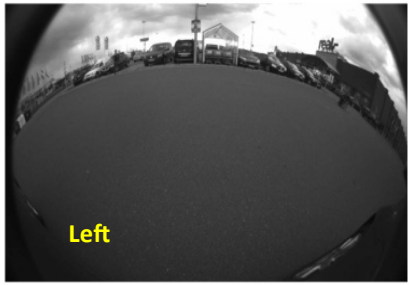










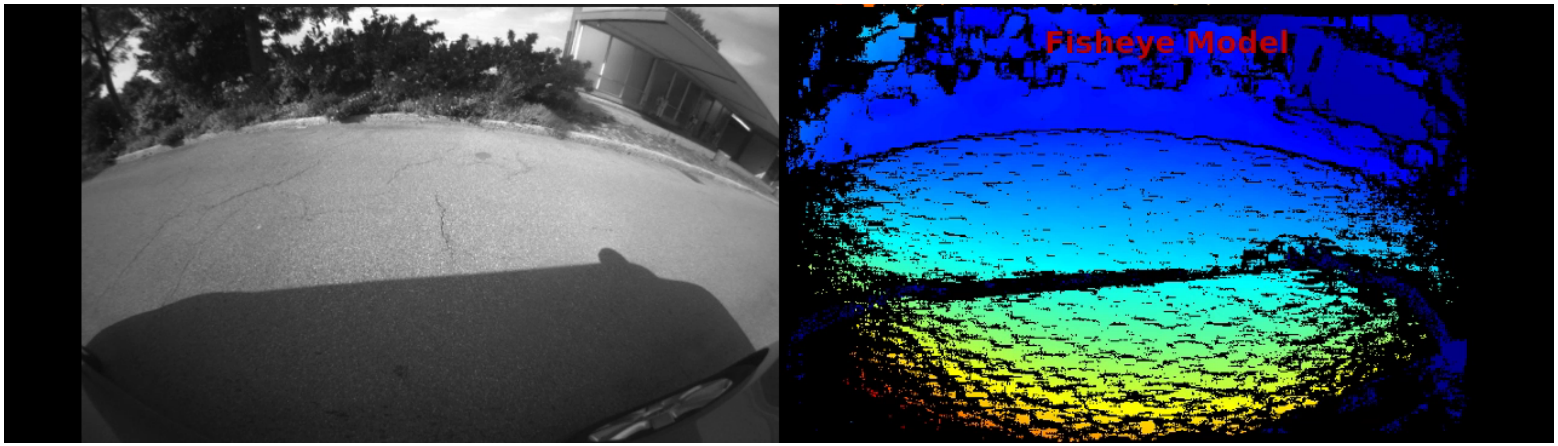






Monocular fish-eye video depth

(Haene et al.)



Provides depth on side of vehicle (e.g. parking spots)
Complementary to stereo forward/backward

Similar technology explored for Project Tango

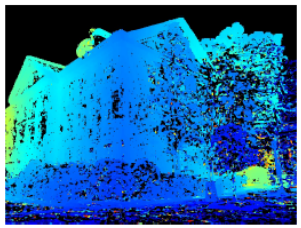


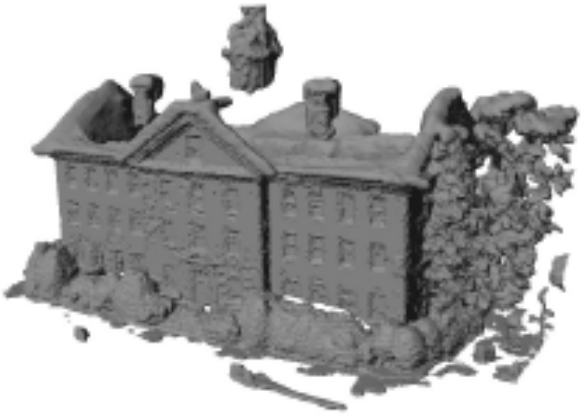


Joint 3D reconstruction and class segmentation

(Haene et al CVPR13)

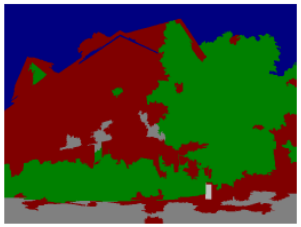


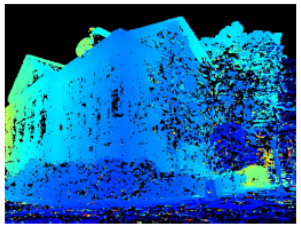


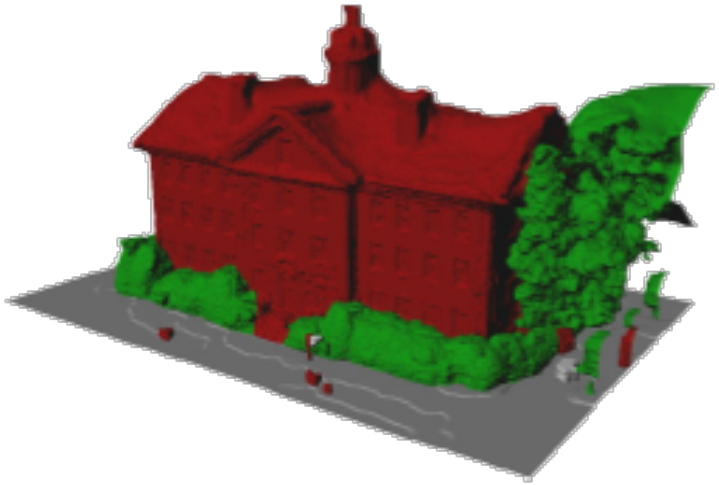


reconstruction only
(isotropic smoothness prior)

joint reconstruction and segmentation
(ground, building, vegetation, stuff)

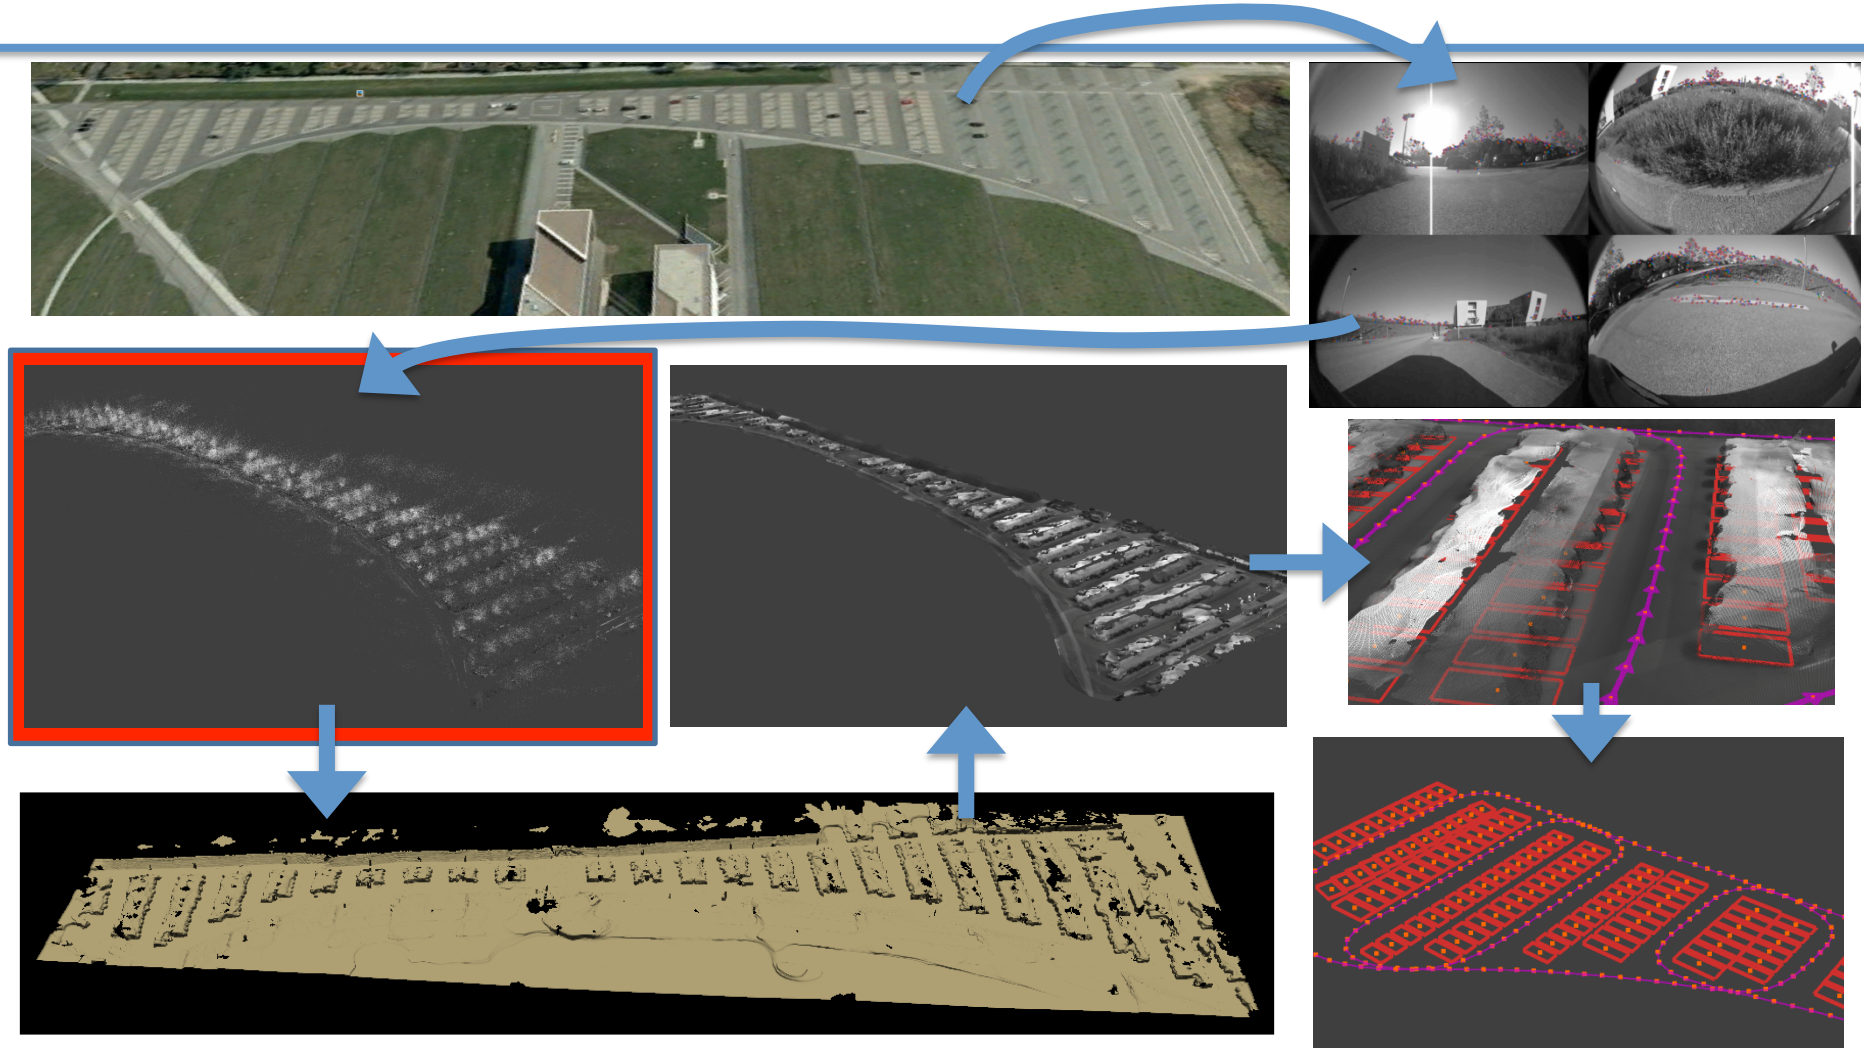




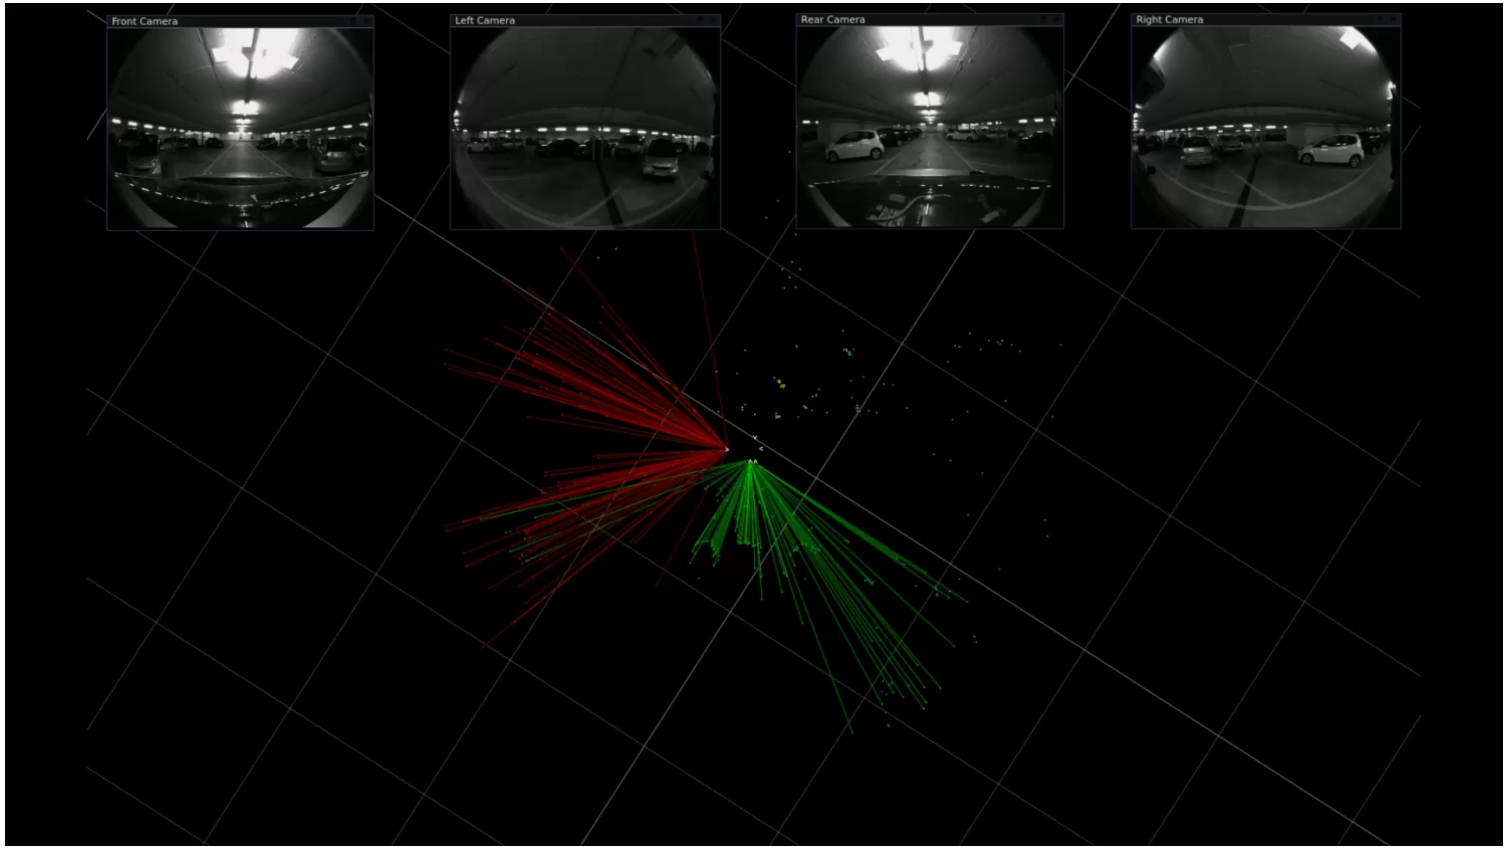


- Building
- Ground
- Vegetation
- Clutter

Workflow: Offline Mapping

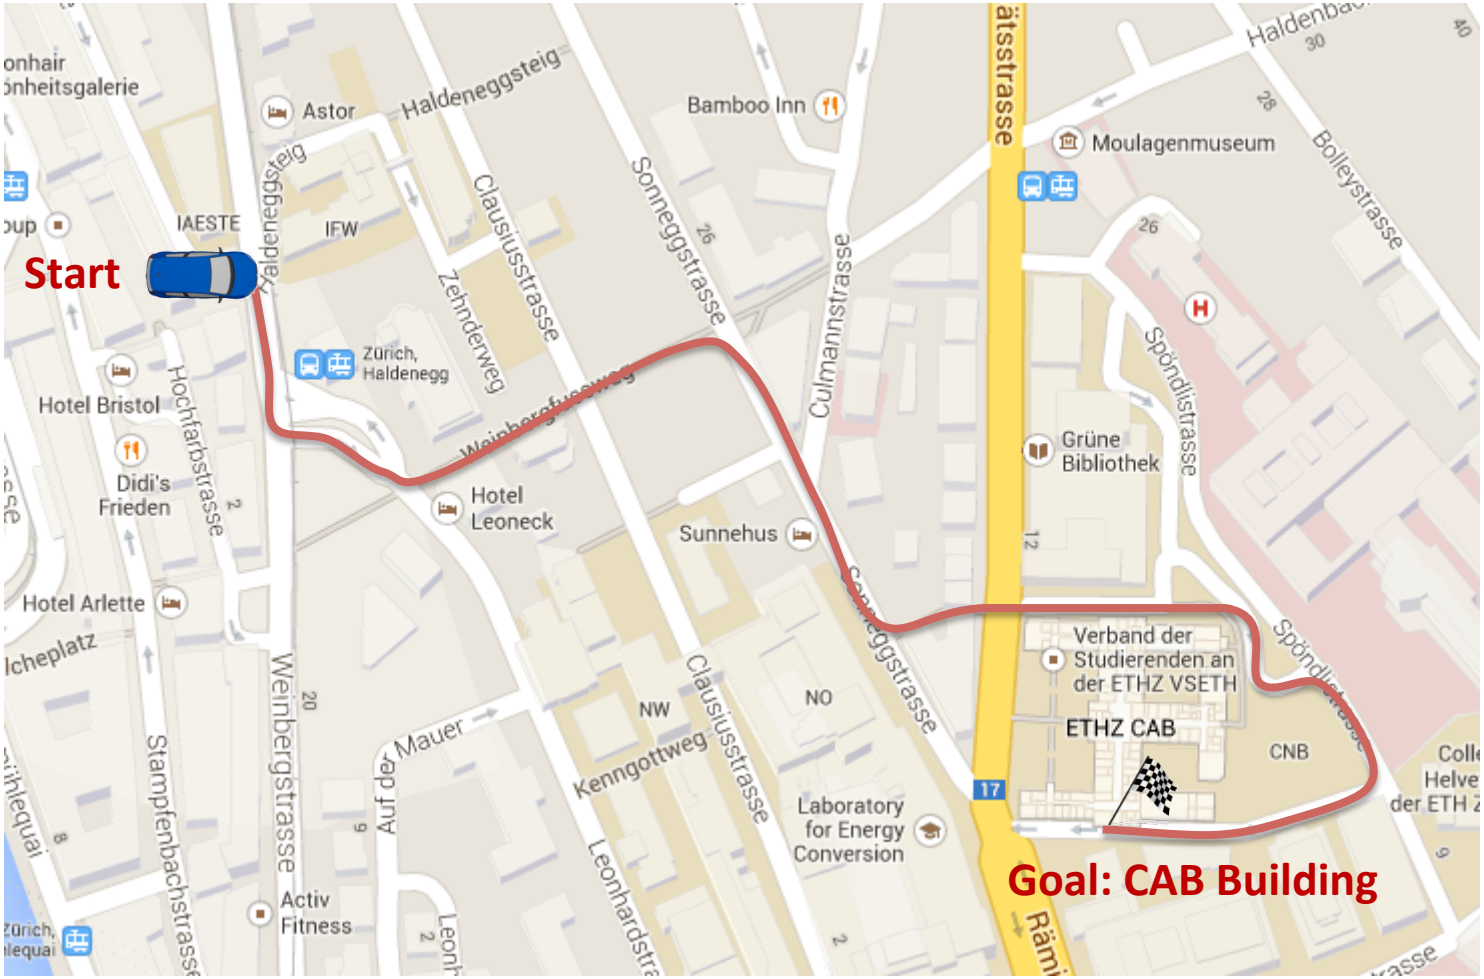


System re-calibration

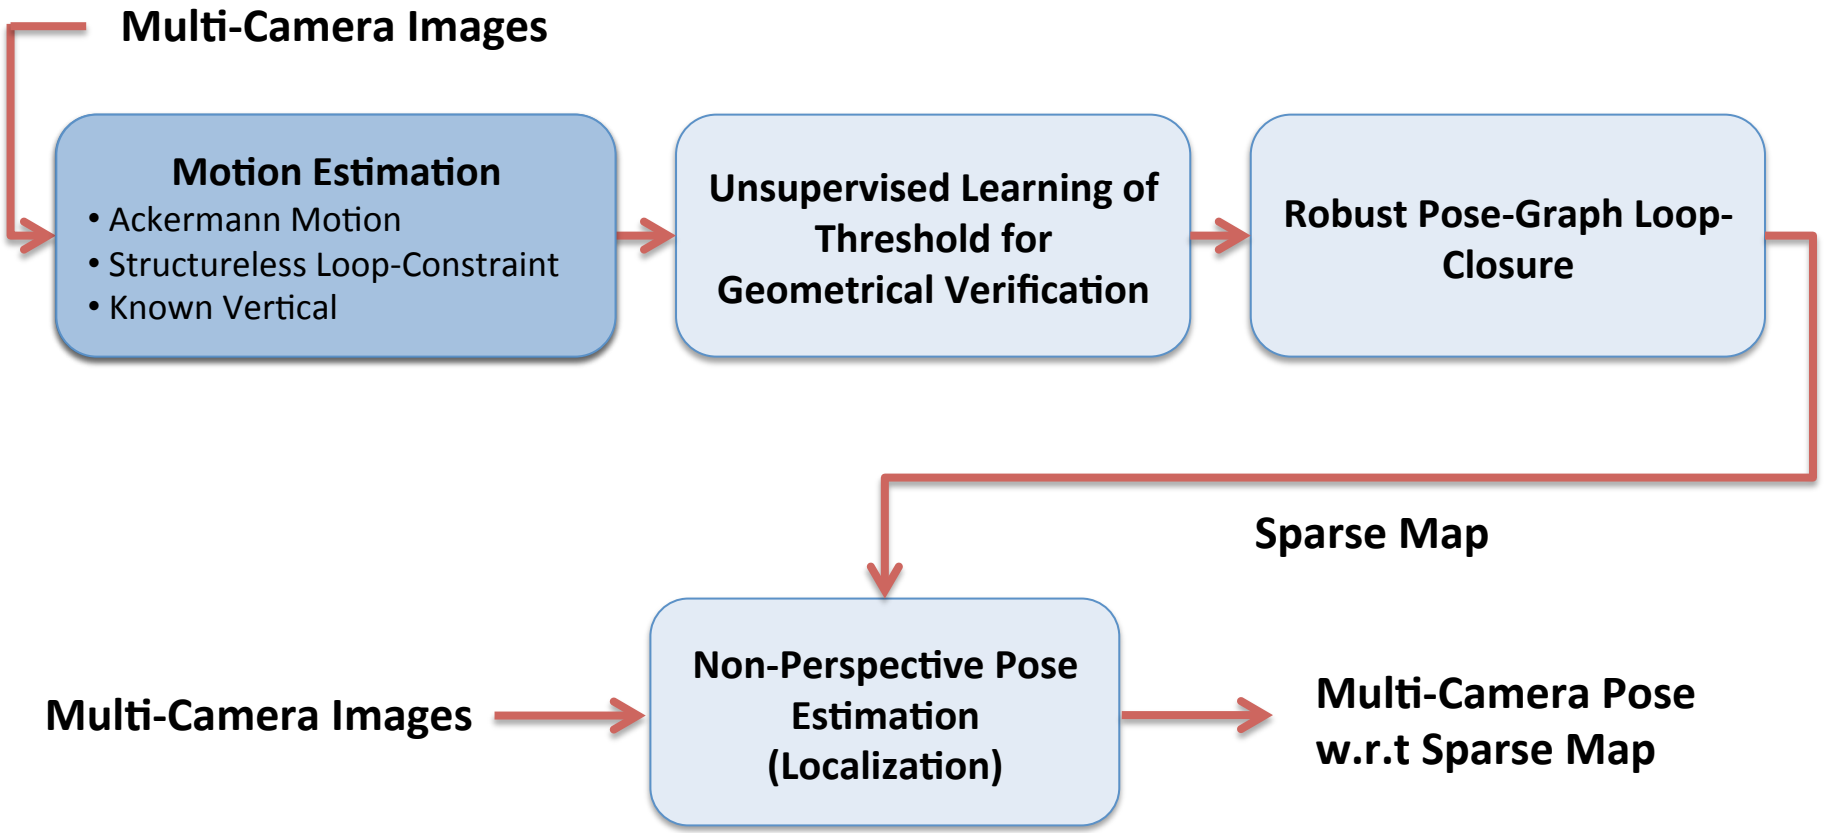


Infrastructure-Based Calibration of a Multi-Camera Rig
Lionel Heng, Mathias Buerki, Gim Hee Lee, Paul Furgale, Roland Siegwart, and Marc Pollefeys
IEEE International Conference on Robotics and Automation (ICRA), 2014.

Why Mapping and Pose Estimation?

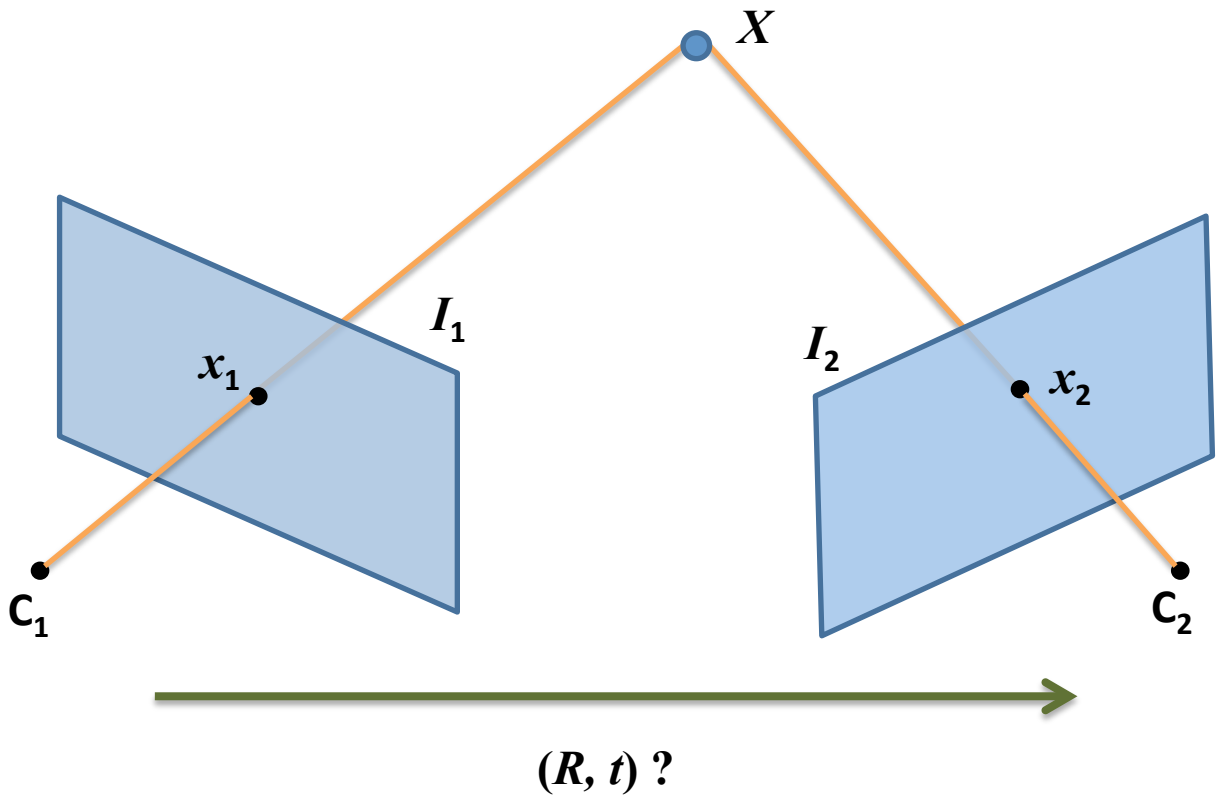


Mapping and Pose Estimation: Overview



Monocular Epipolar Geometry

Given: 2D point correspondences, $x_1 \leftrightarrow x_2$.
Find: R and t that relates the two frames.



Essential Matrix

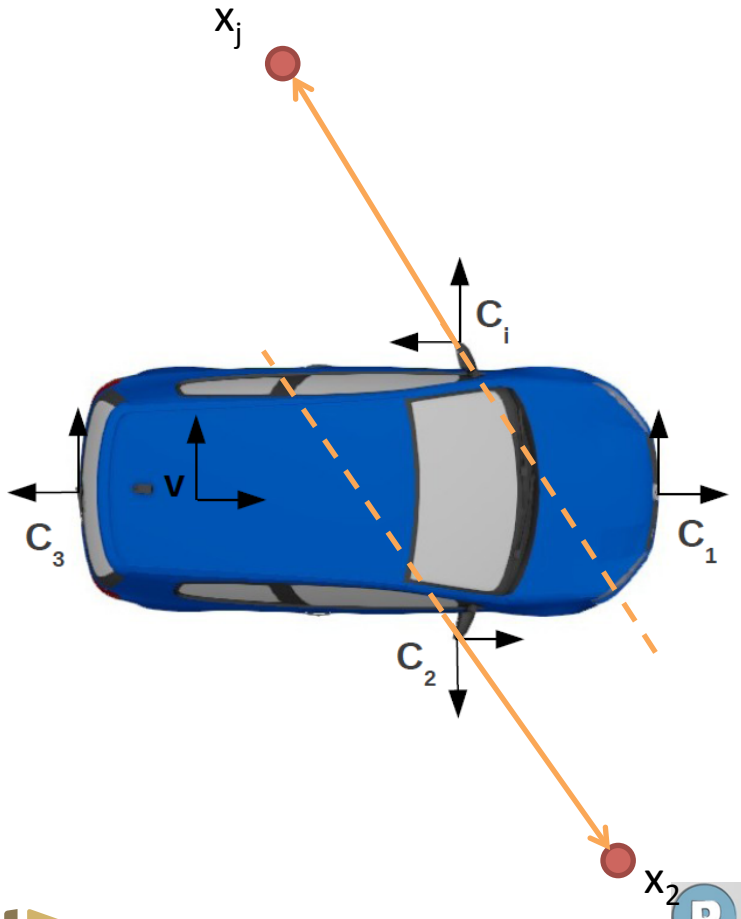
$$x_1^T E x_2 = 0$$

$$E = [t]_x R$$

$$Ae = 0$$

Plücker Line

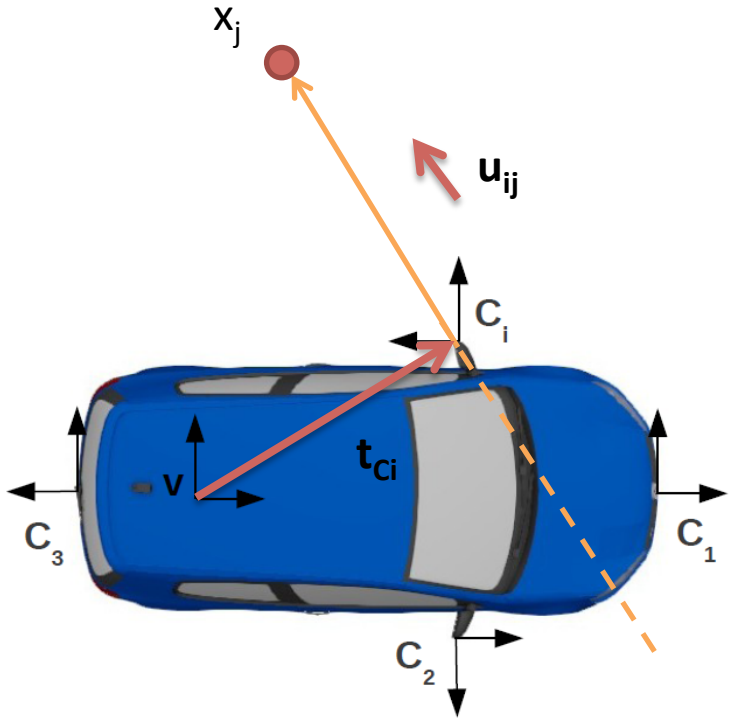
- Light rays do meet at a single center of projection.



Plücker Line

- 6-vector Plücker line to represent the light rays.

$$\mathbf{l}_{ij} = [\mathbf{u}_{ij}^T,\ (t_{C_i} \times \mathbf{u}_{ij})^T]^T$$



V : Reference frame

u_{ij} : Unit direction of ray

T_{ci} : Translation from C_i to V

Generalized Epipolar Constraint

- Generalized Epipolar constraint:

$$\mathbf{l}_{ij}'^{T} \underbrace{\begin{bmatrix} E & R \\ R & 0 \end{bmatrix}}_{E_{GC}} \mathbf{l}_{ij} = 0$$

- $E = [R]_x t$: conventional Essential matrix.
- $[R,\ t]$ is the relative transformation.
- Absolute scale can be obtained!

Using many cameras as one
Robert Pless
IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2003

Generalized Epipolar Constraint

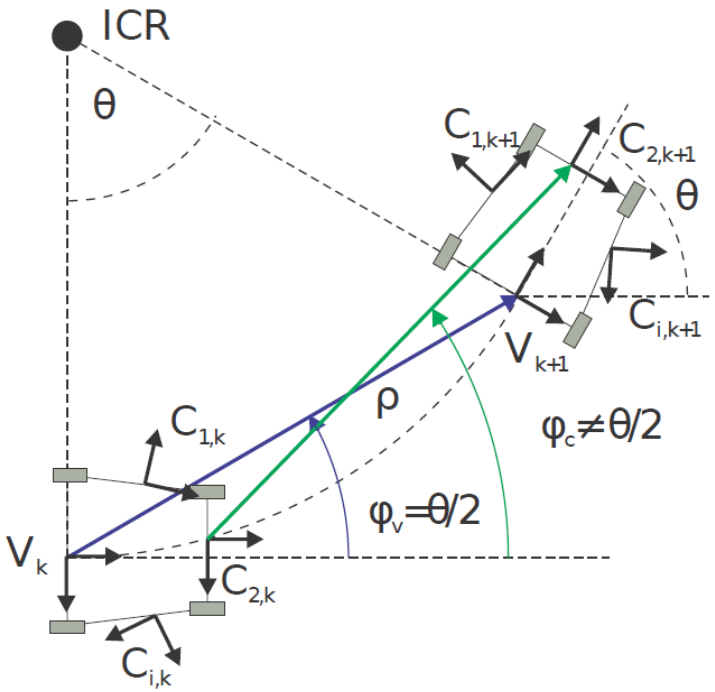
- Problems:
  - Linear solution requires 17-point correspondences.

    → ~600k RANSAC loops needed.
  - Minimal problem needs 6-point correspondences but gives 64 solutions.

Motion Estimation: Ackermann Constraint

- Enforce Ackermann motion constraint:

$$R = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad t = \rho \begin{bmatrix} \cos\varphi_v \\ \sin\varphi_v \\ 0 \end{bmatrix}$$



- 2 degree-of-freedom

  ⟶ 2-point minimal problem.

17 RANSAC Loops!

Motion Estimation for a Self-Driving Car with a Generalized Camera
Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys
IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2013

Motion Estimation: Ackermann Constraint

- Generalized Essential matrix with Ackermann constraint:

$$E_{GC} = \begin{bmatrix} 0 & 0 & \rho \sin\frac{\theta}{2} & \cos\theta & -\sin\theta & 0 \\ 0 & 0 & -\rho\cos\frac{\theta}{2} & \sin\theta & \cos\theta & 0 \\ \rho\sin\frac{\theta}{2} & \rho\cos\frac{\theta}{2} & 0 & 0 & 0 & 1 \\ \cos\theta & -\sin\theta & 0 & 0 & 0 & 0 \\ \sin\theta & \cos\theta & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

- Solve for ρ and θ from the Epipolar constraint.

  → 2 polynomial equations with 2 unknowns.

$$\begin{cases} a_1 \cos\theta + b_1 \sin\theta + c_1 \rho \cos\dfrac{\theta}{2} + e_1 = 0 \\ a_2 \cos\theta + b_2 \sin\theta + c_2 \rho \cos\dfrac{\theta}{2} + e_2 = 0 \end{cases}$$

Motion Estimation: Ackermann Constraint

- Can be solved in closed-form with $\gamma = \sin^2 \frac{\theta}{2}$.

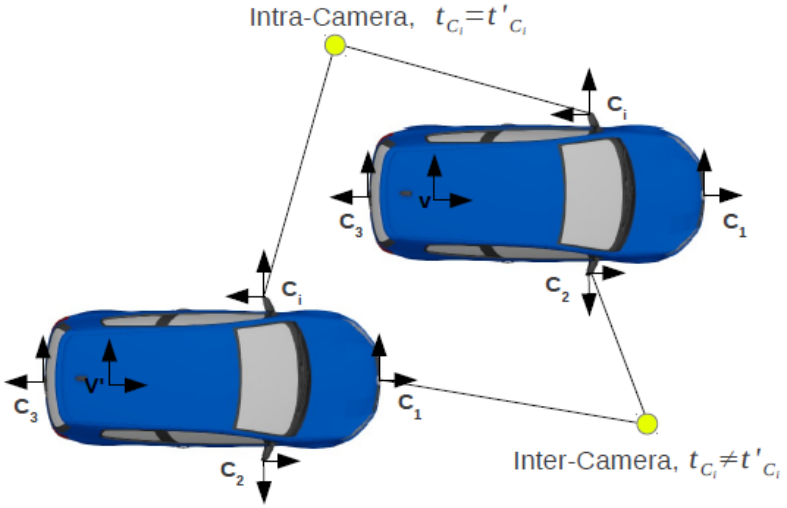
$$A\gamma^3 + B\gamma^2 + C\gamma + D = 0$$

- Up to 6 real solutions.

Degenerated Case

- Pure translation $R = I$ and intra-camera $t_{Ci} = t'_{Ci}$ correspondences:

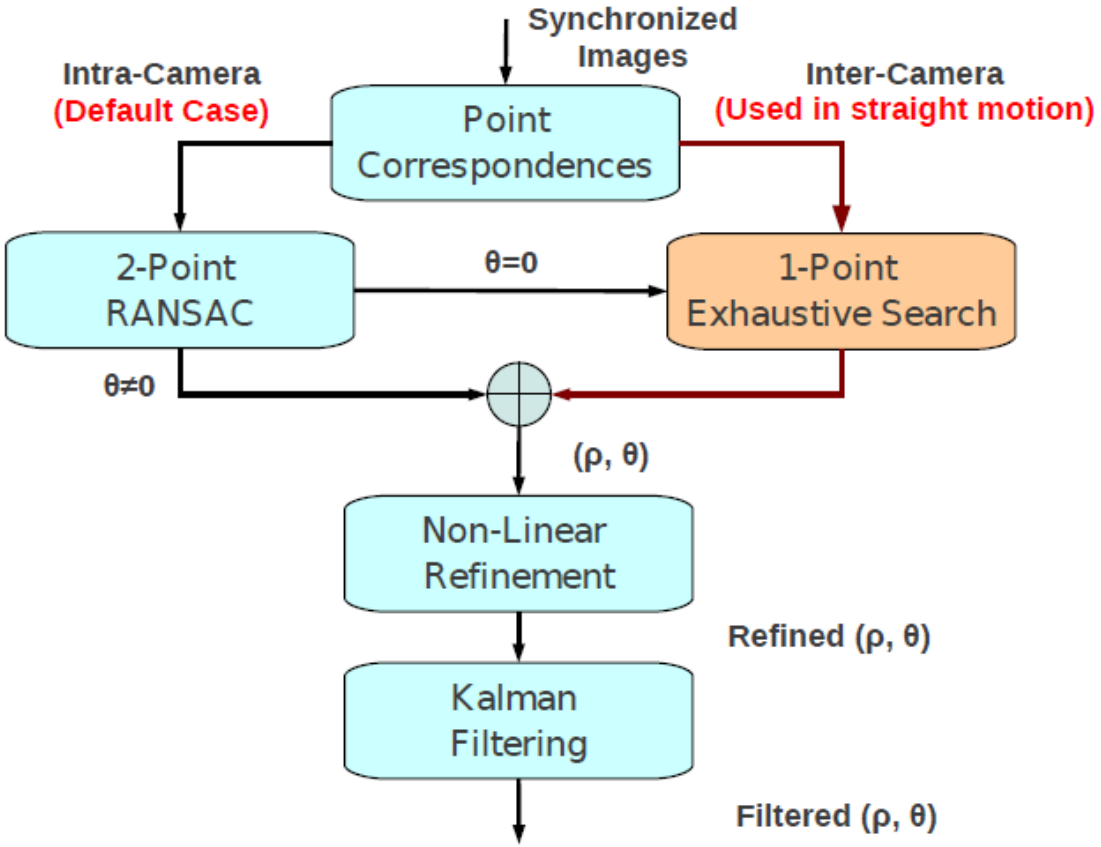
$$\mathbf{l}_{ij}'^{T} \underbrace{\begin{bmatrix} E & R \\ R & 0 \end{bmatrix}}_{E_{GC}} \mathbf{l}_{ij} = 0 \rightarrow u_{ij}'^{T} E u_{ij} + \cancel{(t'_{C_i} \times u'_{ij})^{T} R u_{ij} + u_{ij}'^{T} R (t_{C_i} \times u_{ij})} = 0$$



Solution:
Use 1 additional inter-camera correspondence!

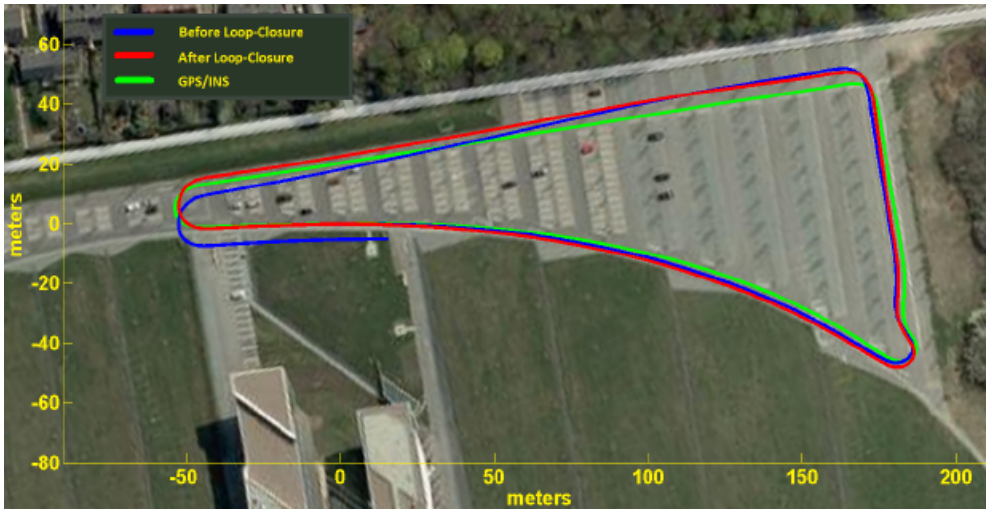
Scale cannot be solved!

Motion Estimation: Ackermann Constraint

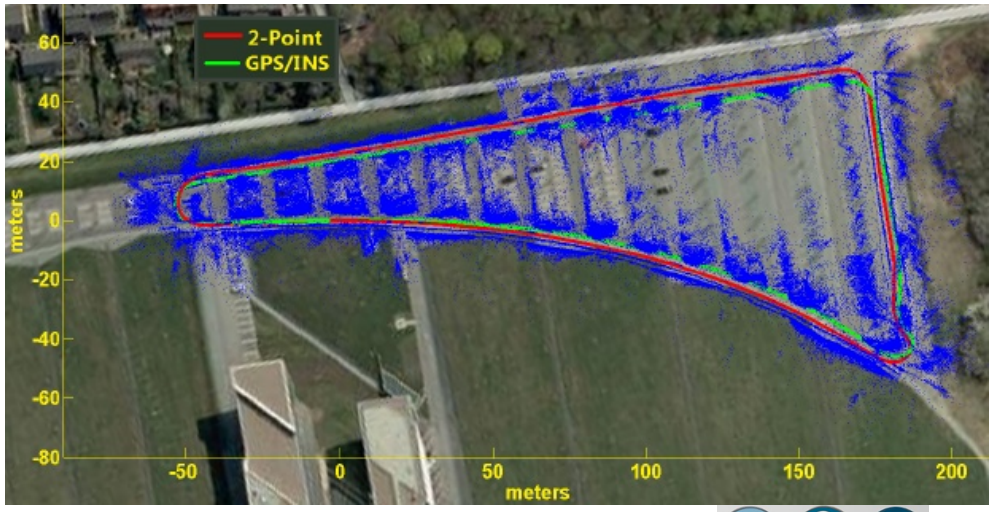


Results

Pose-graph before and after loop-closure.

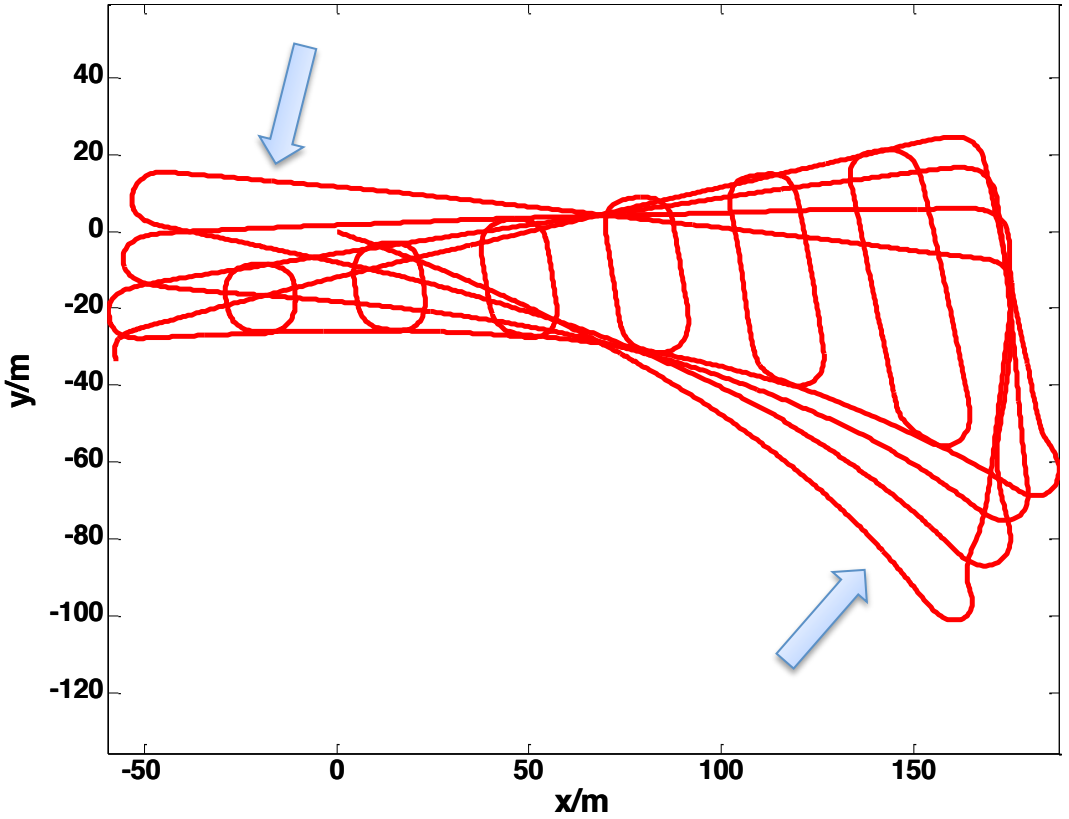


Final map with 3D points after full bundle adjustment.



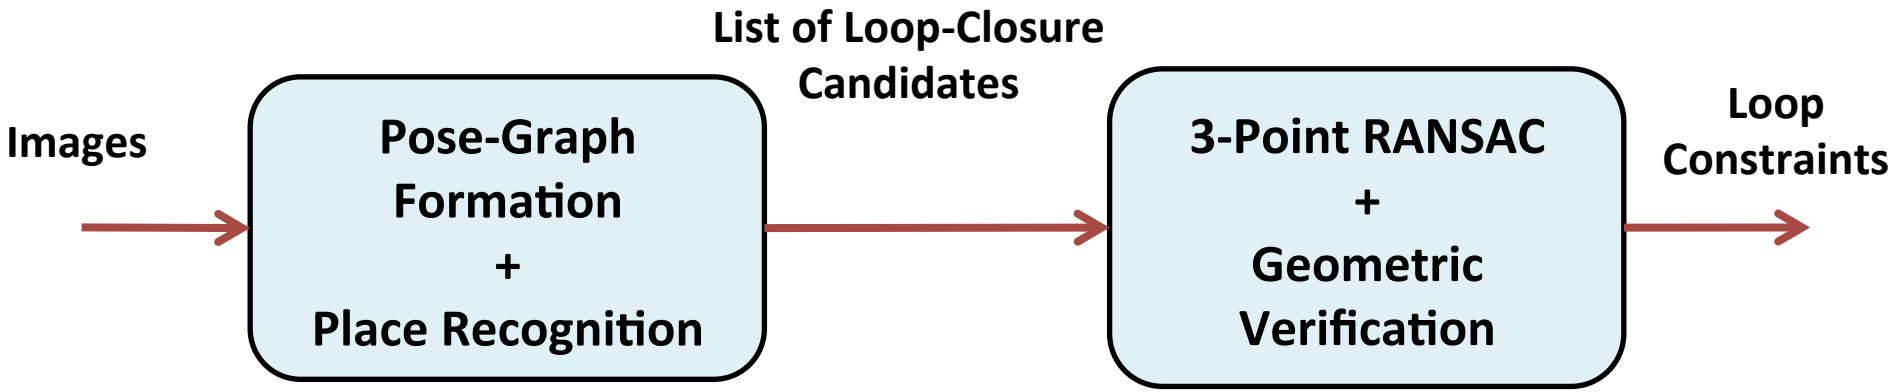
Motion Estimation: Structureless Loop-Constraints

- Drifts from wheel/visual odometry.
- Find the loop-closure constraints with our multi-camera system.
- Absolute scale directly from the epipolar geometry → structureless!



Structureless Pose-Graph Loop-Closure with a Multi-Camera System on a Self-Driving Car
Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2013

Structureless Loop-Constraints: Framework



Our 3-Point Algorithm

- Relax Ackermann to planar constraint:

$$R = \frac{1}{1+q^2}\begin{bmatrix} 1-q^2 & -2q & 0 \\ 2q & 1-q^2 & 0 \\ 0 & 0 & 1+q^2 \end{bmatrix}, \quad t = \begin{bmatrix} x \\ y \\ 0 \end{bmatrix}$$

- 3 degrees-of-freedom ⟶ 3-point minimal problem.

Only 34 RANSAC loops!

Our 3-Point Algorithm

- Generalized Essential matrix with planar constraint:

$$E_{GC} = \begin{bmatrix} 0 & 0 & y & \frac{1-q^2}{1+q^2} & \frac{-2q}{1+q^2} & 0 \\ 0 & 0 & -x & \frac{2q}{1+q^2} & \frac{1-q^2}{1+q^2} & 0 \\ \frac{2xq-y(1-q^2)}{1+q^2} & \frac{2yq+x(1-q^2)}{1+q^2} & 0 & 0 & 0 & 1 \\ \frac{1-q^2}{1+q^2} & \frac{-2q}{1+q^2} & 0 & 0 & 0 & 0 \\ \frac{2q}{1+q^2} & \frac{1-q^2}{1+q^2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

- Solve for x, y and q from the Epipolar constraint.

  $\longrightarrow$ 3 polynomial equations with 3 unknowns.

$$\begin{cases} a_1xq^2 + a_2xq + a_3x + a_4yq^2 + a_5yq + a_6y + a_7q^2 + a_8q + a_9 = 0 \\ b_1xq^2 + b_2xq + b_3x + b_4yq^2 + b_5yq + b_6y + b_7q^2 + b_8q + b_9 = 0 \\ c_1xq^2 + c_2xq + c_3x + c_4yq^2 + c_5yq + c_6y + c_7q^2 + c_8q + c_9 = 0 \end{cases}$$

Hidden Variable Resultant

- "Hide" the variable q to form:

$$\beta(q)\mathbf{X} = 0 \qquad \text{where, } \mathbf{X} = \begin{bmatrix} x & y & 1 \end{bmatrix}^{\mathbf{T}}$$

- From linear algebra, we have $det(\beta(q)) = 0$.
- Univariate 6-degree polynomial:

$$Aq^6 + Bq^5 + Cq^4 + Dq^3 + Eq^2 + Fq + G = 0$$

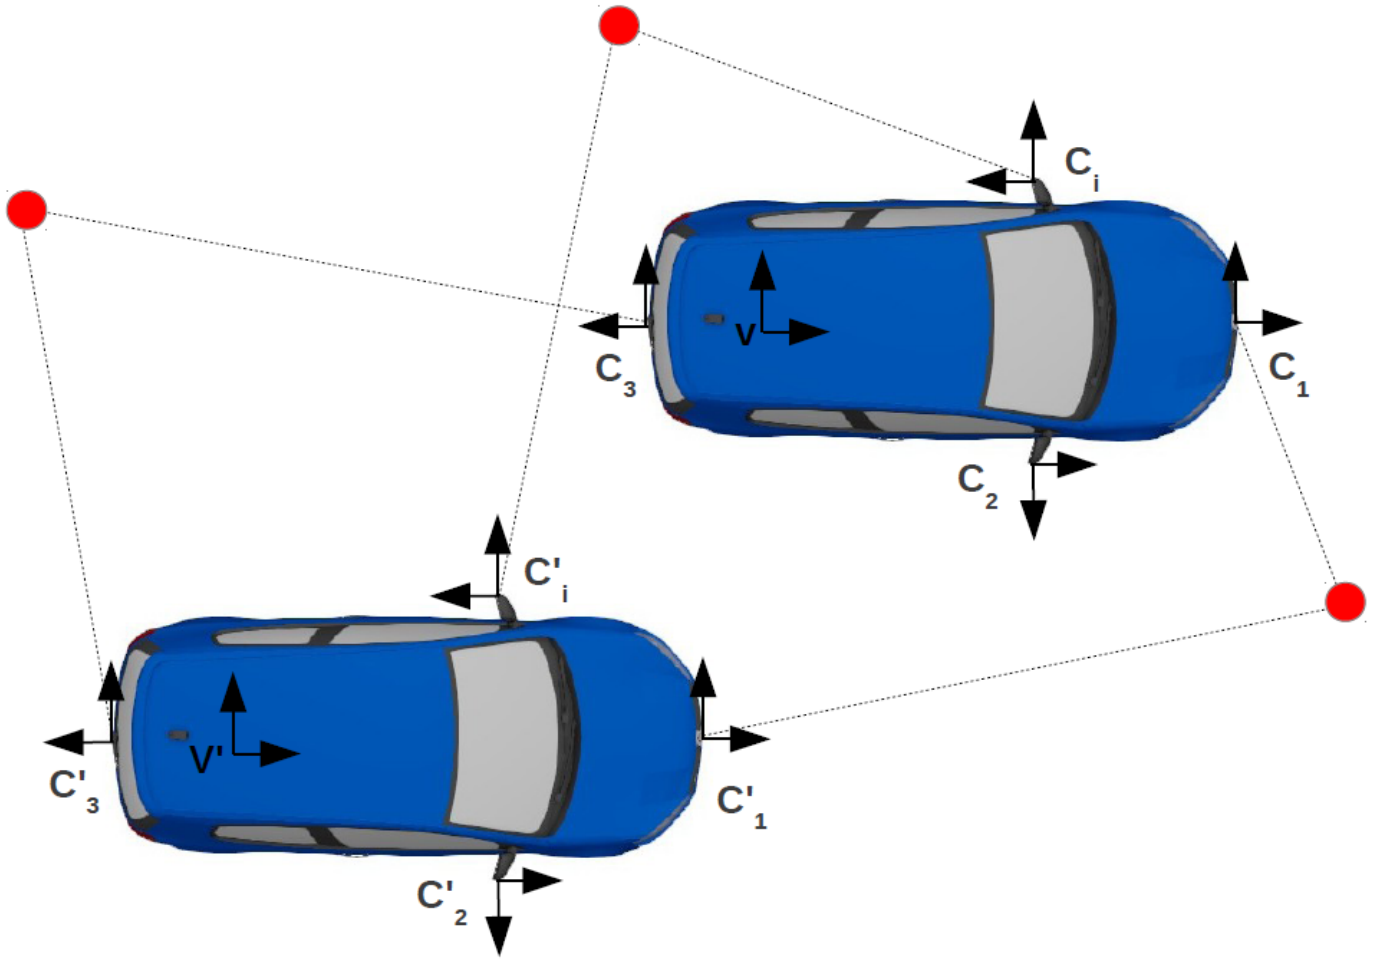
Solving the 6-Degree Polynomial

- Eigen-values of companion matrix are the roots of 6-degree polynomial, i.e. q.

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & -\frac{G}{A} \\ 1 & 0 & 0 & 0 & 0 & -\frac{F}{A} \\ 0 & 1 & 0 & 0 & 0 & -\frac{E}{A} \\ 0 & 0 & 1 & 0 & 0 & -\frac{D}{A} \\ 0 & 0 & 0 & 1 & 0 & -\frac{C}{A} \\ 0 & 0 & 0 & 0 & 1 & -\frac{B}{A} \end{bmatrix}$$

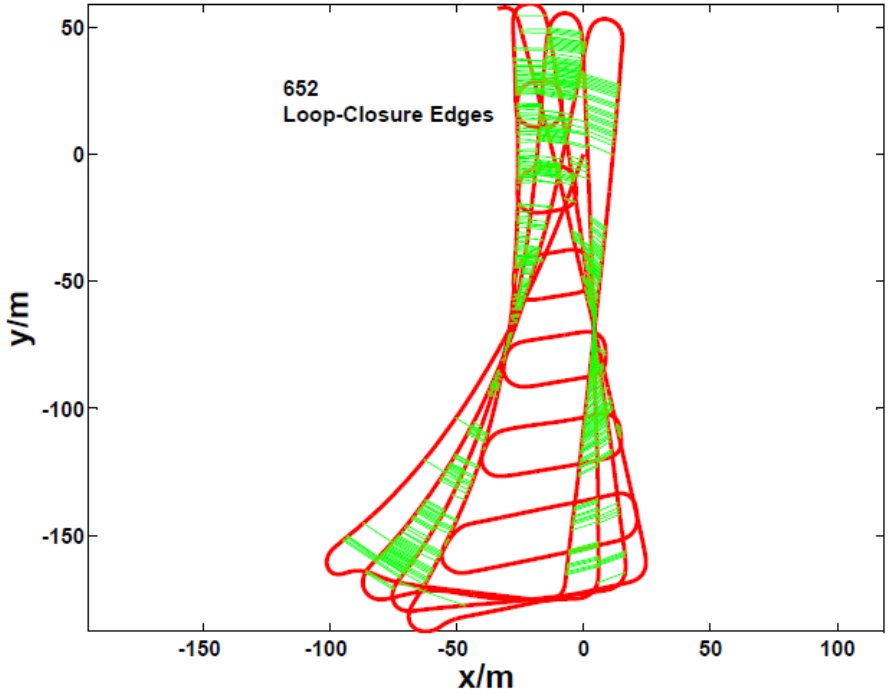
- x and y are computed by back-substitution.

Degenerated Case: No Rotation

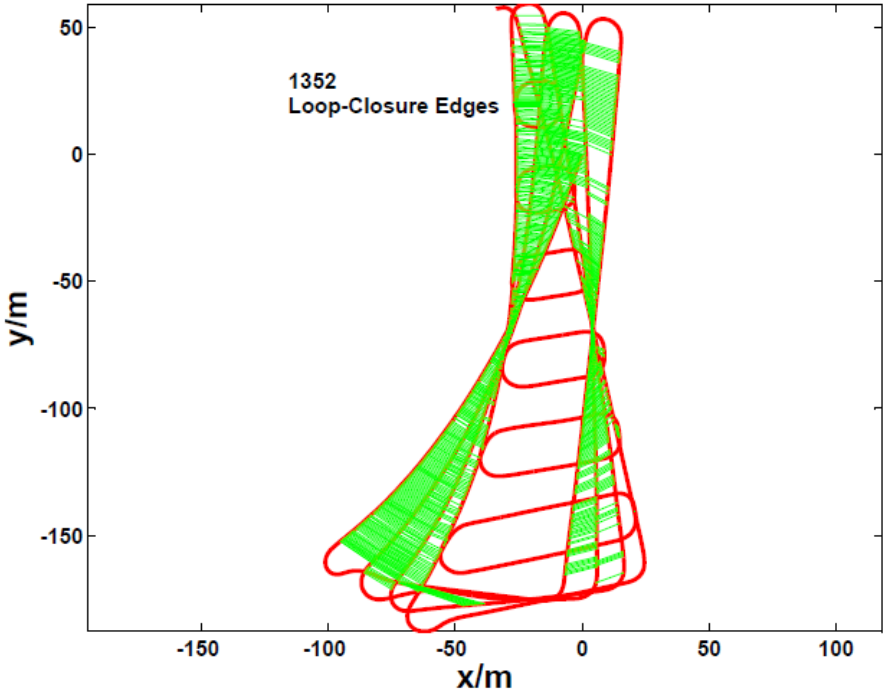


Results: Parking Garage 1

- Wheel odometry

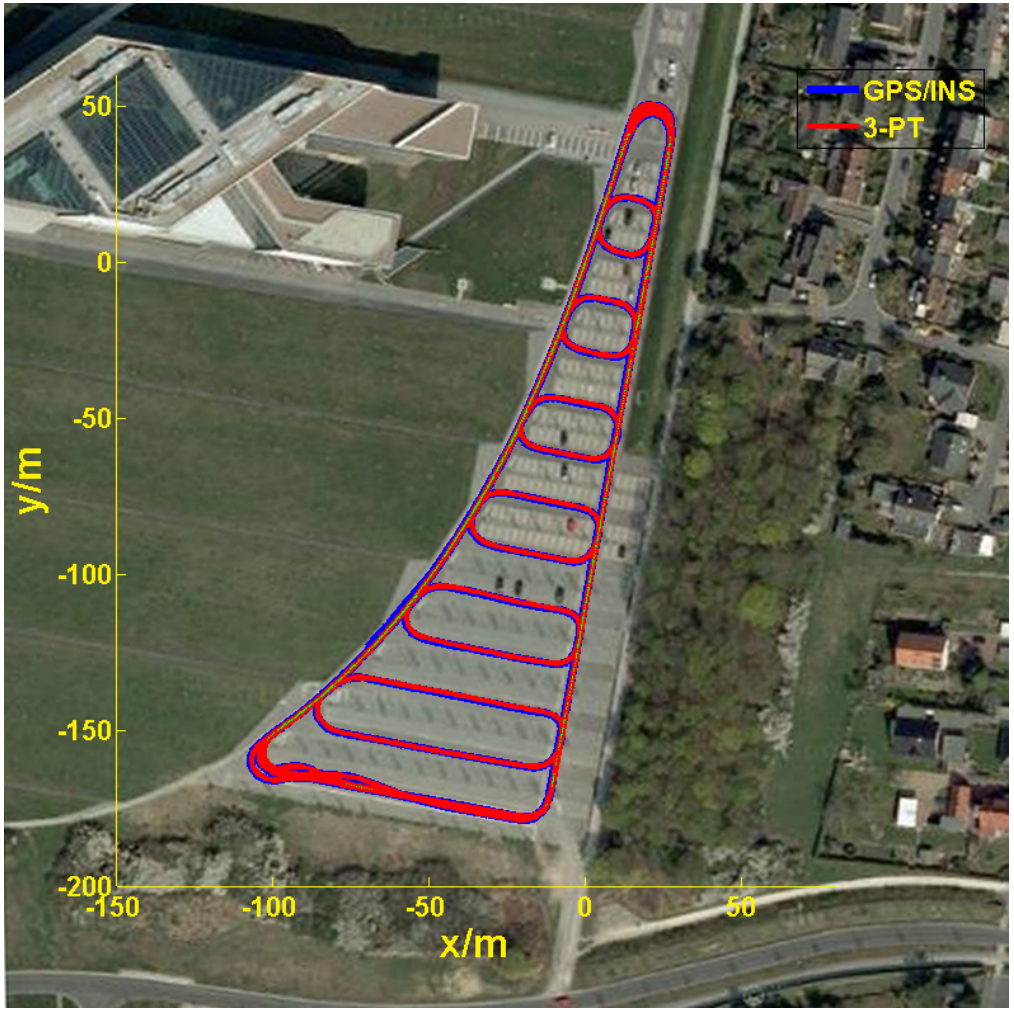


Single front-looking camera



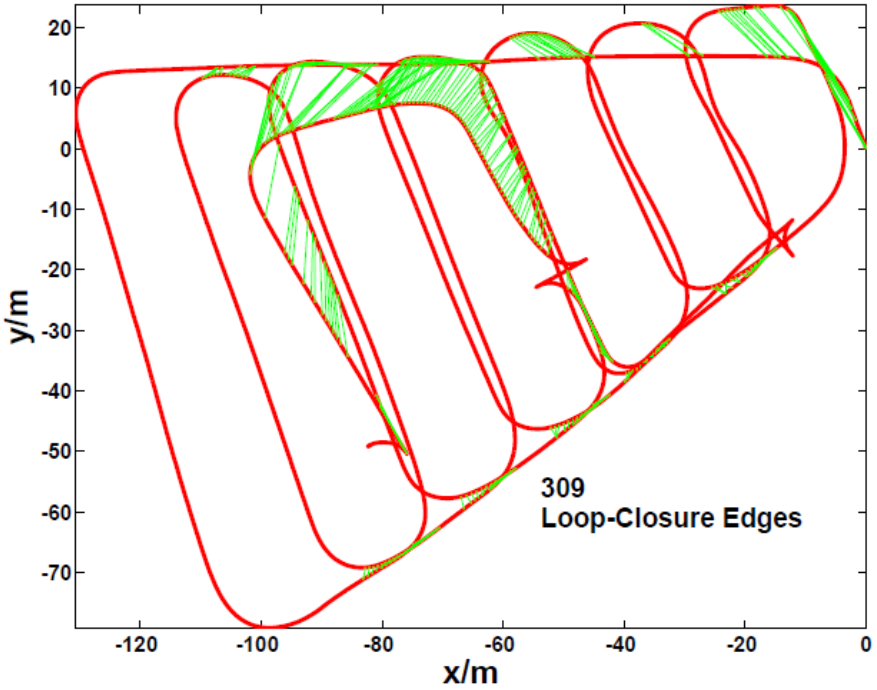
Multi-camera system

Results: Parking Garage 1

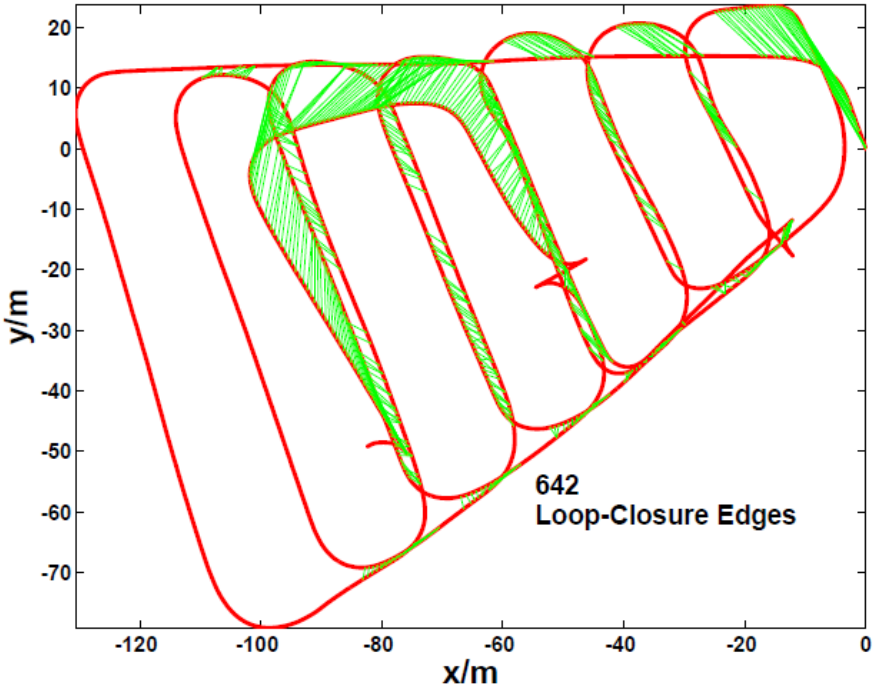


Results: Campus 1

- Wheel Odometry

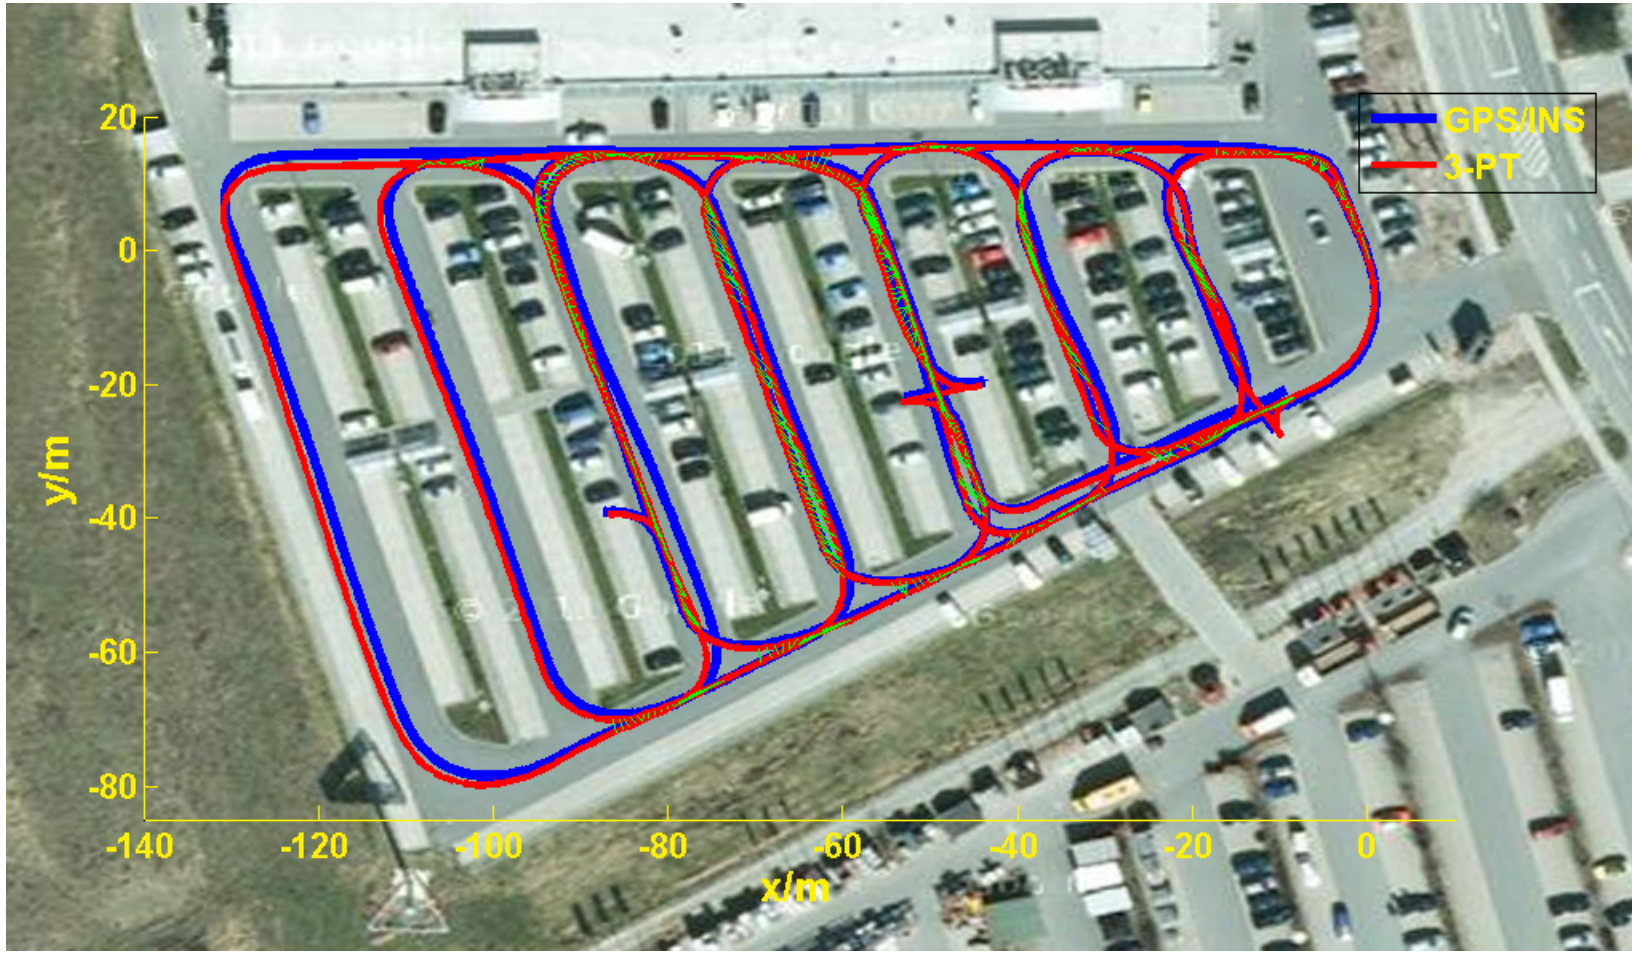


Single front-looking camera



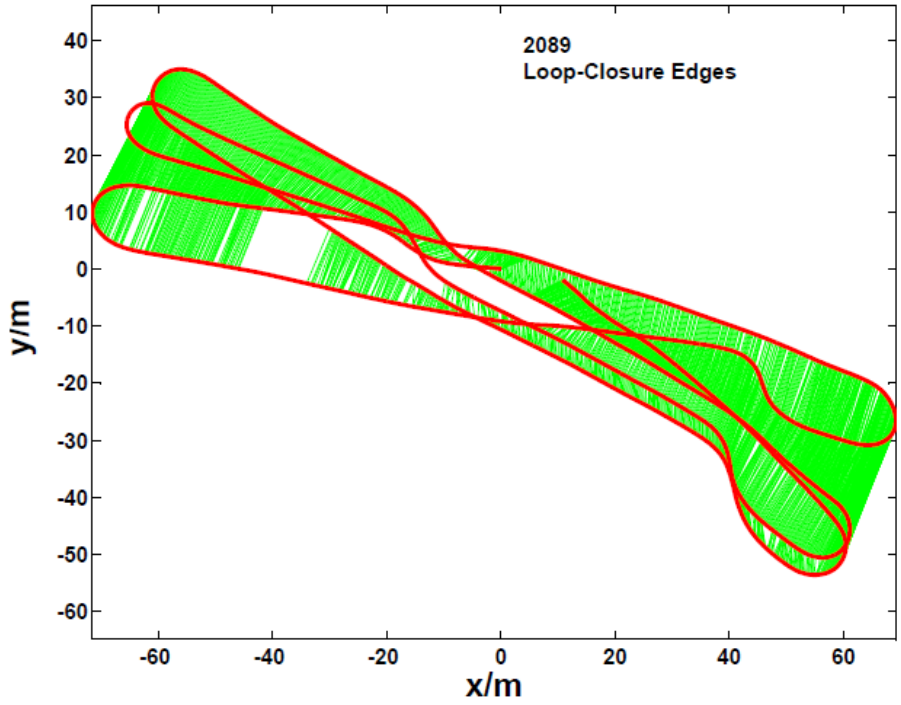
Multi-camera system

Results: Parking Garage 2

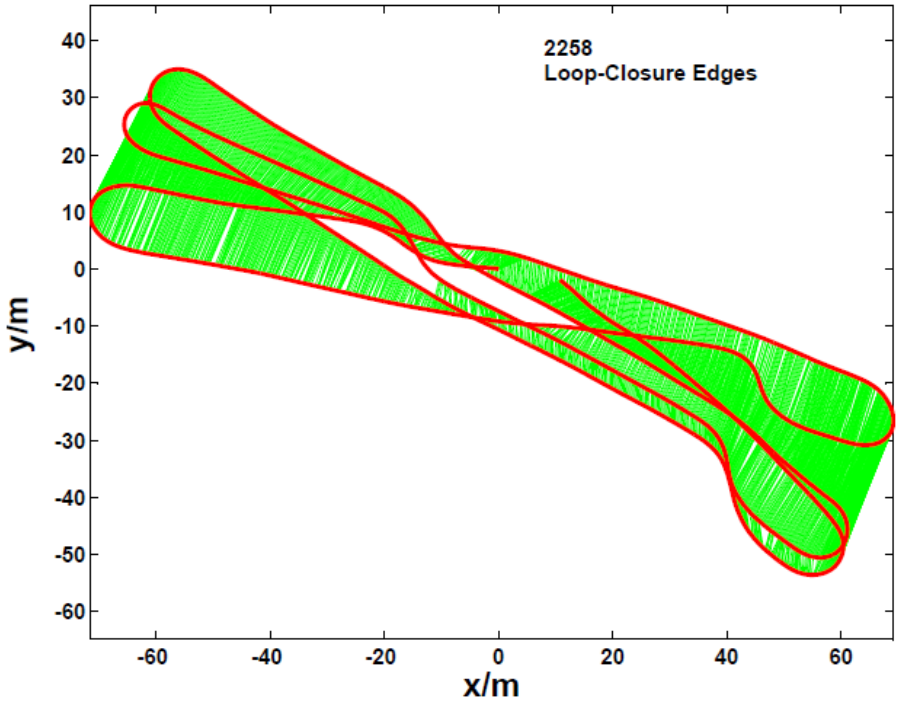


Results: Campus 1

- Visual Odometry

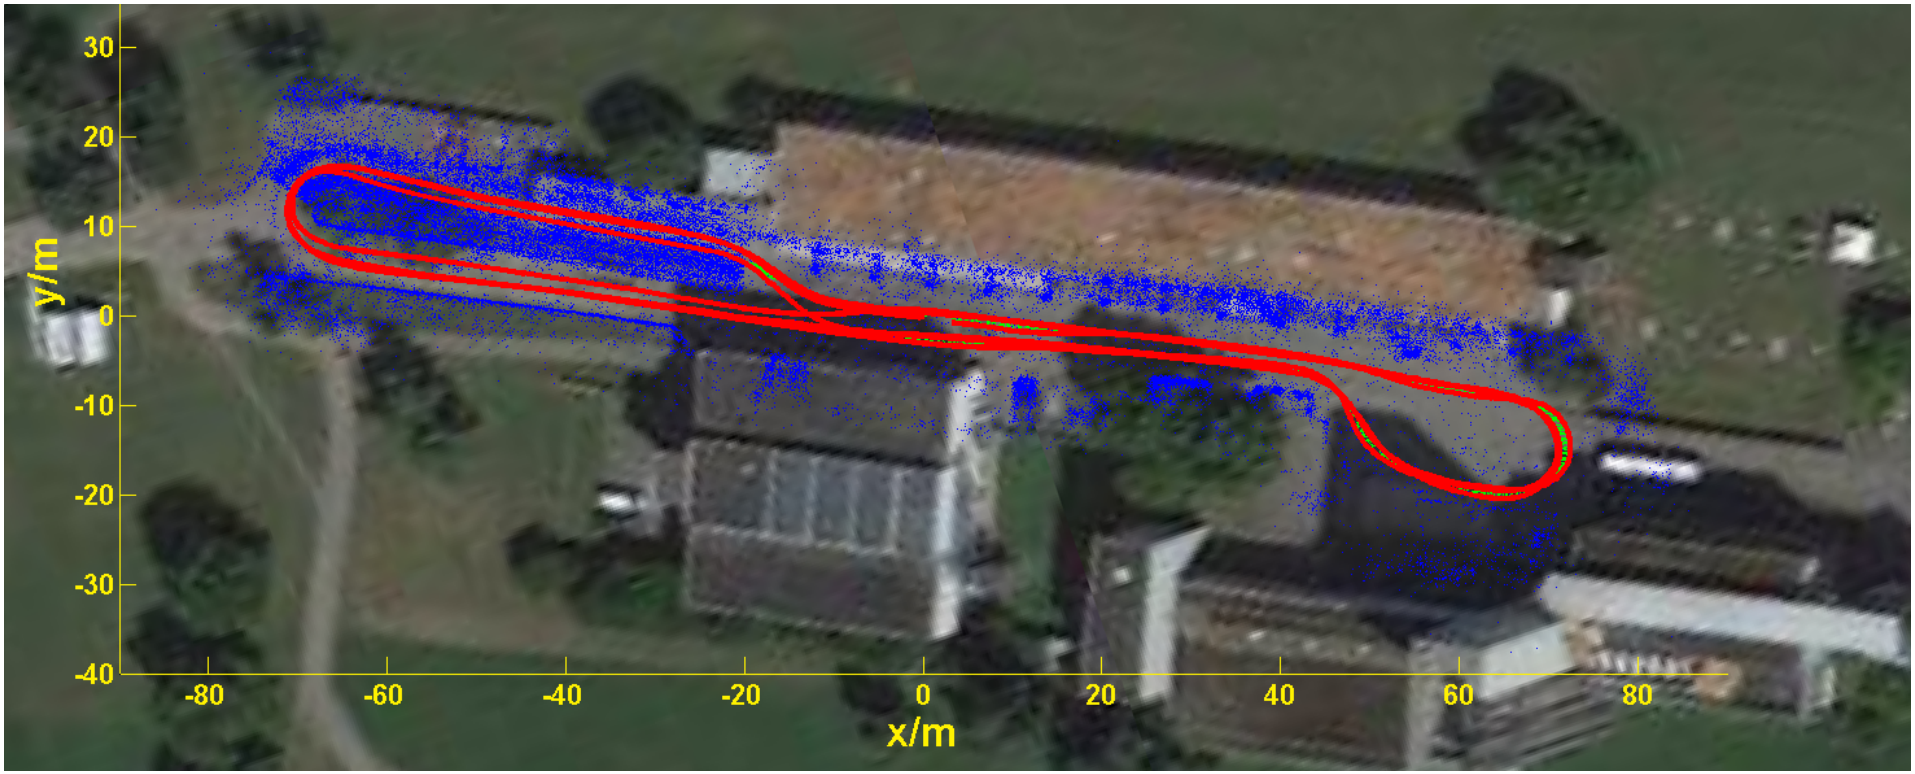


Single front-looking camera

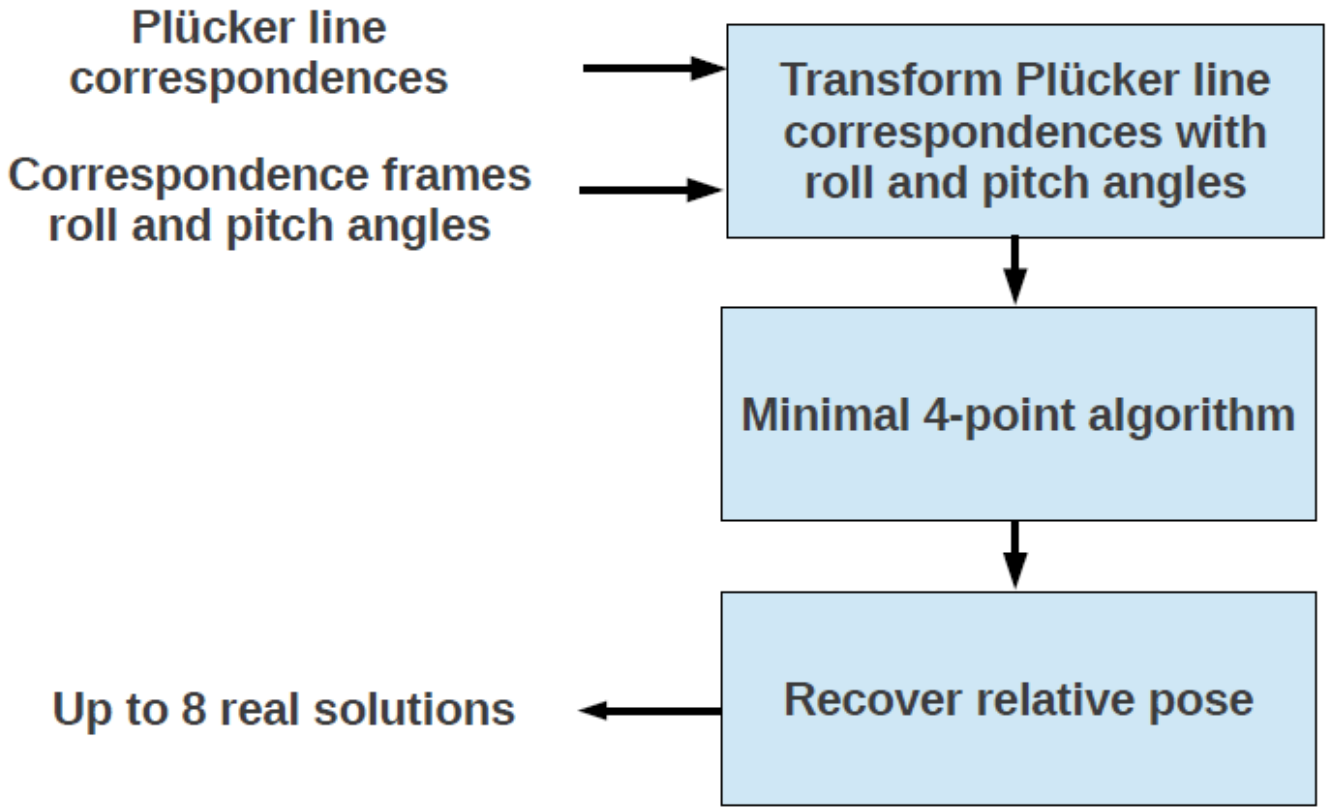


Multi-camera system

Results: Campus 1

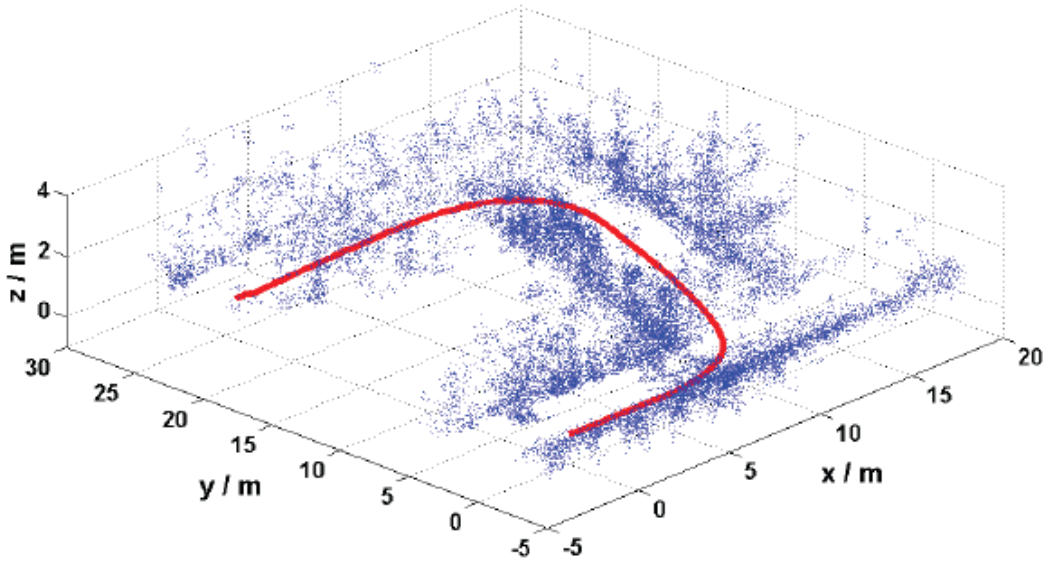


Motion Estimation: Known Vertical

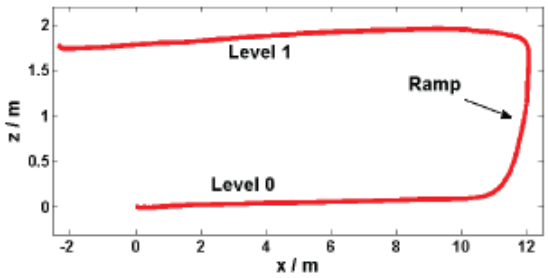


Relative Pose Estimation for a Multi-Camera System with Known Vertical Direction
Gim Hee Lee, Marc Pollefeys, and Friedrich Fraundorfer
Submitted to: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2014

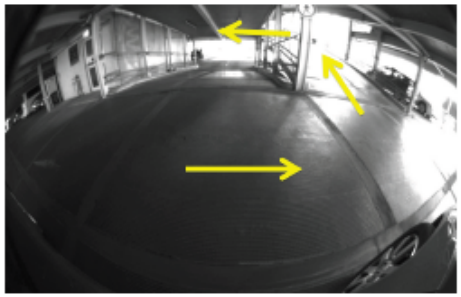
Results: Up a Ramp



(a)

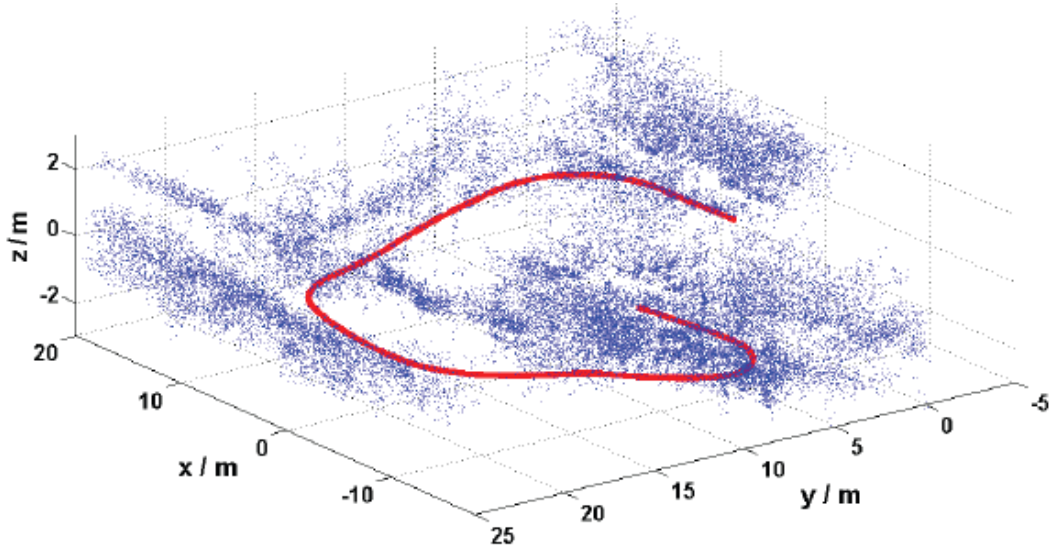


(b)

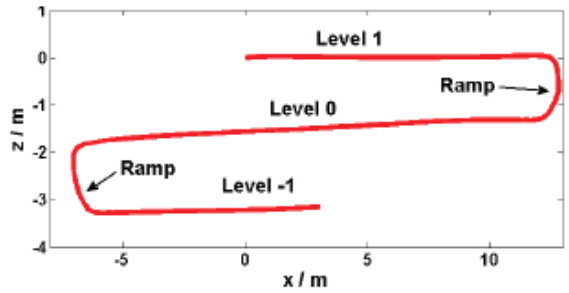


(c)

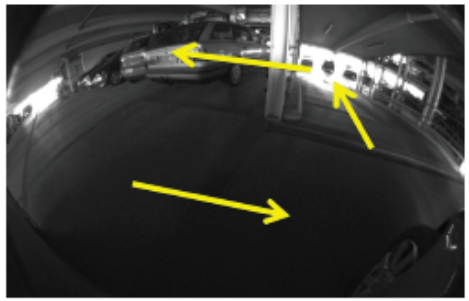
Results: Down 2 Ramps



(a)

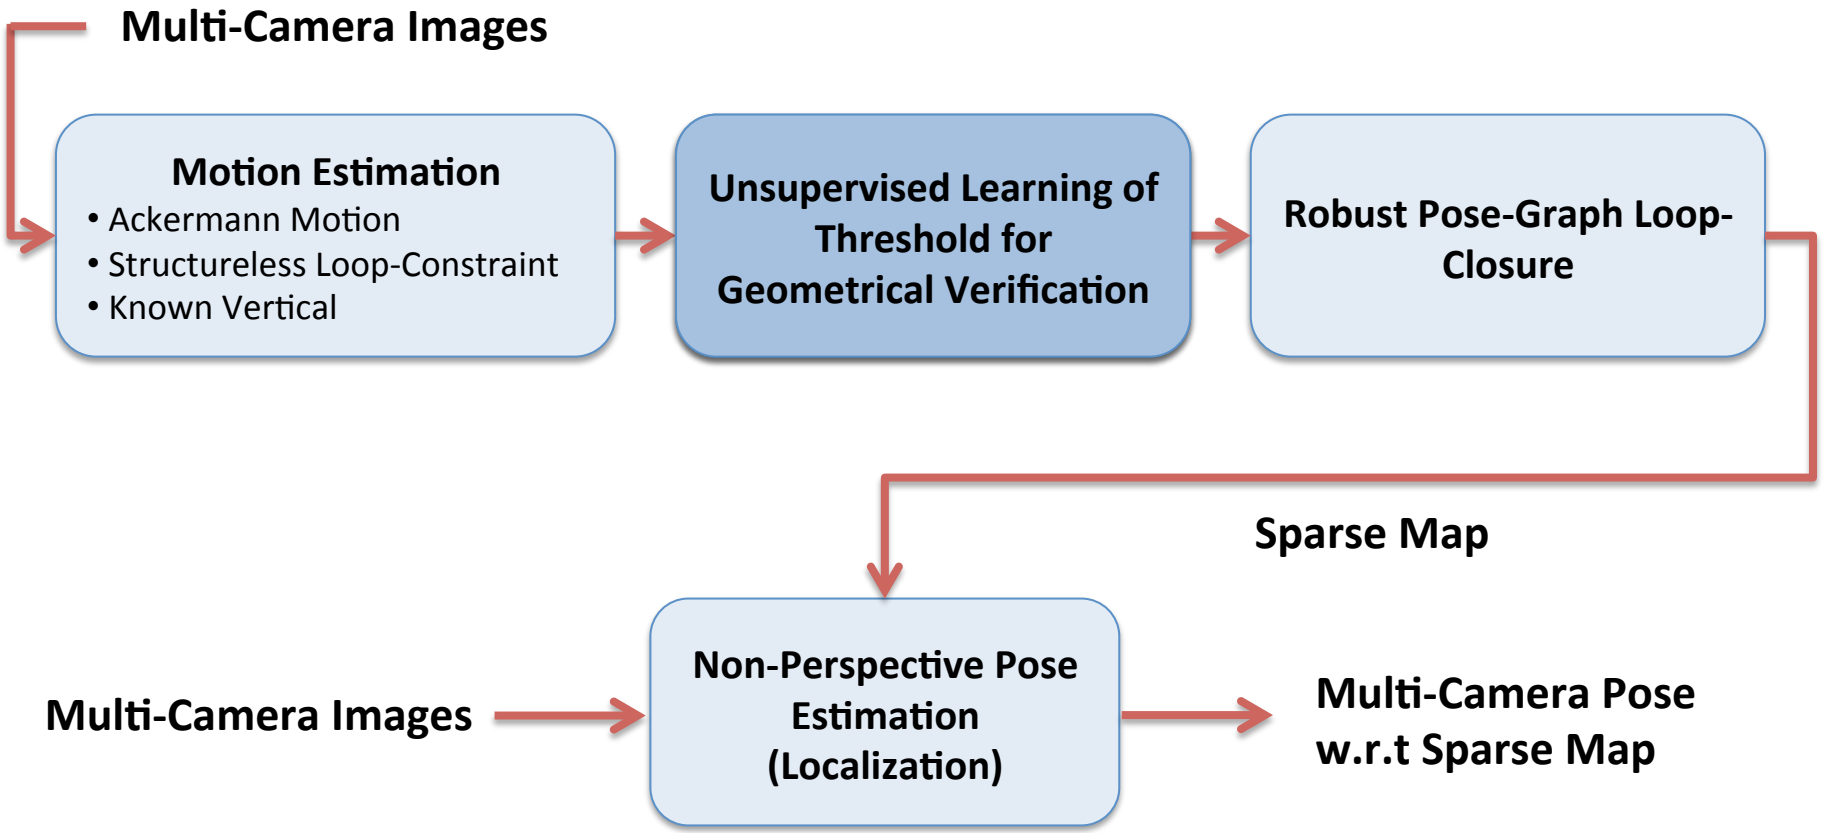


(b)

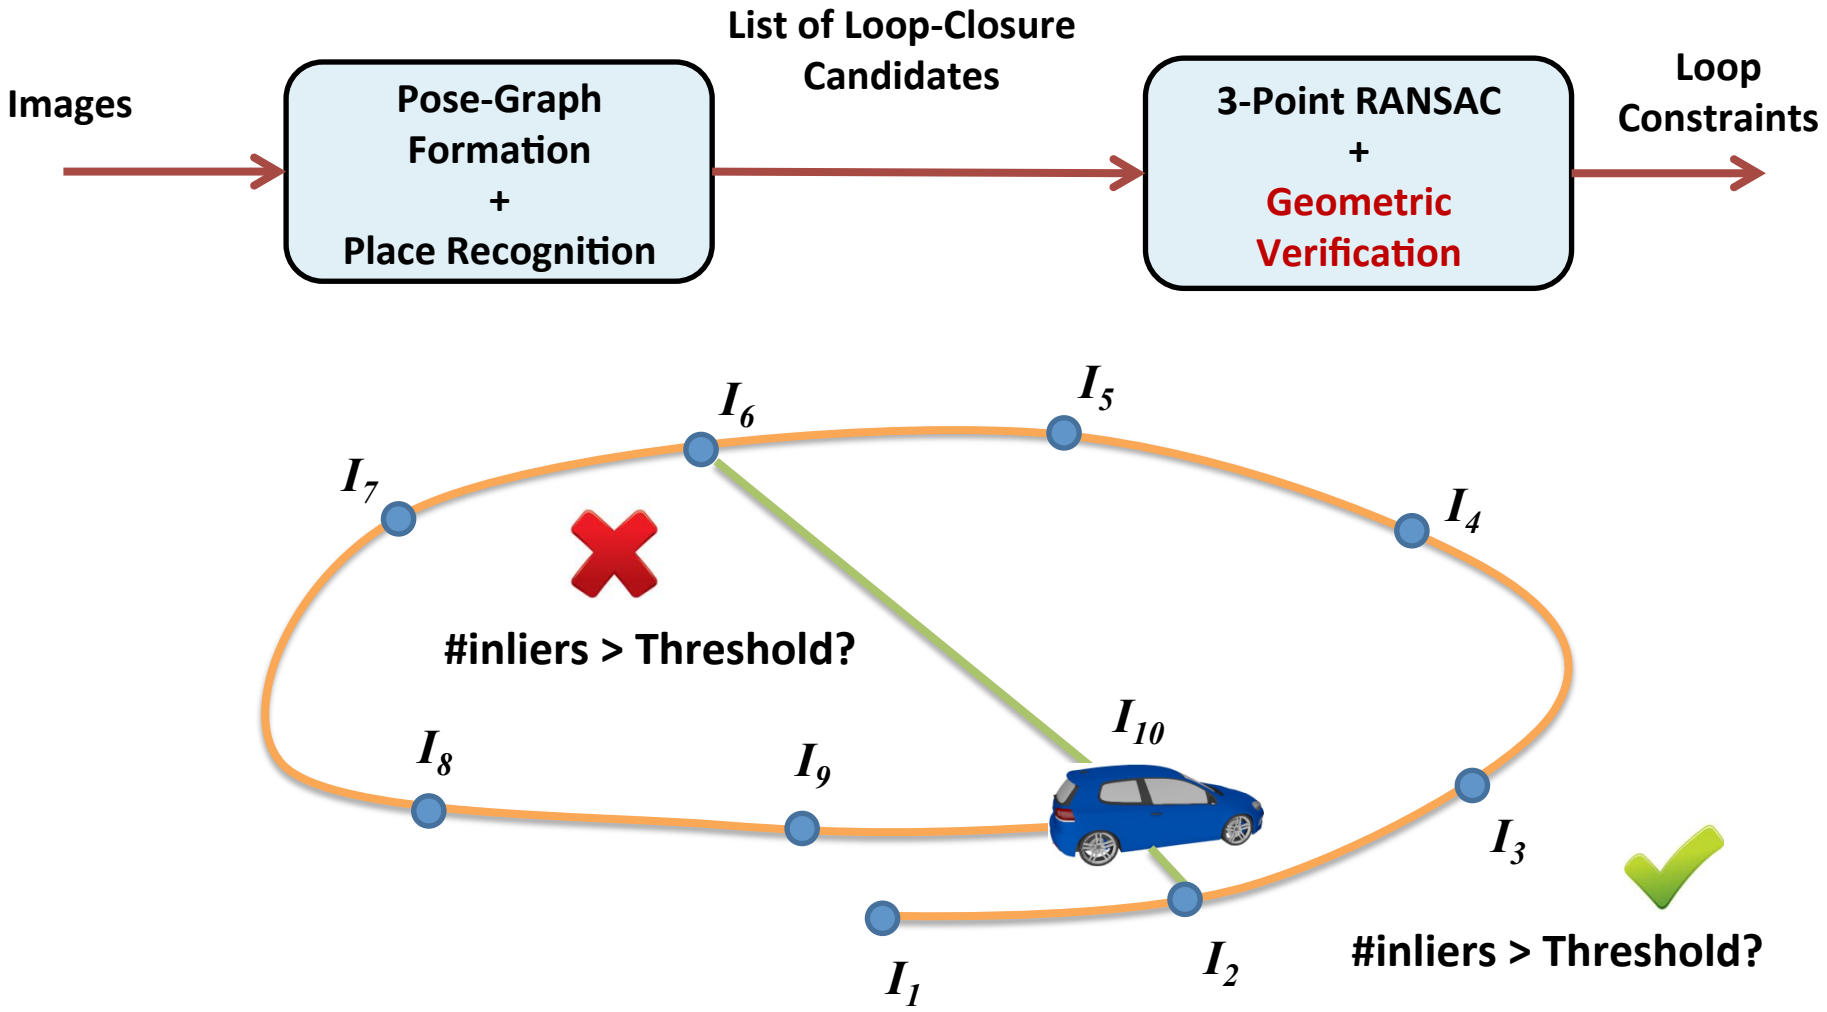


(c)

Mapping and Pose Estimation: Overview



Geometric Verification



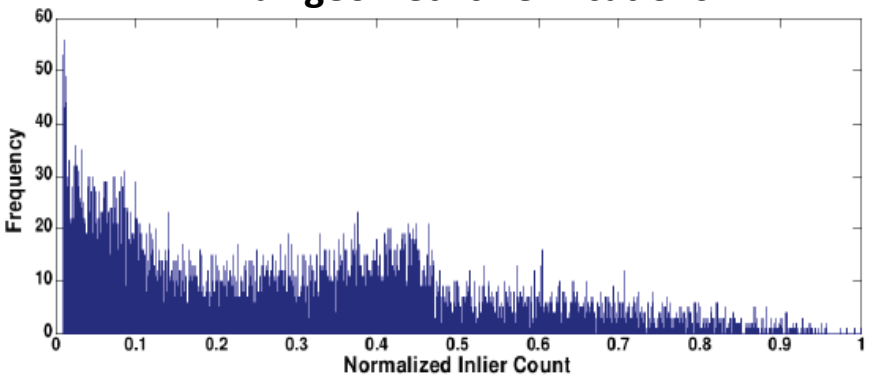
Geometric Verification: Problem

- Threshold is unknown and empirically chosen.
- Too high:
  - many correct loop-closures are omitted.
- Too low:
  - many wrong loop-closures.

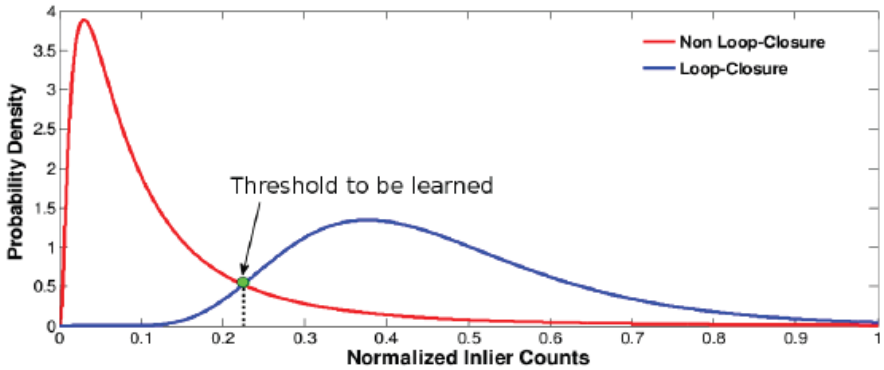
Unsupervised Learning of Threshold for Geometric Verification in Visual-Based Loop-Closure
Gim Hee Lee, and Marc Pollefeys
Submitted to: IEEE International Conference on Robotics and Automation (ICRA) 2014

Unsupervised Learning of Threshold

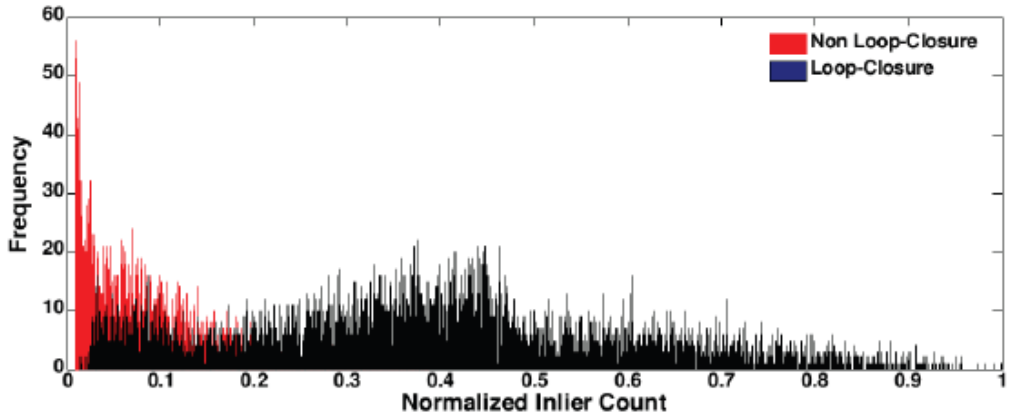
Distribution of inlier counts from all geometric verifications



Modeled as Two-Components Log-Normal Mixture Distribution



Ground Truth from GPS/INS



Unsupervised Learning of Threshold

- Two-Components Log-Normal Mixture Distribution:

$$p(V, L|\mu, \Sigma, \pi) = \prod_{n=1}^{N} \prod_{k=1}^{2} (\pi_k ln\mathcal{N}(v_n|\mu_k, \Sigma_k))^{L_{n,k}}$$

Robot Poses: $X = [x_1, x_2, ..., x_t]^T$

N Potential Loop-Closure Pairs: $\mathcal{X} = [...., \{x_i, x_j\}_N]$

Observed variable

Inlier counts from all N potential loop-closures: $V = [v_1, v_2, ..., v_N]$

Latent variable

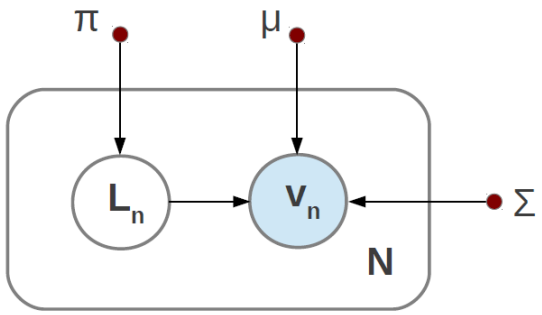
Actual loop-closure state: $L = [l_1, l_2, ..., l_N], \quad l_n \in \{0, 1\}$

Hidden variable

Mean: $\mu = [\mu_1, \mu_2]$ **Standard Deviation:** $\Sigma = [\Sigma_1, \Sigma_2]$

Mixing Coefficients: $\pi = [\pi_1, \pi_2]$

Unsupervised Learning of Threshold



- Find $\phi = [\mu, \Sigma, \pi]$ that maximizes the loglikelihood $\ln p(V|\phi)$

$$\underset{\phi}{\operatorname{argmax}} \ln p(V|\phi) = \underset{\phi}{\operatorname{argmax}} \ln\{\sum_{L} p(V, L|\phi)\}$$

Inference with Expectation Maximization!

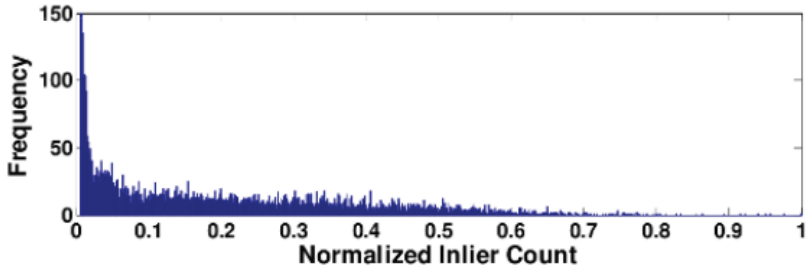
Solve for Threshold

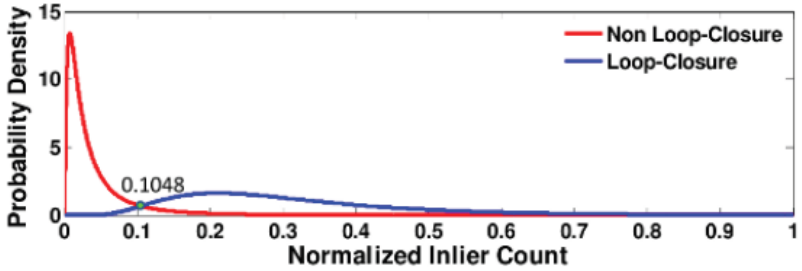
- Threshold is the intersection of the log-normal distributions:

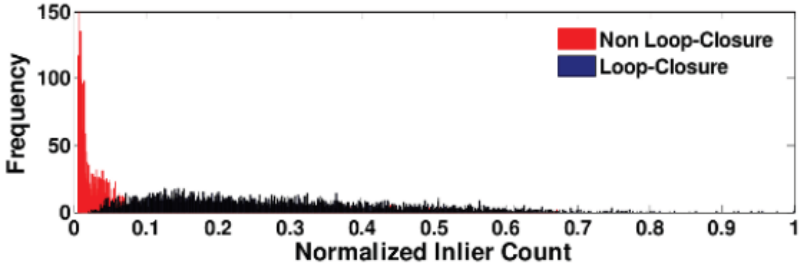
$$f(v) = \pi_1 ln\mathcal{N}(v|\mu_1, \Sigma_1) - \pi_2 ln\mathcal{N}(v|\mu_2, \Sigma_2) = 0$$

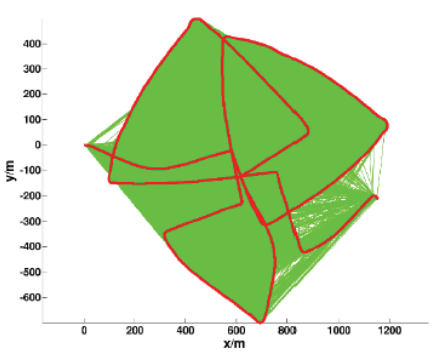
- Degenerated case: Too few loop-closures.
- χ^2 test, i.e. “goodness-of-fit” to detect degeneracy.

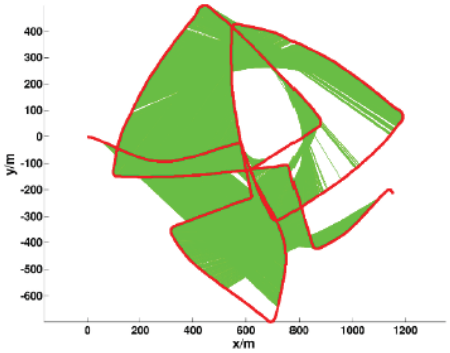
Results: Wolfsburg





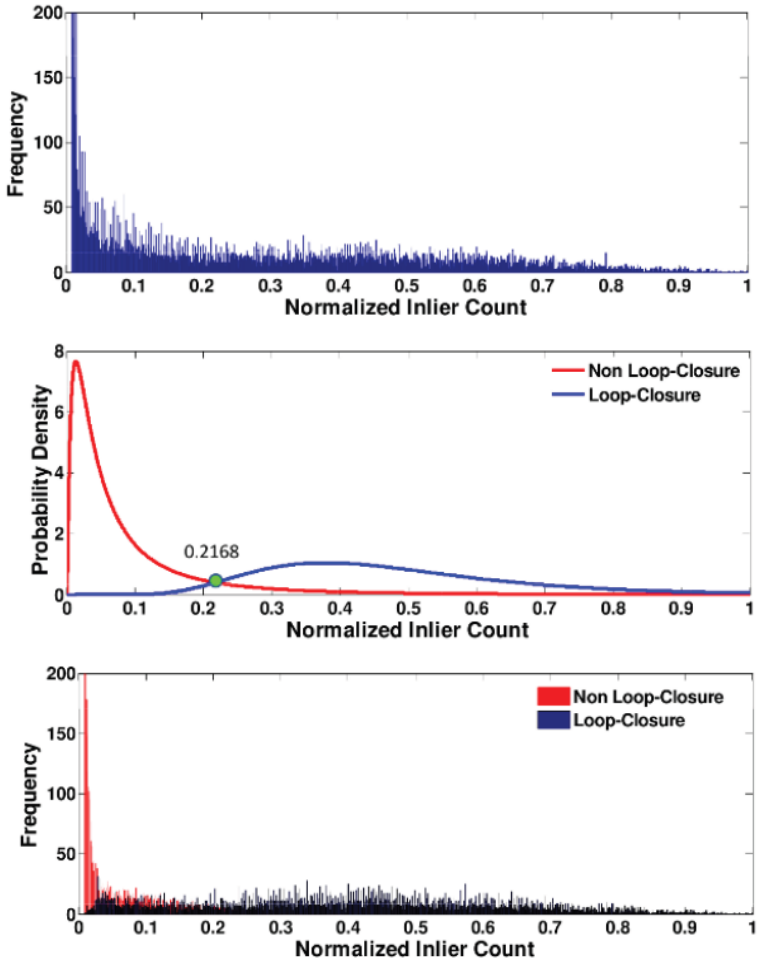


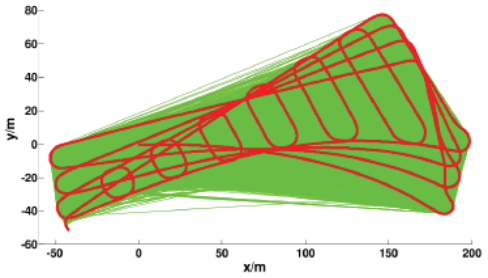


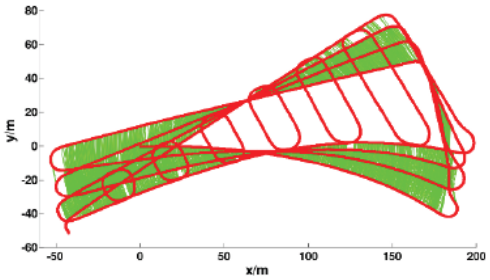


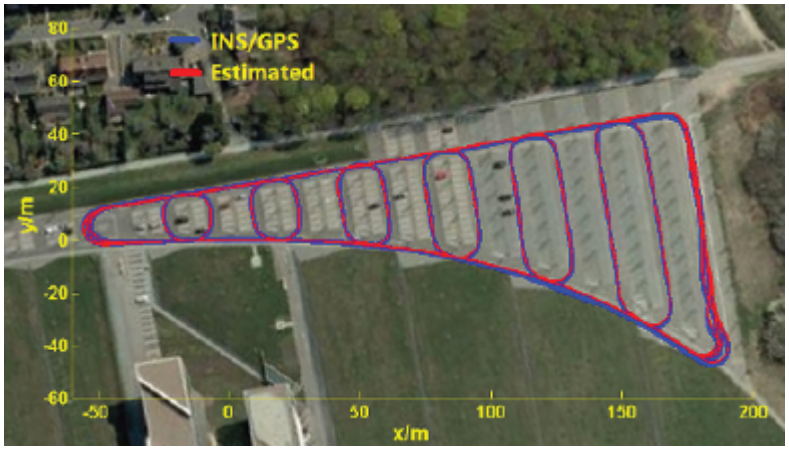


Results: Parking Garage 1

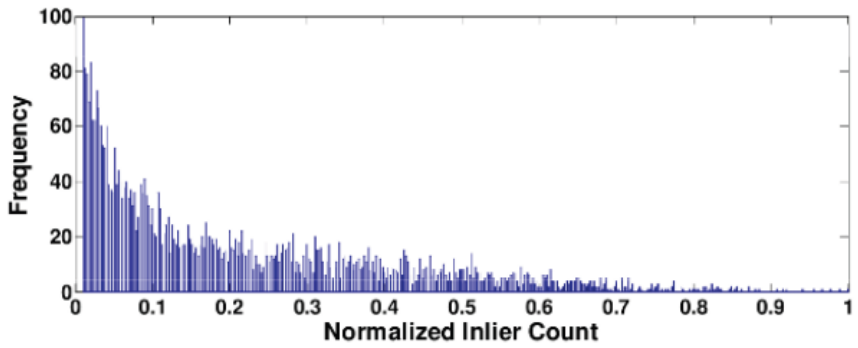


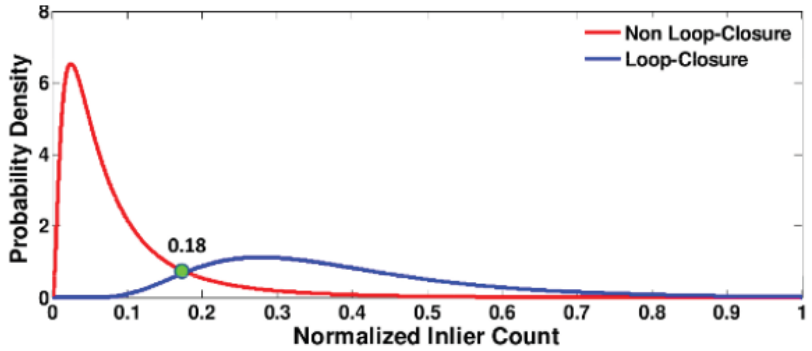


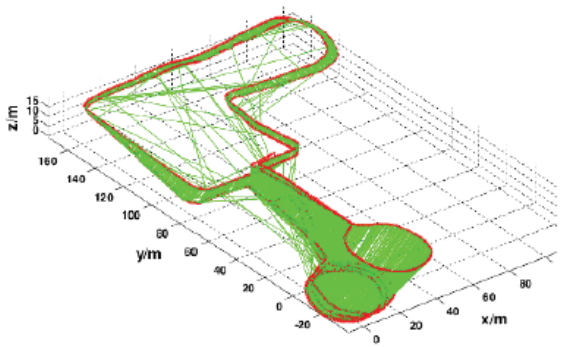


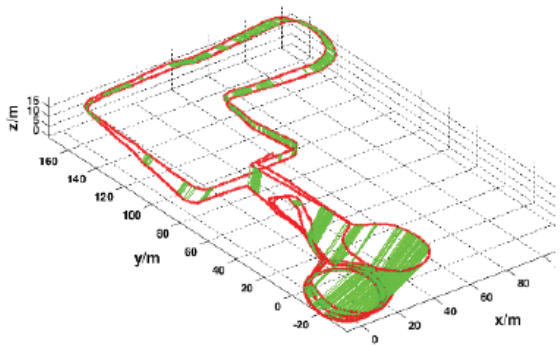


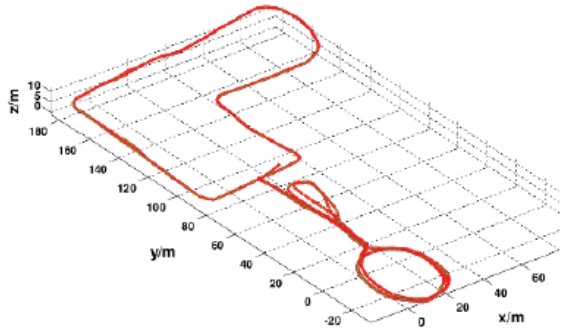
Results: New College (Stereo)





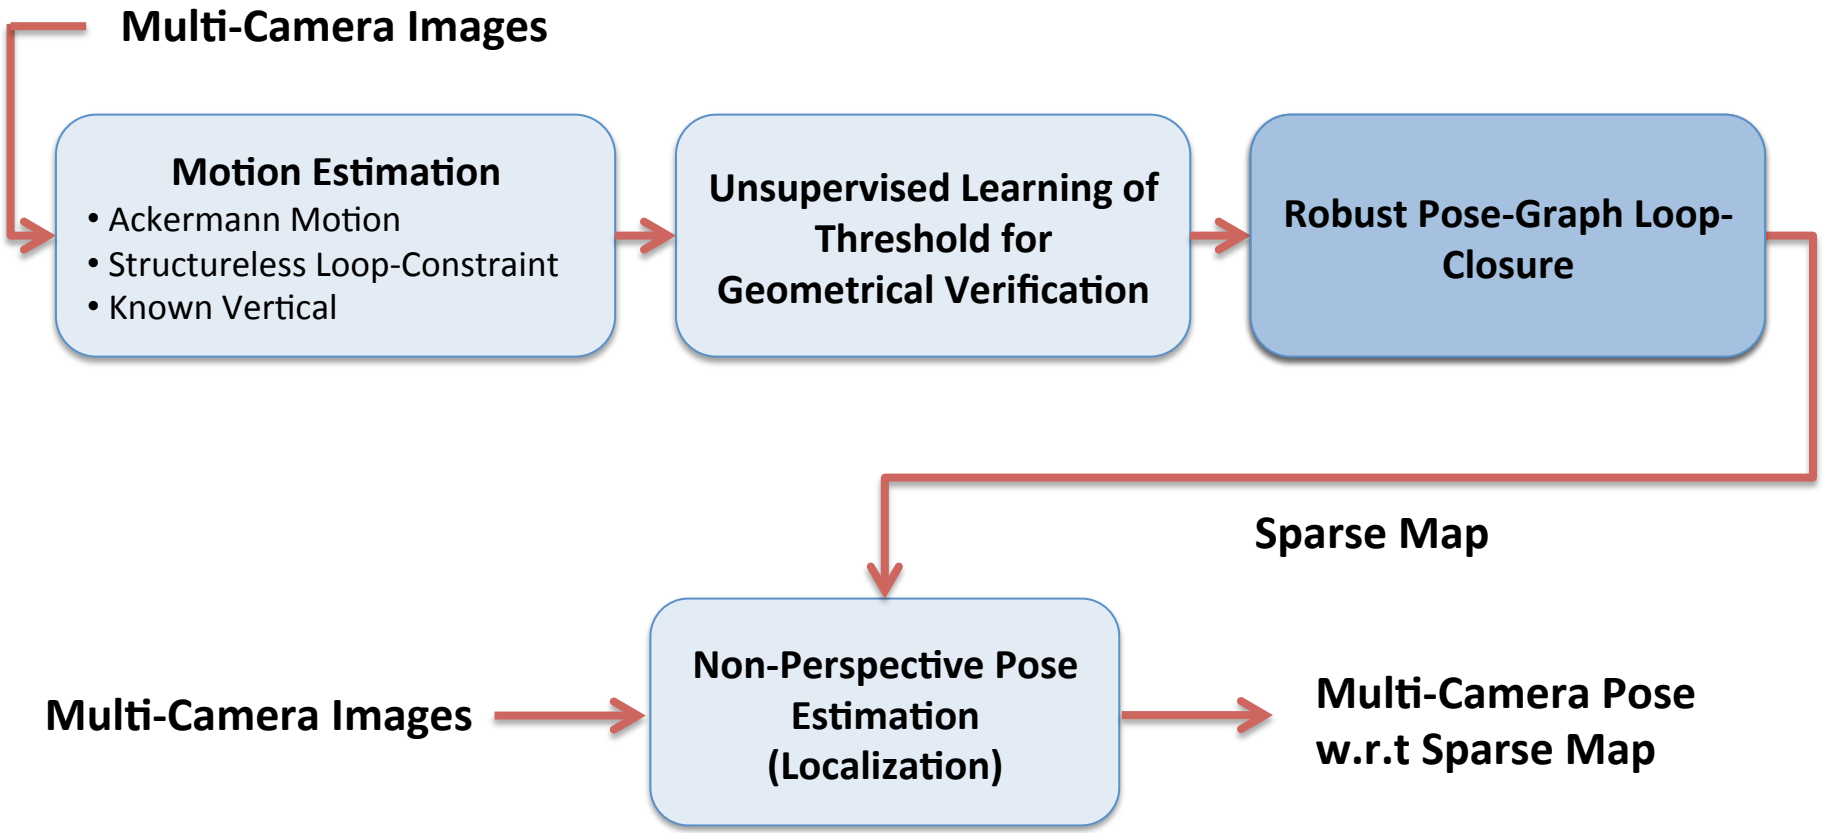






The New College Vision and Laser Dataset
M. Smith, I. Baldwin, W. Churchill, R. Paul, and P. Newman,
International Journal of Robotics Research (IJRR), vol. 28, no. 5, pp. 595–599, May 2009.

Mapping and Pose Estimation: Overview



Robust Pose-Graph Loop-Closure: Motivation

- Pose-graph optimization as non-linear least-squares:

$$\underset{X}{\operatorname{argmin}} \sum_{ij} \left\| z_{ij} - h(x_i, x_j) \right\|^2_{\sum_{ij}}$$

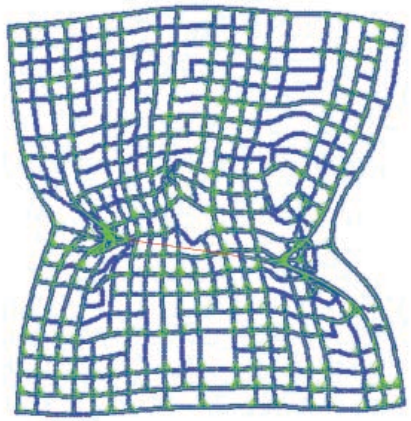
Not robust to outliers!

Robust Pose-Graph Loop-Closures with Expectation-Maximization
Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2013

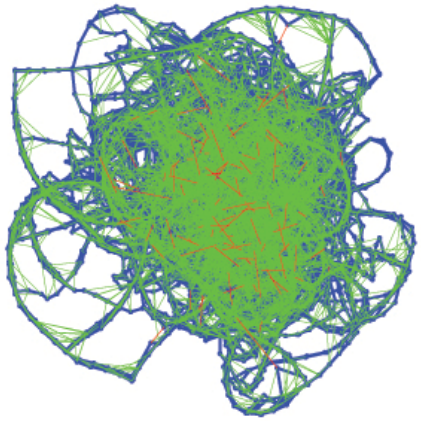
Robust Pose-Graph Loop-Closure: Motivation

City10000 Dataset

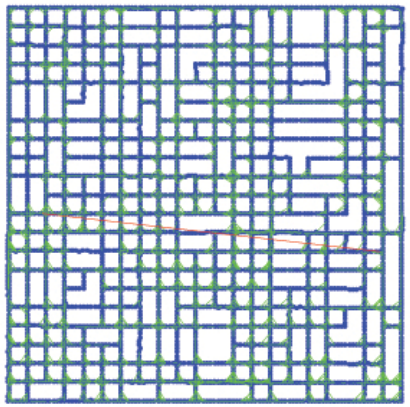
Non-Robust

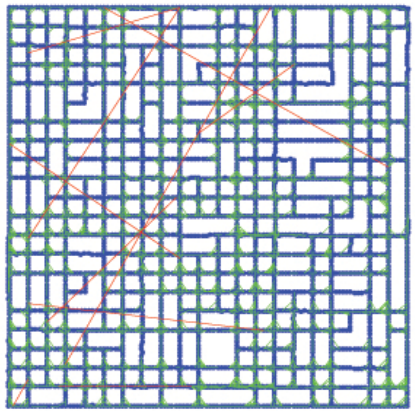


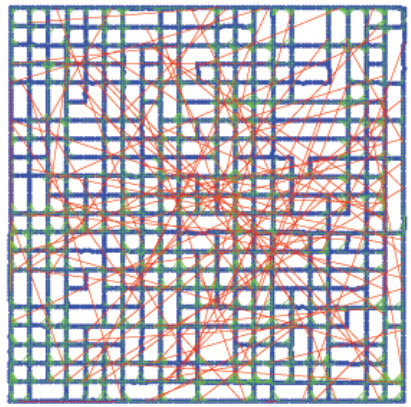




Robust





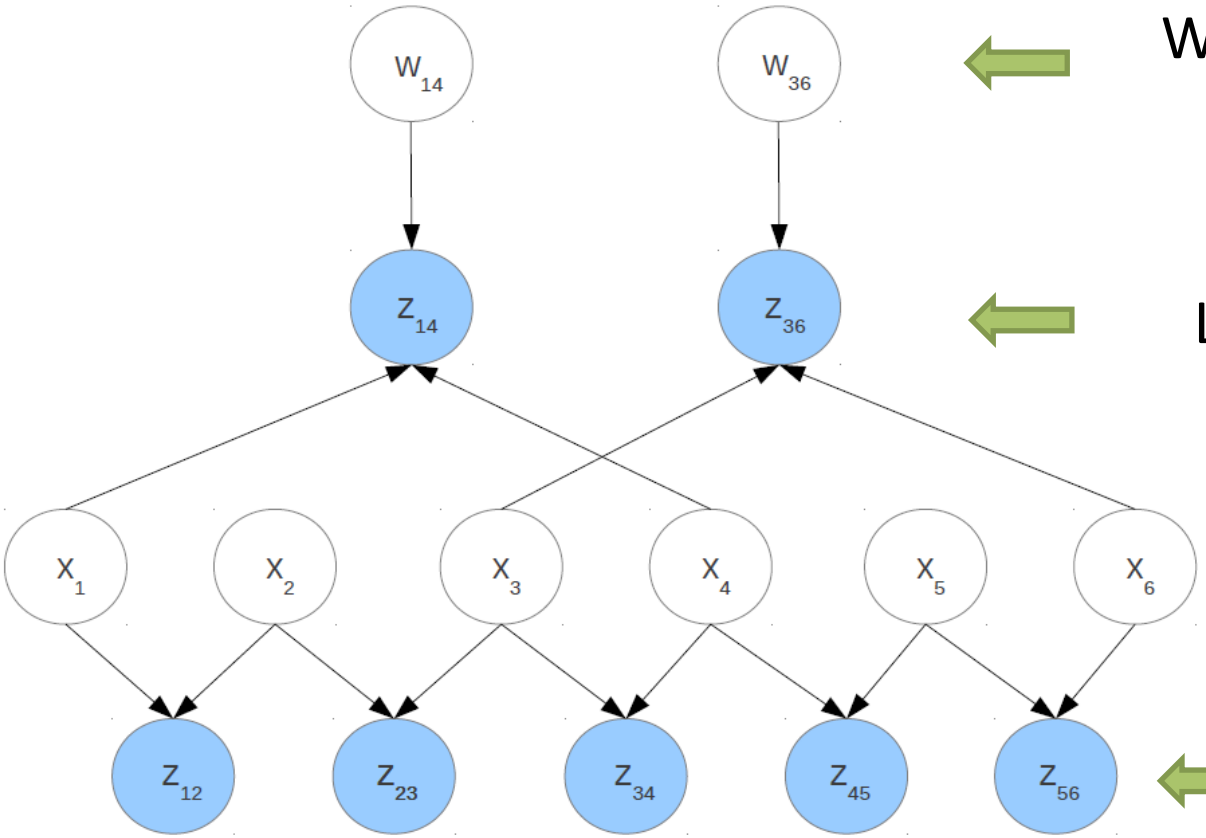


1 Outlier

10 Outliers

100 Outliers

Model Problem as Bayesian Network



Weights for loop-closure constraints, $w \in [0, 1]$

Loop-Closure Constraints

Poses

Odometry constraints

Inference

- Find poses X given the constraints Z (maximum a posterior):

$$\underset{X}{\operatorname{argmax}}\, p(X \mid Z) = \underset{X}{\operatorname{argmax}} \int_{W} p(X, W \mid Z) dW$$

Continuous non-Gaussian latent variable!

- Classification EM:

$$W^{k+1} = \underset{W}{\operatorname{argmax}}\; p(W \mid X^k, Z)$$

$$X^{k+1} = \underset{X}{\operatorname{argmax}}\; \ln p(X \mid W^{k+1}, Z)$$

Iterates until convergence!

Inference: Expectation Step

$$W^{k+1} = \underset{W}{\operatorname{argmin}} \sum_{ij} w_{i,j} ||h(x_i^k, x_j^k) - z_{i,j}||^2_{Q_{i,j}} \quad -C^2(\ln w_{i,j} - w_{i,j})$$

Penalty Term

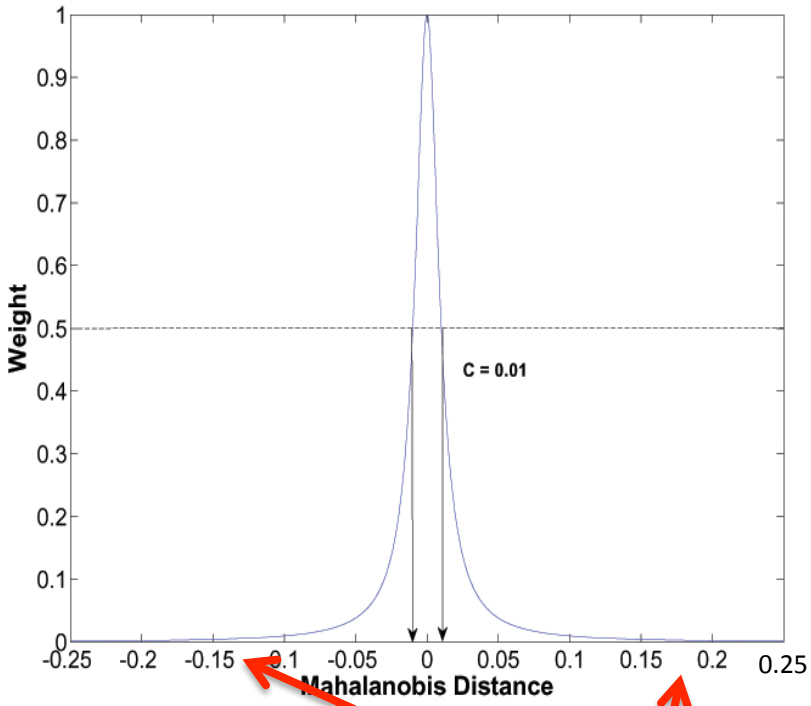
Trivial Solution of 0

- Differentiates and set to 0 gives the Cauchy Fucntion.

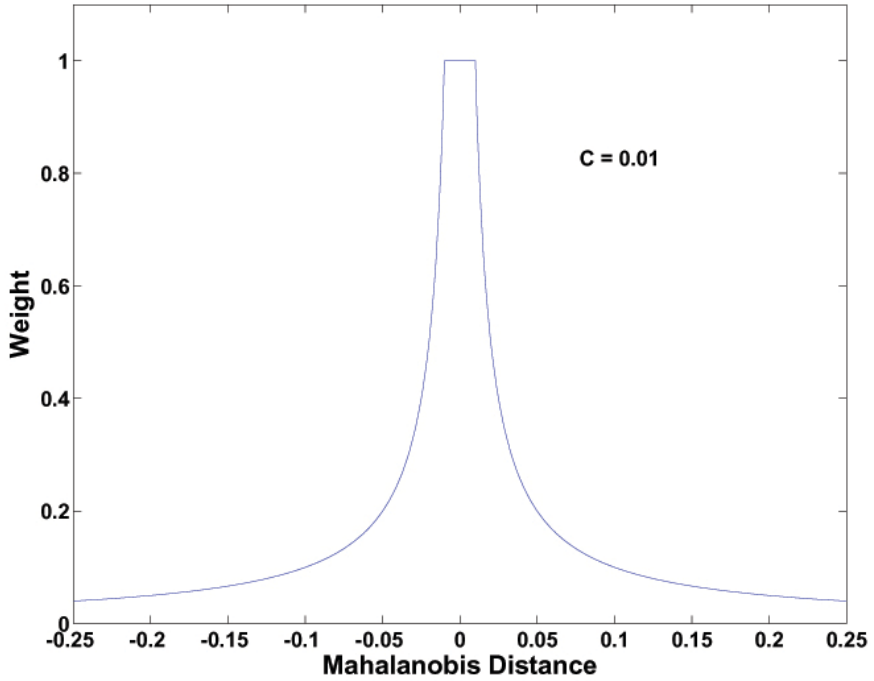
$$w_{i,j}^{k+1} = \frac{C^2}{C^2 + ||h(x_i^k, x_j^k) - z_{i,j}||^2_{Q_{i,j}}}$$

Inference: Expectation Step

Cauchy Function



Huber Function



Near zero weights
for high errors!

Inference: Maximization Step

$$X^{k+1} = \underset{X}{\operatorname{argmax}} \ln p(X|W^{k+1}, Z) = -\underset{X}{\operatorname{argmin}} \ln p(X|W^{k+1}, Z)$$

≡ Pose-Graph SLAM problem!

$$\Rightarrow \quad X^{k+1} = \underset{X}{\operatorname{argmin}} \underbrace{\sum_i ||h(x_i, x_{i+1}) - z_{i,i+1}||^2_{P_{i,i+1}}}_{\text{Odometry cost}} + \underbrace{\sum_{ij} w^{k+1}_{i,j} ||h(x_i, x_j) - z_{i,j}||^2_{Q_{i,j}}}_{\text{Loop-closure cost}}$$

Inference: Maximization Step

- Solve for X^{k+1} :

$$J^T \Lambda J \delta = -J^T \Lambda \epsilon$$
$$X^{k+1} = X^k + \delta$$

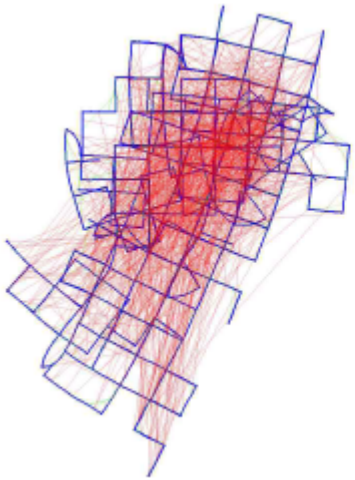
- where Λ is the information matrix formed with:

$$P_{i,i+1}^{-1} \text{ and } \hat{Q}_{i,j}^{-1} = w_{i,j}^{k+1} Q_{i,j}^{-1}$$

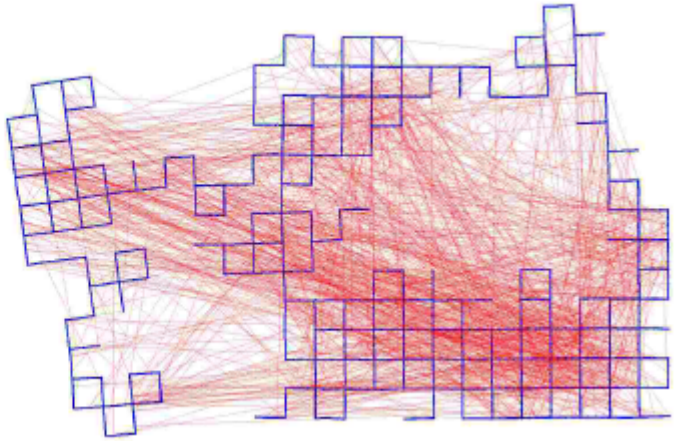
Local Minima

Manhattan3500Olson Dataset

- Prevent local minima:
  - Perform the EM algorithm multiple times.
  - Loop constraints with very low weights are removed after each EM process.



After 1st EM process



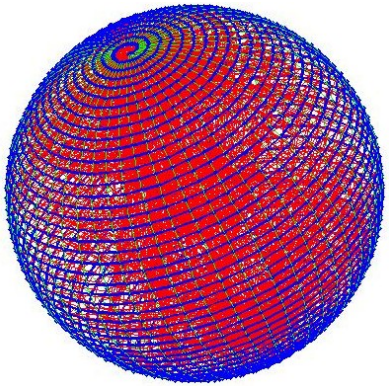
Correct convergence after 3rd EM process

Results: Simulated Datasets

Sphere2500 Dataset

SIMULATED DATASETS USED FOR EVALUATIONS

Dataset	2D/3D	Vertices	Loop-Edges
City10000	2D	10000	10688
Intel	2D	943	894
Manhattan3500 (Olson)	2D	3500	2099
Manhattan3500 (g2o)	2D	3500	2099
Sphere2500	3D	2500	2450
ParkingGarage	3D	1661	4615

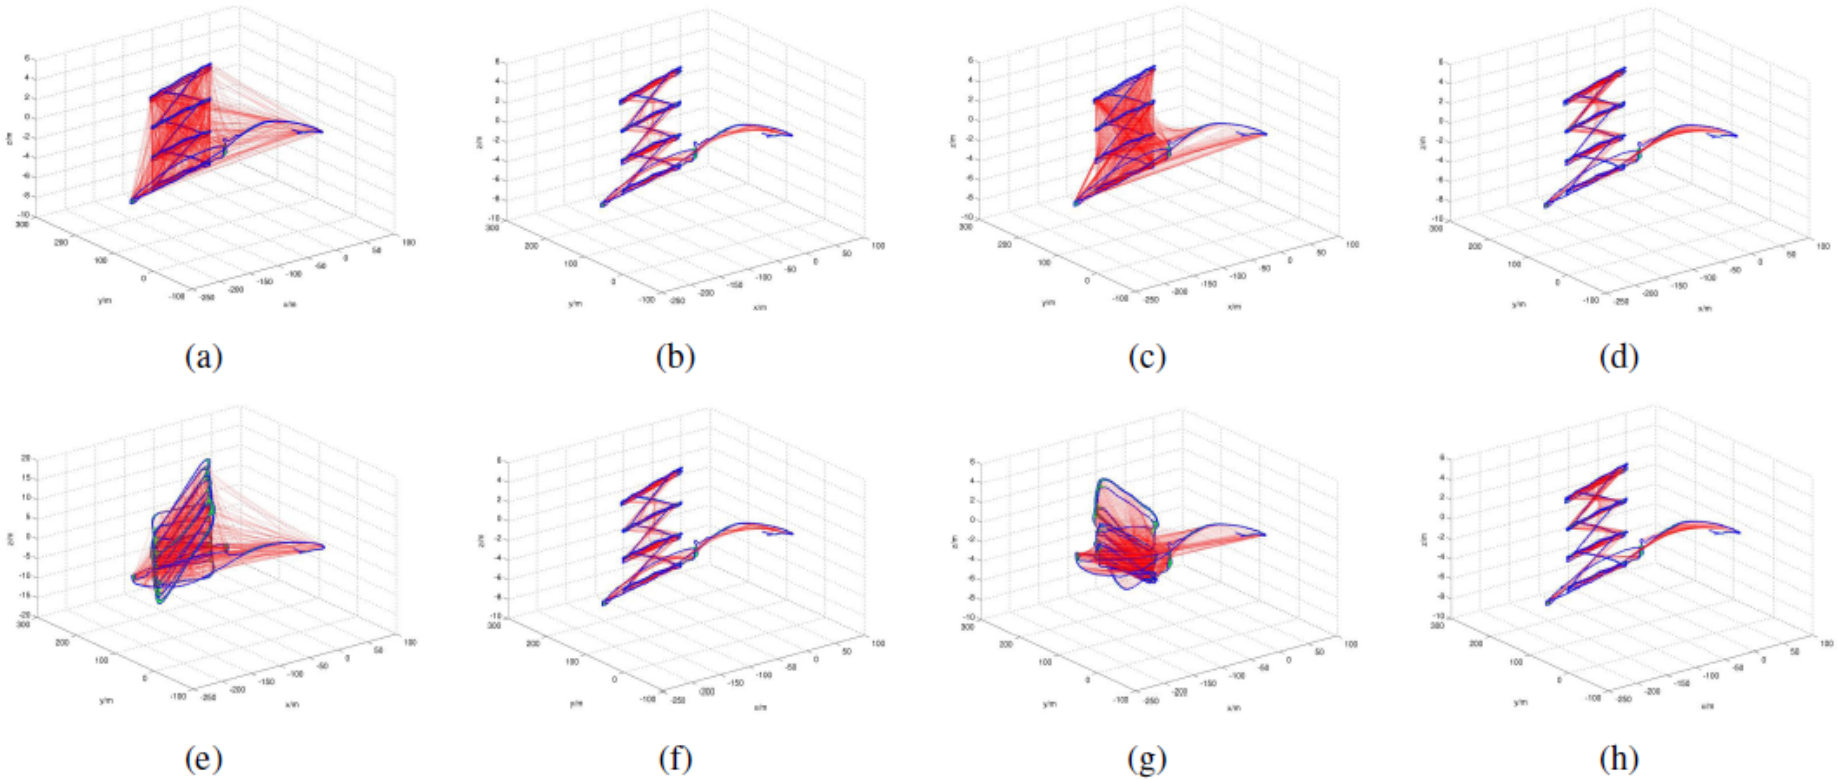


- Random, Local, Randomly Grouped and Locally Grouped constraints.
- Run 10x for each policy with 1000 outliers.
- 100% success rate, maximum relative pose metric error is 3.84×10^{-5}.

Results: Comparison with Switchable Constraint

ParkingGarage Dataset

1000 Random | 1000 Local | 50 sets of 20 Random Grouped | 50 sets of 20 Locally Grouped



(a) (b) (c) (d)

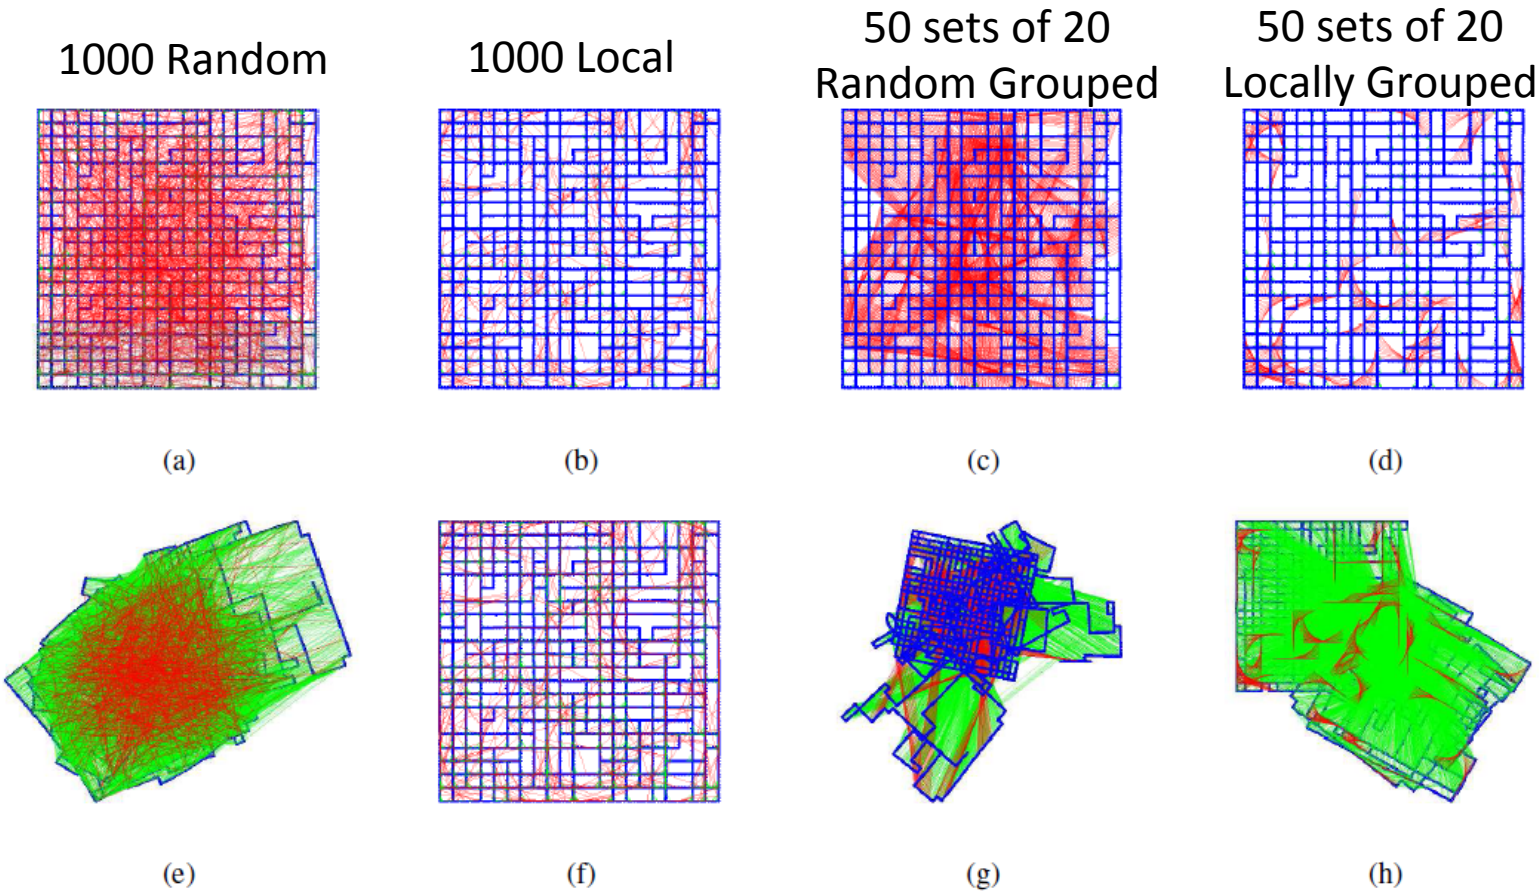
(e) (f) (g) (h)

Towards a robust back-end for pose-graph slam,
N. Suenderhauf and P. Protzel,
International Conference on Robotics and Automation (ICRA), 2012.

(a) – (d) : Our algorithm
(e) – (h) : Switchable Constraints

Results: Comparison with Max-Mixture

City10000 Dataset

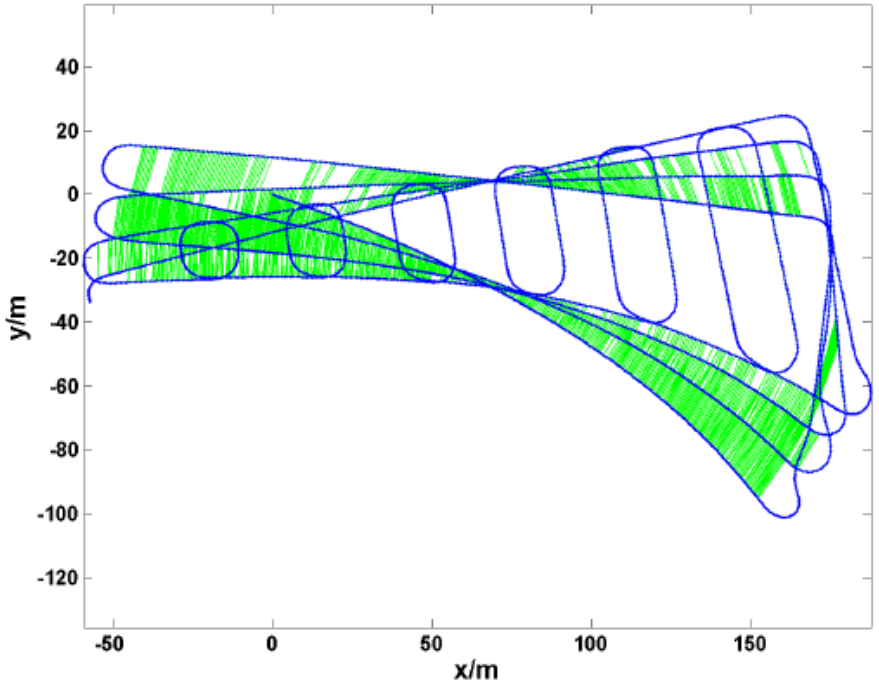


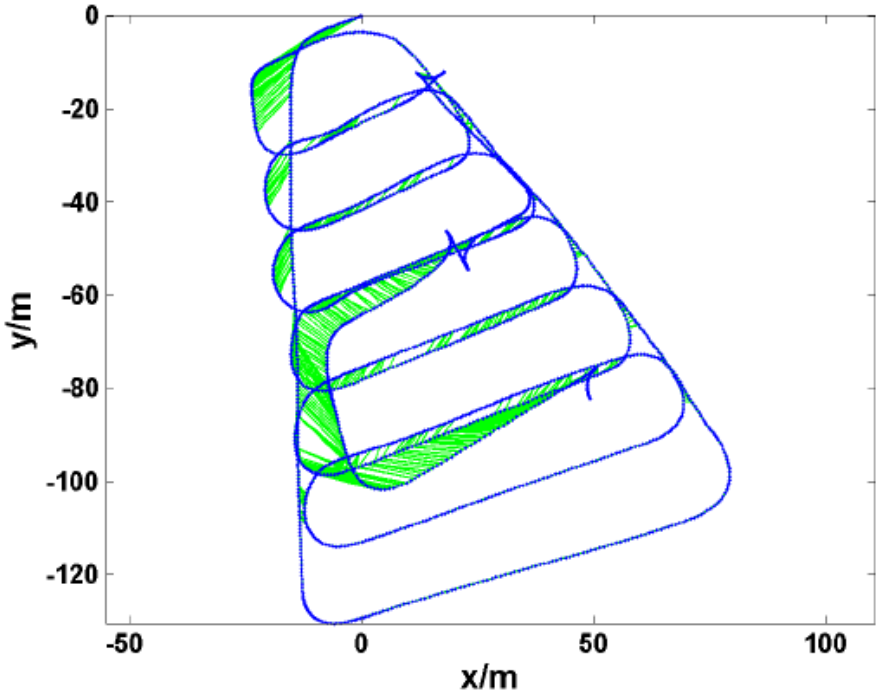
Inference on networks of mixtures for robust robot mapping
E. Olson and P. Agarwal,
Robotics Science and Systems (RSS), 2012.

(a) – (d) : Our algorithm
(e) – (h) : Max-mixture

Results: Real-World Datasets

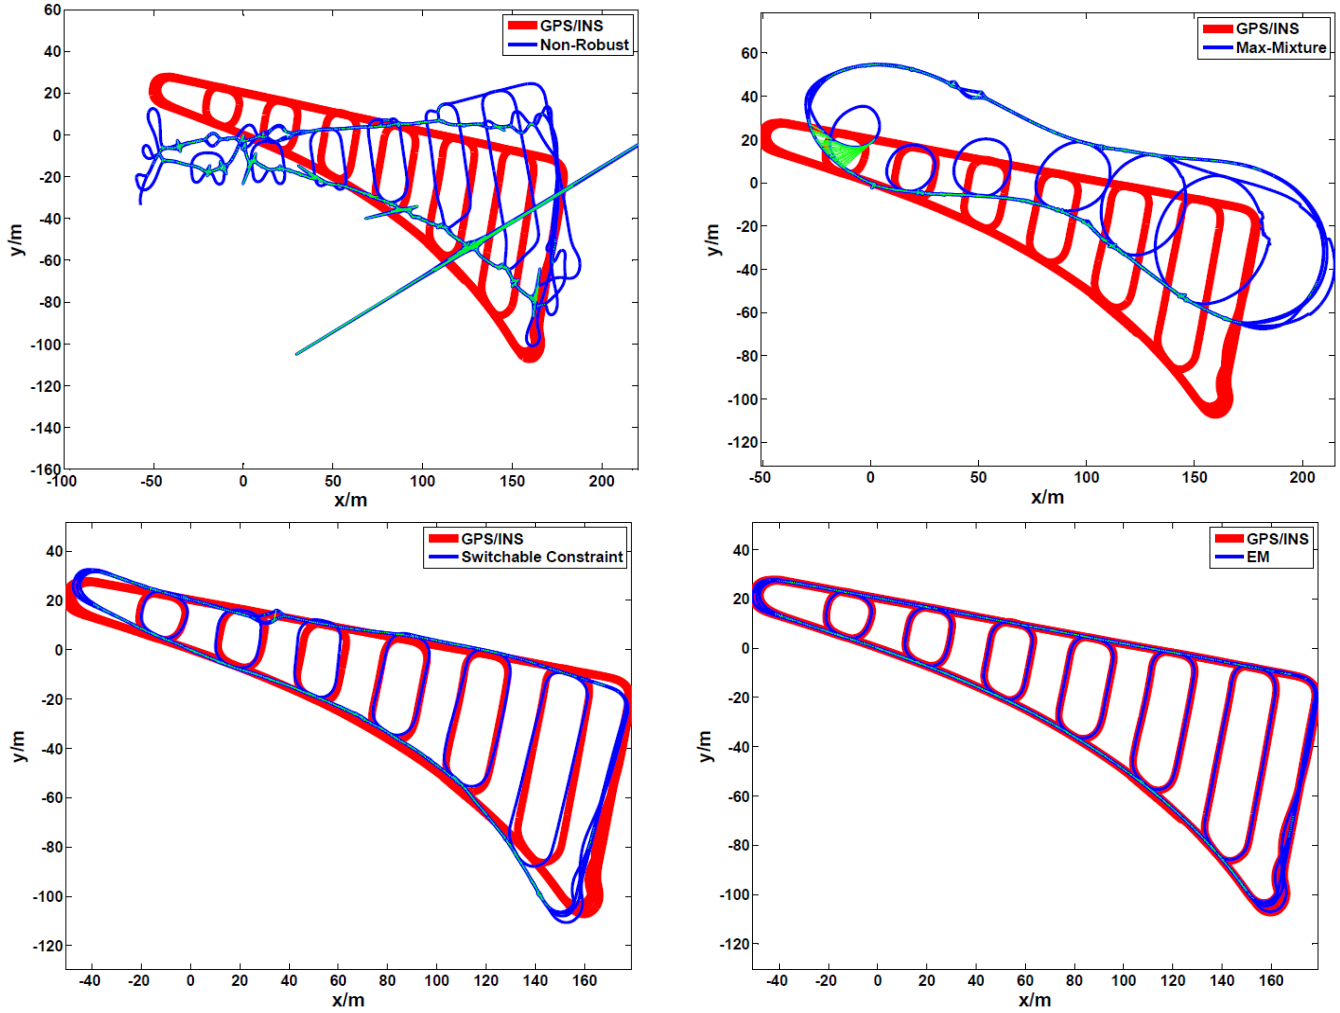
Parking Garage 1 and 2 Datasets





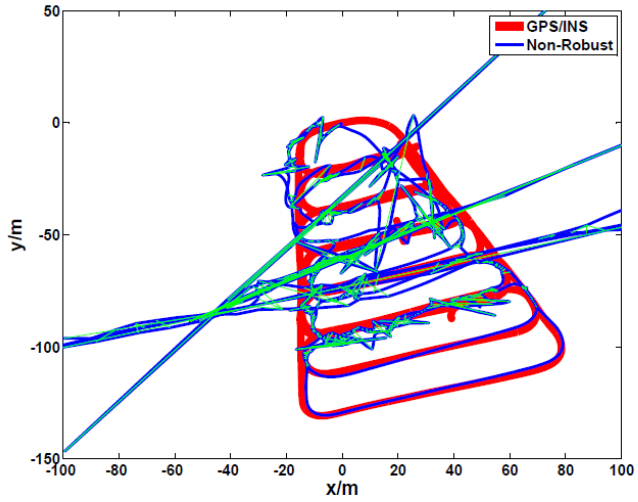
Results: Real-World Datasets

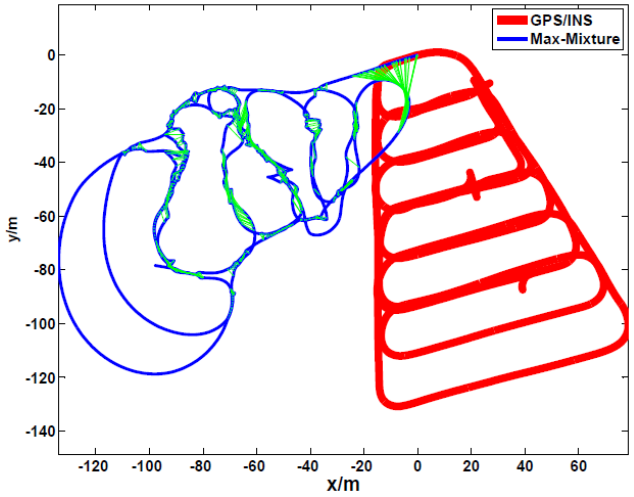
Parking Garage 1 Dataset

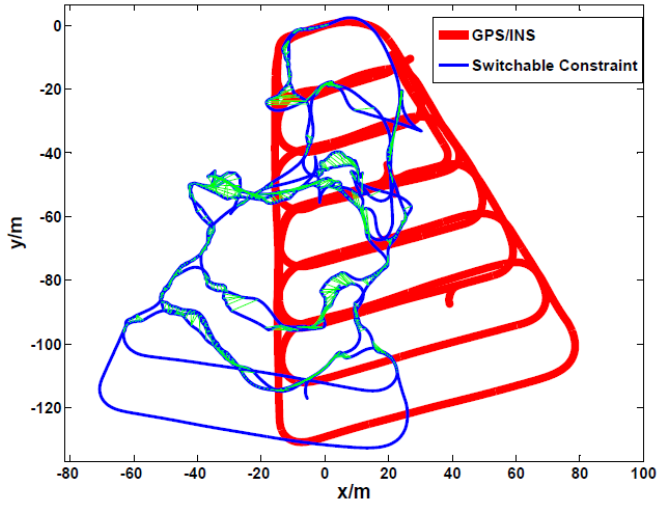


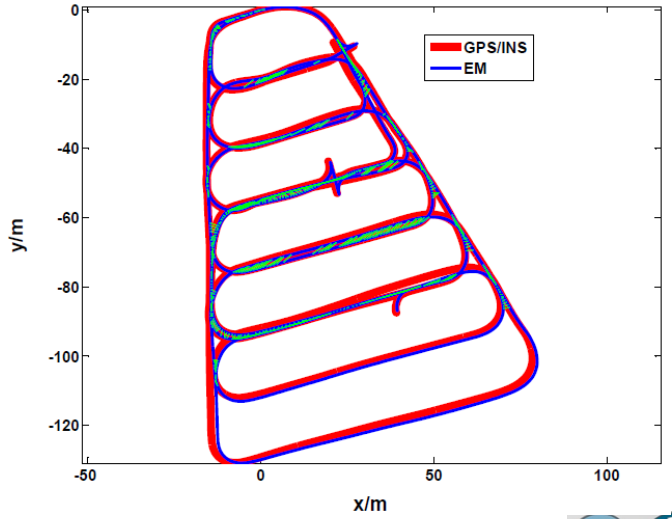
Results: Real-World Datasets

Parking Garage 2 Dataset









Convergence Time

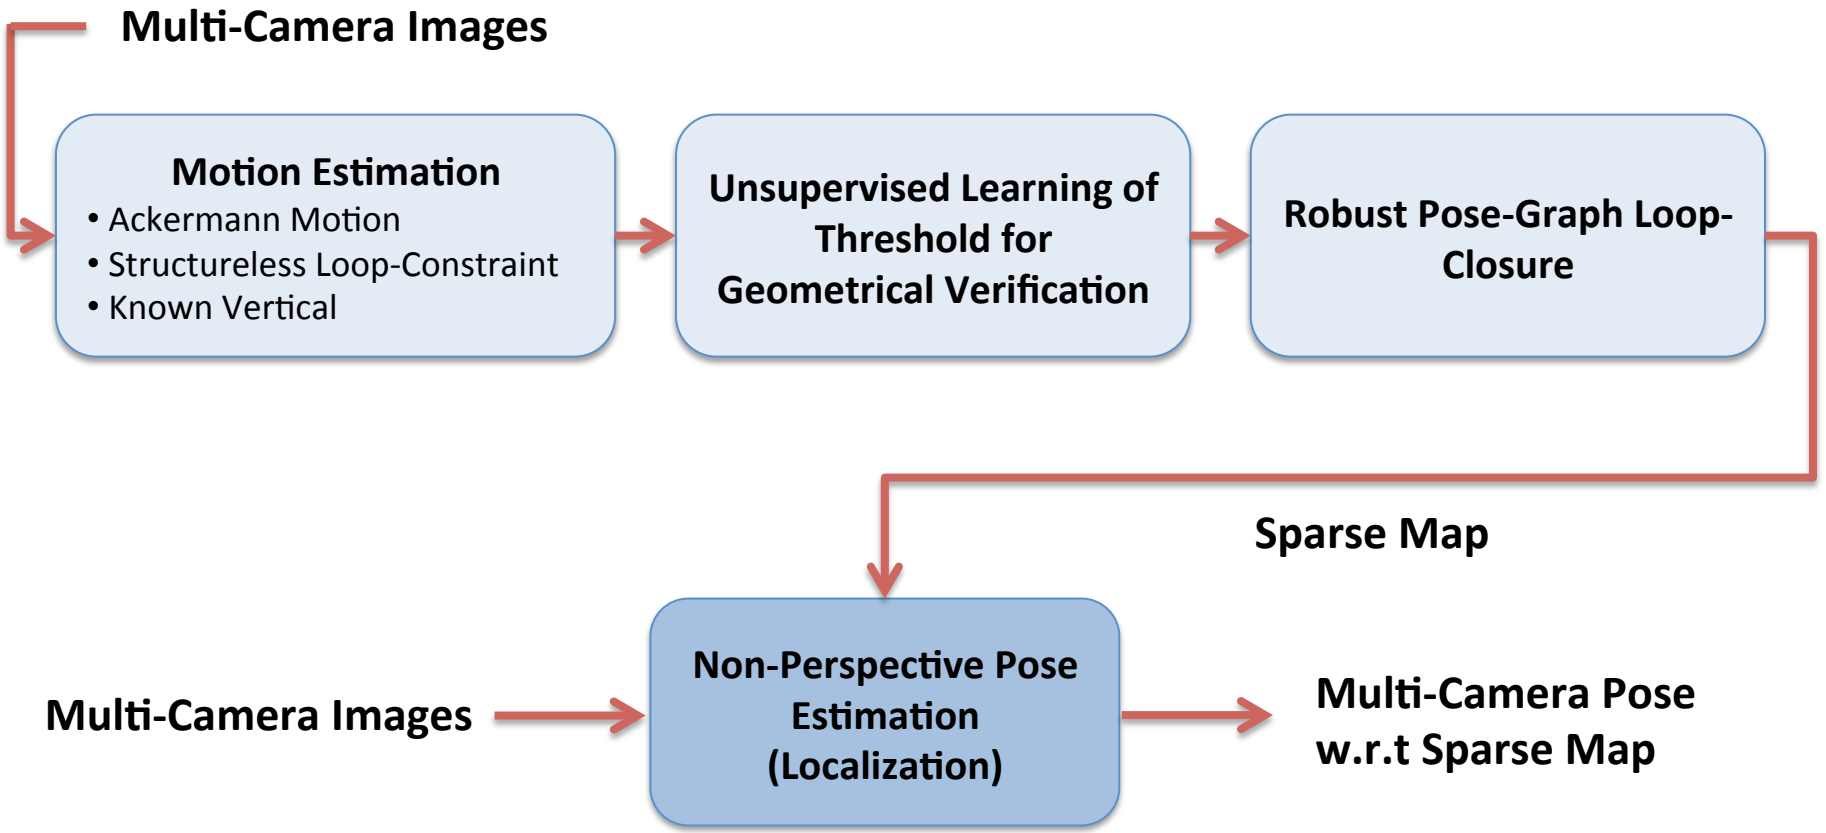
- Implemented with Google Ceres on an Intel Core2 Quad CPU @ 2.40GHz x 4 computer.

CONVERGENCE TIME COMPARISONS

	Runtime (secs)	
Dataset	Non-Robust	EM
City10000	2.2701	2.2926
Intel	0.0813	0.0819
Manhattan3500 (Olson)	0.9900	1.0047
Manhattan3500 (g2o)	0.6662	0.6675
Sphere2500	16.6994	17.1384
ParkingGarage	4.43296	4.49716

Google Ceres Solver
Sameer Agarwal , Keir Mierle and Others
https://code.google.com/p/ceres-solver/

Overview


Multi-Camera Images
Motion Estimation
• Ackermann Motion
• Structureless Loop-Constraint
• Known Vertical
Unsupervised Learning of Threshold for Geometrical Verification
Robust Pose-Graph Loop-Closure
Sparse Map
Multi-Camera Images
Non-Perspective Pose Estimation (Localization)
Multi-Camera Pose w.r.t Sparse Map

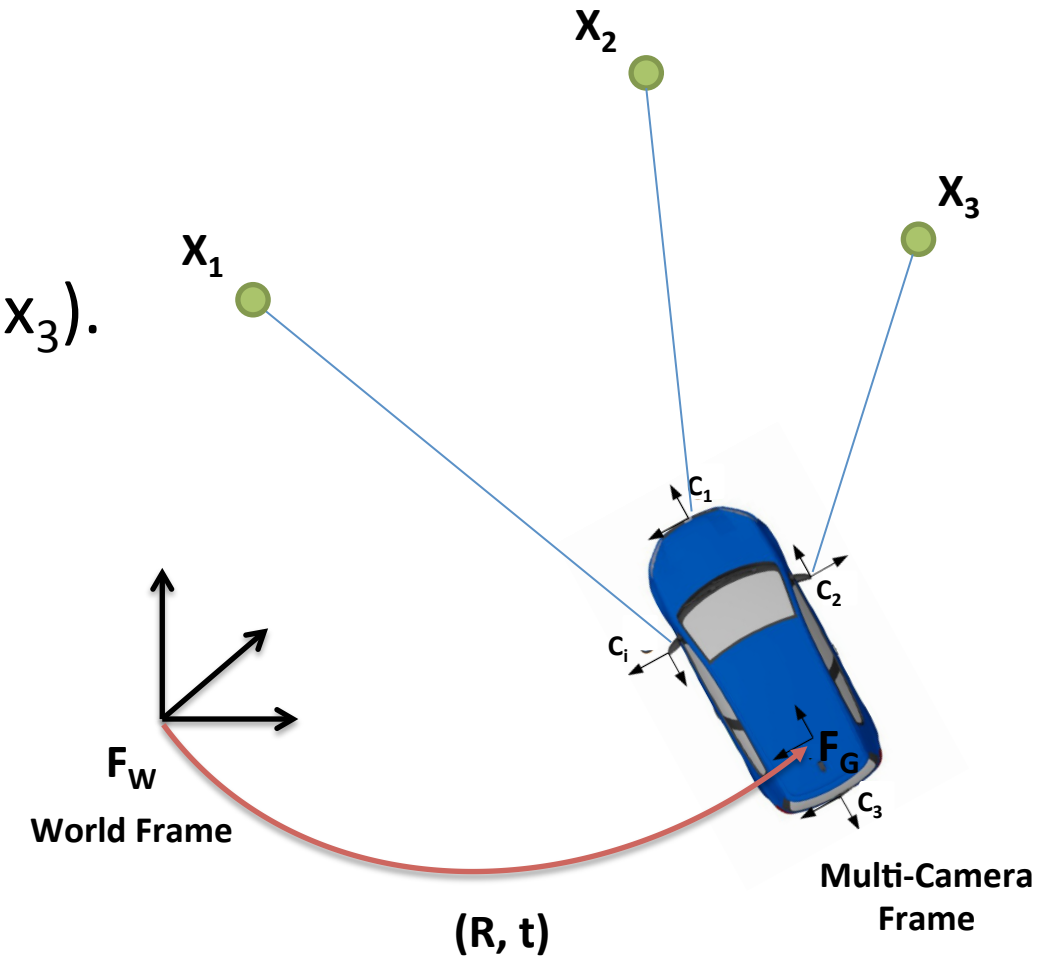
Non-Perspective Pose Estimation: Problem Definition

Given:

1. (X_1, X_2, X_3) in F_W.
2. 2D image points (x_1, x_2, x_3).

Find:

(R, t) that brings F_G to F_W.



Non-Perspective Pose Estimation: Two Steps Approach

1. Find (X_1^G, X_2^G, X_3^G) in multi-camera frame F_G.
2. Use absolute orientation to find (R, t) given

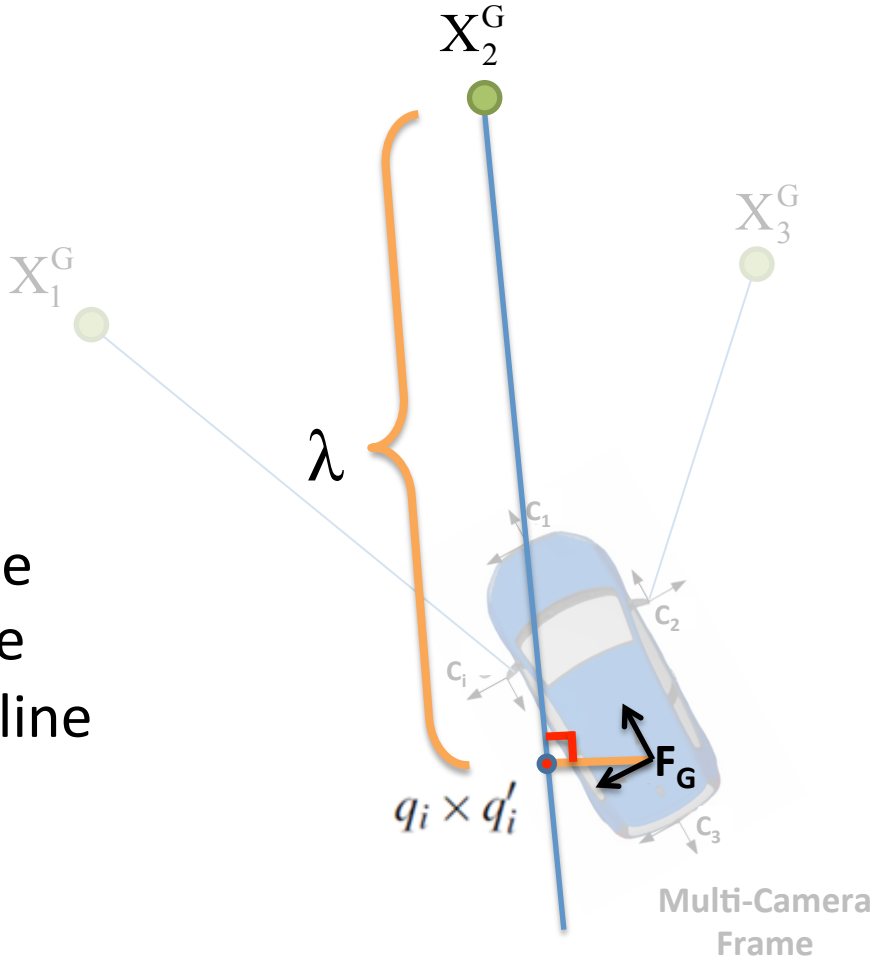
$$(X_1 \leftrightarrow X_1^G, X_2 \leftrightarrow X_2^G, X_3 \leftrightarrow X_3^G)$$

$$\underset{R,t}{\operatorname{argmin}} \left\| X_i - \left(R X_i^G + t \right) \right\|^2$$

Minimal Solutions for Pose Estimation of a Multi-Camera System
Gim Hee Lee, Bo Li, Marc Pollefeys, and Friedrich Fraundorfer
International Symposium on Robotics Research (ISRR) 2013

Point on Plücker line

$$X_i^G = q_i \times q_i' + \lambda_i q_i$$



q_i : first 3 elements of Plücker line

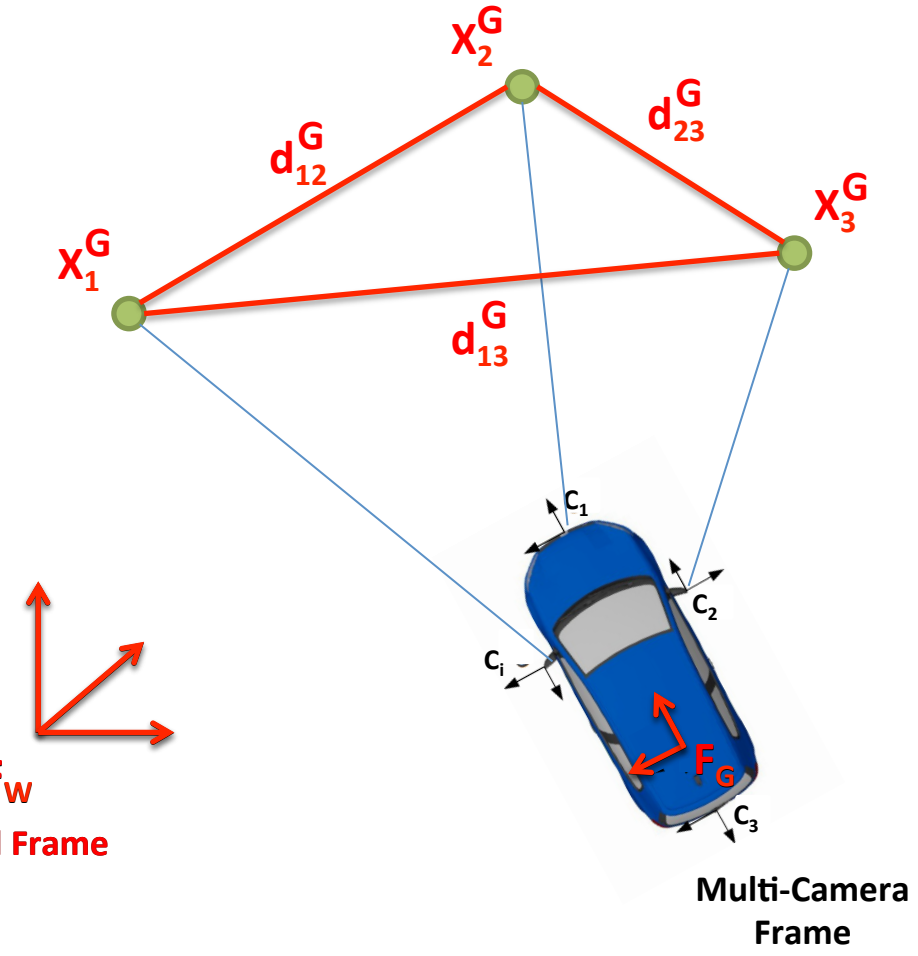
q_i': last 3 elements of Plücker line

λ : Depth of point along Plücker line

Preservation of Distances

$$||X_i - X_j||^2 = ||X_i^G - X_j^G||^2$$

d_{ij} $\qquad$ d_{ij}^G



System of Polynomials

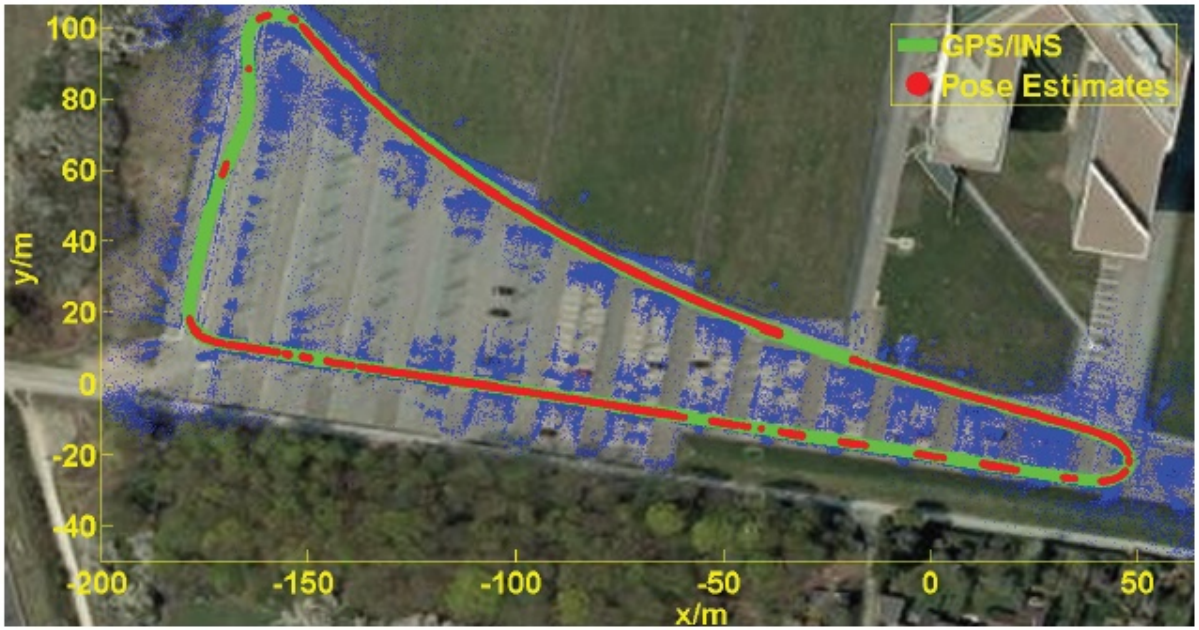
- 3 polynomial equations with 3 unknowns:

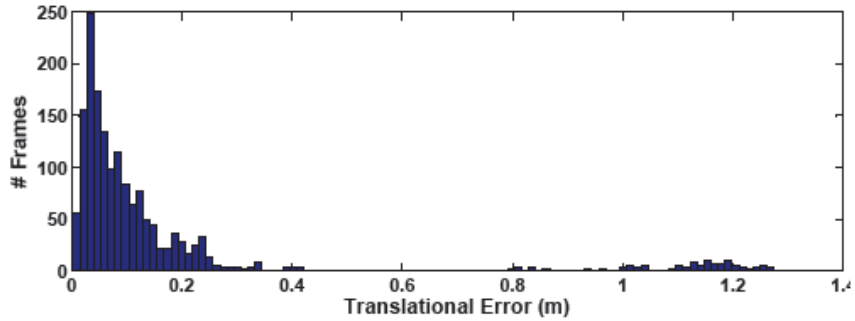
$$k_{11}\lambda_1^2 + (k_{12}\lambda_2 + k_{13})\lambda_1 + (k_{14}\lambda_2^2 + k_{15}\lambda_2 + k_{16}) = 0$$
$$k_{21}\lambda_1^2 + (k_{22}\lambda_3 + k_{23})\lambda_1 + (k_{24}\lambda_3^2 + k_{25}\lambda_3 + k_{26}) = 0$$
$$k_{31}\lambda_2^2 + (k_{32}\lambda_3 + k_{33})\lambda_2 + (k_{34}\lambda_3^2 + k_{35}\lambda_3 + k_{36}) = 0$$

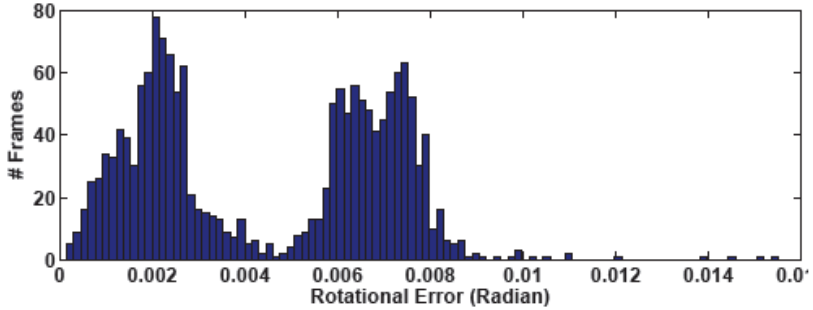
- Solve for λ_1, λ_2, λ_3.
- 8 degree univariate polynomal:

$$A\lambda_3^8 + B\lambda_3^7 + C\lambda_3^6 + D\lambda_3^5 + E\lambda_3^4 + F\lambda_3^3 + G\lambda_3^2 + H\lambda_3 + I = 0$$

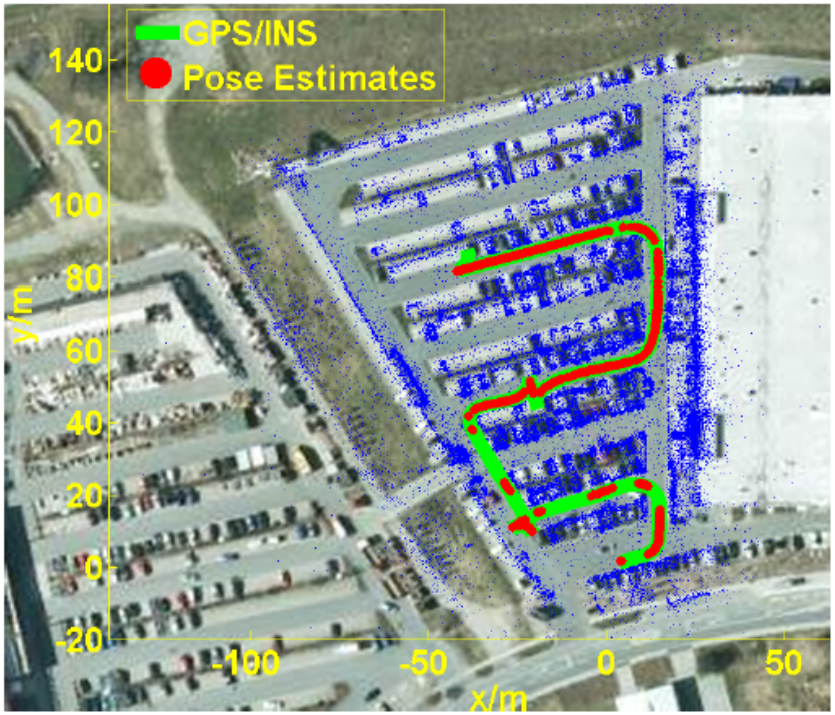
Results: Parking Garage 1

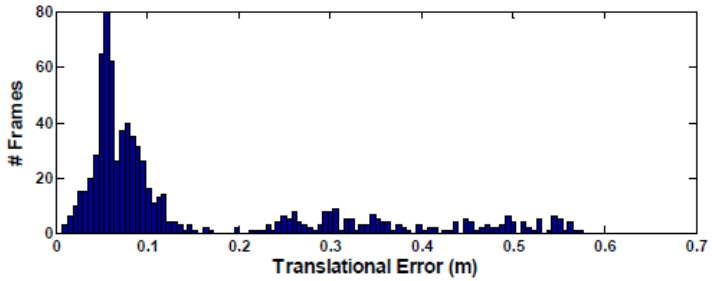


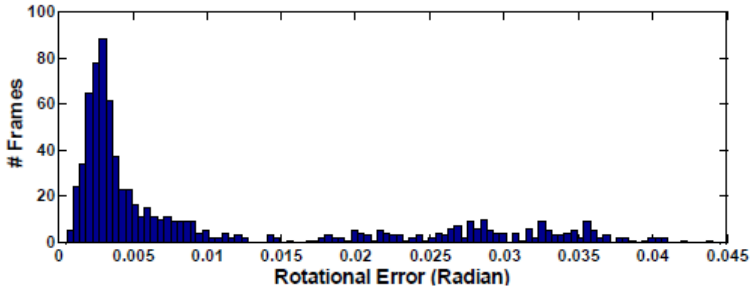




Results: Parking Garage 2







Conclusion

- Sparse mapping:
  1. Motion estimation
  2. Geometric verification
  3. Robust loop-closure
- Non-perspective pose estimation.







Thank You!

Questions?