Robust Camera Location Estimation by Convex Programming

Onur Özyeşil[†] and Amit Singer[‡]

[†]INTECH Investment Management LLC[1]

[‡]PACM and Department of Mathematics, Princeton University

DIMACS Workshop on Distance Geometry

07/28/2016, DIMACS Center, Rutgers University

[1]Work conducted as part of Özyeşil's Ph.D. work at Princeton University

Table of contents

Structure from Motion (SfM) Problem

Given a **collection of** 2**D photos** of a 3D object, **recover the 3D structure by estimating the camera motion**, i.e. **camera locations and orientations**

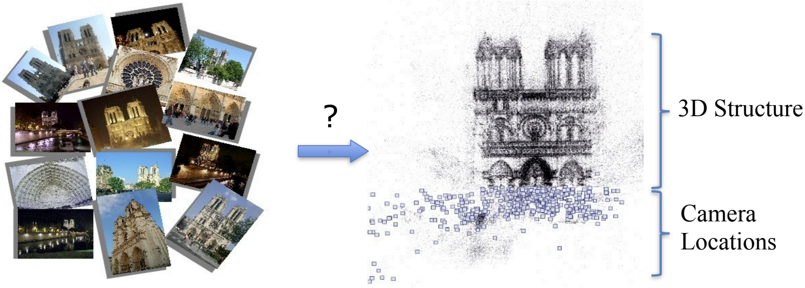

$\Longrightarrow$ We are primarily interested in the **camera location** estimation part

Structure from Motion

Classical Approach

- Find **corresponding points** between images, estimate **relative poses**
- Estimate camera **orientations** and **locations**, i.e. camera motion
- Estimate the **3D structure** (e.g., by reprojection error minimization)

Structure from Motion

Classical Approach

- Find **corresponding points** between images, estimate **relative poses**
- Estimate camera **orientations** and **locations**, i.e. camera motion
- Estimate the **3D structure** (e.g., by reprojection error minimization)

Previous Methods

- **Incremental methods**: Incorporate images one by one (or in small groups) to maintain efficiency $\Rightarrow$ prone to **accumulation of errors**
- **Joint structure and motion estimation**: Computationally hard, usually non-convex methods, no guarantees of convergence to global optima
- **Orientation estimation methods**: Relatively stable and efficient solvers

Structure from Motion

Classical Approach

- Find **corresponding points** between images, estimate **relative poses**
- Estimate camera **orientations** and **locations**, i.e. camera motion
- Estimate the **3D structure** (e.g., by reprojection error minimization)

Previous Methods

- **Incremental methods**: Incorporate images one by one (or in small groups) to maintain efficiency $\Rightarrow$ prone to **accumulation of errors**
- **Joint structure and motion estimation**: Computationally hard, usually non-convex methods, no guarantees of convergence to global optima
- **Orientation estimation methods**: Relatively stable and efficient solvers

Global Location Estimation

- **Ill-conditioned** problem (because of **undetermined relative scales**)
- Current methods: Usually **not well-formulated, not stable, inefficient**

Problem: Location Estimation from Pairwise Directions

- Estimate the locations $\mathbf{t}_1, \mathbf{t}_2, \ldots, \mathbf{t}_n \in \mathbb{R}^d$, for arbitrary $d \geq 2$, from a subset of (noisy) measurements of the pairwise directions, where the direction between $\mathbf{t}_i$ and $\mathbf{t}_j$ is given by the unit norm vector γ_{ij}:

$$\gamma_{ij} = \frac{\mathbf{t}_i - \mathbf{t}_j}{\|\mathbf{t}_i - \mathbf{t}_j\|}$$

Pairwise Directions **Locations**

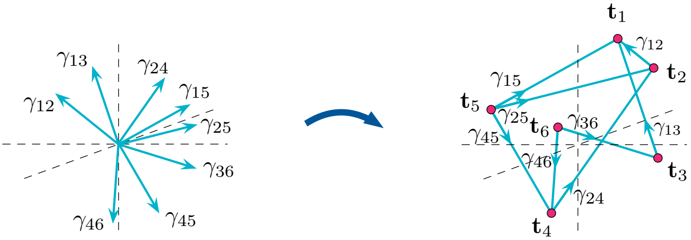

A (noiseless) instance in $\mathbb{R}^3$, with $n = 6$ locations and $m = 8$ pairwise directions.

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Fundamental Questions

- Is the problem well-posed, i.e. do we have enough information to estimate the locations $\{\mathbf{t}_i\}_{i \in V_t}$ stably (or, exactly in the noiseless case) ?

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Fundamental Questions

- Is the problem well-posed, i.e. do we have enough information to estimate the locations $\{\mathbf{t}_i\}_{i \in V_t}$ stably (or, exactly in the noiseless case) ?
- What does well-posedness depend on, is it a (generic) property of G_t?

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Fundamental Questions

- Is the problem well-posed, i.e. do we have enough information to estimate the locations $\{\mathbf{t}_i\}_{i \in V_t}$ stably (or, exactly in the noiseless case) ?
- What does well-posedness depend on, is it a (generic) property of G_t?
- Can it be decided efficiently?

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Fundamental Questions

- Is the problem well-posed, i.e. do we have enough information to estimate the locations $\{\mathbf{t}_i\}_{i \in V_t}$ stably (or, exactly in the noiseless case) ?
- What does well-posedness depend on, is it a (generic) property of G_t?
- Can it be decided efficiently?
- What can we do if an instance is not well-posed?

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Fundamental Questions

- Is the problem well-posed, i.e. do we have enough information to estimate the locations $\{\mathbf{t}_i\}_{i \in V_t}$ stably (or, exactly in the noiseless case) ?
- What does well-posedness depend on, is it a (generic) property of G_t?
- Can it be decided efficiently?
- What can we do if an instance is not well-posed?

Well-posedness of the Location Estimation Problem

• We represent the total pairwise information using a graph $G_t = (V_t, E_t)$ and endow each edge (i, j) with the direction measurement γ_{ij}

Fundamental Questions

- Is the problem well-posed, i.e. do we have enough information to estimate the locations $\{\mathbf{t}_i\}_{i \in V_t}$ stably (or, exactly in the noiseless case) ?
- What does well-posedness depend on, is it a (generic) property of G_t?
- Can it be decided efficiently?
- What can we do if an instance is not well-posed?

$\Rightarrow$ Well-posedness was previously studied in various contexts, under the general title of **parallel rigidity theory**.

Well-posedness: Simple Examples

Consider the following noiseless instance:

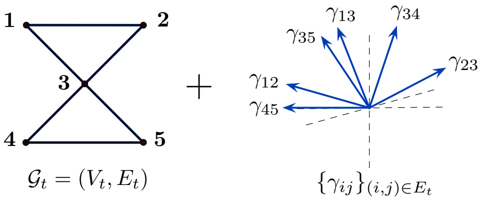

Well-posedness: Simple Examples

Consider the following noiseless instance:

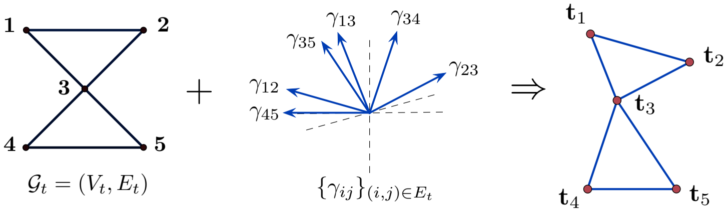

Well-posedness: Simple Examples

Consider the following noiseless instance:

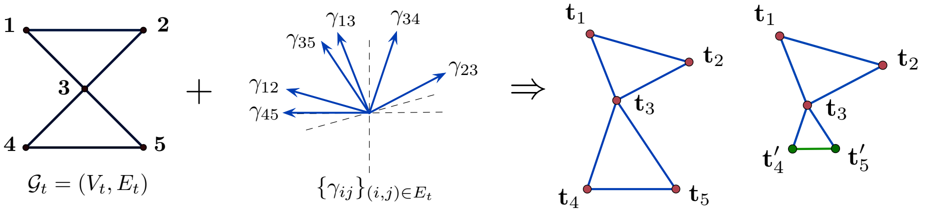

Well-posedness: Simple Examples

Consider the following noiseless instance:

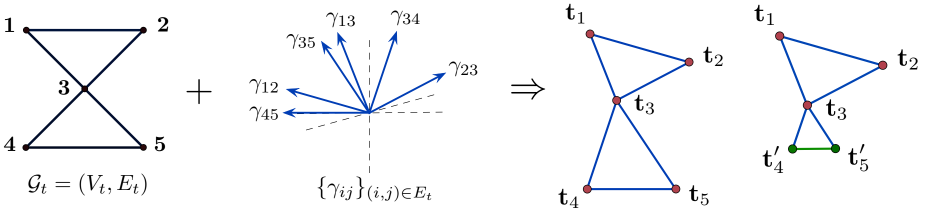

Well-posedness depends on the dimension:

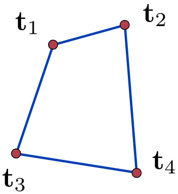

$\Rightarrow$ Well-posed in $\mathbb{R}^3$, but not in $\mathbb{R}^2$

Main Results of Parallel Rigidity Theory

Unique Realization of Locations

Unique solution exists if and only if the formation is **parallel rigid**.

Main Results of Parallel Rigidity Theory

Unique Realization of Locations

Unique solution exists if and only if the formation is **parallel rigid**.

Rigidity is Generic

Parallel rigidity is a generic property, i.e. it is **a function of G_t only**.

Main Results of Parallel Rigidity Theory

Unique Realization of Locations

Unique solution exists if and only if the formation is **parallel rigid**.

Rigidity is Generic

Parallel rigidity is a generic property, i.e. it is **a function of G_t only**.

Theorem (Efficient Decidability, Whiteley, 1987)

For a graph $G = (V, E)$, let $(d-1)E$ denote the set consisting of $(d-1)$ copies of each edge in E. Then, G is generically parallel rigid in $\mathbb{R}^d$ if and only if there exists a nonempty set $D \subseteq (d-1)E$, with $|D| = d|V| - (d+1)$, such that for all subsets D' of D,

$$|D'| \leq d|V(D')| - (d+1) \ ,$$

where $V(D')$ denotes the vertex set of the edges in D'.

Main Results of Parallel Rigidity Theory

Unique Realization of Locations

Unique solution exists if and only if the formation is **parallel rigid**.

Rigidity is Generic

Parallel rigidity is a generic property, i.e. it is **a function of G_t only**.

Theorem (Efficient Decidability, Whiteley, 1987)

For a graph $G = (V, E)$, let $(d-1)E$ denote the set consisting of $(d-1)$ copies of each edge in E. Then, G is generically parallel rigid in $\mathbb{R}^d$ if and only if there exists a nonempty set $D \subseteq (d-1)E$, with $|D| = d|V| - (d+1)$, such that for all subsets D' of D,

$$|D'| \leq d|V(D')| - (d+1) \ ,$$

where $V(D')$ denotes the vertex set of the edges in D'.

Maximally Parallel Rigid Components

If G_t is **not parallel rigid**, we can efficiently find **maximally parallel rigid subgraphs**.

Robust Location Estimation from Pairwise Directions

Objective

Minimize the effects of direction measurements with large errors, i.e. **outlier directions**, in location estimation, by maintaining **computational efficiency**

Robust Location Estimation from Pairwise Directions

Objective

Minimize the effects of direction measurements with large errors, i.e. **outlier directions**, in location estimation, by maintaining **computational efficiency**

Recipe

Formulation based on a **robust cost function**, approximate by **convex programming**

Robust Location Estimation from Pairwise Directions

Objective

Minimize the effects of direction measurements with large errors, i.e. **outlier directions**, in location estimation, by maintaining **computational efficiency**

Recipe

Formulation based on a **robust cost function**, approximate by **convex programming**

Linearization of Pairwise Directions

$$\gamma_{ij} = \frac{\mathbf{t}_i - \mathbf{t}_j}{\|\mathbf{t}_i - \mathbf{t}_j\|} + \epsilon_{ij}^{\gamma} \;\Leftrightarrow\; \epsilon_{ij}^{\mathbf{t}} = \mathbf{t}_i - \mathbf{t}_j - d_{ij}\gamma_{ij}$$

where $d_{ij} = \|\mathbf{t}_i - \mathbf{t}_j\|$ and $\epsilon_{ij}^{\mathbf{t}} = \|\mathbf{t}_i - \mathbf{t}_j\|\epsilon_{ij}^{\gamma}$

Robust Location Estimation from Pairwise Directions

Objective

Minimize the effects of direction measurements with large errors, i.e. **outlier directions**, in location estimation, by maintaining **computational efficiency**

Recipe

Formulation based on a **robust cost function**, approximate by **convex programming**

Linearization of Pairwise Directions

$$\gamma_{ij} = \frac{\mathbf{t}_i - \mathbf{t}_j}{\|\mathbf{t}_i - \mathbf{t}_j\|} + \epsilon_{ij}^{\gamma} \;\Leftrightarrow\; \epsilon_{ij}^{\mathbf{t}} = \mathbf{t}_i - \mathbf{t}_j - d_{ij}\gamma_{ij}$$

where $d_{ij} = \|\mathbf{t}_i - \mathbf{t}_j\|$ and $\epsilon_{ij}^{\mathbf{t}} = \|\mathbf{t}_i - \mathbf{t}_j\|\epsilon_{ij}^{\gamma}$

Idea

Large ϵ_{ij}^{γ} induces large $\epsilon_{ij}^{\mathbf{t}}$ $\Rightarrow$ exchange robustness to ϵ_{ij}^{γ} with robustness to $\epsilon_{ij}^{\mathbf{t}}$

Least Unsquared Deviations (LUD) formulation

Relax the non-convex constraint $d_{ij} = \|\mathbf{t}_i - \mathbf{t}_j\|$ to obtain:

$$\begin{aligned} \underset{\substack{\{\mathbf{t}_i\}_{i \in V_t} \subseteq \mathbb{R}^d \\ \{d_{ij}\}_{(i,j) \in E_t}}}{\text{minimize}} \quad & \sum_{(i,j) \in E_t} \|\mathbf{t}_i - \mathbf{t}_j - d_{ij}\gamma_{ij}\| \\ \text{subject to} \quad & \sum_{i \in V_t} \mathbf{t}_i = \mathbf{0} \; ; \; d_{ij} \geq c, \; \forall (i,j) \in E_t \end{aligned}$$

- Inspired by convex programs developed for robust signal recovery in the presence of outliers, exact signal recovery from partially corrupted data
- Prevents collapsing solutions via the constraint $d_{ij} \geq c$

IRLS for LUD

Initialize: $w_{ij}^0 = 1, \forall (i,j) \in E_t$

for $r = 0, 1, \ldots$ **do**

- Compute $\{\hat{\mathbf{t}}_i^{r+1}\}, \{\hat{d}_{ij}^{r+1}\}$ by solving the QP:

$$\underset{\left\{\substack{\sum \mathbf{t}_i = \mathbf{0}, \\ d_{ij} \geq 1}\right\}}{\text{minimize}} \sum_{(i,j) \in E_t} w_{ij}^r \left\| \mathbf{t}_i - \mathbf{t}_j - d_{ij}\gamma_{ij} \right\|^2$$

- $w_{ij}^{r+1} \leftarrow \left(\left\| \hat{\mathbf{t}}_i^{r+1} - \hat{\mathbf{t}}_j^{r+1} - \hat{d}_{ij}^{r+1}\gamma_{ij} \right\|^2 + \delta \right)^{-1/2}$

Robust Pairwise Direction Estimation

Pinhole Camera Model

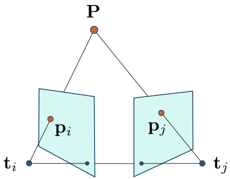

For a camera $C_i = (R_i, \mathbf{t}_i, f_i)$ and $\mathbf{P} \in \mathbb{R}^3$:

- Represent $\mathbf{P}$ in i'th coordinate system:
  $\mathbf{P}_i = R_i^T(\mathbf{P} - \mathbf{t}_i) = (P_i^x, P_i^y, P_i^z)^T$
- Project onto i'th image plane:
  $\mathbf{q}_i = (f_i/P_i^z)(P_i^x, P_i^y)^T \in \mathbb{R}^2$

Robust Pairwise Direction Estimation

Pinhole Camera Model

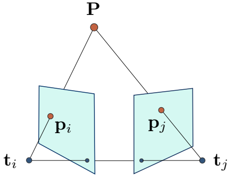

For a camera $C_i = (R_i, \mathbf{t}_i, f_i)$ and $\mathbf{P} \in \mathbb{R}^3$:

- Represent $\mathbf{P}$ in i'th coordinate system:
  $\mathbf{P}_i = R_i^T(\mathbf{P} - \mathbf{t}_i) = (P_i^x, P_i^y, P_i^z)^T$
- Project onto i'th image plane:
  $\mathbf{q}_i = (f_i/P_i^z)(P_i^x, P_i^y)^T \in \mathbb{R}^2$

Pairwise Directions from Epipolar Constraints

For a pair (C_i, C_j) with given (R_i, R_j), (f_i, f_j) and a set $\{(\mathbf{q}_i^k, \mathbf{q}_j^k)\}_{k=1}^{m_{ij}}$ of corresponding points:

- Fact: $\mathbf{P}, \mathbf{t}_i, \mathbf{t}_j$ are **coplanar** (equiv. to epipolar constraint)

$$[(\mathbf{P} - \mathbf{t}_i) \times (\mathbf{P} - \mathbf{t}_j)]^T(\mathbf{t}_i - \mathbf{t}_j) = 0$$

$$\Leftrightarrow (R_i\eta_i^k \times R_j\eta_j^k)^T(\mathbf{t}_i - \mathbf{t}_j) = 0 \quad (\text{for } \eta_i^k := \begin{bmatrix} \mathbf{q}_i^k/f_i \\ 1 \end{bmatrix})$$

$$\Leftrightarrow \boxed{(\nu_{ij}^k)^T(\mathbf{t}_i - \mathbf{t}_j) = 0, \quad (\text{for } \nu_{ij}^k := \frac{R_i\eta_i^k \times R_j\eta_j^k}{\|R_i\eta_i^k \times R_j\eta_j^k\|})}$$

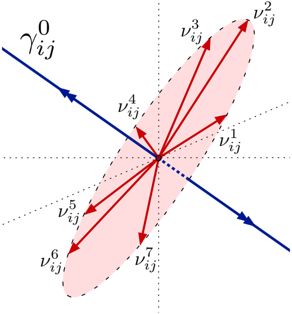

Robust Directions from Noisy 2D Subspace Samples

- Given noisy R_i's, f_i's and $\mathbf{q}_i^k$'s, we essentially obtain **noisy samples** $\hat{\nu}_{ij}^k$'s from the 2D subspace orthogonal to $\mathbf{t}_i - \mathbf{t}_j$

Robust Directions from Noisy $2D$ Subspace Samples

- Given noisy R_i's, f_i's and $\mathbf{q}_i^k$'s, we essentially obtain **noisy samples** $\hat{\nu}_{ij}^k$'s from the $2D$ subspace orthogonal to $\mathbf{t}_i - \mathbf{t}_j$

- To maintain **robustness to outliers** among $\hat{\nu}_{ij}^k$'s, we estimate the lines γ_{ij}^0 using the non-convex problem (solved via a heuristic IRLS method):

$$\begin{aligned} \underset{\gamma_{ij}^0 \in \mathbb{R}^3}{\text{minimize}} \quad & \sum_{k=1}^{m_{ij}} |\langle \gamma_{ij}^0, \hat{\nu}_{ij}^k \rangle| \\ \text{subject to} \quad & \|\gamma_{ij}^0\| = 1 \end{aligned}$$

Robust Directions from Noisy 2D Subspace Samples

- Given noisy R_i's, f_i's and $\mathbf{q}_i^k$'s, we essentially obtain **noisy samples** $\hat{\nu}_{ij}^k$'s from the 2D subspace orthogonal to $\mathbf{t}_i - \mathbf{t}_j$
- To maintain **robustness to outliers** among $\hat{\nu}_{ij}^k$'s, we estimate the lines γ_{ij}^0 using the non-convex problem (solved via a heuristic IRLS method):

$$\underset{\gamma_{ij}^0 \in \mathbb{R}^3}{\text{minimize}} \sum_{k=1}^{m_{ij}} |\langle \gamma_{ij}^0, \hat{\nu}_{ij}^k \rangle| \quad \text{subject to} \quad \|\gamma_{ij}^0\| = 1$$

- The heuristic IRLS method is not guaranteed to converge (because of non-convexity)
- However, we empirically observed **high quality line estimates**
- Computationally much more efficient (relative to previous methods with similar accuracy)
- Pairwise directions are computed from the lines using the fact that the 3D points should lie in front of the cameras.

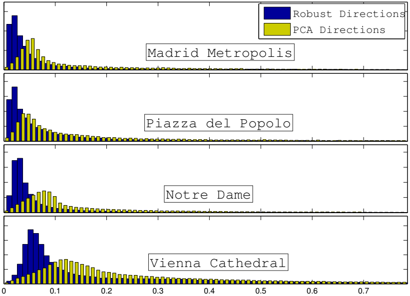

Histogram plots of direction errors as compared to simpler PCA method (for datasets from [WS14])

[WS14] K. Wilson and N. Snavely, *Robust global translations with 1DSfM*, ECCV 2014.

Synchetic Data Experiments

Measurement Model

- Measurement graphs $G_t = (V_t, E_t)$ are random graphs drawn from Erdös-Rényi model, i.e. each (i, j) is in the edge set E_t with probability q, independently of all other edges.
- Noise model: Given a set of locations $\{\mathbf{t}_i\}_{i=1}^n \subseteq \mathbb{R}^d$ and $G_t = (V_t, E_t)$, for each $(i, j) \in E_t$, we let $\gamma_{ij} = \tilde{\gamma}_{ij} / \|\tilde{\gamma}_{ij}\|$, where

$$\tilde{\gamma}_{ij} = \begin{cases} \gamma_{ij}^U \quad , & \text{w.p.} \quad p \\ (\mathbf{t}_i - \mathbf{t}_j)/\|\mathbf{t}_i - \mathbf{t}_j\| + \sigma \gamma_{ij}^G & \text{w.p.} \ 1 - p \end{cases}$$

Here, $\{\gamma_{ij}^U\}_{(i,j)\in E_t}$ are i.i.d. Unif(S^{d-1}), $\{\gamma_{ij}^G\}_{(i,j)\in E_t}$ and $\mathbf{t}_i$'s are i.i.d. $\mathcal{N}(\mathbf{0}, I_3)$.

Synthetic Data: Relatively High Robustness to Outliers

Performance Measure

Normalized root mean square error (NRMSE) of the estimates $\hat{\mathbf{t}}_i$ w.r.t. the original locations $\mathbf{t}_i$ (after the removal of global scale and translation, and for $\mathbf{t}_0$ denoting the center of $\mathbf{t}_i$'s.)

$$\text{NRMSE}(\{\hat{\mathbf{t}}_i\}) = \sqrt{\frac{\sum_i \|\hat{\mathbf{t}}_i - \mathbf{t}_i\|^2}{\sum_i \|\mathbf{t}_i - \mathbf{t}_0\|^2}}$$

Synthetic Data: Relatively High Robustness to Outliers

Performance Measure

Normalized root mean square error (NRMSE) of the estimates $\hat{\mathbf{t}}_i$ w.r.t. the original locations $\mathbf{t}_i$ (after the removal of global scale and translation, and for $\mathbf{t}_0$ denoting the center of $\mathbf{t}_i$'s.)

$$\text{NRMSE}(\{\hat{\mathbf{t}}_i\}) = \sqrt{\frac{\sum_i \|\hat{\mathbf{t}}_i - \mathbf{t}_i\|^2}{\sum_i \|\mathbf{t}_i - \mathbf{t}_0\|^2}}$$

Methods for comparison:

- **Least Squares** (LS) method
[AKK12] M. Arie-Nachimson, S. Kovalsky, I. Kemelmacher-Shlizerman, AS, and R. Basri, *Global motion estimation from point matches*, 3DimPVT, 2012.
[BAT04] M. Brand, M. Antone, and S. Teller, *Spectral solution of large-scale extrinsic camera calibration as a graph embedding problem*, ECCV, 2004.
- **Constrained Least Squares** (CLS) method
[TV14] R. Tron and R. Vidal, *Distributed 3-D localization of camera sensor networks from 2-D image measurements*, IEEE Trans. on Auto. Cont., 2014.
- **Semidefinite Relaxation** (SDR) method
[OSB15] OÖ, AS, and R. Basri, *Stable camera motion estimation using convex programming*, SIAM J. Imaging Sci., 8(2):1220-1262, 2015.

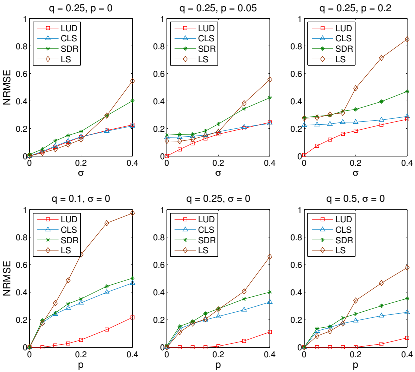

Exact Recovery with Partially Corrupted Directions

- Empirical observation: LUD can recover locations exactly with partially corrupted data

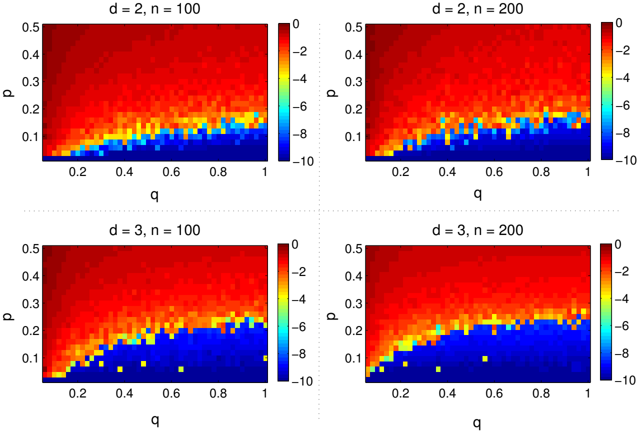

The color intensity of each pixel represents $\log_{10}$(NRMSE), depending on the edge probability q (x-axis), and the outlier probability p (y-axis). NRMSE values are averaged over 10 trials.

Real Data Experiments: Internet Photo Collections

We tested our methods on internet photo collections from [WS14] and observed that:

- The LUD solver is **more robust to outliers** as compared to existing methods, and is **more efficient** as compared to methods with similar accuracy
- The robust direction estimation by IRLS further improves robustness to outliers

Montreal Notre Dame

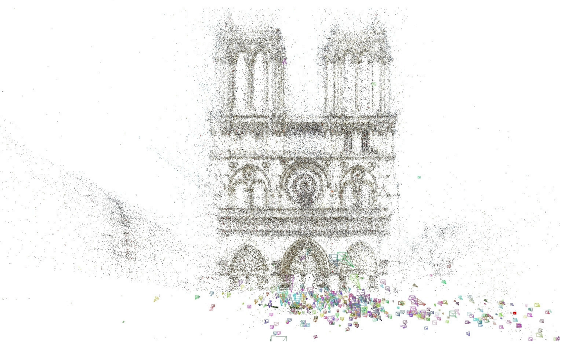

Notre Dame

Snapshots of selected 3D structures computed using the LUD solver (before bundle adjustment)

Real Datasets: Comparison of Estimation Accuracy and Efficiency

Estimation errors in meters (Bundler sequential SfM [SSS06] is taken as ground truth):

Dataset			LUD							CLS [TV14]					SDR [OSB15]					1DSfM [WS14]				LS [G01]	
Name	Size		Initial				After BA			Initial		After BA			Initial		After BA			Init.	After BA			After BA	
			PCA		Robust		Robust			Robust		Robust			Robust		Robust								
	m	N_c	$\tilde{e}$	$\hat{e}$	$\tilde{e}$	$\hat{e}$	N_c	$\tilde{e}$	$\hat{e}$	$\tilde{e}$	$\hat{e}$	N_c	$\tilde{e}$	$\hat{e}$	$\tilde{e}$	$\hat{e}$	N_c	$\tilde{e}$	$\hat{e}$	$\tilde{e}$	N_c	$\tilde{e}$	$\hat{e}$	N_c	$\tilde{e}$
Piazza del Popolo	60	328	3.0	7	**1.5**	**5**	305	**1.0**	**4**	3.5	6	305	1.4	5	1.9	8	305	1.3	7	3.1	308	2.2	200	93	16
NYC Library	130	332	4.9	9	**2.0**	**6**	320	1.4	7	5.0	8	320	3.9	8	5.0	8	320	4.6	8	2.5	295	**0.4**	**1**	271	1.4
Metropolis	200	341	4.3	8	**1.6**	**4**	288	1.5	**4**	6.4	10	288	3.1	7	4.2	8	288	3.1	7	9.9	291	**0.5**	70	240	18
Yorkminster	150	437	5.4	10	**2.7**	**5**	404	1.3	**4**	6.2	9	404	2.9	8	5.0	10	404	4.0	10	3.4	401	**0.1**	500	345	6.7
Tower of London	300	572	12	25	**4.7**	**20**	425	3.3	**10**	16	30	425	15	30	20	30	425	17	30	11	414	**1.0**	40	306	44
Montreal N. D.	30	450	1.4	2	**0.5**	**1**	435	**0.4**	**1**	1.1	2	435	0.5	**1**	–	–	–	–	–	2.5	427	**0.4**	**1**	357	9.8
Notre Dame	300	553	1.1	2	**0.3**	**0.8**	536	**0.2**	**0.7**	0.8	2	536	0.3	0.9	–	–	–	–	–	10	507	1.9	7	473	2.1
Alamo	70	577	1.5	3	**0.4**	**2**	547	**0.3**	**2**	1.3	3	547	0.6	**2**	–	–	–	–	–	1.1	529	**0.3**	2e7	422	2.4
Vienna Cathedral	120	836	7.2	12	**5.4**	**10**	750	4.4	**10**	8.8	**10**	750	8.2	**10**	–	–	–	–	–	6.6	770	**0.4**	2e4	652	12

Running times in seconds:

Dataset				LUD			CLS [TV14]			SDR [OSB15]			1DSfM [WS14]					[G01]	[SSS06]
	T_R	T_G	T_γ	$T_\mathbf{t}$	T_{BA}	T_{tot}	$T_\mathbf{t}$	T_{BA}	T_{tot}	$T_\mathbf{t}$	T_{BA}	T_{tot}	T_R	T_γ	$T_\mathbf{t}$	T_{BA}	T_{tot}	T_{tot}	T_{tot}
Piazza del Popolo	35	43	18	35	31	162	9	106	211	358	39	493	14	9	35	191	249	138	1287
NYC Library	27	44	18	57	54	200	7	47	143	462	52	603	9	13	54	392	468	220	3807
Metropolis	27	37	13	27	38	142	6	23	106	181	45	303	15	8	20	201	244	139	1315
Yorkminster	19	46	33	51	148	297	10	133	241	648	75	821	11	18	93	777	899	394	3225
Tower of London	24	54	23	41	86	228	8	202	311	352	170	623	9	14	55	606	648	264	1900
Montreal N. D.	68	115	91	112	167	553	21	270	565	–	–	–	17	22	75	1135	1249	424	2710
Notre Dame	135	214	325	247	126	1047	52	504	1230	–	–	–	53	42	59	1445	1599	1193	6154
Alamo	103	232	96	186	133	750	40	339	810	–	–	–	56	29	73	752	910	1403	1654
Vienna Cathedral	267	472	265	255	208	1467	46	182	1232	–	–	–	98	60	144	2837	3139	2273	10276

[SSS06] N. Snavely, S. M. Seitz, and R. Szeliski, *Photo tourism: exploring photo collections in 3D*, SIGGRAPH, 2006.
[G01] V. M. Govindu, *Combining two-view constraints for motion estimation*, CVPR, 2001.

THANK YOU!!!

Acknowledgements:

AFOSR
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

SIMONS
FOUNDATION