

Visual SLAM for Mobile

Instructor - Simon Lucey

16-623 - Designing Computer Vision Apps

Example of SLAM for AR



Our method with rotational velocity estimation



ORB-SLAM



PTAM



RDSLAM

Example of SLAM for AR



Our method with rotational velocity estimation



ORB-SLAM



PTAM



RDSLAM

Example of SLAM for AR



Our method with rotational velocity estimation



ORB-SLAM



PTAM



RDSLAM

What is SLAM??

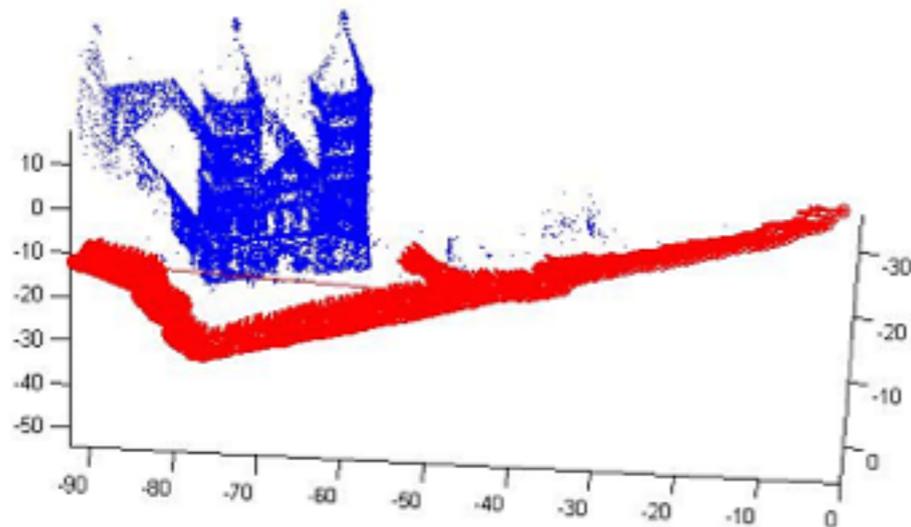
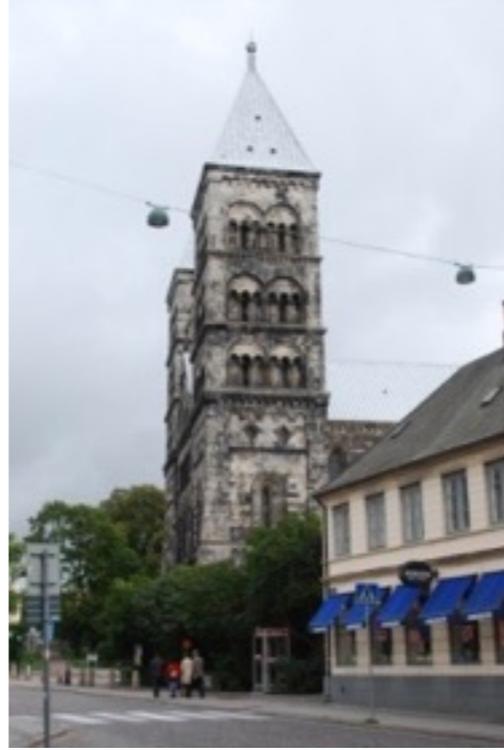
- **S**imultaneous **L**ocalization **a**nd **M**apping.
- On mobile interested primarily in **V**isual **SLAM** (VSLAM).
- Sometimes called Mono SLAM if there is only one camera.
- Can be viewed as an online SfM problem.



Today

- SfM - Bundle Adjustment
- VSLAM - Keyframe vs. Filtering
- Visual Odometry
- Loop Closure

Reminder - Bundle Adjustment



The cathedral dataset:

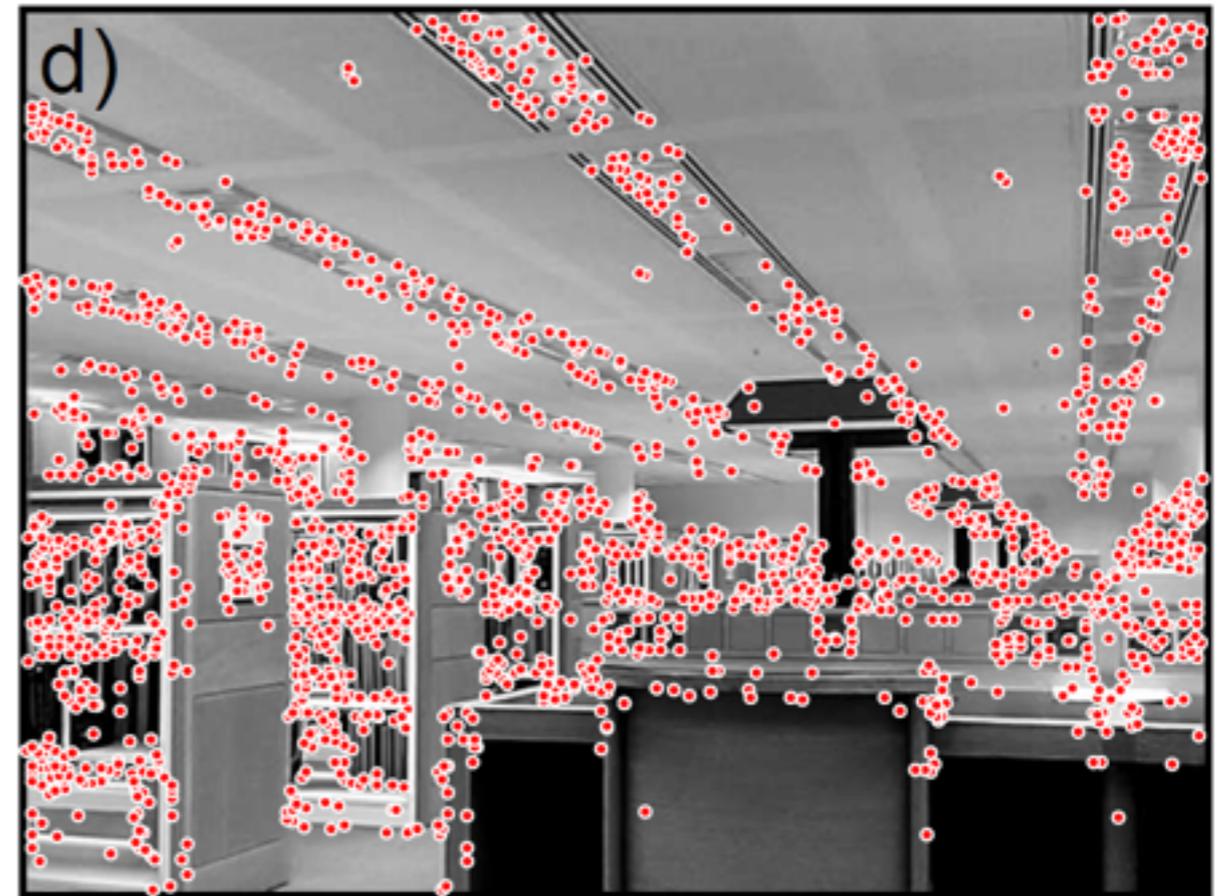
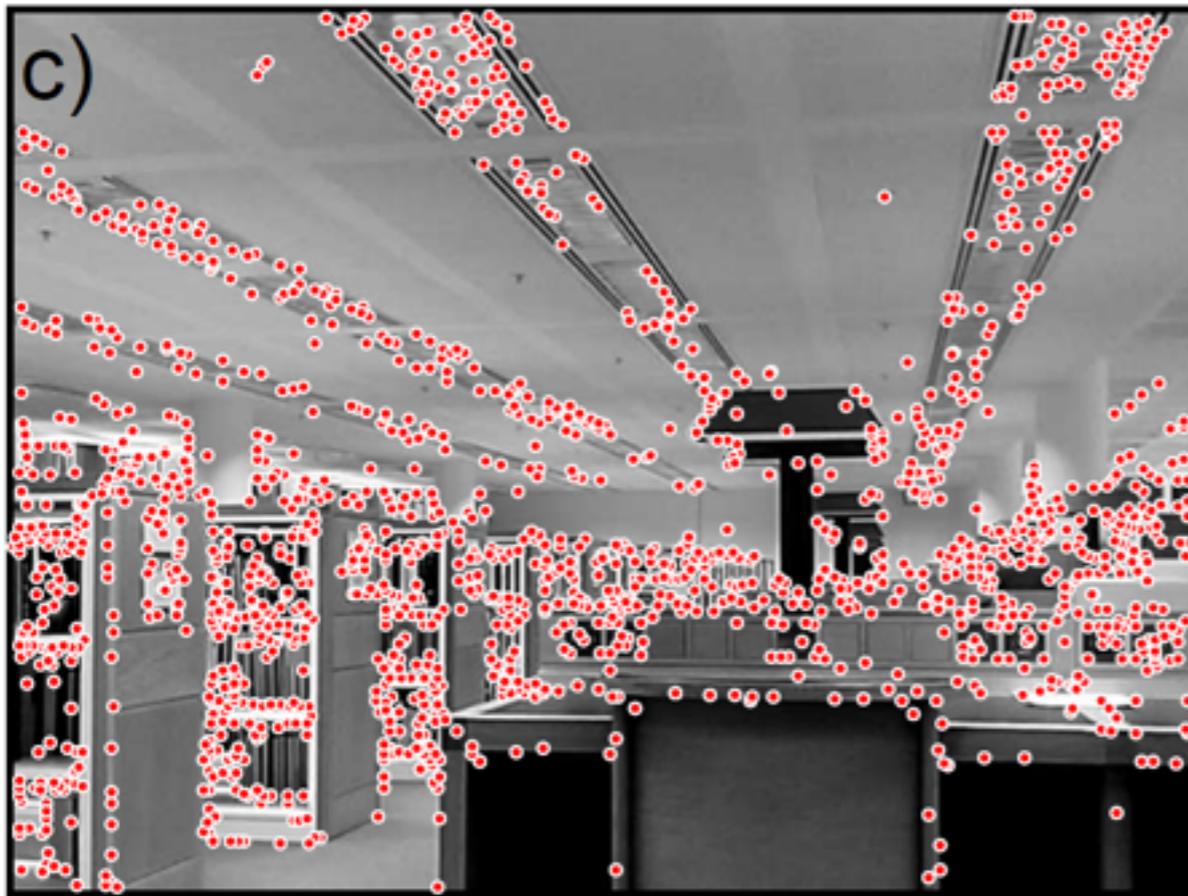
- 480 camera matrices $[\Omega_i, \tau_i]$
- Total dof = $480 \times (3 + 3) = 2880$
- 91178 3D points.
- Total dof = $91178 \times 3 = 273543$

Reminder - Two view reconstruction



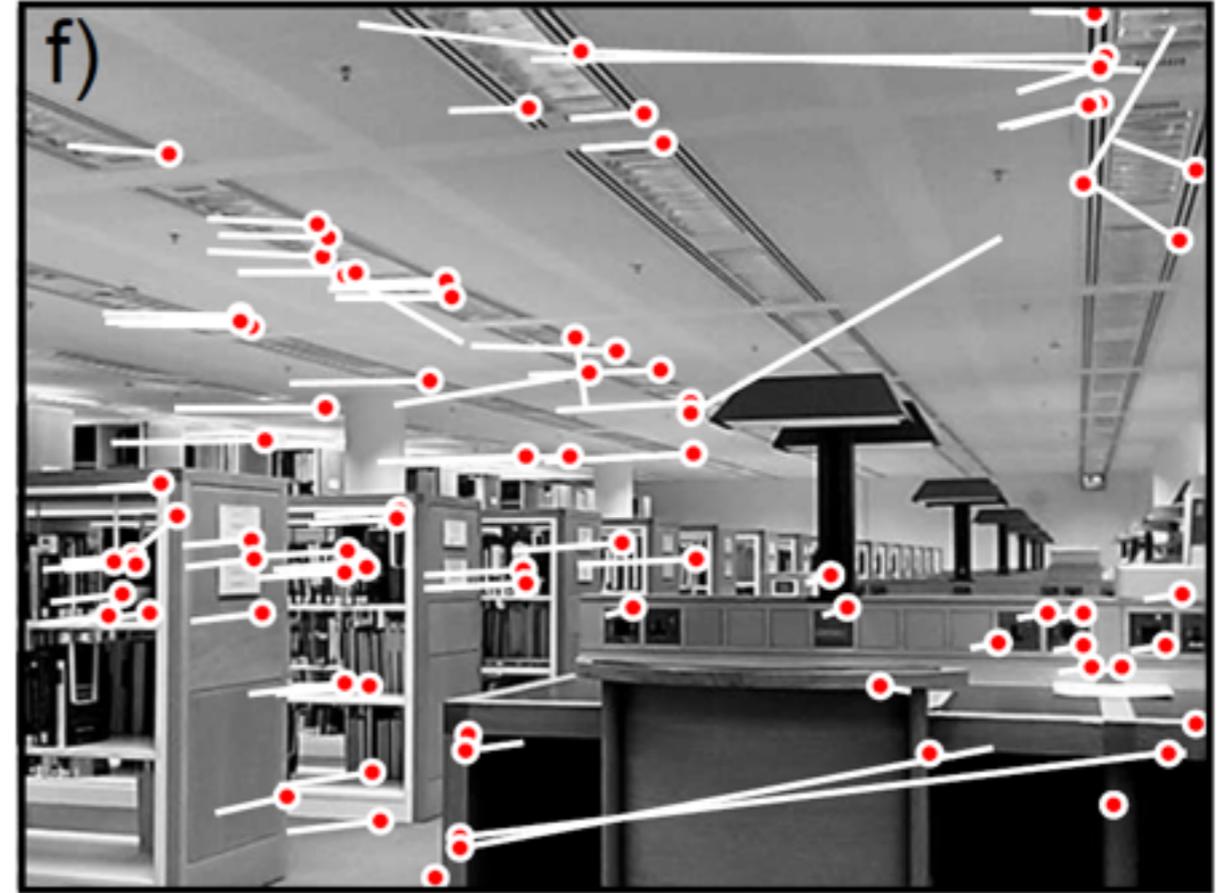
Start with pair of images taken from slightly different viewpoints

Reminder - Two view reconstruction



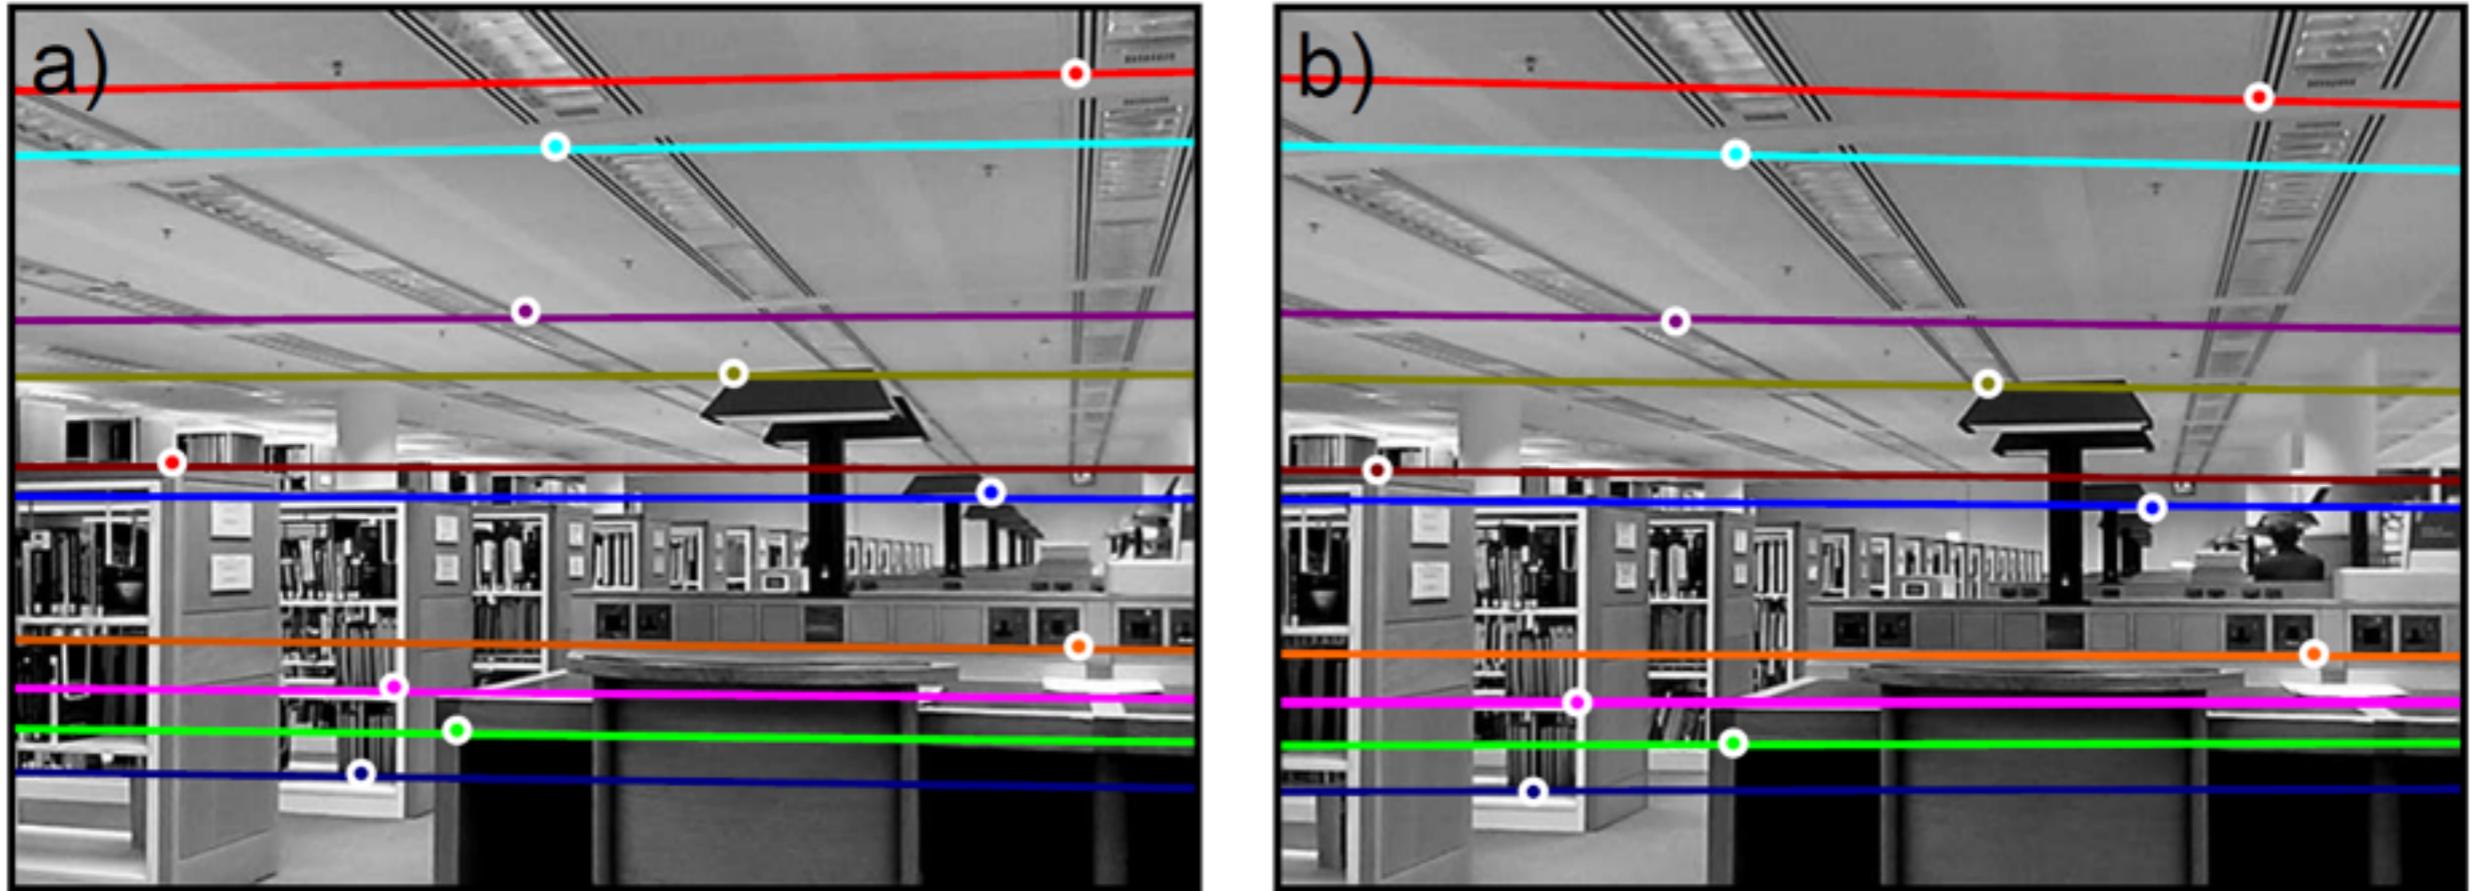
Find features using a corner detection algorithm

Reminder - Two view reconstruction



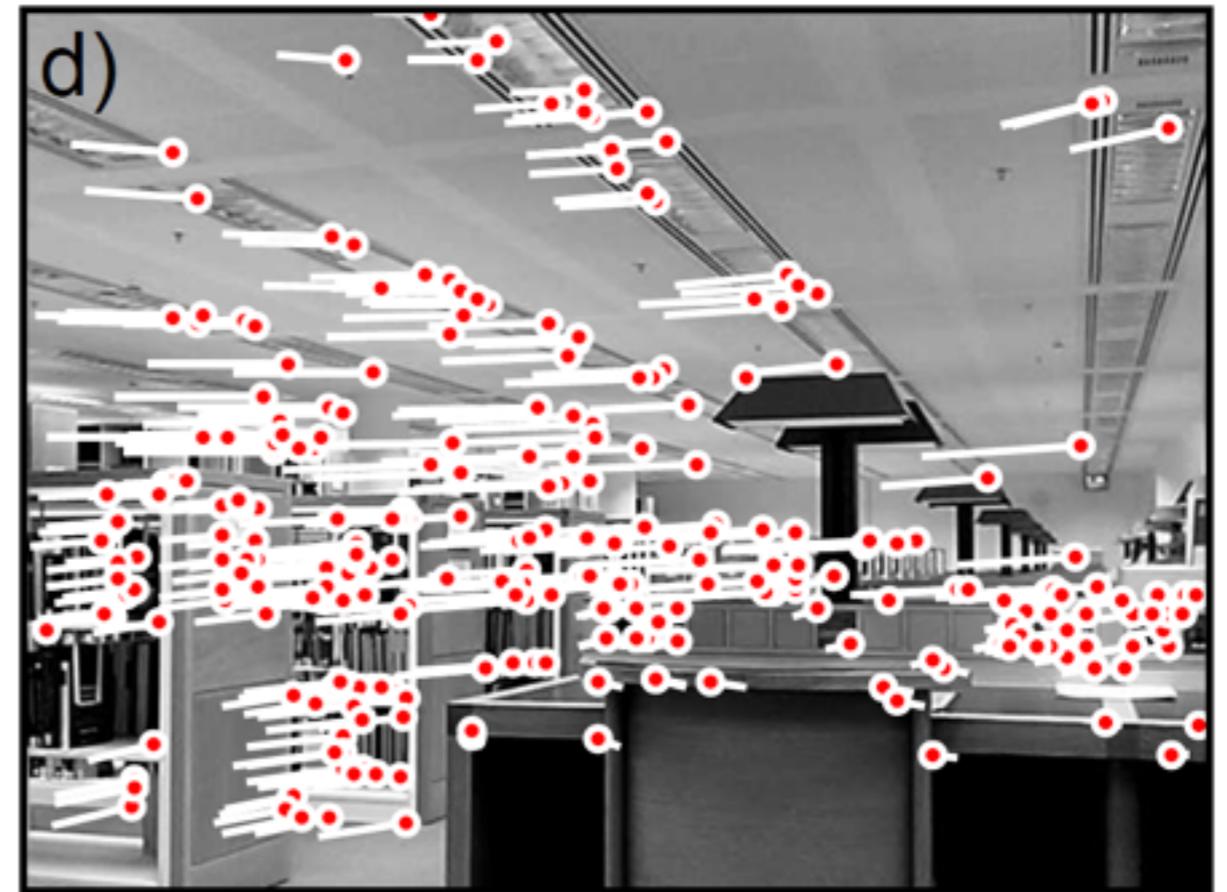
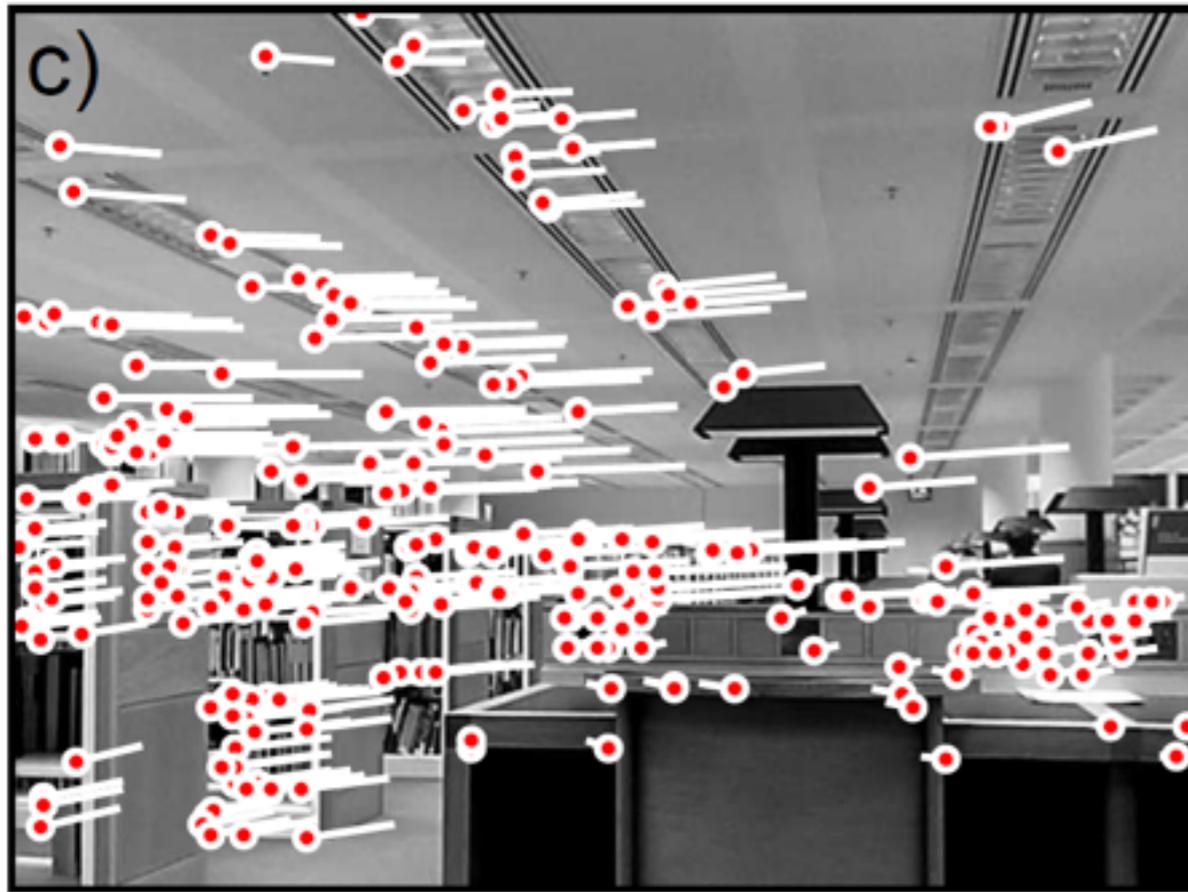
Match features using a greedy algorithm

Reminder - Two view reconstruction



Fit fundamental matrix using robust algorithm such as RANSAC

Reminder - Two view reconstruction



Find matching points that agree with the fundamental matrix

Reminder - Two view reconstruction

- Extract essential matrix from fundamental matrix.
- Extract Ω rotation and τ translation from essential matrix.
- Reconstruct the 3D positions \mathbf{w} of points.

$$\lambda \tilde{\mathbf{x}} = \Omega \mathbf{w} + \tau$$

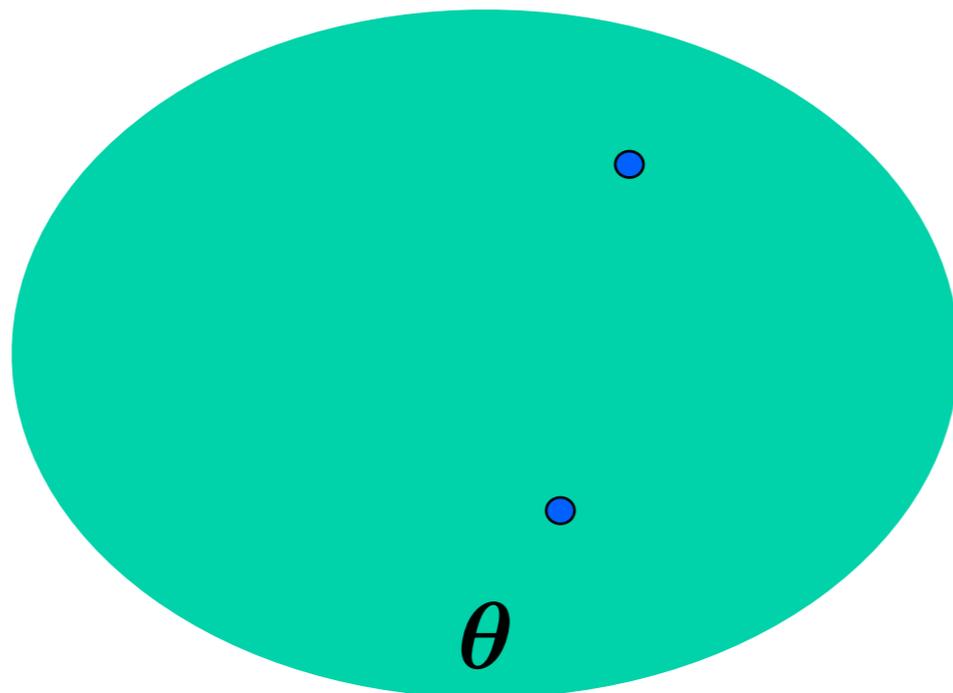
- We refer to these matrices as belonging to the **S**pecial **E**uclidean **G**roup - SE(3).

$$\mathbf{T} = \begin{bmatrix} \Omega & \tau \\ \mathbf{0}^T & 1 \end{bmatrix} \in \text{SE}(3)$$

Reminder: Lie Algebra

- Exponential maps on the $SO(3)$, $SL(3)$ and $SE(3)$ groups are related to the much broader topic of Lie Algebra.
- More details on this topic can be found at in Murray et al. 1994.

$$\mathbf{T} = \begin{bmatrix} \boldsymbol{\Omega} & \boldsymbol{\tau} \\ \mathbf{0}^T & 1 \end{bmatrix} \in SE(3)$$

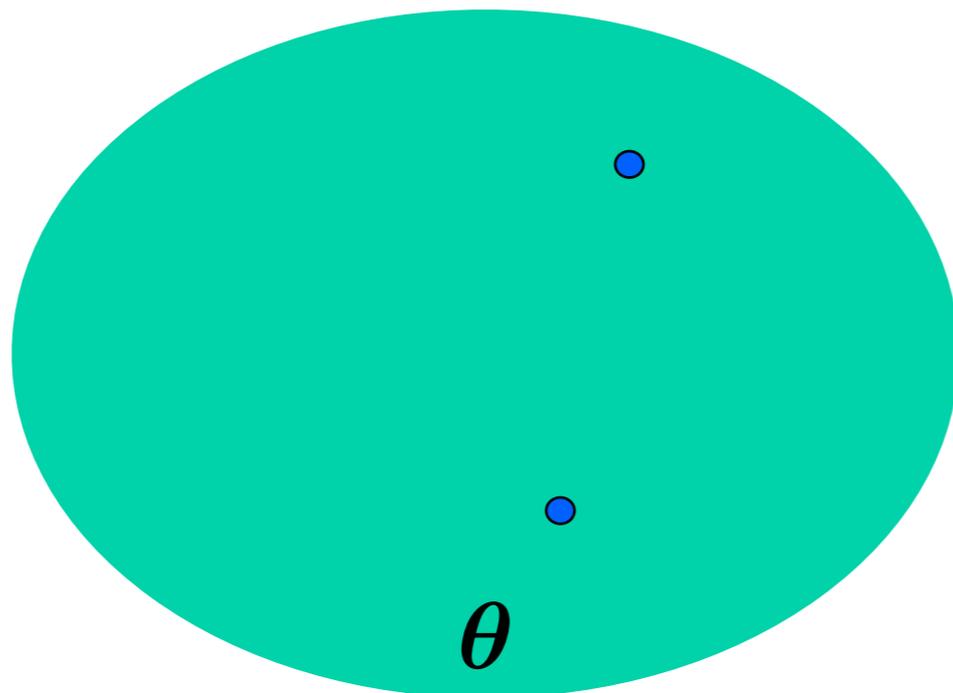


“Sophus Lie”

Reminder: Lie Algebra

- Exponential maps on the $SO(3)$, $SL(3)$ and $SE(3)$ groups are related to the much broader topic of Lie Algebra.
- More details on this topic can be found at in Murray et al. 1994.

$$\mathbf{T}(\boldsymbol{\theta}) = \exp \left(\sum_{i=1}^6 \theta_i \mathbf{A}_i \right) \in SE(3)$$

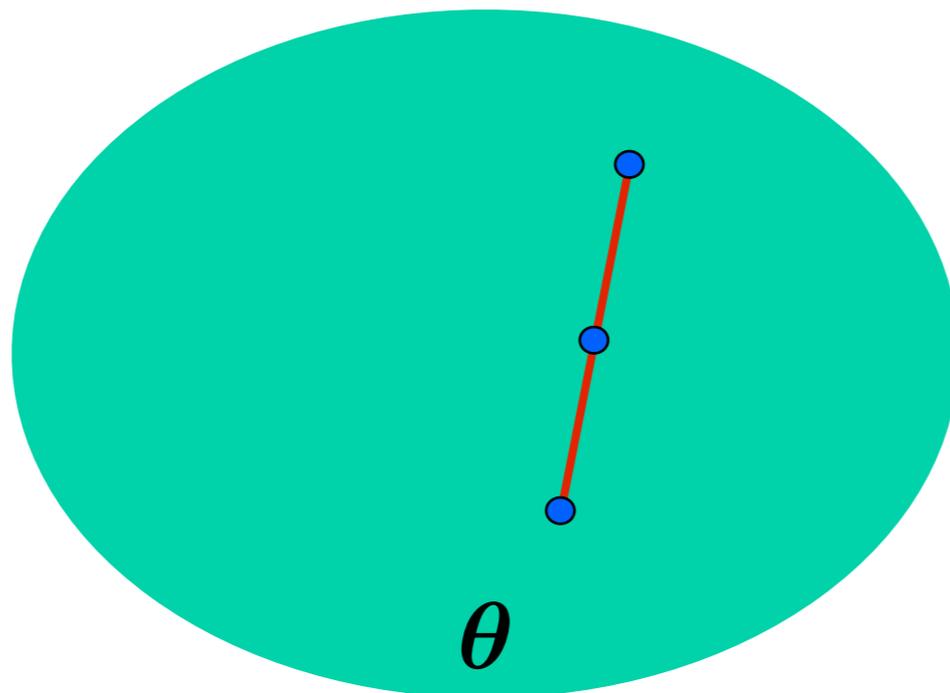


“Sophus Lie”

Reminder: Lie Algebra

- Exponential maps on the $SO(3)$, $SL(3)$ and $SE(3)$ groups are related to the much broader topic of Lie Algebra.
- More details on this topic can be found at in Murray et al. 1994.

$$\mathbf{T}(\boldsymbol{\theta}) = \exp \left(\sum_{i=1}^6 \theta_i \mathbf{A}_i \right) \in SE(3)$$



“Sophus Lie”

SfM - Bundle Adjustment

$$\arg \min_{\mathbf{w}, \boldsymbol{\theta}} \sum_{f=1}^F \sum_{n=1}^N \left\| \mathbf{x}_n^f - \pi(\mathbf{w}_n; \boldsymbol{\theta}^f) \right\|_2^2$$

$\mathbf{x} \leftarrow$ 2D projection $\mathbf{w} \leftarrow$ 3D point

$\boldsymbol{\theta} \leftarrow$ extrinsics $N \leftarrow$ no. of points

$\pi \leftarrow$ projection function

$F \leftarrow$ no. of frames

SfM - Linearization

$$\pi(\mathbf{w}_n + \Delta\mathbf{w}_n; \boldsymbol{\theta}_f \circ \Delta\boldsymbol{\theta}_f) \approx \pi(\mathbf{w}_n; \boldsymbol{\theta}_f) + \mathbf{J}_n^f \begin{bmatrix} \Delta\boldsymbol{\theta}_f \\ \Delta\mathbf{w}_n \end{bmatrix}$$

SfM - Linearization

$$\pi(\mathbf{w}_n + \Delta\mathbf{w}_n; \theta_f \odot \Delta\theta_f) \approx \pi(\mathbf{w}_n; \theta_f) + \mathbf{J}_n^f \begin{bmatrix} \Delta\theta_f \\ \Delta\mathbf{w}_n \end{bmatrix}$$

why not additive??

SfM - Linearization

$$\pi(\mathbf{w}_n + \Delta\mathbf{w}_n; \boldsymbol{\theta}_f \circ \Delta\boldsymbol{\theta}_f) \approx \pi(\mathbf{w}_n; \boldsymbol{\theta}_f) + \mathbf{J}_n^f \begin{bmatrix} \Delta\boldsymbol{\theta}_f \\ \Delta\mathbf{w}_n \end{bmatrix}$$



$$\arg \min_{\Delta\boldsymbol{\theta}, \Delta\mathbf{w}} \sum_{f=1}^F \sum_{n=1}^N \left\| \mathbf{x}_n^f - \pi(\mathbf{w}_n; \boldsymbol{\theta}_f) - \mathbf{J}_n^f \begin{bmatrix} \Delta\boldsymbol{\theta}_f \\ \Delta\mathbf{w}_n \end{bmatrix} \right\|_2^2$$

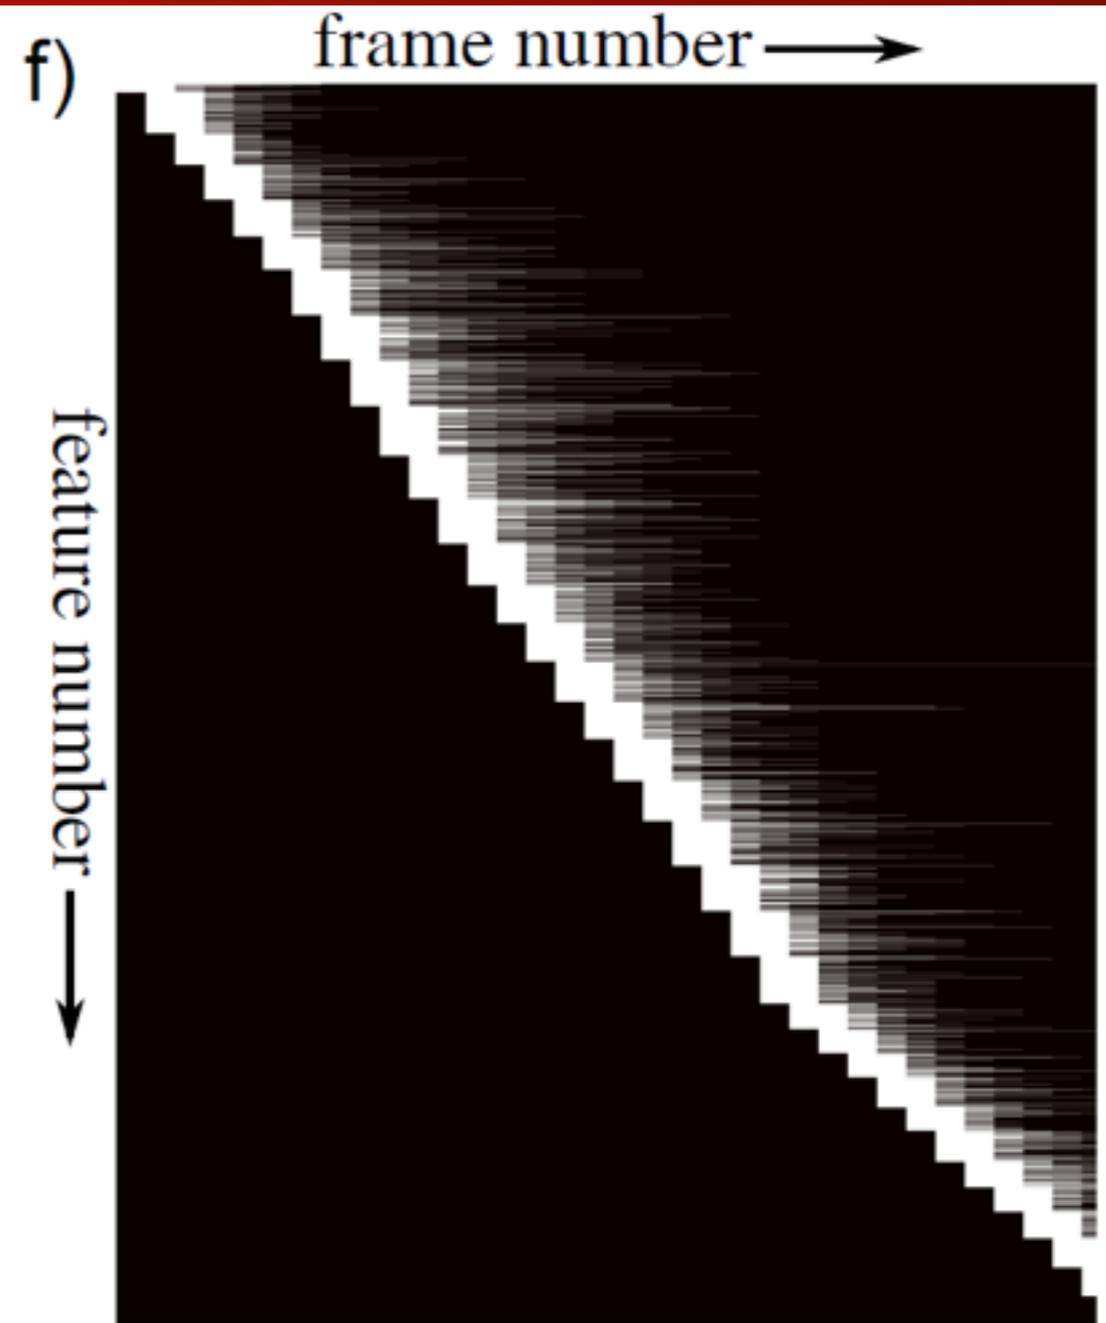
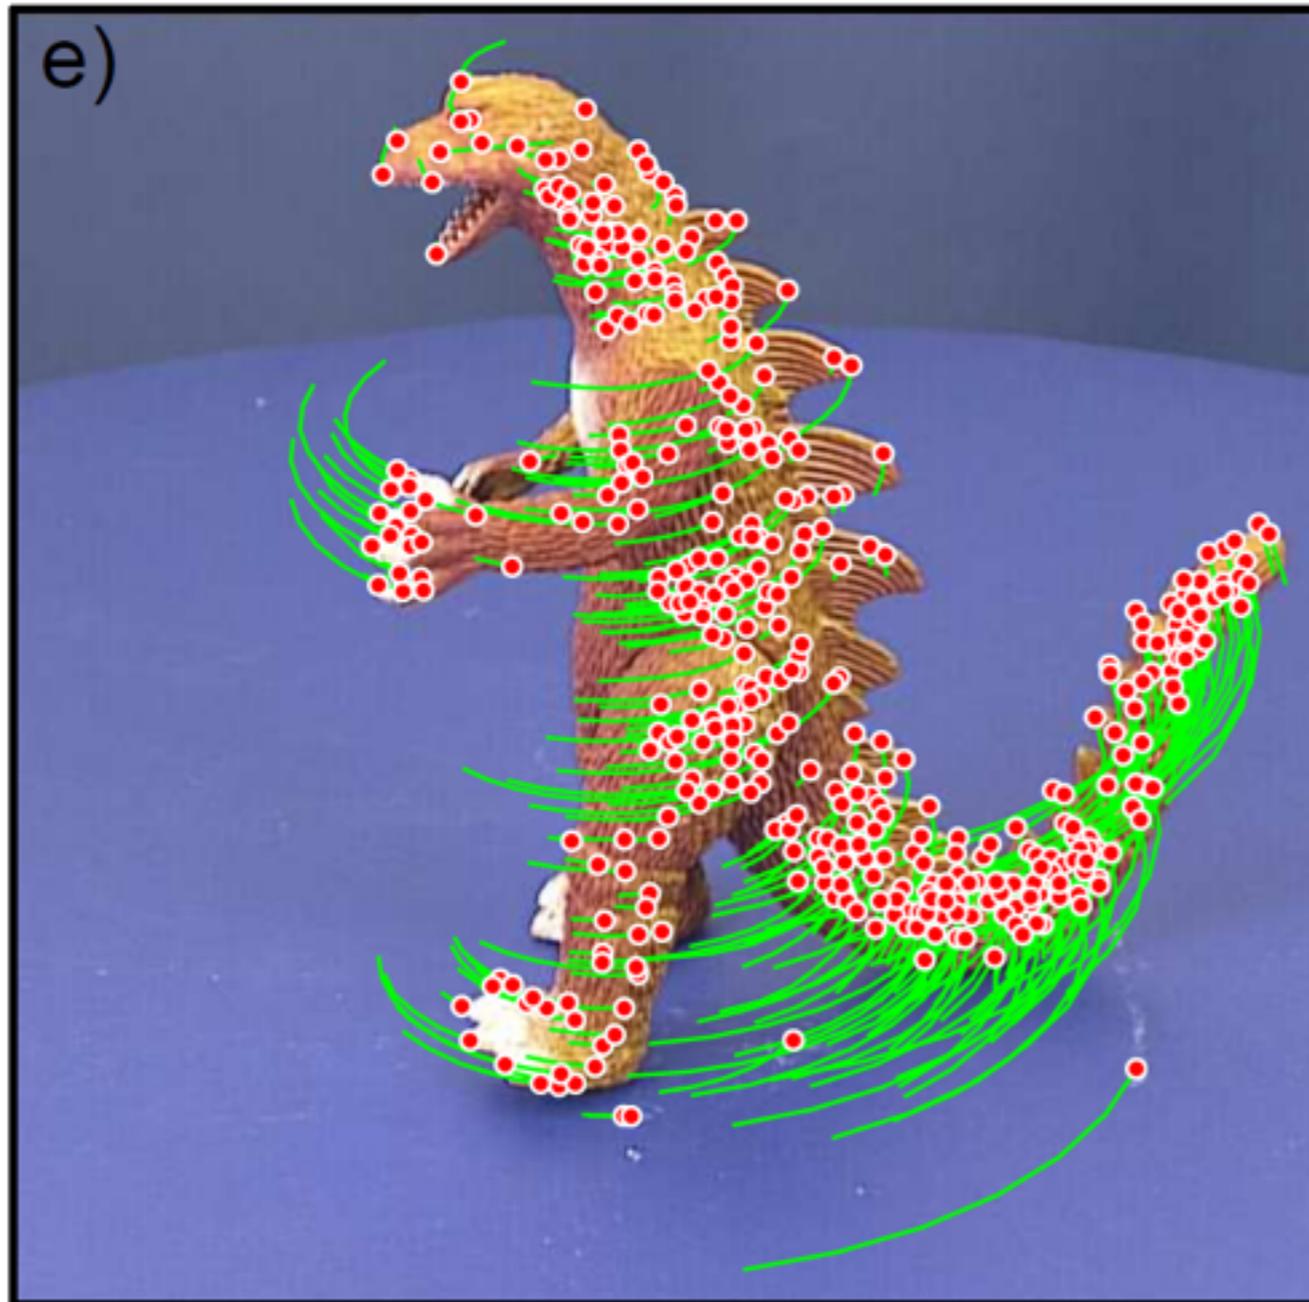
$\mathbf{x} \leftarrow$ 2D projection $\mathbf{w} \leftarrow$ 3D point

$\boldsymbol{\theta} \leftarrow$ extrinsics $N \leftarrow$ no. of points

$\pi \leftarrow$ projection function

$F \leftarrow$ no. of frames

Visibility of Points



$$\Upsilon = \begin{bmatrix} \rho_1^1 & \dots & \rho_1^F \\ \vdots & \ddots & \vdots \\ \rho_N^1 & \dots & \rho_N^F \end{bmatrix}$$

“visibility matrix”

SfM - Bundle Adjustment

$$\pi(\mathbf{w}_n + \Delta \mathbf{w}_n; \boldsymbol{\theta}_f \circ \Delta \boldsymbol{\theta}_f) \approx \pi(\mathbf{w}_n; \boldsymbol{\theta}_f) + \mathbf{J}_n^f \begin{bmatrix} \Delta \boldsymbol{\theta}_f \\ \Delta \mathbf{w}_n \end{bmatrix}$$



$$\arg \min_{\Delta \boldsymbol{\theta}, \Delta \mathbf{w}} \sum_{f=1}^F \sum_{n=1}^N \rho_n^f \|\mathbf{x}_n^f - \pi(\mathbf{w}_n; \boldsymbol{\theta}_f) - \mathbf{J}_n^f \begin{bmatrix} \Delta \boldsymbol{\theta}_f \\ \Delta \mathbf{w}_n \end{bmatrix}\|_2^2$$

$\mathbf{x} \leftarrow$ 2D projection $\mathbf{w} \leftarrow$ 3D point

$\boldsymbol{\theta} \leftarrow$ extrinsics $N \leftarrow$ no. of points

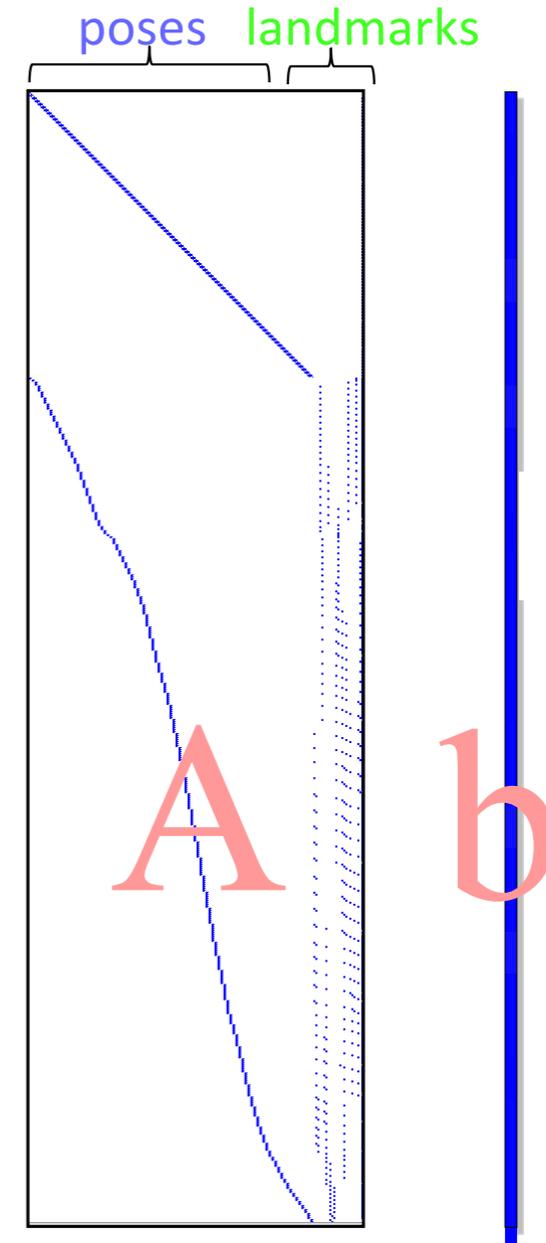
$\pi \leftarrow$ projection function

$F \leftarrow$ no. of frames $\rho \rightarrow$ visibility $\in [0, 1]$

SfM - Bundle Adjustment

$$\arg \min_{\Delta \boldsymbol{\theta}, \Delta \mathbf{w}} \|\mathbf{b} - \mathbf{A} \begin{bmatrix} \Delta \boldsymbol{\theta} \\ \Delta \mathbf{w} \end{bmatrix}\|_2^2$$

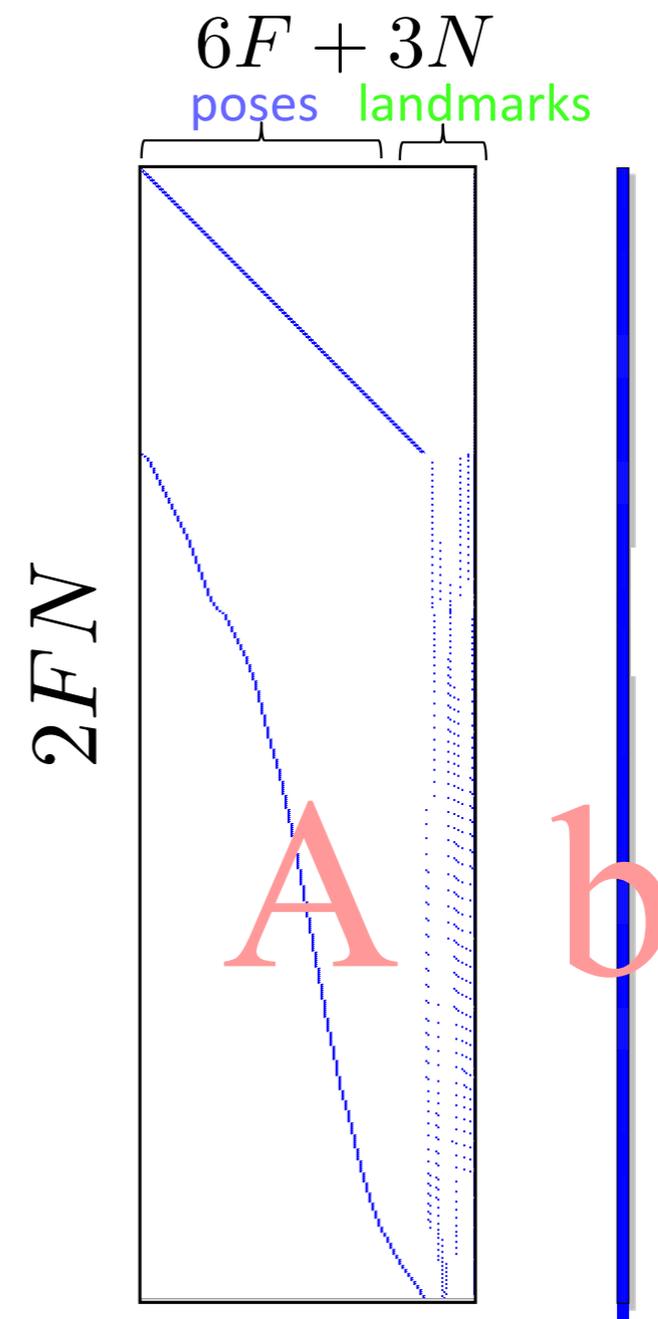
- Can be solved efficiently using sparse linear solvers such as,
 - Google Ceres Solver - <http://ceres-solver.org>
 - G2o - <https://openslam.org/g2o.html> .
- Then iteratively apply GN or LM algorithm.



SfM - Bundle Adjustment

$$\arg \min_{\Delta \boldsymbol{\theta}, \Delta \mathbf{w}} \left\| \mathbf{b} - \mathbf{A} \begin{bmatrix} \Delta \boldsymbol{\theta} \\ \Delta \mathbf{w} \end{bmatrix} \right\|_2^2$$

- Can be solved efficiently using sparse linear solvers such as,
 - Google Ceres Solver - <http://ceres-solver.org>
 - G2o - <https://openslam.org/g2o.html> .
- Then iteratively apply GN or LM algorithm.



Reminder: Gauss-Newton Algorithm

- Gauss-Newton (GN) algorithm common strategy for optimizing non-linear least-squares problems.

$$\arg \min_{\mathbf{y}} \|\mathbf{x} - \mathcal{F}(\mathbf{y})\|_2^2$$
$$\text{s.t. } \mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}^M$$

Step 1:

$$\arg \min_{\Delta \mathbf{y}} \left\| \mathbf{x} - \mathcal{F}(\mathbf{y}) - \frac{\partial \mathcal{F}(\mathbf{y})}{\partial \mathbf{y}^T} \Delta \mathbf{y} \right\|_2^2$$

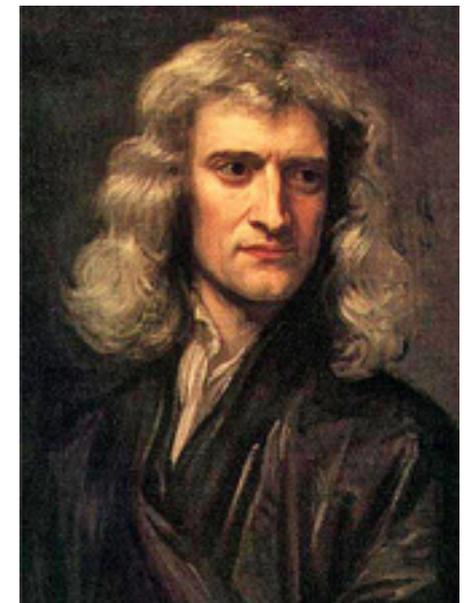
Step 2:

$$\mathbf{y} \rightarrow \mathbf{y} + \Delta \mathbf{y}$$

keep applying steps until $\Delta \mathbf{y}$ converges.



“Carl Friedrich Gauss”



“Isaac Newton”

Reminder: Gauss-Newton Algorithm

- Gauss-Newton (GN) algorithm common strategy for optimizing non-linear least-squares problems.

$$\arg \min_{\mathbf{y}} \|\mathbf{x} - \mathcal{F}(\mathbf{y})\|_2^2$$
$$\text{s.t. } \mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}^M$$



“Carl Friedrich Gauss”

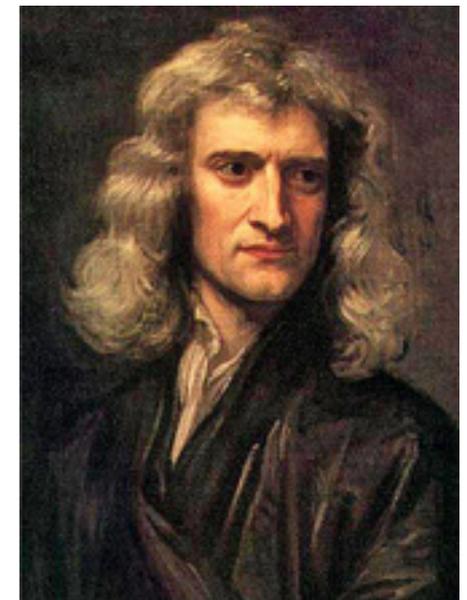
Step 1:

$$\arg \min_{\Delta \mathbf{y}} \left\| \mathbf{x} - \mathcal{F}(\mathbf{y}) - \frac{\partial \mathcal{F}(\mathbf{y})}{\partial \mathbf{y}^T} \Delta \mathbf{y} \right\|_2^2$$

Step 2:

$\mathbf{y} \rightarrow \mathbf{y} + \Delta \mathbf{y}$ “Is the update additive?”

keep applying steps until $\Delta \mathbf{y}$ converges.



“Isaac Newton”

Today

- SfM - Bundle Adjustment
- **VSLAM - Keyframe vs. Filtering**
- Visual Odometry
- Loop Closure

Mono SLAM = Online SFM

- Monocular SLAM is just another name for “online” SFM.
- If computation was not an issue, one would just apply Bundle Adjustment after every new frame

$$\arg \min_{\mathbf{w}, \boldsymbol{\theta}} \sum_{f=1}^F \sum_{n=1}^N \left\| \mathbf{x}_n^f - \pi(\mathbf{w}_n; \boldsymbol{\theta}^f) \right\|_2^2$$

$\mathbf{x} \leftarrow$ 2D projection $\mathbf{w} \leftarrow$ 3D point

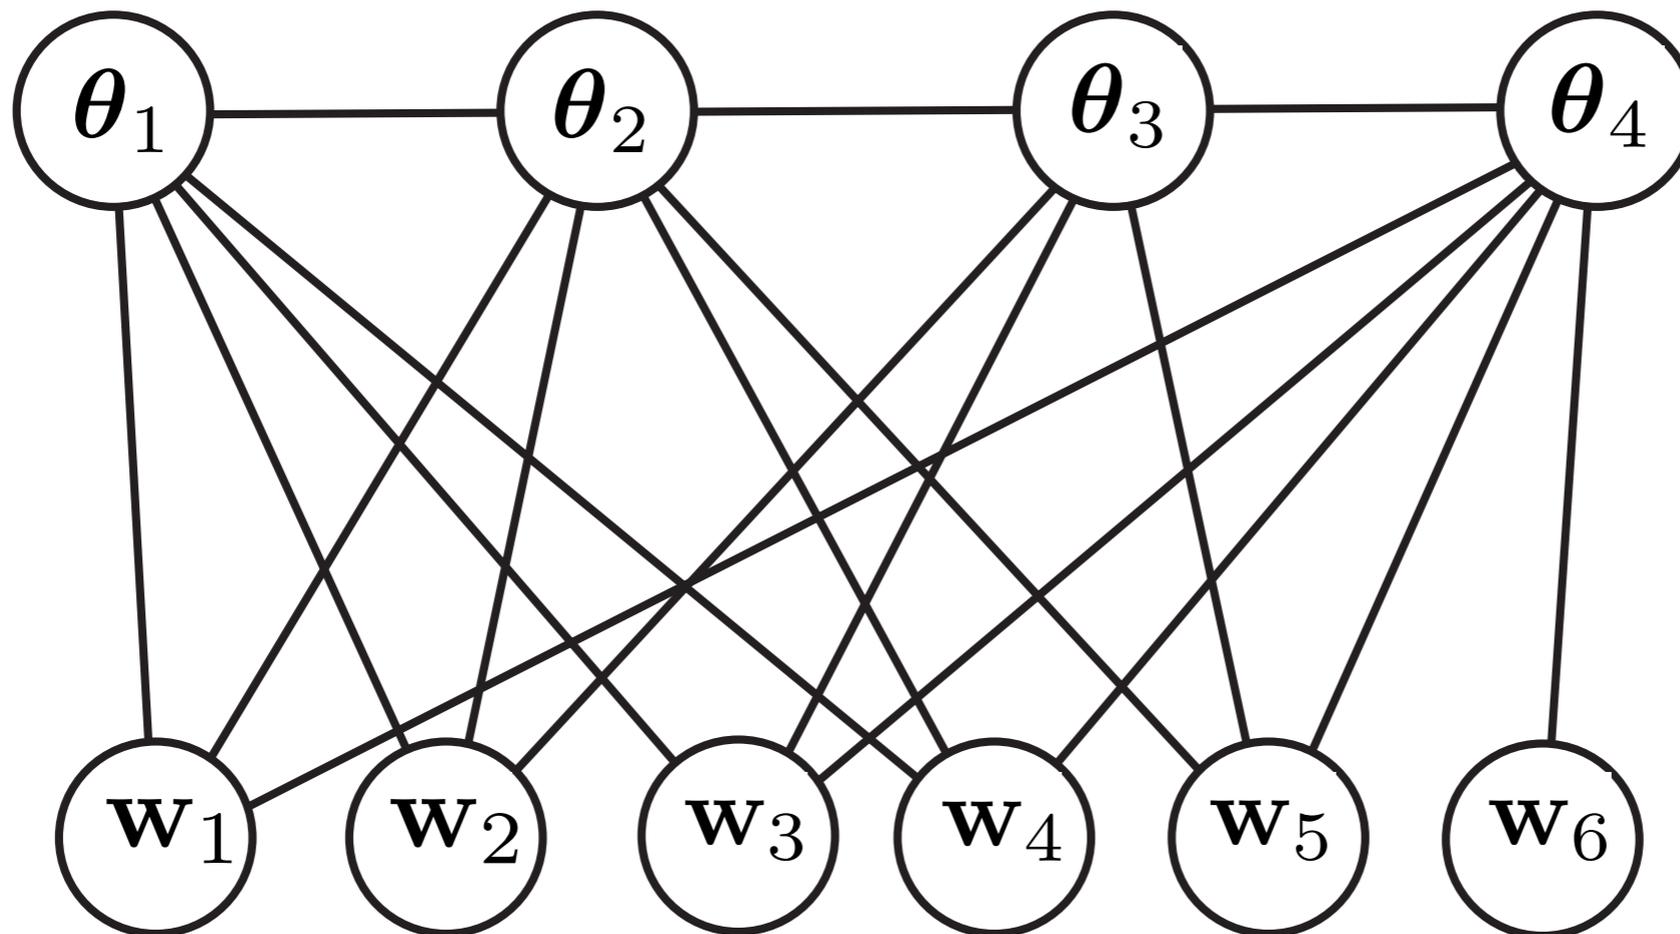
$\boldsymbol{\theta} \leftarrow$ extrinsics $N \leftarrow$ no. of points

$\pi \leftarrow$ projection function

$F \leftarrow$ no. of frames

Mono SLAM - MRF

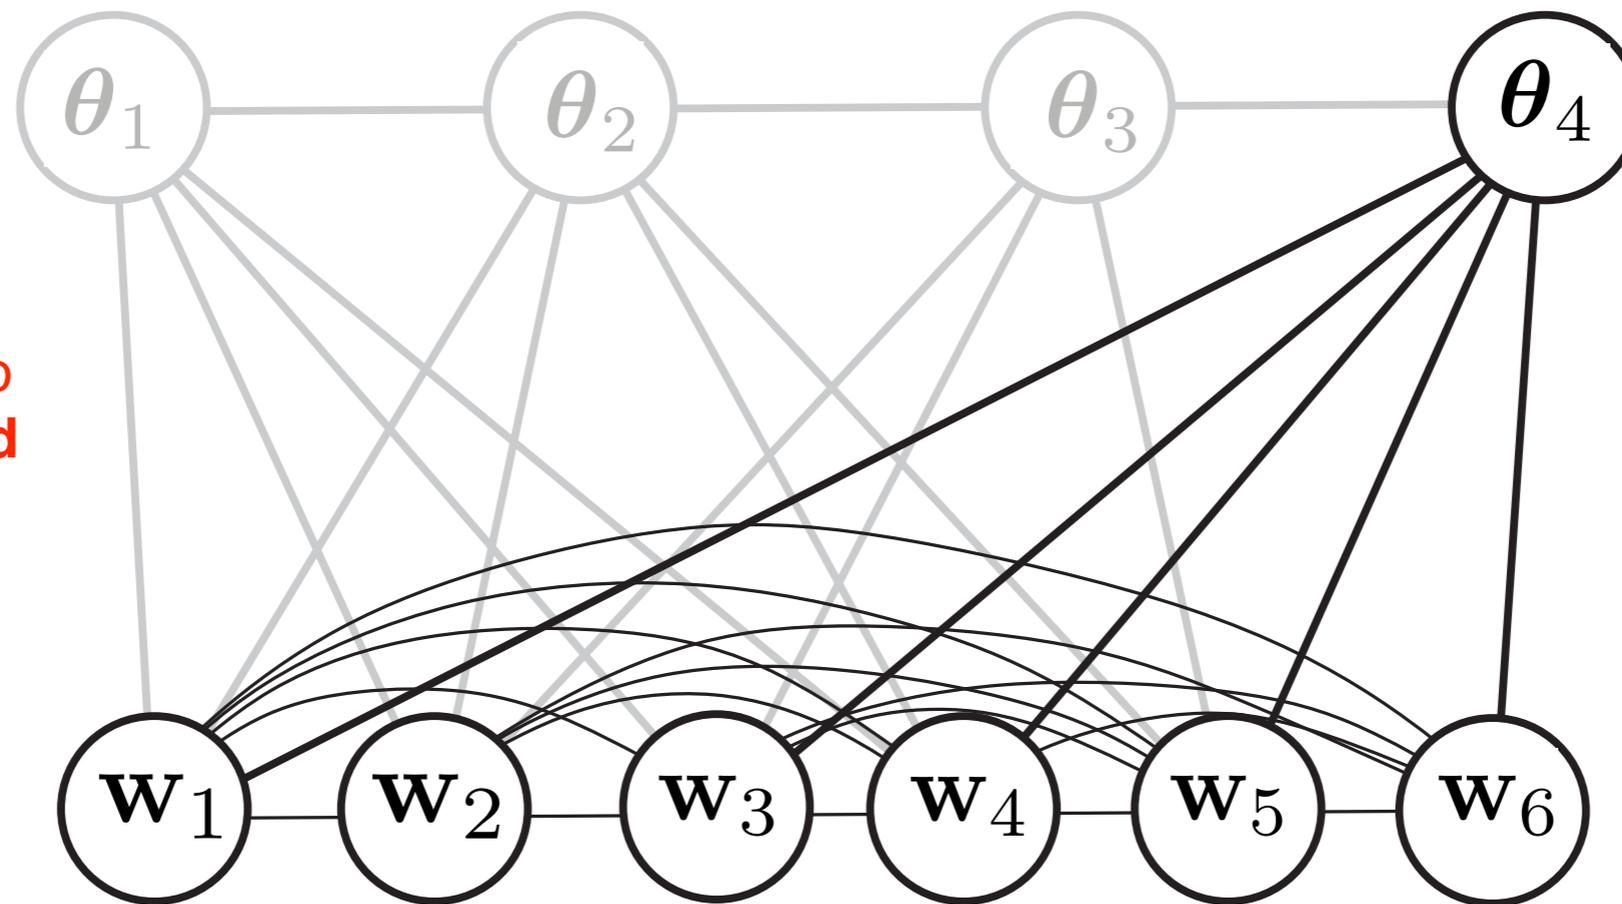
- One can view the problem of SfM - Bundle Adjustment as doing inference on a Markov Random Field (MRF).
- **Problem** - becomes exponentially harder as times goes on.



“edges based on visibility” ρ

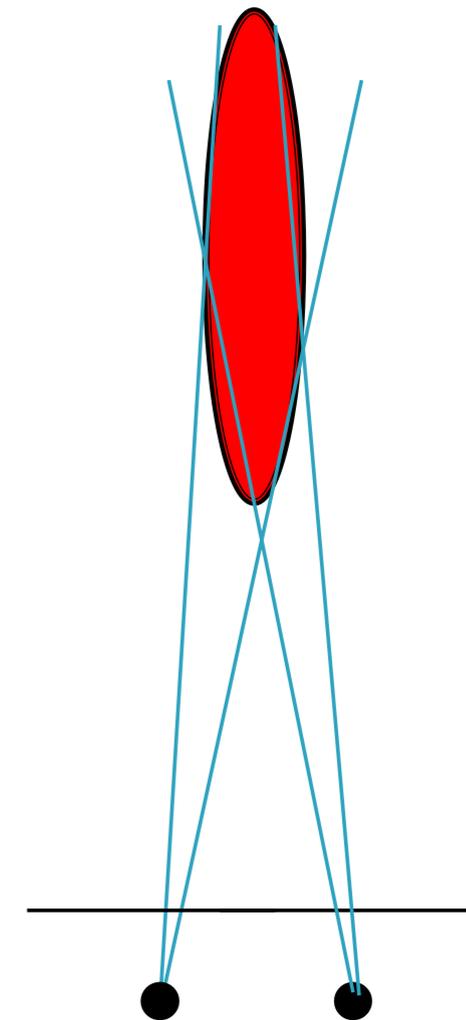
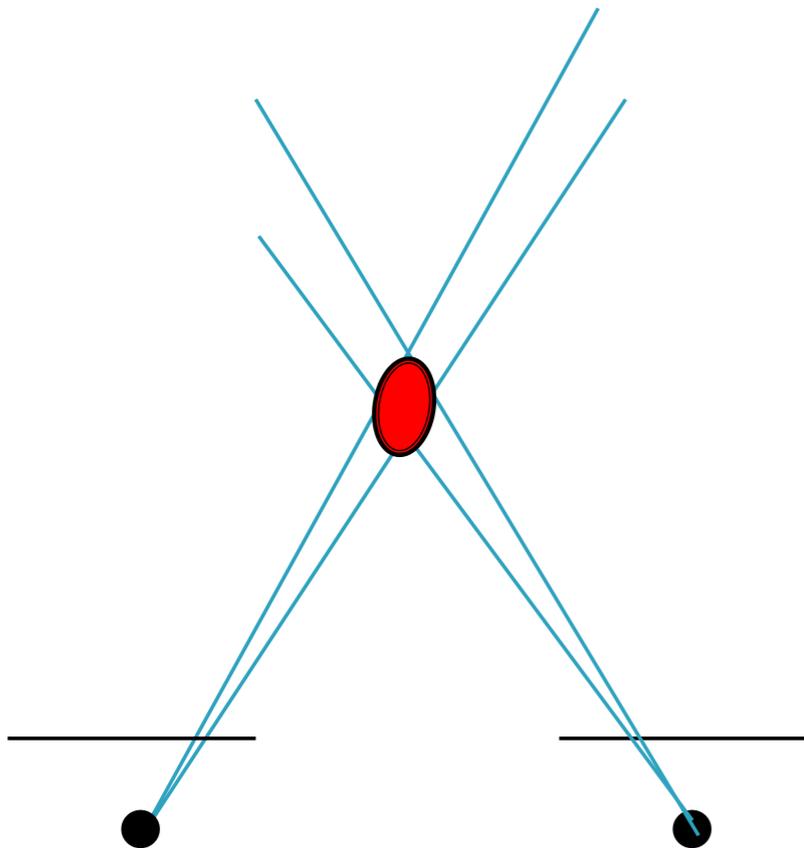
Mono SLAM - Filtering

- Classic way of resolving this was to pose BA problem as a filter - such as an Extended Kalman Filter (EKF).
- **Problem** - Wastes processing time on frames that added very little information.



Mono SLAM - Filtering

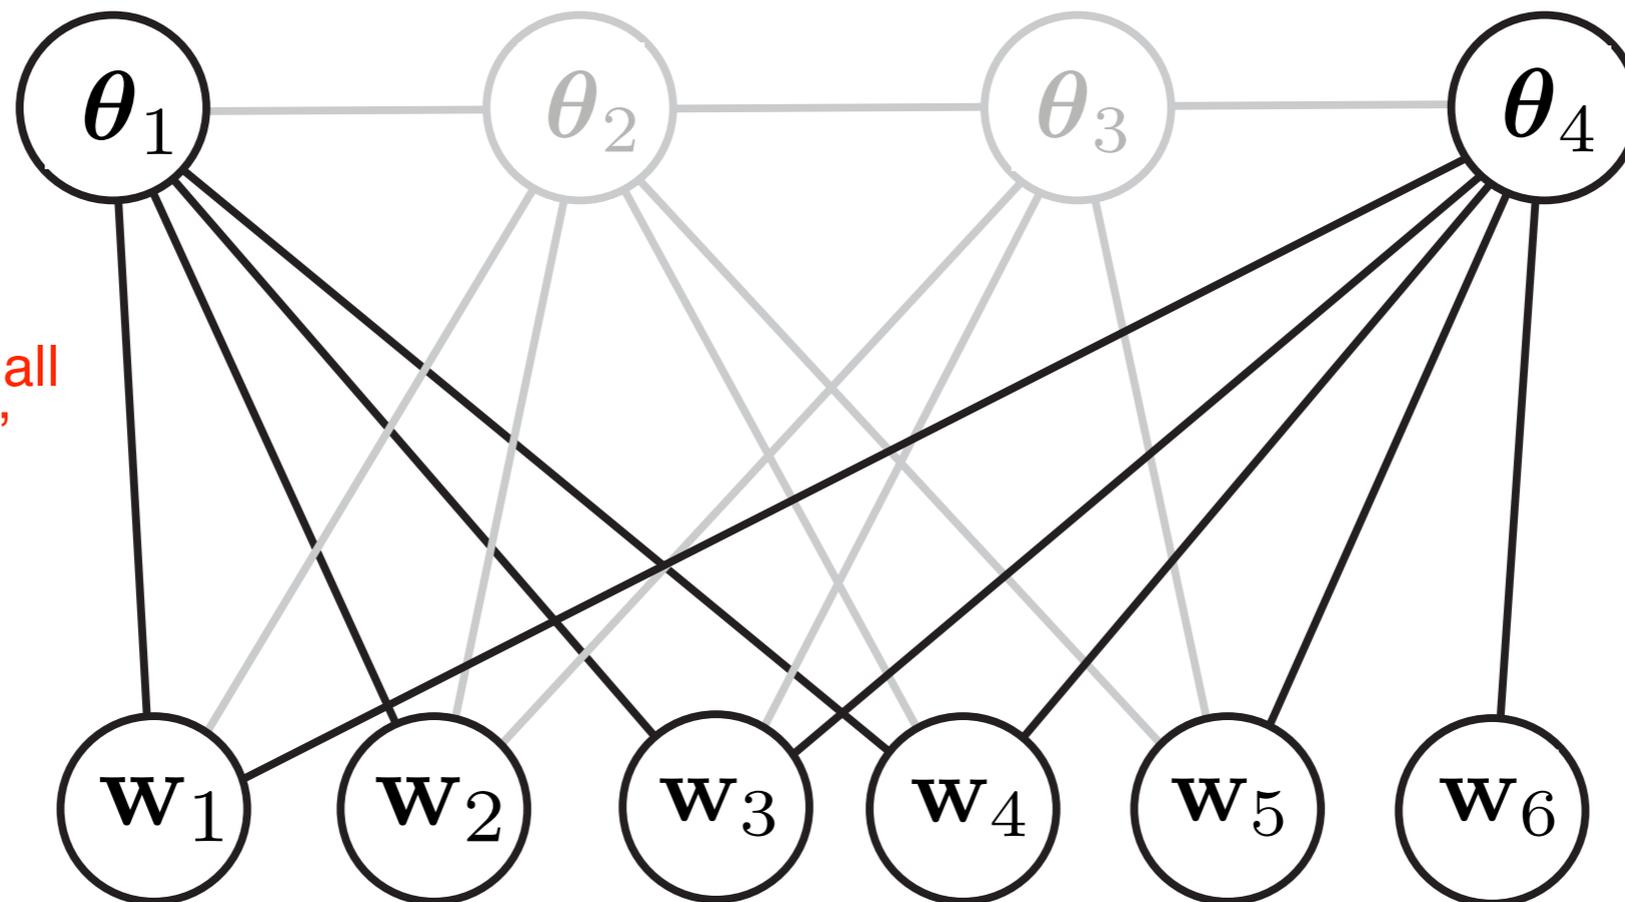
- Filtering approaches are often times problematic (e.g. think when the device stops moving).
- When frames are taken at nearby positions compared to the scene distance, 3D points will exhibit large uncertainty.



Mono SLAM - Keyframe

- A better strategy is to employ **keyframe BA**.
- Made popular by Klein & Murray's - Parallel Tracking and Mapping (PTAM) algorithm.

“remove **all** but a small subset of keyframes”



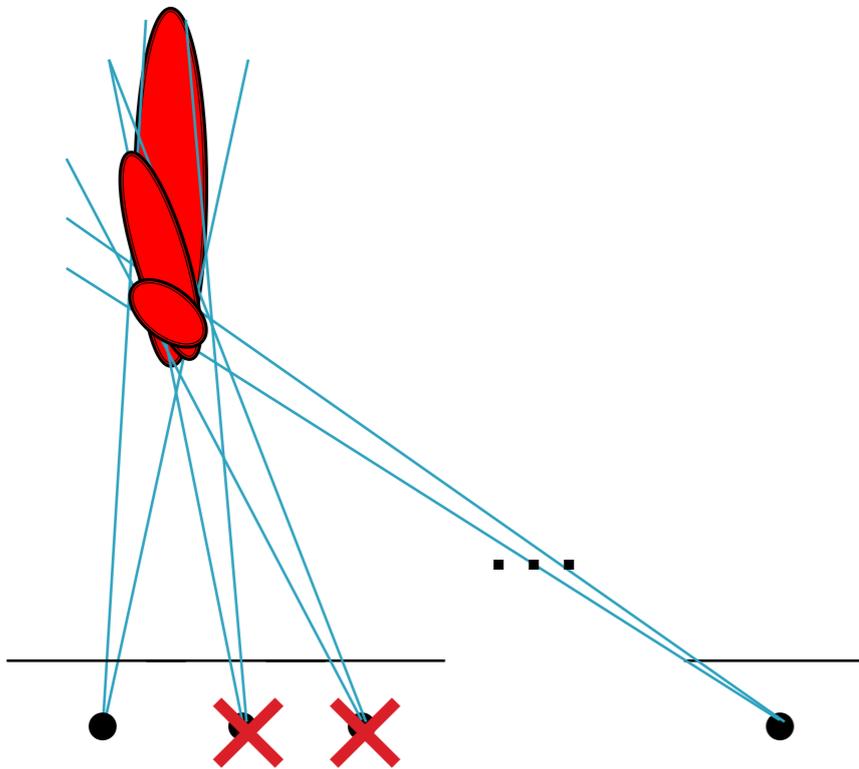
G. Klein and D. Murray, “Parallel tracking and mapping for small AR workspaces”, ISMAR 2007.

H. Strasdat, J. M. M. Montiel, and A. J. Davison, “Visual SLAM: Why filter?” Image and Vision Computing, vol. 30, no. 2, pp. 65–77, 2012.

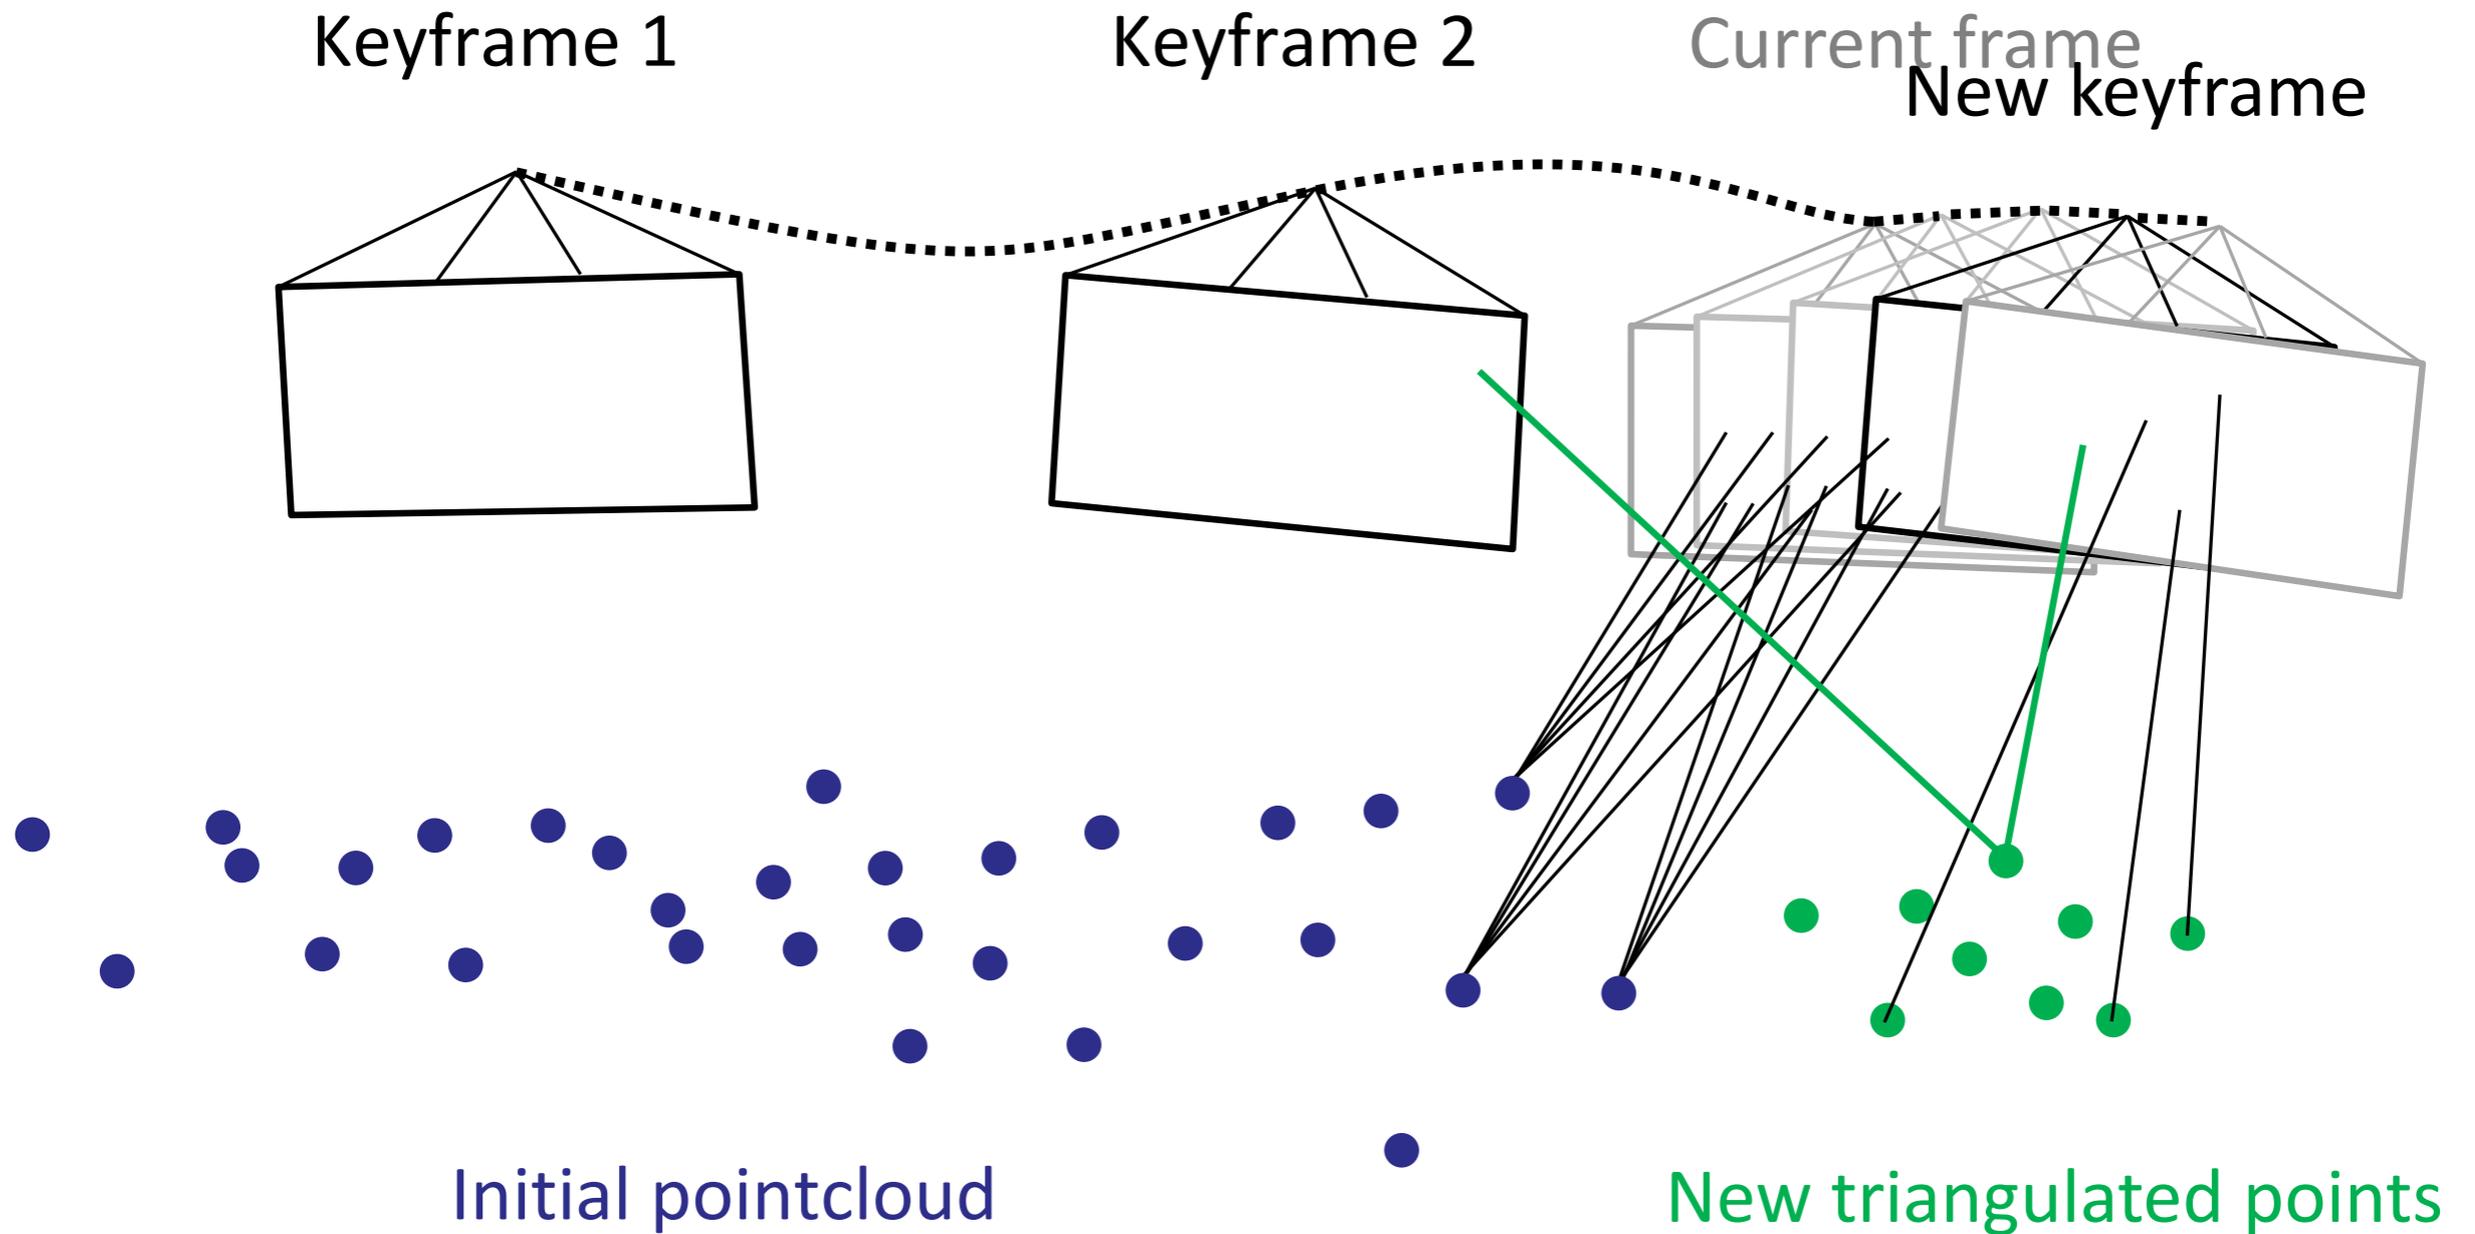
Keyframe Selection

- One way to avoid this consists of skipping frames until the average uncertainty of the 3D points decreases below a certain threshold. The selected frames are called **keyframes**.
- Rule of thumb: add a keyframe when,

$$\frac{\textit{keyframe distance}}{\textit{average-depth}} > \textit{threshold} (\sim 10-20 \%)$$



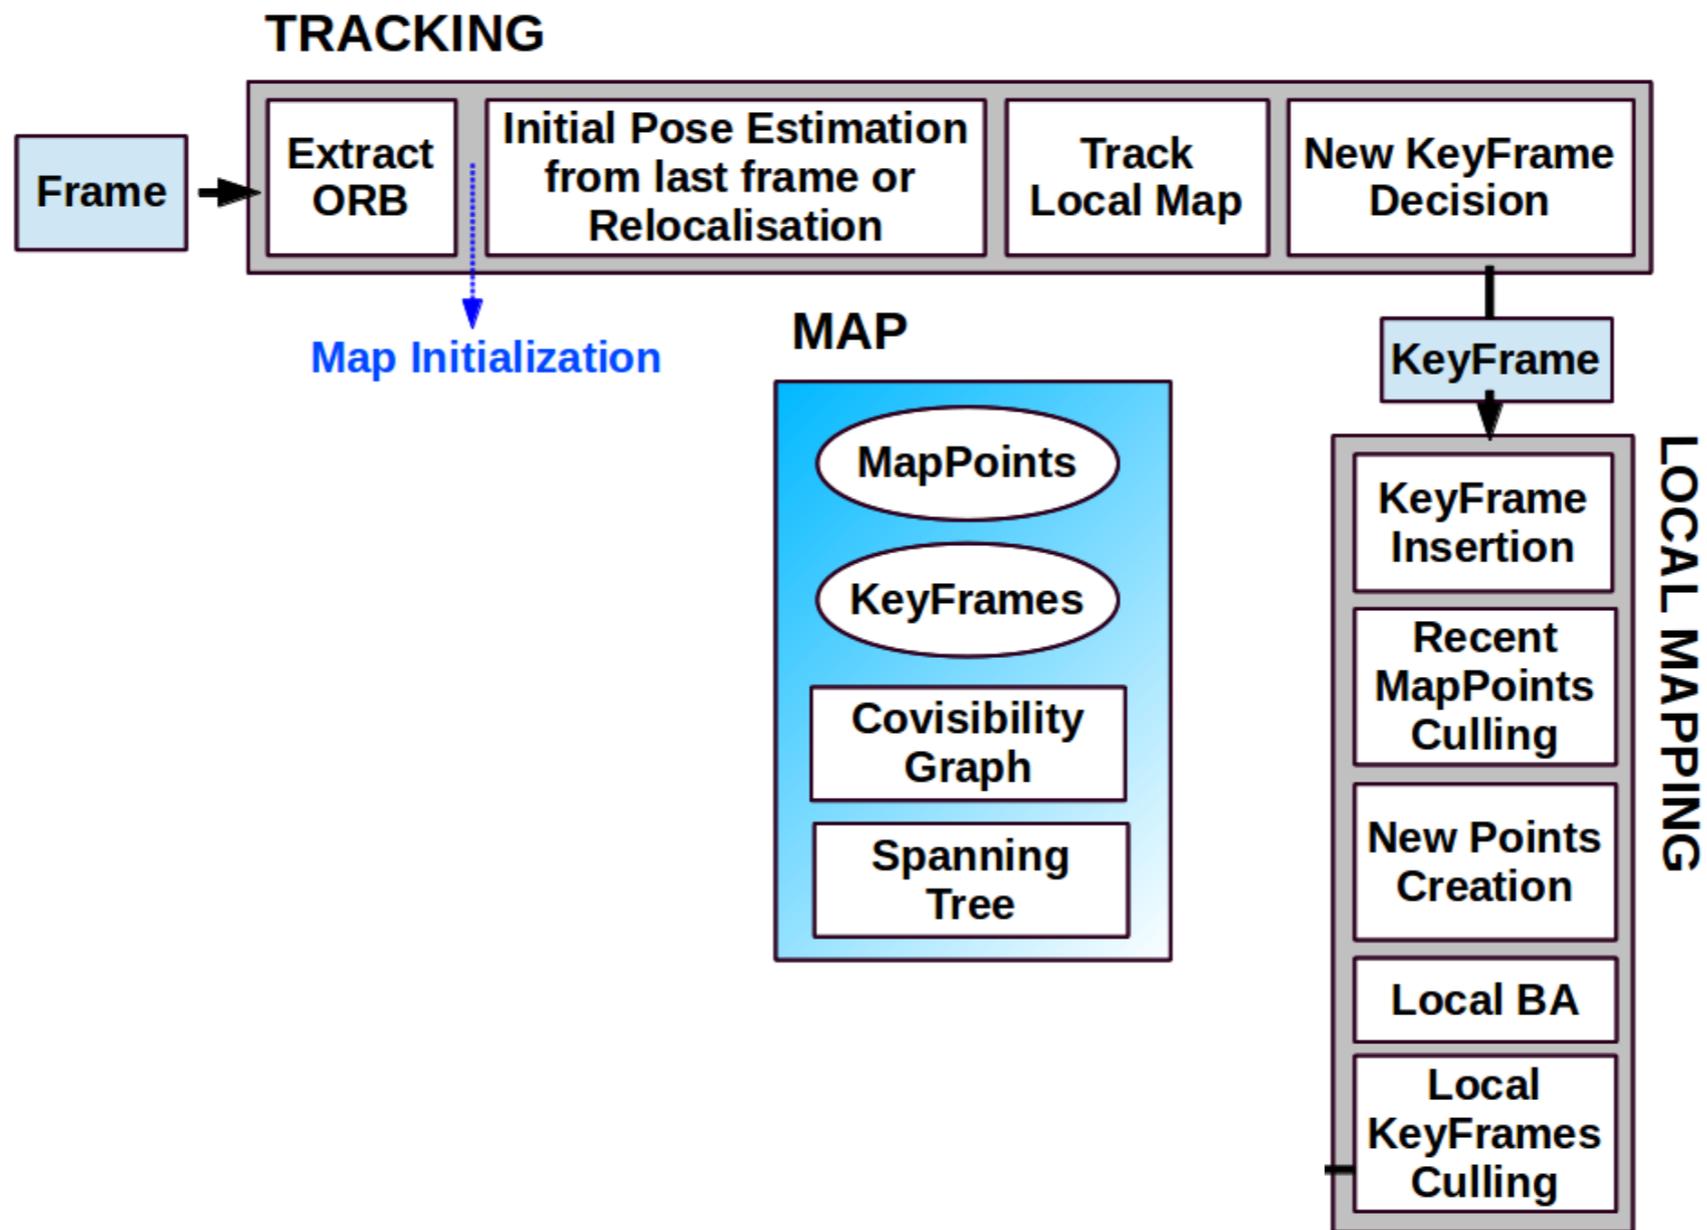
Keyframe-based SLAM



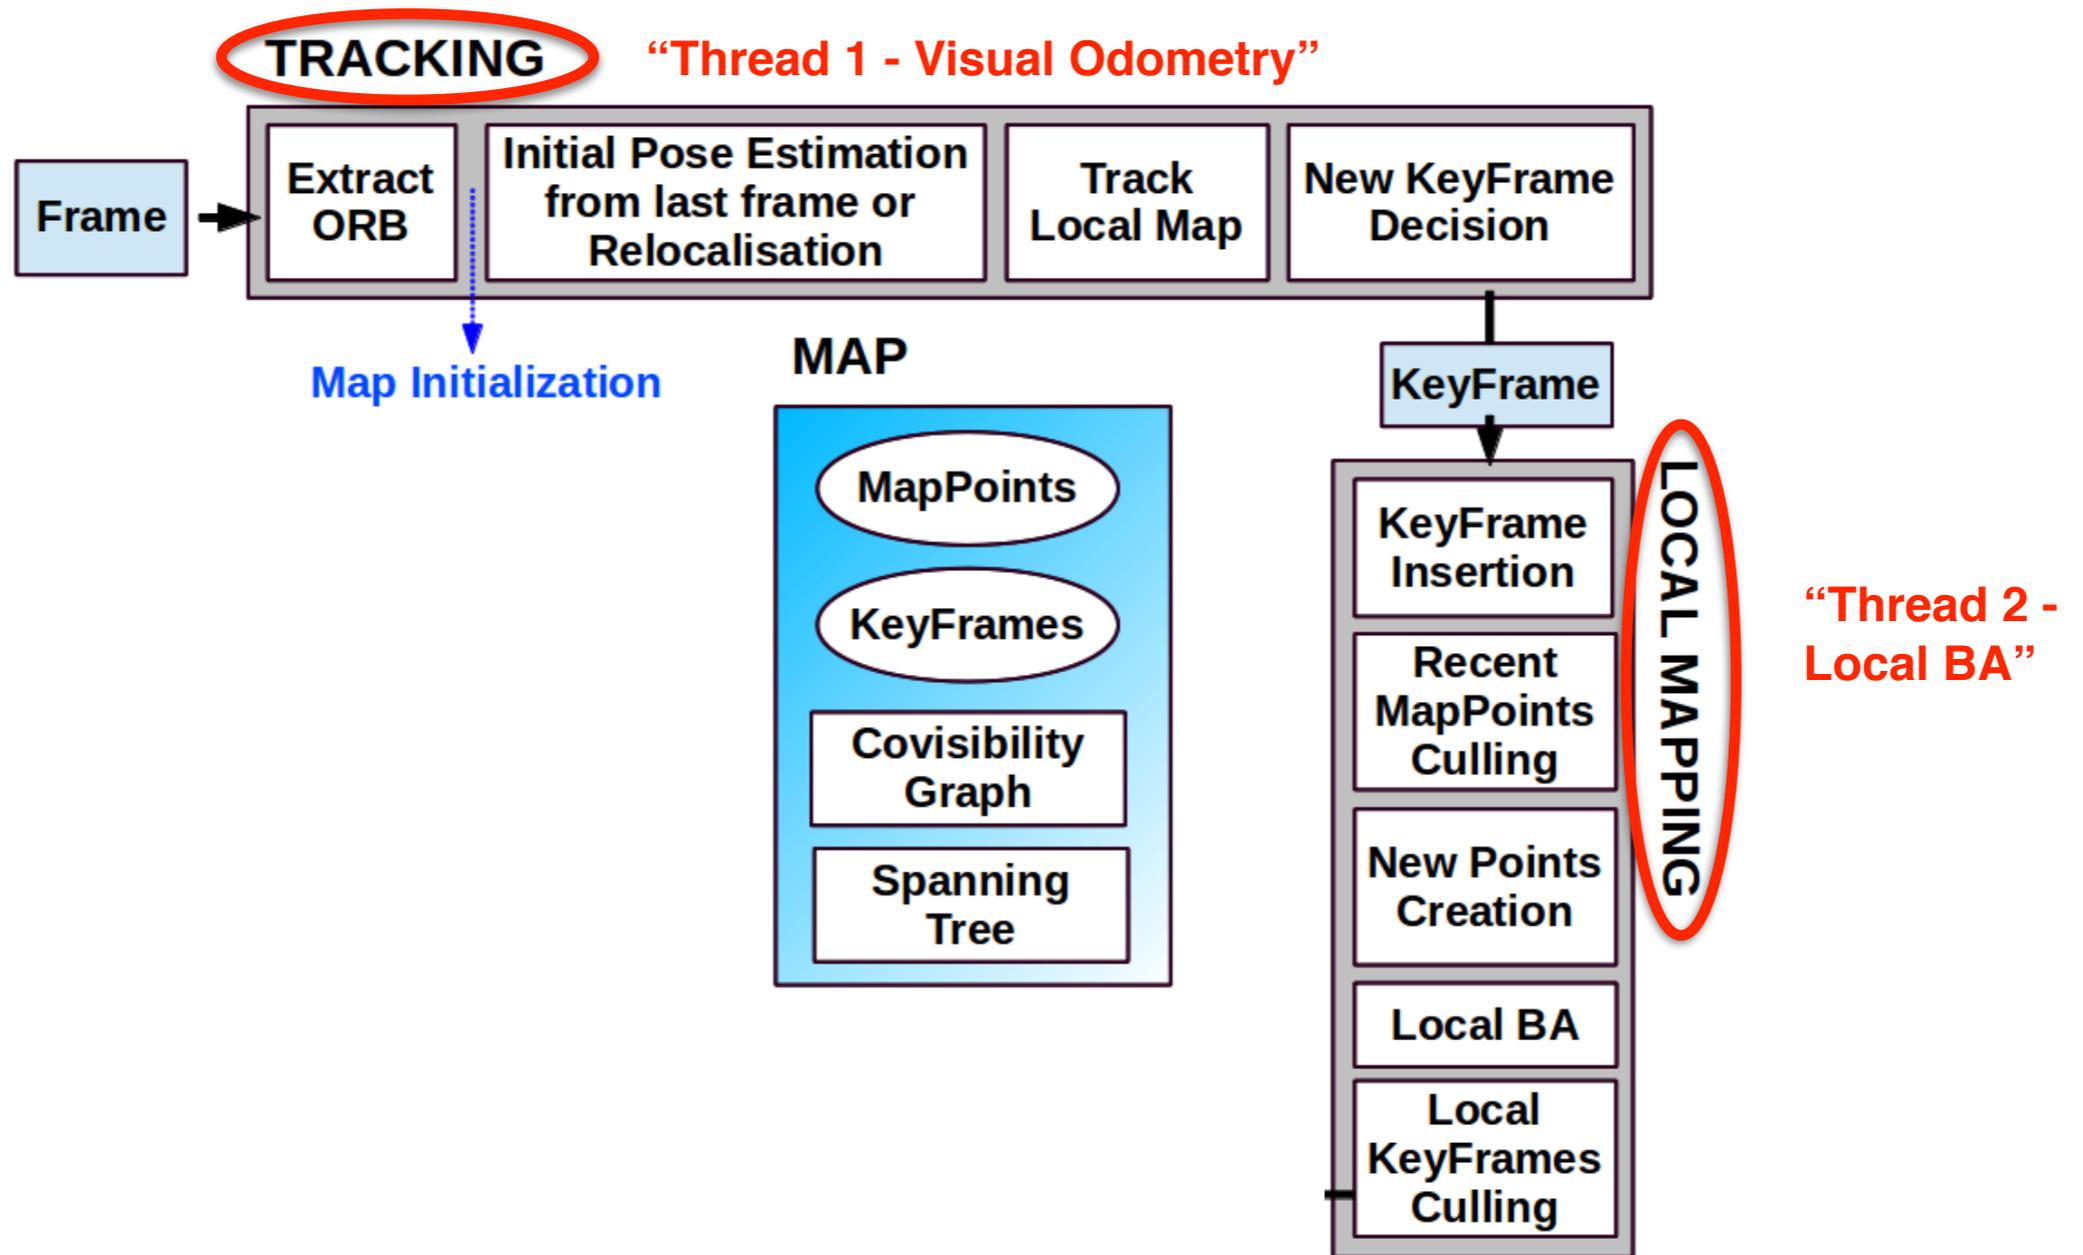
PTAM - Separate Threads

- An innovation of Klein & Murray's PTAM method was the separation of the camera tracking (θ) and map estimation (w) tasks.
- Camera Tracking or **visual odometry (VO)** runs on one thread in real-time.
- Map estimation runs on a separate thread (not having to run in real-time) allowing for bundle adjustment.
- Same idea is still being utilized in current state of the art visual SLAM algorithms (e.g. ORB SLAM).

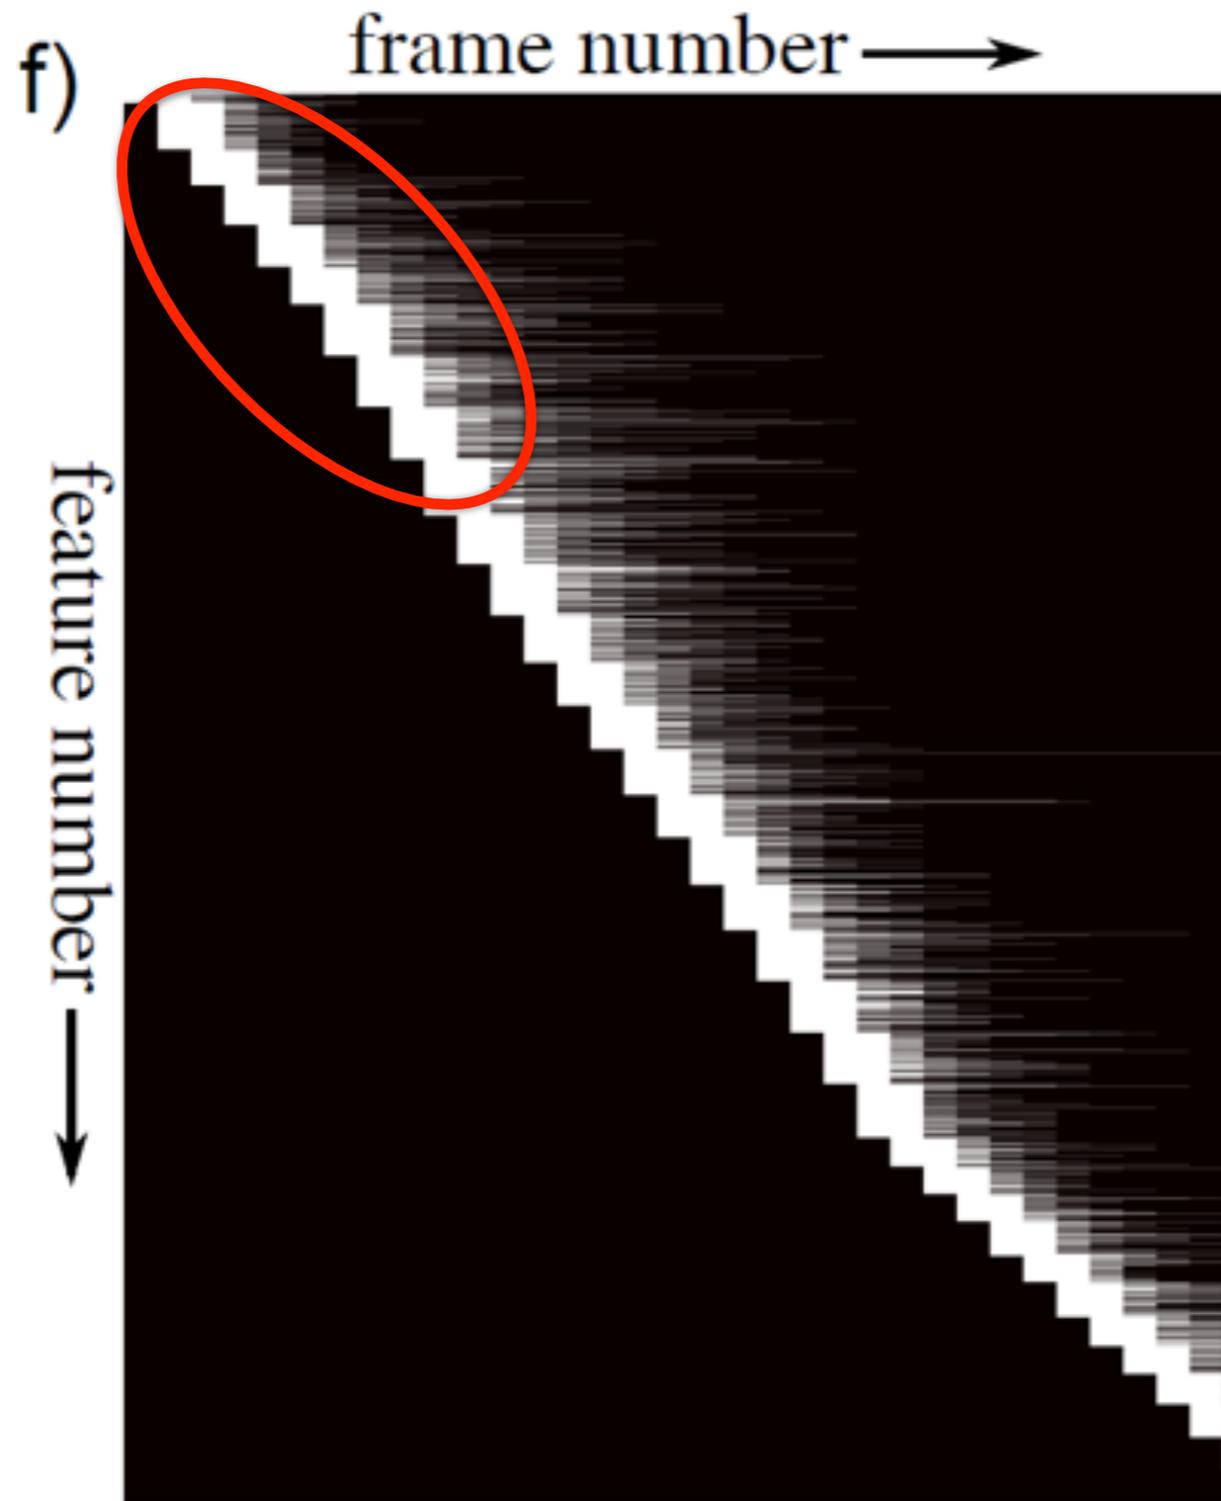
Example - ORB SLAM



Example - ORB SLAM



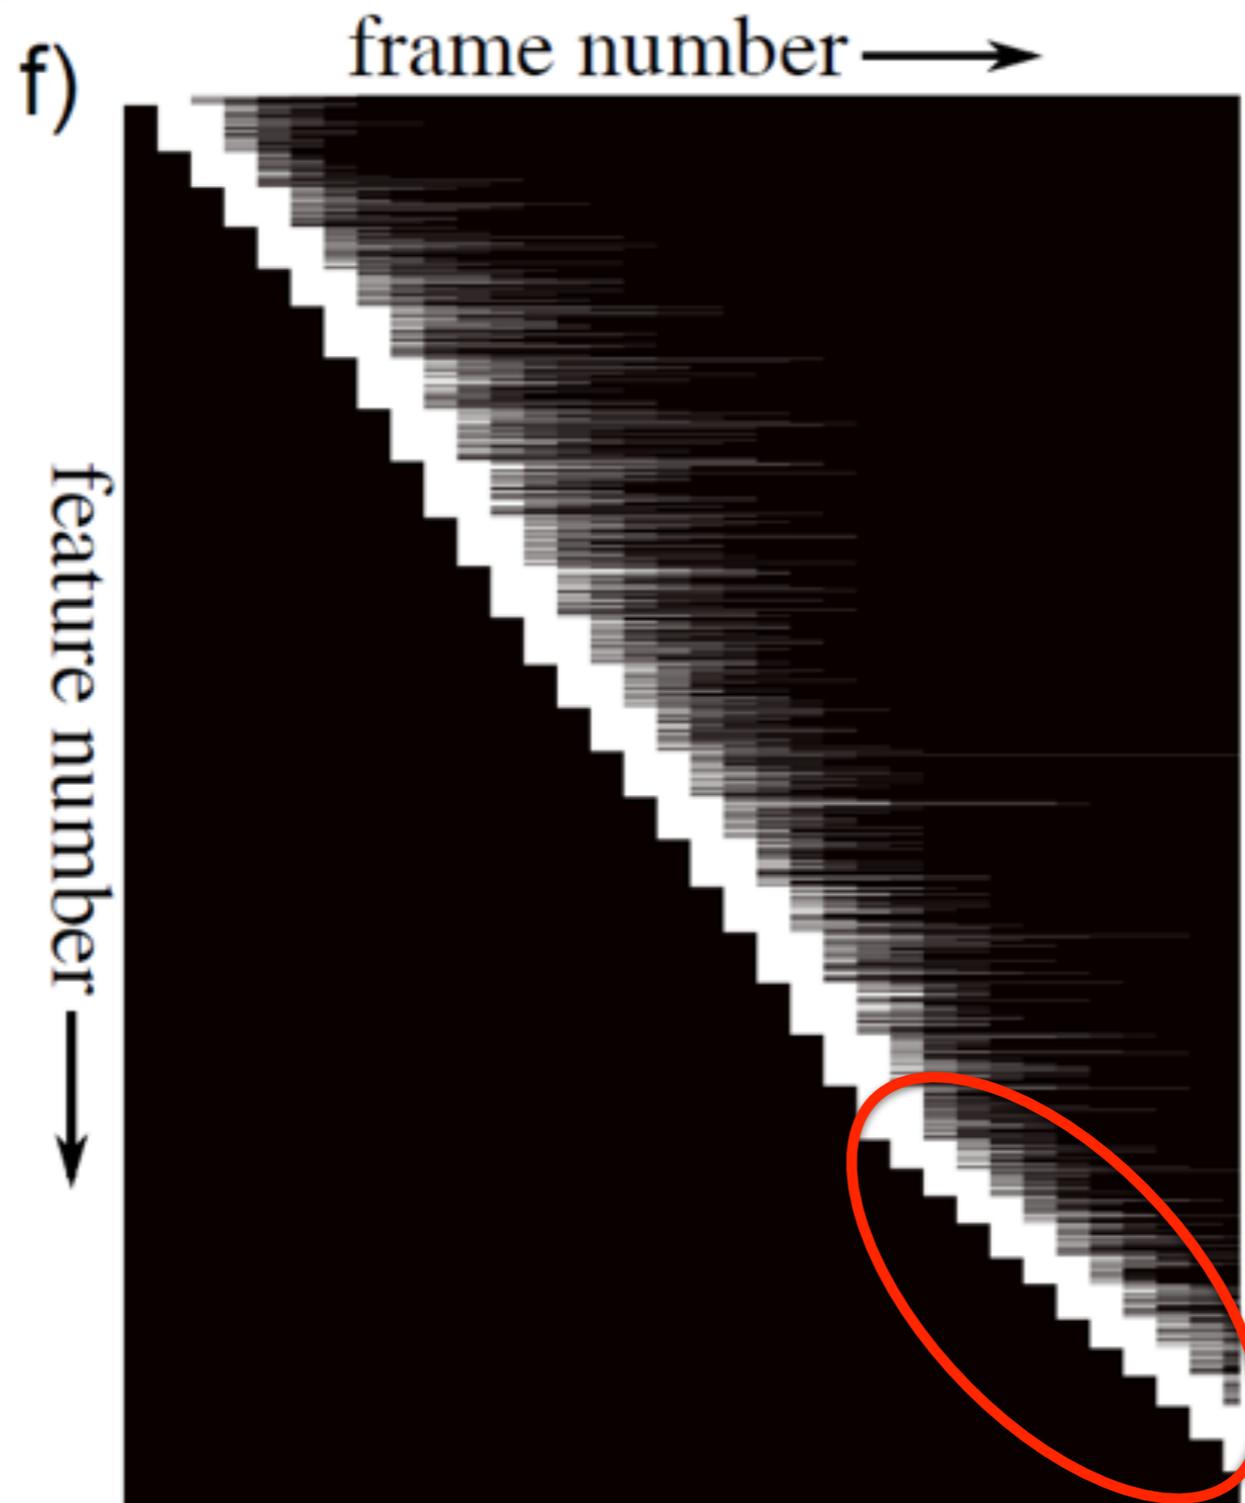
Local Bundle Adjustment



$$\Upsilon = \begin{bmatrix} \rho_1^1 & \dots & \rho_1^F \\ \vdots & \ddots & \vdots \\ \rho_N^1 & \dots & \rho_N^F \end{bmatrix}$$

“visibility matrix”

Local Bundle Adjustment



$$\Upsilon = \begin{bmatrix} \rho_1^1 & \dots & \rho_1^F \\ \vdots & \ddots & \vdots \\ \rho_N^1 & \dots & \rho_N^F \end{bmatrix}$$

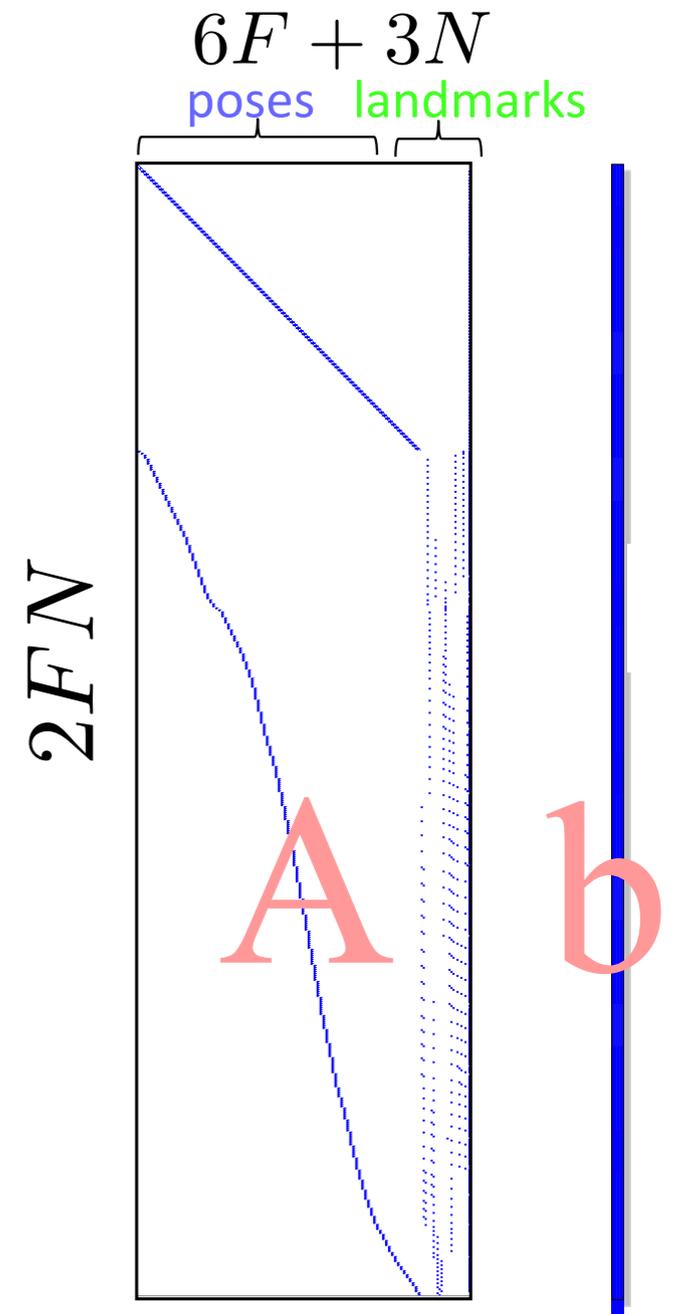
“visibility matrix”

iSAM: Incremental Smoothing & Mapping

$$\arg \min_{\Delta \theta, \Delta \mathbf{w}} \|\mathbf{b} - \mathbf{A} \begin{bmatrix} \Delta \theta \\ \Delta \mathbf{w} \end{bmatrix}\|_2^2$$

- iSAM can be entertained using QR Factorization.

$$\mathbf{A} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix}$$



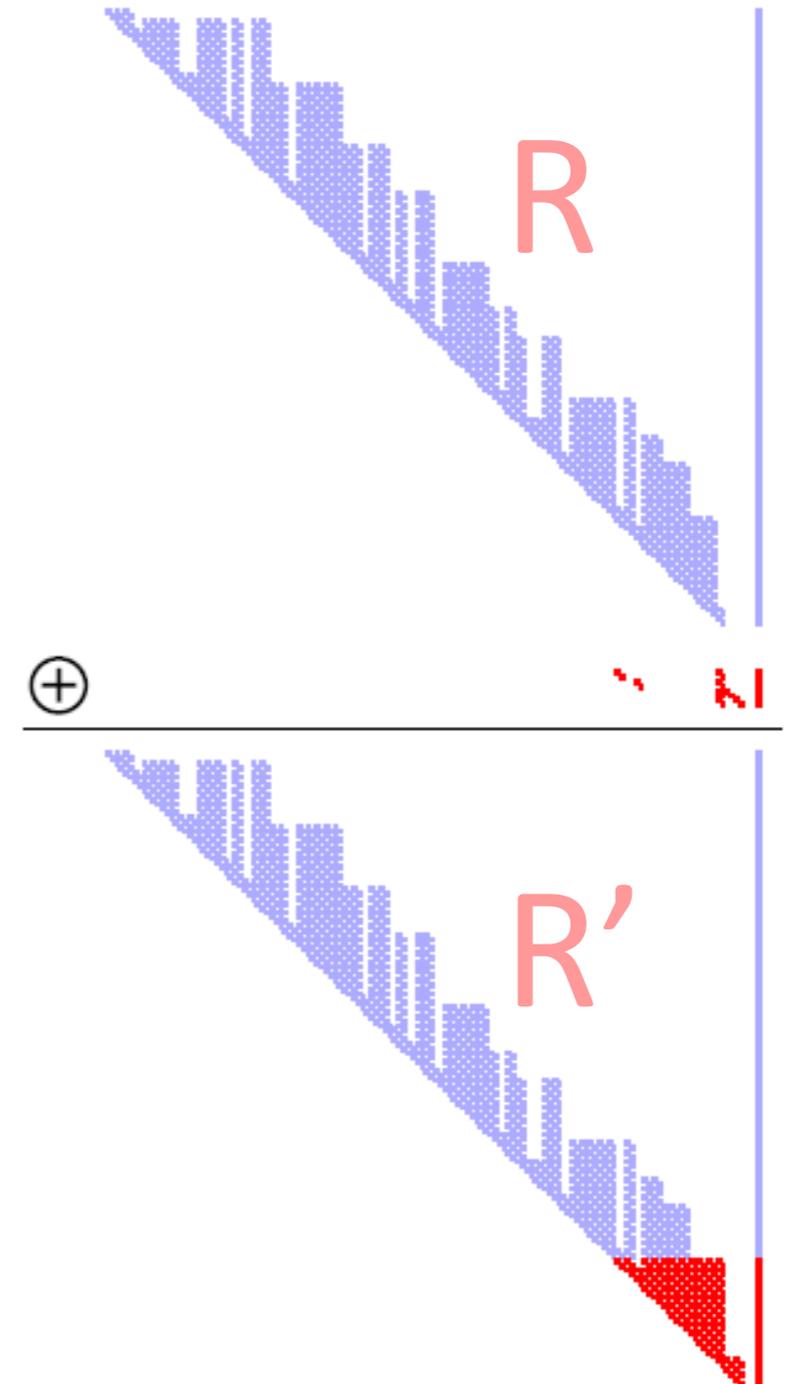
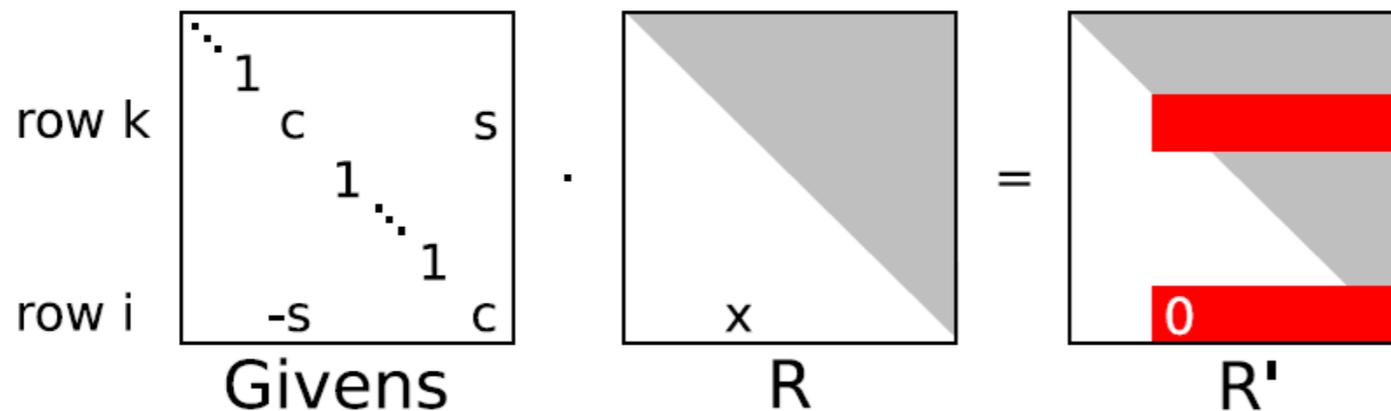
iSAM: Incremental Smoothing & Mapping

Solving a growing system:

- R factor from previous step
- How do we add new measurements?

Key idea:

- Append to existing matrix factorization
- “Repair” using Givens rotations

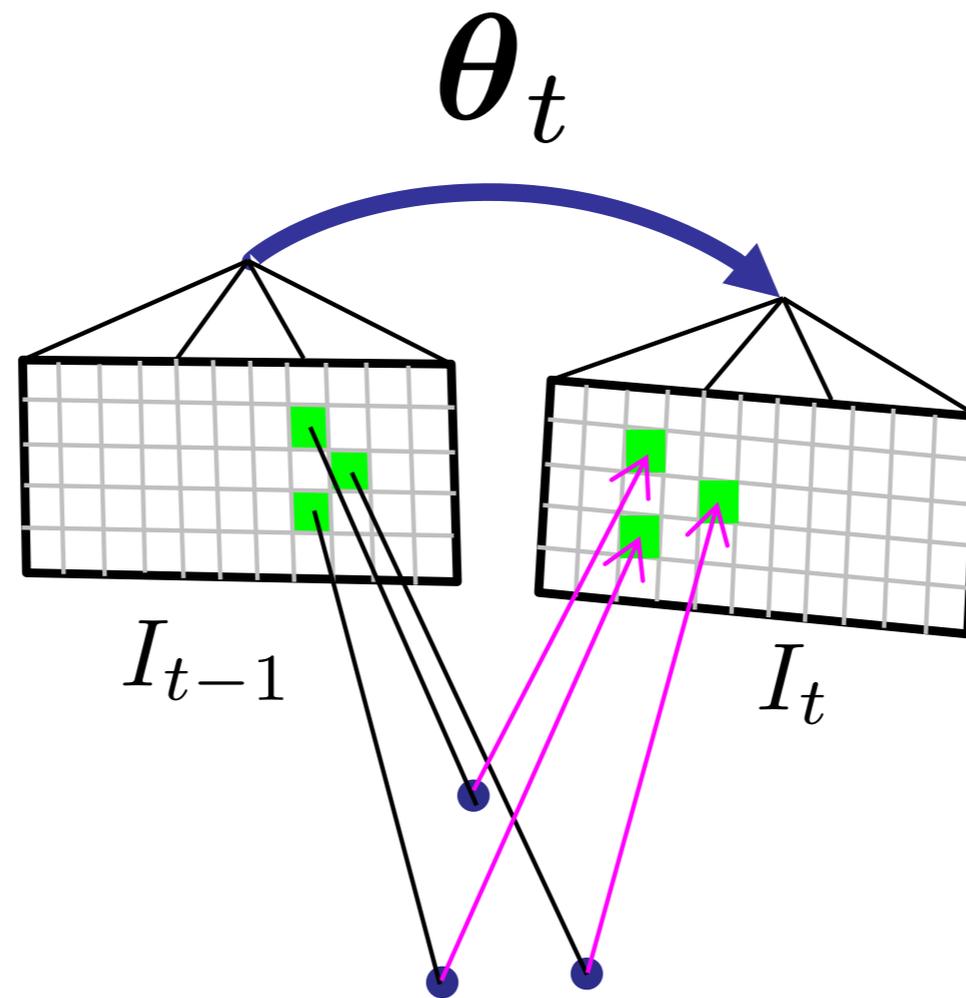


Today

- SfM - Bundle Adjustment
- VSLAM - Keyframe vs. Filtering
- **Visual Odometry**
- Loop Closure

VO vs SFM

- VO is a **particular case** of SFM
- VO focuses on estimating the 3D motion of the camera **sequentially** (as a new frame arrives) and in **real time**.
- Terminology: sometimes SFM is used as a synonym of VO.

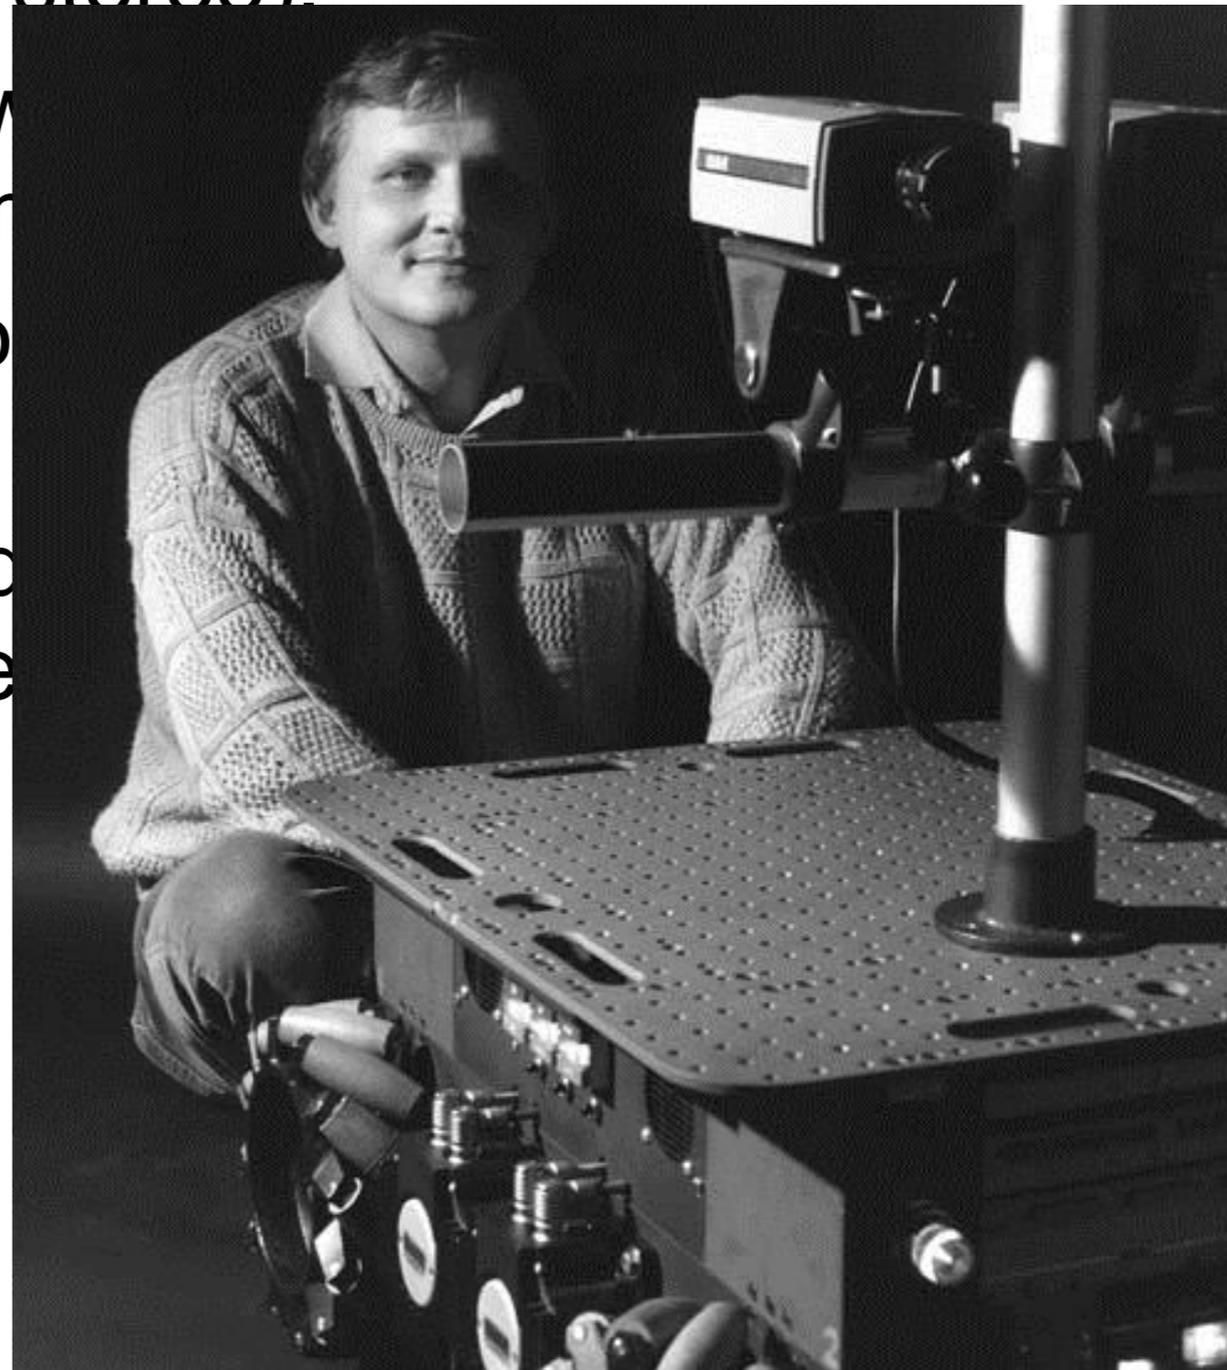


A Brief History of VO

- **1980:** First known VO real-time implementation on a robot by Hans Moravec PhD thesis (NASA/JPL) for Mars rovers using one sliding camera (sliding stereo).
- **1980 to 2000:** The VO research was dominated by NASA/JPL in preparation of 2004 Mars mission.
- **2004:** VO used on a robot on another planet: Mars rovers Spirit and Opportunity
- **2004.** VO was revived in the academic environment by Nister et al. The term VO became popular.

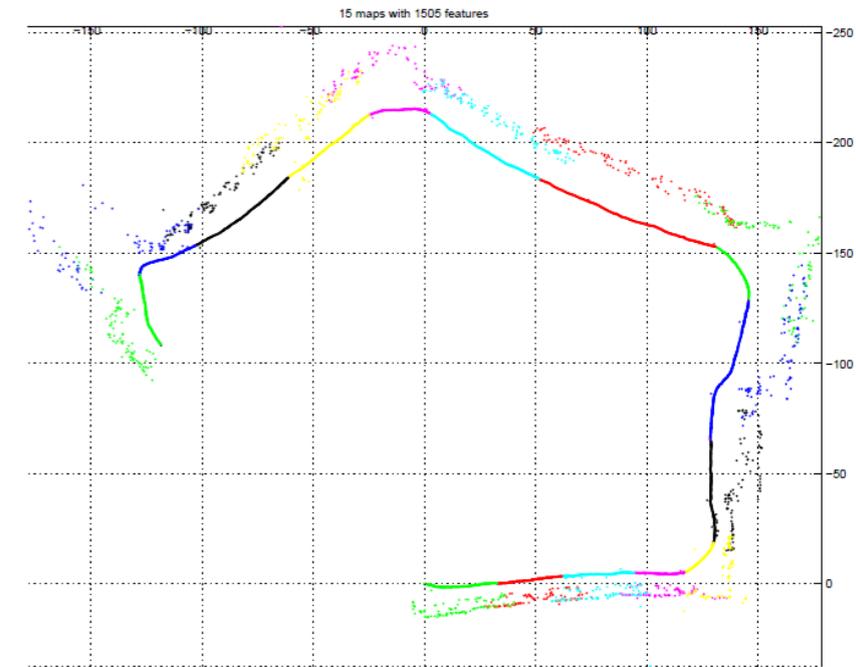
A Brief History of VO

- **1980:** First known VO real-time implementation on a robot by Hans Moravec PhD thesis (NASA/JPL) for Mars rovers using one sliding camera (sliding stereo).
- **1980 to 2000:** The VO research was done at JPL in preparation of 2004 Mars rovers.
- **2004:** VO used on a robot on another planet Spirit and Opportunity
- **2004.** VO was revived in the academic community by Nister et al. The term VO became

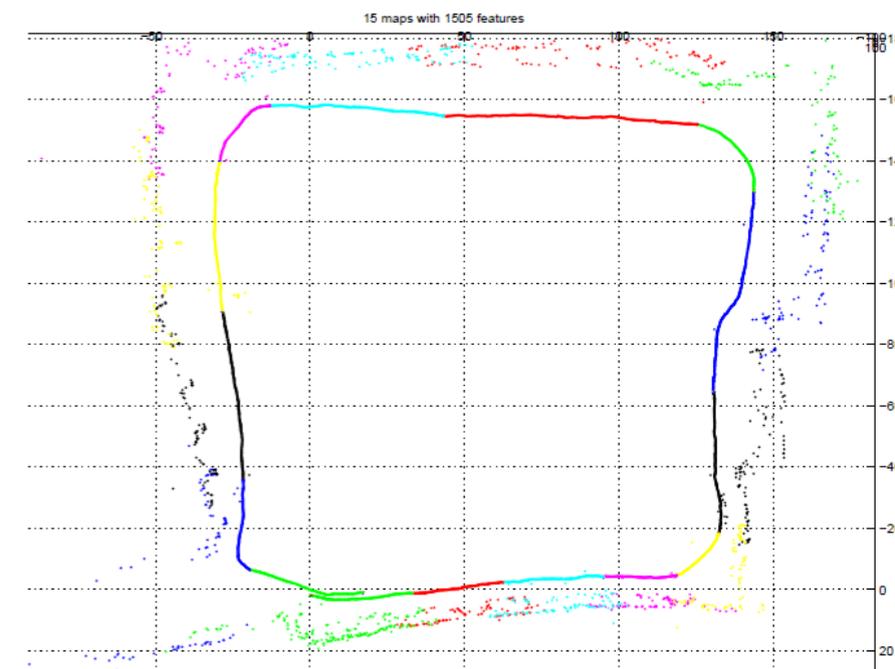


VO vs Visual SLAM

- VO only aims to the local consistency of the trajectory.
- SLAM aims to the global consistency of the trajectory and of the map.
- VO can be used as a building block of SLAM.
- VO is SLAM before **closing the loop!**
- The choice between VO and V-SLAM depends on the tradeoff between performance and consistency, and simplicity in implementation.
- VO trades off consistency for real-time performance, without the need to keep track of all the previous history of the camera.

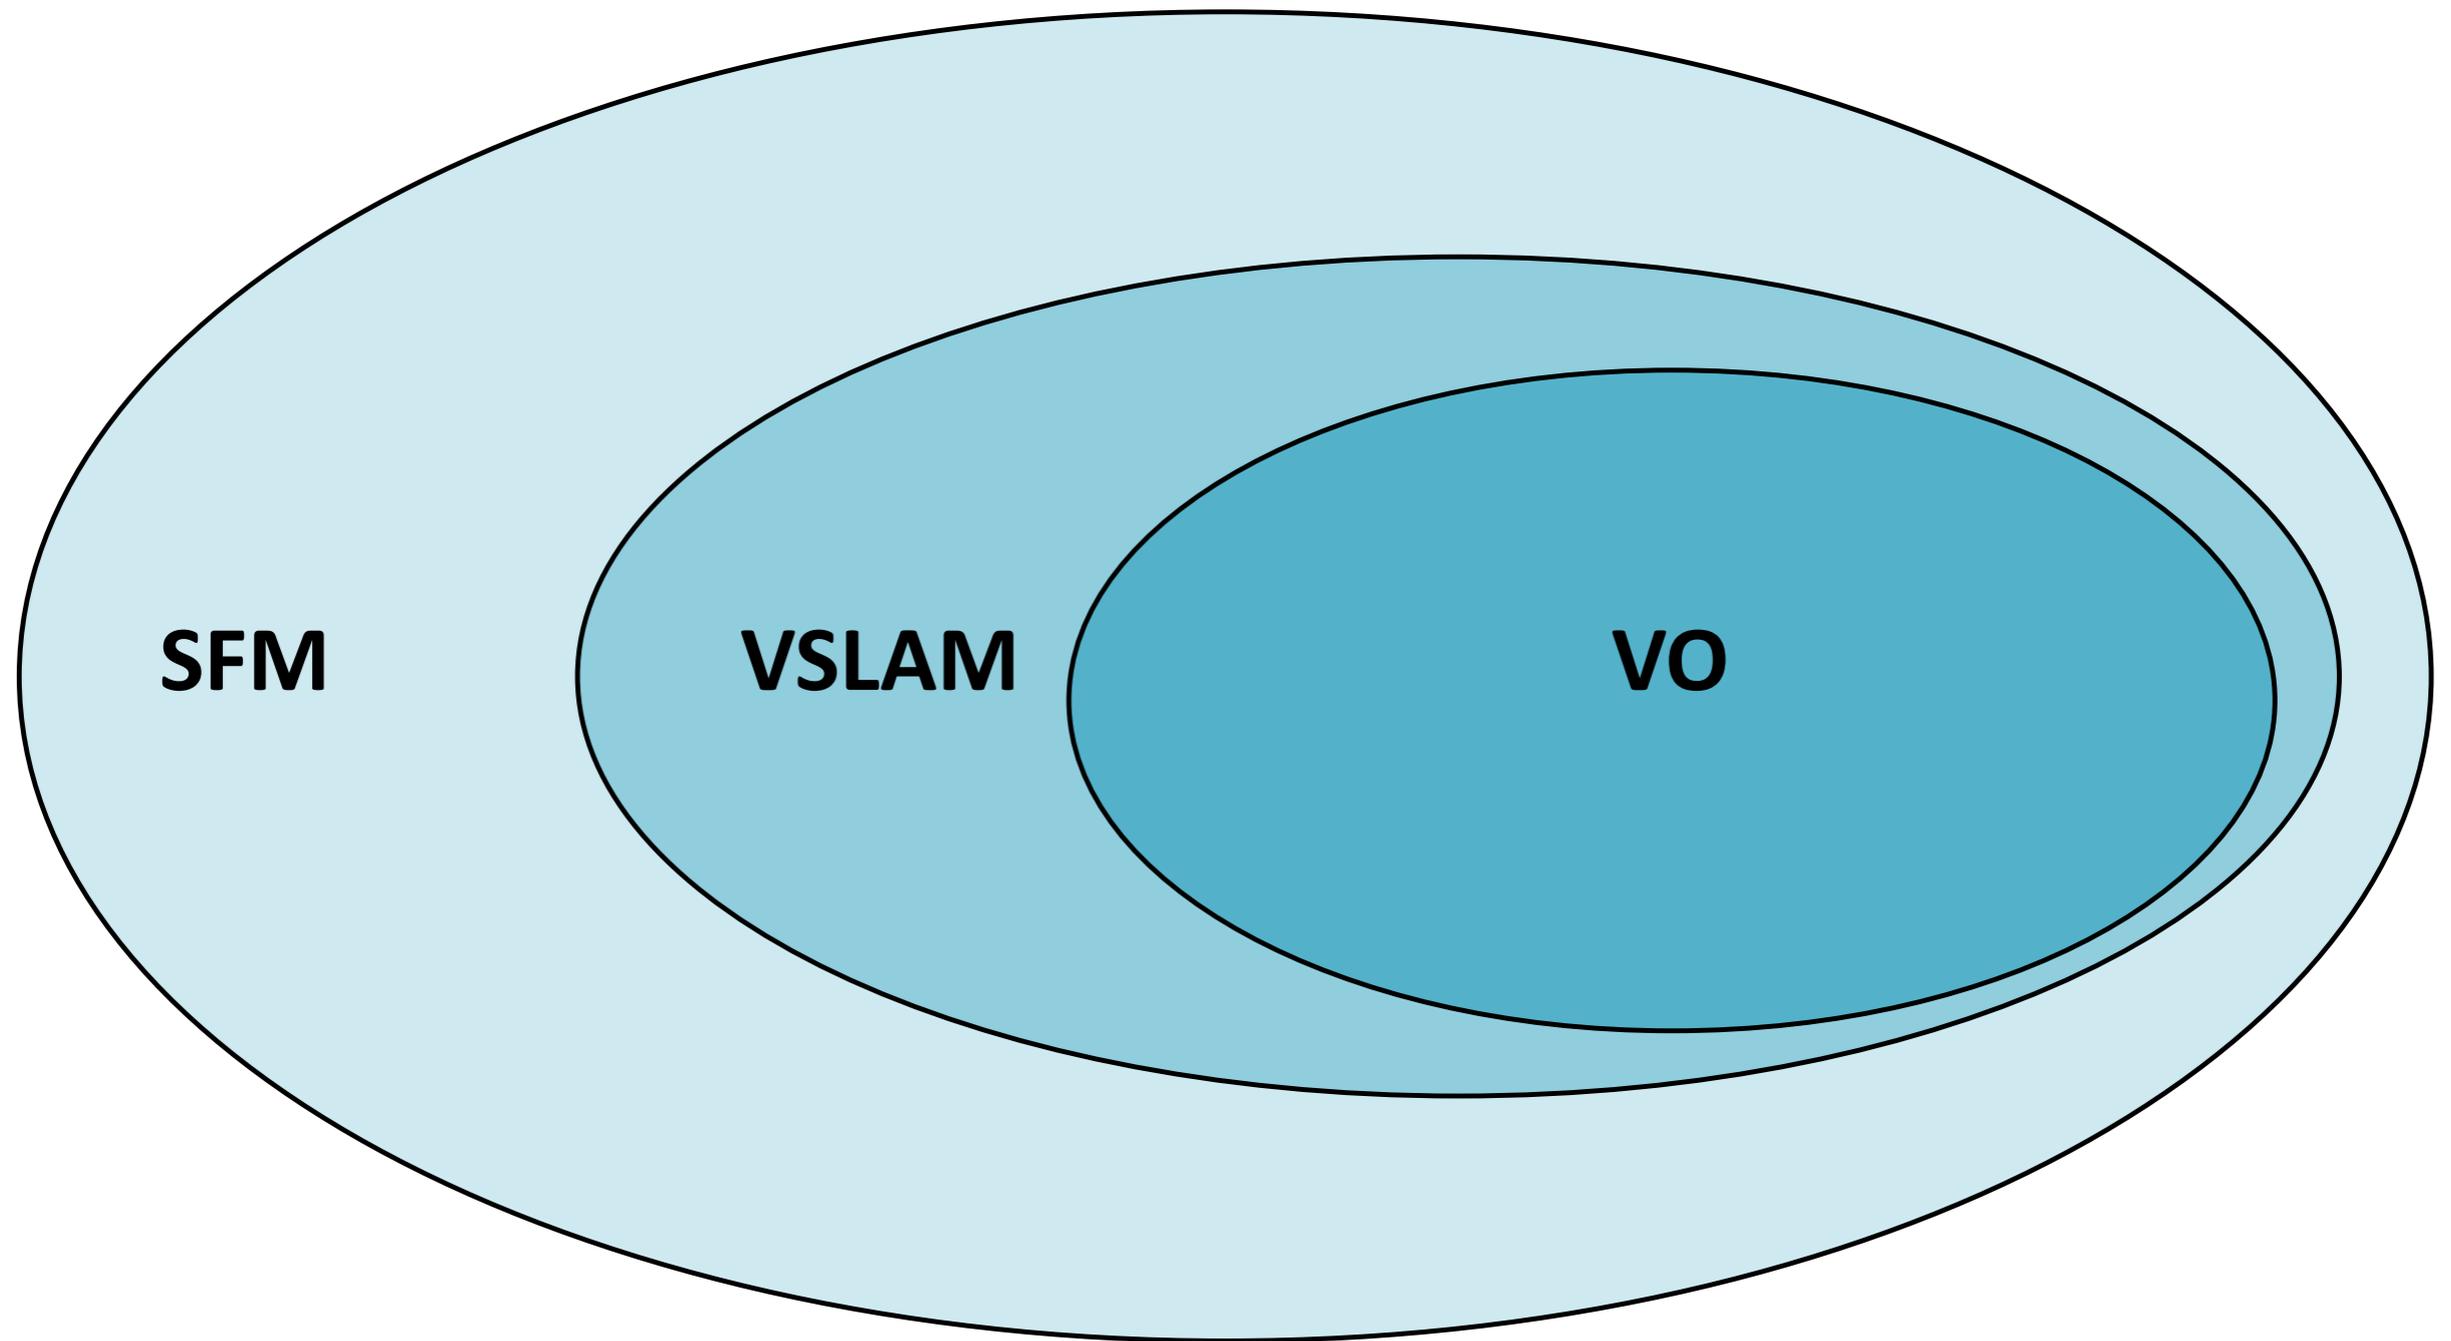


Visual odometry



Visual SLAM

VO vs VSLAM vs SFM



Today

- SfM - Bundle Adjustment
- VSLAM - Keyframe vs. Filtering
- Visual Odometry
- **Loop Closure**

Loop Closure Detection

- Loop constraints are very valuable constraints for local BA.
- Loop constraints can be found by evaluating visual similarity between the current camera images and past camera images.
- Visual similarity can be computed using global image descriptors (GIST descriptors) or local image descriptors (e.g., ORB features).
- Image retrieval is the problem of finding the most similar image of a template image in a database of billion images (image retrieval).
- This can be solved efficiently with Bag of Words.

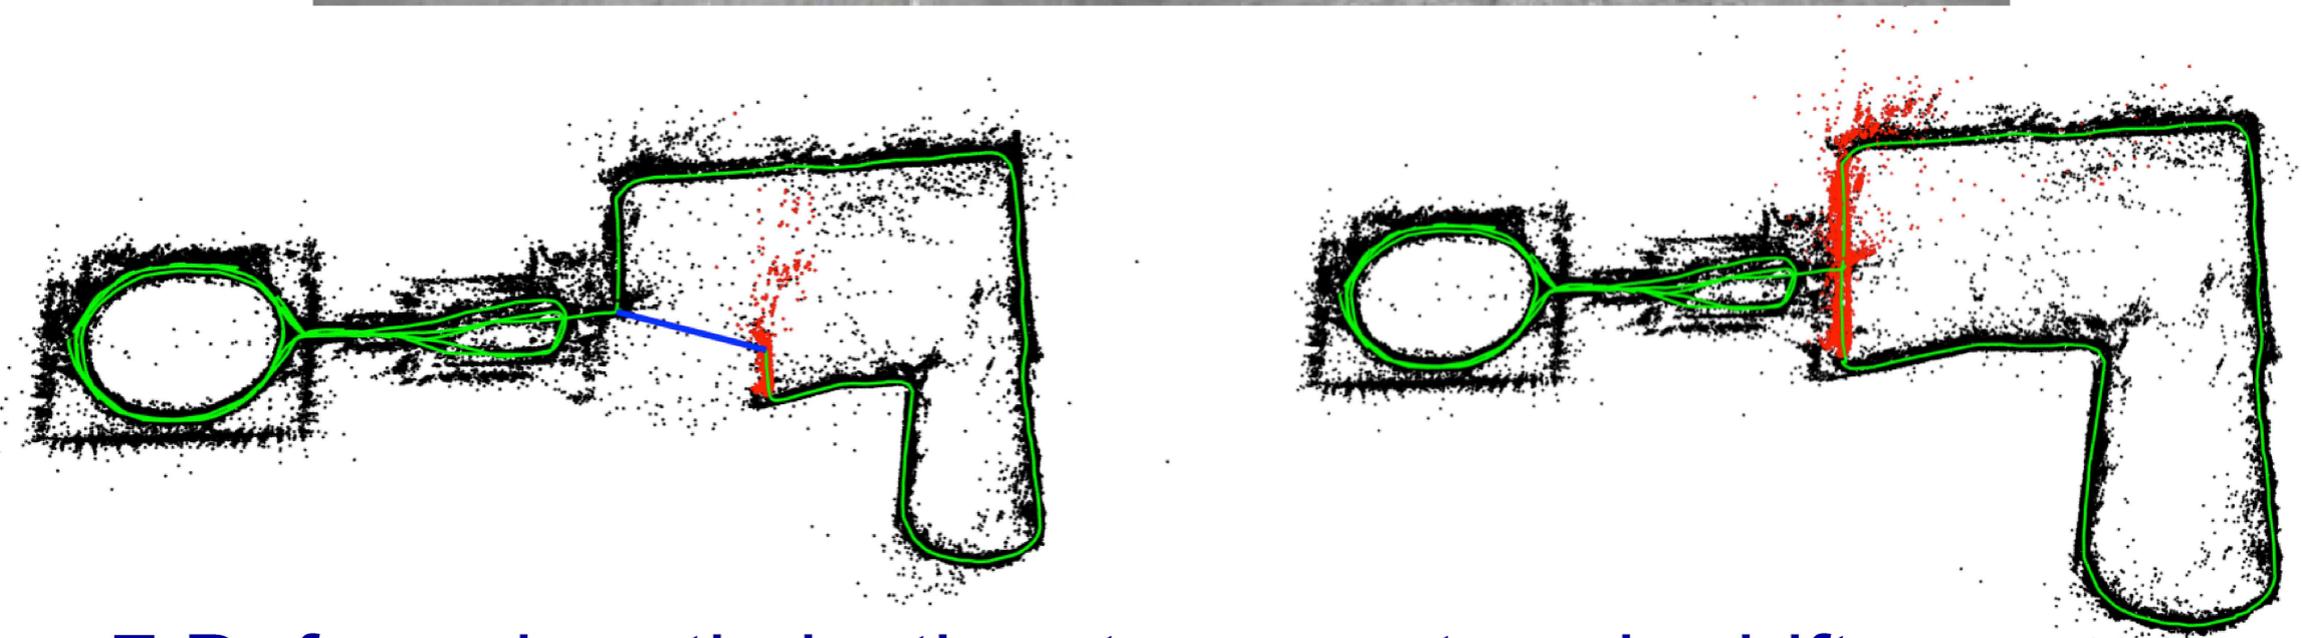
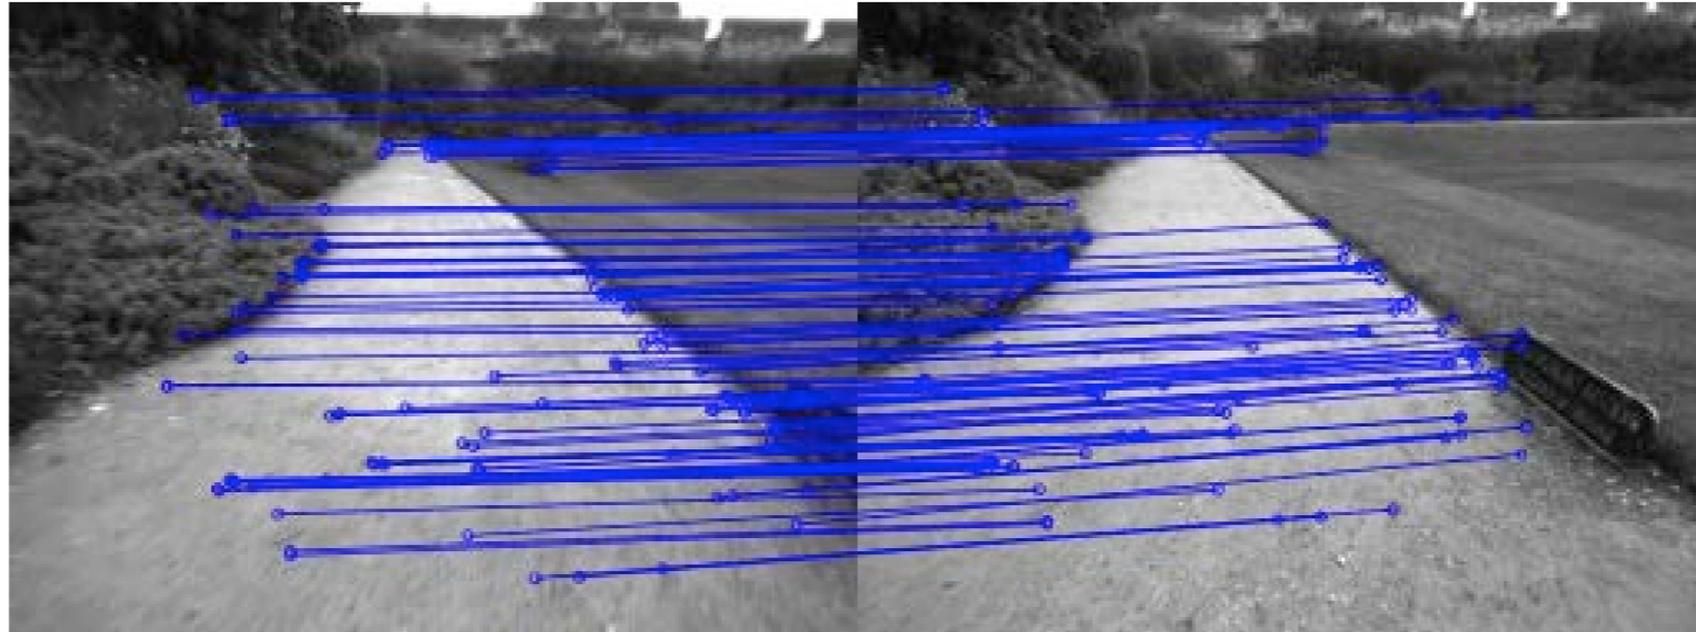


First observation



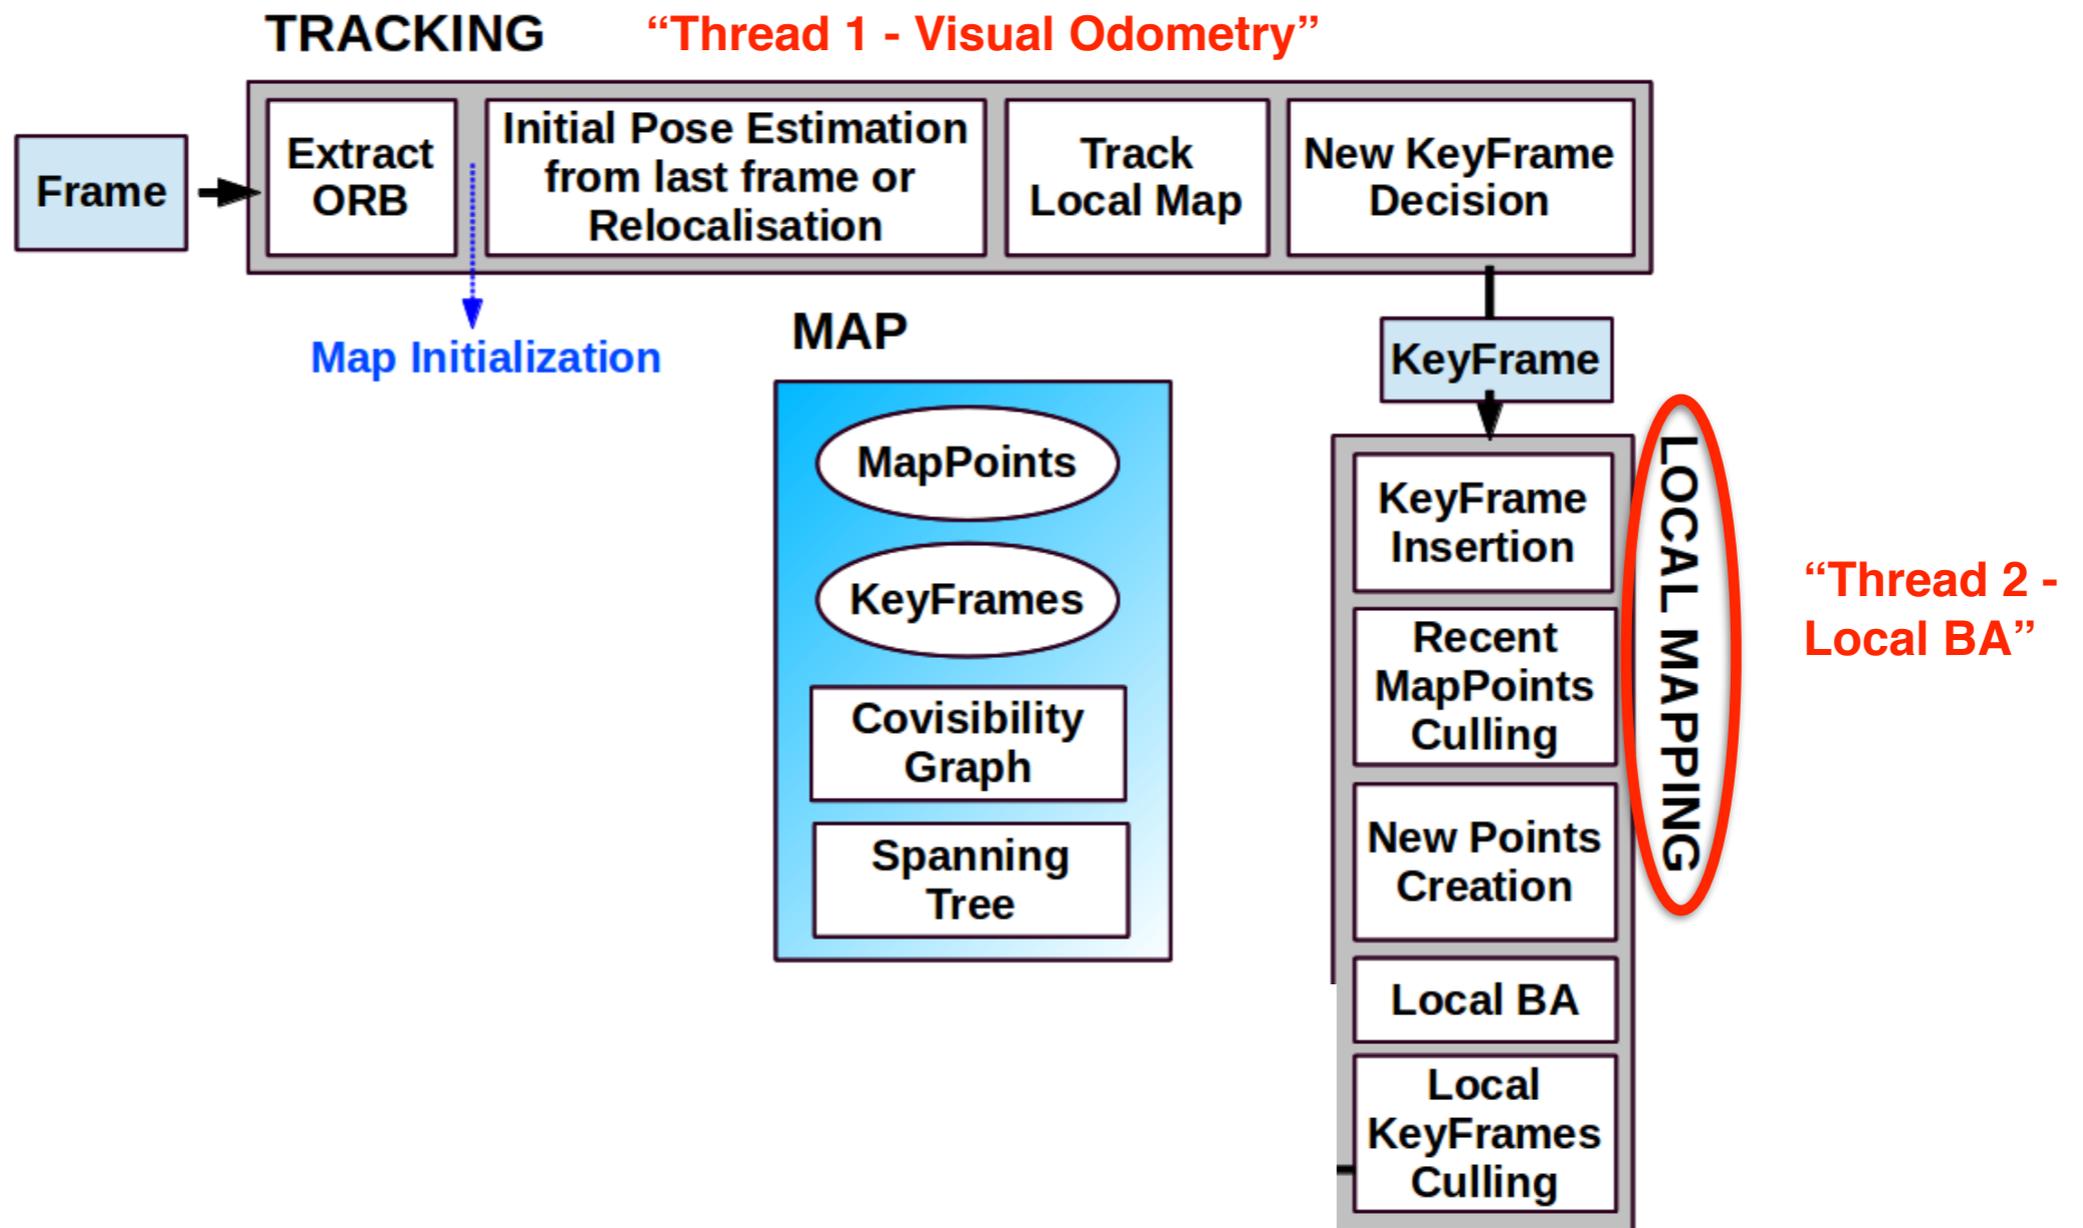
Second observation after a loop

Loop Closure Detection

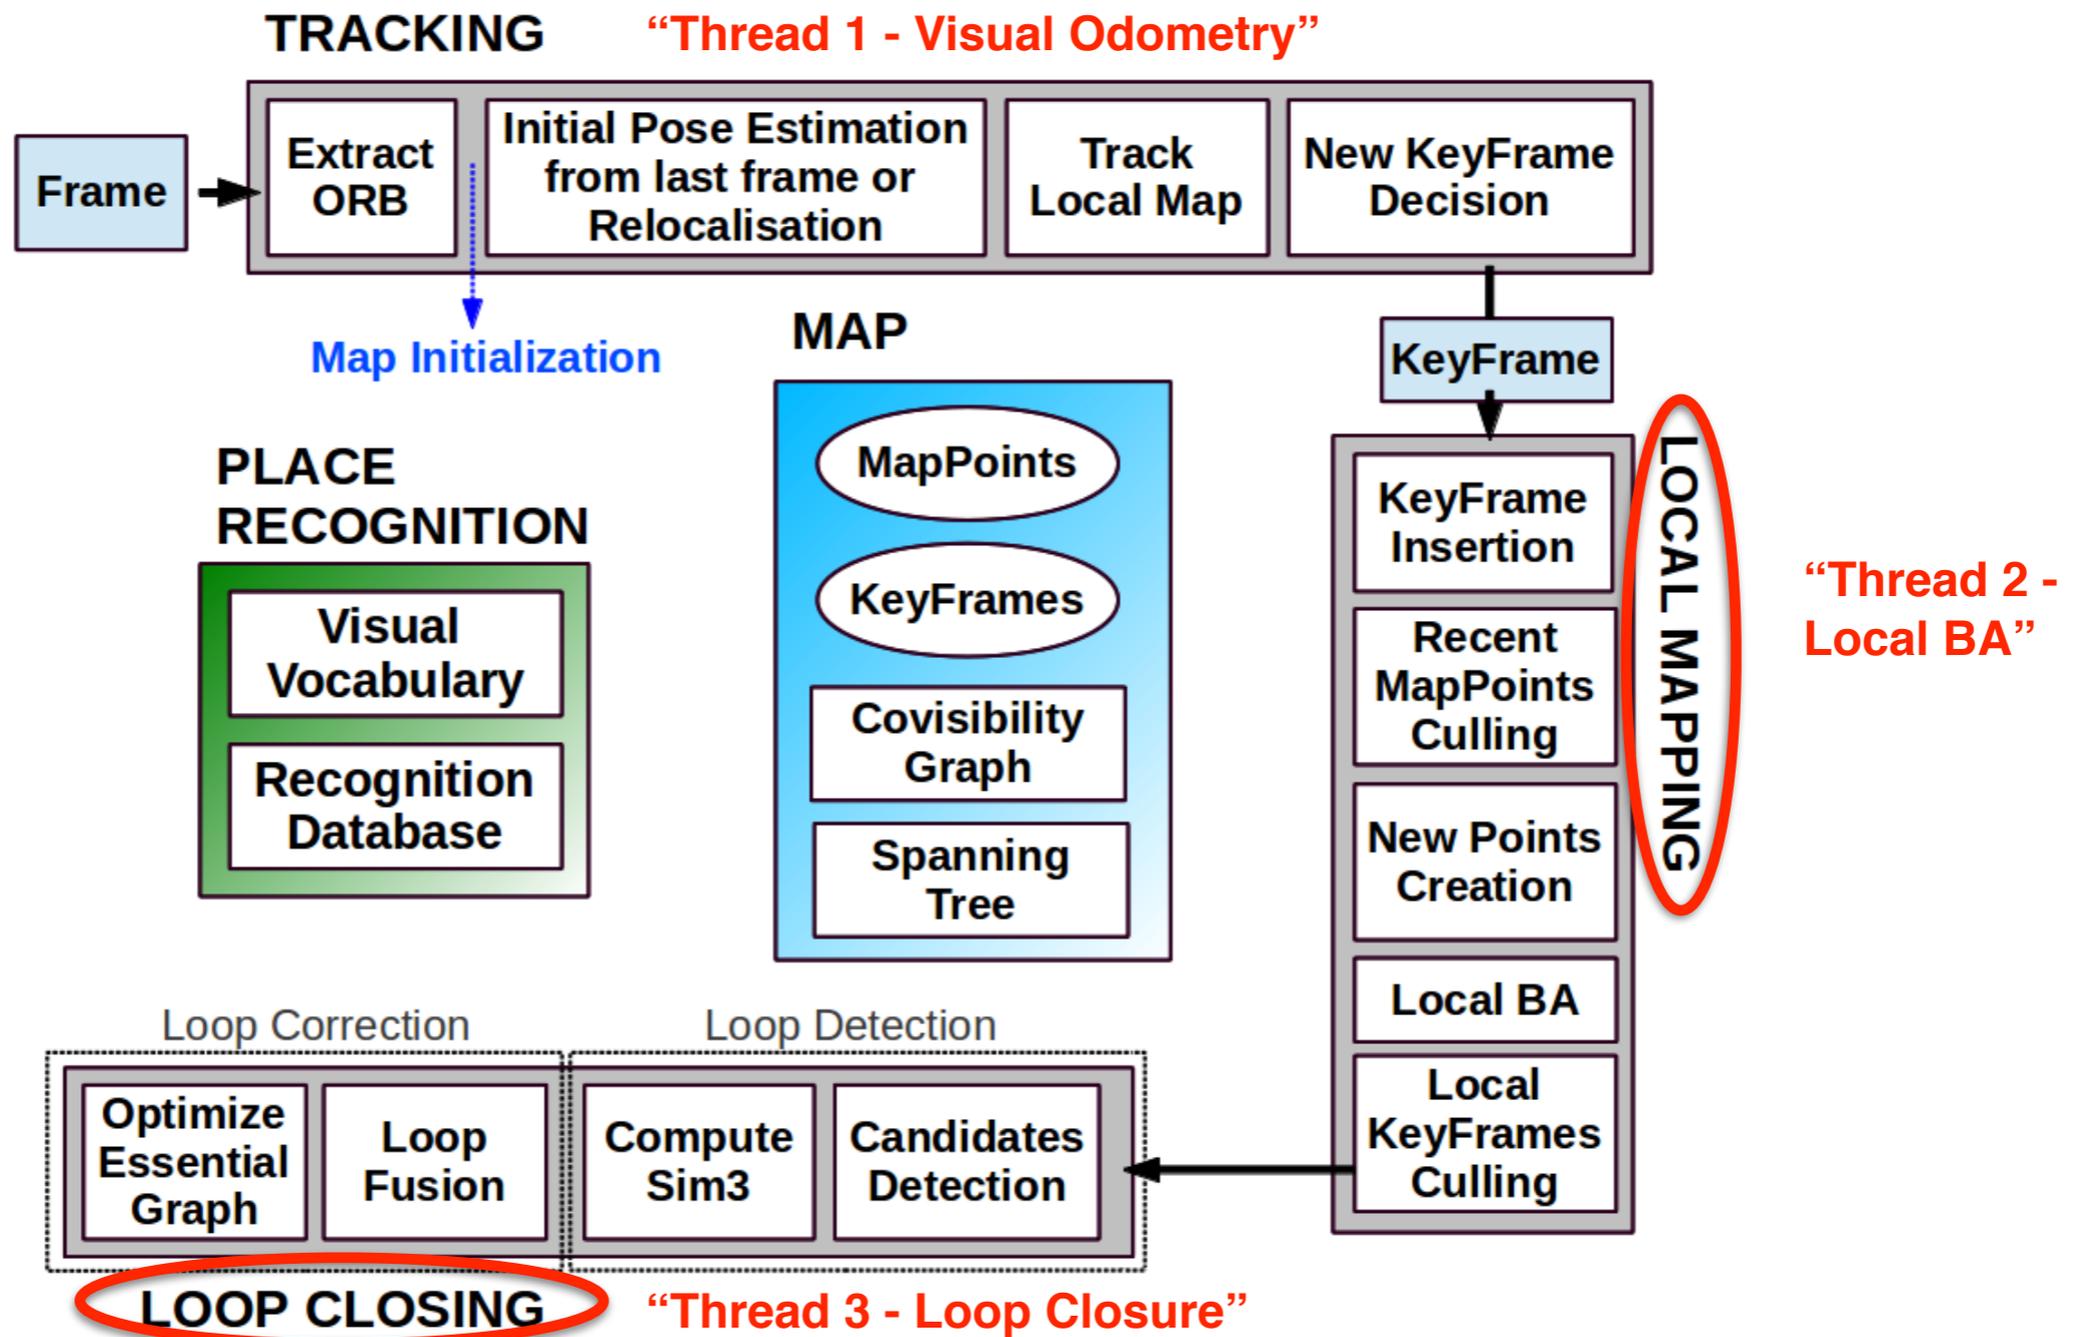


- 7 Dof graph optimization, to correct scale drift
- And optionally Full BA (little improvement, much slower)

Example - ORB SLAM



Example - ORB SLAM



ORB SLAM

- Essentially “greatest hits” in terms of work in VSLAM.
- Uses the same features (i.e. ORB) for,
 - Tracking
 - Mapping
 - Loop closure
- Real-time, large scale operation.
- Survival of the fittest for points and keyframes.
- Further information can be found at,
 - <http://webdiis.unizar.es/~raulmur/orbslam/>
 - Source code available under GPLv3.

ORB SLAM

WAITING FOR IMAGES



ORB SLAM

WAITING FOR IMAGES

