

# Lecture 28, Mar 18, 2026

## Batch SLAM

- Filtering-based SLAM keeps only the most recent pose, leading to issues over long-term
  - Filters lock in their decisions; since we don't do global optimization, new information doesn't always allow us to correct for earlier mistakes, e.g. large loop closures
    - \* EKF-SLAM needs uncertainties to be kept perfects to backpropagate the loop closure uncertainty
    - \* Particle filters need lots of particles to keep track of the growing uncertainty and map distribution
  - By marginalizing out previous states (i.e. removing  $\mathbf{x}_{k-1}$ ), EKF-SLAM destroys the sparse arrowhead structure of the information matrix as all landmark poses start being correlated
- In batch SLAM, we keep previous poses around and perform global optimization upon loop closure, similar to bundle adjustment
  - We build a graph using odometry measurements and landmark observations, with nodes representing poses/landmarks and edges representing measurements
  - The SLAM problem reduces to finding the configuration of the poses that minimizes a cost
  - Analogous to a network of springs where the stiffness of each spring is the confidence of the measurement, and finding the minimum energy configuration
- Consider landmarks with location  $\mathbf{m} = \begin{bmatrix} x \\ y \end{bmatrix}$ , initial pose  $\mathbf{x}_0 = \mathbf{0}$ , with odometry measurements  $\mathbf{u} = \begin{bmatrix} u \\ \omega \end{bmatrix}$

and landmark measurements  $\mathbf{y} = \begin{bmatrix} r \\ \phi \end{bmatrix}$ ; we want to find robot poses  $\mathbf{x} = \begin{bmatrix} x \\ y \\ \theta \end{bmatrix}$

- The full state consists of stacking all the poses and landmark locations into  $\mathbf{z} = [\mathbf{x}_1 \cdots \mathbf{x}_k \mathbf{m}_{1,1} \cdots \mathbf{m}_{k,l}]^T$
- Consider a motion model  $\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_k, \mathbf{w}_k)$ ,  $\mathbf{w}_k \sim (\mathbf{0}, \mathbf{Q})$  and an observation model  $\mathbf{y}_{kl} = \mathbf{g}(\mathbf{x}_k, \mathbf{m}_l, \mathbf{n}_k)$ ,  $\mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ 
  - \* Often we use a unicycle model and a range and bearing to landmarks:
  - \*  $\mathbf{x}_k = \mathbf{x}_{k-1} + T \begin{bmatrix} \cos \theta_{k-1} & 0 \\ \sin \theta_{k-1} & 0 \\ 0 & 1 \end{bmatrix} \mathbf{u}_k + \mathbf{w}_k$
  - \*  $\mathbf{y}_{kl} = \begin{bmatrix} \|\mathbf{x}_l - \mathbf{x}_k\| \\ \text{atan2}(y_l - y_k, x_l - x_k) \end{bmatrix} + \mathbf{n}_k$
- We form a motion model error  $\mathbf{e}_k = \mathbf{x}_k - \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_k, \mathbf{0}) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$  and observation model error  $\mathbf{e}_{kl} = \mathbf{y}_{kl} - \mathbf{g}(\mathbf{x}_k, \mathbf{m}_l, \mathbf{0}) \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$
- Define a quadratic loss:  $J = \frac{1}{2} \sum_{k=1}^K \mathbf{e}_k^T \mathbf{Q}^{-1} \mathbf{e}_k + \frac{1}{2} \sum_{k=1}^K \sum_{l=1}^L \mathbf{e}_{kl}^T \mathbf{R}^{-1} \mathbf{e}_{kl}$ 
  - \* This is equivalent to finding the  $\mathbf{z}$  that maximizes the likelihood of the measurements, assuming independence of measurements with each other and inputs, and using Gaussians
- To optimize the cost, we linearize the cost for a small step  $\delta \mathbf{x}$ , find the optimal step, and repeat
  - $\mathbf{e}_k = \mathbf{x}_k - \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_k, \mathbf{0})$ 

$$= \bar{\mathbf{x}}_k + \delta \mathbf{x}_k - \mathbf{f}(\bar{\mathbf{x}}_{k-1} + \delta \mathbf{x}_{k-1}, \mathbf{u}_k, \mathbf{0})$$

$$= \bar{\mathbf{x}}_k - \mathbf{f}(\bar{\mathbf{x}}_{k-1} + \delta \mathbf{x}_{k-1}, \mathbf{u}_k, \mathbf{0}) + \delta \mathbf{x}_k - \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \delta \mathbf{x}_{k-1}$$

$$= \bar{\mathbf{e}}_k + \delta \mathbf{x}_k - \mathbf{F}_{k-1} \delta \mathbf{x}_{k-1}$$
  - $\mathbf{e}_{kl} = \mathbf{y}_k - \mathbf{g}(\mathbf{x}_k, \mathbf{m}_l, \mathbf{0})$ 

$$= \mathbf{y}_k - \mathbf{g}(\bar{\mathbf{x}}_k + \delta \mathbf{x}_k, \bar{\mathbf{m}}_l + \delta \mathbf{m}_l, \mathbf{0})$$

$$= \bar{\mathbf{y}}_k - \mathbf{g}(\bar{\mathbf{x}}_k, \bar{\mathbf{m}}_l, \mathbf{0}) - \frac{\partial \mathbf{g}}{\partial \mathbf{x}} \delta \mathbf{x}_k - \frac{\partial \mathbf{g}}{\partial \mathbf{m}} \delta \mathbf{m}_l$$

$$= \bar{\mathbf{e}}_{kl} - \mathbf{G}_{x,kl} \delta \mathbf{x}_k - \mathbf{G}_{m,kl} \delta \mathbf{m}_l$$




Figure 1: Cost terms for a simple batch SLAM problem.

- Perturbed cost: 
$$J = \frac{1}{2} \sum_{k=1}^K (\bar{\mathbf{e}}_k + \delta \mathbf{x}_k - \mathbf{F}_{k-1} \delta \mathbf{x}_{k-1})^T \mathbf{Q}^{-1} (\bar{\mathbf{e}}_k + \delta \mathbf{x}_k - \mathbf{F}_{k-1} \delta \mathbf{x}_{k-1})$$

$$+ \frac{1}{2} \sum_{k=1}^K \sum_{l=1}^L (\bar{\mathbf{e}}_{kl} - \mathbf{G}_{x,kl} \delta \mathbf{x}_k - \mathbf{G}_{m,kl} \delta \mathbf{m}_l)^T \mathbf{R}^{-1} (\bar{\mathbf{e}}_{kl} - \mathbf{G}_{x,kl} \delta \mathbf{x}_k - \mathbf{G}_{m,kl} \delta \mathbf{m}_l)$$

$$= \frac{1}{2} (\bar{\mathbf{e}} + \mathbf{E} \delta \mathbf{z})^T \mathbf{T}^{-1} (\bar{\mathbf{e}} + \mathbf{E} \delta \mathbf{z})$$

\* Here  $\mathbf{T}$  is a large block diagonal matrix with the individual covariance matrices  $\mathbf{Q}$ ,  $\mathbf{R}$  repeated as many times as necessary along the diagonal

- To get the optimal update we solve  $\mathbf{E}^T \mathbf{T}^{-1} \mathbf{E} \delta \mathbf{z}^* = -\mathbf{E}^T \mathbf{T}^{-1} \bar{\mathbf{e}}$ ; starting with an initial guess (odometry and first observed landmark locations), we solve for  $\delta \mathbf{z}^*$  and update  $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{z}} + \delta \mathbf{z}^*$  and iterate until  $\delta \mathbf{z}^*$  is sufficiently small
- Since this method assumes locally quadratic costs, in practice we use the following techniques to improve robustness:
  1. *Levenberg-Marquardt* iteration:  $(\mathbf{E}^T \mathbf{T}^{-1} \mathbf{E} + \mu \text{diag}(\mathbf{E}^T \mathbf{T}^{-1} \mathbf{E})) \delta \mathbf{z}^* = -\mathbf{E}^T \mathbf{T}^{-1} \bar{\mathbf{e}}$  where  $\mu$  is a tunable parameter and  $\text{diag}(\mathbf{A})$  means a diagonal matrix obtained by setting all off-diagonal terms of  $\mathbf{A}$  to 0
    - \* Helps when  $\mathbf{A}$  is noninvertible and prevents large jumps for highly nonlinear functions
  2. *Line search*: After obtaining  $\delta \mathbf{z}^*$ , do a line search along its direction to make sure we don't go too far; do  $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{z}} + \alpha \delta \mathbf{z}^*$  where we pick the largest  $\alpha \in [0, 1]$  that still decreases  $J$
- $\mathbf{A} = \mathbf{E}^T \mathbf{T}^{-1} \mathbf{E}$  is the precision or inverse covariance of the final estimate of  $\bar{\mathbf{z}}$ ; it is sparse with an arrowhead structure, so the Schur complement can be used for inversion

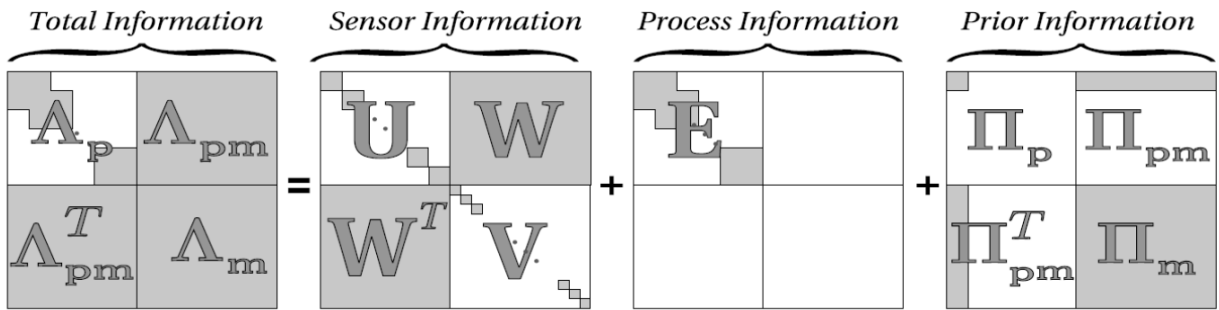


Figure 2: Structure of the inverse covariance matrix. With no prior, we obtain the sparse arrowhead structure.

- Running batch SLAM online requires changes to limit the growth of the problem

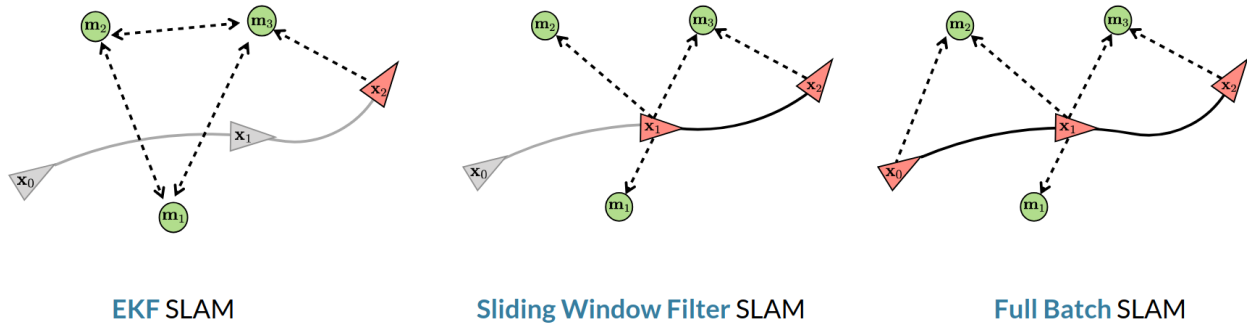


Figure 3: Sliding window filter SLAM is a hybrid between EKF-SLAM and full batch SLAM.

- Sliding window filter: marginalize out old poses and invisible landmarks beyond the sliding window, but keeps poses within the window
  - \* This is similar to a hybrid between EKF-SLAM/filtering based methods and batch SLAM
  - \* However loop closure is still a problem
- Adaptive relative bundle adjustment: represent poses relatively, and use an adaptive local window of active landmarks and poses for optimization
  - \* This takes advantage of the fact that in a large scene, typically most of the map changes very little, so most of it does not need to be optimized over
  - \* Landmark poses are represented relative to the robot pose that observed them, and robot poses are represented relative to the last pose
    - This removes the need for a single privileged global coordinate frame, which makes optimization costly for very large problems, since with a global frame a small change in the pose of the first steps leads to huge swings in the poses of subsequent steps
  - \* Only landmarks and robot poses within the active window are optimized
  - \* The active window adapts based on which poses have seen the current landmarks and the current error
    - On loop closure, it starts by activating the most recent poses, and activates more as needed until the error drops below a threshold
    - This avoids having to optimize over every landmark and pose for a loop closure