A Provably Consistent Method for Imposing Sparsity in Feature-based SLAM Information Filters

October 13, 2005

Matthew Walter, Ryan Eustice*, & John Leonard

Massachusetts Institute of Technology
*Johns Hokpins University

Roadmap

Recent related work (Information Form of SLAM)

- Sparse Extended Information Filters (Thrun, Koller, Ghahramani, Durrant-Whyte, and Ng, 2002)
- Thin Junction Tree Filter (Paskin, 2003)
- Treemap (Frese, 2004)
- Exactly Sparse Delayed State Filters (Eustice, 2005)
- Graphical SLAM (Folkesson and Christensen, 2005)
- Graph SLAM (Thrun, Burgard, and Fox, 2005, Chapter 11)
- Square Root SAM (Daellart, 2005)
- D-SLAM (Wang, Huang, and Dissanayake, 2005)

Recent related work (Information Form of SLAM)

- Sparse Extended Information Filters (Thrun, Koller, Ghahramani, Durrant-Whyte, and Ng, 2002)
- Thin Junction Tree Filter (Paskin, 2003)
- Treemap (Frese, 2004)
- Exactly Sparse Delayed State Filters (Eustice, 2005)
- Graphical SLAM (Folkesson and Christensen, 2005)
- Graph SLAM (Thrun, Burgard, and Fox, 2005, Chapter 11)
- Square Root SAM (Daellart, 2005)
- D-SLAM (Wang, Huang, and Dissanayake, 2005)

Canonical Gaussian Parameterization for SLAM

$$\boldsymbol{\xi}_t \sim \mathcal{N}\left(\boldsymbol{\mu}_t, \Sigma_t\right)$$
$$\sim \mathcal{N}^{-1}\left(\boldsymbol{\eta}_t, \Lambda_t\right) \longleftarrow$$

$\Lambda_t = \Sigma_t^{-1}$	information matrix
$\boldsymbol{\eta}_t = \Lambda_t \boldsymbol{\mu}_t$	information vector

- Encodes Markov random field

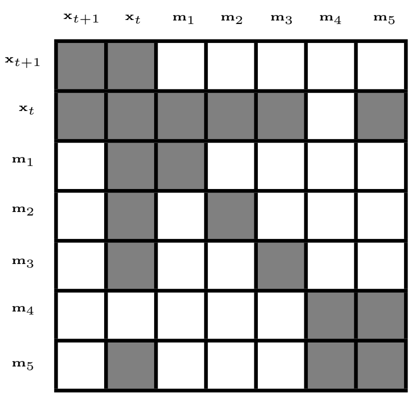

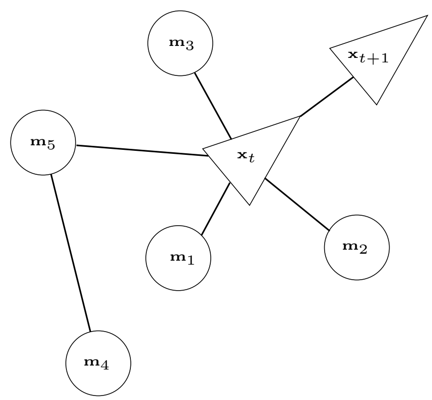

Represents independence relationships

$$p\left(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_3, \mathbf{m}_4, \mathbf{m}_5\right) = p\left(\mathbf{x}_{t+1} | \mathbf{x}_t\right)$$

Duality of standard and information forms

$$p\left(\boldsymbol{\alpha},\boldsymbol{\beta}\right)=\mathcal{N}\left(\begin{bmatrix}\boldsymbol{\mu}_\alpha\\ \boldsymbol{\mu}_\beta\end{bmatrix},\begin{bmatrix}\Sigma_{\alpha\alpha} & \Sigma_{\alpha\beta}\\ \Sigma_{\beta\alpha} & \Sigma_{\beta\beta}\end{bmatrix}\right)=\mathcal{N}^{-1}\left(\begin{bmatrix}\boldsymbol{\eta}_\alpha\\ \boldsymbol{\eta}_\beta\end{bmatrix},\begin{bmatrix}\Lambda_{\alpha\alpha} & \Lambda_{\alpha\beta}\\ \Lambda_{\beta\alpha} & \Lambda_{\beta\beta}\end{bmatrix}\right)$$

	Covariance Form	Information Form
Marginalization $p\left(\boldsymbol{\alpha}\right)=\int p\left(\boldsymbol{\alpha},\boldsymbol{\beta}\right)d\boldsymbol{\beta}$	$\boldsymbol{\mu}'=\boldsymbol{\mu}_\alpha$ $\Sigma'=\Sigma_{\alpha\alpha}$ easy	$\boldsymbol{\eta}'=\boldsymbol{\eta}_\alpha-\Lambda_{\alpha\beta}\Lambda_{\beta\beta}^{-1}\boldsymbol{\eta}_\beta$ $\Lambda'=\Lambda_{\alpha\alpha}-\Lambda_{\alpha\beta}\Lambda_{\beta\beta}^{-1}\Lambda_{\beta\alpha}$ hard
Conditioning $p\left(\boldsymbol{\alpha}\mid\boldsymbol{\beta}\right)=\frac{p\left(\boldsymbol{\alpha},\boldsymbol{\beta}\right)}{p\left(\boldsymbol{\beta}\right)}$	$\boldsymbol{\mu}'=\boldsymbol{\mu}_\alpha+\Sigma_{\alpha\beta}\Sigma_{\beta\beta}^{-1}(\boldsymbol{\beta}-\boldsymbol{\mu}_\beta)$ $\Sigma'=\Sigma_{\alpha\alpha}-\Sigma_{\alpha\beta}\Sigma_{\beta\beta}^{-1}\Sigma_{\beta\alpha}$	$\boldsymbol{\eta}'=\boldsymbol{\eta}_\alpha-\Lambda_{\alpha\beta}\boldsymbol{\beta}$ $\Lambda'=\Lambda_{\alpha\alpha}$

SEIF's key insight 1:

Information matrix is "relatively" sparse

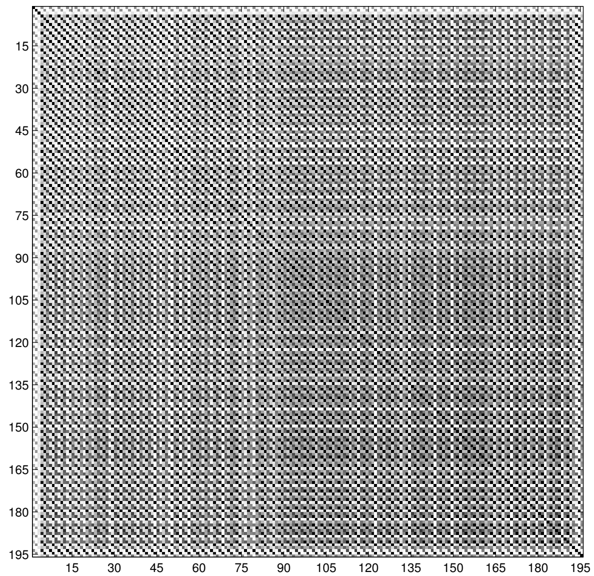

$$\Lambda_t = \Sigma_t^{-1}$$

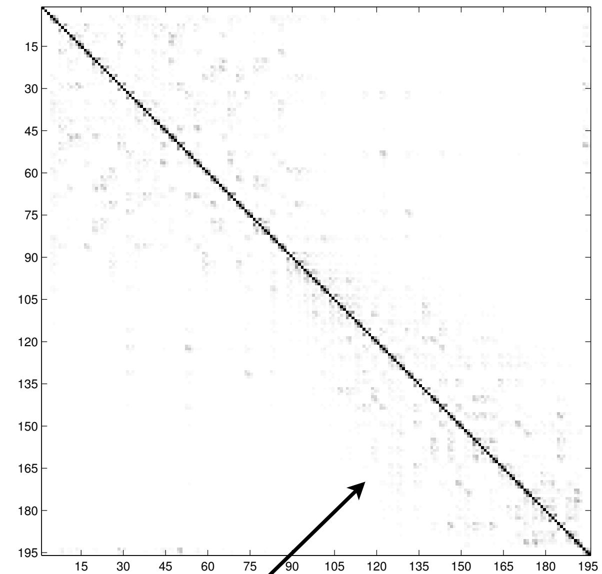

small but not zero!

Problem: time projection causes fill-in

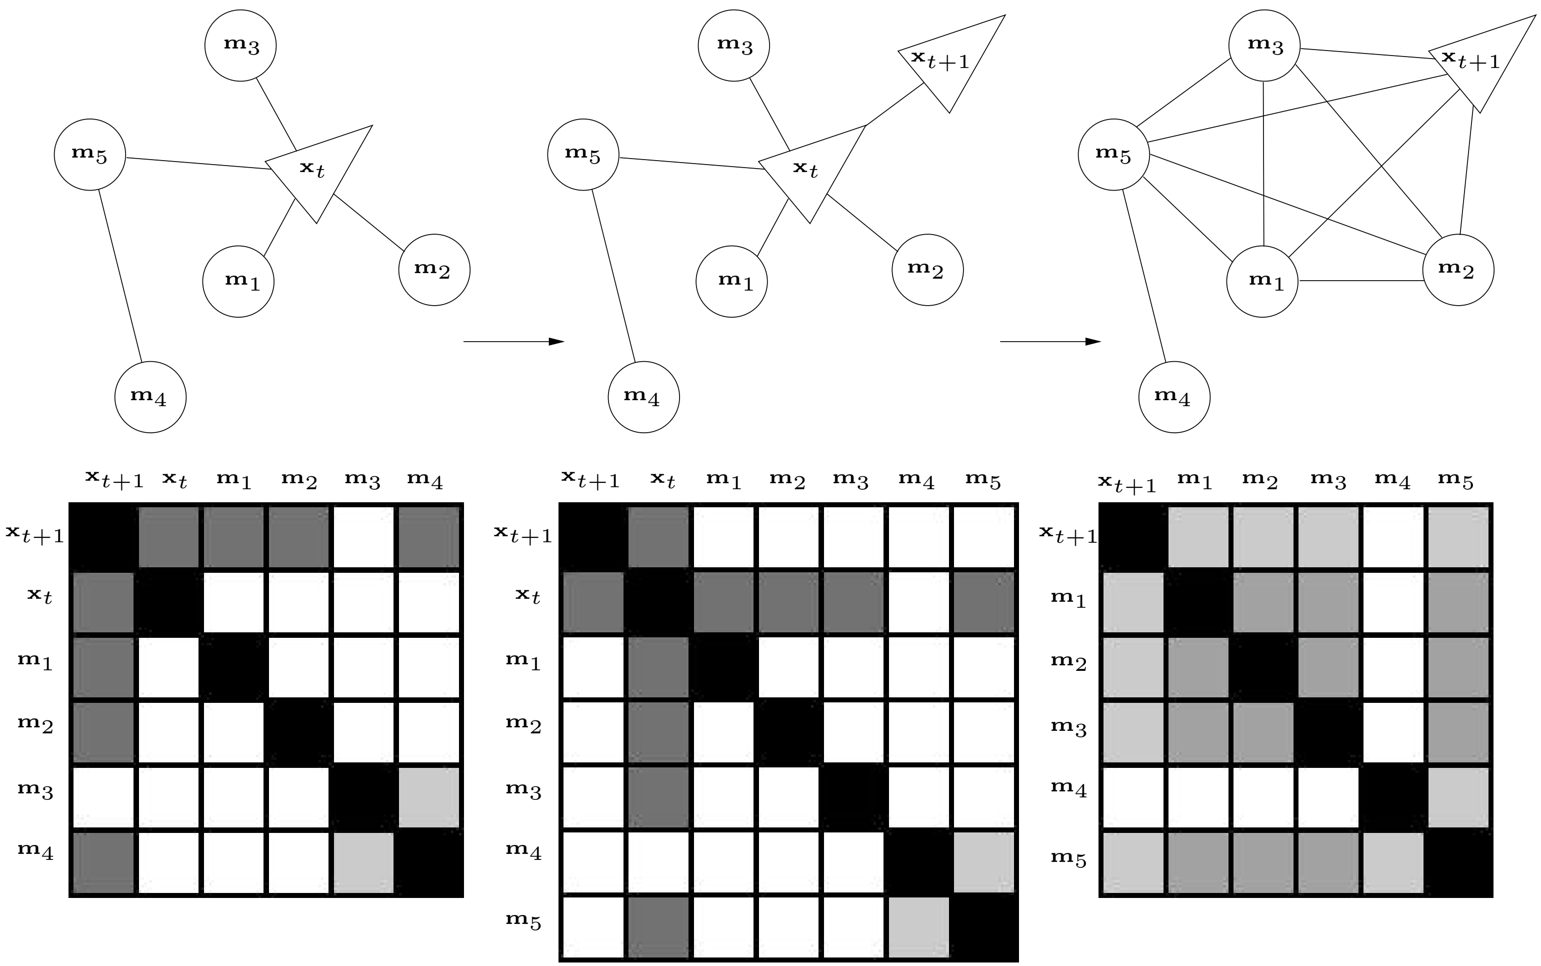

SEIF's Key Insight II: It is possible to limit this fill-in by bounding the number of "active" links

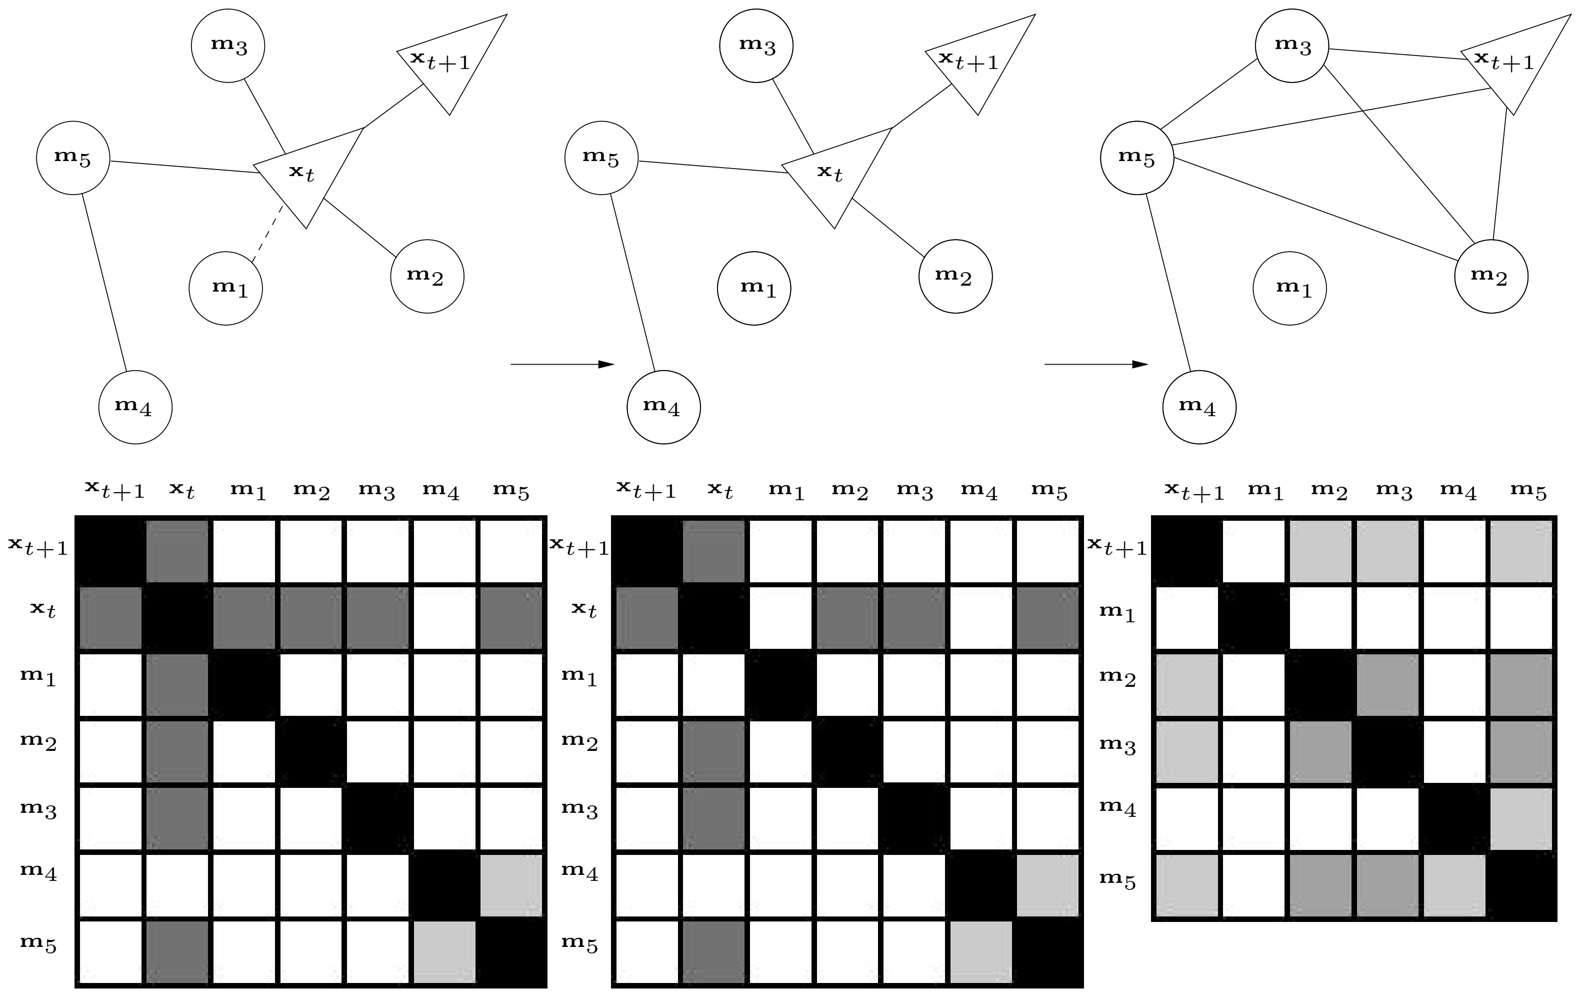

Roadmap

The SEIF algorithm cuts links to the weakest active features

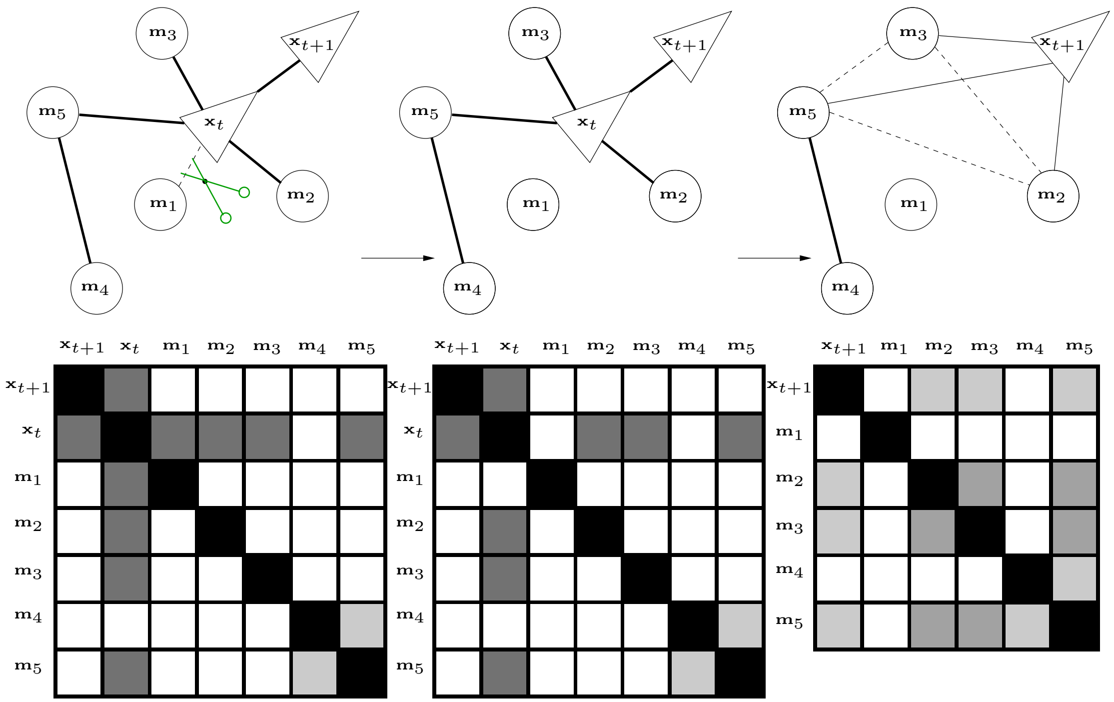

Link deactivation imposes conditional independence

Issue: this leads to overconfidence

$$
\begin{aligned}
p(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) &= p\left(\mathbf{x}_t \mid \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right) p\left(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right) \\
&= p\left(\mathbf{x}_t \mid \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^- = \boldsymbol{\alpha}\right) p\left(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right)
\end{aligned}
$$

ok

ok

$$
\begin{aligned}
\tilde{p}_{\text{SEIFs}}(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) &= p\left(\mathbf{x}_t \mid \mathbf{m}^+, \mathbf{m}^- = \boldsymbol{\alpha}\right) p\left(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right) \\
&= \frac{p_B(\mathbf{x}_t, \mathbf{m}^+ \mid \mathbf{m}^- = \boldsymbol{\alpha})}{p_C(\mathbf{m}^+ \mid \mathbf{m}^- = \boldsymbol{\alpha})} p_D(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-)
\end{aligned}
$$

NOT OK: can't ignore $\mathbf{m}^0$

Roadmap

I Introduction

II Sparsification in SEIFS

III ESEIFS

IV Results

V Conclusions

Exactly Sparse Extended Information Filters

- Reformulation of SEIFs that avoids approximation to enforce sparseness
- Key idea: periodically marginalize out the vehicle pose and “relocate” the vehicle into the map
- This step occurs at most at the same frequency as sparsification in SEIFs
- Compute a slightly modified posterior, using all of the feature measurements, and nearly all of the odometry information
- Prevents “fill-in” of the inverse covariance matrix for terms relating active and in-active features

Instead of breaking "weak" links, we control how active links are formed to maintain sparsity

activates m_3

$$\mathbf{Z}_t = \{\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_5\}$$

$$\mathbf{z}_\alpha = \{\mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_5\}$$

$$\mathbf{z}_\beta = \mathbf{z}_1$$

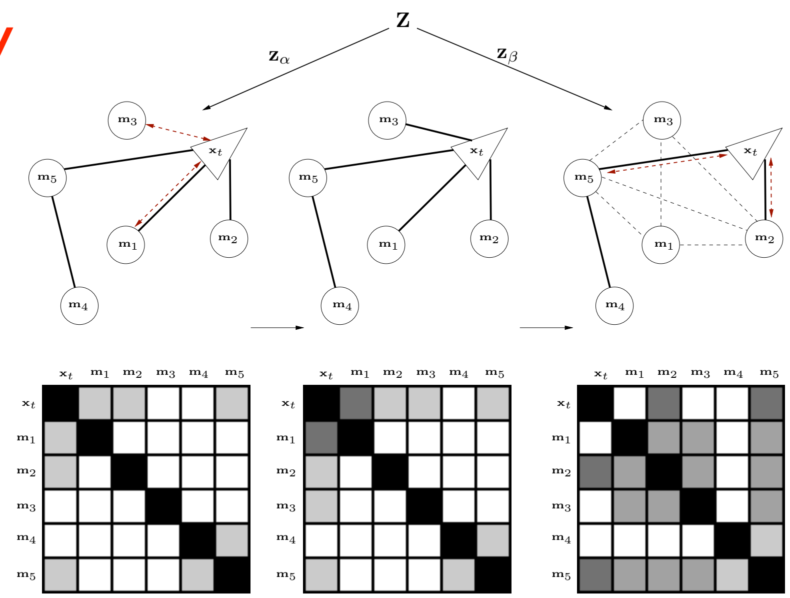

1) Update using $\mathbf{z}_\alpha$: $p(\boldsymbol{\xi}_t | \mathbf{z}^{t-1}, \mathbf{u}^t) \xrightarrow{\mathbf{z}_\alpha = \{\mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_5\}} p_1\big(\boldsymbol{\xi}_t | \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big)$

2) Marginalize robot : $p_2\big(\mathbf{M}_t | \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big) = \int_{\mathbf{x}_t} p_1\big(\boldsymbol{\xi}_t | \{\mathbf{z}^{t-1}, \mathbf{z}_\alpha\}, \mathbf{u}^t\big) d\mathbf{x}_t$

3) Relocate with $\mathbf{z}_\beta$: $\mathbf{x}_t = \mathbf{g}\big(\mathbf{m}_\beta, \mathbf{z}_\beta\big) + \mathbf{w}_t$

$$\approx \mathbf{g}\big(\breve{\boldsymbol{\mu}}_{m_\beta}, \mathbf{z}_\beta\big) + \mathrm{G}\big(\mathbf{m} - \breve{\boldsymbol{\mu}}_t\big) + \mathbf{w}_t$$

$\longrightarrow p_{\text{ESEIF}}(\boldsymbol{\xi}_t | \mathbf{z}^t, \mathbf{u}^t) = \mathcal{N}^{-1}\big(\boldsymbol{\xi}_t; \breve{\boldsymbol{\eta}}_t, \breve{\boldsymbol{\Lambda}}_t\big)$

Example

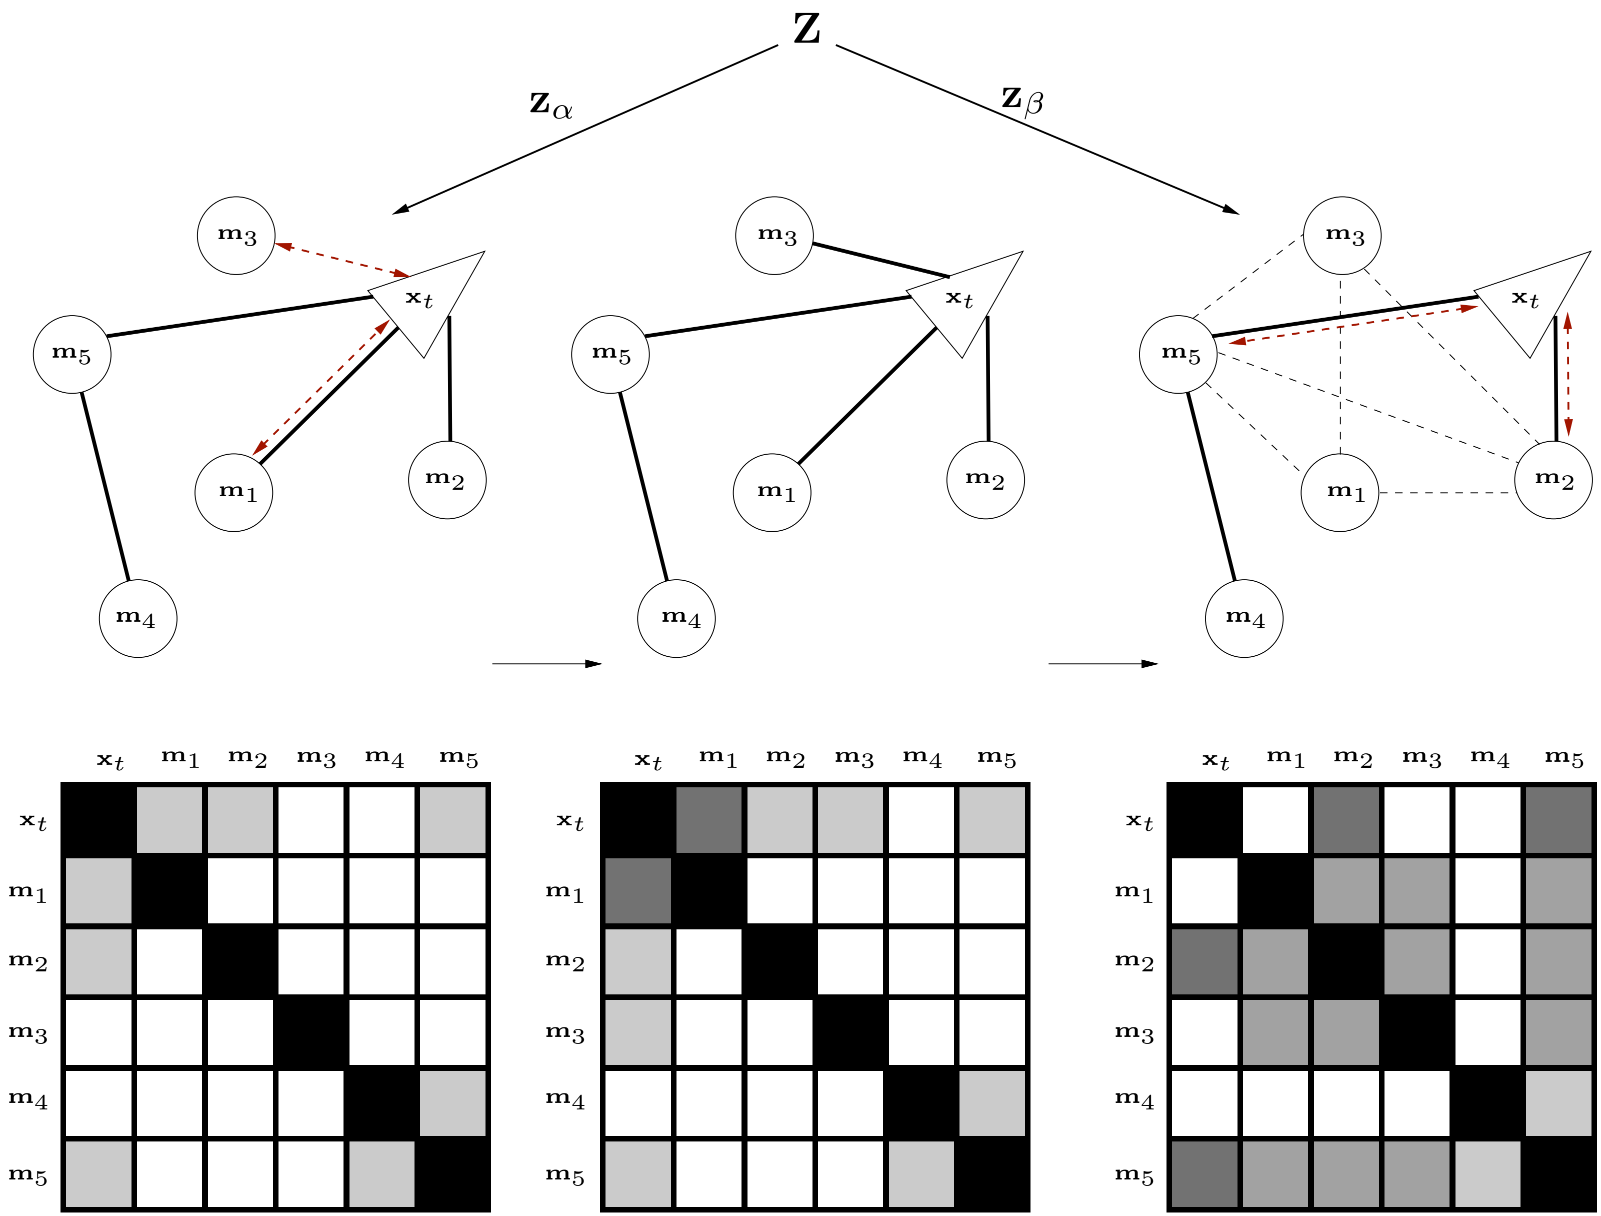

Z
$\mathbf{z}_\alpha$
$\mathbf{z}_\beta$
$\mathbf{x}_t$
$\mathbf{m}_1$
$\mathbf{m}_2$
$\mathbf{m}_3$
$\mathbf{m}_4$
$\mathbf{m}_5$

Complexity

(1) Update (2) marginalization and (3) relocalization are constant time (w/ estimate of mean*)

$$\check{\Lambda}_t = \mathrm{S}^{\top}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-}\bar{\Lambda}_t\mathrm{S}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-} - \mathrm{S}^{\top}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-}\bar{\Lambda}_t\mathrm{S}_{\mathbf{x}_t}\left(\mathrm{S}^{\top}_{\mathbf{x}_t}\bar{\Lambda}_t\mathrm{S}_{\mathbf{x}_t}\right)^{-1}\mathrm{S}^{\top}_{\mathbf{x}_t}\bar{\Lambda}_t\mathrm{S}_{\mathbf{m}^0,\mathbf{m}^+,\mathbf{m}^-}$$

↑ inverse of matrix O(robot pose)

Joint posterior exactly represents desired conditional independence

$$p_{\mathrm{ESEIF}}(\boldsymbol{\xi}_t|\mathbf{z}^t,\mathbf{u}^t) = p(\mathbf{x}_t|\mathbf{m}_\beta,\mathbf{z}_\beta)\,p_2(\mathbf{M}_t|\{\mathbf{z}^{t-1},\mathbf{z}_\alpha\},\mathbf{u}^t)$$

Ignoring odometry information ⟶ Exact sparsity

(Identical to the relocated-EKF)

*Linearization requires mean robot and update feature states. Efficiently accessible via multilevel relaxation.

Roadmap

Linear Gaussian simulation results (KF optimal)

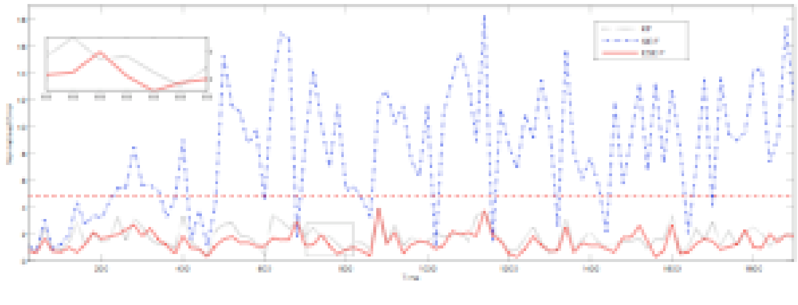

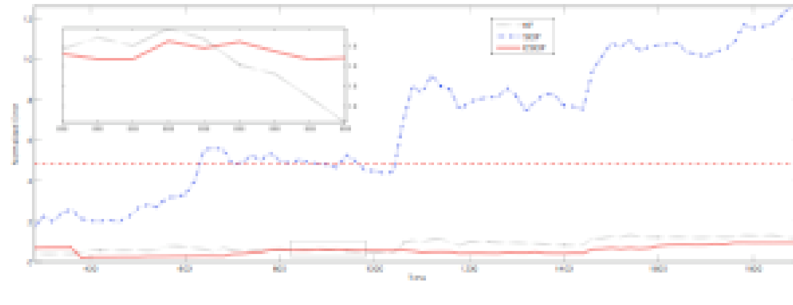

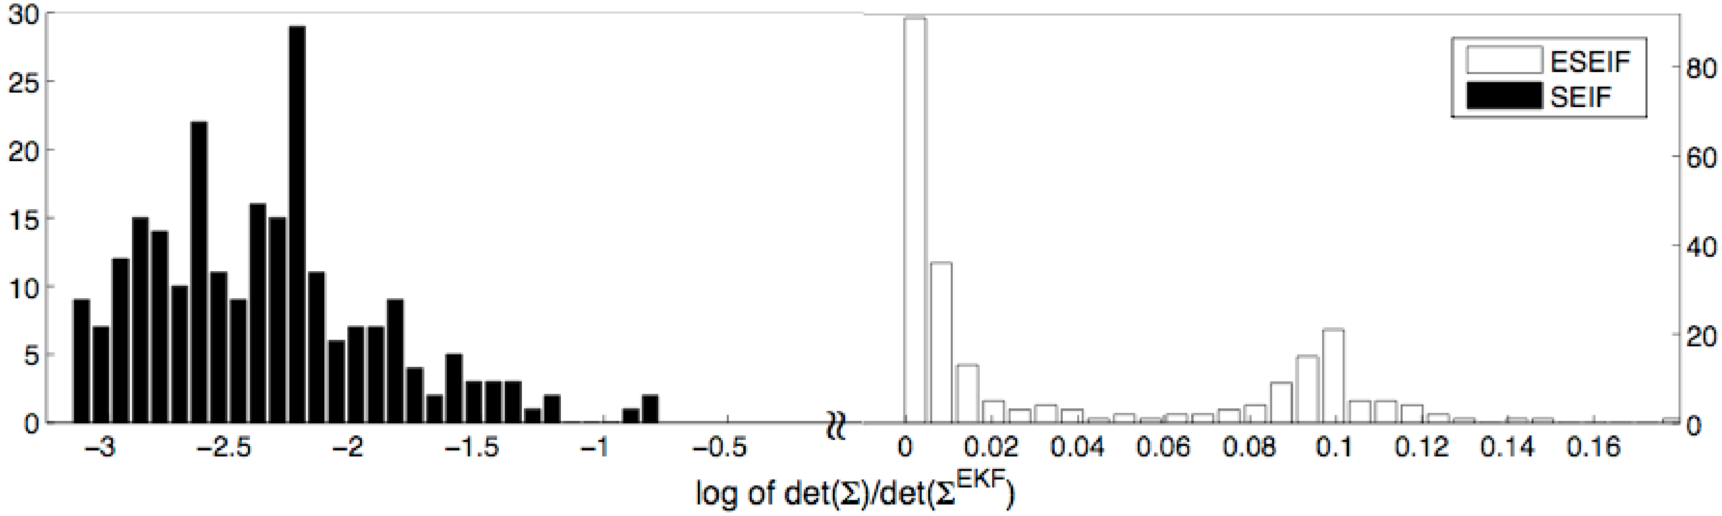

Nonlinear dataset 1: MIT Hurdles

MIT Indoor Track

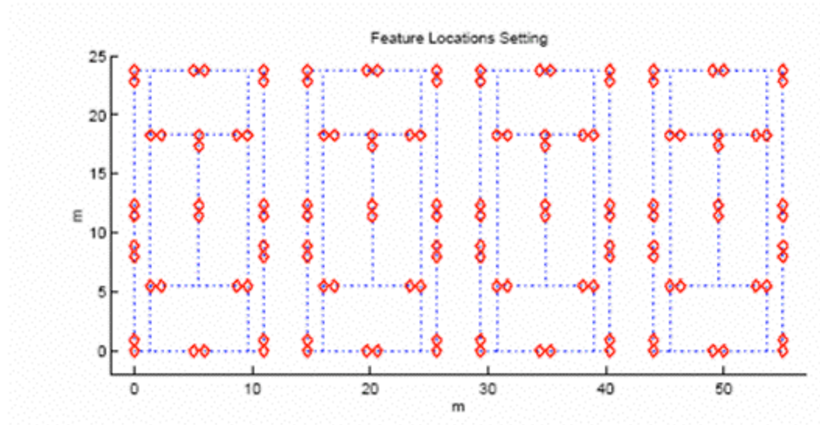

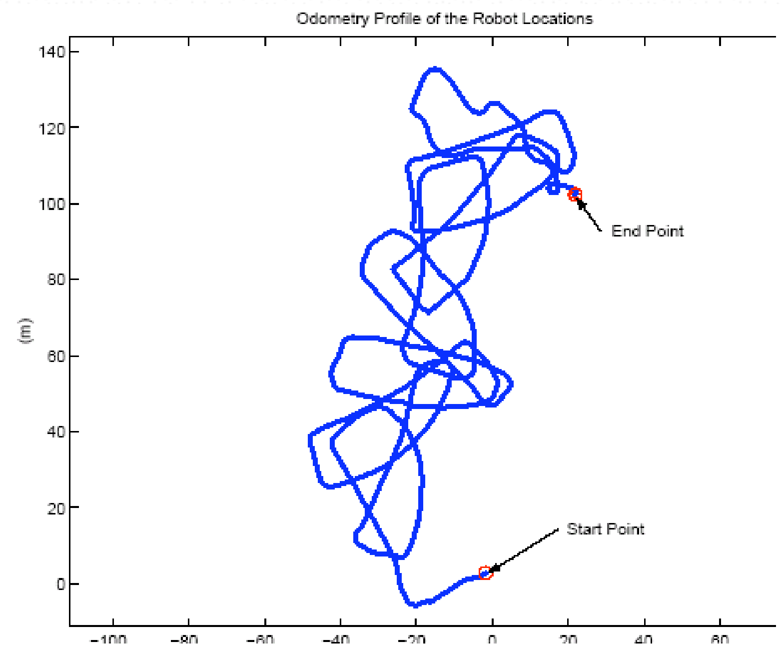

Results: Globally referenced maps

ESEIFS

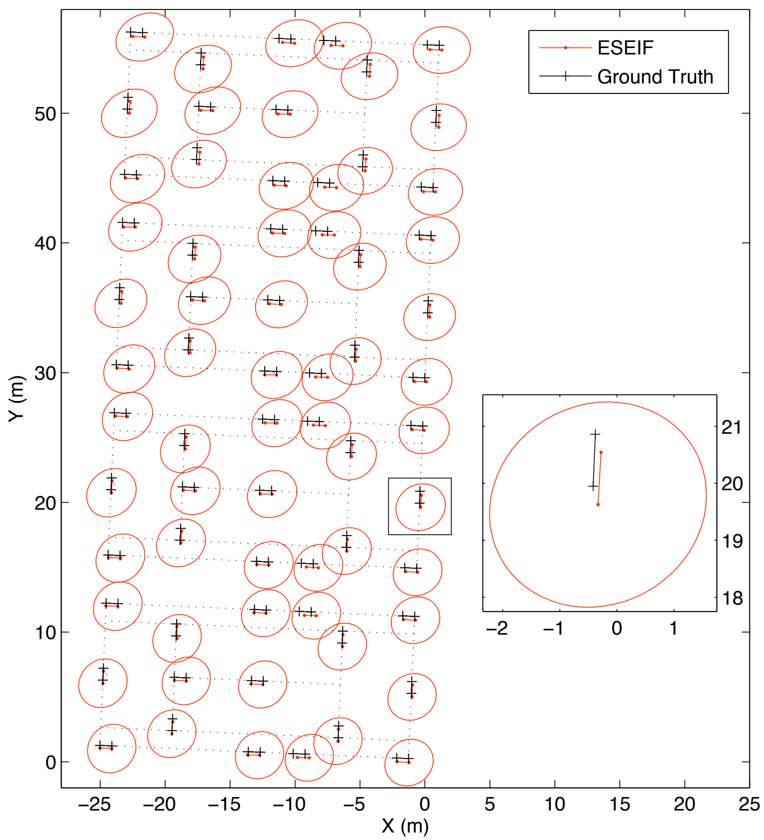

SEIFS

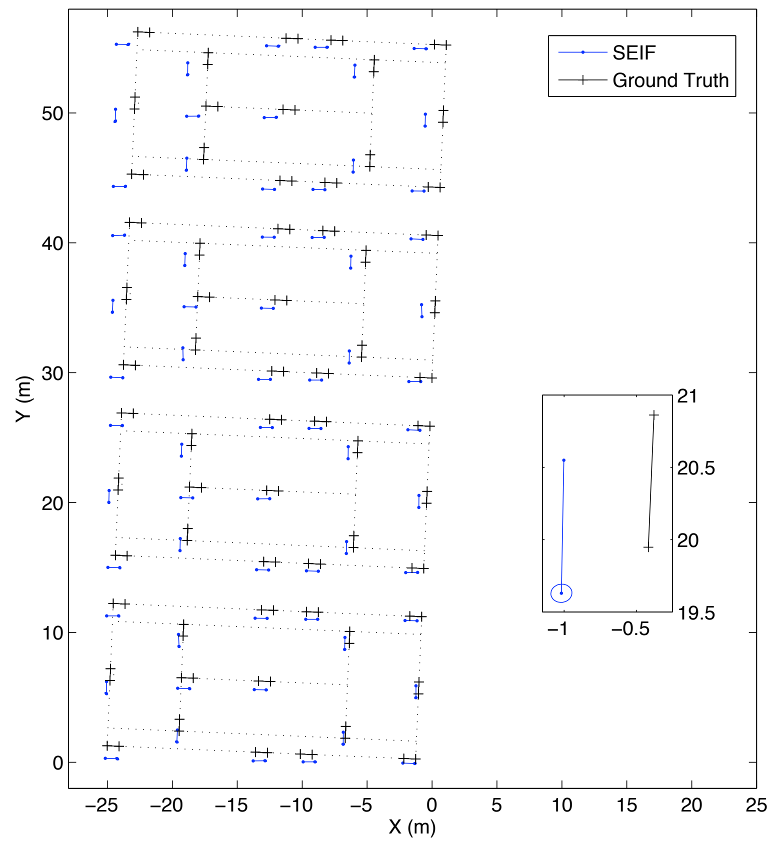

ESEIF estimates are conservative yet nearly identical to EKF

SEIF estimates are significantly over-confident

Nonlinear dataset 2: Sydney Victoria Park (Data courtesy Eduardo Nebot)

Results: Victoria Park Dataset

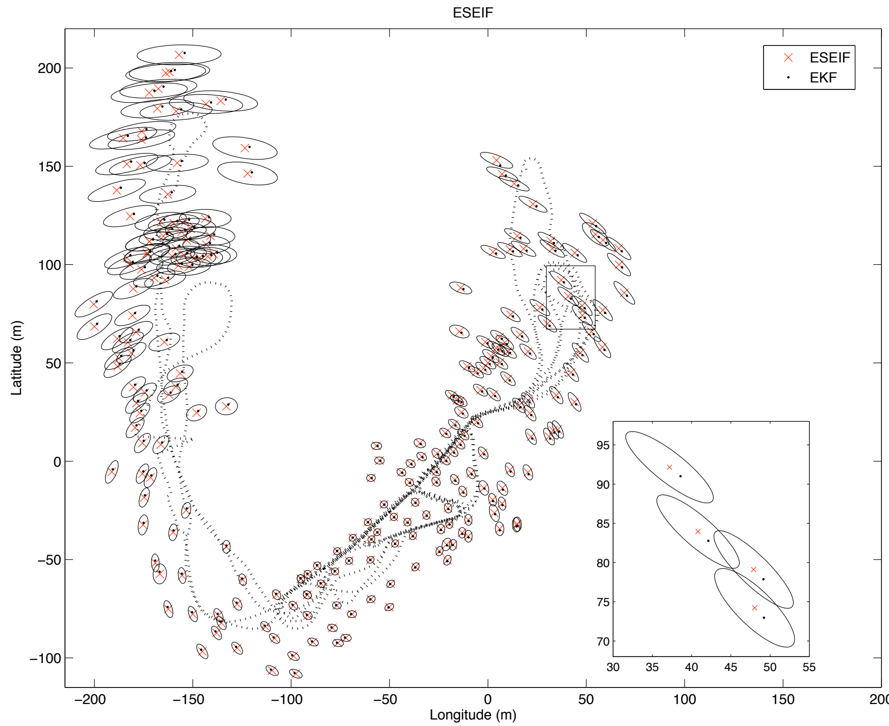

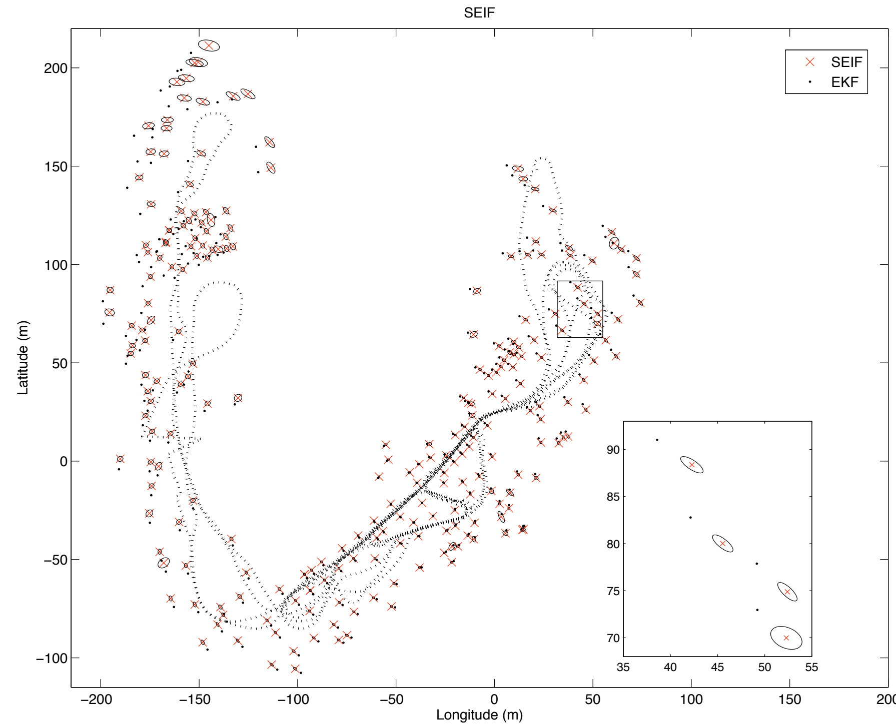

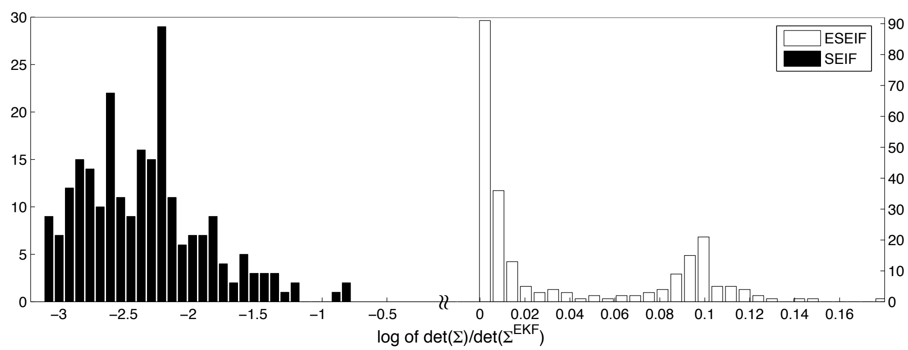

Histogram of feature covariance determinants relative to EKF estimates
> 0 Conservative   < 0   Over-confident

ESEIF estimates are conservative yet nearly identical to EKF

SEIF estimates are significantly over-confident

Roadmap

I Introduction

II Sparsification in SEIFS

III SEIFS

IV Results

V Conclusions

Summary

- New alternative to the original Sparse Extended Information Filter (SEIF) formulation (Thrun et al, WAFR 2002)
- Avoids approximation to enforce sparseness
- Key idea: periodically marginalize out the vehicle pose and relocate the vehicle into the map
- Prevents fill-in of the inverse covariance matrix for terms relating active and in-active features
- Exact sparseness is achieved, ensuring consistency for the Linear Gaussian case
- Experimental analysis demonstrates good performance on two real-world nonlinear data sets

Compare and Contrast with D-SLAM:

- Both papers enforce *exact* sparsity in the SLAM Information Matrix
- Our work is much closer to the original SEIF formulation
- Our methods uses all the feature measurements and very nearly all the odometry measurements
- Our results demonstrate close agreement with the full covariance solution (extremely “tight” bound)
- Several other significant differences .... talk to me, Matt, and/or Ryan for discussion off-line

Any Questions?

Challenges for the future:

3D

Persistent autonomy

Limit as t -> infinity

Long-term memory

RobotGoogle

“wget the world”

Backups

Why is Concurrent Mapping and Localization Difficult? (ISRR 1999)

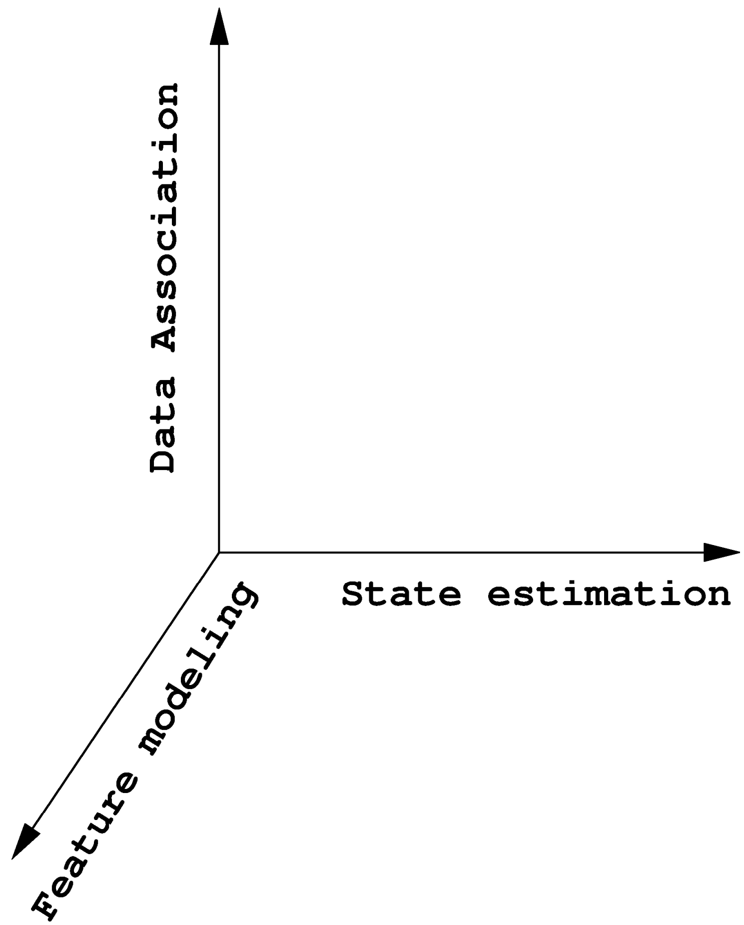

Why is Concurrent Mapping and Localization Difficult? (ISRR 1999)

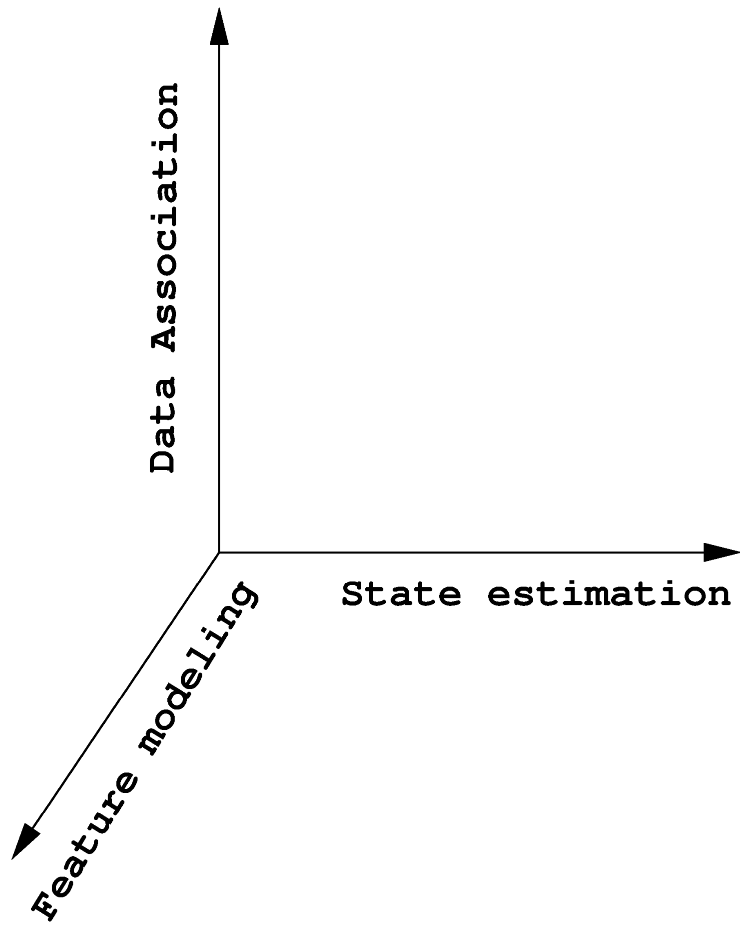

"The key scientific and technological issue in robotics is that of coping with uncertainty ... In fact, the uncertainty is such that one of the most challenging activities for a mobile robot is simply going from point A to point B."

Tomas Lozano-Perez, 1990

Consequences for EKF and EIF

	Covariance Form	Information Form
Marg.	$\boldsymbol{\mu}' = \boldsymbol{\mu}_\alpha$ $\Sigma' = \Sigma_{\alpha\alpha}$	$\boldsymbol{\eta}' = \boldsymbol{\eta}_\alpha - \Lambda_{\alpha\beta}\Lambda_{\beta\beta}^{-1}\boldsymbol{\eta}_\beta$ $\Lambda' = \Lambda_{\alpha\alpha} - \Lambda_{\alpha\beta}\Lambda_{\beta\beta}^{-1}\Lambda_{\beta\alpha}$
Cond.	$\boldsymbol{\mu}' = \boldsymbol{\mu}_\alpha + \Sigma_{\alpha\beta}\Sigma_{\beta\beta}^{-1}(\boldsymbol{\beta} - \boldsymbol{\mu}_\beta)$ $\Sigma' = \Sigma_{\alpha\alpha} - \Sigma_{\alpha\beta}\Sigma_{\beta\beta}^{-1}\Sigma_{\beta\alpha}$	$\boldsymbol{\eta}' = \boldsymbol{\eta}_\alpha - \Lambda_{\alpha\beta}\boldsymbol{\beta}$ $\Lambda' = \Lambda_{\alpha\alpha}$

II. Time projection step: state augmentation + marginalization

$$p\left(\mathbf{x}_{t+1}, \mathbf{M}\right) = \int p\left(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{M}\right) d\mathbf{x}_t \qquad \boldsymbol{\alpha} = \{\mathbf{x}_{t+1}, \mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_3, \mathbf{m}_4, \mathbf{m}_5\} \quad \boldsymbol{\beta} = \mathbf{x}_t$$

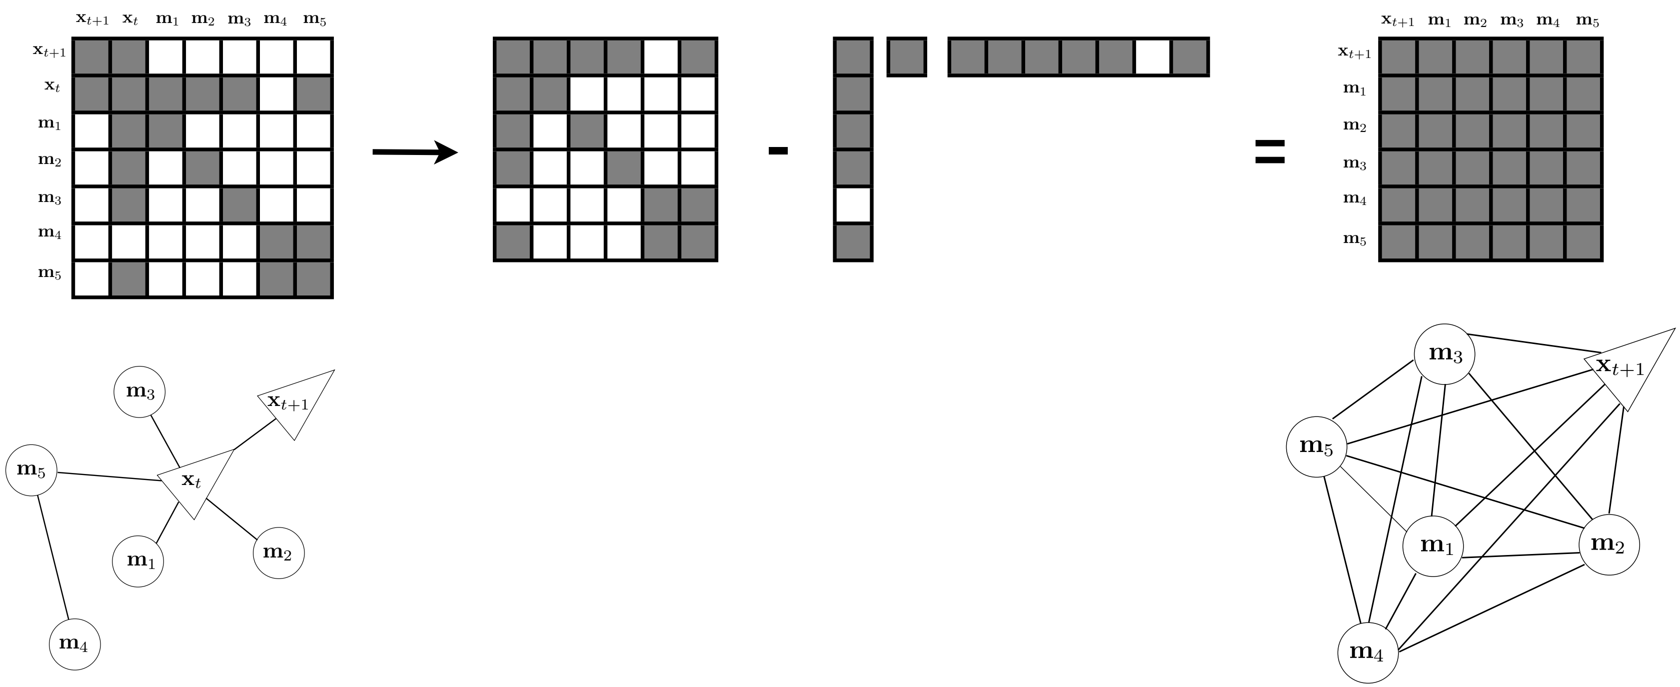

Robot's conditional independence from passive features

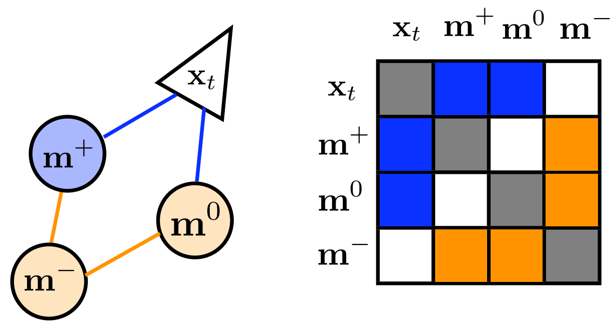

$$p\left(\mathbf{x}_t, \mathbf{m}^- \mid \mathbf{m}^0, \mathbf{m}^+\right) = \mathcal{N}^{-1}\left(S^\top_{x_t,m^-}\boldsymbol{\xi}_t; \boldsymbol{\eta}_{x_t,m^-|m^0,m^+}, \Lambda_{x_t,m^-|m^0,m^+}\right)$$

$$\Lambda_{x_t,m^-|m^0,m^+} = S^\top_{x_t,m^-}\Lambda_t S_{x_t,m^-}$$ ← extract $\mathbf{x}_t$ $\mathbf{m}^-$ sub-block

$$\boldsymbol{\eta}_{x_t,m^-|m^0,m^+} = S^\top_{x_t,m^-}\boldsymbol{\eta}_t - \left(S^\top_{x_t,m^-}\Lambda_t S_{m^0,m^+}\right)S^\top_{m^0,m^+}\boldsymbol{\xi}_t$$

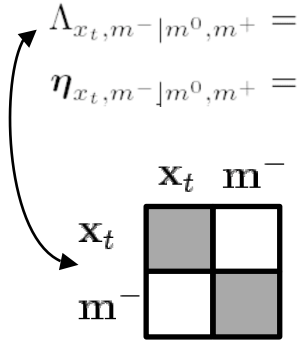

The conditional information matrix is block-diagonal and the inverse of a block-diagonal matrix is also block-diagonal. This implies that the robot and passive features are uncorrelated which for a Gaussian RV implies independence. QED

Sparsification in SEIFs

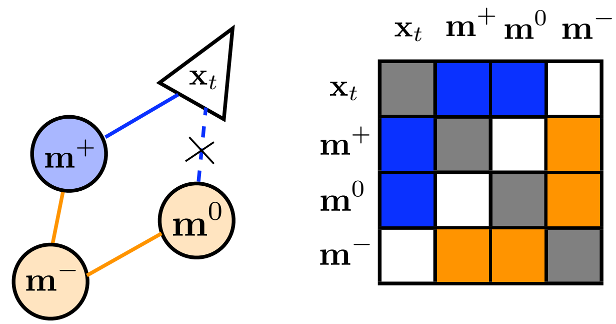

Bayes rule: $$p(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) = p\left(\mathbf{x}_t \mid \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right) p\left(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right)$$
$$= p\left(\mathbf{x}_t \mid \mathbf{m}^0, \mathbf{m}^+, \underbrace{\mathbf{m}^- = \boldsymbol{\alpha}}_{\text{C.I.}}\right) p\left(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right)$$

SEIF's rule deactivates link by forcing conditional independence to feature we want to deactivate

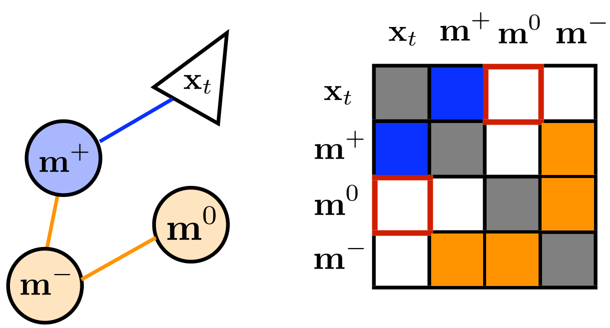

$$\tilde{p}_{\text{SEIFs}}(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) = p\left(\mathbf{x}_t \mid \mathbf{m}^+, \mathbf{m}^- = \boldsymbol{\alpha}\right) p\left(\mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-\right)$$

Sparsification in SEIFs

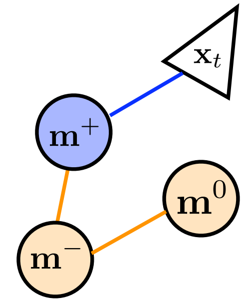

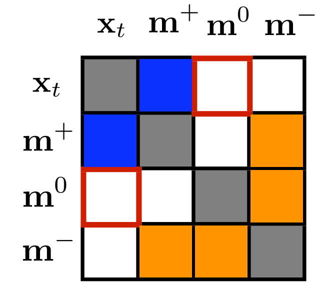

$$\tilde{p}_{\text{SEIFs}}(\mathbf{x}_t, \mathbf{m}^0, \mathbf{m}^+, \mathbf{m}^-) = \mathcal{N}^{-1}(\tilde{\boldsymbol{\eta}}_t, \tilde{\Lambda}_t)$$

$$\begin{aligned}\tilde{\Lambda}_t &= S_{x_t,m^+} \Lambda_B S_{x_t,m^+}^\top \\ &\quad - S_{m^+} \Lambda_C S_{m^+}^\top + S_{m^0,m^+,m^-} \Lambda_D S_{m^0,m^+,m^-}^\top \\ \tilde{\boldsymbol{\eta}}_t &= S_{x_t,m^+} \boldsymbol{\eta}_B - S_{m^+} \boldsymbol{\eta}_C + S_{m^0,m^+,m^-} \boldsymbol{\eta}_D\end{aligned}$$

$$\begin{aligned}\boldsymbol{\eta}_\alpha &= \Sigma_t S_{m^-} \boldsymbol{\alpha} \\ \Lambda_B &= S_{x_t,m^+}^\top \left(\mathrm{I} - \Lambda_t S_{m^0} \left(S_{m^0}^\top \Lambda_t S_{m^0} \right)^{-1} S_{m^0}^\top \right) \Lambda_t S_{x_t,m^+} \\ \boldsymbol{\eta}_B &= S_{x_t,m^+}^\top \left(\mathrm{I} - \Lambda_t S_{m^0} \left(S_{m^0}^\top \Lambda_t S_{m^0} \right)^{-1} S_{m^0}^\top \right) (\boldsymbol{\eta}_t - \boldsymbol{\eta}_\alpha) \\ \Lambda_C &= S_{m^+}^\top \left(\mathrm{I} - \Lambda_t S_{x_t,m^0} \left(S_{x_t,m^0}^\top \Lambda_t S_{x_t,m^0} \right)^{-1} S_{x_t,m^0}^\top \right) \Lambda_t S_{m^+} \\ \boldsymbol{\eta}_C &= S_{m^+}^\top \left(\mathrm{I} - \Lambda_t S_{x_t,m^0} \left(S_{x_t,m^0}^\top \Lambda_t S_{x_t,m^0} \right)^{-1} S_{x_t,m^0}^\top \right) (\boldsymbol{\eta}_t - \boldsymbol{\eta}_\alpha) \\ \Lambda_D &= S_{m^0,m^+,m^-}^\top \left(\mathrm{I} - \Lambda_t S_{x_t} \left(S_{x_t}^\top \Lambda_t S_{x_t} \right)^{-1} S_{x_t}^\top \right) \Lambda_t S_{m^0,m^+,m^-} \\ \boldsymbol{\eta}_D &= S_{m^0,m^+,m^-}^\top \left(\mathrm{I} - \Lambda_t S_{x_t} \left(S_{x_t}^\top \Lambda_t S_{x_t} \right)^{-1} S_{x_t}^\top \right) \boldsymbol{\eta}_t\end{aligned}$$

only requires matrix inversion on the order of the number of links we are breaking

Links are weak, yet Markov network is fully connected

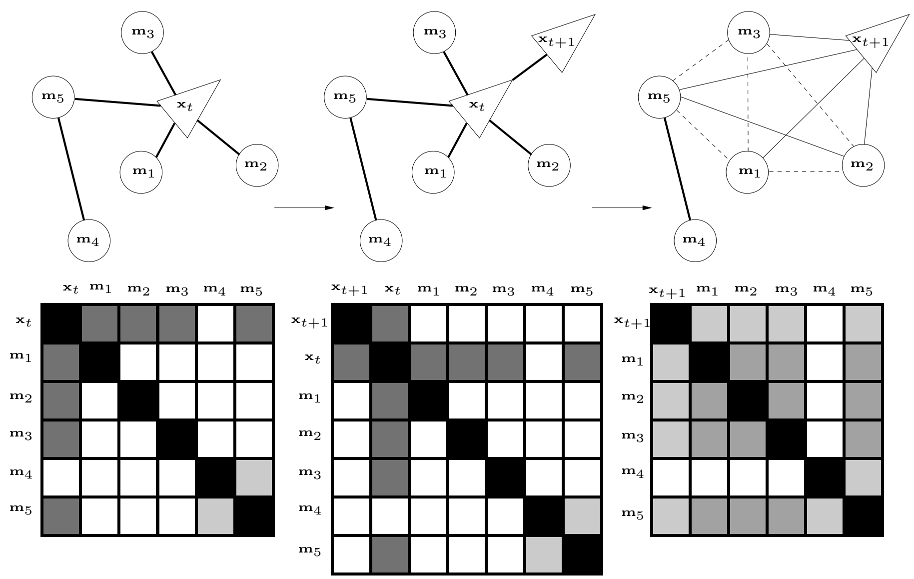

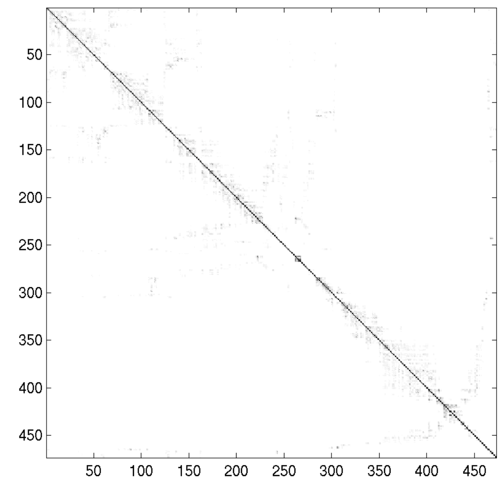

Victoria Park information matrix

Active features naturally lead to fill-in → To control matrix population, deactivate features

PROBLEM: Active features never become passive

Controlling density: limit number of active features

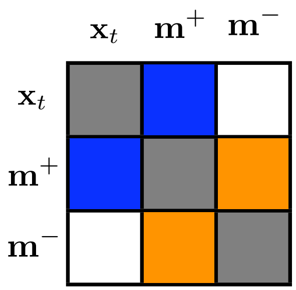

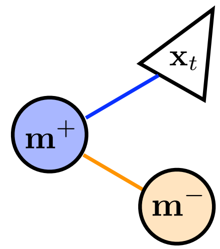

$\mathbf{m}^-$ active features

$\mathbf{m}^+$ passive features

Controlling density: limit number of active features

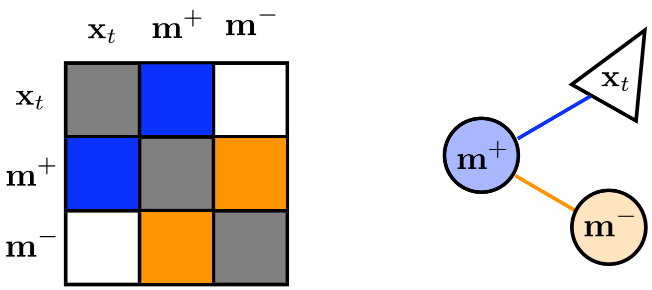

$\mathbf{m}^+$ active features

$\mathbf{m}^-$ passive features

Pacify active landmarks by breaking weak links

$\mathbf{m}^0$ active features to be made passive

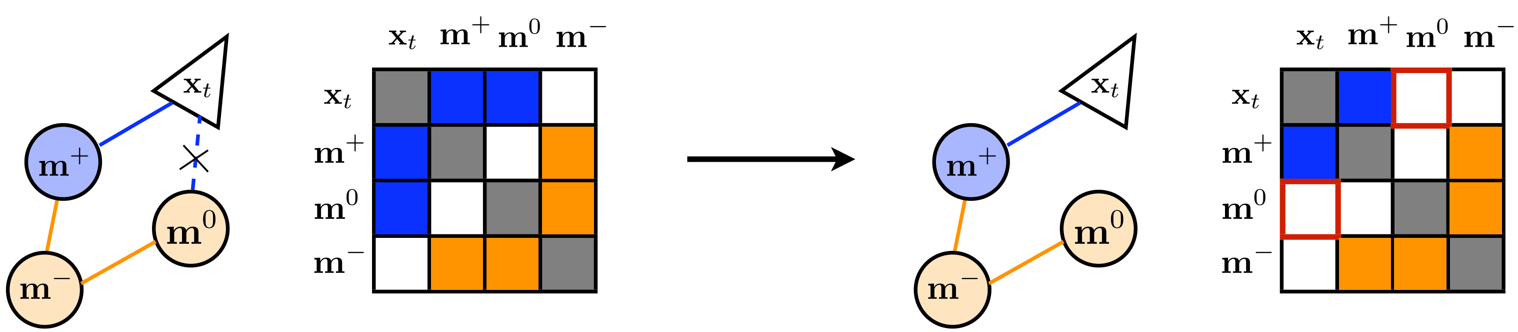

Deactivation: imposing conditional independence

$\mathbf{m}^+$	active features
$\mathbf{m}^0$	active features to be made passive
$\mathbf{m}^-$	passive features

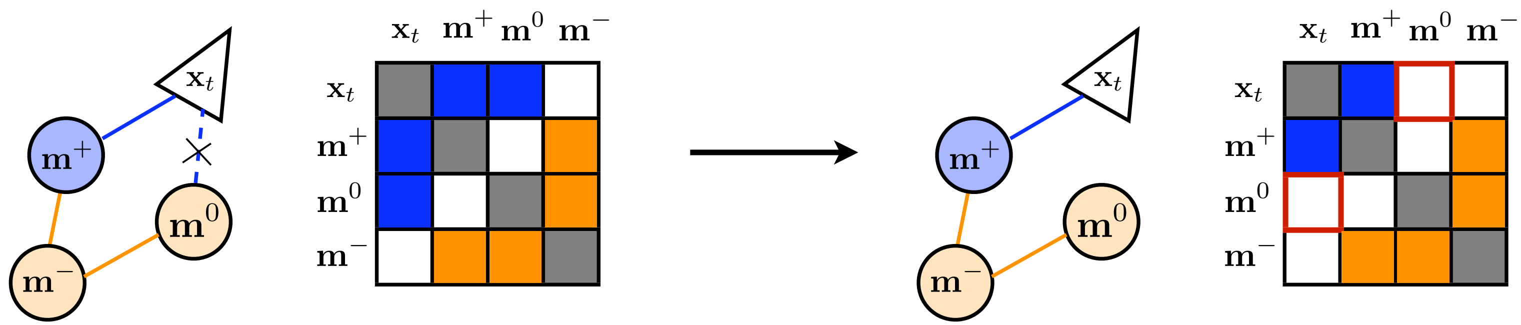

$$p(\boldsymbol{\xi}_t) = \underline{p(\mathbf{x}_t|\mathbf{m}^+, \mathbf{m}^0)}\, p(\mathbf{m}^+, \mathbf{m}^0, \mathbf{m}^-)$$

$$\tilde{p}(\boldsymbol{\xi}_t) = \underline{p(\mathbf{x}_t|\mathbf{m}^+)}\, p(\mathbf{m}^+, \mathbf{m}^0, \mathbf{m}^-)$$

How do we force $p(\mathbf{x}_t|\mathbf{m}^+, \mathbf{m}^0) \longrightarrow p(\mathbf{x}_t|\mathbf{m}^+)$?

How we sparsify is nontrivial!