Towards Information-Theoretic Decision Making in a *Conservative* Information Space

Vadim Indelman

TECHNION
Israel Institute of Technology

American Control Conference (ACC), July 2015

Introduction

- Decision making under uncertainty - fundamental problem in autonomous systems and artificial intelligence
- Examples
  - Informative planning, active sensing
  - Sensor selection, sensor deployment
  - Belief space planning
  - Active simultaneous localization and mapping (SLAM)
  - Multi-agent informative planning and active SLAM
  - Target tracking

Introduction

- Information-theoretic decision making
  - **Objective**: find action(s) that minimizes an information-theoretic objective function (e.g. entropy)
  - Extensively investigated, e.g., in the context of sensor selection

- Decision making over **high-dimensional** state spaces is expensive!

  State vector: $\mathbf{x} \in \mathbb{R}^n$

  Covariance matrix: $\Sigma \doteq \mathbb{E}\left[\left(\mathbf{x} - \mathbb{E}\left[\mathbf{x}\right] \right) \left(\mathbf{x} - \mathbb{E}\left[\mathbf{x}\right] \right)^T \right] \in \mathbb{R}^{n \times n}$

- Evaluating impact of a candidate action typically involves determinant calculation - $O(n^3)$ in the general case

Introduction – Motivating Example I

- **Belief Space Planning, Active SLAM**
  - Robot operates in unknown/uncertain environments
  - Concurrently infers its own state and the observed environment

	Recursive	Smoothing
State:	$\mathbf{x}_k \doteq \begin{bmatrix} x_k^T & l_1^T & \cdots & l_m^T \end{bmatrix}^T$	$\mathbf{x}_{0:k} \doteq \begin{bmatrix} x_0^T & \cdots & x_k^T & l_1^T & \cdots & l_m^T \end{bmatrix}^T$
pdf:	$p\left(\mathbf{x}_k \middle\vert z_{0:k}, u_{0:k-1}\right)$	$p\left(\mathbf{x}_{0:k} \middle\vert z_{0:k}, u_{0:k-1}\right)$

  - How to autonomously determine future action(s)?
  - Involves reasoning, for different candidate actions, about belief evolution

$$p\left(\mathbf{x}_{k+L} \middle\vert z_{0:k}, u_{0:k-1}, u_{k:k+L-1}, z_{k+1:k+L}\right)$$

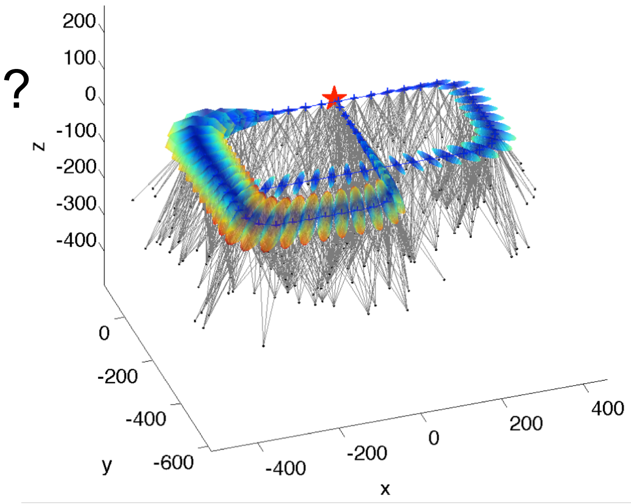

Introduction – Motivating Example II

- **Sensor Deployment**
  - **Objective**: deploy k sensors in an $N \times N$ area
  - e.g., provide localization, monitor spatial-temporal field

Prior uncertainty field:

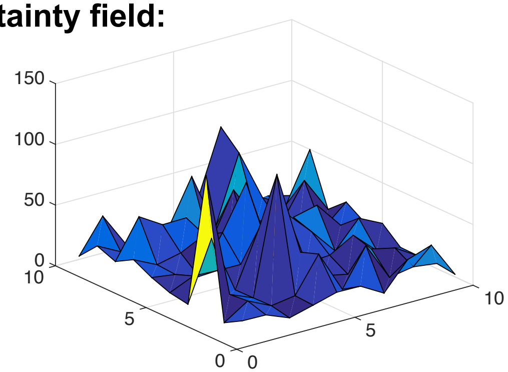

A priori joint covariance (with correlations between cells)

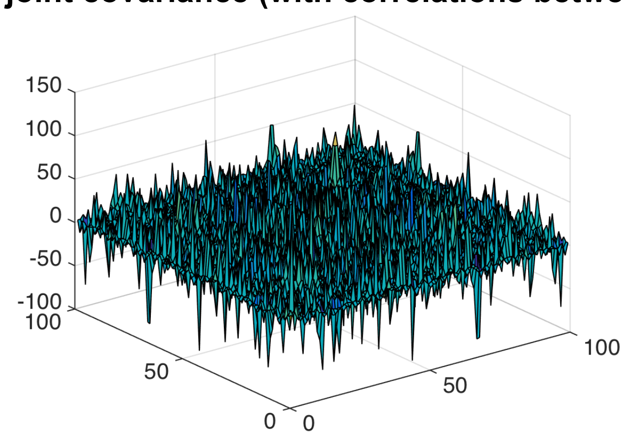

After deploying 8 sensors:

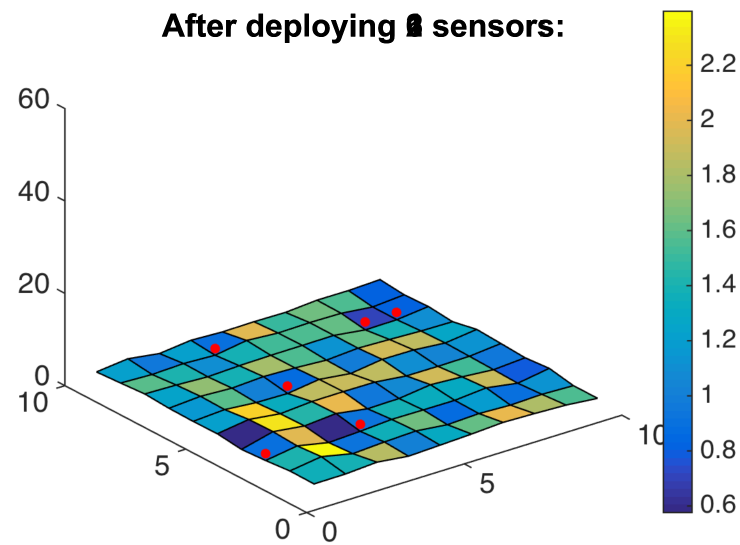

Introduction

- More generally, decision making over multiple look-ahead steps
  - A partially observable Markov decision process (POMDP), NP-hard
  - Different sub-optimal approaches exist (greedy, sampling, …)
- **This work:**
  - Resort to **conservative information fusion** techniques for **information-theoretic decision making**
- **Conservative information fusion** approaches
  - Allow to fuse information from multiple correlated sources, **without** knowing the correlation
  - Guarantee **consistent** estimation
  - Pioneered by Julier & Uhlmann [ACC 1997]: **Covariance intersection**

Introduction

- More generally, decision making over multiple look-ahead steps
  - A partially observable Markov decision process (POMDP), NP-hard
  - Different sub-optimal approaches exist (greedy, sampling, …)

- **This work:**
  - Resort to **conservative information fusion** techniques for **information-theoretic decision making**
  - **Motivation**: these techniques allow correlation terms to be unknown!
  - **Key idea:**
    - Reduce computational complexity by (appropriately) dropping correlations
    - Extreme case: drop all correlations; computational complexity becomes

$$O(n^3) \longrightarrow O(n)$$

  - **Do we get the same performance??**

Problem Formulation

- Probability distribution function (pdf) at time t_k : $p\left(x_k | z_{0:k}, u_{0:k-1}\right)$
- Transition/motion model $p\left(x_{k+1} | x_k, u_k\right)$
- Observation model $p\left(z_k | x_k\right)$

- Given control u_k and new observation(s) z_{k+1}, pdf becomes

$$p\left(x_{k+1} | z_{0:k+1}, u_{0:k}\right) = \eta p\left(z_k | x_k\right) \cdot \int p\left(x_k | z_{0:k}, u_{0:k-1}\right) p\left(x_{k+1} | x_k, u_k\right) dx_k$$

- Entropy: $\mathcal{H}\left(p\left(x\right)\right) = -\mathbb{E}\left[\log p\left(x\right)\right] = -\int p\left(x\right) \log p\left(x\right) dx$

- **Information-theoretic objective function** (single look-ahead step):

$$J\left(u_k\right) = \mathbb{E}_{z_{k+1}}\left[\mathcal{H}\left(p\left(x_{k+1} | z_{0:k+1}, u_{0:k}\right)\right)\right]$$

- **Optimal control:** $u_k^\star = \arg\min_{u_k} J\left(u_k\right).$

Problem Formulation

- Assumptions:
  - Gaussian distributions $p\left(x_k|z_{0:k}, u_{0:k-1}\right) = N\left(\mu_k, I_k^{-1}\right)$
  - Deterministic control (for now) $z_i = h_i\left(x_i\right) + v_i \quad , \quad v_i \sim N\left(0, \Sigma_{vi}\right)$
- Entropy becomes $\mathcal{H}\left(p\left(x_{k+1}|z_{0:k+1}, u_{0:k}\right)\right) = -\frac{1}{2}\log\left[\left(2\pi e\right)^n \left|I_{k+1}^{+}\right|\right]$
- A posteriori information matrix: $I_{k+1}^{+} = I_k + H^T \Sigma_v^{-1} \underbrace{H}_{\text{Jacobian}}$
- Best action = highest information gain
- Impact evaluation for a candidate action is in the general case: $O(n^3)$

Conservative Information Space

- Conservative approximation of a pdf – sufficient conditions **[Bailey et al. 2012 Fusion]**:
  - Entropy: $\mathcal{H}\left(p\left(x\right)\right) \leq \mathcal{H}\left(p_c\left(x\right)\right)$
  - Order preserving (same shape):

$$\forall x_i, x_j \quad p_c(x = x_i) \leq p_c(x = x_j) \text{ iff } p(x = x_i) \leq p(x = x_j)$$

- Gaussian case:

$$|I_c| \leq |I|$$

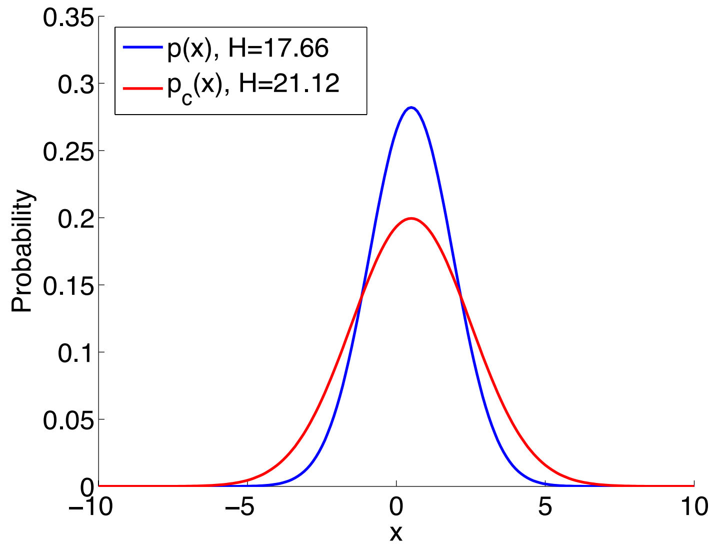

Concept

Decision Making Over a Conservative Information Space - **1D Case**

- Consider some two actions **a** and **b** with measurement models

$$z_a = h_a(x) + v_a \qquad z_b = h_b(x) + v_b$$

- **Theorem** - for the 1D case:

$$\boxed{\left|I^{a+}\right| \leq \left|I^{b+}\right| \text{ iff } \left|I_c^{a+}\right| \leq \left|I_c^{b+}\right|}$$

– where the a posteriori information matrices are calculated using

	Action **a**	Action **b**
original information matrix:	$I^{a+} = I + H_a^T \Sigma_v^{-1} H_a$	$I^{b+} = I + H_b^T \Sigma_v^{-1} H_b$
conservative information matrix:	$I_c^{a+} = I_c + H_a^T \Sigma_v^{-1} H_a$	$I_c^{b+} = I_c + H_b^T \Sigma_v^{-1} H_b$

Concept

Decision Making Over a Conservative Information Space - **1D Case**

- Consider some two actions **a** and **b** with measurement models

$$z_a = h_a(x) + v_a \qquad z_b = h_b(x) + v_b$$

- **Theorem** - for the 1D case:

$$\boxed{\left|I^{a+}\right| \leq \left|I^{b+}\right| \text{ iff } \left|I_c^{a+}\right| \leq \left|I_c^{b+}\right|}$$

- In words:

  the impact of any two candidate actions has the same trend regardless if it is calculated based on the original or conservative information space

- Therefore: decision making can be done considering a conservative information space

Basic Example – 1D Case

$$\left|I^{a+}\right| \leq \left|I^{b+}\right| \text{ iff } \left|I_c^{a+}\right| \leq \left|I_c^{b+}\right|$$

Entropy values are shown in legend

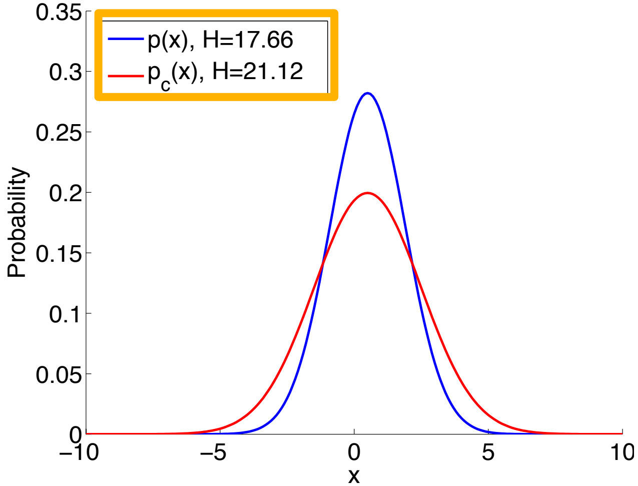

I, I_c

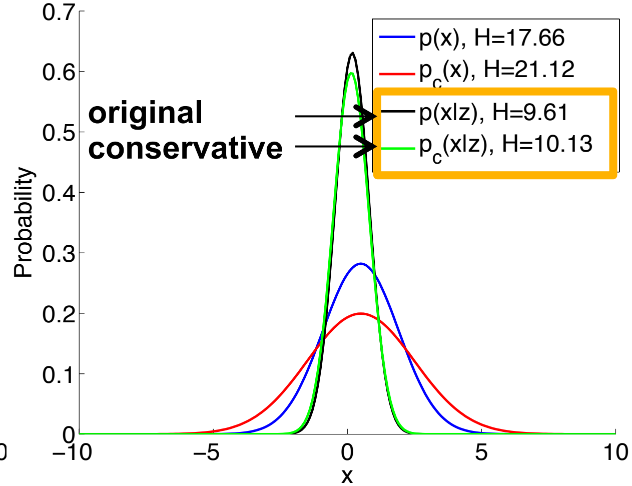

Action a

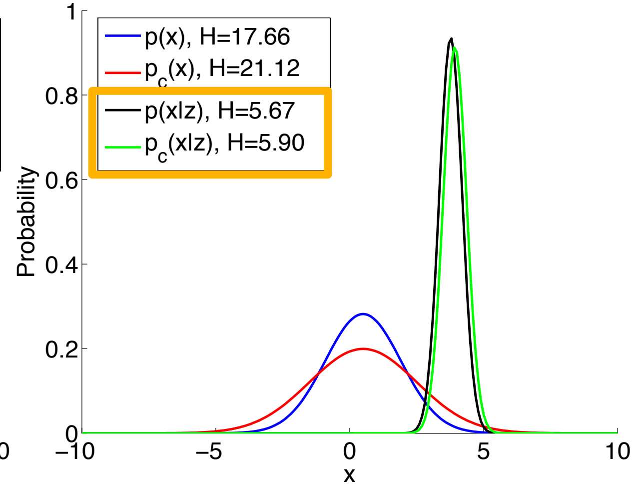

Action b

$$z_a = h_a(x) + v_a \qquad z_b = h_b(x) + v_b$$

$$\Sigma_v = 0.5^2 \qquad \Sigma_v = 0.2^2$$

High Dimensional State Space

Recall: $\mathcal{H}\left(p\left(x_{k+1}|z_{0:k+1}, u_{0:k}\right)\right) = -\frac{1}{2}\log\left[(2\pi e)^n \left|I^+_{k+1}\right|\right]$

- Is the concept valid also for high-dimensional spaces?
- Why is it interesting?
  - Consider an information matrix $I \in \mathbb{R}^{n\times n}$
  - Calculating $|I|$ is often expensive ($O(n^3)$, in the general case)
  - **Instead**
    - Calculate a conservative **sparse** information matrix I_c
    - Evaluating $|I_c|$ can be done very efficiently
    - If concept applies, same performance is guaranteed!
- Next: Going to the extreme – appropriately **drop all correlation terms**
  - I_c is diagonal
  - Complexity is reduced to $O(n)$

“Decoupled” Conservative PDF

- Definition:

$$p_c\left(X\right) \doteq \eta \prod_i p^{w_i}\left(x_i\right) \qquad \forall x_i \in X \qquad \sum_i w_i = 1$$

- 2D case:

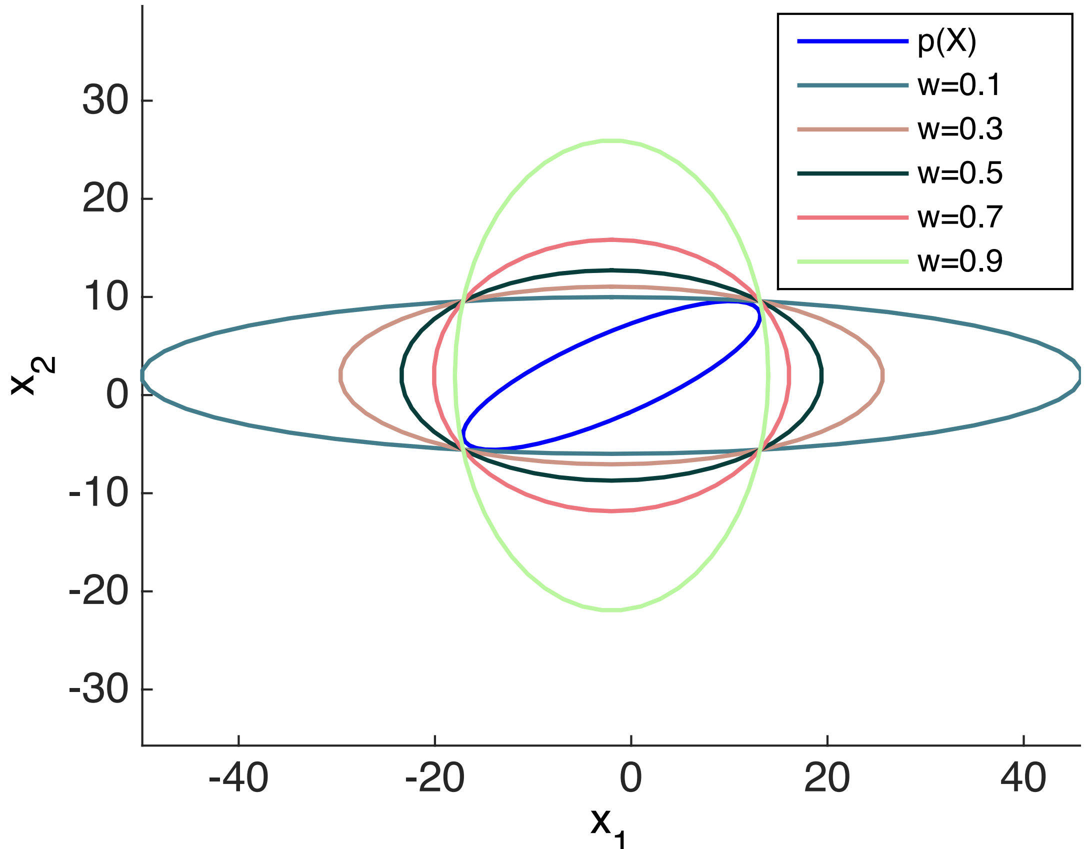

"Decoupled" Conservative PDF

- Definition:

$$p_c\left(X\right) \doteq \eta \prod_i p^{w_i}\left(x_i\right) \qquad \forall x_i \in X \qquad \sum_i w_i = 1$$

- Example - $X \in \mathbb{R}^4$:

$$p\left(X\right) = p\left(x_0\right) p\left(x_1|x_0\right) p\left(x_2|x_1,x_0\right) p\left(x_3|x_2,x_1\right) \Longrightarrow p_c\left(X\right) = p_c\left(x_0\right) p_c\left(x_1\right) p_c\left(x_2\right) p_c\left(x_3\right).$$

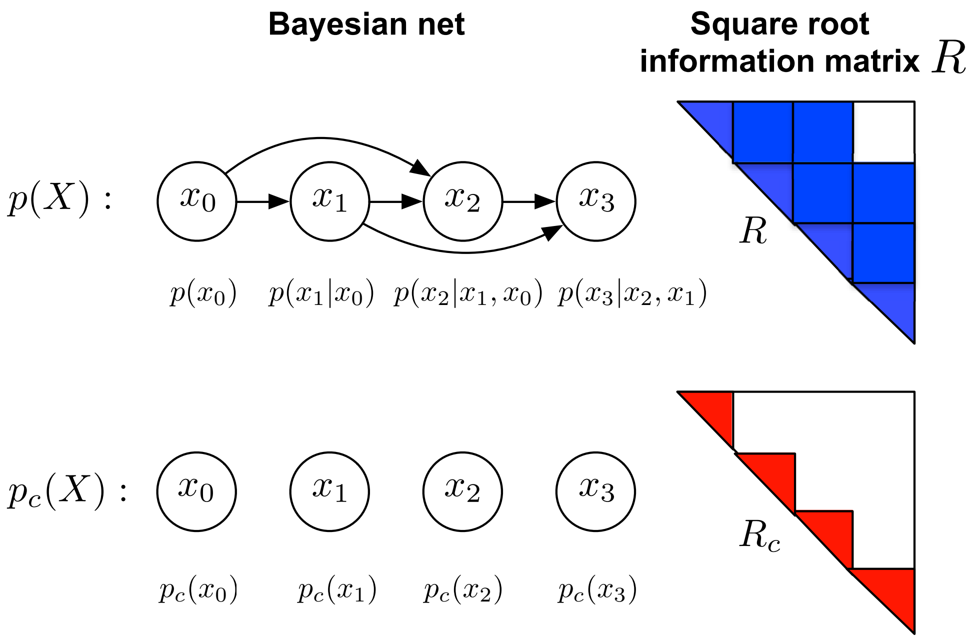

$$I = R^T R$$

High Dimensional State Space

- Is the concept valid also for high-dimensional spaces?
  - In particular, in conjunction with the **decoupled** conservative pdf

- Valid (at least) in the following cases:
  - Observation models include the same arbitrary states, possibly with different measurement noise covariance

$$z_i = h\left(X'\right) + v_i \qquad X' \subset X$$

  - Unary observation models, possibly involving different states

$$z_i = h_i\left(x_i\right) + v_i \qquad x_i \in X$$

Example I

  - Binary observation models with the same uncorrelated state

$$z_i = h_i\left(x, x_i\right) + v_i \qquad x, x_i \in X$$

    - Here, x is not correlated with other states

Example II

Example I

$$|I^{a+}| \leq |I^{b+}| \text{ iff } |I_c^{a+}| \leq |I_c^{b+}|$$

- Unary observation models, possibly involving different states $z_i = h_i(x_i) + v_i$

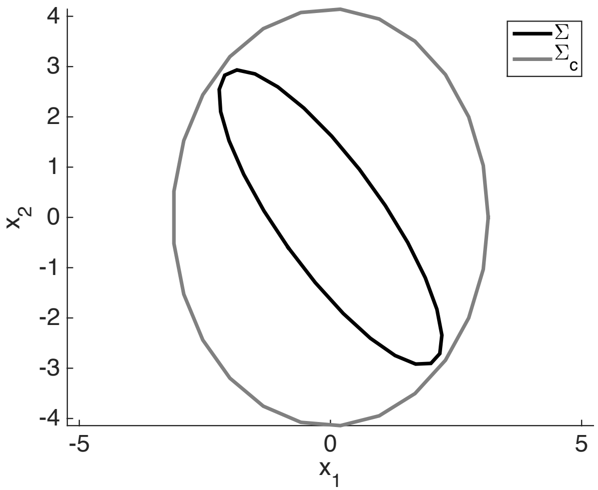

- Original covariance: $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}^T & \Sigma_{22} \end{bmatrix}$

- Conservative covariance: $\Sigma_c = \begin{bmatrix} \Sigma_{c,11} & 0 \\ 0 & \Sigma_{c,22} \end{bmatrix}$

- Consider two actions/sensors:
  - Action **a**: 2nd state is measured
  - Action **b**: 1st state is measured

- Recall - a posteriori information matrix:

$$I^+ = I + H^T \Sigma_v^{-1} H$$

Example I

$$|I^{a+}| \leq |I^{b+}| \text{ iff } |I_c^{a+}| \leq |I_c^{b+}|$$

- Unary observation models, possibly involving different states $z_i = h_i(x_i) + v_i$

Action **a**

Action **b**

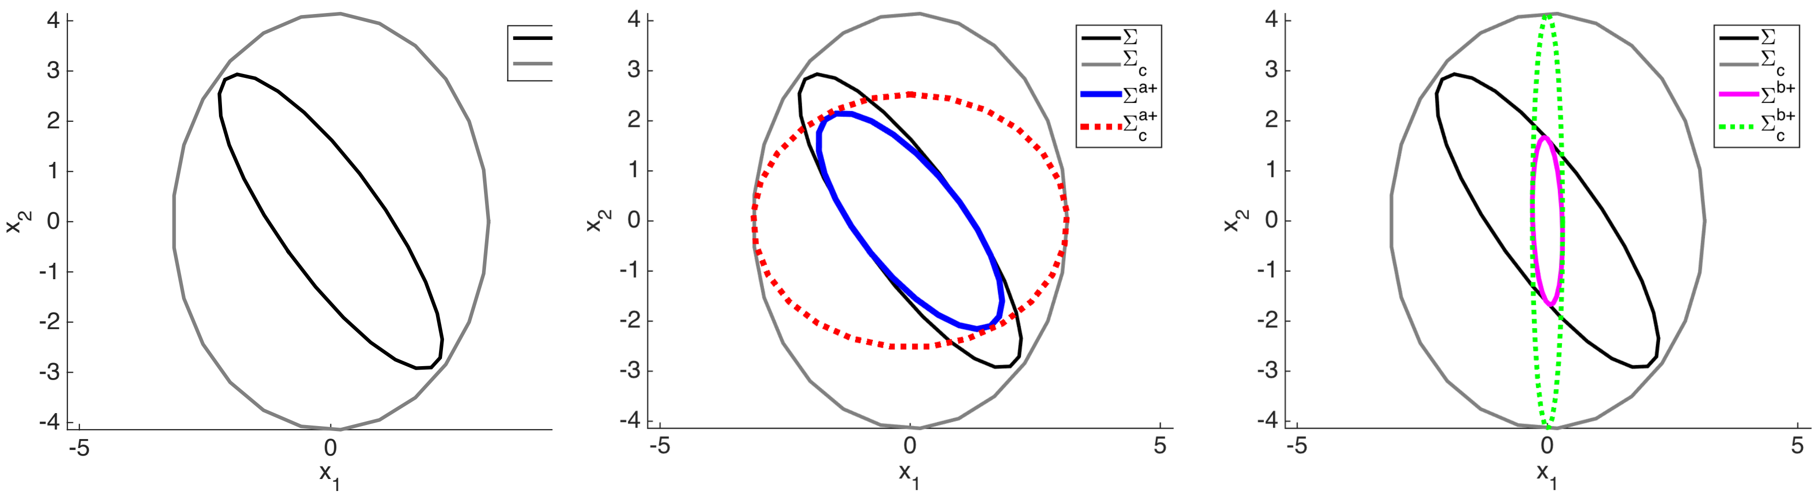

- Consider two actions/sensors:
  - Action **a**: 2nd state is measured
  - Action **b**: 1st state is measured
- Recall - a posteriori information matrix:

$$I^{+} = I + H^{T}\Sigma_v^{-1}H$$

Example I

$$|I^{a+}| \leq |I^{b+}| \ \text{ iff } \ |I_c^{a+}| \leq |I_c^{b+}|$$

Action **a** Action **b**

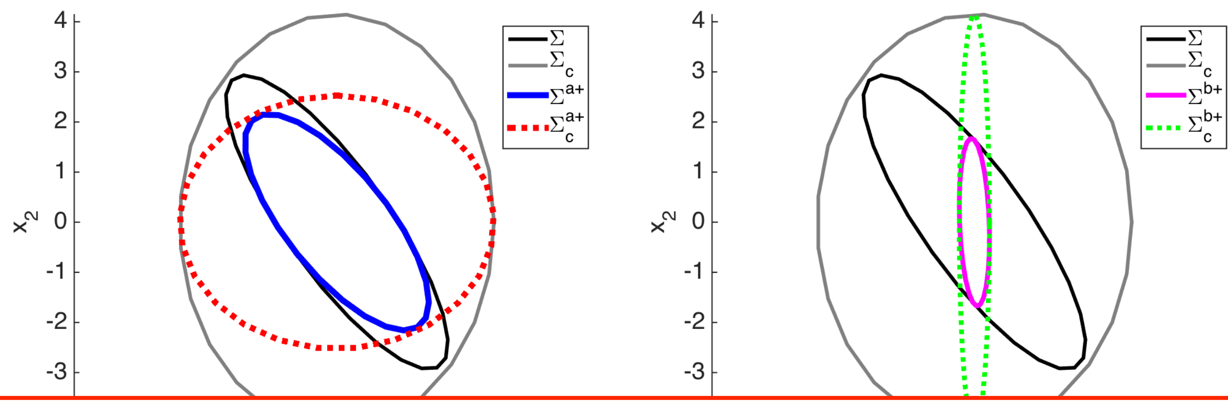

Do **not** need correlations to decide which action is better

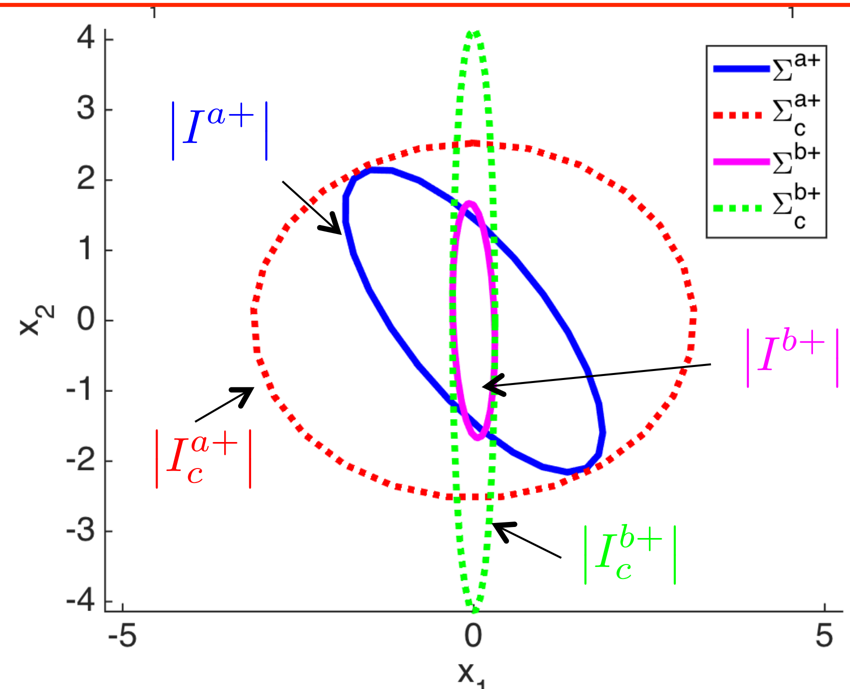

Example II

$$\left|I^{a+}\right| \leq \left|I^{b+}\right| \text{ iff } \left|I_c^{a+}\right| \leq \left|I_c^{b+}\right|$$

- Binary observation models with the same uncorrelated state

$$z_i = h_i\left(x, x_i\right) + v_i$$

- Aerial visual SLAM scenario

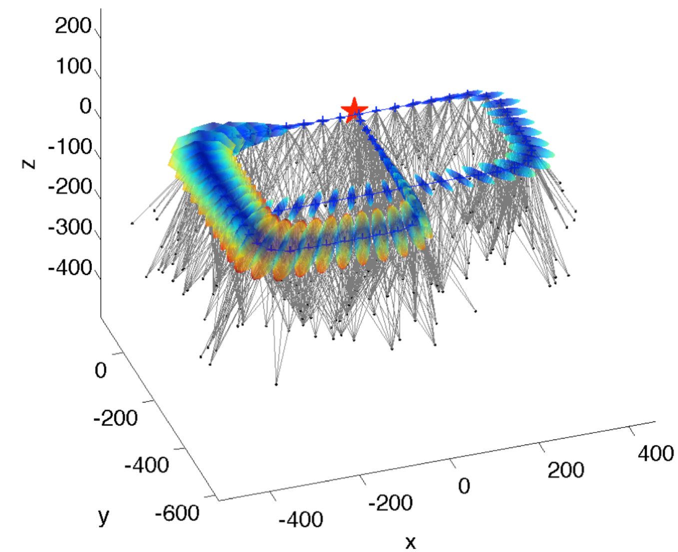

- Objective - each time a new image is received:
  - Decide what image observations to use
  - Identify most informative visual observations

- Remarks:
  - New camera pose x remains uncorrelated as long as no image observations have been incorporated
  - Note: can still add a prior $p(x)$

Example II

$$|I^{a+}| \leq |I^{b+}| \;\text{ iff }\; |I_c^{a+}| \leq |I_c^{b+}|$$

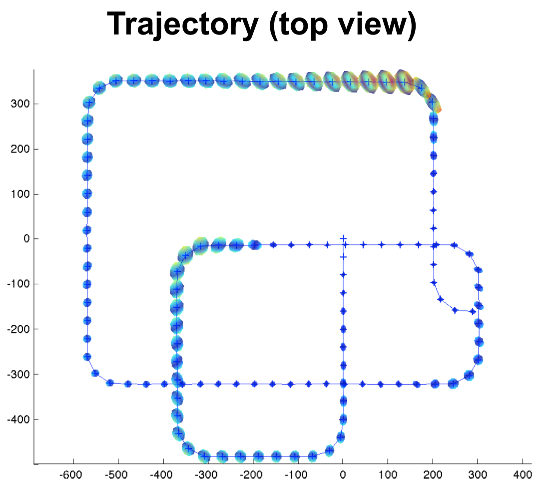

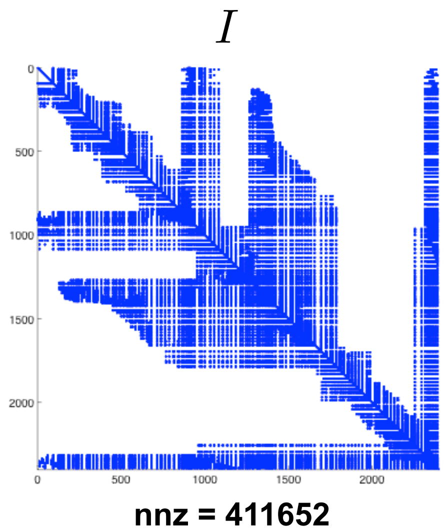

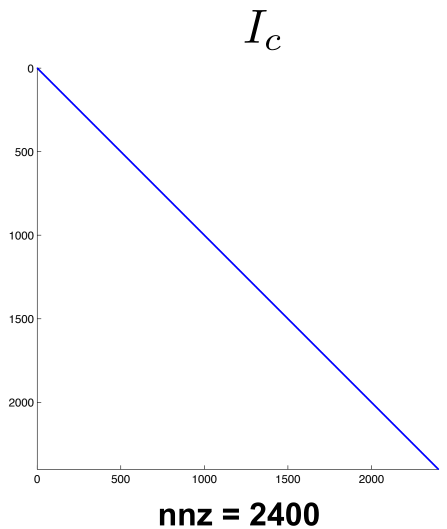

- Same trend!

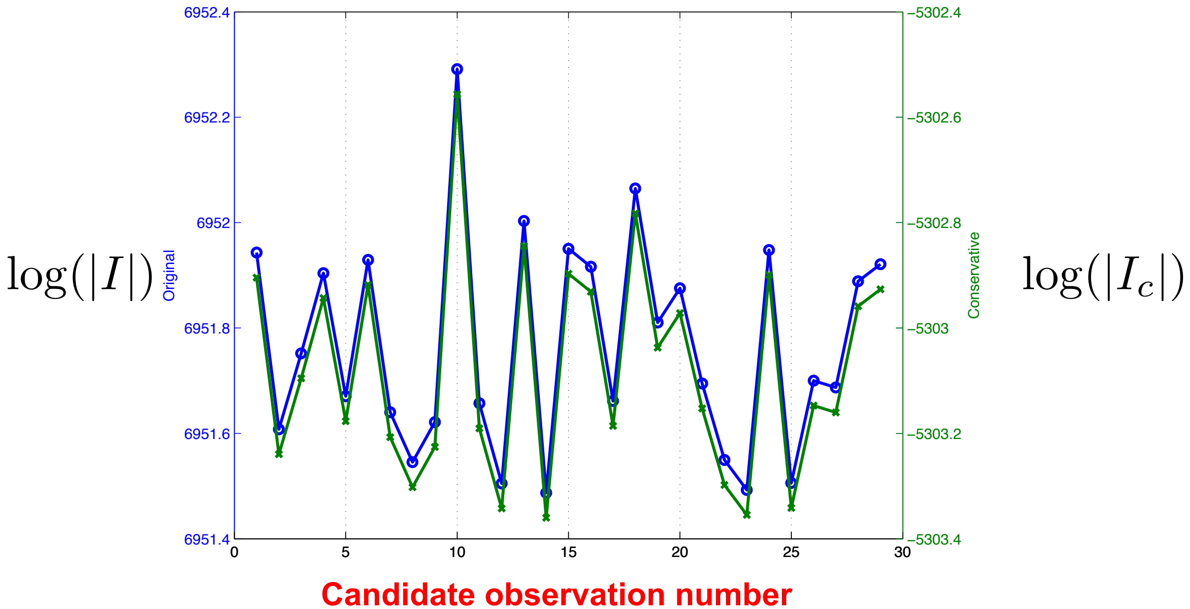

Conclusions

- **Decision making in the conservative information space**
  - Use a **sparse conservative** information matrix to greatly reduce computational complexity
  - In particular:
    - **Decoupled** conservative pdf – diagonal information matrix
    - Computational complexity is reduced by 2 orders of magnitude
    - Concept was proved to yield **the same performance** (decisions) in several scenarios of interest
- Multiple extensions to be investigated

TECHNION
Israel Institute
of Technology