TECHNION
Israel Institute
of Technology

involve-MI: Informative Planning with High-Dimensional Non-Parametric Beliefs

Gilad Rotman

under the supervision of Assoc. Prof. Vadim Indelman

M.Sc. Seminar, January 2022

ANPL | Autonomous Navigation and Perception Lab

Introduction

Autonomous Systems

What is an autonomous system?

- **System:** can span from a small drone to a city and beyond…
- **Autonomous:** there are many sub-problems in making a system fully autonomous, e.g.:
  - Learning
  - Perception / Inference
  - Decision making / Planning

The focus of this work

Informative Planning

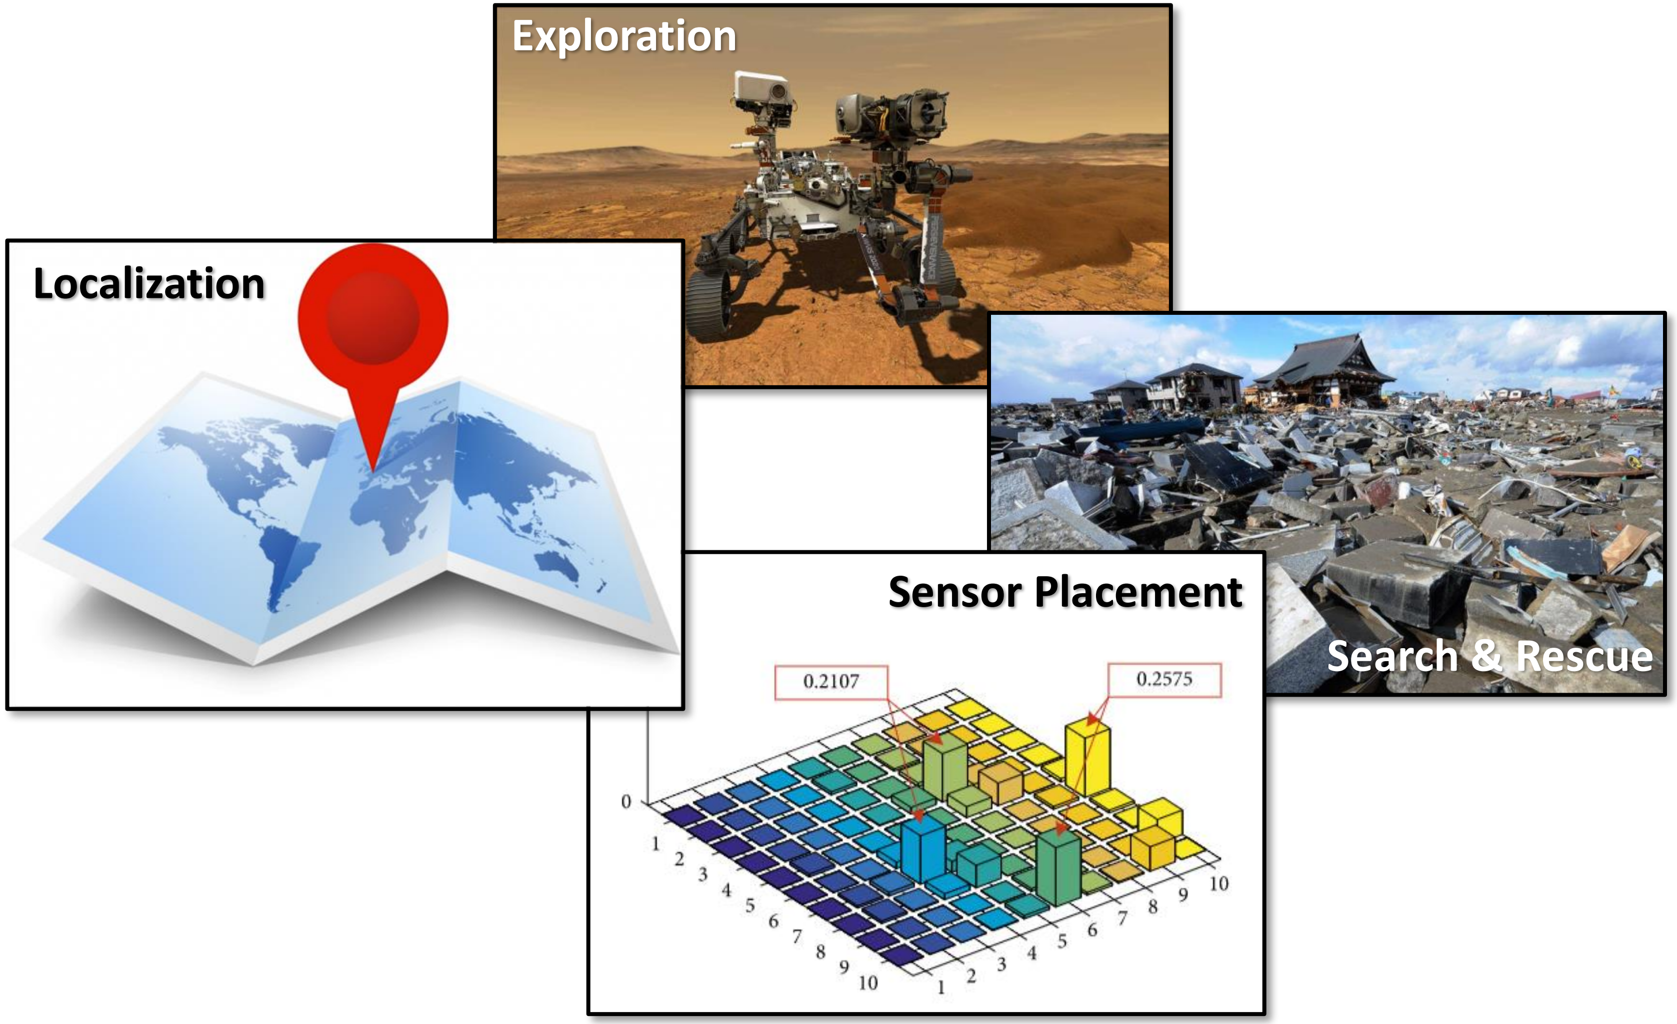

Problem setting & contributions

- What is the problem we aim to solve?
  - Informative Planning
  - (Under uncertainty)
  - Belief-space is high-dimensional
  - Beliefs are non-parametric

- What are the contributions?
  I. Dimensionality reduction for evaluating uncertainty (**involve-MI**)
     - Non-augmented & augmented
  II. Avoiding the reconstruction of future belief's surfaces (**MI-SMC**)
  III. Applicability to belief trees

Related Work

- ~~Informative~~ Planning with ~~High-Dimensional~~ Non-Parametric Beliefs:
  - Kurniawati et al., RSS'04 (SARSOP)
  - Silver and Veness, NIPS'10 (POMCP)
  - Somani et al., NIPS'13 (DESPOT)
  - Garg et al., RSS'19 (DESPOT-α)
- Informative Planning with ~~High-Dimensional~~ Non-Parametric Beliefs:
  - Sunberg and Kochenderfer, ICAPS'18 (PFT-DPW)
  - Fischer and Tas, ICML'20 (IPFT)
  - Platt et al., ISRR'11
- Informative Planning with High-Dimensional ~~Non-Parametric~~ Beliefs:
  - Kopitkov and Indelman, IJRR'17
  - Elimelech and Indelman, IJRR'21 (Accepted)
- Related to specific planning/decision making tasks:
  - Chli, Ph.D. Thesis'09 (CLAM)
  - Zhang et al., IJRR'20 (FSMI)
  - Stachniss et al., RSS'05

Problem Formulation

Beliefs: accounting for uncertainty

- State at time t:

$$X_t \in \mathbb{R}^D$$

  - D is the dimension of the state

- In Full SLAM, for example, the state is composed of the robot’s trajectory and the map:

$$X_t = \{x_{1:t}, M\}$$

Beliefs: accounting for uncertainty

- State at time t: X_t

- Probabilistic **transition** model:
  - $x_t = g(X_{t-1}, a_{t-1}, \text{noise})$
  - $x_t \sim \mathbb{P}_T(x_t \mid X_{t-1}, a_{t-1})$

Beliefs: accounting for uncertainty

- State at time t: X_t
- Probabilistic transition model: $x_t \sim \mathbb{P}_T(x_t \mid X_{t-1}, a_{t-1})$

- Probabilistic **observation** model:
  - $Z_t = h(X_t, \text{noise})$
  - $Z_t \sim \mathbb{P}_Z(Z_t \mid X_t)$

Beliefs: accounting for uncertainty

- State at time t: X_t
- Probabilistic transition model: $x_t \sim \mathbb{P}_T(x_t \mid X_{t-1}, \boldsymbol{a_{t-1}})$
- Probabilistic observation model: $Z_t \sim \mathbb{P}_O(\boldsymbol{Z_t} \mid X_t)$

- History up to time t:

$$h_t = \{\boldsymbol{Z_{1:t}}, \boldsymbol{a_{0:t-1}}\}$$

- **Belief** at time t:

$$b[X_t] \triangleq \mathbb{P}(X_t \mid h_t)$$

Beliefs: accounting for uncertainty

Beliefs can be represented as factor graphs

- Bi-partite graphs
- Variable nodes
- Factor nodes – probabilistic constraints

For example:

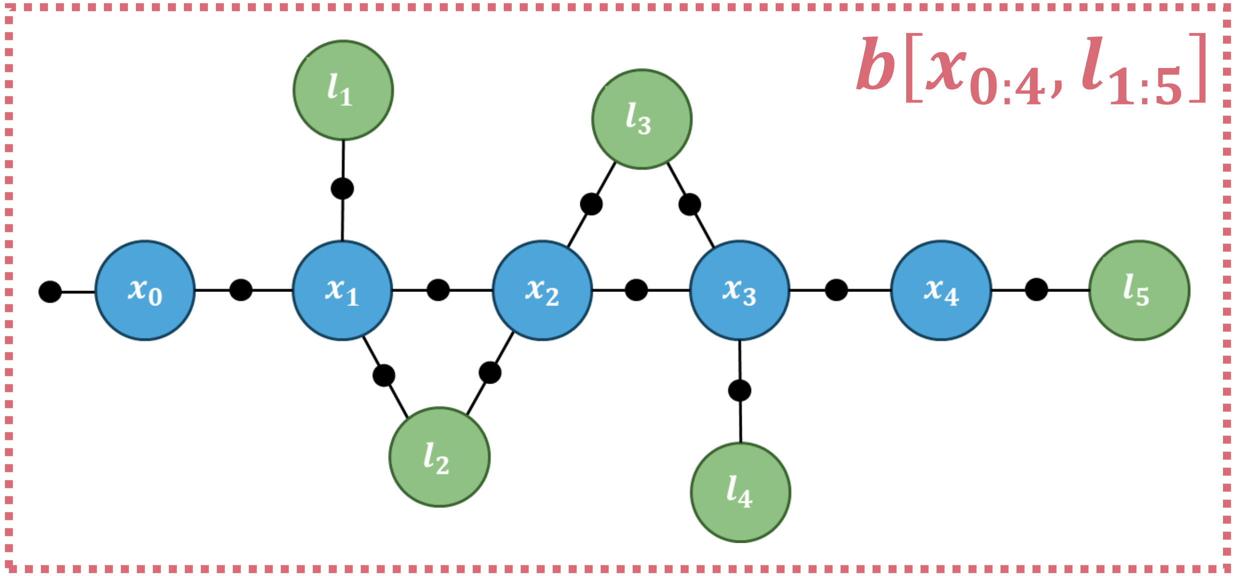

Planning

- Toy example:
- A drone (✦) needs to choose where to go next
- **What is the best action?**

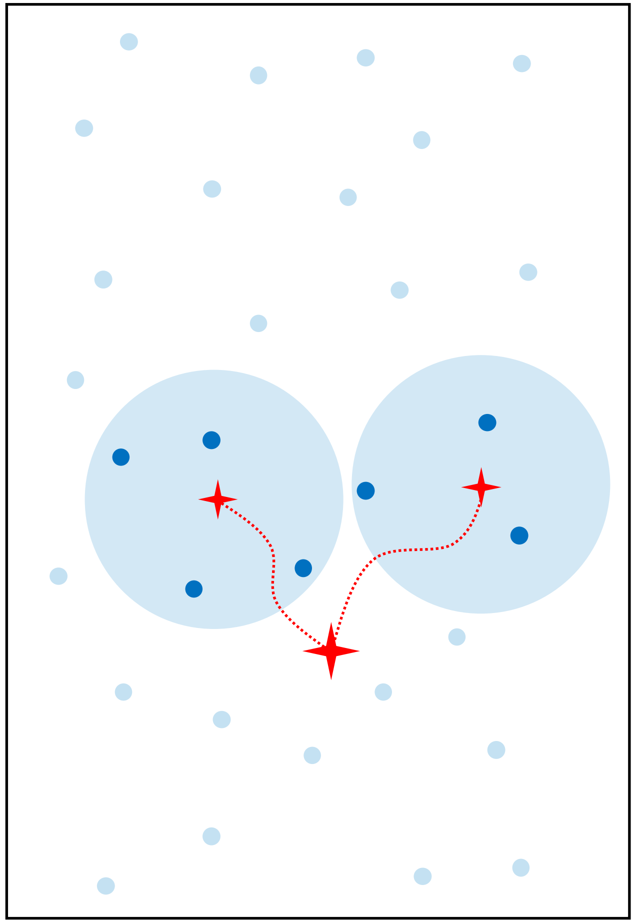

Planning framework – POMDPs

- A Partially Observable Markov Decision Process (POMDP):

$$\langle \mathcal{X}, \mathcal{A}, \mathcal{Z}, \boxed{b[X_0]}, \mathbb{P}_T, \mathbb{P}_Z, r \rangle$$

$$\Downarrow$$

Belief at planning time = prior

- Rewards:

$$r_t = r(X_t, a_t), r_T = r(X_T)$$

- Objective function:

$$J(b[X_0], a_{0:T-1}) = \mathbb{E}_{Z_{1:T}}\left[\sum_{t=1}^{T-1} r_t + r_T\right]$$

- Choosing the best action (sequence):

$$a_{0:T-1}{}^{*} = \underset{a_{0:T-1}\in\mathcal{A}}{\operatorname{argmax}}\, J(b[X_0], a_{0:T-1})$$

Our approach also supports policy formulation

Planning framework – $\boldsymbol{\rho}$-POMDPs

- For **Informative Planning** we use $\boldsymbol{\rho}$-POMDP:

$$\langle \mathcal{X}, \mathcal{A}, \mathcal{Z}, b[X_0], \mathbb{P}_T, \mathbb{P}_Z, \boldsymbol{\rho} \rangle$$

- Info-theoretic rewards (allows measuring uncertainty):

$$\rho_t = \rho(\boldsymbol{b}[\boldsymbol{X_t}], a_t),\ \rho_T = \rho(\boldsymbol{b}[\boldsymbol{X_t}])$$

- Objective function:

$$J(b[X_0], a_{0:T-1}) = \mathbb{E}_{\mathcal{Z}_{1:T}}\left[\sum_{t=1}^{T-1} \boldsymbol{\rho_t} + \boldsymbol{\rho_T}\right]$$

- Choosing the best action (sequence):

$$a_{0:T-1}{}^{*} = \underset{a_{0:T-1} \in \mathcal{A}}{\operatorname{argmax}}\, J(b[X_0], a_{0:T-1})$$

POMDPs VS. $\boldsymbol{\rho}$-POMDPs

POMDP:

- Reward: $r(X_t, a_t)$

ρ-POMDP:

- Reward: $\rho(\boldsymbol{b}[\boldsymbol{X_t}], a_t)$

Reward over the belief

$\Downarrow$

Evaluation of the distribution's value

$\Downarrow$

Higher complexity

Solving a planning problem

The solution is comprised of many building blocks, e.g.:

- Optimization method (belief tree, gradient descent)
- Inference method (factor graph methods, particle filter)
- Reward evaluation

The focus of this work

Info-theoretic rewards

- (Negative) Differential entropy:
  - Before getting an observation:

$$-\mathcal{H}[X] \triangleq \int_{\mathcal{X}} \mathbb{P}(X) \log \mathbb{P}(X) \, dX$$

  - After getting an observation:

$$-\mathcal{H}[X \mid Z = z] \triangleq \int_{\mathcal{X}} \mathbb{P}(X \mid Z = z) \log \mathbb{P}(X \mid Z = z) \, dX$$

High-dim. Integration

- Information Gain (IG):

$$IG[X; Z = z] \triangleq \mathcal{H}[X] - \mathcal{H}[X \mid Z = z]$$

The amount of information gained by this observation

Augmentation

- In this work, we use a **smoothing** formulation
  - Reminder: $x_t = g(X_{t-1}, a_{t-1}, \text{noise})$
  - When transitioning the state, the new state is **augmented** to the previous:
    $$X_t = \{X_{t-1}, x_t\}$$
  - In filtering: a subset of X_{t-1} might be marginalized out
  - Necessary only in the context of planning

- We might encounter this, for example, in active SLAM

- In short, we will denote:
  $$X = X_0, X_{new} = x_{1:t}, X' = X_t$$

Info-theoretic rewards

With augmentation:

- (Negative) Differential entropy:

$$-\mathcal{H}[X' \mid Z = z] \triangleq \int_{\mathcal{X}'} b[X'] \log b[X'] \, dX' \; ^*$$

- **Augmented** IG:

$$IG_{aug}[X \boxplus X_{new}; Z = z] \triangleq \mathcal{H}[X] - \mathcal{H}[X, X_{new} \mid Z = z]$$

Measures also the uncertainty introduced by transitioning

- Augmented IG is a generalization of IG
  - For $X_{new} = \emptyset \Rightarrow IG_{aug} = IG$

$^*X' = \{X, X_{new}\}$

Problem Formulation

Info-theoretic rewards

- Which reward function should we use?

- For ρ-POMDPs:
  - The prior entropy $\mathcal{H}[X]$ is equal for any action
  - $-\mathcal{H}[X, X_{new} \mid Z = z] \Leftrightarrow \underbrace{IG_{aug}[X \boxplus X_{new}; Z = z]}_{= \mathcal{H}[X] - \mathcal{H}[X, X_{new} \mid Z = z]}$

- We will work with (augmented) IG

Non-parametric beliefs

- Usually represented with a weighted particle set:

$$\{X^{(i)}, w^{(i)}\}_{i=1}^{N}$$

- In order to have sufficient resolution: $N \propto \alpha^D$ $(\alpha > 1)$

Curse of Dimensionality

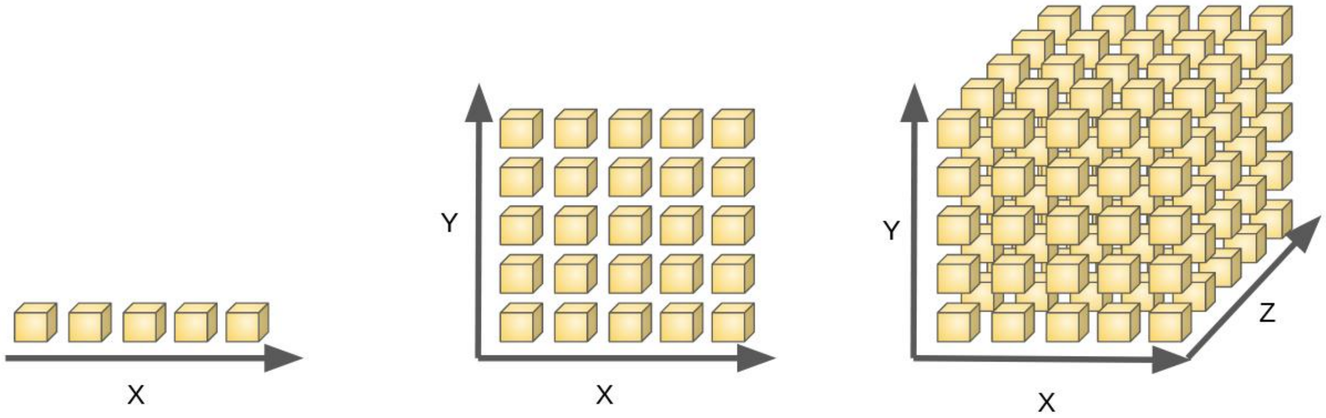

Approach

Approach

Dimensionality reduction for reward evaluation

Dimensionality reduction

- Getting back to the toy example:
  - A drone (✦) needs to choose between 2 actions
  - The belief is high-dimensional – many landmarks (• •)
  - It might observe only a subset of the landmarks set (•)
    - The involved variables

- Can we solve the informative planning problem considering only the **involved** variables?

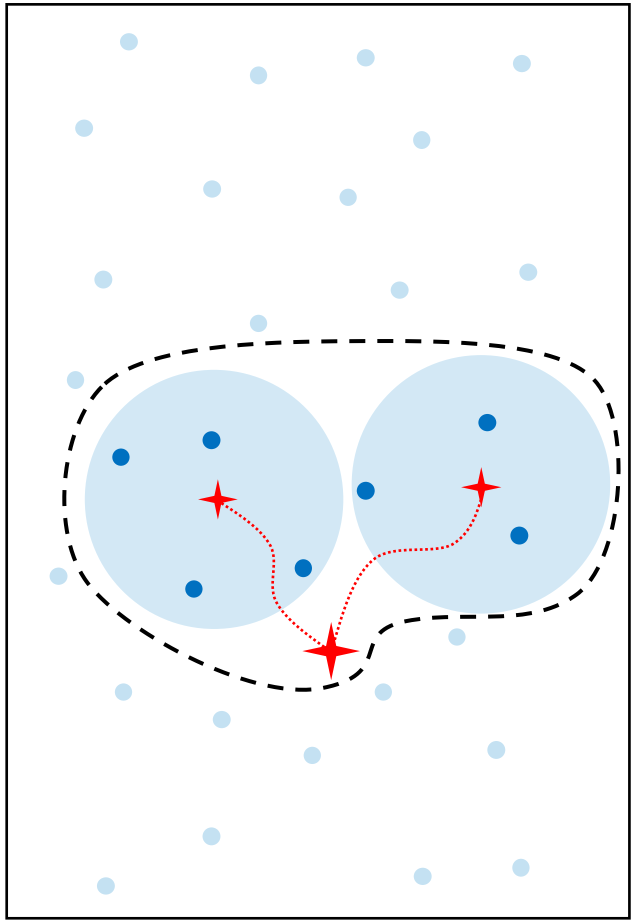

Dimensionality reduction

The involved variables

- Represented with a factor graph:

$X = \{x_{0:4}, l_{1:5}\}$

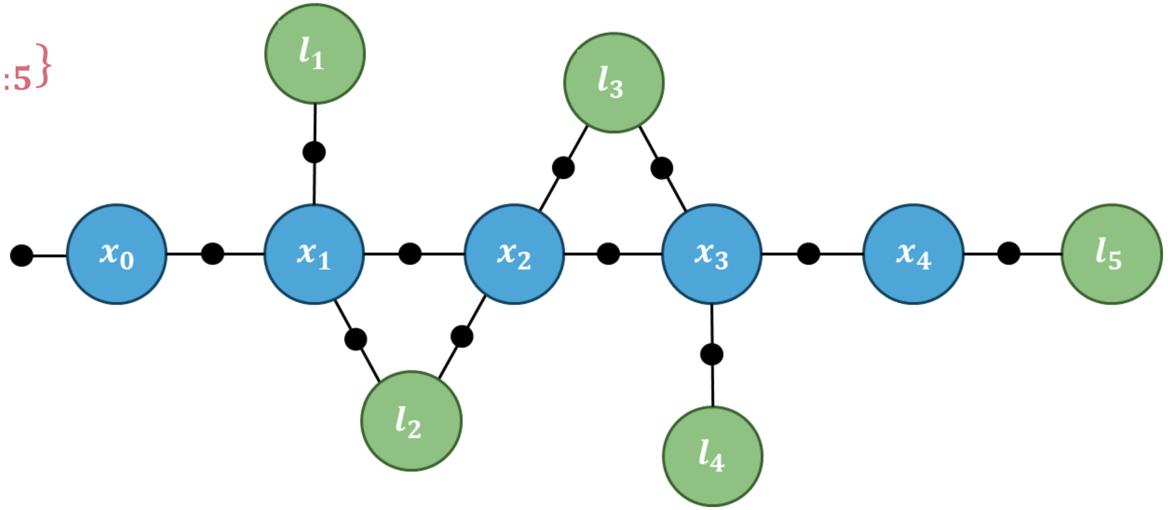

Dimensionality reduction

The involved variables

- Represented with a factor graph:

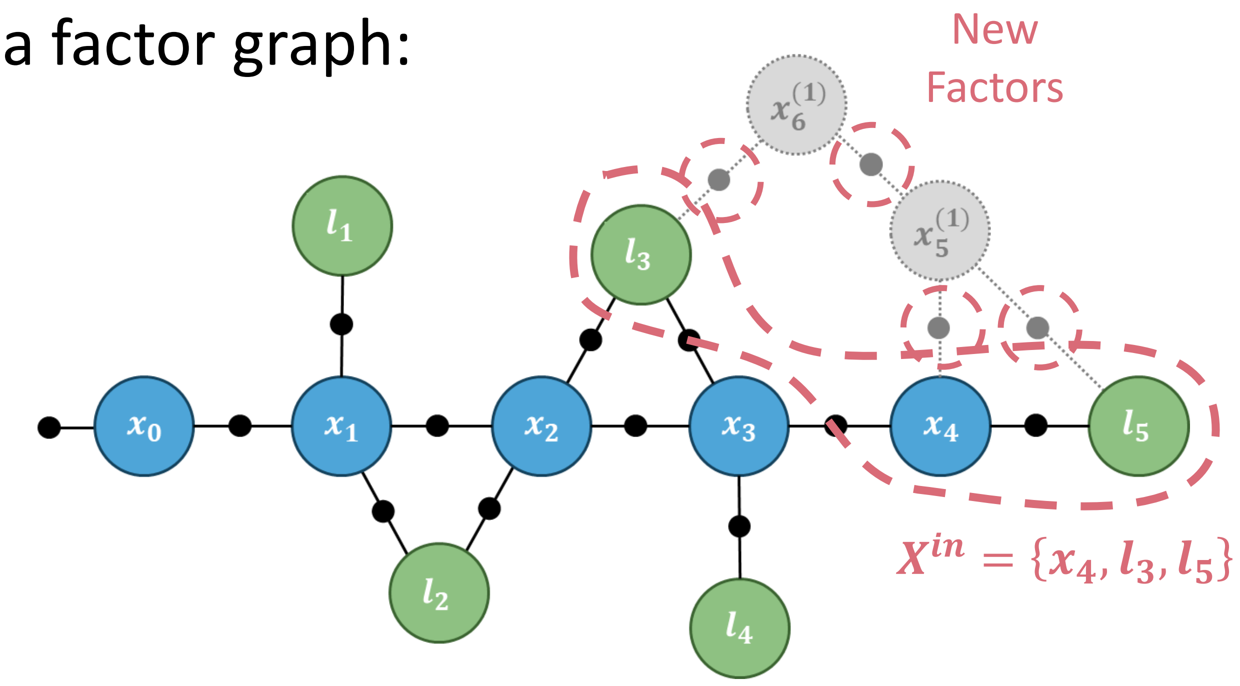

Dimensionality reduction

The involved variables

- Represented with a factor graph:

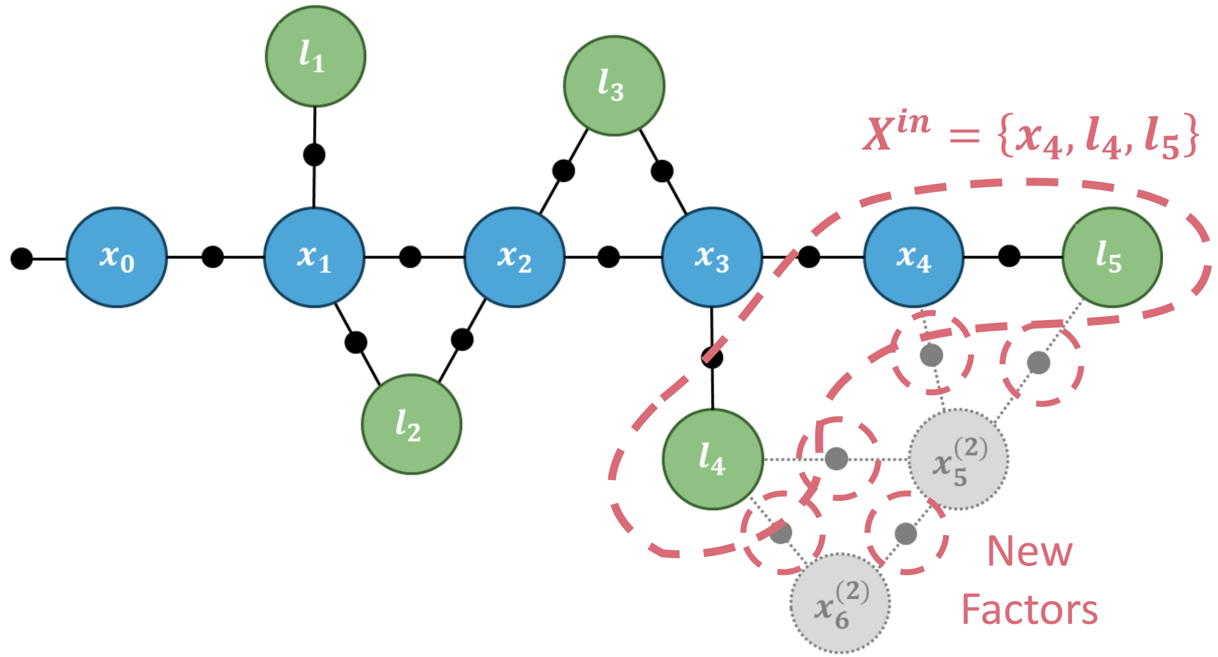

Dimensionality reduction

The involved variables

- Directly participate in generating future transitions and observations
  - $\mathbb{P}_T(x_t \mid X_{t-1}, a_{t-1}) = \mathbb{P}_T(x_t \mid \boldsymbol{X_{t-1}^{tr}}, a_{t-1})$
  - $\mathbb{P}_Z(Z_t \mid X_t) = \mathbb{P}_Z\left(Z_t \mid \boldsymbol{X_t^{obs}}\right)$

$\boldsymbol{X_{t-1}^{tr}}, \boldsymbol{X_t^{obs}} \in \boldsymbol{X^{in}}$

- A subset of the entire prior state: $X^{in} \subseteq X$

- Might be much lower-dimensional: $X^{in} \in \mathbb{R}^d$, $d \ll D$

Approach

Dimensionality reduction

The involved variables

- Given the involved variables $\Leftrightarrow$ exploiting structure
- For Gaussian beliefs:
  - **IG over the involved variables is <u>exactly</u> IG over the entire state:**
    $$IG_{aug}[X \boxplus X_{new}; Z = z] = IG_{aug}\left[X^{in} \boxplus X_{new}; Z = z\right]$$
    - Kopitkov and Indelman IJRR'17 (marginalization)
    - Elimelech and Indelman IJRR'21 (sparsification)
- **For Non-Gaussian – not necessarily true**

Dimensionality reduction

Info-theoretic EXPECTED rewards

- Reminder - objective function:

$$J = \mathbb{E}_{\mathcal{Z}_{1:T}}\left[\sum_{t=1}^{T-1} \rho_t + \rho_T\right]$$

- Can also be written as sum of expected rewards:

$$J = \sum_{t=1}^{T-1} \left[\underbrace{\mathbb{E}_{\mathcal{Z}_{1:T}}[\rho_t]}_{\text{Expected rewards}}\right] + \underbrace{\mathbb{E}_{\mathcal{Z}_{1:T}}[\rho_T]}_{\text{Expected rewards}}$$

- Can we reduce the dimensionality for the evaluation of the **expected** rewards?

Dimensionality reduction

- The expected reward of IG – Mutual Information (MI):

$$I[X;Z] \triangleq \mathbb{E}_{\mathcal{Z}}\big[IG[X;Z=z]\big]$$
$$\triangleq \mathcal{H}[X] - \mathcal{H}[X \mid Z]$$

- **MI over the involved variables is <u>exactly</u> MI over the entire state:**

$$\boldsymbol{I[X;Z] = I\left[X^{in};Z\right]}$$

  - For any distribution
  - Integration is done over a smaller subset of state
  - Underlying assumption: data association is assumed to be solved (at planning time)

Dimensionality reduction

$$I[X; Z] = I[X^{in}; Z]$$

Information diagram

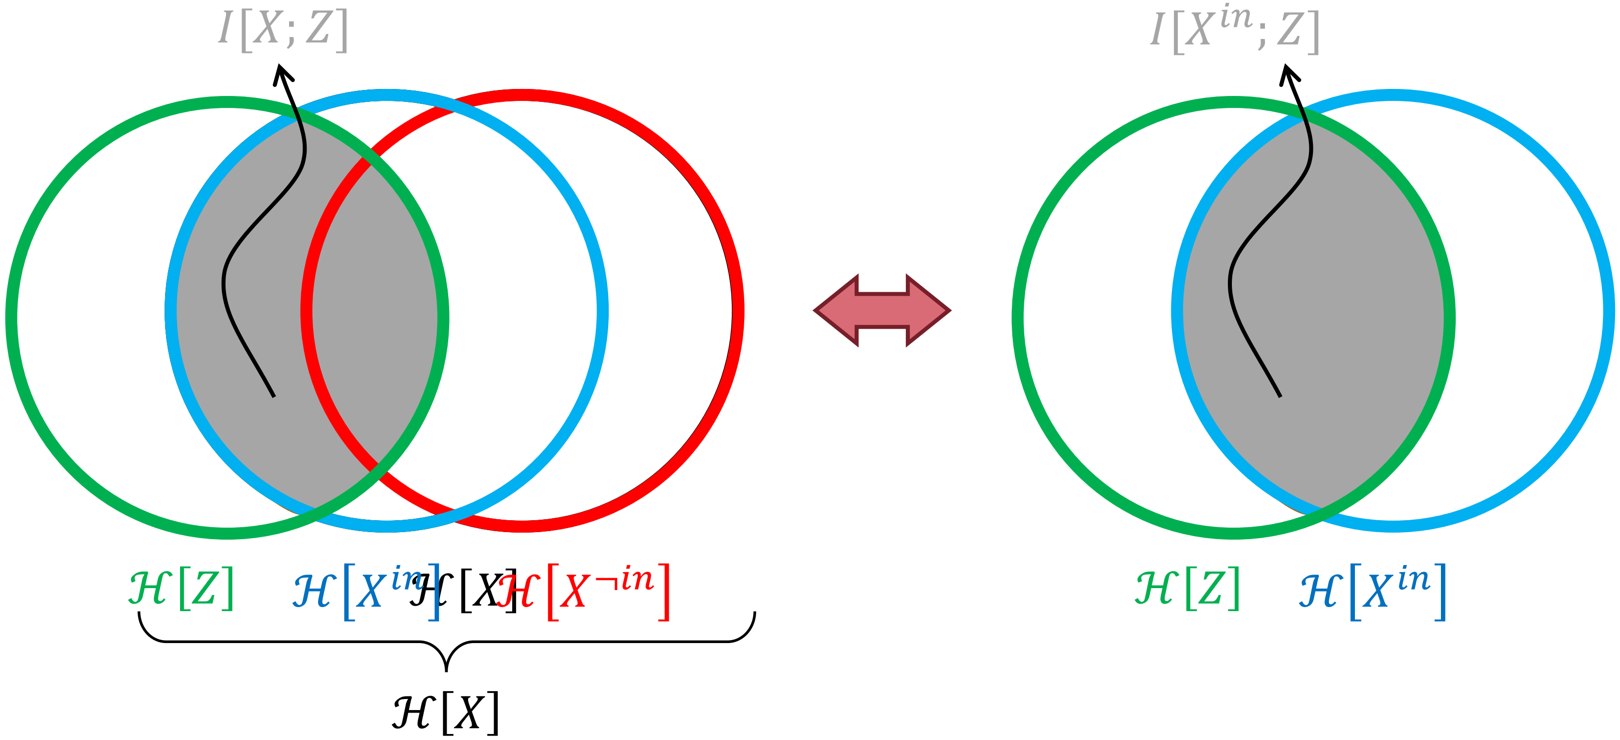

Approach

Dimensionality reduction

Augmentation

- **Augmented** mutual information:

$$I_{aug}[X \boxplus X_{new}; Z] \triangleq \mathbb{E}_Z\left[IG_{aug}[X \boxplus X_{new}; Z = z]\right]$$
$$\triangleq \mathcal{H}[X] - \mathcal{H}[X, X_{new} \mid Z]$$

- Dimensionality reduction for the **augmented** **MI** calculation:

$$\boldsymbol{I_{aug}[X \boxplus X_{new}; Z] = I_{aug}\left[X^{in} \boxplus X_{new}; Z\right]}$$

- Relation to MI:

$$I_{aug}[X \boxplus X_{new}; Z] = \underbrace{I\left[X^{in}, X_{new}; Z\right]}_{\text{Original MI}} - \underbrace{\mathcal{H}\left[X_{new} \mid X^{in}\right]}_{\text{Expected uncertainty from transition model}}$$

Approach

Dimensionality reduction

$$I_{aug}[X \boxplus X_{new}; Z] = I_{aug}\left[X^{in} \boxplus X_{new}; Z\right]$$
$$= I\left[X^{in}, X_{new}; Z\right] - \mathcal{H}\left[X_{new} \mid X^{in}\right]$$

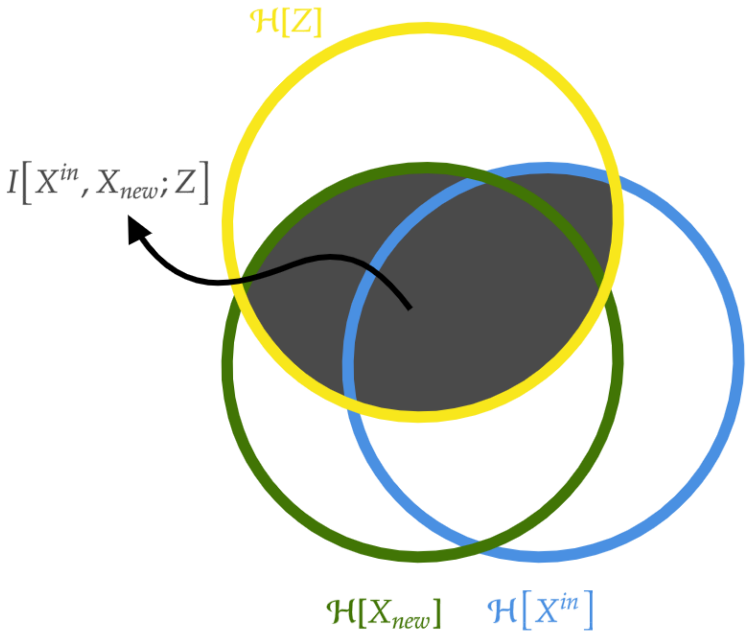

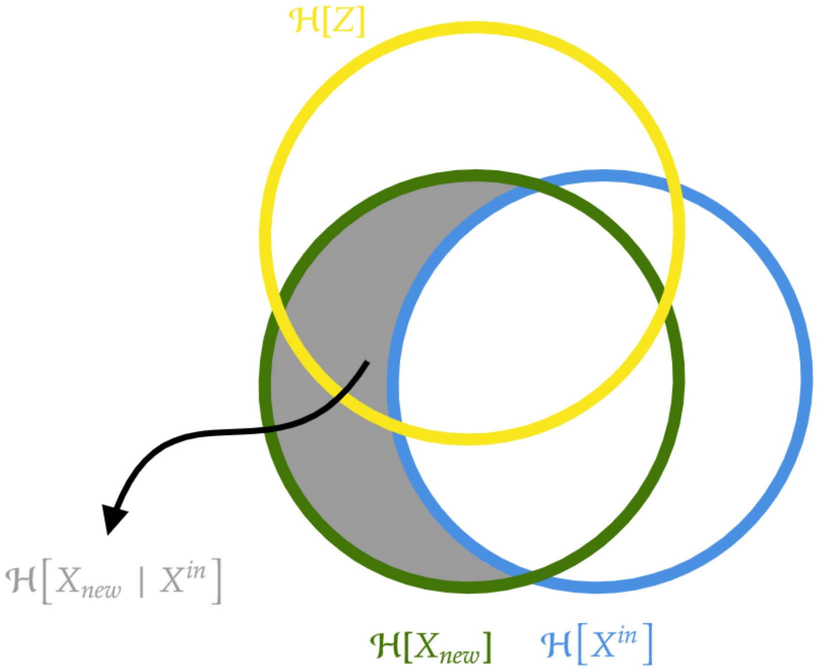

Dimensionality reduction

involve-MI

- An algorithm which uses:

$$I_{aug}[X \boxplus X_{new}; Z] = I_{aug}\left[X^{in} \boxplus X_{new}; Z\right]$$

- Basic framework:
  - Determine involved (with some heuristic)
  - Marginalize out uninvolved variables
  - Calculate MI over the involved variables

- It can be used for any task involving MI calculation between two multi-dimensional variables

Dimensionality reduction

involve-MI - extension

- Marginalization might entail heavy costs
- We can choose a bigger subset $\boldsymbol{X^{in+}}$, which follows
$$X^{in} \subseteq \boldsymbol{X^{in+}} \subseteq X$$
  - e.g. for one-time marginalization: $\boldsymbol{X^{in+}} = \bigcup_{a \in \mathcal{A}} X^{in(a)}$
- A more general relation:
$$I_{aug}[X \boxplus X_{new}; Z] = I_{aug}\left[\boldsymbol{X^{in+}} \boxplus X_{new}; Z\right]$$

Approach

Dimensionality reduction

involve-MI - extension

- Represented with a factor graph:

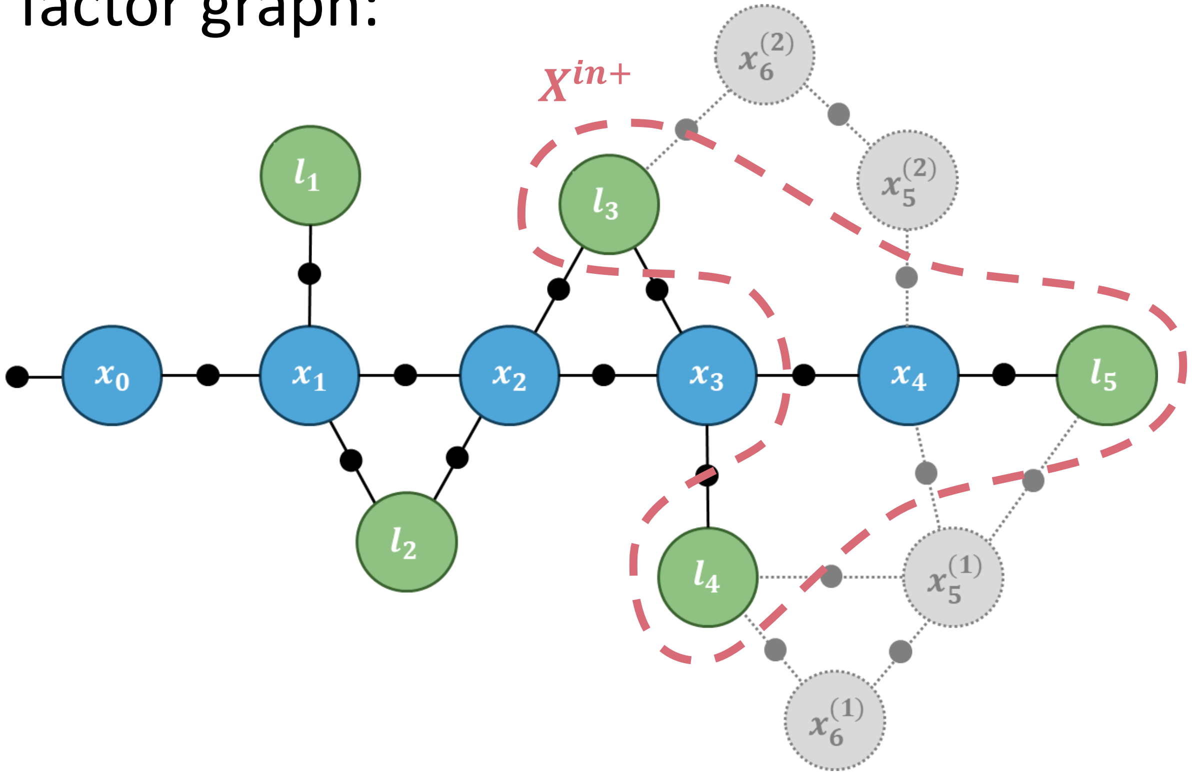

Approach

Dimensionality reduction

- involve-MI breaks the relation between the dimensionality D of the state to the accuracy and complexity
- These are now dependent on $d \ll D$
- Only $n \propto \alpha^d$ samples are needed, and $n \ll N$

Approach

Approach

Avoiding the reconstruction of future belief's surfaces

Avoiding beliefs reconstruction

- The common approach to calculate MI is to go through the entropy terms
  - Reminder: $I_{aug}[X \boxplus X_{new}; Z] \triangleq \mathcal{H}[X] - \mathcal{H}[X, X_{new} \mid Z]$

- The entropy terms are calculated through the evaluation of the posterior beliefs
  - Reminder: $-\mathcal{H}[X' \mid Z = z] \triangleq \int_{\mathcal{X}'} b[X'] \log b[X'] \, dX'$

- Some estimators wish to first reconstruct these beliefs, such as the re-substitution estimator:

$$\hat{\mathcal{H}}\,[X' \mid Z = z] = \sum_{i=1}^{N} w'^{(i)} \log \boxed{\hat{b}\left[X'^{(i)}\right]}$$

- $\hat{b}[X']$ is an estimation of the belief with e.g. KDE

Approach

Avoiding beliefs reconstruction

- We would like to avoid future beliefs reconstruction:
  - It might add to the estimation error
  - It might entail another level of complexity (hyperparameters)

- The MI can be calculated using known models:

$$I_{aug}\left[\boldsymbol{X}^{in} \boxplus X^{new}; Z\right] = \underbrace{-\mathcal{H}\left[X^{new} \mid \boldsymbol{X}^{in}\right]}_{\text{Transition uncertainty}} - \underbrace{\mathcal{H}\left[Z \mid \boldsymbol{X}^{in}, X^{new}\right]}_{\text{Observation uncertainty}} + \underbrace{\mathcal{H}[Z]}_{\text{Normalization uncertainty}}$$

  - $\mathcal{H}[Z]$ is calculated using known models as well
  - Integration over involved variables only

No need to reconstruct future beliefs surfaces

Approach

Avoiding beliefs reconstruction

MI-SMC

- An estimator which uses the relation

$$I_{aug}\left[X^{in} \boxplus X^{new}; Z\right] = -\mathcal{H}\left[X^{new} \mid X^{in}\right] - \mathcal{H}\left[Z \mid X^{in}, X^{new}\right] + \mathcal{H}[Z]$$

- General idea:
  - Propagate state samples in a Sequential Monte Carlo (SMC) manner
  - Generate possible future observations
  - Evaluate the models at these sampled instances
    - Inherently we need samples of only the involved variables

Related to involve-MI

Avoiding beliefs reconstruction

MI-SMC

- Complexity: $O(nmd)$
  - n – number of state samples
  - m – number of observation samples
  - d – dimensionality of involved variables subset
  - RS-KDE complexity (for example): $\boxed{O(n^2md)}$

For using involve-MI
Otherwise: $O(N^2mD)$

- Anytime algorithm

Approach

Approach

Applicability to belief trees

Applicability to belief trees

- The solution to the planning problem is obtained by maximizing the objective function:

$$J_0^* = \max_{a_{0:t-1}} \left\{ \mathbb{E}_{\mathcal{Z}_{1:T}} \left[\sum_{t=1}^{T-1} \rho_t + \rho_T \right] \right\}$$

- Recursively – the Bellman optimality equation:

$$J_t^* = \max_{a_t} \{ \rho_t + \mathbb{E}_{\mathcal{Z}_{t+1}} [J_{t+1}^*] \}$$

- Commonly solved with a belief tree:

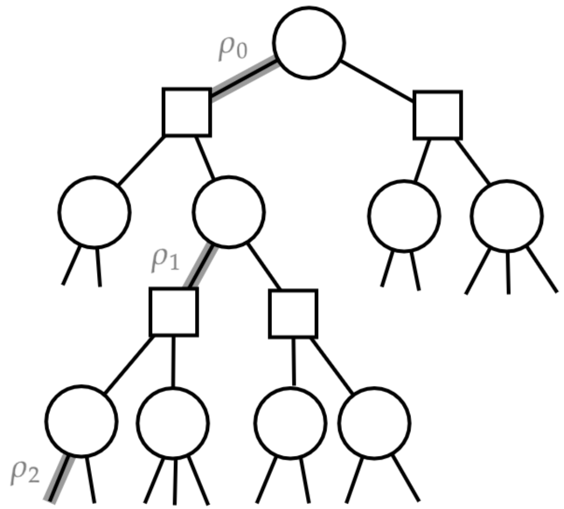

Approach

Applicability to belief trees

- The Bellman optimality equation:

$$J_t^* = \max_{a_t}\{\boldsymbol{\rho_t} + \mathbb{E}_{z_{t+1}}[J_{t+1}^*]\}$$

- Reminder: our approach deals with EXPECTED reward, $\mathbb{E}[\boldsymbol{\rho_t}]$

- Goal: our approach should cope with belief tree solvers

Not trivial

Applicability to belief trees

How can ***involve-MI*** *cope with belief trees?*

- Reminder: $\underbrace{I_{aug}[X \boxplus X_{new}; Z]}_{\mathbb{E}[\boldsymbol{\rho}_t]} = \underbrace{I_{aug}\left[X^{in} \boxplus X_{new}; Z\right]}_{\mathbb{E}[\boldsymbol{\rho}_t^{in}]}$

- We cannot state in general that $\boldsymbol{\rho}_t$ and $\boldsymbol{\rho}_t^{in}$ are equal (without expectation)

- However: solving the optimization problem with any of these rewards is equivalent!

Approach

Applicability to belief trees

How can ***involve-MI*** *cope with belief trees?*

- Answer: solve the planning problem with the involved reward:

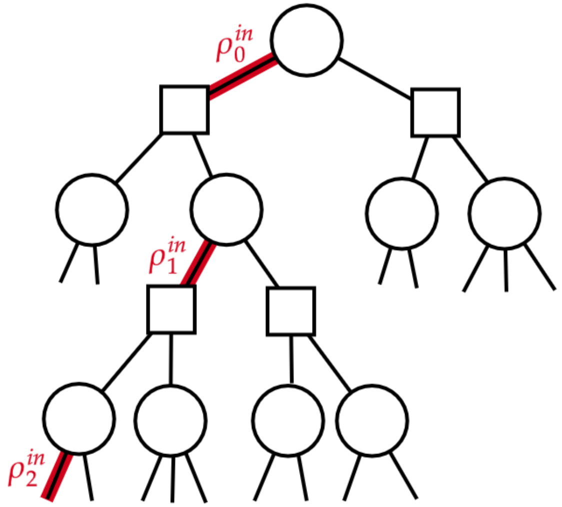

- Using the formulation with $\boldsymbol{X^{in+}} \Longrightarrow$ maintaining a low-dimensional belief during the planning process

Applicability to belief trees

How can ***MI-SMC*** *cope with belief trees?*

- Reminder:
  - MI-SMC calculates MI directly (without going through IG)
  - The MI was treated in general as **sequential**:

$$I_{aug}\left[X^{in} \boxplus X_{new}; Z\right] = \boxed{I_{aug}\left[X_0^{in} \boxplus x_{1:t}; Z_{1:t}\right]} \; \mathbb{E}_{Z_{1:t}}[\cdot]$$

- Naively calculating the sequential MI yields a degenerate belief tree:

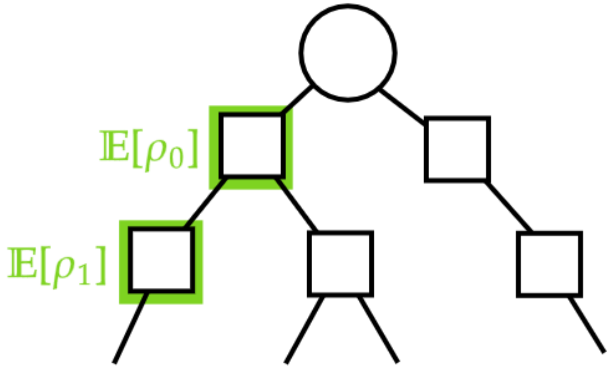

Approach

Applicability to belief trees

How can ***MI-SMC*** *cope with belief trees?*

- Denote:
  - Sequential MI: $I_0^{t^{in}} \triangleq I_{aug}\left[X_0^{in} \boxplus x_{1:t}; Z_{1:t}\right] \Rightarrow \mathbb{E}_{\mathcal{Z}_{1:t}}[\cdot]$
  - Consecutive MI: $I_{t-1}^{t^{in}} \triangleq I_{aug}\left[X_{t-1}^{in} \boxplus x_t; Z_t\right] \Rightarrow \mathbb{E}_{\mathcal{Z}_t}[\cdot]$
- Define a new reward over the consecutive MI values:

$$\rho'_t = \sum_{i=1}^{t+1} \underbrace{\boxed{\left[I_{i-1}^{i^{in}}\right]}}_{\text{MI-SMC}} = \rho'_{t-1} + I_t^{t+1^{in}}$$

- Solving the optimization problem with the new reward is equivalent to solving it with the original reward!

Applicability to belief trees

How can ***MI-SMC*** *cope with belief trees?*

- <u>Answer</u>: solve the planning problem with the new reward:

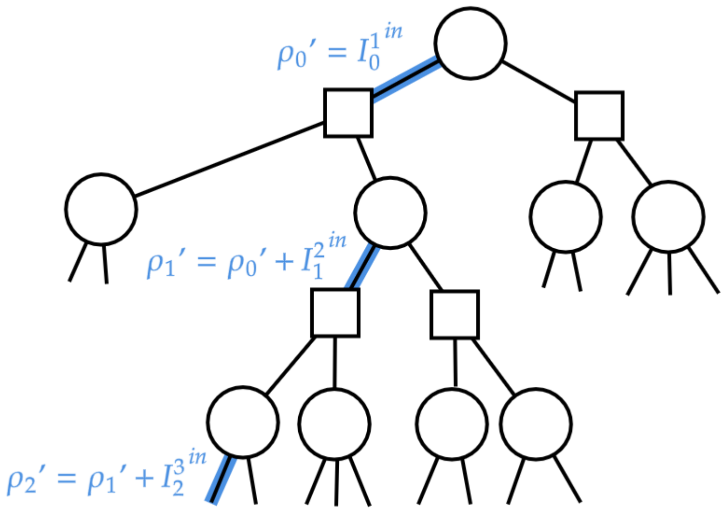

- Using the formulation with $X^{in+} \Longrightarrow$ maintaining a low-dimensional belief during the planning process

Results

Scenario

- 2D SLAM
  - Created synthetically by a factor graph
- Gaussian distributions
  - For analytical solution
- 3 calculation methods:
  - Naïve KDE
  - involve-MI-KDE
  - involve-MI-SMC
- *Note:* KDEs were implemented with perfect inference

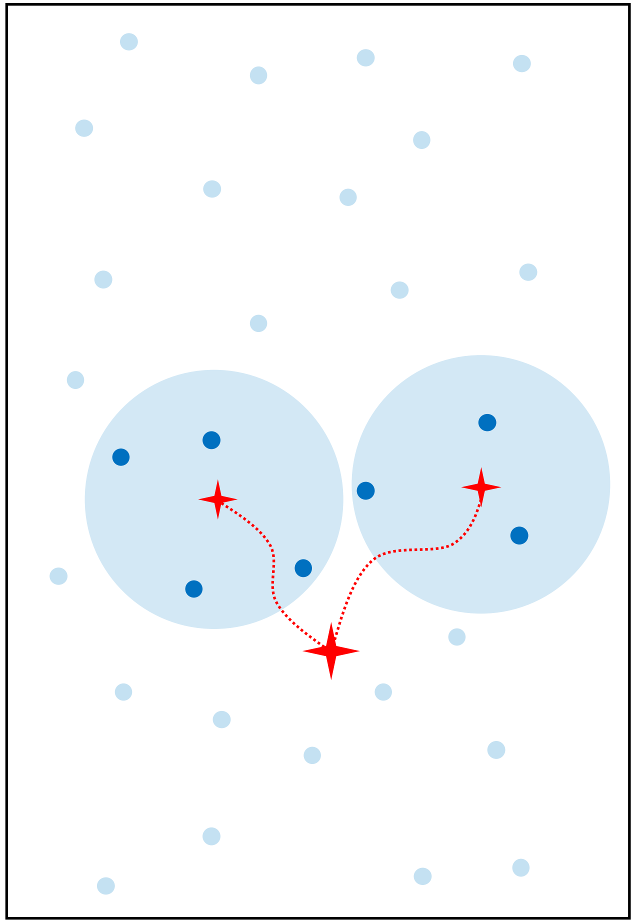

Dimensionality ⟹ Choosing an action

- 4 different actions
- Dimension of X: ~ 150
- Dimension of X^{in}: 4
- Samples: 300
- Trials: 100

- High variance for the **naïve** approach
  - Impacts action to be chosen
- Low variance for **involve-MI**
- Bias of **MI-SMC**

True optimal action ⇓

Chosen action for naive ⇓

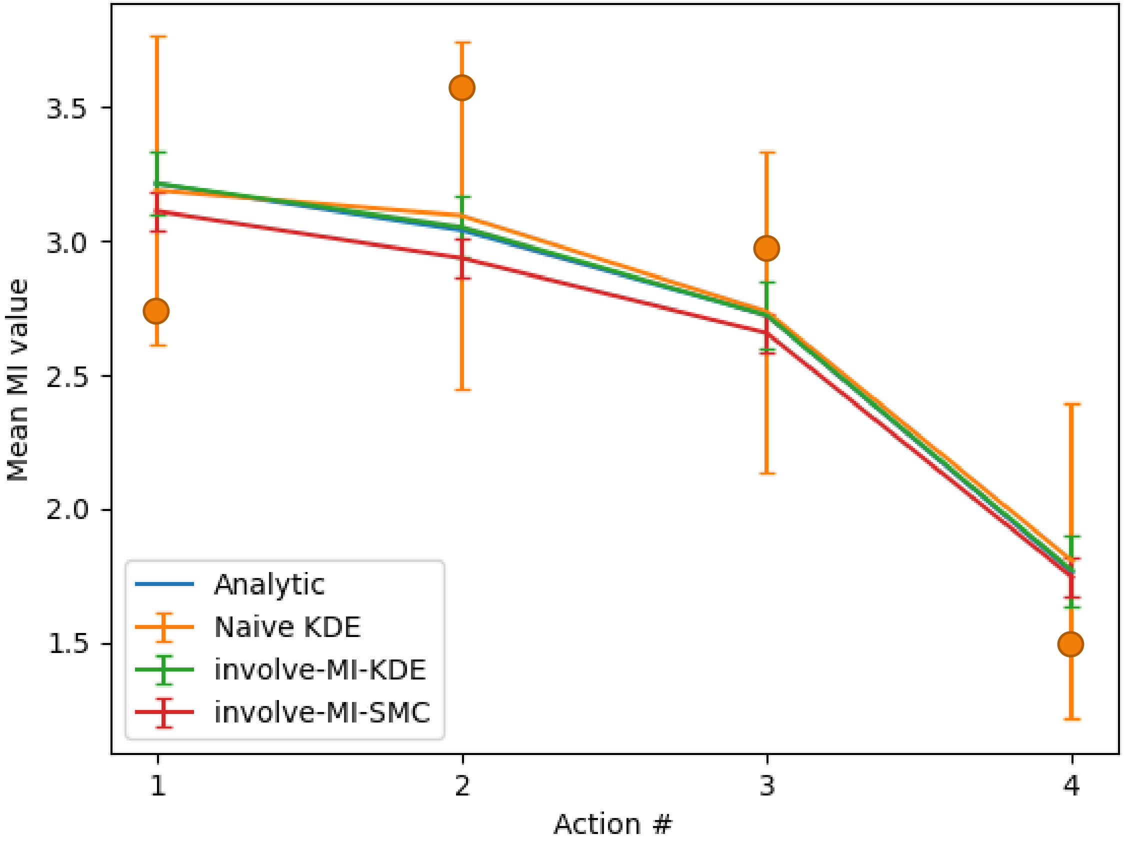

Results

Dimensionality $\Longrightarrow$ Accuracy

- One action
- Increasing dimensionality of X
- Dimension of X^{in}: 4
- Samples: 300
- Trials: 100

- Increasing variance for the **naïve** approach
- Constant variance for **involve-MI**
- Smallest variance for **MI-SMC**

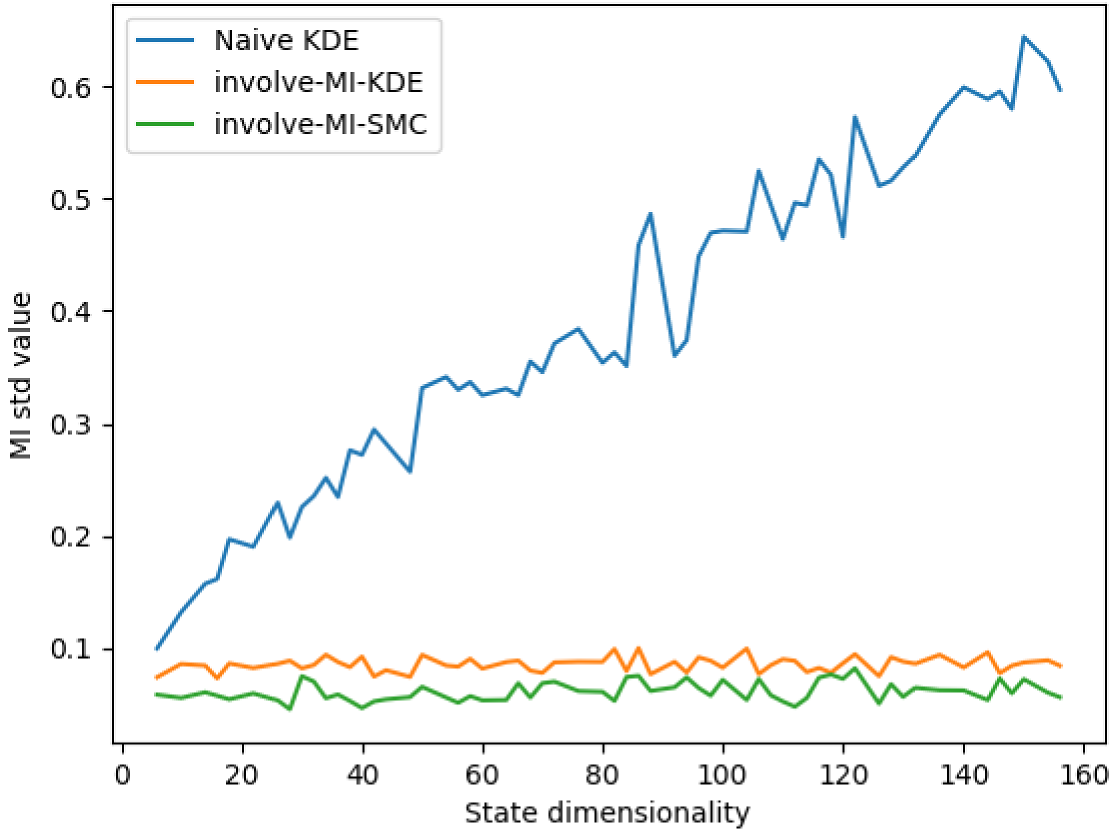

Results

Dimensionality $\Longrightarrow$ Timing

- One action
- Increasing dimensionality of X
- Dimension of X^{in}: 4
- Samples: 300
- Trials: 100

- Increasing calculation time for the **naïve** approach
- Constant calculation time for **involve-MI**
- Smallest calculation time for **MI-SMC**

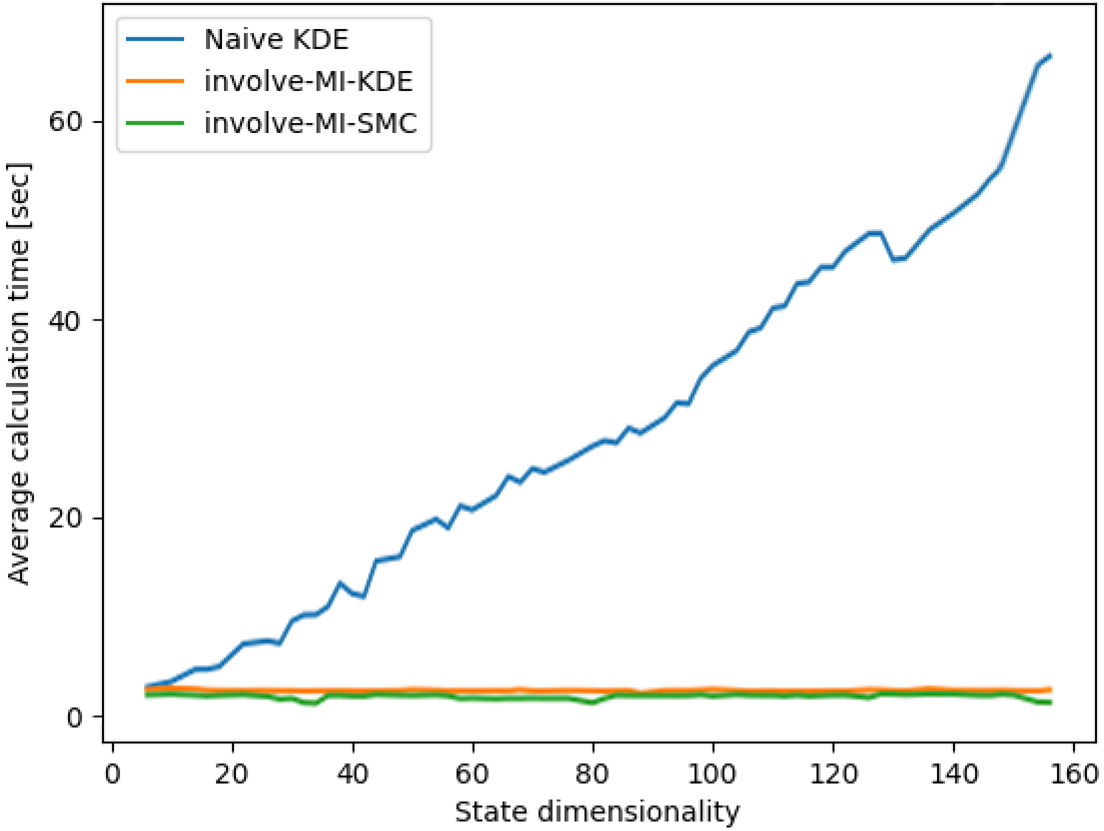

Results

Summary

Summary

- What is the problem we aimed to solve?
  - Informative Planning
  - (Under uncertainty)
  - Belief-space is high-dimensional
  - Beliefs are non-parametric

- What are the contributions?
  I. Dimensionality reduction for evaluating uncertainty (**involve-MI**)
    - Non-augmented & augmented
  II. Avoiding the reconstruction of future belief's surfaces (**MI-SMC**)
  III. Applicability to belief trees

Future research

- Focused case:
  - Quantifying the uncertainty over a subset of the entire state $X^F \subseteq X$
- Non-parametric inference impact:
  - Huang et al., arXiv'21 (NF-iSAM)
  - Relevant for efficient and accurate non-parametric marginalization

Thank you!

Questions?