Structure Aware Probabilistic Inference and Belief Space Planning with Performance Guarantees

Moshe Shienman

Supervisor: Associate Professor Vadim Indelman

Technion – Israel Institute of Technology, Israel

April 2024

Introduction

- A new era: Intelligent autonomous agents and robots

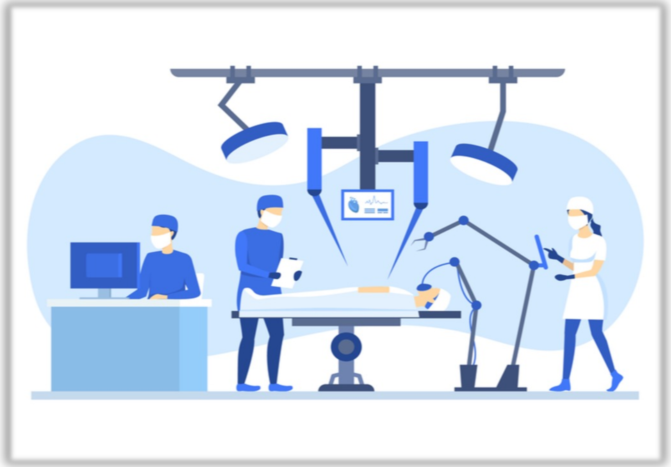

Robotic Surgery

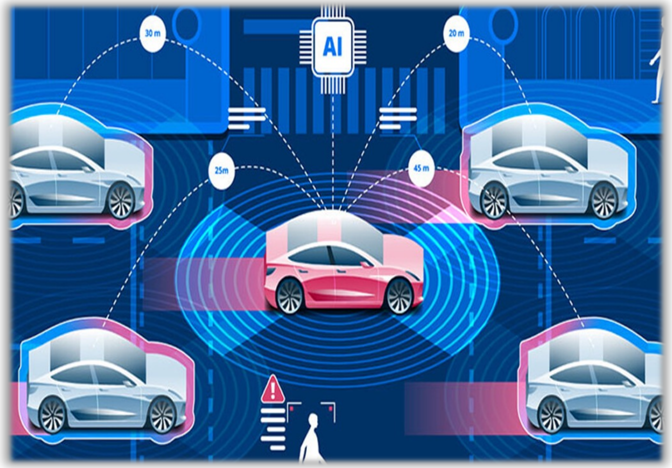

Autonomous Vehicles

Drones

Introduction

- Required to operate reliably and efficiently under different sources of uncertainty

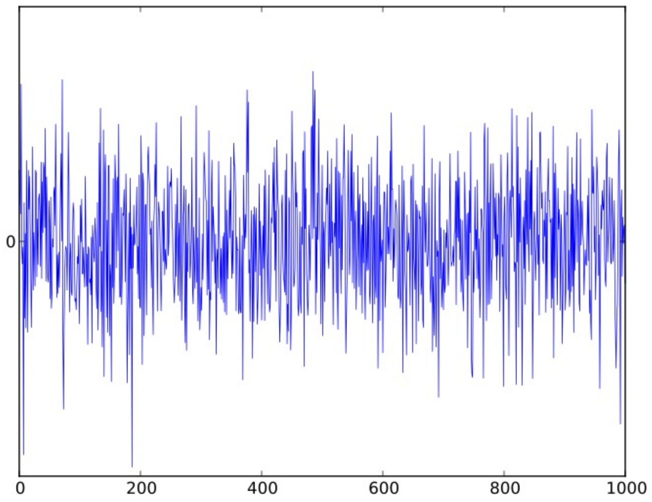

Noisy measurements

Imprecise actions

Dynamic environments

Introduction

- Reason over high dimensional probabilistic states known as *beliefs* for both:

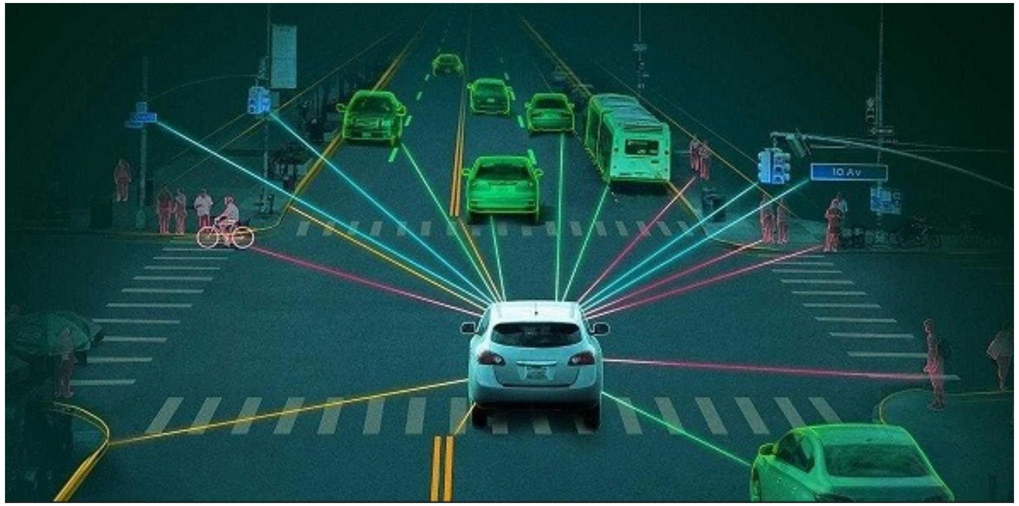

Inference

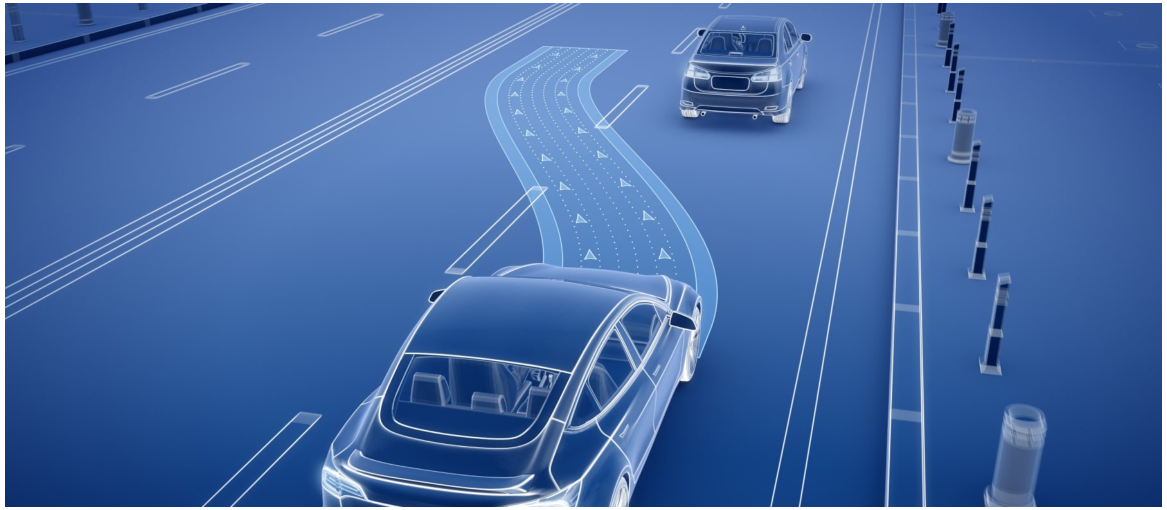

Decision making under uncertainty
aka Belief Space Planning (BSP)

Introduction

- In each discrete time step *k*, the autonomous agent:

Takes action u_k

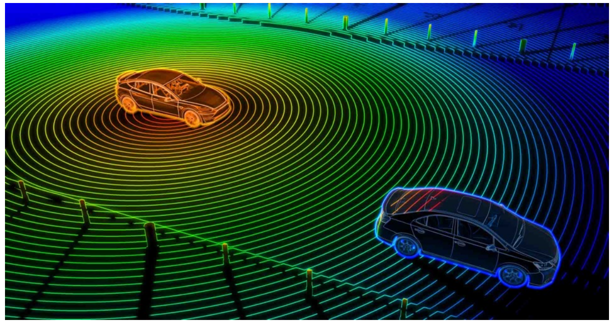

Acquires observation z_k

Notations

- x_k : state at time k (e.g., position and orientation) $X_k = \{x_0, x_1, ..., x_k\}$: the joint state

- $x_{k+1} = f(x_k, u_k, w_k)$, $z_k = h(x_k, v_k)$ motion and observation models with known noise terms, e.g., Gaussian

- $b[X_k] \doteq \mathbb{P}(X_k | z_{0:k}, u_{0:k-1})$: posterior probability density function over the joint state – the **belief**

Probabilistic Inference

- Maximum a Posteriori (MAP) estimate $X_k^* = \underset{X_k}{\operatorname{argmax}}\ b\left[X_k\right] \doteq \underset{X_k}{\operatorname{argmax}}\ \mathbb{P}\left(X_k | z_{0:k}, u_{0:k-1}\right)$

- Gaussian case - a nonlinear least squares problem $\underset{X_k}{\operatorname{argmin}} -\log\left(b\left[X_k\right]\right) = \underset{X_k}{\operatorname{argmin}} \frac{1}{2} \sum_i \|h_i\left(x_i\right) - z_i\|_\Sigma^2$

- Solved with nonlinear optimization methods such as Gauss-Newton, where each iteration solves a *linear* least squares problen $\|\mathcal{A}\Delta X - b\|^2$

$\mathcal{A} =$ 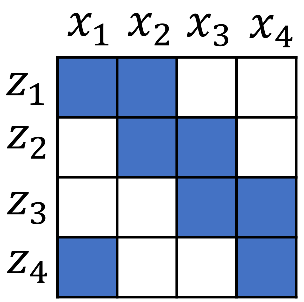

$$\Lambda = \mathcal{A}^T \mathcal{A} = \Sigma^{-1}$$

Measurement Jacobian matrix

Information matrix

Belief Space Planning

- Determine optimal actions with respect to a given objective $\mathcal{U}^* = \underset{\mathcal{U}}{\operatorname{argmin}}\, J\left(b_k, \mathcal{U}\right)$

- $c\left(\cdot\right) \rightarrow \mathbb{R}$: a cost (reward) function

- A general objective function

$$J\left(b_k, \mathcal{U}\right) = \underset{Z_{k+1:k+L}}{\mathbb{E}} \left[\sum_{l=0}^{L-1} c_l\left(b\left[X_{k+l}\right], u_{k+l}\right) + c_L\left(b\left[X_{k+L}\right]\right) \right]$$

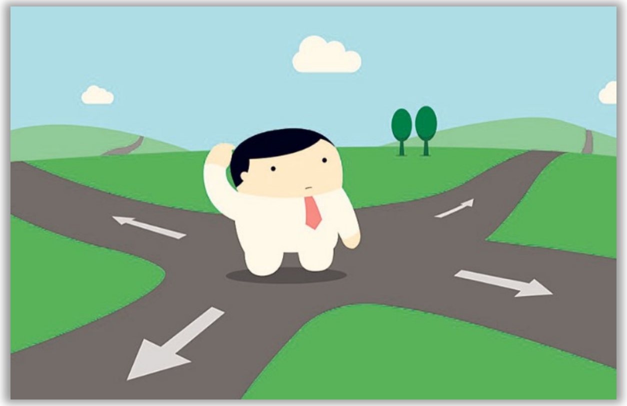

Computational Challenge

- Both inference and BSP are computationally very hard in high dimensional state spaces!

- A challenge in real-world autonomous systems where the agent is required to operate in real time, often using inexpensive hardware

Research Goal

- Leverage structures and solve simplified problems while providing performance guarantees

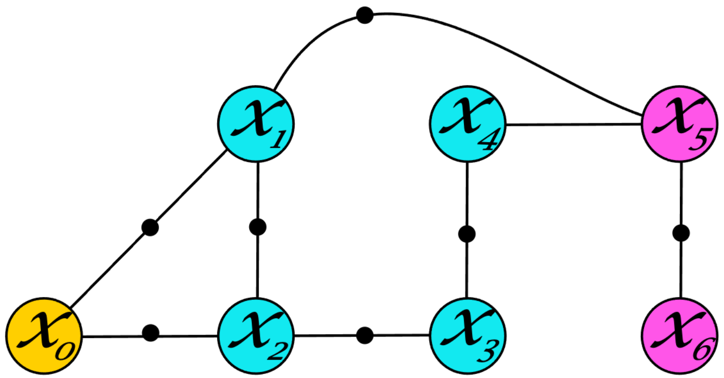

Topological Structures

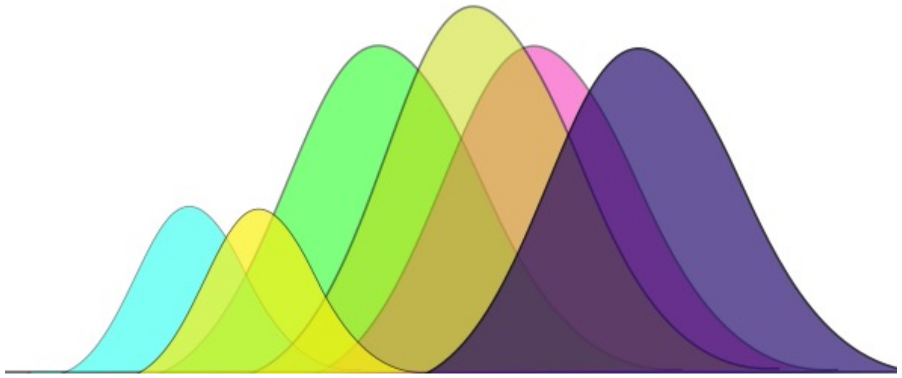

Posterior *Beliefs* Structures

Belief Space Planning - Simplification

- Solve an alternative problem with respect to the same set of candidate actions

1. Be less expensive to compute

2. Should discriminate between candidate actions

3. Would ideally yield a solution which is consistent with the optimal solution of the original problem

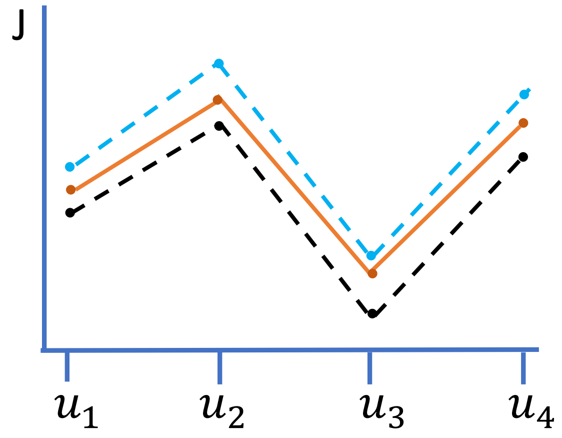

$$\mathcal{LB}\left[J\left(b_k,u_k\right)\right]\leq J\left(b_k,u_k\right)\leq\mathcal{UB}\left[J\left(b_k,u_k\right)\right]$$

- M. Shienman, A. Kitanov and V. Indelman [2021 IEEE RA-L]
  "Focused Topological Belief Space Planning"

- M. Shienman and V. Indelman [2022 ICRA] ***Outstanding Paper Award Finalist***
  "Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints"

- M. Shienman and V. Indelman [2022 ISRR]
  "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints"

- M. Shienman, O. Levy Or, M. Kaess and V. Indelman [2024 IROS - Submitted]
  "A Slices Perspective for Incremental Nonparametric Inference in High Dimensional State Spaces"

Factor Graph

- [Kschischang et al. 2001] probabilistic graphical model

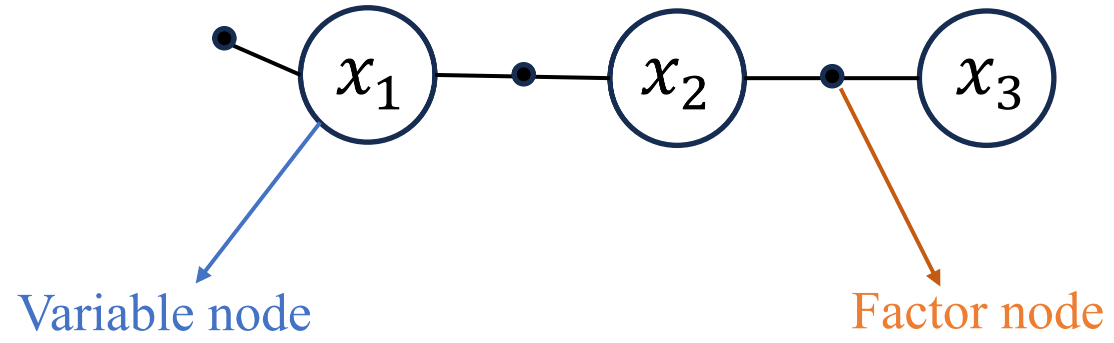

- Represents a factorization of the joint belief

$$b\left[X_3\right] = \eta \cdot \mathbb{P}\left(x_1\right) \cdot \mathbb{P}\left(x_2 | x_1, u_1\right) \cdot \mathbb{P}\left(x_3 | x_2, u_2\right) \propto f\left(x_1\right) \cdot f\left(x_1, x_2\right) \cdot f\left(x_2, x_3\right)$$

Motivation

- Leverage probabilistic graphical models of the underlying problem
- Use topological aspects (signatures) induced from the connectivity of variables

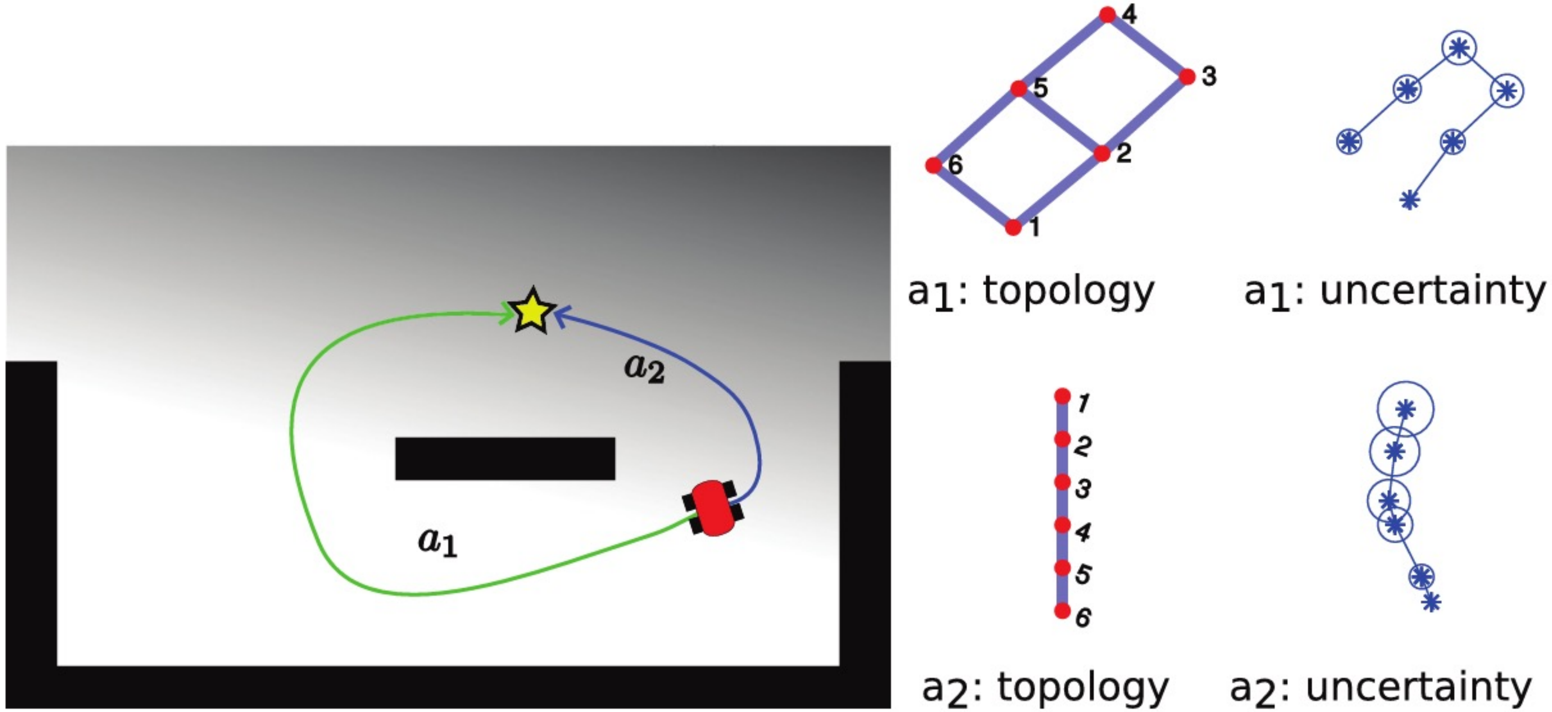

Belief Topology

Factor Graph

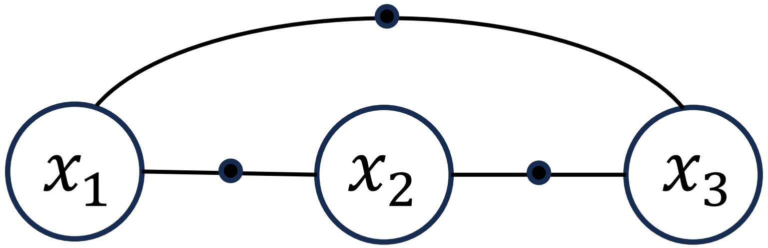

Topological Graph

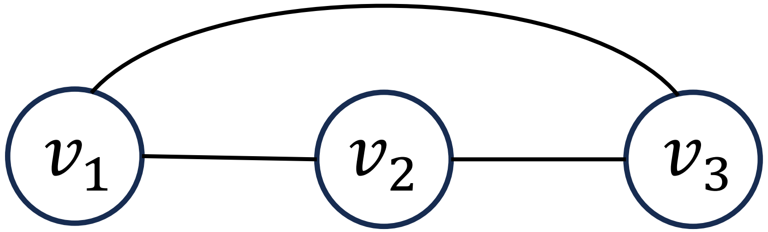

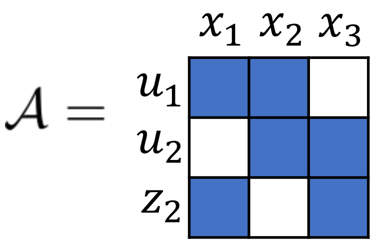

Measurement Jacobian matrix

$$\Lambda = \mathcal{A}^T \mathcal{A} = \Sigma^{-1}$$

Information matrix

$$\begin{array}{c|ccc} & v_1 & v_2 & v_3 \\ \hline v_1 & 2 & -1 & -1 \\ v_2 & -1 & 2 & -1 \\ v_3 & -1 & -1 & 2 \end{array}$$

Laplacian matrix

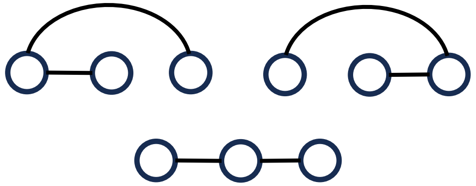

Kirchoff's theorem
spanning trees

$$\tau\left(G\right) = \frac{1}{n}\lambda_1\lambda_2\cdots\lambda_{n-1}$$

Belief Topology

- [Khosoussi et al. 2015 RSS] identified Laplacian structures in $\Lambda = \begin{bmatrix} L_{w_p} \otimes \cdots & \cdots \\ \vdots & L_{w_\theta} + \cdots \end{bmatrix}$
- Used the topological signature to bound the determinant of the information matrix

$$\tau_w(\mathcal{G}) \leq \log|\Lambda| \leq \tau_w(\mathcal{G}) + n \cdot \log(1 + \delta/\lambda_1)$$

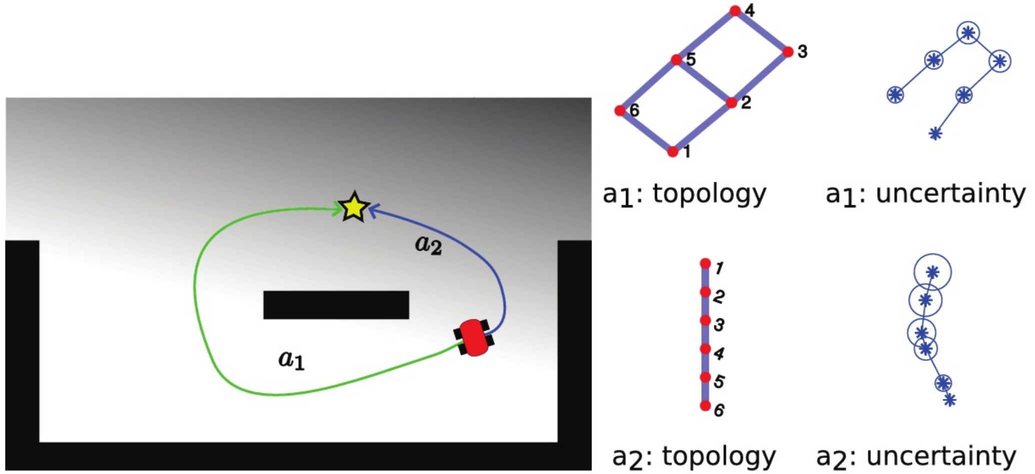

- [Kitanov and Indelman 2018 ICRA] first to extend to BSP problems

The *Focused* Case

- *unfocused* BSP - the objective function considers all variables

- *focused* BSP – the objective function prioritizes or only considers a predefined subset of focused variables

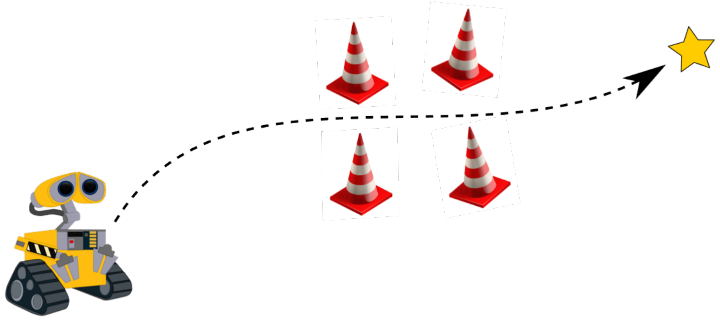

collision avoidance

focused reconstruction task

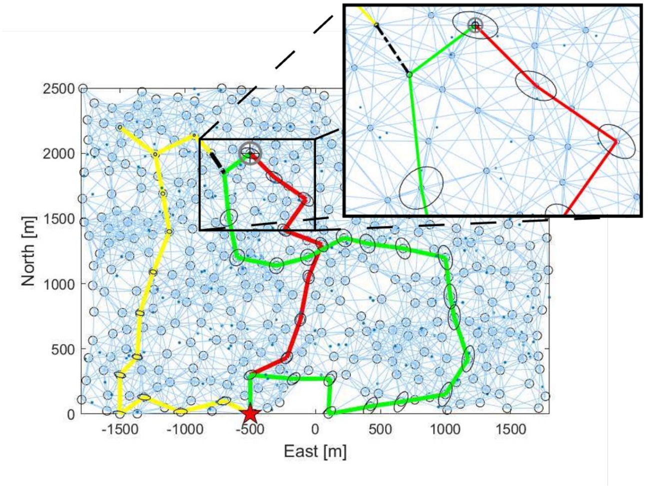

focused vs *unfocused*

- $X_k^F \subseteq X_k$ a *focused* subset of states $(X_k^F \bigcup X_k^U = X_k)$

Our Contributions

- The first to consider utilizing topological aspects in a *focused* BSP problem
- Derive two topological signatures to approximate a *focused* cost function
- Prove asymptotic convergence and develop bounds for one of the signatures

The *Focused* Objective Function

- Information theoretic cost – Differential entropy

- Objective function for the *focused* case, considering only the terminal marginal belief

$$J_{\mathcal{H}}^{F}\left(\mathcal{U}\right)=\mathcal{H}\left(b\left[X_{k+L}^{F}\right]\right)=\frac{n^{F}}{2}\log\left(2\pi e\right)-\frac{1}{2}\log\left|\Lambda_{k+L}^{M,F}\right|$$

- Problem : the set of *focused* can be very small with respect to the entire problem

→ calculating the marginal information matrix via expensive Schur complement operation

$$J_{\mathcal{H}}^{F}\left(\mathcal{U}\right)=\frac{n^{F}}{2}\log\left(2\pi e\right)-\frac{1}{2}\log\left|\Lambda_{k+L}\right|+\frac{1}{2}\log\left|\Lambda_{k+L}^{U}\right|$$

The *Unfocused* Augmented Graph

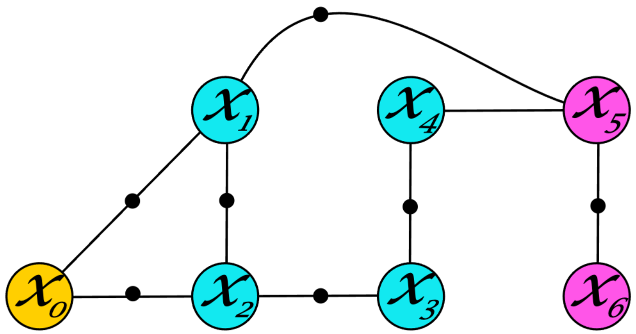

Factor Graph

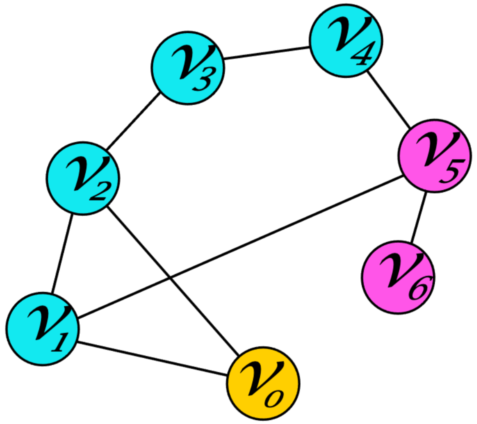

Topological Graph

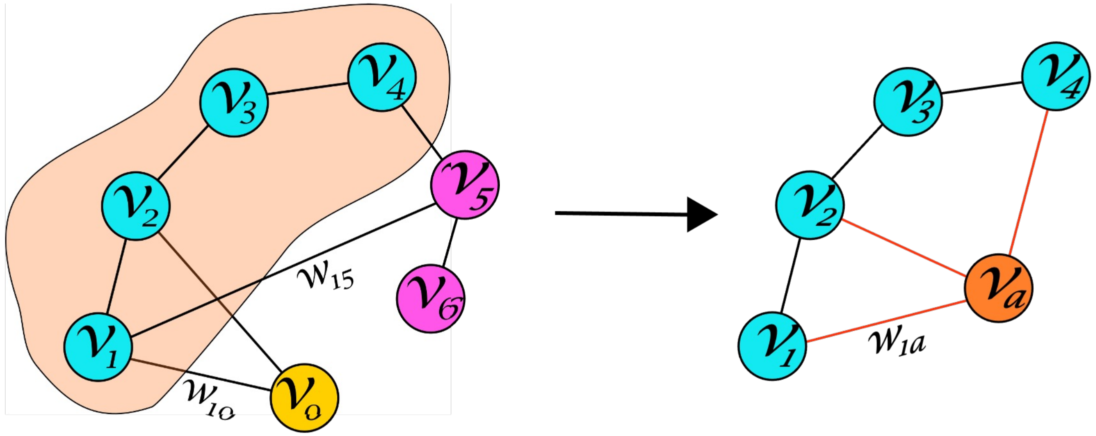

focused topological signatures

- Weighted Tree Connectivity Difference (WTCD) $\boxed{S_{WTCD} = \frac{n^F}{2}\log\left(2\pi e\right) - \frac{1}{2}\left[\tau_w - \tau_w^{U,A}\right]}$

- The approximation error $\epsilon(J_{\mathcal{H}}^F) \doteq J_{\mathcal{H}}^F - S_{WTCD}$ is bounded by $-\frac{n}{2}\log\left(1+\frac{\delta}{\lambda_1}\right) \leq \epsilon(J_{\mathcal{H}}^F) \leq \frac{n^U}{2}\log\left(1+\frac{\delta^U}{\lambda_1^U}\right)$

- Requires calculating the determinants of the associated Laplacian matrices
  Can we do better (computationally)?

- The Von Neumann Difference signature $\boxed{S_{VND} = \frac{n^F}{2}\log\left(2\pi e\right) - \frac{1}{2}\left[h_w - h_w^{U,A}\right]}$

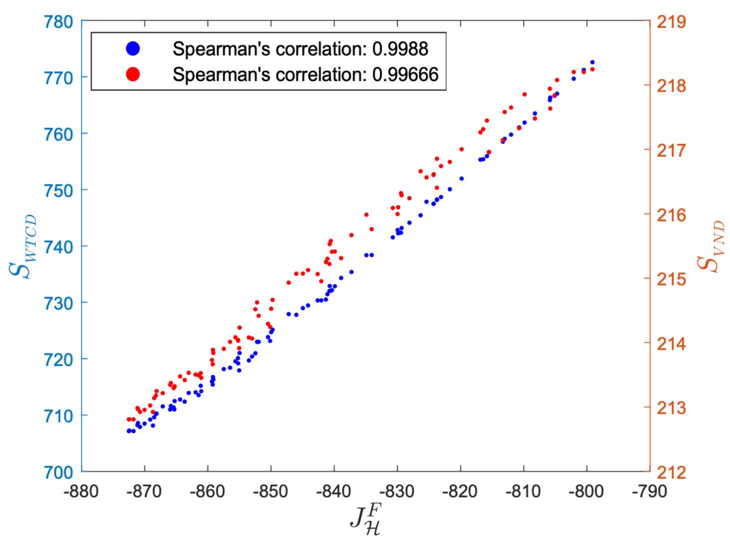

Measurement Selection

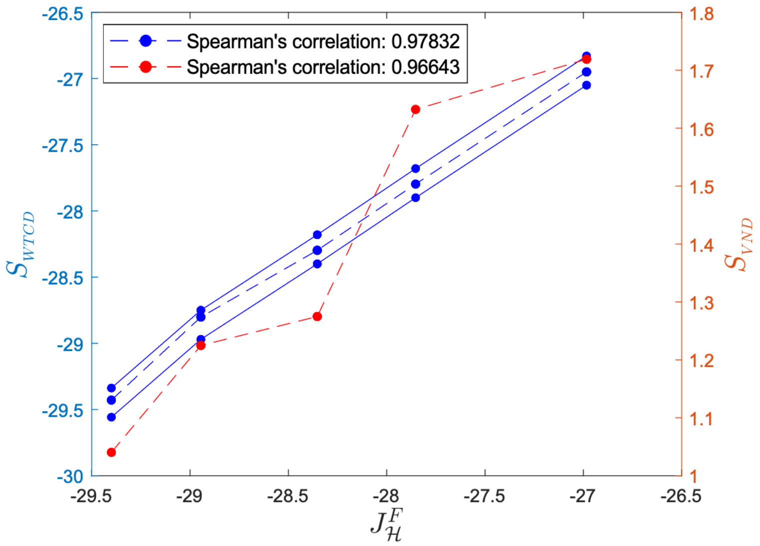

Active 2D Pose SLAM

signature	measurement selection	active SLAM
S_{WTCD}	18.88	1.21
S_{VND}	**12.02**	**0.14**
$J_{\mathcal{H}}^{F}$	146.24	6.34

Average running time experiments in *ms*

- M. Shienman, A. Kitanov and V. Indelman [2021 IEEE RA-L]
  "Focused Topological Belief Space Planning"

- M. Shienman and V. Indelman [2022 ICRA] ***Outstanding Paper Award Finalist***
  "Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints"

- M. Shienman and V. Indelman [2022 ISRR]
  "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints"

- M. Shienman, O. Levy Or, M. Kaess and V. Indelman [2024 IROS - Submitted]
  "A Slices Perspective for Incremental Nonparametric Inference in High Dimensional State Spaces"

Motivation

- Ambiguity - when a certain observation has more than one possible interpretation

Slip and Grip

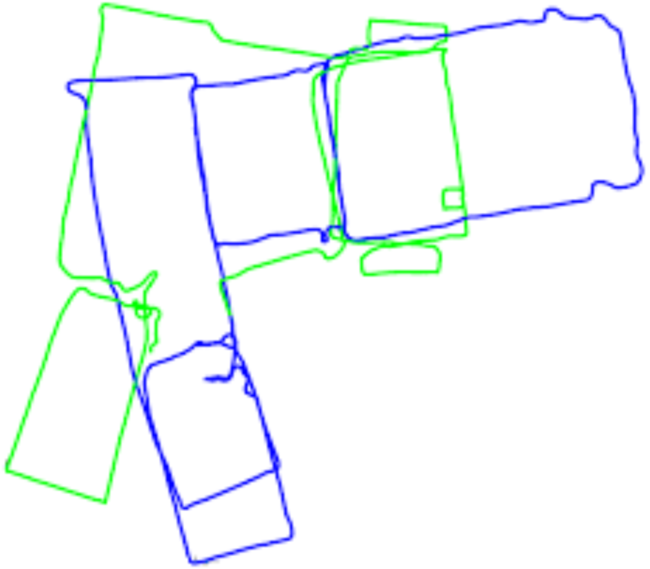

loop closures

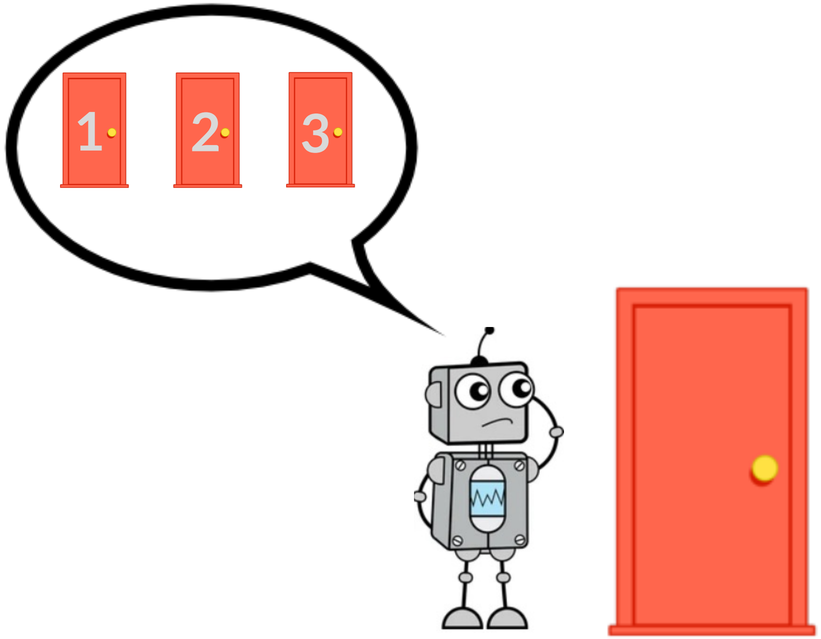

unresolved data associations

→ Multi-modal distributions

Motivation

- Number of hypotheses grows exponentially (in both inference and planning)

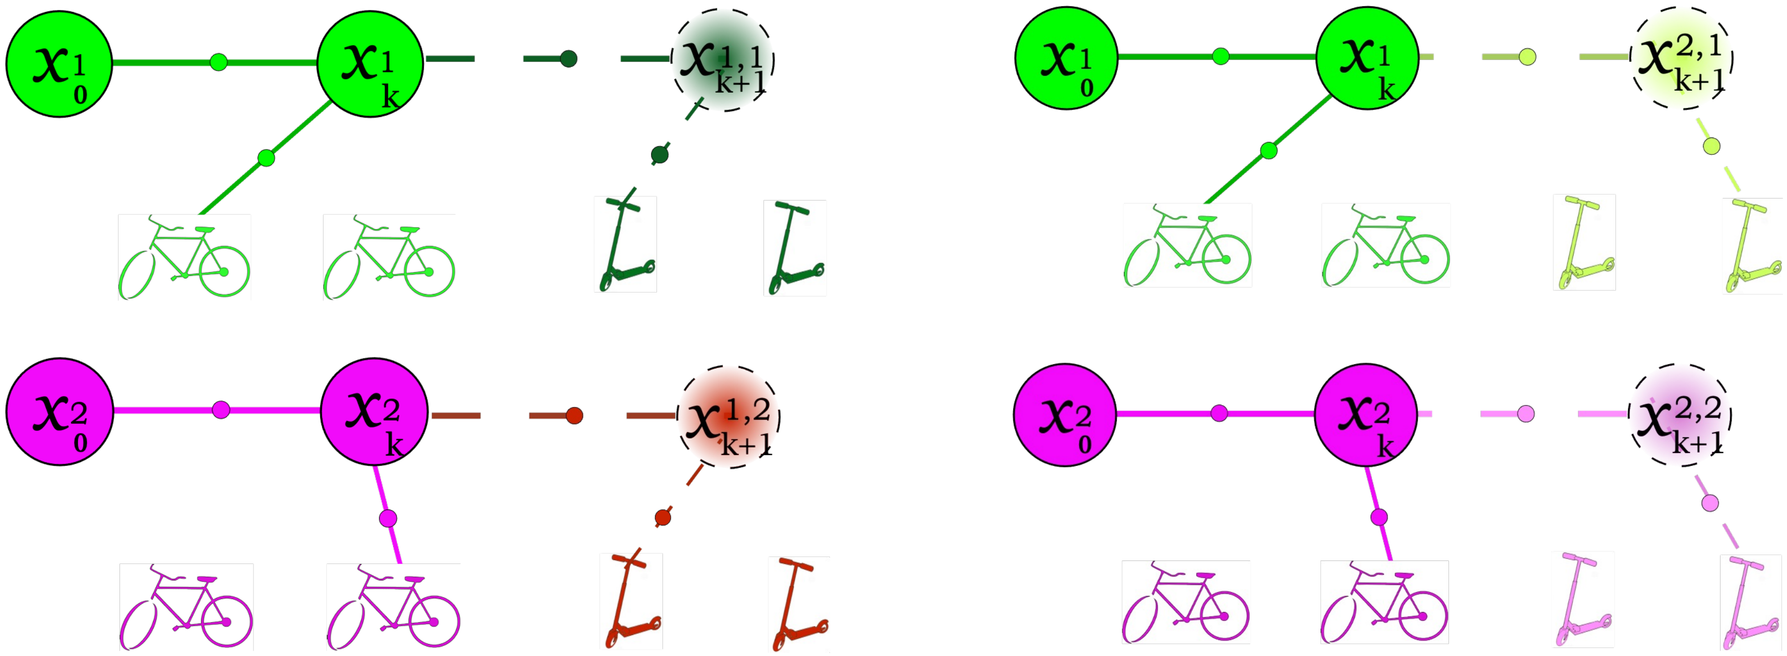

Motivation

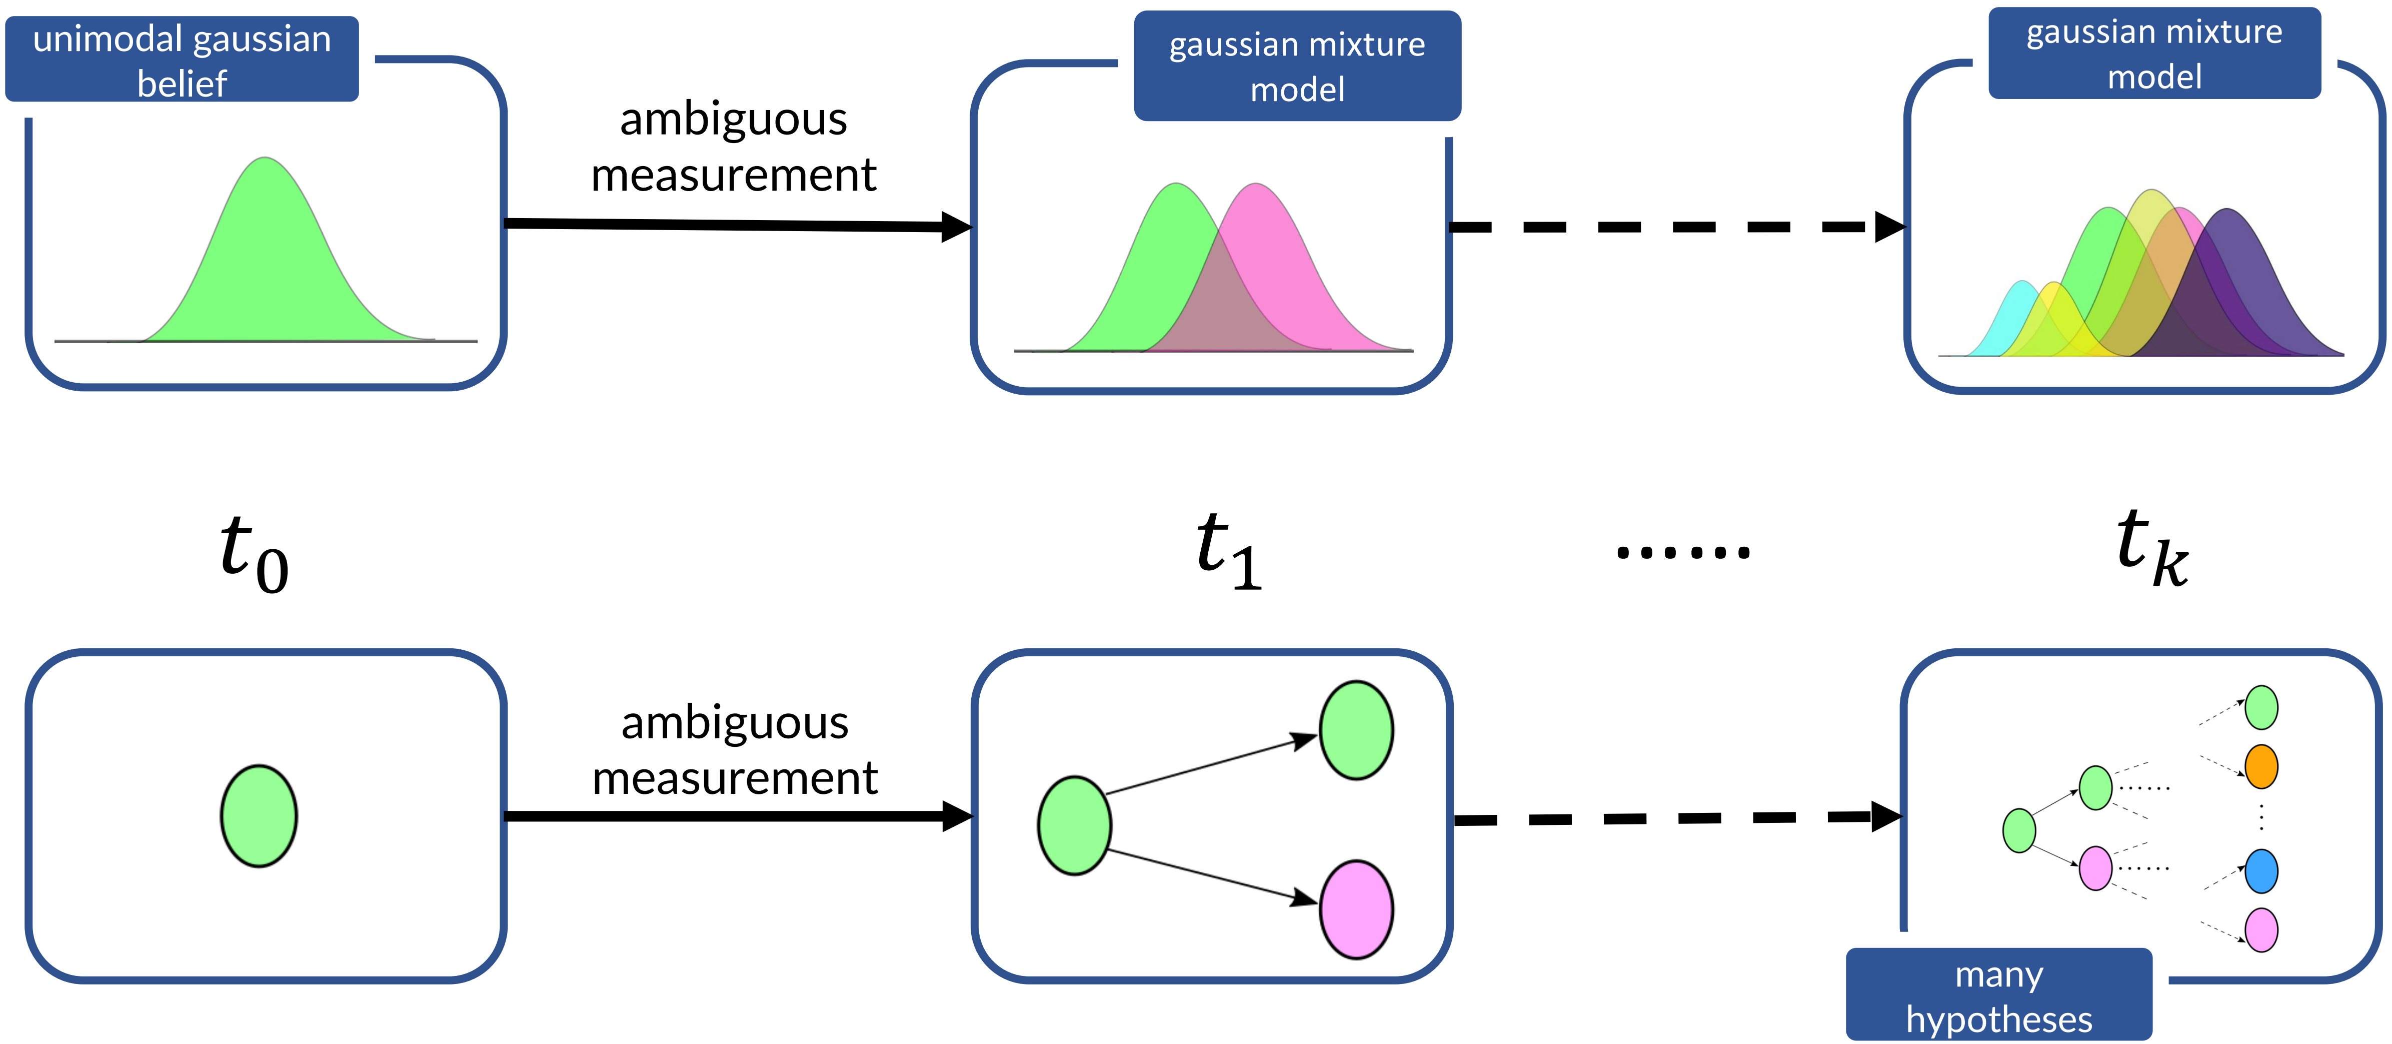

unimodal gaussian belief
ambiguous measurement
gaussian mixture model
gaussian mixture model
t_0
t_1
……
t_k
ambiguous measurement
many hypotheses

Problem Formulation

- The belief at time k is over both discrete and continues random variables, expressed as a linear combination

$$b_k = \sum_{j=1}^{M_k} \underbrace{\mathbb{P}\left(x_k | \beta^j_{1:k}, H_k\right)}_{b^j_k} \underbrace{\mathbb{P}\left(\beta^j_{1:k} | H_k\right)}_{w^j_k}$$

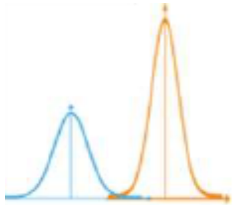

- $\beta_k \in \mathbb{N}^{n_k}$ data association realization vector at time k

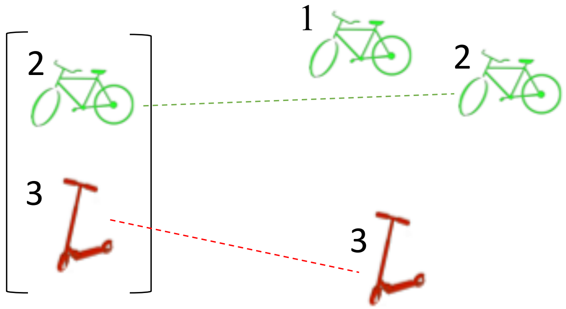

Our Contributions

- Utilize a distilled subset of hypotheses in planning to reduce computational complexity

- Develop a connection between our approach and the true analytical solution, owing to every possible data association, for the myopic case

- Derive bounds over the true analytical solution and prove they convergence

- Address the challenging setting of hard budget constraints, and show, for the first time, how these bounds provide performance guarantees

A Simplified Belief

- Use only a distilled subset of hypotheses $M_k^s \subseteq M_k$ from time k based on weights w_k^j

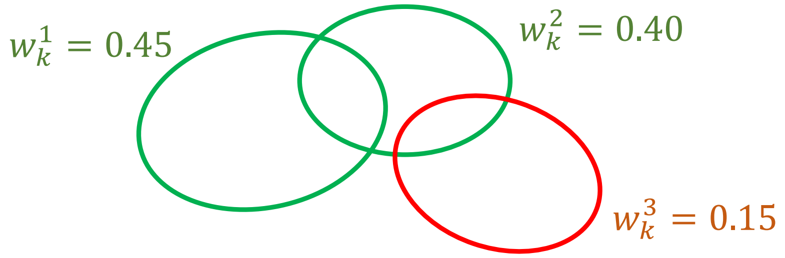

- A simplified belief is formally defined as $b_k^s \triangleq \sum_{j=1}^{M_k^s} w_k^{s,j} b_k^j \quad , \quad w_k^{s,j} \triangleq \frac{w_k^j}{w_k^{m,s}},$

Cost Function

- Information theoretic cost function

$$\mathcal{H}(x,\beta) = -\int \sum \mathbb{P}(x,\beta)\, log\mathbb{P}(x,\beta) =$$

$$-\int \sum \mathbb{P}(x|\beta)\,\mathbb{P}(\beta)\, log\mathbb{P}(\beta) - \int \sum \mathbb{P}(\beta)\,\mathbb{P}(x|\beta)\, log\mathbb{P}(x|\beta) = \mathcal{H}(\beta) + \underset{\beta\sim\mathbb{P}(\beta)}{\mathbb{E}} \mathcal{H}(x|\beta)$$

Entropy over posterior weights
(hypotheses disambiguation)

focus on this term

Conditional entropy over x
(uncertainty of each hypothesis)

Objective Function

- A myopic setting $J\left(b_k, u_k\right) = \int\limits_{Z_{k+1}} \eta_{k+1} c\left(b_{k+1}\right) dZ_{k+1} \quad , \qquad \eta_{k+1} \triangleq \mathbb{P}(Z_{k+1} | H_{k+1}^{-})$

- For performance guarantees, we bound the objective function for each candidate action

$$\int\limits_{Z_{k+1}} \mathcal{LB}\left[\eta\right] \mathcal{LB}\left[c\left(b_{k+1}\right)\right] dZ_{k+1} \leq J\left(b_k, u_k\right) \leq \int\limits_{Z_{k+1}} \mathcal{UB}\left[\eta\right] \mathcal{UB}\left[c\left(b_{k+1}\right)\right] dZ_{k+1}$$

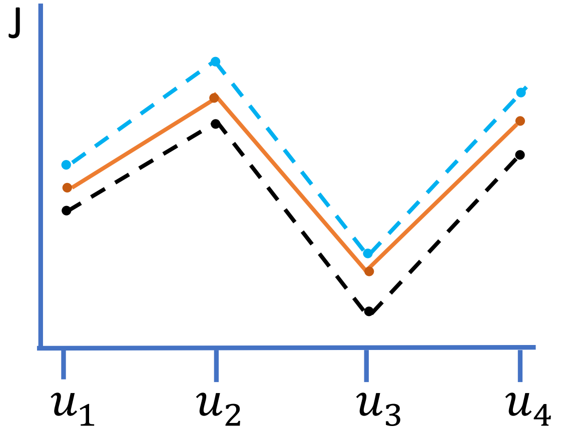

$\mathcal{LB}\left[J\left(b_k, u_k\right)\right] \leq J\left(b_k, u_k\right) \leq \mathcal{UB}\left[J\left(b_k, u_k\right)\right]$

Budget Free Scenario

- Goal : fully disambiguate between all prior hypotheses

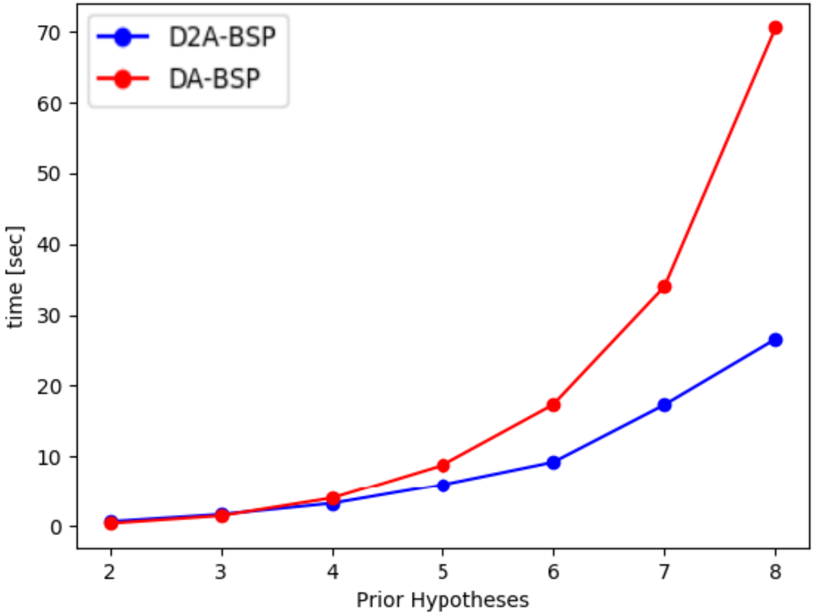

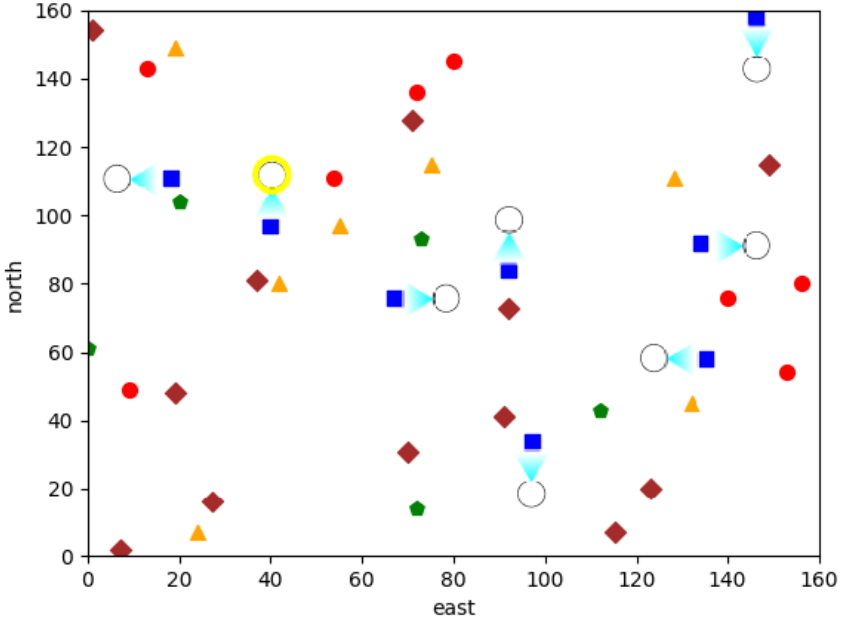

Budget Constrained Scenario

- Goal : disambiguate between hypotheses (weighted equally) under budget constraints
- Budget constraint : agent can only use **one** hypothesis in planning

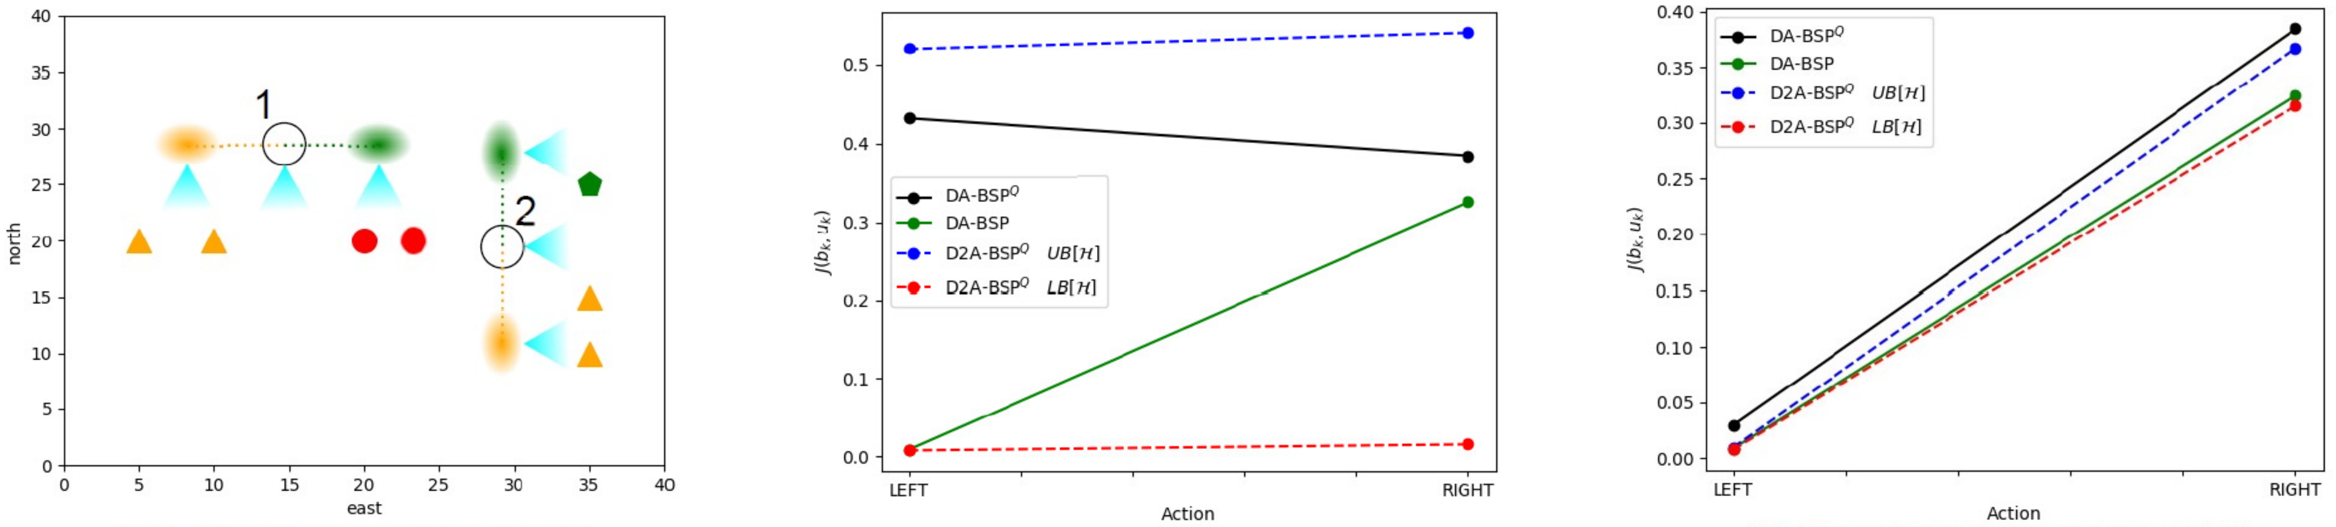

(a) BSP with 2 prior hypotheses (b) Planning using component #1 (c) Planning using component #2

Our Contributions - Recap

D2A-BSP — A novel planning approach that utilizes only a distilled subset of hypotheses in a myopic setting

Bounds — Over the true analytical solution considering all possible hypotheses

Budget Free — Use bounds to reduce computational complexity while preserving action selection

Budget Constraints — Use bounds to provide performance guarantees

- M. Shienman, A. Kitanov and V. Indelman [2021 IEEE RA-L]
"Focused Topological Belief Space Planning"

- M. Shienman and V. Indelman [2022 ICRA] ***Outstanding Paper Award Finalist***
"Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints"

- M. Shienman and V. Indelman [2022 ISRR]
"Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints"

- M. Shienman, O. Levy Or, M. Kaess and V. Indelman [2024 IROS - Submitted]
"A Slices Perspective for Incremental Nonparametric Inference in High Dimensional State Spaces"

A Non-Myopic Setting

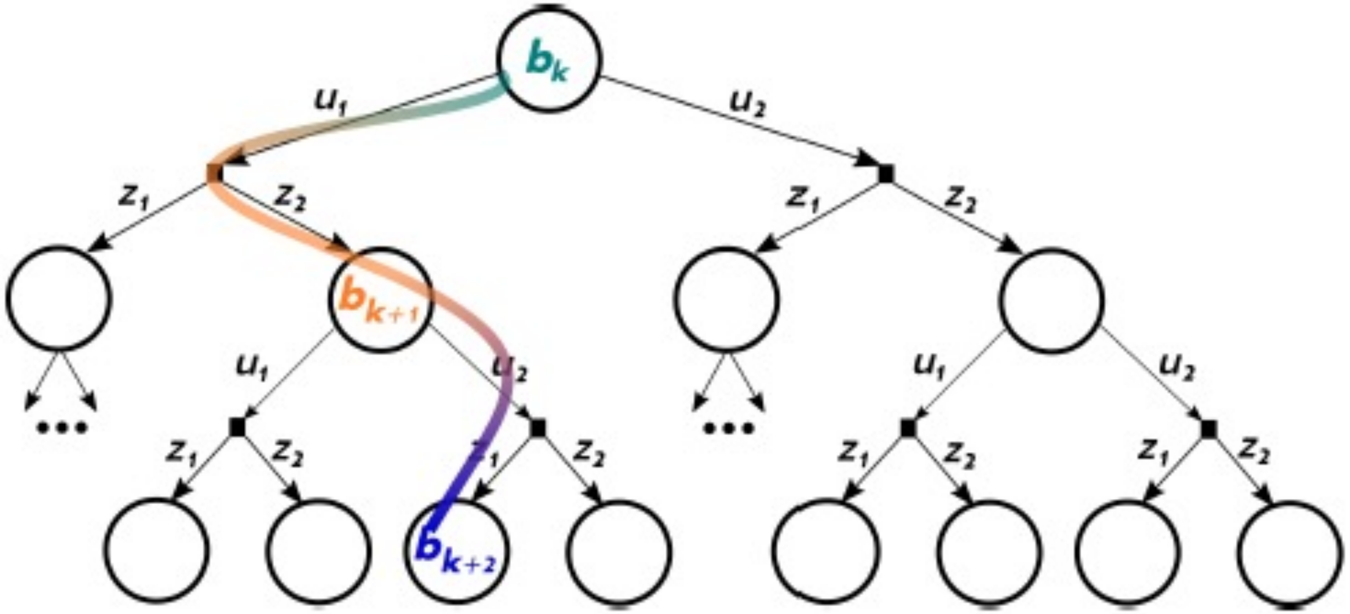

planning tree

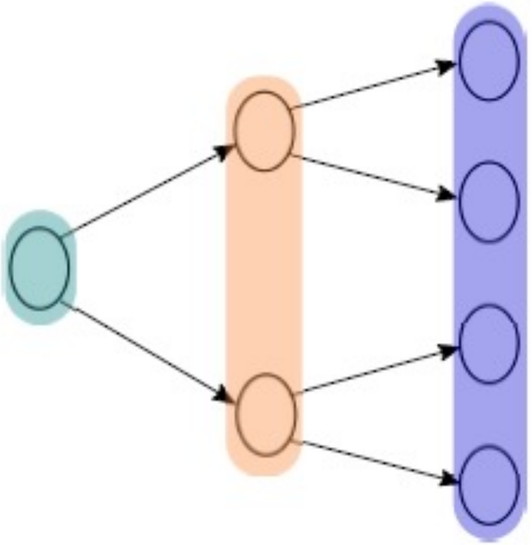

exponential growth
of hypotheses

How To Construct a Planning Tree?

- Problem: under budget constraints, simplifying a belief in a specific tree node affects all children node

- Should consider the implications of simplification in both inference and planning

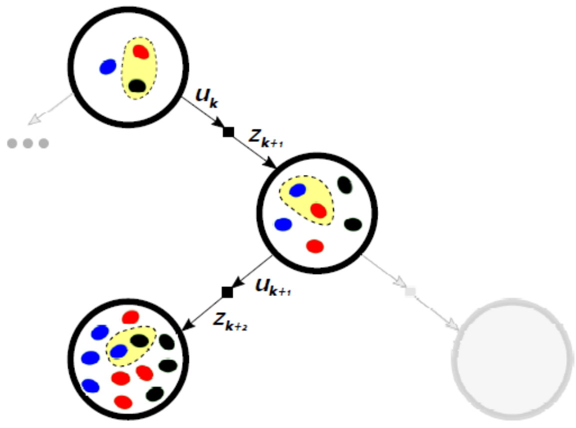

Inference: no budget constraints
Planning: under budget constraints

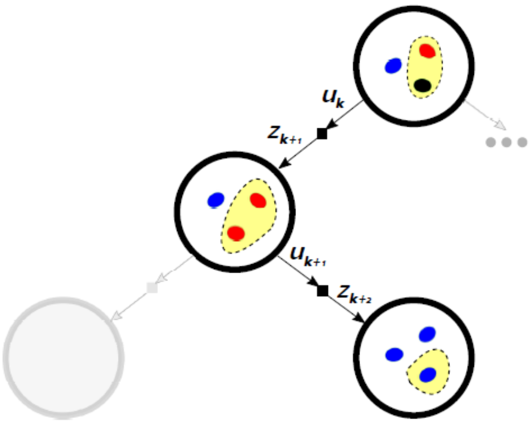

Inference: under budget constraints
Planning: no budget constraints

Our Contributions

- Extend previous work to a non-myopic setting
- Derive bounds for the true analytical solution based on simplified beliefs and prove they convergence
- Thoroughly study, for the first time, the impacts of hard budget constraints in either planning and/or inference

A Simplified Belief

- For the nonmyopic case $b^{s}_{k+n} \triangleq \sum_{r \in M^{s}_{k+n}} w^{s,r}_{k+n} b^{r}_{k+n} \quad , \quad w^{s,r}_{k+n} \triangleq \frac{w^{r}_{k+n}}{w^{m,s}_{k+n}},$

- Cost function : entropy over posterior weights (hypotheses disambiguation)

- For each belief tree node, representing a belief b_{k+n} with components M_{k+n} and a subset $M^{s}_{k+n} \subseteq M_{k+n}$ the cost can be expressed as

$$\mathcal{H}_{k+n} = \frac{w^{m,s}_{k+n}}{\eta_{k+n}} \left[\mathcal{H}^{s}_{k+n} + log \left(\frac{\eta_{k+n}}{w^{m,s}_{k+n}} \right) \right] - \sum_{r \in \neg M^{s}_{k+n}} \frac{w^{r}_{k+n}}{\eta_{k+n}} log \left(\frac{w^{r}_{k+n}}{\eta_{k+n}} \right)$$

Bounds In The Non-Myopic Case

- Cost function : entropy over posterior weights (hypotheses disambiguation)

- The cost term in each belief tree node is bounded by

$$\mathcal{LB}\left[\mathcal{H}_{k+n}\right] = \frac{w_{k+n}^{m,s}}{\mathcal{UB}\left[\eta_{k+n}\right]}\left[\mathcal{H}_{k+n}^{s} + log\left(\frac{\mathcal{LB}\left[\eta_{k+n}\right]}{w_{k+n}^{m,s}}\right)\right],$$

$$\mathcal{UB}\left[\mathcal{H}_{k+n}\right] = \frac{w_{k+n}^{m,s}}{\mathcal{LB}\left[\eta_{k+n}\right]}\left[\mathcal{H}_{k+n}^{s} + log\left(\frac{\mathcal{UB}\left[\eta_{k+n}\right]}{w_{k+n}^{m,s}}\right)\right] - \bar{\gamma} log\left(\frac{\bar{\gamma}}{\left|\neg M_{k+n}\right|}\right)$$

- Convergence $\lim_{M_{k+n}^{s} \to M_{k+n}} \mathcal{LB}\left[\mathcal{H}_{k+n}\right] = \mathcal{H}_{k+n} = \mathcal{UB}\left[\mathcal{H}_{k+n}\right]$

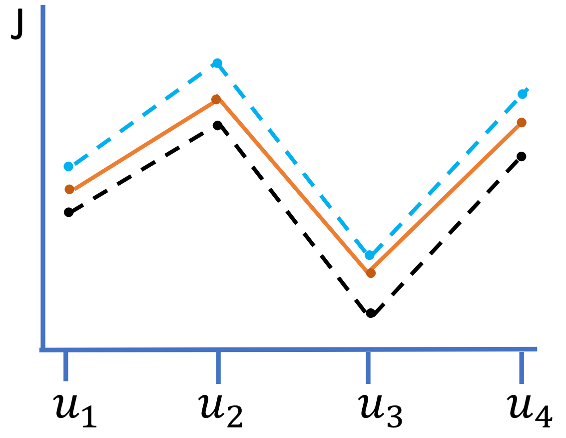

$$\mathcal{LB}\left[J\left(b_k, u_k\right)\right] \leq J\left(b_k, u_k\right) \leq \mathcal{UB}\left[J\left(b_k, u_k\right)\right]$$

Kidnapped Robot Scenario – No Budget Constraints

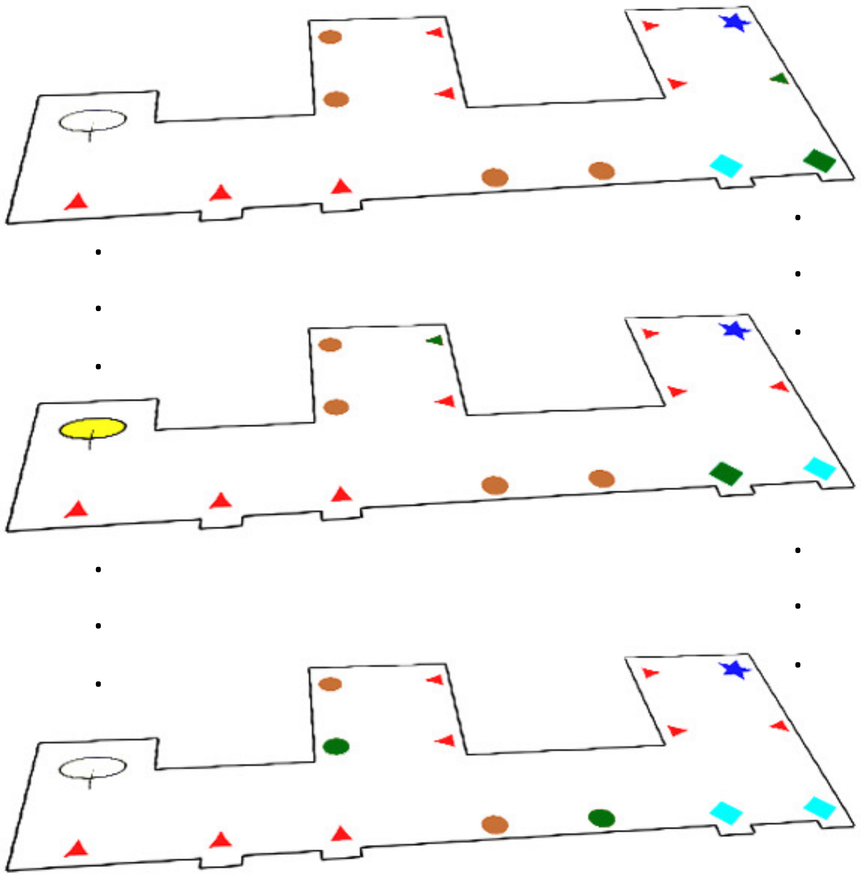

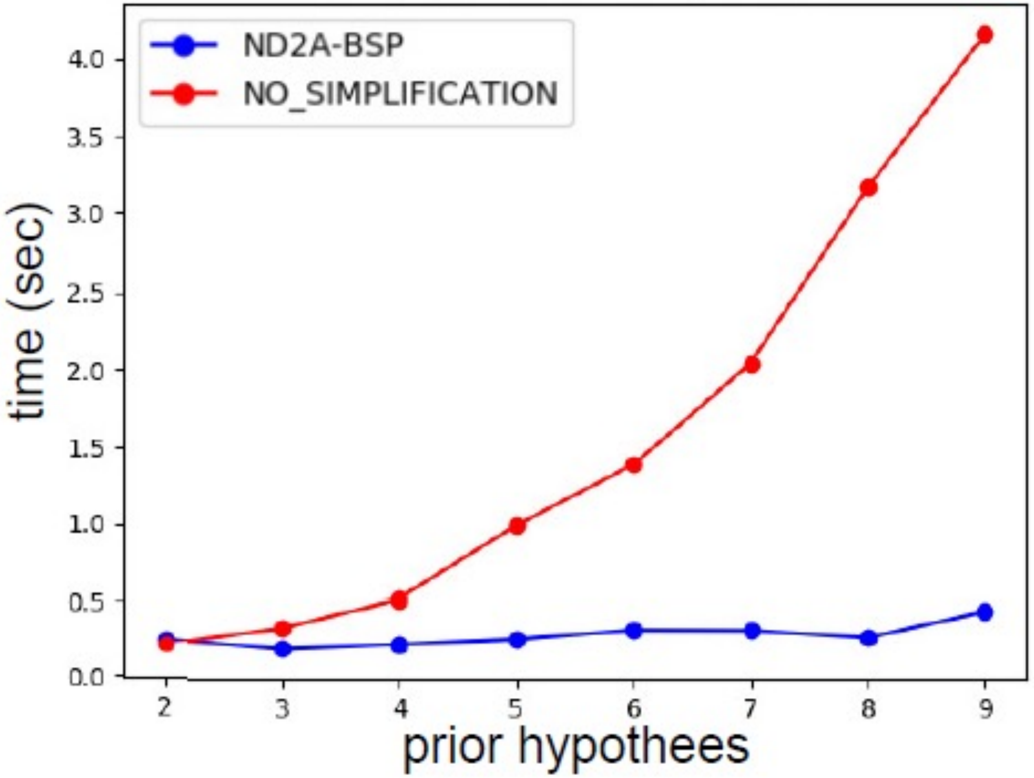

Under Budget Constraints in Planning

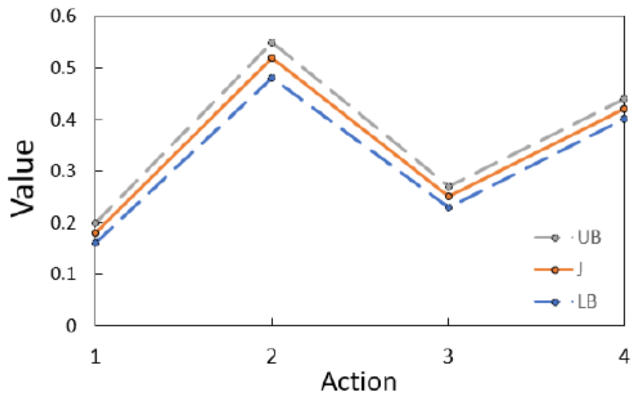

(a) No overlap between bounds

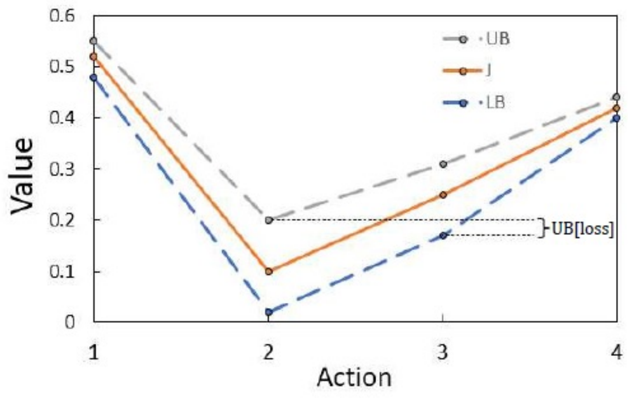

(b) Bounds overlap

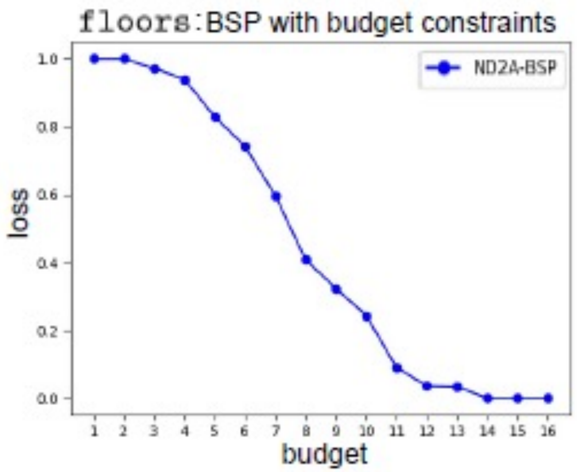

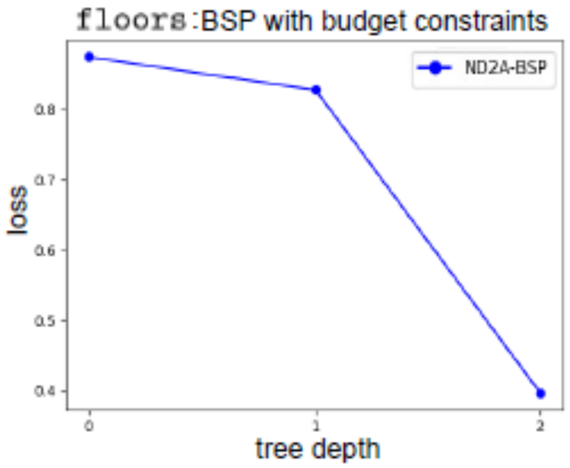

- M. Shienman, A. Kitanov and V. Indelman [2021 IEEE RA-L]
  "Focused Topological Belief Space Planning"

- M. Shienman and V. Indelman [2022 ICRA] ***Outstanding Paper Award Finalist***
  "Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints"

- M. Shienman and V. Indelman [2022 ISRR]
  "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints"

- M. Shienman, O. Levy Or, M. Kaess and V. Indelman [2024 IROS - Submitted]
  "A Slices Perspective for Incremental Nonparametric Inference in High Dimensional State Spaces"

Motivation

- In real-world problems, the posterior distribution is often non-Gaussian, having multiple modes or a nonparametric structure

- Due to the complex, non-Gaussian nature of such posterior distributions, obtaining closed-form analytical solutions is challenging and frequently impractical

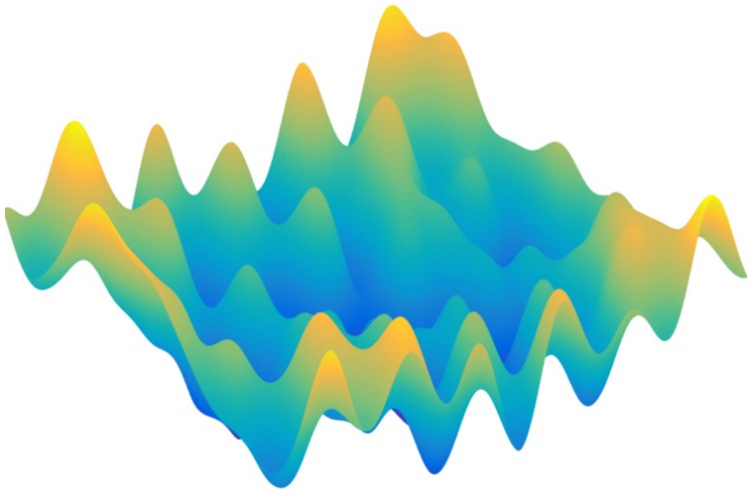

- Notice : models are still assumed to be given (but not Gaussians..)

Background - The Gaussian Case

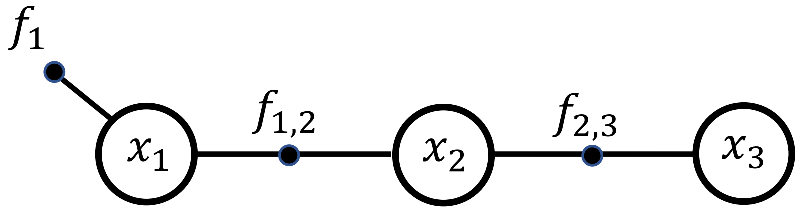

- Solved using the *forward-backward* algorithm

- *forward* pass: in each step a single variable is eliminated from the graph

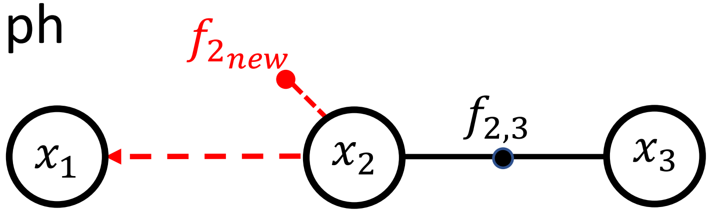

$$f_{2_{new}}(x_2) = p(x_2; f_1, f_{1,2}) = \underset{x_1 \sim f_1}{\mathbb{E}}[f_{1,2}] = \int f_1(x_1) f_{1,2}(x_1, x_2) dx_1$$

- Once the *forward* pass is completed, the joint distribution is expressed via conditionals

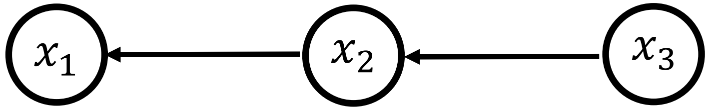

$$p(x_1, x_2, x_3) = p(x_3)\, p(x_2|x_3)\, p(x_1|x_2)$$

- The marginal distributions are calculated using backsubstitution

Non-Parametric High Dimensional Settings – Previous Works

- [Fourie et al. IROS 2016]

Using intermediate reconstructions with KDE

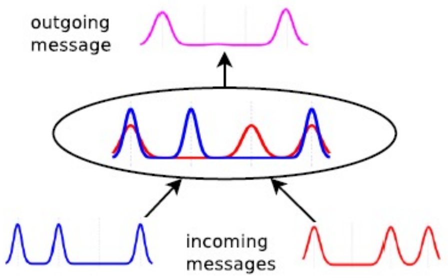

Fig. 1. Illustration of a Bayesian clique operation as part of a larger multi-modal belief propagation on a Bayes tree. Two incoming messages are *combined* with local potentials to produce one outgoing message during the upward pass procedure towards the root. Multi-modality is allowed to exist amongst cliques, rather than selecting a single mode as a maximum-product type algorithm would.

- [Huang et al. ICRA 2021]

Using intermediate reconstruction with learned transformations

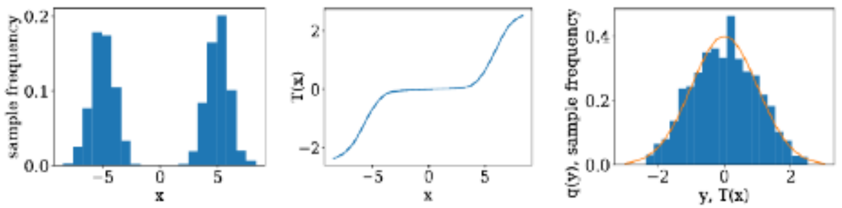

Fig. 2. A one-dimensional example of normalizing flow: histogram of sample x (left), transformation function $T(x)$ (middle), and histogram of transformed samples and reference variable $y \sim N(0,1)$ (right).

Key Observation

- Use a *slices* perspective and avoid intermediate reconstructions

joint: $\mathbb{P}(X,Y) = \mathbb{P}(X|Y)\cdot\mathbb{P}(Y)$

marginal: $\mathbb{P}(X) = \int_Y \mathbb{P}(X|Y)\cdot\mathbb{P}(Y)\,dY$

estimated marginal:

$$\hat{\mathbb{P}}(X) = \underset{y\sim\mathbb{P}(Y)}{\hat{\mathbb{E}}}[\mathbb{P}(X|Y=y)] = \frac{1}{N}\sum_{i=1}^{N}\mathbb{P}(X|Y=y^i)$$

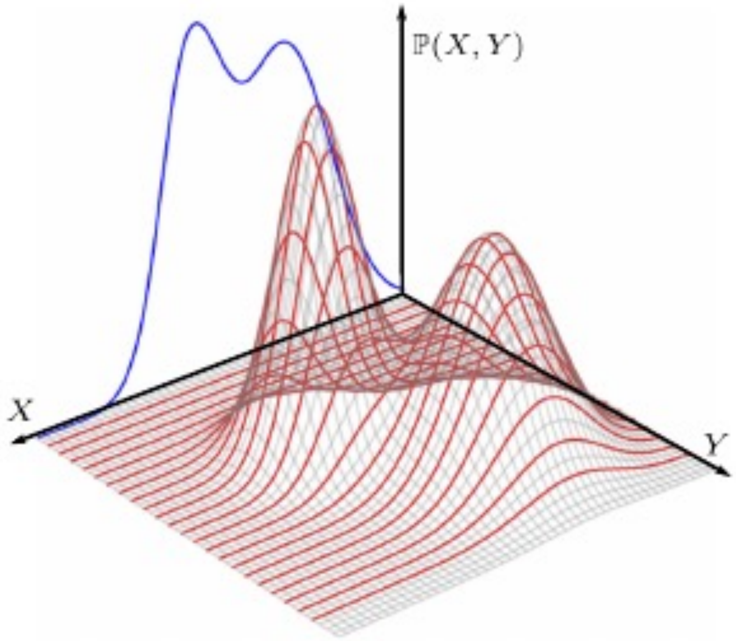

Our Contributions

- Leverage *slices* from high-dimensional surfaces to approximate joint and marginal posterior distributions without any further intermediate reconstructions

- A novel early stopping heuristic criteria (*backward* pass) to further speed up calculations

- Requires less samples and consistently outperforms state-of-the-art nonparametric inference algorithms in terms of accuracy and computational complexity

Slices For Non-Parametric Inference

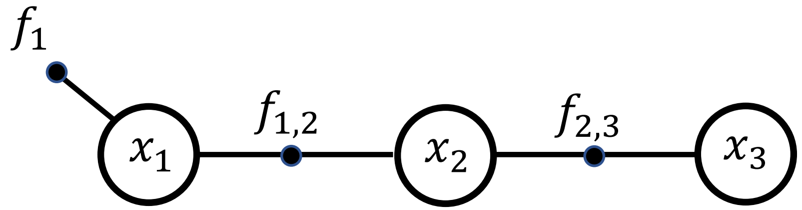

- Follow the *forward-backward* approach

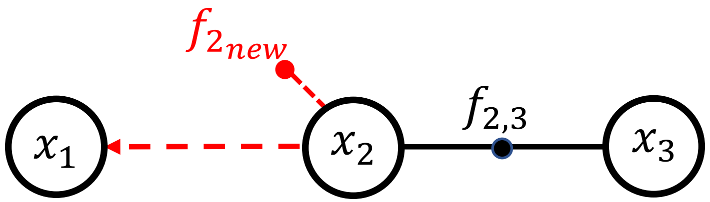

- *Forward* pass $f_{new}(x_2) = \eta^{-1} \int_{x_1} f_1(x_1) f_{1,2}(x_1, x_2)\, dx_1$

$$\hat{f}_{new}(x_2) = \frac{\eta^{-1}}{N} \sum_{n=1}^{N} f_{1,2}(x_1^n, x_2)$$

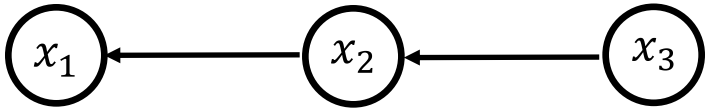

- *Backward* pass $\hat{\mathbb{P}}(x_2) = \frac{1}{N} \sum_{n=1}^{N} \hat{\mathbb{P}}(x_2 | x_3^n)$

Early stopping heuristic

- Incremental settings - perform inference whenever new data is present

- Early stopping heuristic in the Gaussian case based on variables estimate change [iSAM2 , M. Kaess et al. IJRR 2012]

- For general distributions we propose Maximum Mean Discrepancy (MMD)

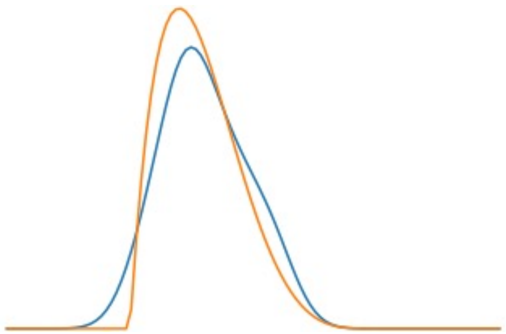

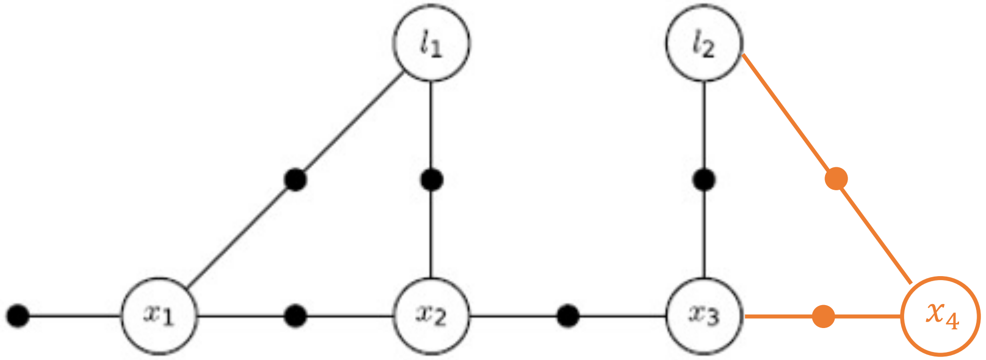

Results

- Plaza

a real-world dataset with range measurements

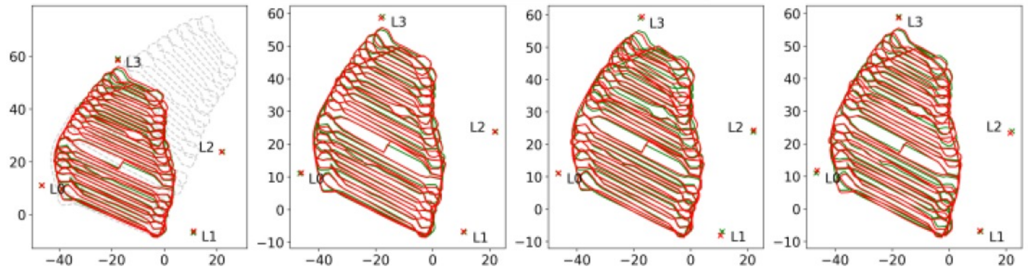

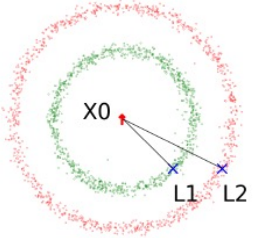

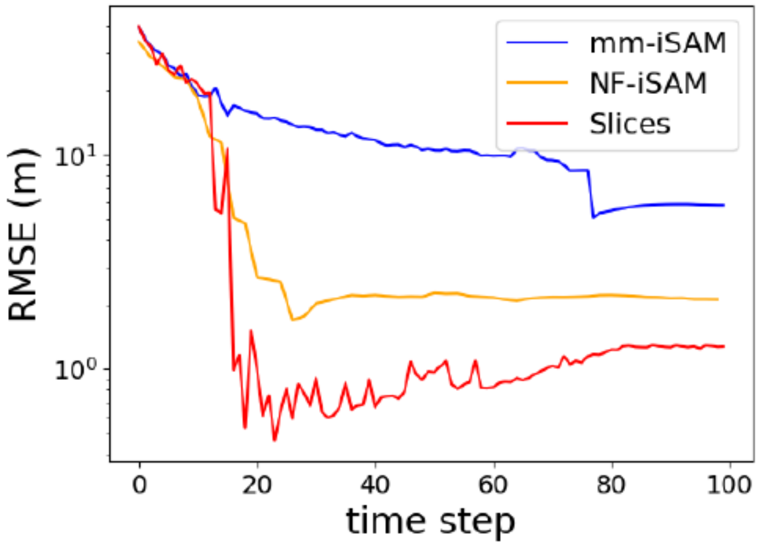

(a)

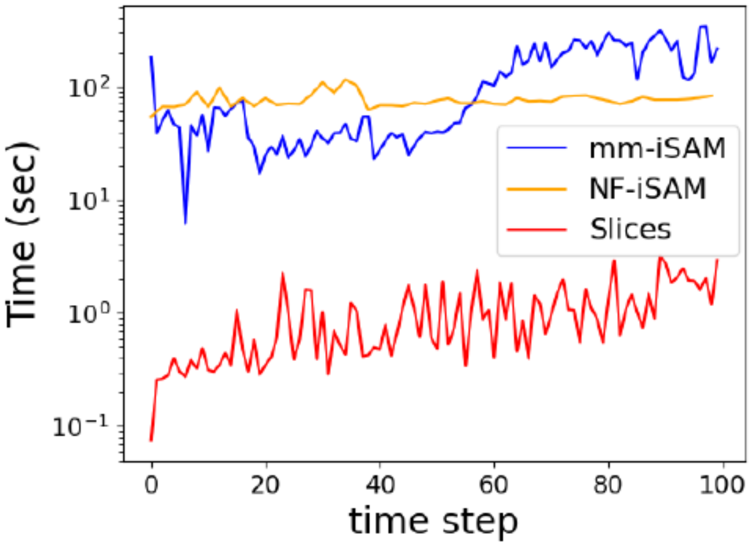

(b)

Summary

- Leverage structures in both inference and BSP problems

- Reduce computational complexity by solving simplified problems to handle real world scenarios under budget constraints

- Providing performance guarantees (in planning) by bounding the error between the original (computationally hard) problem and the simplified problem

Thank you!