Towards Scalable Online Decision Making Under Uncertainty in Partially Observable Environments

Vadim Indelman

TECHNION Israel Institute of Technology

ANPL Autonomous Navigation and Perception Lab

December 29, 2025

Advanced Autonomy

Involves autonomous navigation, active SLAM, informative gathering, active sensing, etc.

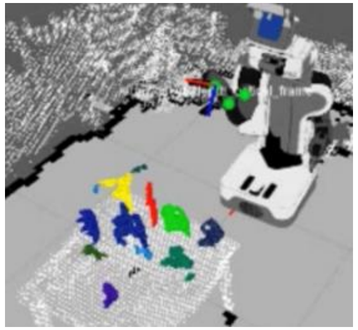

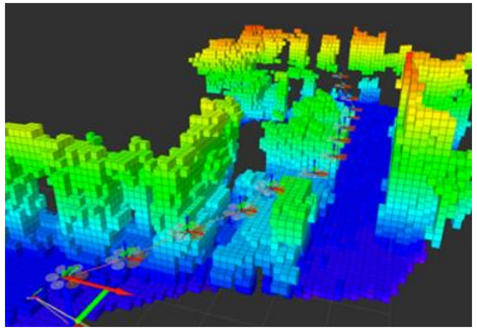

Advanced Autonomy

Perception and Inference

Where am I? What is the surrounding environment?

Key required capabilities

Decision-Making Under Uncertainty

What should I be doing next?

Determine best action(s) to accomplish a task, account for different sources of uncertainty

Perception and Inference ⟷ **Decision-Making Under Uncertainty**

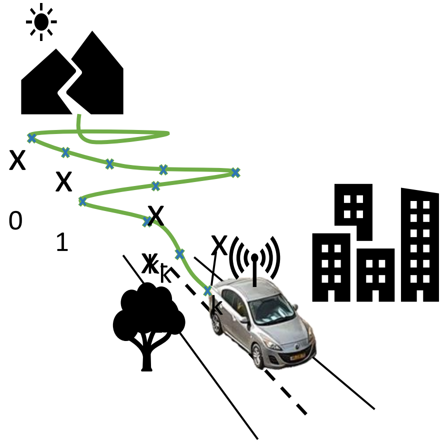

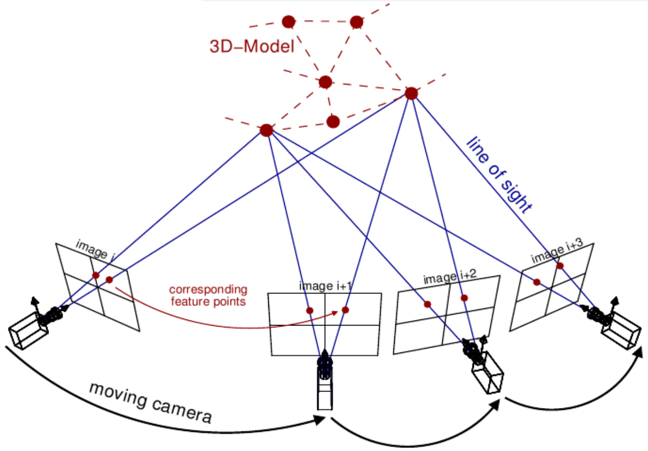

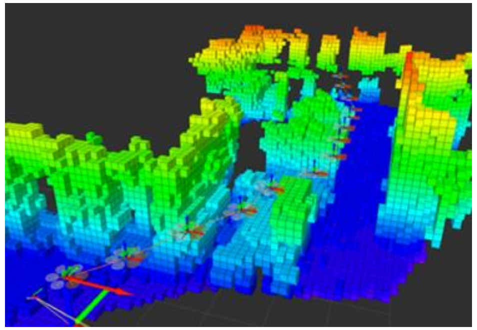

Perception and Inference

- Posterior belief at time instant k: $b_k \triangleq b[X_k] = \mathbb{P}(X_k \mid a_{0:k-1}, z_{1:k})$

state at time instant k

actions

observations

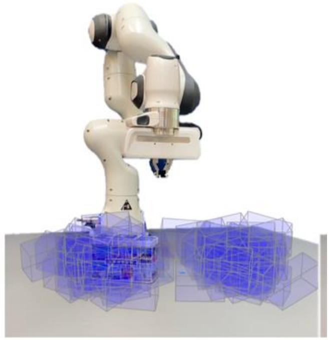

$$X_k \doteq \{x_0, \ldots, x_k, L_k\}$$

Past & current robot states

Environment representation, e.g. landmarks

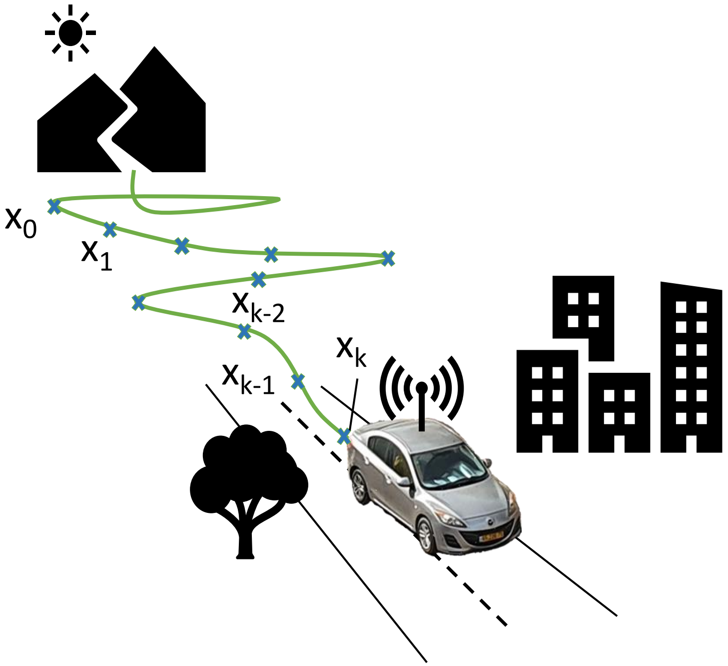

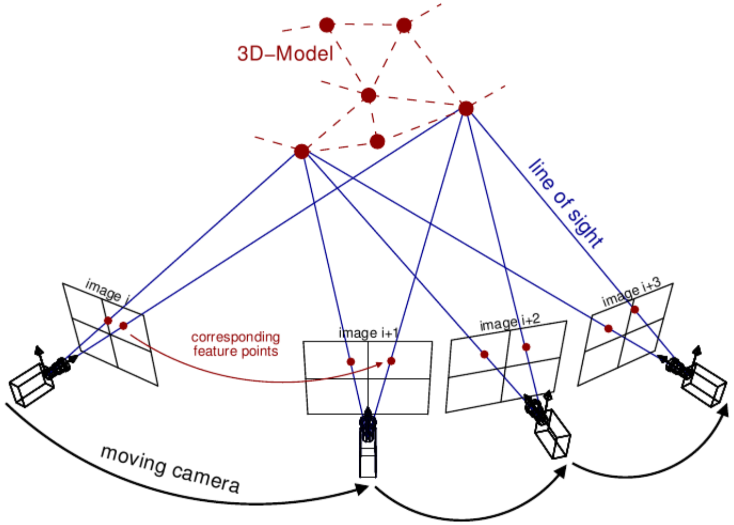

Can be represented with graphical models, e.g. a Factor Graph

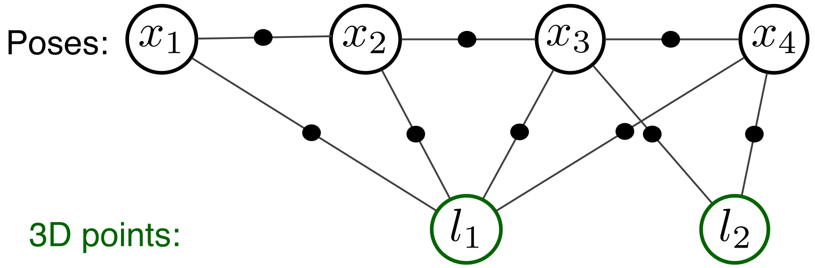

Partially Observable Markov Decision Process (POMDP)

- POMDP tuple:

$$\langle \mathcal{X}, \mathcal{Z}, \mathcal{A}, T, O, \rho, b_k \rangle$$

State, observation, and action spaces

Transition and observation models

Belief-dependent reward function

Belief at planning time instant k

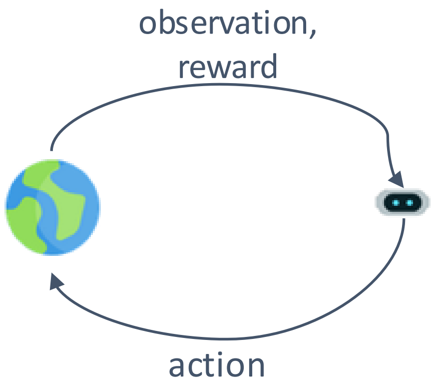

- Value function

$$V^{\pi}(b_k) = \mathbb{E}_{z_{k+1:k+L}|b_k,\pi}[\sum_{\ell=0}^{L} \rho(b_{k+\ell}, \pi_{k+\ell}(b_{k+\ell}))]$$

Belief-dependent reward function

$$Q^{\pi}(b_k, a_k) = \rho(b_k, a_k) + \mathbb{E}_{z_{k+1}|b_k,a_k}[V^{\pi}(b_{k+1})]$$

Partially Observable Markov Decision Process (POMDP)

- Value function $V^{\pi}(b_k) = \mathbb{E}_{z_{k+1:k+L}|b_k,\pi}[\sum_{\ell=0}^{L} \rho(b_{k+\ell}, \pi_{k+\ell}(b_{k+\ell}))]$

Belief-dependent reward function

- Belief at the ℓ-th look-ahead step: $b_{k+\ell} \triangleq b[X_{k+\ell}] = \mathbb{P}(X_{k+\ell} \mid a_{0:k-1}, z_{0:k}, a_{k:k+\ell-1}, z_{k+1:k+\ell})$

State at the *l*-th look ahead step

Past actions & observations

Future actions & observations

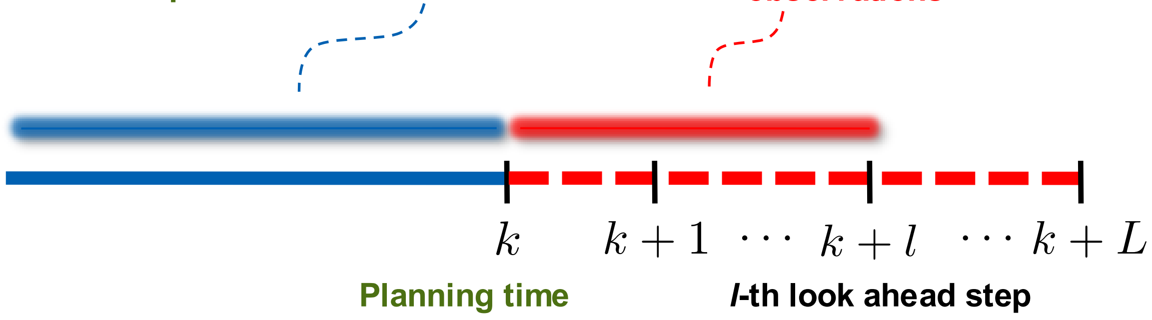

- Examples for reward function $\rho(b, a)$:
  - Expected distance to goal (navigate to a goal)
  - Information theoretic reward (reduce uncertainty)
  - …

Challenge

Probabilistic Inference

Maintain a distribution over the state given data

$$b_k \triangleq b[X_k] = \mathbb{P}(X_k \mid a_{0:k-1}, z_{1:k})$$

state — actions — observations

Decision-making under uncertainty

Involves reasoning about the entire observation and action spaces along planning horizon

Computationally intractable

More so, in high dimensional settings

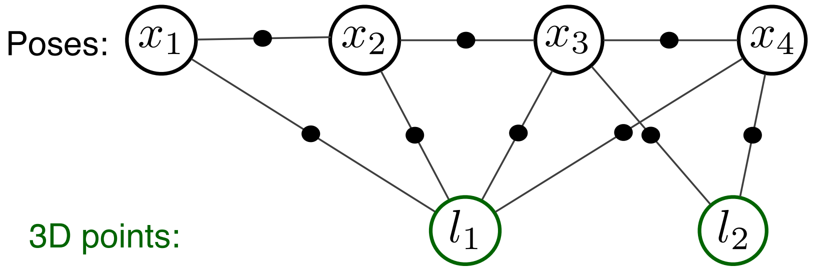

Example – grid world

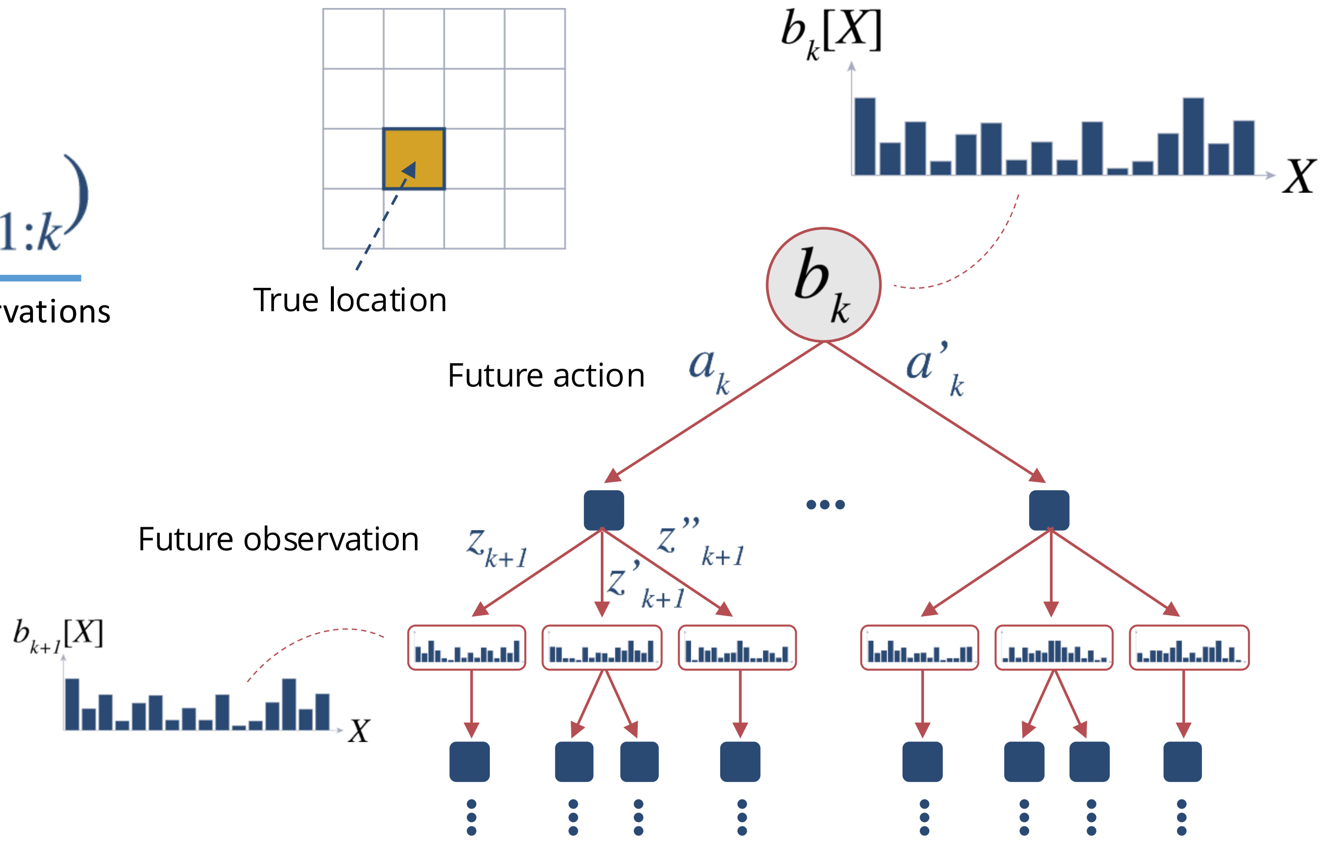

ct autonomously online and efficiently tasks in a safe and reliable fashion??

Agenda

Experience Reuse in POMDP Planning

POMDP Planning with Hybrid Beliefs

Simplification of POMDP with Formal Guarantees

Multi-agent POMDP Planning with Inconsistent Beliefs

Agenda

- Experience Reuse in POMDP Planning
- POMDP Planning with Hybrid Beliefs
- Simplification of POMDP with Formal Guarantees
- Multi-agent POMDP Planning with Inconsistent Beliefs

Experience Reuse in POMDP Planning

- Consider POMDPs with continuous state, action, and observation spaces

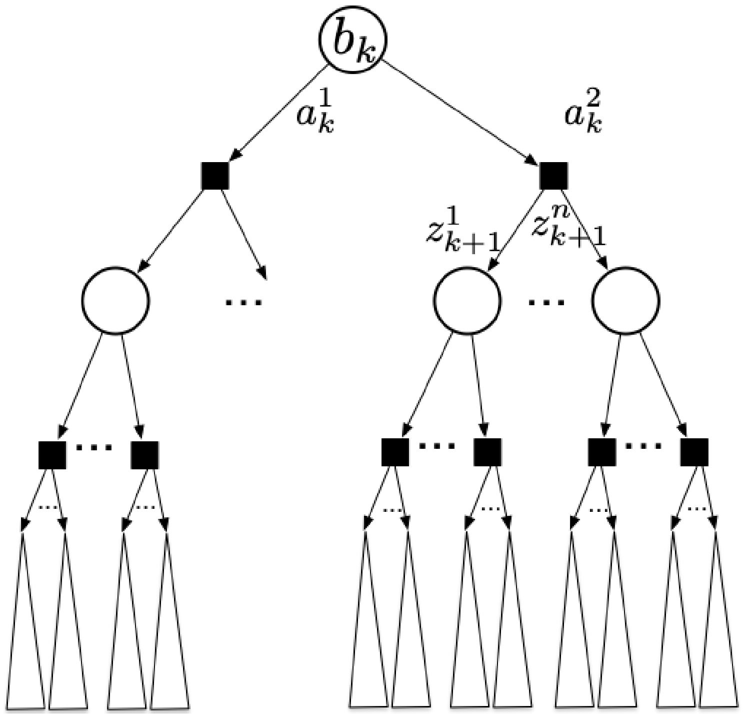

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Experience Reuse in POMDP Planning

- Consider POMDPs with continuous state, action, and observation spaces
- The probability of sampling the same belief/observation twice is zero

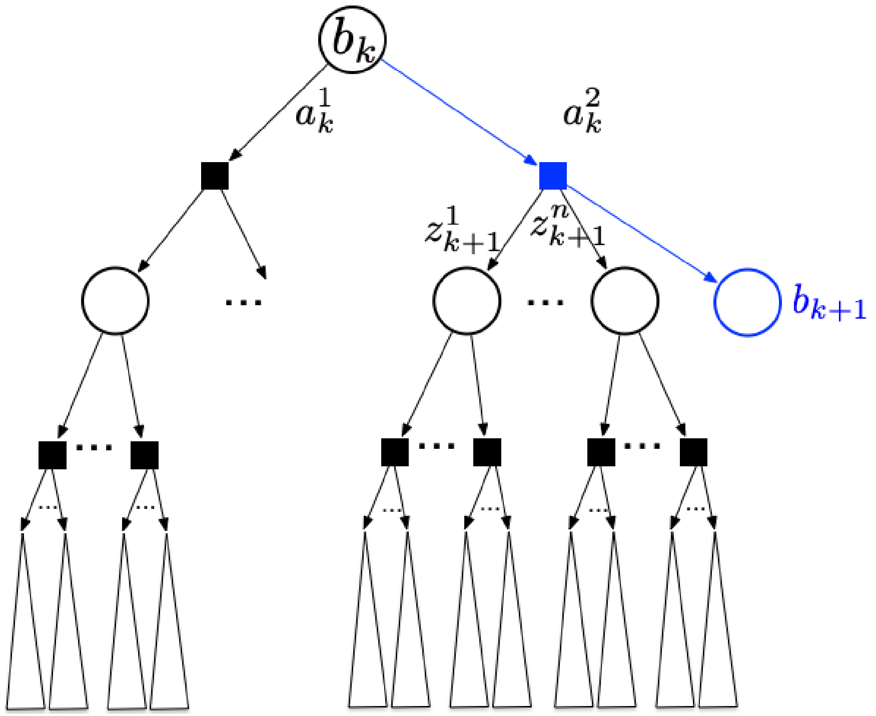

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Experience Reuse in POMDP Planning

- Consider POMDPs with continuous state, action, and observation spaces
- The probability of sampling the same belief/observation twice is zero

Online SOTA POMDP solvers typically perform calculations from **scratch at each planning session**

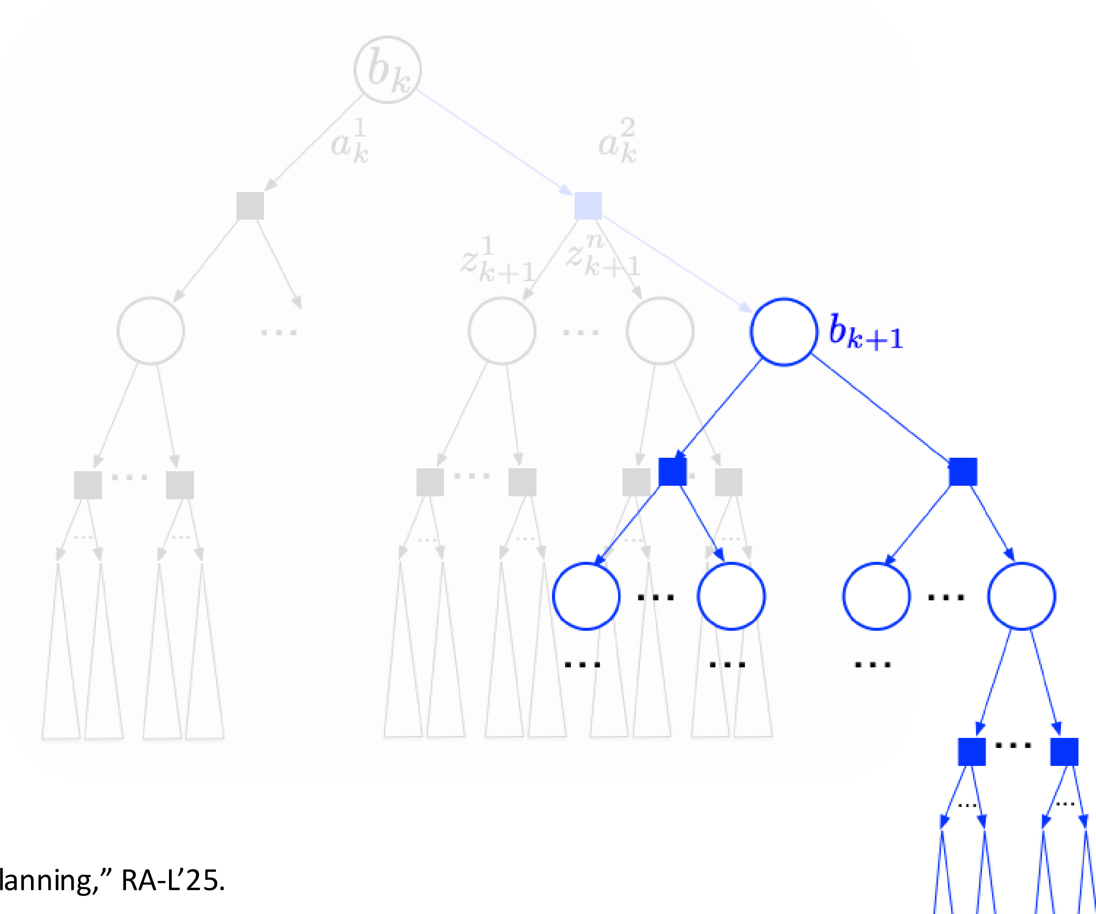

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Experience Reuse in POMDP Planning

- Consider POMDPs with continuous state, action, and observation spaces
- The probability of sampling the same belief/observation twice is zero
- Previous planning sessions (experience) can provide useful information in the current planning session

Online SOTA POMDP solvers typically perform calculations from **scratch at each planning session**

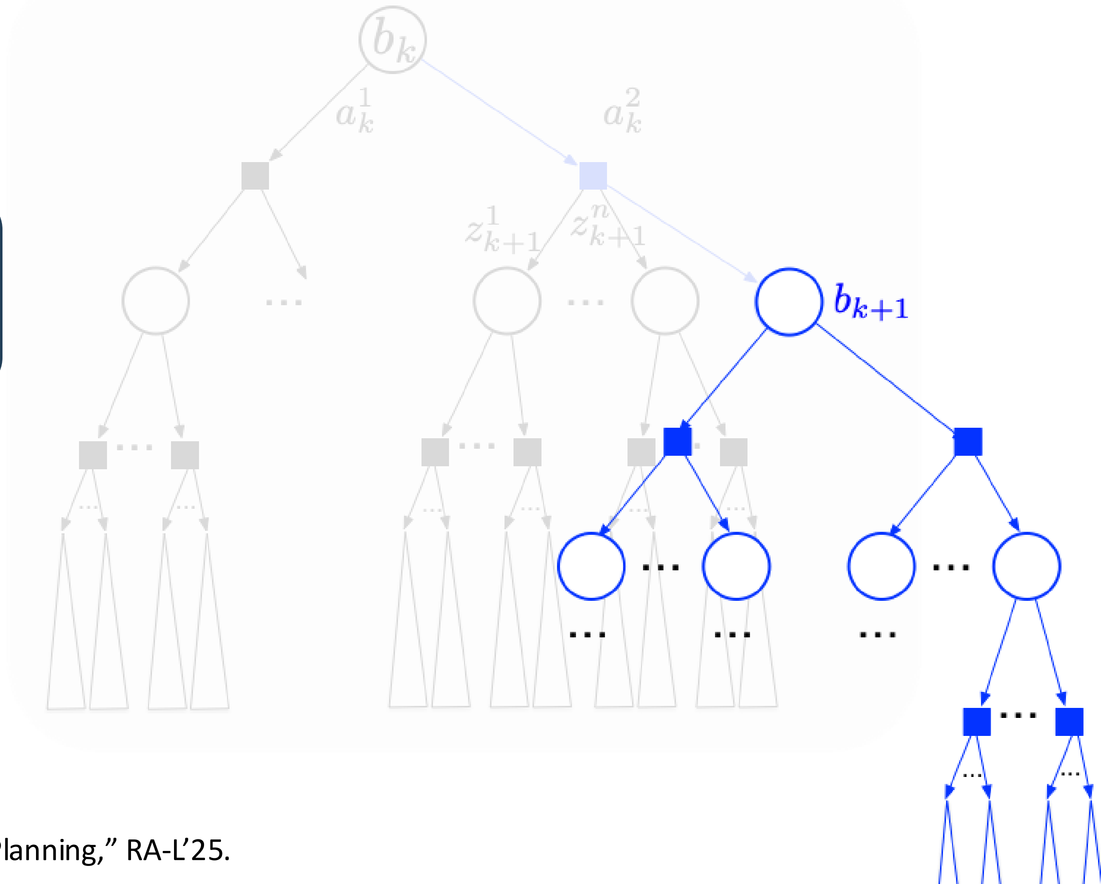

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Experience Reuse in POMDP Planning

- Consider POMDPs with continuous state, action, and observation spaces
- The probability of sampling the same belief/observation twice is zero
- Previous planning sessions (experience) can provide useful information in the current planning session

Key idea: Reuse previous planning session(s) to get an efficient estimation of

$$Q^{\pi}(b,a) = \mathbb{E}_{\pi}[\sum_{i=k}^{k+L-1} \gamma^{i-k} r(b_i, \pi_i(b_i), b_{i+1}) \mid b_k = b, a_k = a] \triangleq \mathbb{E}_{\pi}[G \mid b_k = b, a_k = a]$$

- Instead of calculating each planning session from scratch (state of the art)

E. Farhi and V. Indelman, “iX-BSP: Incremental Belief Space Planning,” ICRA’19, arXiv’21.
M. Novitsky, M. Barenboim, and V. Indelman, “Previous Knowledge Utilization In Online Anytime Belief Space Planning,” RA-L’25.

Experience Reuse in POMDP Planning

- Consider a planning session at time instant k

$$Q^{\pi}(b_k, a_k)$$

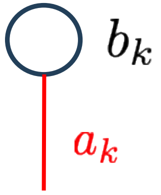

Current time

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Experience Reuse in POMDP Planning

- Consider a planning session at time instant k

$$Q^{\pi}(b_k, a_k)$$

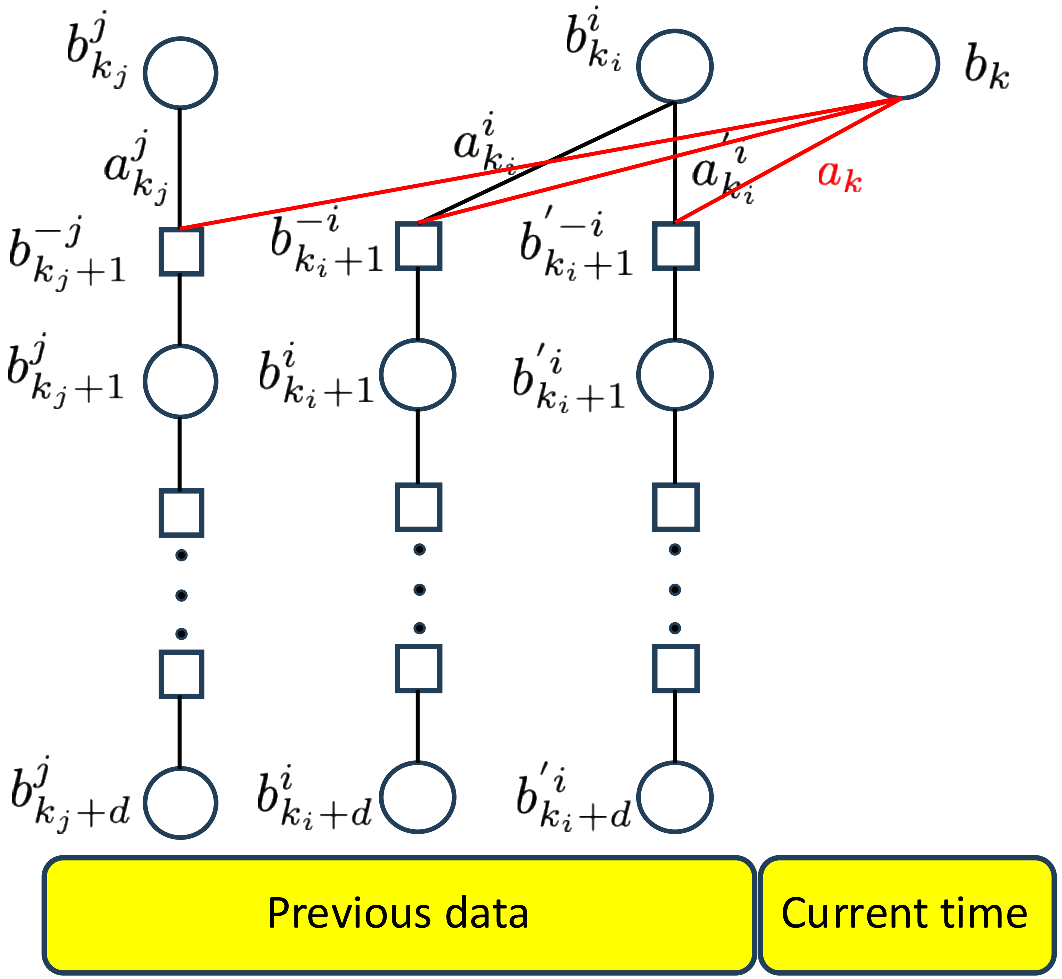

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Experience Reuse in POMDP Planning

- **Key idea**: multiple importance sampling (MIS) estimator

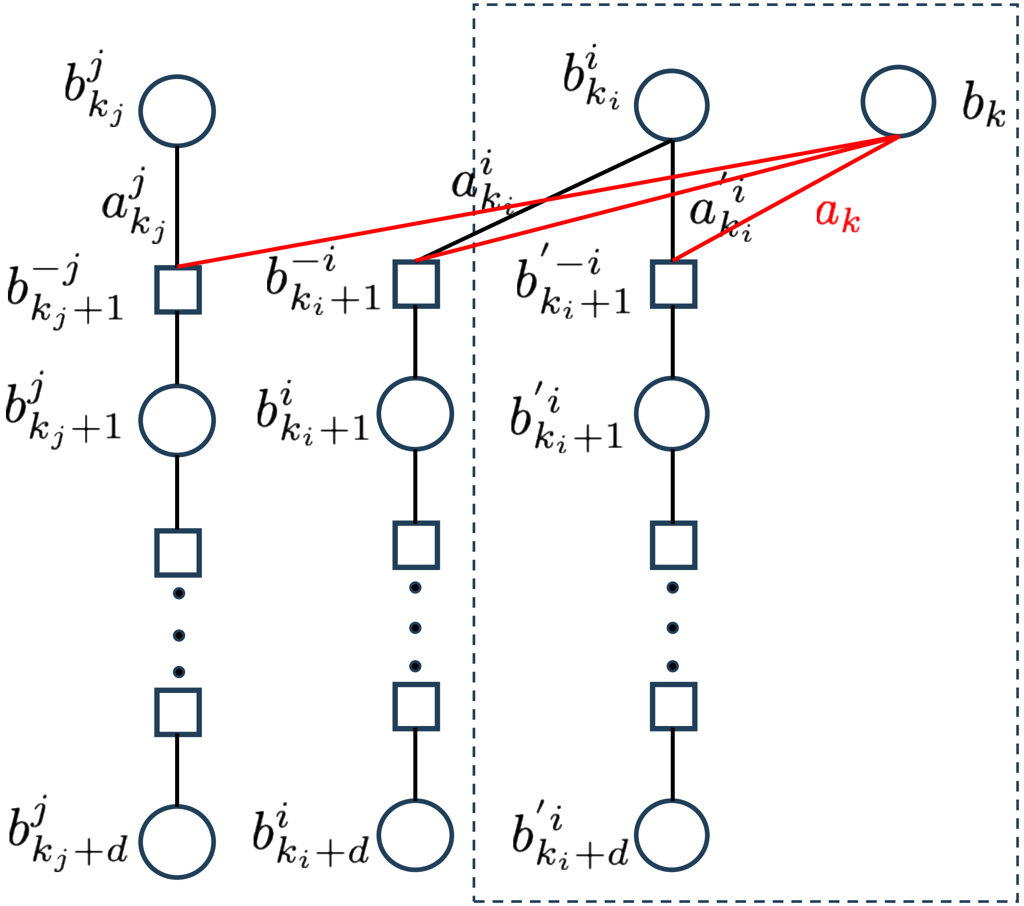

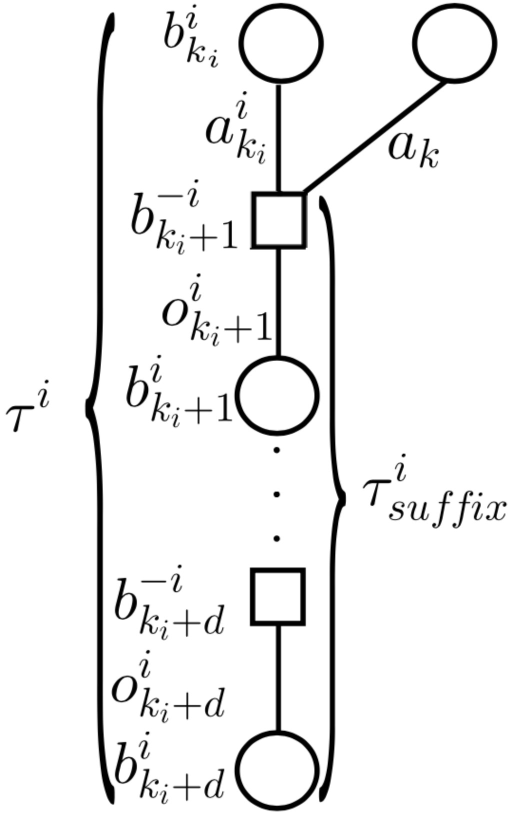

$$w_i \triangleq \frac{\mathbb{P}(\tau^i_{suffix} | b_k, a_k, \pi)}{\mathbb{P}(\tau^i_{suffix} | b^i_{k_i}, a^i_{k_i}, \pi)}$$

E. Farhi and V. Indelman, “iX-BSP: Incremental Belief Space Planning,” ICRA’19, arXiv’21.

M. Novitsky, M. Barenboim, and V. Indelman, “Previous Knowledge Utilization In Online Anytime Belief Space Planning,” RA-L’25.

Experience Reuse in POMDP Planning

- **<u>Key idea</u>**: multiple importance sampling (MIS) estimator

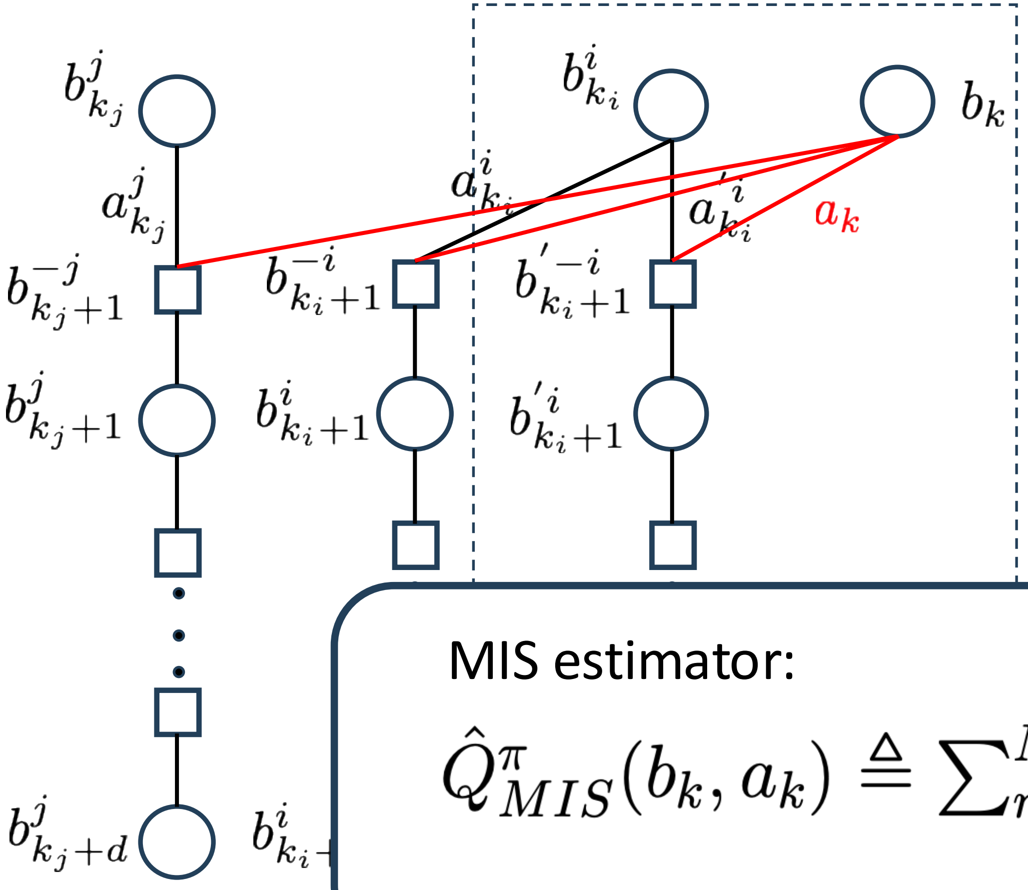

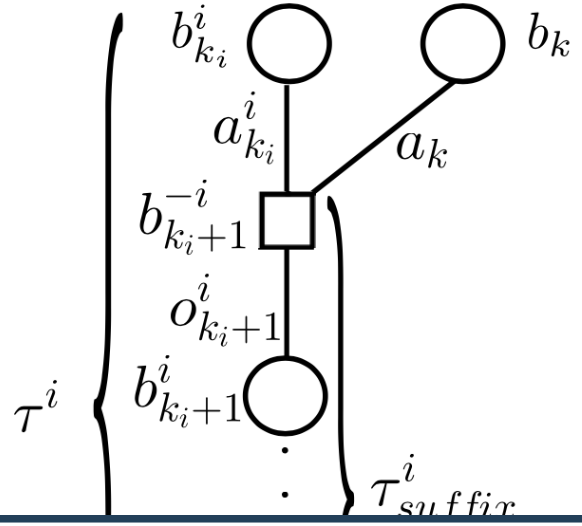

$$w_i \triangleq \frac{\mathbb{P}(\tau^i_{suffix} | b_k, a_k, \pi)}{\mathbb{P}(\tau^i_{suffix} | b^i_{k_i}, a^i_{k_i}, \pi)}$$

MIS estimator:

$$\hat{Q}^{\pi}_{MIS}(b_k, a_k) \triangleq \sum_{m=1}^{M} \sum_{l=1}^{n_m} \frac{\mathbb{P}(\tau^{l,m}_{suffix} | b_k, a_k, \pi) \tilde{G}^{l,m}}{\sum_{j=1}^{M} n_j \cdot \mathbb{P}(\tau^{l,m}_{suffix} | b^j_{k_j}, a^j_{k_j}, \pi)}.$$

return – reused!

E. Farhi and V. Indelman, “iX-BSP: Incremental Belief Space Planning,” ICRA’19, arXiv’21.
M. Novitsky, M. Barenboim, and V. Indelman, “Previous Knowledge Utilization In Online Anytime Belief Space Planning,” RA-L’25.

Incremental Belief Space Planning

Basic simulation – autonomous navigation in unknown environments:

ML-BSP: BSP with ML observations (one sample per look ahead step)

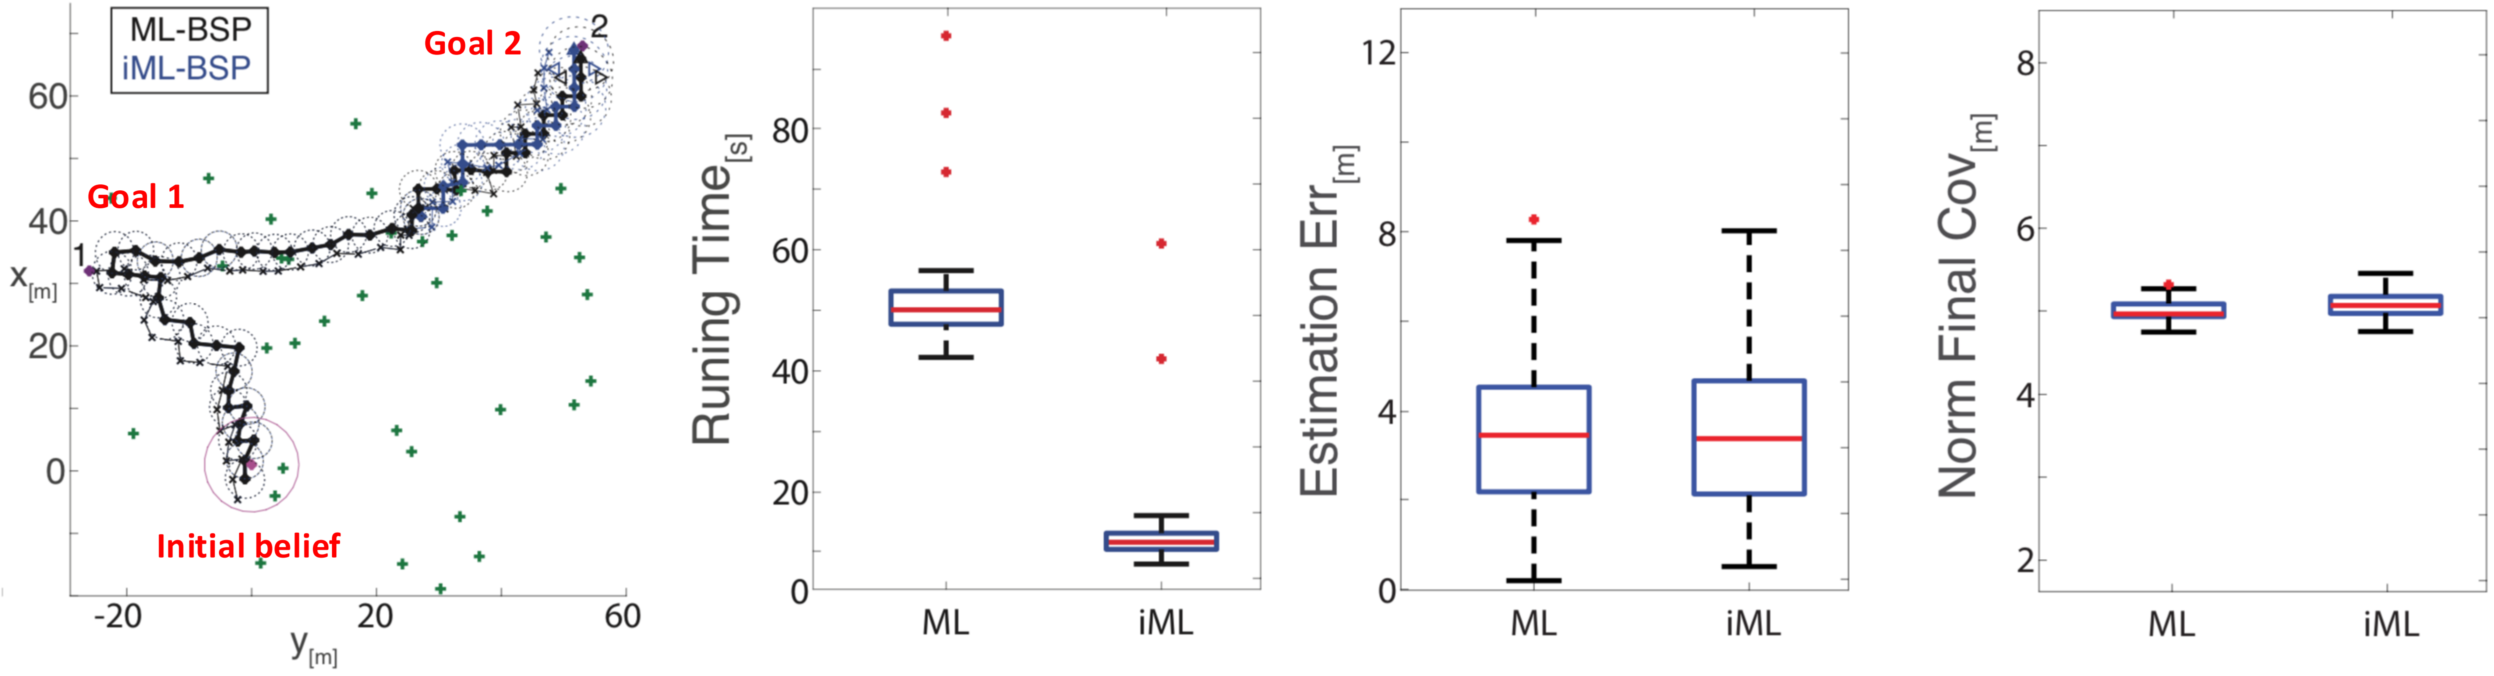

E. Farhi and V. Indelman, “iX-BSP: Incremental Belief Space Planning,” ICRA’19, arXiv’21.
M. Novitsky, M. Barenboim, and V. Indelman, “Previous Knowledge Utilization In Online Anytime Belief Space Planning,” RA-L’25.

Incremental Reuse Particle Filter Tree (IR-PFT)

- Extend PFT-DPW[1] , incorporating trajectories from previous planning sessions for fast estimation of $Q(b_k, a_k)$

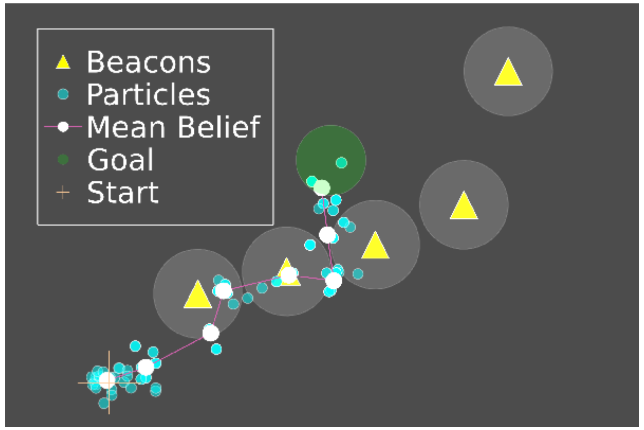

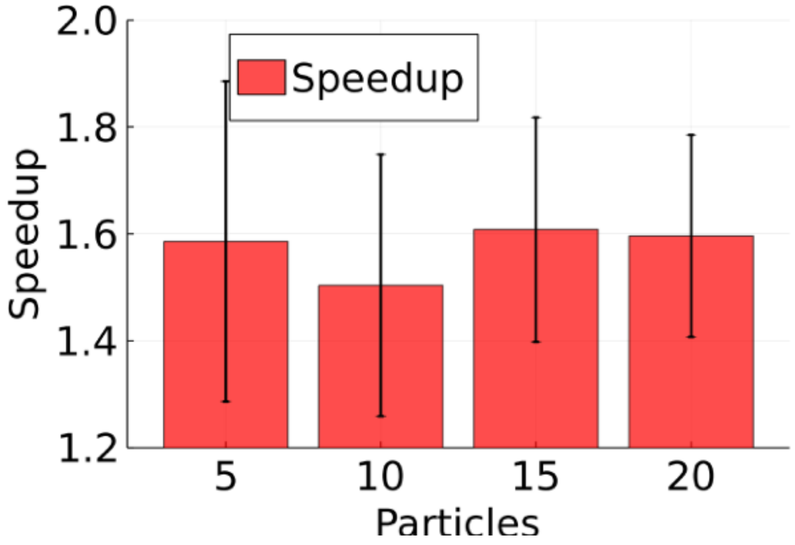

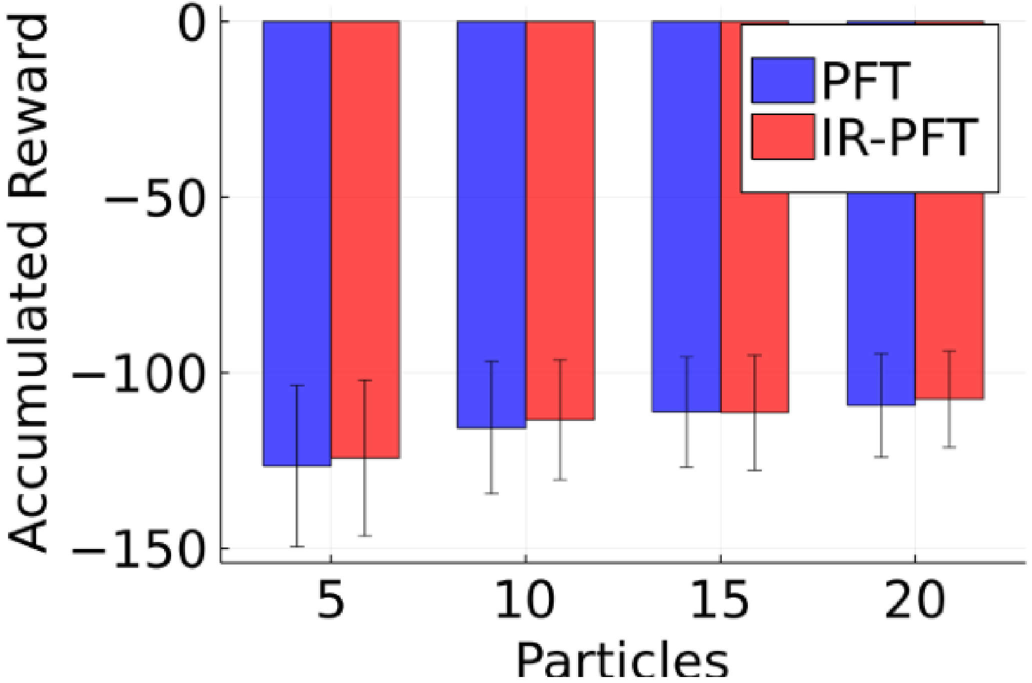

[1]Z. Sunberg and M. Kochenderfer. "Online algorithms for POMDPs with continuous state, action, and observation spaces." ICAPS, 2018.

E. Farhi and V. Indelman, "iX-BSP: Incremental Belief Space Planning," ICRA'19, arXiv'21.
M. Novitsky, M. Barenboim, and V. Indelman, "Previous Knowledge Utilization In Online Anytime Belief Space Planning," RA-L'25.

Action-Gradient Monte Carlo Tree Search for Non-Parametric Continuous (PO)MDPs

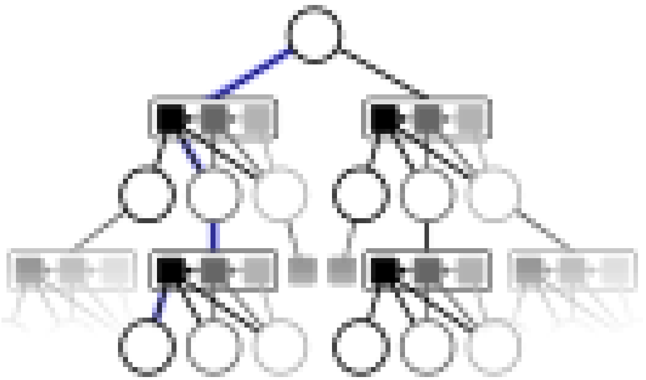

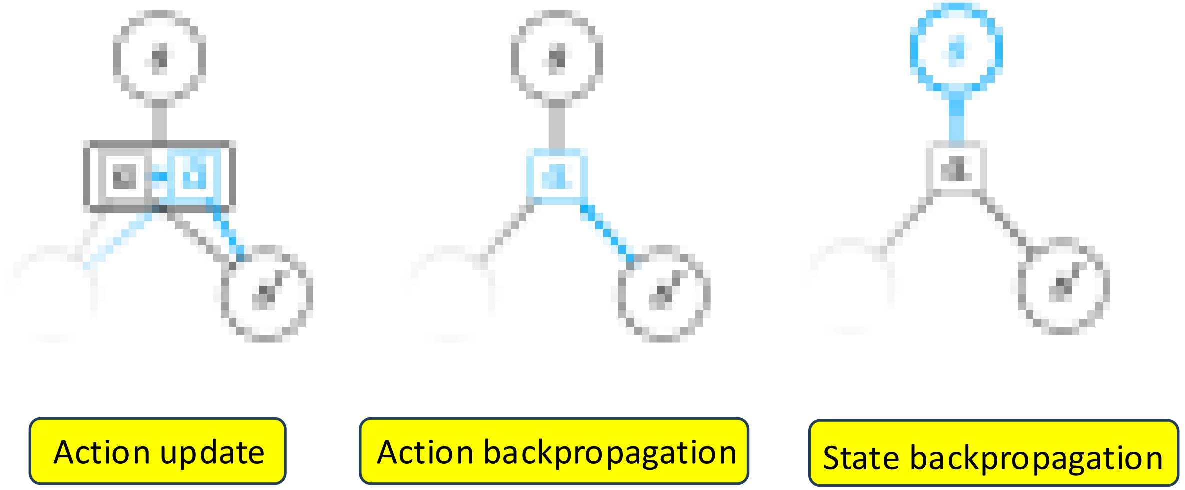

I. Lev-Yehudi, M. Novitsky, M. Barenboim, R. Benchetrit, and V. Indelman, “Action-Gradient Monte Carlo Tree Search for Non-Parametric Continuous (PO)MDPs,” arXiv’25.

Agenda

Experience Reuse in POMDP Planning

POMDP Planning with Hybrid Beliefs

Simplification of POMDP with Formal Guarantees

Multi-agent POMDP Planning with Inconsistent Beliefs

POMDP Planning with Hybrid Beliefs

Autonomous Semantic Perception

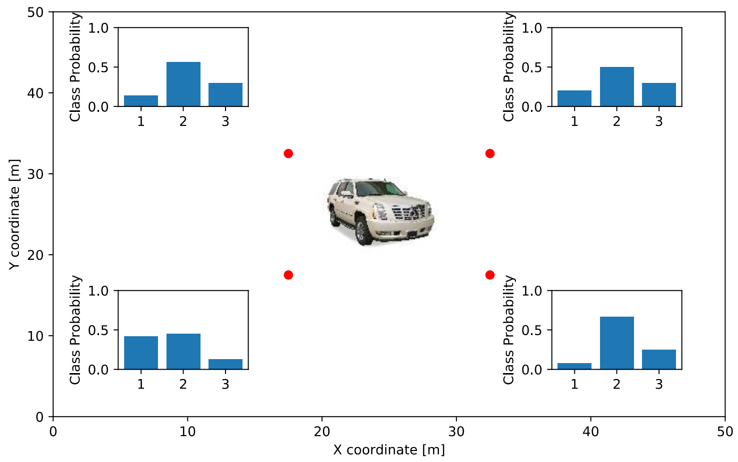

Viewpoint-dependent semantic models

Ambiguous Environments

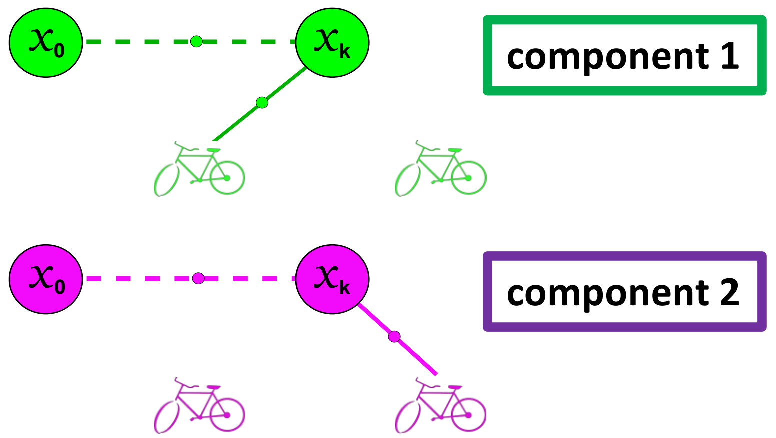

Data association hypotheses

- Hybrid beliefs (over continuous and discrete random variables)
- The number of hypotheses can grow exponentially
- How do we do probabilistic inference and POMDP planning?

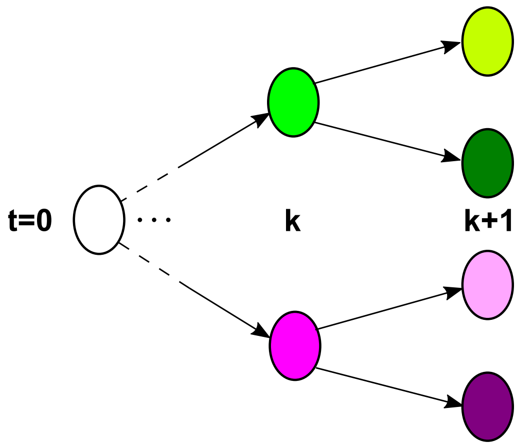

TECHNION Israel Institute of Technology

ANPL Autonomous Navigation and Perception Lab

Semantic Perception & SLAM

- Usually, semantics and geometry are considered **separately**
- Cannot use coupled observation models or priors
- Can lead to absurd results

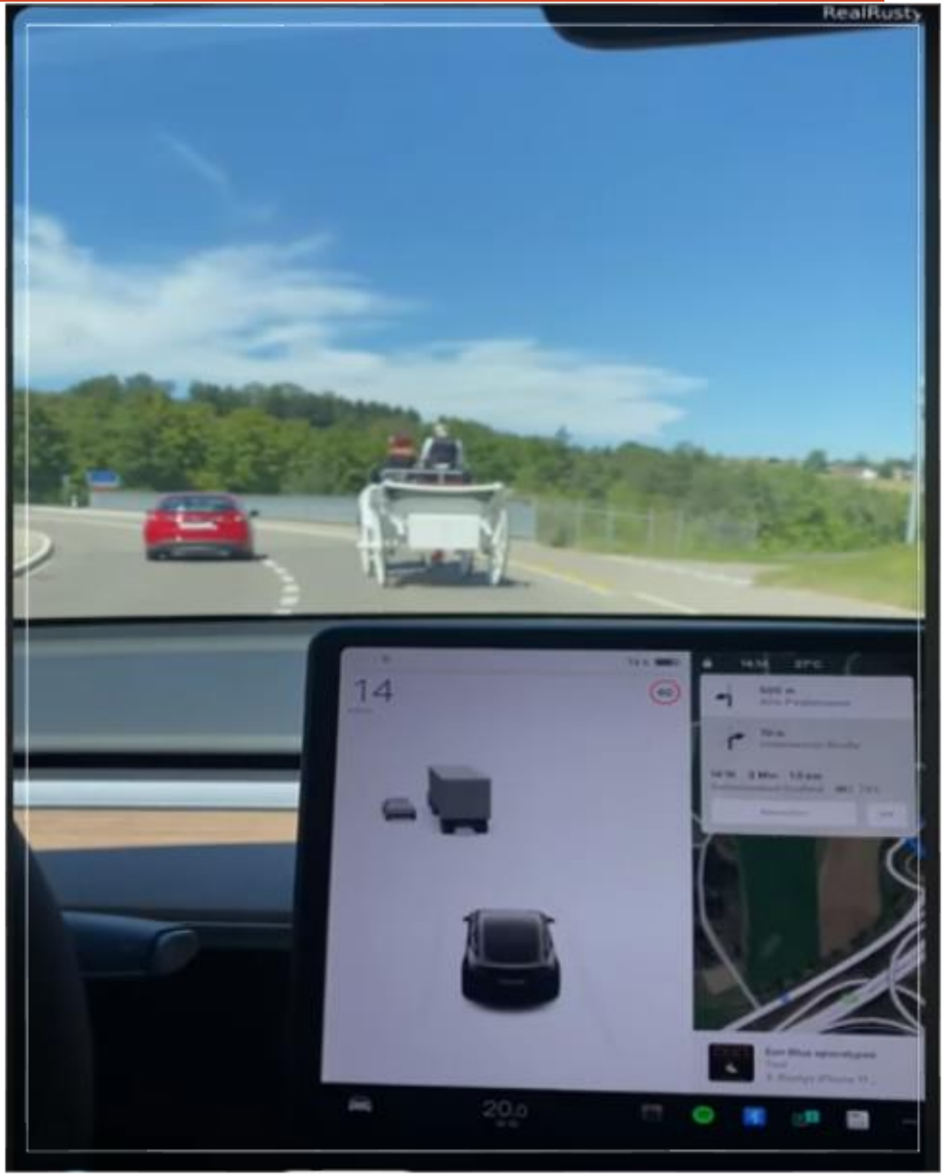

TECHNION Israel Institute of Technology

ANPL Autonomous Navigation and Perception Lab

Class- and Viewpoint-Dependency

- Is it a floor or a roof?
- Depending on the viewpoint of the viewer!
  - Looking on the people below - it's a floor
  - Looking on the people above - it's a roof

- How do we know the viewpoint?

Class- and Viewpoint-Dependency

- View-dependent semantic observation model:

$$\mathbb{P}(z^s \mid c, \mathcal{X}^{rel})$$

Object class

Semantic observation (from a classifier)

Agent's viewpoint relative to object

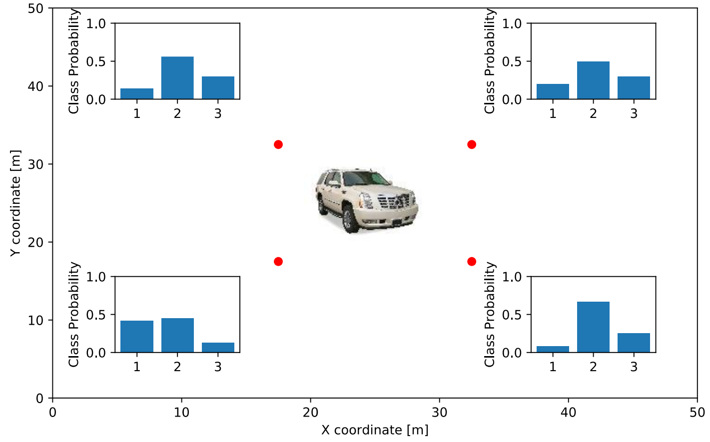

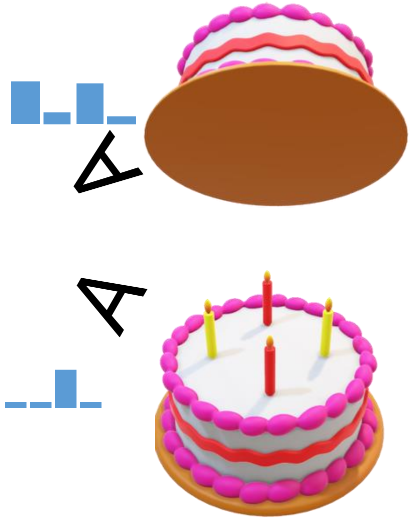

- Classes and poses can be coupled via learned prior probabilities.
- Reward/constraint can depend on both classes and poses (e.g., object search, search & rescue)

Y. Feldman and V. Indelman, "Bayesian Viewpoint-Dependent Robust Classification under Model and Localization Uncertainty," ICRA'18, AURO'20.
V. Tchuiev, Y. Feldman, and V. Indelman, "Data Association Aware Semantic Mapping and Localization via a Viewpoint Dependent Classifier Model," IROS'19.
V. Tchuiev and V. Indelman, "Epistemic Uncertainty Aware Semantic Localization and Mapping for Inference and Belief Space Planning," Artificial Intelligence, 2023.
T. Lemberg and V. Indelman, "Online Hybrid-Belief POMDP with Coupled Semantic-Geometric Models and Semantic Safety Awareness", arXiv'25.

Hybrid Belief

- **Hybrid Belief** at time instant k:

$$b[X_k, C] = \mathbb{P}(X_k, C \mid \mathcal{H}_k)$$

Robot’s and objects’ poses — Objects’ classes — History (actions, geometric & semantic observations)

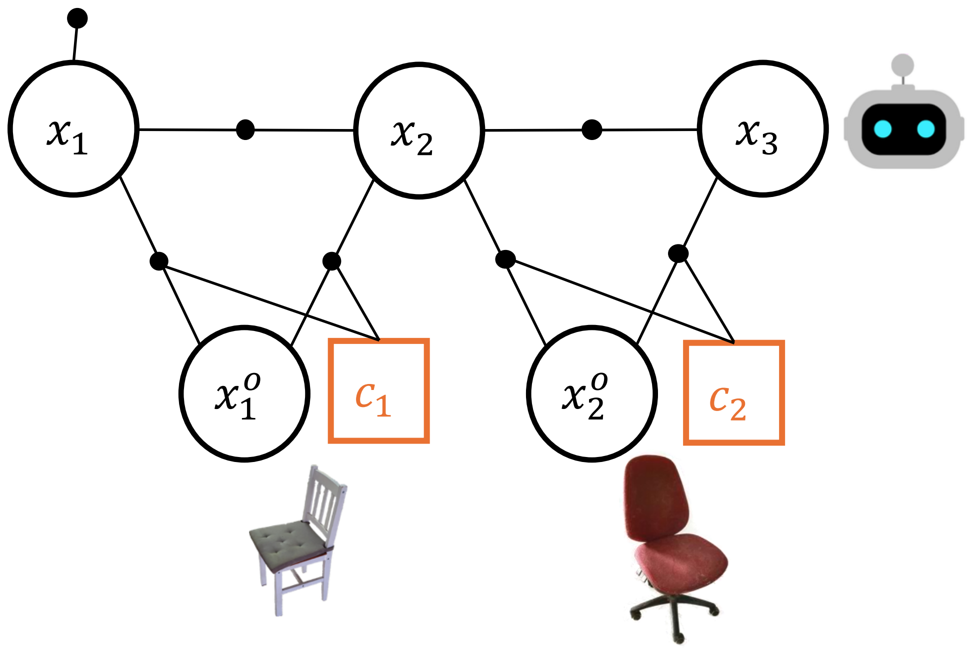

- Classes and agent poses are dependent
- Classes of different objects are dependent

Y. Feldman and V. Indelman, “Bayesian Viewpoint-Dependent Robust Classification under Model and Localization Uncertainty,” ICRA’18 , AURO’20.
V. Tchuiev, Y. Feldman, and V. Indelman, “Data Association Aware Semantic Mapping and Localization via a Viewpoint Dependent Classifier Model,” IROS’19.
V. Tchuiev and V. Indelman, “Epistemic Uncertainty Aware Semantic Localization and Mapping for Inference and Belief Space Planning,” Artificial Intelligence, 2023.
T. Lemberg and V. Indelman, “Online Hybrid-Belief POMDP with Coupled Semantic-Geometric Models and Semantic Safety Awareness”, arXiv’25.

POMDP Planning with Hybrid Semantic-Geometric Beliefs

- Value function

$$V^{\pi}(b_k) = \mathbb{E}_{z_{k+1:k+L}}\left[\sum_{l=k}^{k+L-1} \rho(b_l, \pi_l(b_l), b_{l+1})\right]$$

- Semantic Risk Awareness

$$\mathbb{P}_{safe} \triangleq \mathbb{P}(\{\wedge_{t=k+1}^{L} x_t \notin \mathcal{X}_{unsafe}(C, X^o)\} \mid b_k[x_k, C, X^o], \pi)$$

Objects' classes

Objects' poses

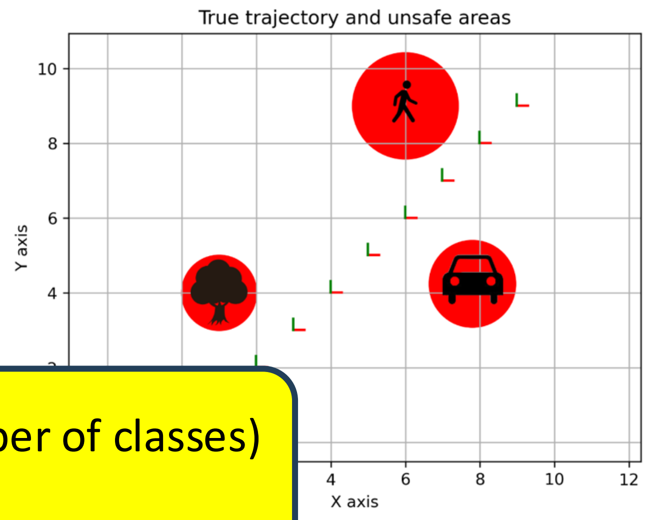

The number of classification hypotheses is M^N (N: number of objects, M: number of classes)
How to sample w/o pruning hypotheses? How to estimate $\mathbb{P}_{safe}$?

T. Lemberg and V. Indelman, "Online Hybrid-Belief POMDP with Coupled Semantic-Geometric Models and Semantic Safety Awareness," arXiv'25.

POMDP Planning with Hybrid Beliefs

Autonomous Semantic Perception

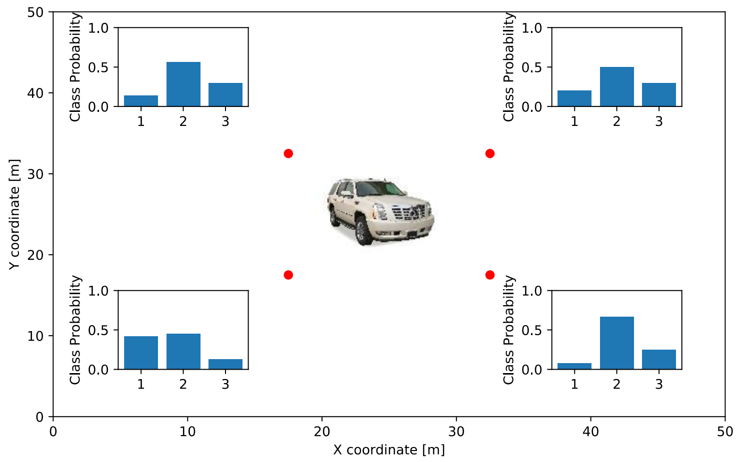

Viewpoint-dependent semantic models

Ambiguous Environments

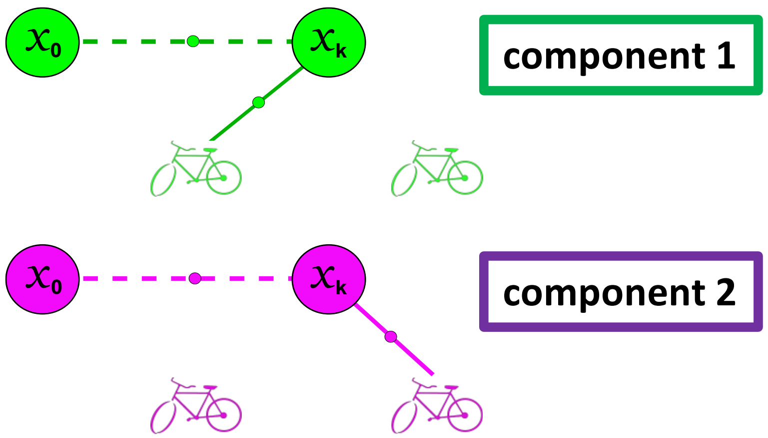

Data association hypotheses

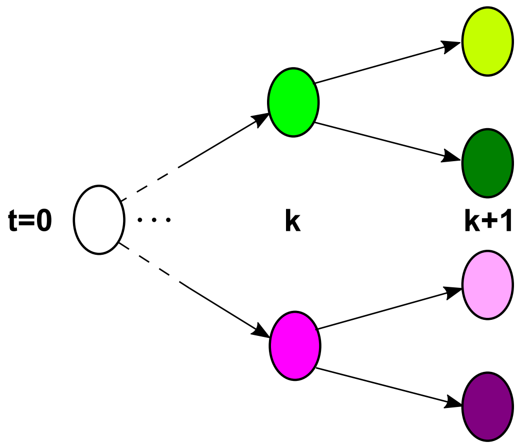

- Hybrid beliefs (over continuous and discrete random variables)
- The number of hypotheses can grow exponentially
- How do we do probabilistic inference and POMDP planning?

Ambiguous Scenarios

- Have to reason about data association hypotheses within inference and planning

An observation:
(e.g. LIDAR)

How should the agent act?

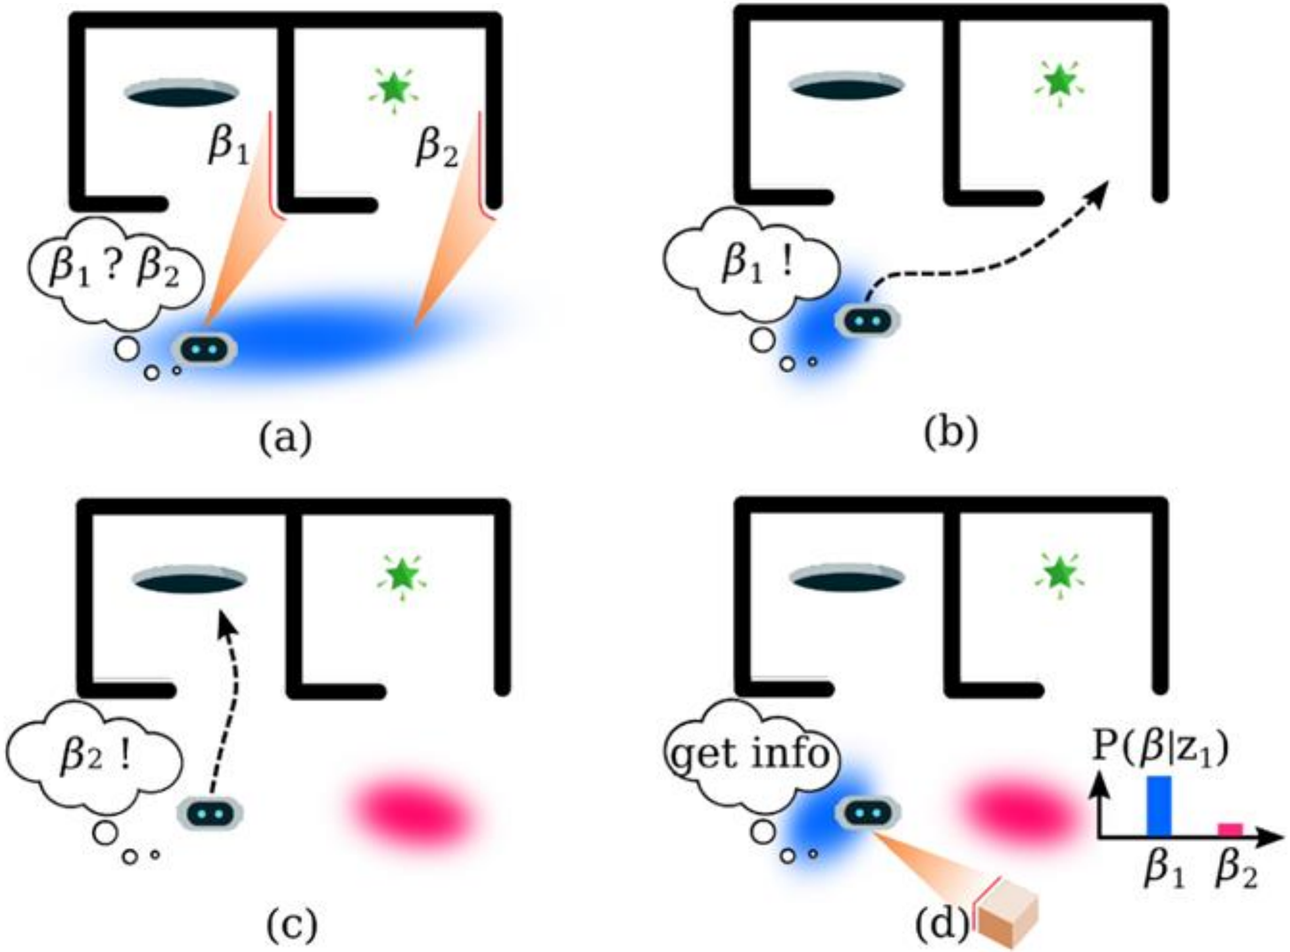

(a) (b) (c) (d)

S. Pathak, A. Thomas, and V. Indelman, "A Unified Framework for Data Association Aware Belief Space Planning and Perception", IJRR'18.

Continuous-Discrete State Spaces

- The number of hypotheses may grow **exponentially** with the planning horizon!

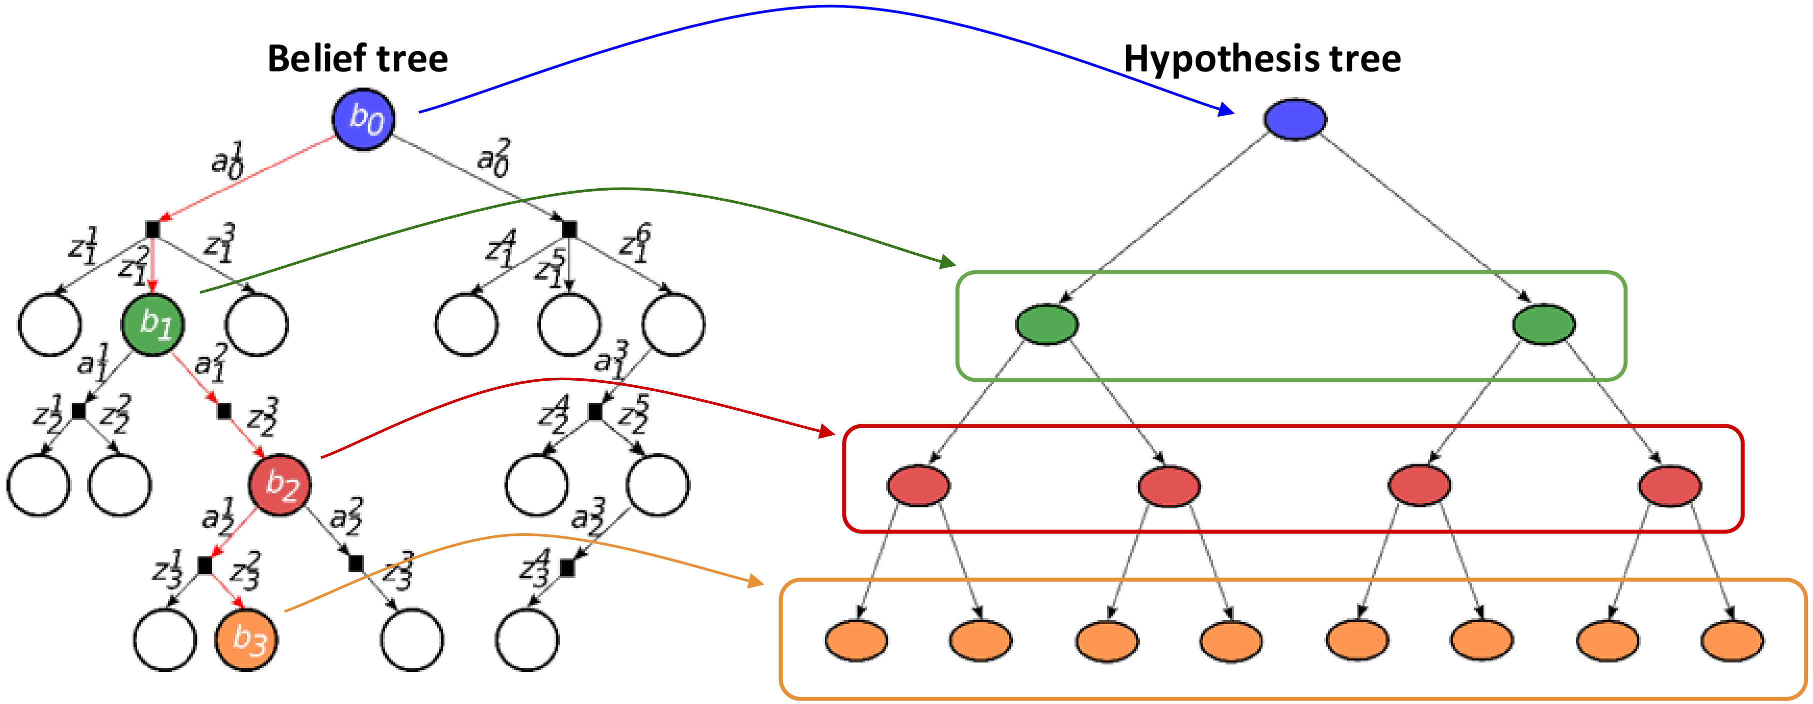

M. Barenboim, M. Shienman, and V. Indelman, "Monte Carlo Planning in Hybrid Belief POMDPs," IEEE RA-L'23.

Continuous-Discrete State Spaces

- The number of hypotheses may grow **exponentially** with the planning horizon!

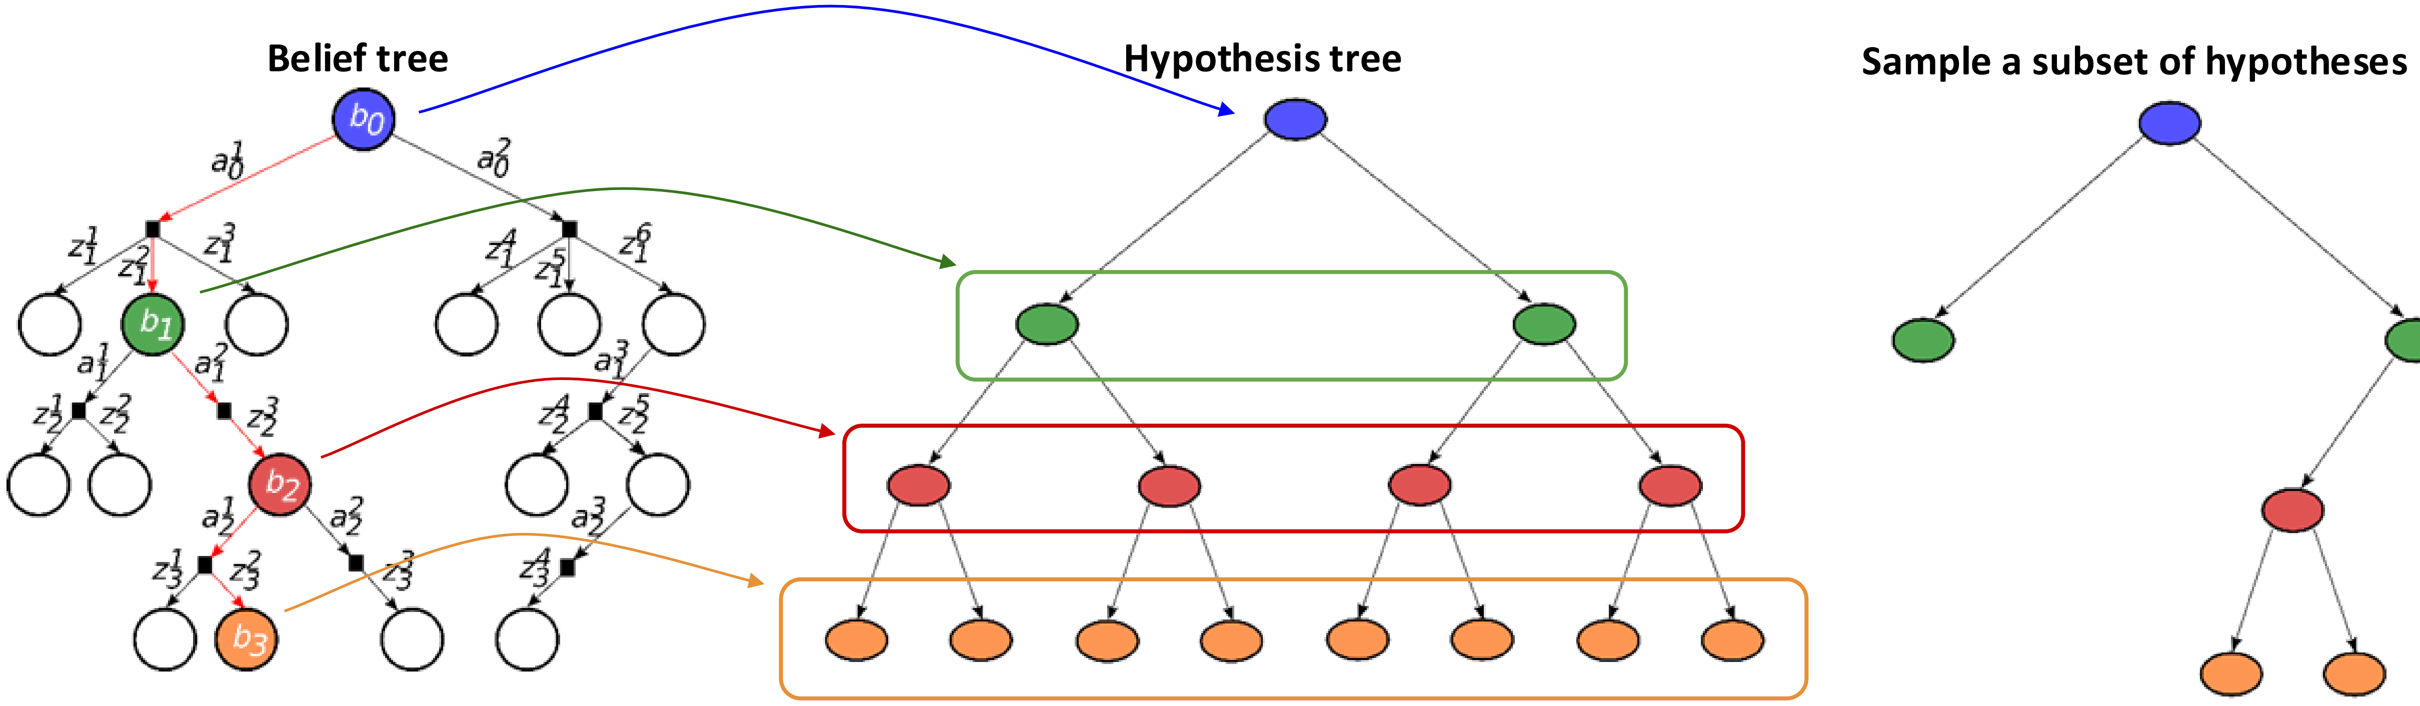

Impact on decision making?

M. Barenboim, M. Shienman, and V. Indelman, “Monte Carlo Planning in Hybrid Belief POMDPs,” IEEE RA-L’23.

Agenda

Experience Reuse in POMDP Planning

POMDP Planning with Hybrid Beliefs

Simplification of POMDP with Formal Guarantees

Multi-agent POMDP Planning with Inconsistent Beliefs

Simplification Framework

Accelerate decision making by adaptive simplification while providing performance guarantees

Break original problem into (computationally) easier problems

Provide performance guarantees

Relate between the original & simplified problems

Tradeoff complexity & performance

Develop online adaptation mechanisms

[Indelman RAL'16; Elimelech & Indelman IJRR'22; Sztyglic & Indelman IROS'22, Zhitnikov & Indelman AIJ'22, TRO'24; Shienman & Indelman ICRA'22; Barenboim & Indelman NIPS'23; Kitanov & Indelman IJRR'24; Zhitnikov et al. IJRR'24; Lev-Yehudi, Barenboim & Indelman AAAI'24, Yotam & Indelman TRO'24, Da & Indelman ISRR'24]

Simplification of Decision-Making Problems

- Each element of the decision-making problem can be simplified
- **Action-consistent** simplification <u>preserves order</u> between actions w.r.t. original problem

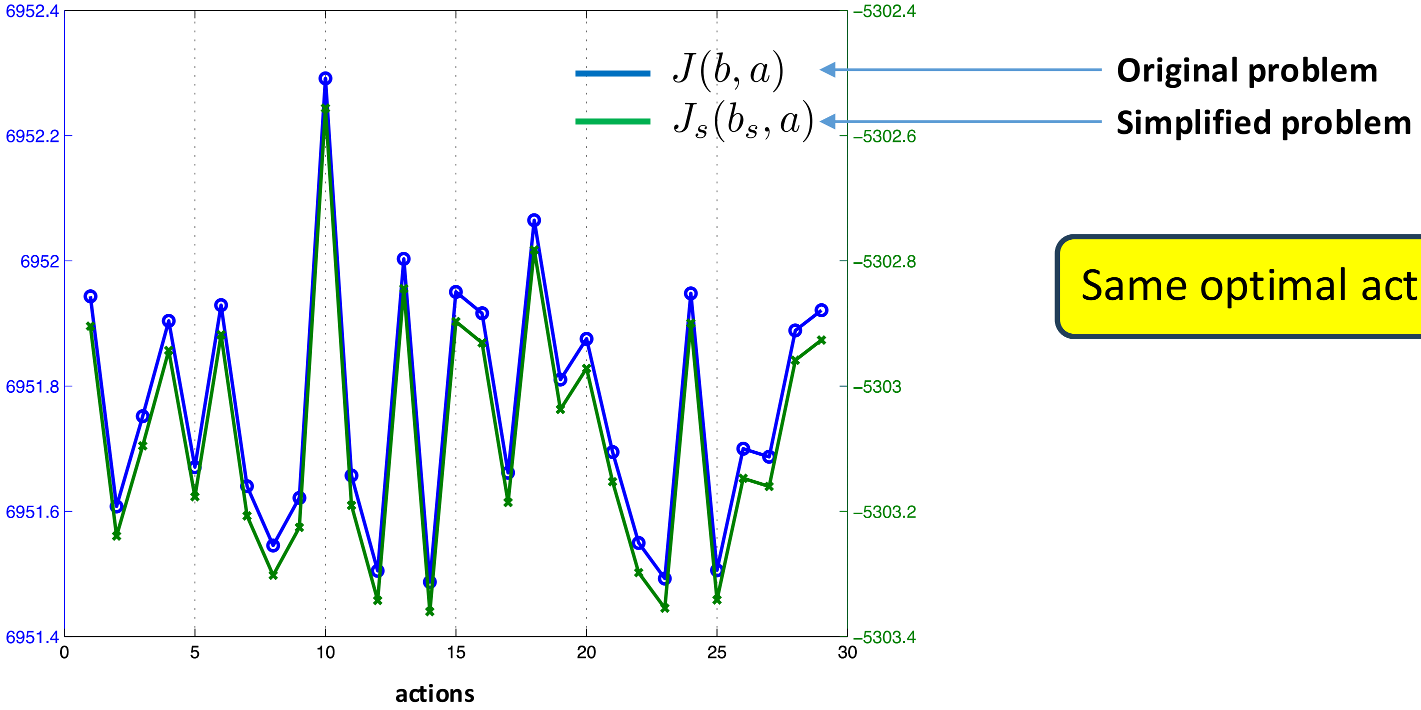

Same optimal action!

V. Indelman, "No Correlations Involved: Decision Making Under Uncertainty in a Conservative Sparse Information Space," IEEE RA-L'16.
K. Elimelech and V. Indelman, "Simplified decision making in the belief space using belief sparsification," IJRR'22.

Simplification of Decision-Making Problems

$$\mathcal{LB}(b,a) \leq Q(b,a) \leq \mathcal{UB}(b,a)$$

Computationally cheap(er) bounds

V. Indelman, "No Correlations Involved: Decision Making Under Uncertainty in a Conservative Sparse Information Space," IEEE RA-L'16.
K. Elimelech and V. Indelman, "Simplified decision making in the belief space using belief sparsification," IJRR'22.

Simplification of Decision-Making Problems

Concept:

- Identify and solve a **simplified (computationally) easier** decision-making problem
- Provide performance guarantees

Specific simplifications include:

- Sparsification of Gaussian beliefs (high dim. state)
- Topological metric for Gaussian beliefs (high dim. state)
- Utilize a subset of samples (nonparametric beliefs)
- Utilize a subset of hypotheses (hybrid beliefs)
- Simplified models and spaces
- Simplification of policy space
- Simplification of Risk-Averse & Robust Planning
- Simplification in a multi-agent setting

[Indelman RAL'16; Elimelech & Indelman IJRR'22; Sztyglic & Indelman IROS'22, Zhitnikov & Indelman AIJ'22, TRO'24; Shienman & Indelman ICRA'22; Kitanov & Indelman IJRR'24; Zhitnikov et al. IJRR'24, Barenboim & Indelman NIPS'23; Lev-Yehudi, Barenboim & Indelman AAAI'24, Kong & Indelman ISRR 2024]

Simplification of POMDPs with Nonparametric Beliefs

- Value function

$$V^{\pi}(b_k) = \mathbb{E}_{z_{k+1:k+L}}\left[\sum_{\ell=0}^{L} \rho(b_{k+\ell}, \pi_{k+\ell}(b_{k+\ell}))\right]$$

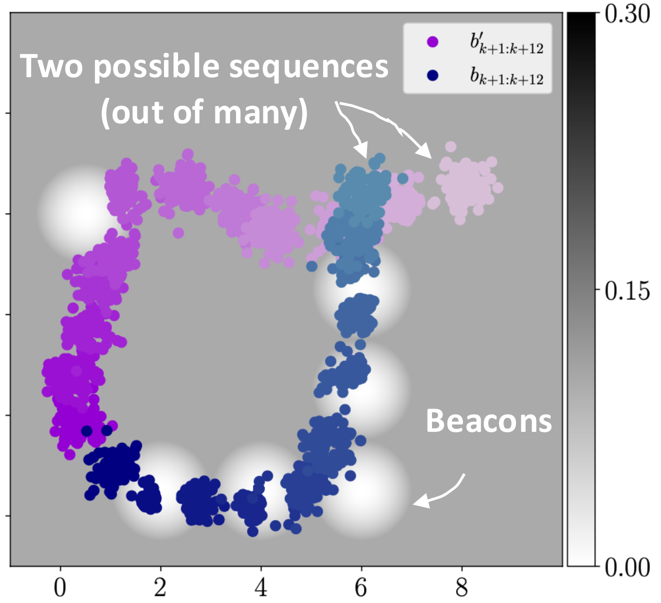

Simplification:

- Utilize a **subset** of samples for planning
- Information-theoretic reward (entropy)
- Analytical (**cheaper**) bounds over the reward

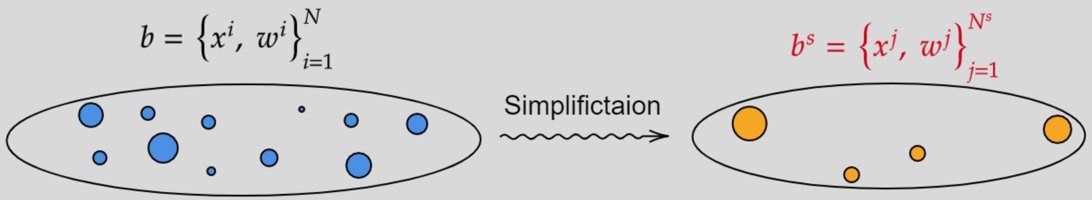

$$lb(b, b^s, a) \leq \rho(b, a) \leq ub(b, b^s, a)$$

O. Sztyglic and V. Indelman, “Speeding up POMDP Planning via Simplification”, IROS’22.
A. Zhitnikov, O. Sztyglic, and V. Indelman, “No Compromise in Solution Quality: Speeding Up Belief-dependent Continuous POMDPs via Adaptive Multilevel Simplification”, IJRR’24.

Simplification of POMDPs with Nonparametric Beliefs

- **Adaptive multi-level** simplification in a Sparse Sampling setting:

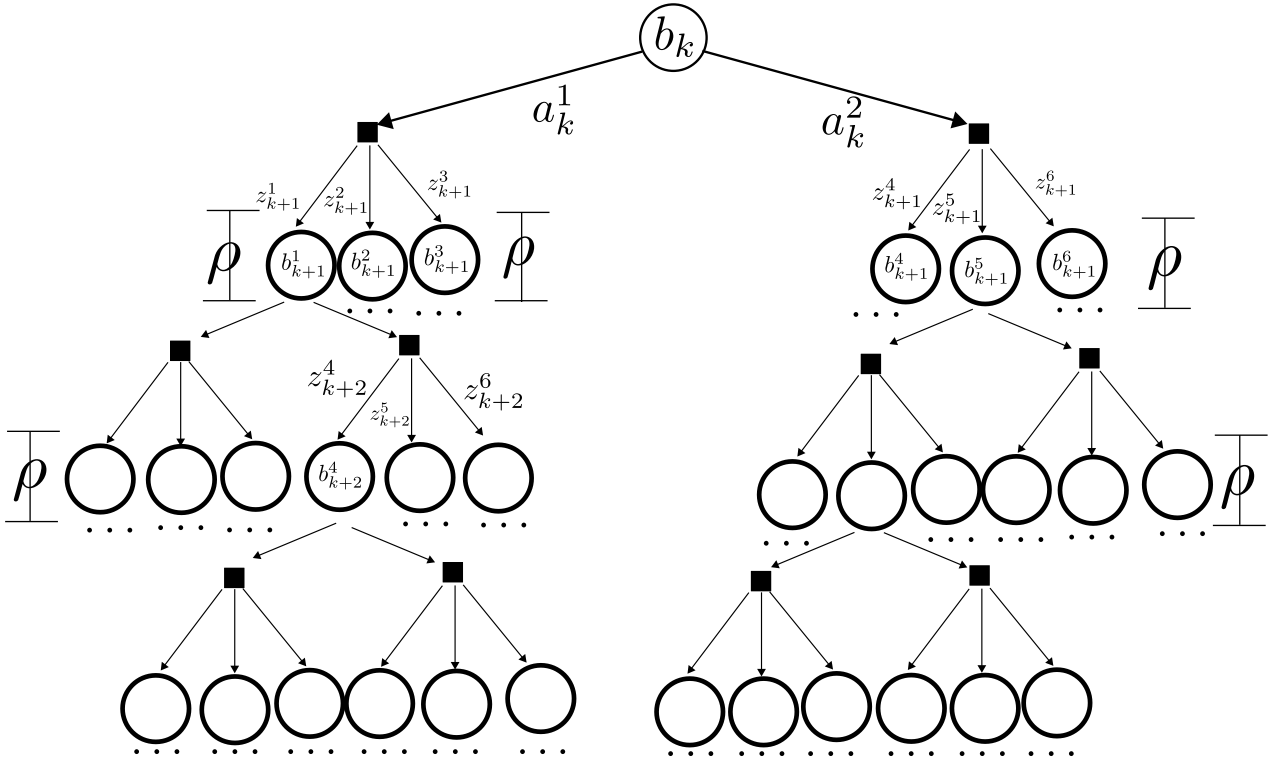

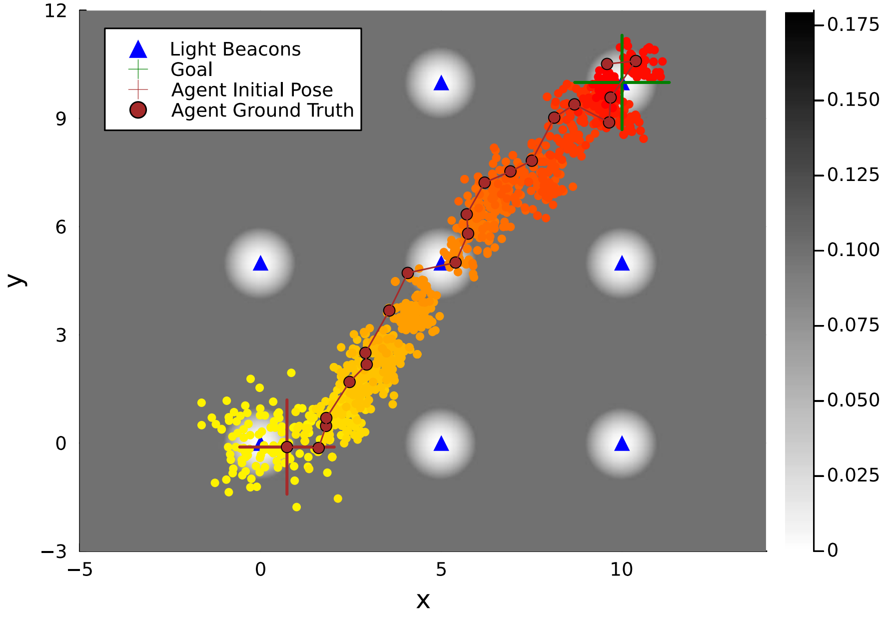

Typical speedup of 20% - 50%,
Same performance!

O. Sztyglic and V. Indelman, “Speeding up POMDP Planning via Simplification”, IROS’22.
A. Zhitnikov, O. Sztyglic, and V. Indelman, “No Compromise in Solution Quality: Speeding Up Belief-dependent Continuous POMDPs via Adaptive Multilevel Simplification”, IJRR’24.

Simplification of POMDPs with Nonparametric Beliefs

- **Adaptive multi-level** simplification in an MCTS setting:

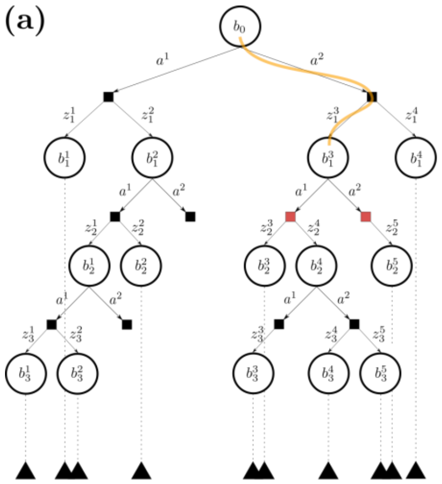

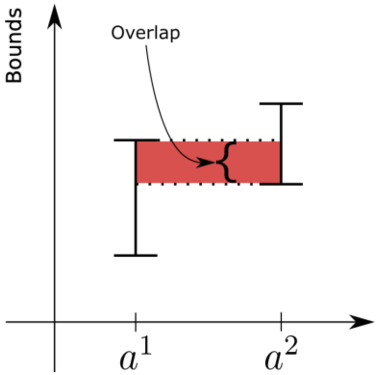

O. Sztyglic and V. Indelman, “Speeding up POMDP Planning via Simplification”, IROS’22.
A. Zhitnikov, O. Sztyglic, and V. Indelman, “No Compromise in Solution Quality: Speeding Up Belief-dependent Continuous POMDPs via Adaptive Multilevel Simplification”, IJRR’24.

Simplification of Decision-Making Problems

Concept:

- Identify and solve a **simplified (computationally) easier** decision-making problem
- Provide performance guarantees

Specific simplifications include:

- Sparsification of Gaussian beliefs (high dim. state)
- Topological metric for Gaussian beliefs (high dim. state)
- Utilize a subset of samples (nonparametric beliefs)
- Utilize a subset of hypotheses (hybrid beliefs)
- Simplified models and spaces
- Simplification of policy space
- Simplification of Risk-Averse & Robust Planning
- Simplification in a multi-agent setting

[Indelman RAL'16; Elimelech & Indelman IJRR'22; Sztyglic & Indelman IROS'22, Zhitnikov & Indelman AIJ'22, TRO'24; Shienman & Indelman ICRA'22; Kitanov & Indelman IJRR'24; Zhitnikov et al. IJRR'24, Barenboim & Indelman NIPS'23; Lev-Yehudi, Barenboim & Indelman AAAI'24, Kong & Indelman ISRR 2024]

Simplification of BSP/POMDP with Hybrid Beliefs

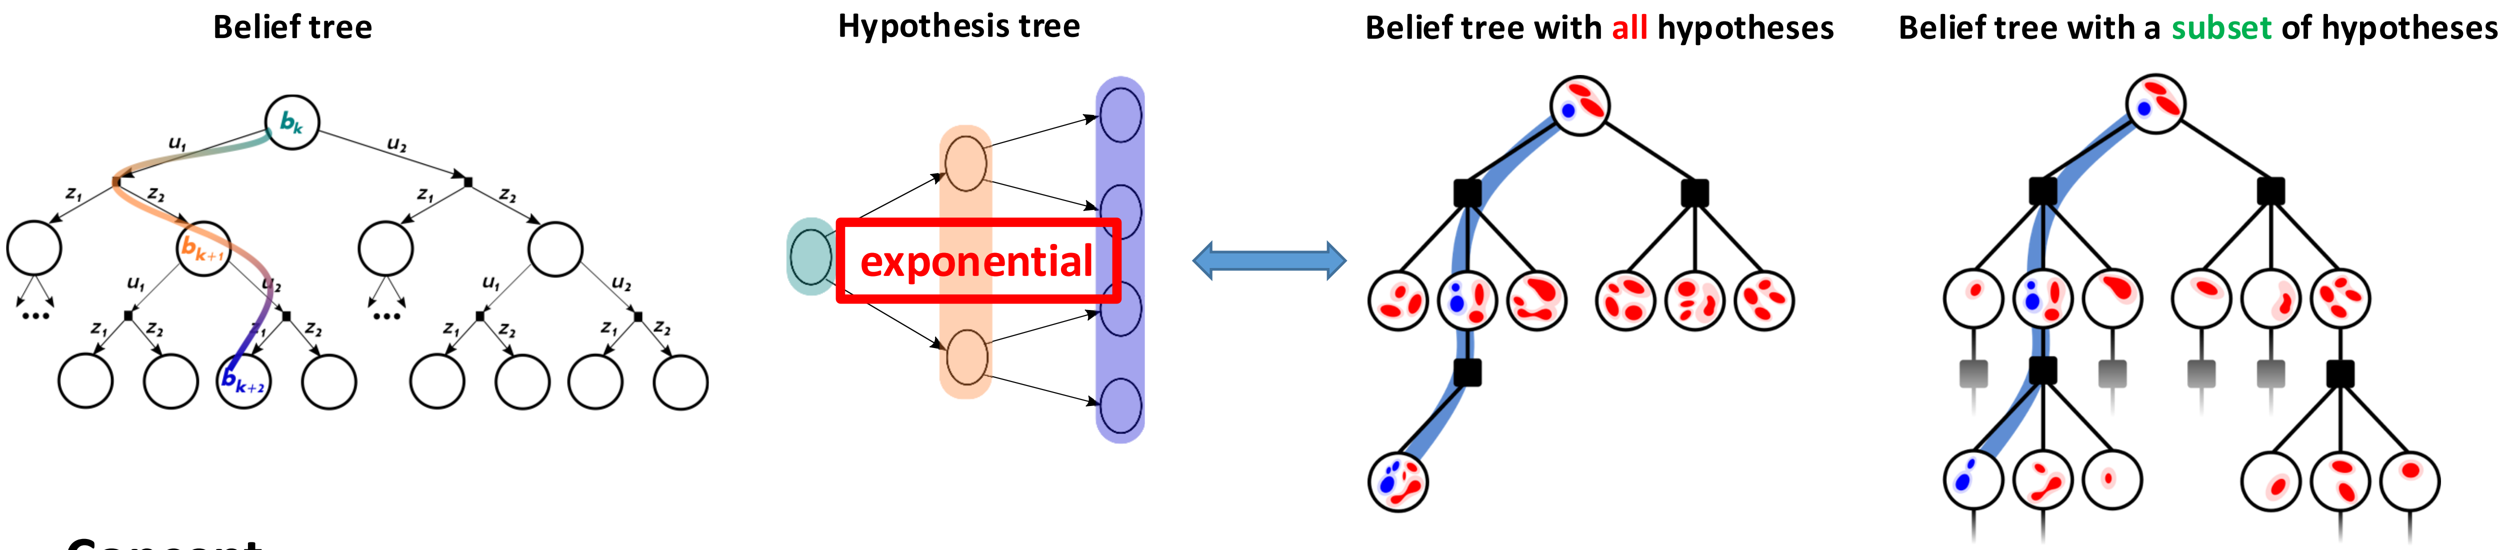

Concept:
- Instead, utilize only a **subset** of hypotheses
- Derive reward bounds, given planning task (reward)
  - Disambiguate between hypotheses
  - Navigate to a goal
  - ..

$$\mathcal{LB}(b_k, \pi) \leq V^{\pi}(b_k) \leq \mathcal{UB}(b_k, \pi)$$

M. Shienman and V. Indelman, "D2A-BSP: Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints," ICRA'22, **Outstanding Paper Award Finalist**.
M. Shienman and V. Indelman, "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints," ISRR'22.
M. Barenboim, M. Shienman, and V. Indelman, "Monte Carlo Planning in Hybrid Belief POMDPs," IEEE RA-L'23.
M. Barenboim, I. Lev-Yehudi, and V. Indelman, "Data Association Aware POMDP Planning with Hypothesis Pruning Performance Guarantees," IEEE RA-L'23.

Simplification of BSP/POMDP with Hybrid Beliefs

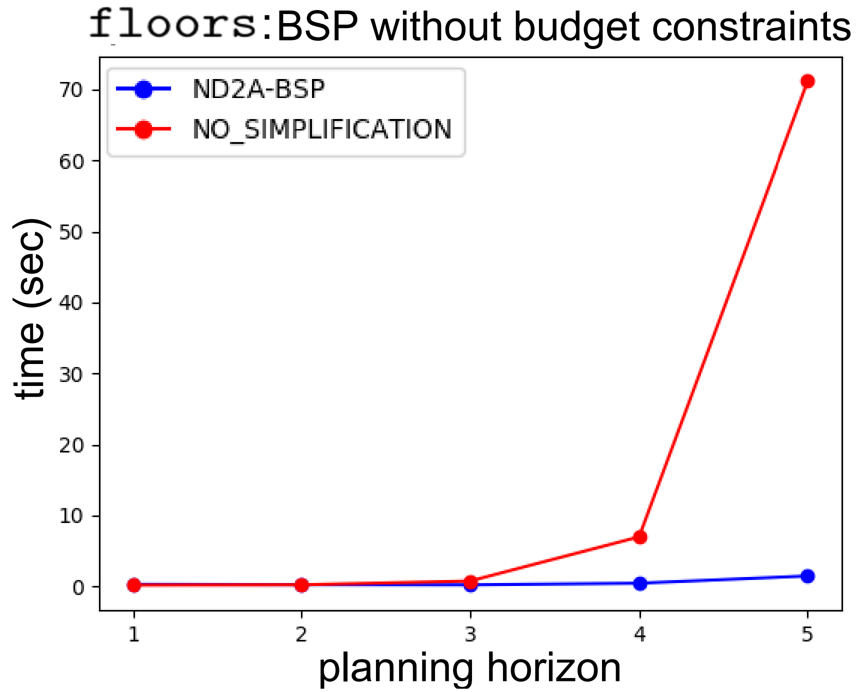

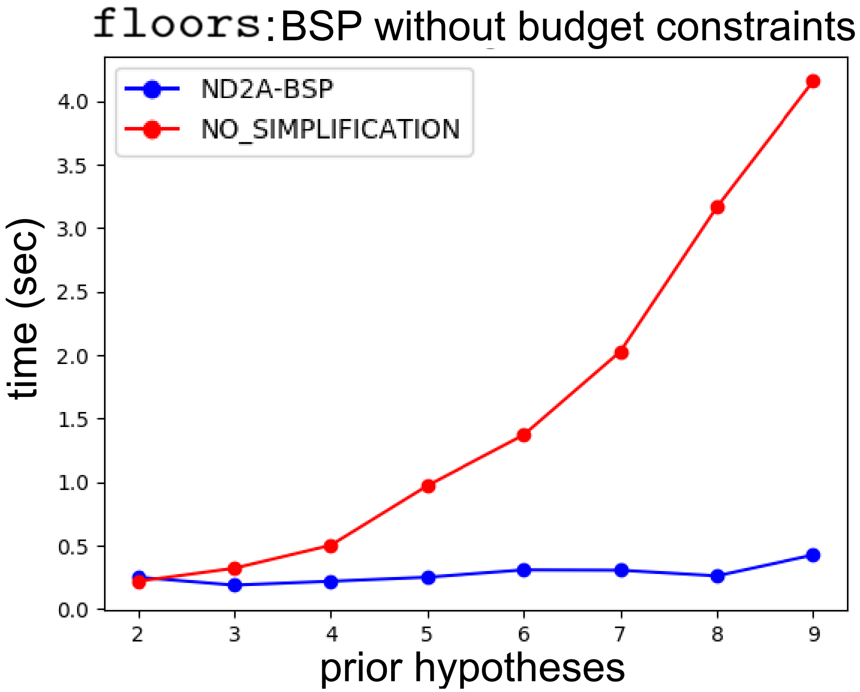

- Significant speed-up in planning
- Same planning performance is **guaranteed** (no overlap between bounds)

M. Shienman and V. Indelman, "D2A-BSP: Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints," ICRA'22, **Outstanding Paper Award Finalist**.
M. Shienman and V. Indelman, "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints," ISRR'22.

Simplification of BSP/POMDP with Hybrid Beliefs

- Derived a deterministic bound to relate the full set of hypotheses to a subset thereof,

Corollary

For any policy π*, and selection of hypotheses set* $\{\beta^i_{0:\mathcal{T}}\}_{i=0}^{|\mathcal{B}|}$ *the following holds,*

$$\left|V^\pi(b_0) - \bar{V}^\pi(\bar{b}_0)\right| \le \mathcal{R}_{max}\left[\mathcal{T}\delta_0^\beta + \sum_{k=1}^{\mathcal{T}}\sum_{\tau=1}^{k}\overline{\mathbb{E}}_{z_{1:\tau}}\left[\delta_\tau^\beta\right]\right].$$

Full tree

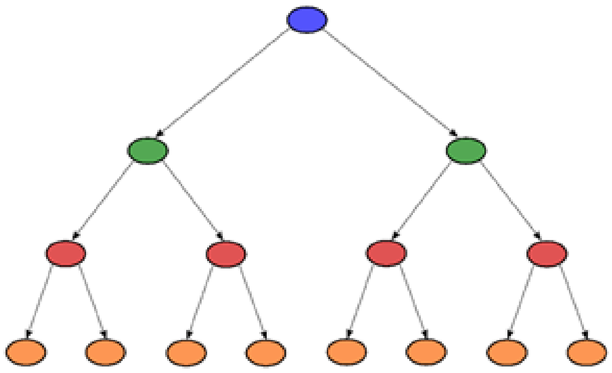

Any subset

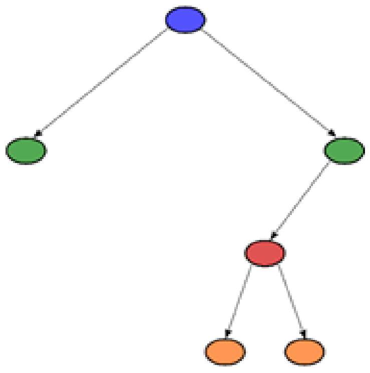

Importantly, the bound relies on the available hypotheses

Can bound the theoretical value with access only to the simplified tree

Bounds can be evaluated online

M. Barenboim, I. Lev-Yehudi, and V. Indelman, "Data Association Aware POMDP Planning with Hypothesis Pruning Performance Guarantees," IEEE RA-L'23.

Simplification of Decision-Making Problems

Concept:

- Identify and solve a **simplified (computationally) easier** decision-making problem
- Provide performance guarantees

Specific simplifications include:

- Sparsification of Gaussian beliefs (high dim. state)
- Topological metric for Gaussian beliefs (high dim. state)
- Utilize a subset of samples (nonparametric beliefs)
- Utilize a subset of hypotheses (hybrid beliefs)
- Simplified models and spaces
- Simplification of policy space
- Simplification of Risk-Averse & Robust Planning
- Simplification in a multi-agent setting

[Indelman RAL'16; Elimelech & Indelman IJRR'22; Sztyglic & Indelman IROS'22, Zhitnikov & Indelman AIJ'22, TRO'24; Shienman & Indelman ICRA'22; Kitanov & Indelman IJRR'24; Zhitnikov et al. IJRR'24, Barenboim & Indelman NIPS'23; Lev-Yehudi, Barenboim & Indelman AAAI'24, Kong & Indelman ISRR 2024]

POMDPs with Deterministic Guarantees

SOTA sampling based approaches come with probabilistic theoretical guarantees

DESPOT

AdaOPS

POMCP

…

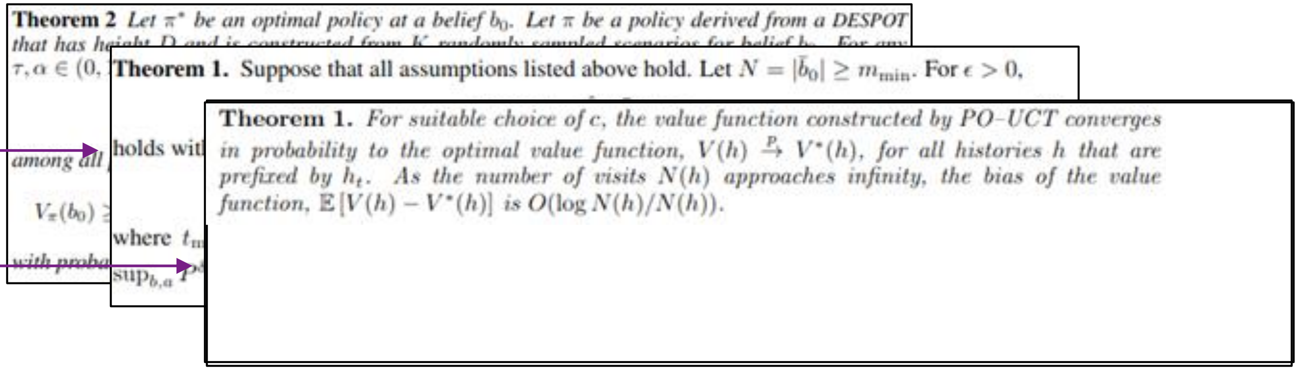

Can we get deterministic guarantees?

We show that deterministic guarantees are indeed possible!

M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Guarantees," NeurIPS'23.
M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Optimality Guarantees," Artificial Intelligence 2026.

Online POMDP Planning with Anytime Deterministic Guarantees

Concept:

Instead of solving the original POMDP, consider a simplified version of that POMDP.

$$\mathcal{M} \rightsquigarrow \bar{\mathcal{M}}$$

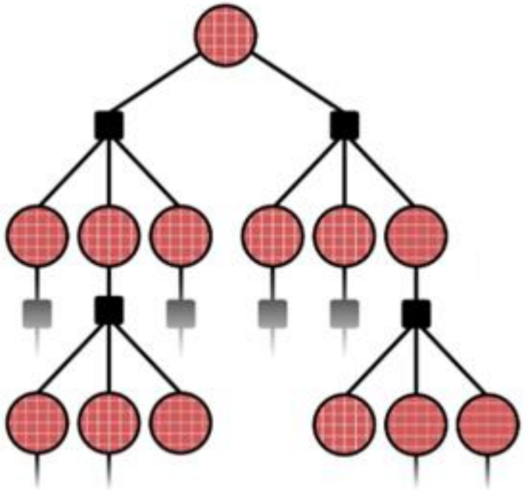

Derive a mathematical relationship between the solution of the simplified, and the theoretical POMDP.

$$V(b) \quad ? \quad \bar{V}(b)$$

M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Guarantees," NeurIPS'23.
M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Optimality Guarantees," Artificial Intelligence 2026.

Online POMDP Planning with Anytime Deterministic Guarantees

- Given a POMDP: $\mathcal{M} = \langle \mathcal{X}, \mathcal{Z}, \mathcal{A}, b_0, \mathcal{P}_T, \mathcal{P}_Z, \rho, \gamma \rangle$
- Define a **simplified** POMDP,

$$\bar{\mathcal{M}} = \langle \underbrace{\bar{\mathcal{X}}, \bar{\mathcal{Z}}}, \mathcal{A}, \underbrace{\bar{b}_0, \bar{\mathcal{P}}_T, \bar{\mathcal{P}}_Z}, \rho, \gamma \rangle$$

$$\bar{\mathcal{X}}(H_t) \subset \mathcal{X}$$
$$\bar{\mathcal{Z}}(H_t) \subset \mathcal{Z}$$

$$\bar{b}_0(x) \triangleq \begin{cases} b_0(x) & , x \in \bar{\mathcal{X}}_0 \\ 0 & , otherwise \end{cases}$$

$$\bar{\mathbb{P}}(x_{t+1} \mid x_t, a_t) \triangleq \begin{cases} \mathbb{P}(x_{t+1} \mid x_t, a_t) & , x_{t+1} \in \bar{\mathcal{X}}(H^-_{t+1}) \\ 0 & , otherwise \end{cases}$$

$$\bar{\mathbb{P}}(z_t \mid x_t) \triangleq \begin{cases} \mathbb{P}(z_t \mid x_t) & , z_t \in \bar{\mathcal{Z}}(H_t) \\ 0 & , otherwise \end{cases}$$

- Simplified value function

$$\bar{V}^{\pi}(\bar{b}_t) \triangleq r(\bar{b}_t, \pi_t) + \bar{\mathbb{E}}_{z_{t+1:\mathcal{T}}} \left[\bar{V}^{\pi}(\bar{b}_{t+1}) \right]$$

M. Barenboim and V. Indelman, “Online POMDP Planning with Anytime Deterministic Guarantees,” NeurIPS’23.
M. Barenboim and V. Indelman, “Online POMDP Planning with Anytime Deterministic Optimality Guarantees,” Artificial Intelligence 2026.

Online POMDP Planning with Anytime Deterministic Guarantees

- Deterministic guarantees (assuming discrete spaces)

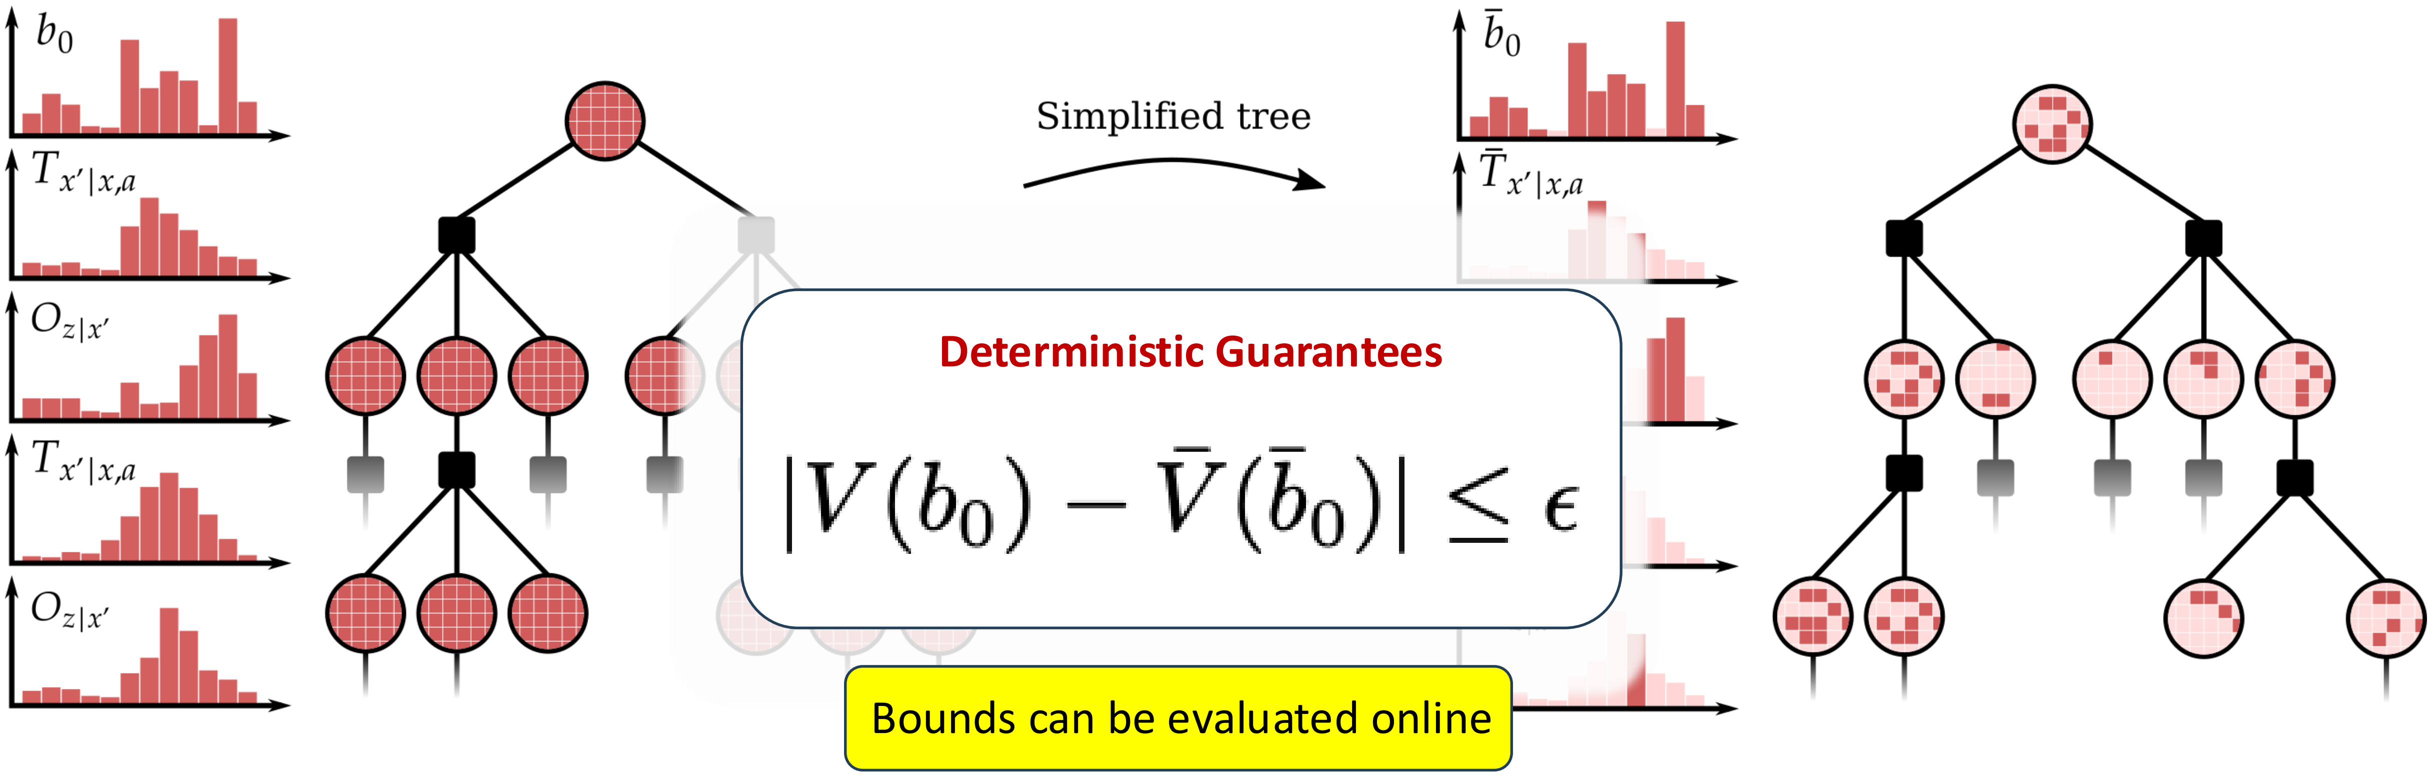

M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Guarantees," NeurIPS'23.
M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Optimality Guarantees," Artificial Intelligence 2026.

Online POMDP Planning with Anytime Deterministic Guarantees

Deterministic Guarantees

$$|V(b_0) - \bar{V}(\bar{b}_0)| \leq \epsilon$$

Importantly, the bounds can be calculated during planning.

How can we use them?

- **Pruning of sub-optimal branches**

  → Made possible by the deterministic guarantees

- **Stopping criteria for the planning phase**

  → Made possible by the deterministic guarantees

- **Finding the optimal solution in finite time**

  → Without recovering the theoretical tree

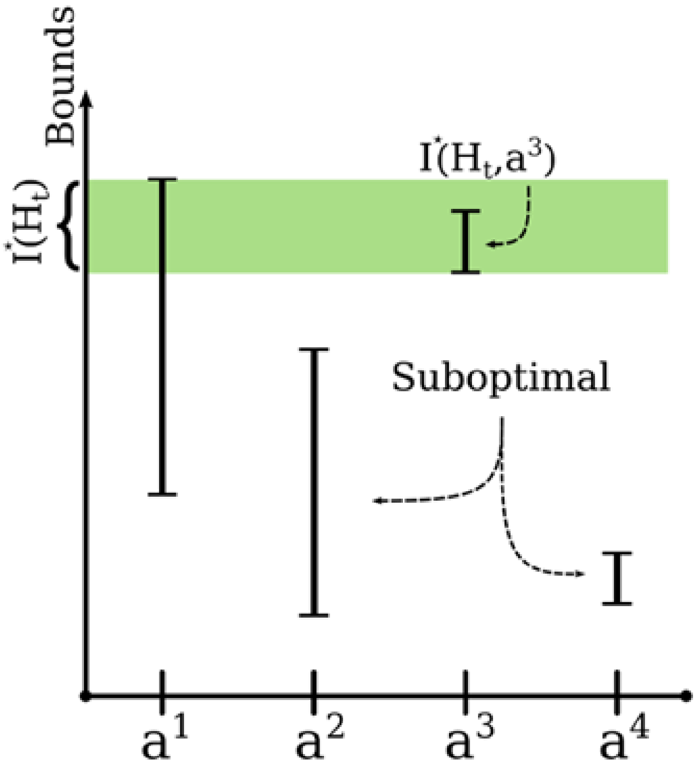

M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Guarantees," NeurIPS'23.
M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Optimality Guarantees," Artificial Intelligence 2026.

Simplifying Complex Observation Models with Probabilistic Guarantees

- We replace the (learned) observation model p_Z with a cheaper model q_Z
  - Simpler GMM, Shallower Neural Network, etc.
  - Example:

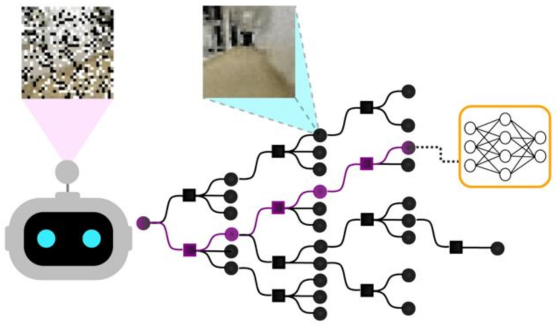

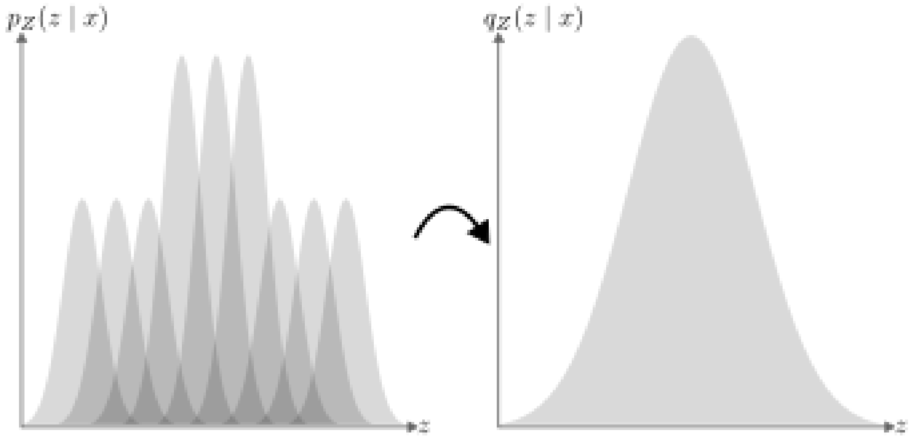

Simplified models

$p_\theta(z \mid x)$

Original, **expensive**

$q_\phi(z \mid x)$

Simplified, **cheap**

Can we simplify the learned models?
What is the impact on planning performance?

I. Lev-Yehudi, M. Barenboim, and V. Indelman, "Simplifying Complex Observation Models in Continuous POMDP Planning with Probabilistic Guarantees and Practice," AAAI'24.

Simplifying Complex Observation Models with Probabilistic Guarantees

- We replace the (learned) observation model p_Z with a cheaper model q_Z
- Simplified action-value function: $Q_{\mathbf{P}}^{q_Z}$

Corollary 3

For arbitrary $\varepsilon, \delta > 0$ there exists a number of particles for which

$$|Q_{\mathbf{P}}^{p_Z}(b_t, a) - \hat{Q}_{\mathbf{M_P}}^{q_Z}(\bar{b}_t, a)| \leq \hat{\Phi}_{\mathbf{M_P}}(\bar{b}_t, a) + \varepsilon$$

with probability of at least $1 - \delta$ for any guaranteed planner

Theoretical Q function of the POMDP, with **original** models

Estimator of the Q function of a **particle-belief** POMDP, with **simplified** models

I. Lev-Yehudi, M. Barenboim, and V. Indelman, "Simplifying Complex Observation Models in Continuous POMDP Planning with Probabilistic Guarantees and Practice," AAAI'24.

Partitioning of a Multivariate Observation Space

- Consider a **multivariate** random variable $Z \in \mathcal{Z}$, that represents future observations:

$$Z = (Z^1, Z^2, \ldots, Z^m)$$

- Examples:

Raw measurement of an image sensor

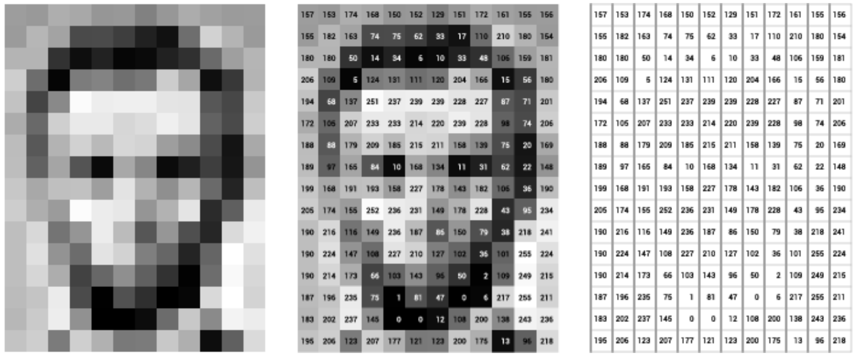

Factor graph

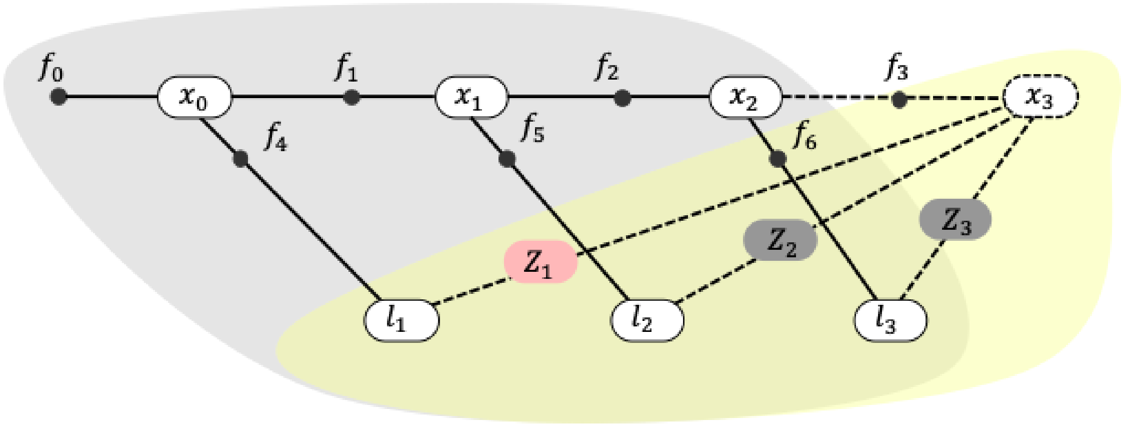

T. Yotam and V. Indelman, “Measurement Simplification in ρ-POMDP with Performance Guarantees,” IEEE T-RO’24.

Partitioning of a Multivariate Observation Space

- Consider a **multivariate** random variable $Z \in \mathcal{Z}$, that represents future observations:

$$Z = (Z^1, Z^2, \ldots, Z^m)$$

- We can partition $Z \in \mathcal{Z}$ into different subsets/components, e.g.

$$Z^s = \{Z^1, Z^2, \ldots, Z^n\}$$
$$Z^{\bar{s}} = \{Z^{n+1}, Z^{n+2}, \ldots, Z^m\}$$

$$Z = Z^s \cup Z^{\bar{s}}$$

- Hierarchical Partitioning:

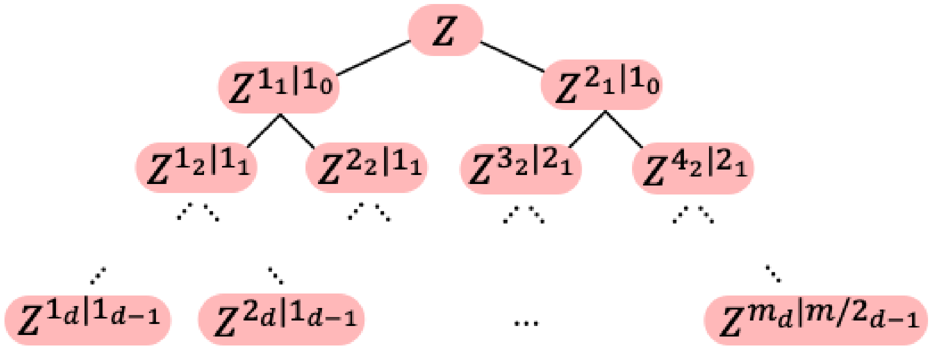

T. Yotam and V. Indelman, “Measurement Simplification in ρ-POMDP with Performance Guarantees,” IEEE T-RO’24.

Partitioning of a Multivariate Observation Space

$$\mathcal{LB} \leq \mathcal{H}(X|Z) \leq \mathcal{UB}$$

$$\mathcal{LB} \triangleq \mathcal{H}(Z^s \mid X) + \mathcal{H}(Z^{\bar{s}}|X) - \mathcal{H}(Z^s) - \mathcal{H}(Z^{\bar{s}}) + \mathcal{H}(X)$$

$$\mathcal{UB} \triangleq \mathcal{H}\left(Z^s|X\right) + \mathcal{H}\left(X\right) - \mathcal{H}\left(Z^s\right)$$

$$\mathcal{H}(Z^s, Z^{\bar{s}}) = \mathcal{H}(Z^s) + \mathcal{H}(Z^{\bar{s}}) - \mathcal{I}(Z^s; Z^{\bar{s}})$$

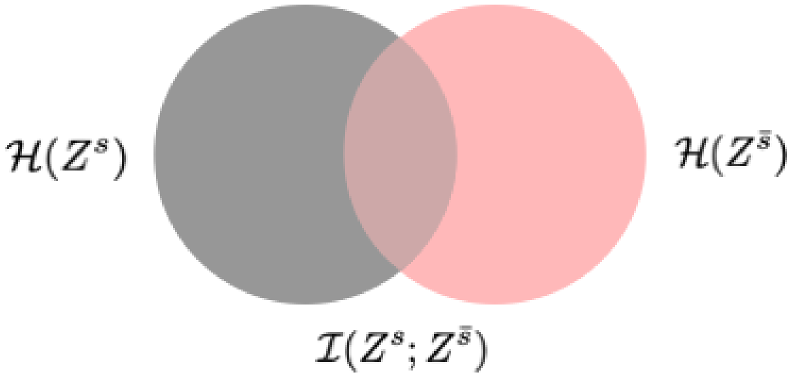

T. Yotam and V. Indelman, “Measurement Simplification in ρ-POMDP with Performance Guarantees,” IEEE T-RO’24.

Partitioning of a Multivariate Observation Space

Application to Active SLAM

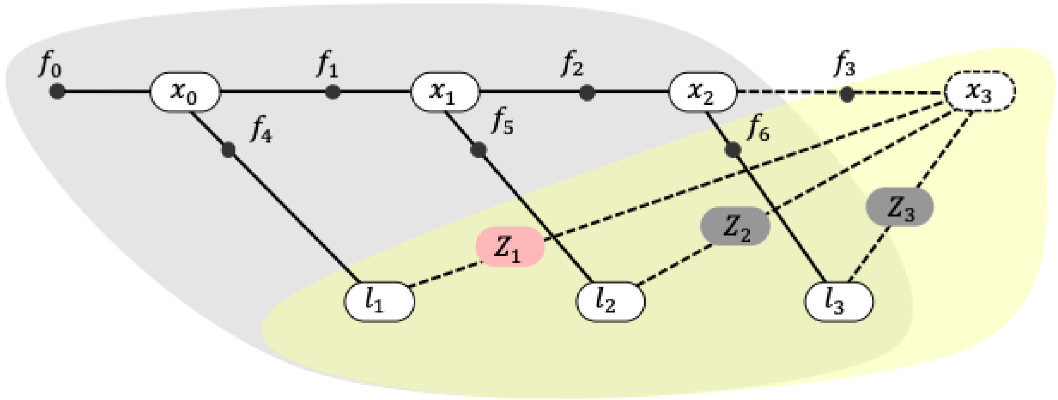

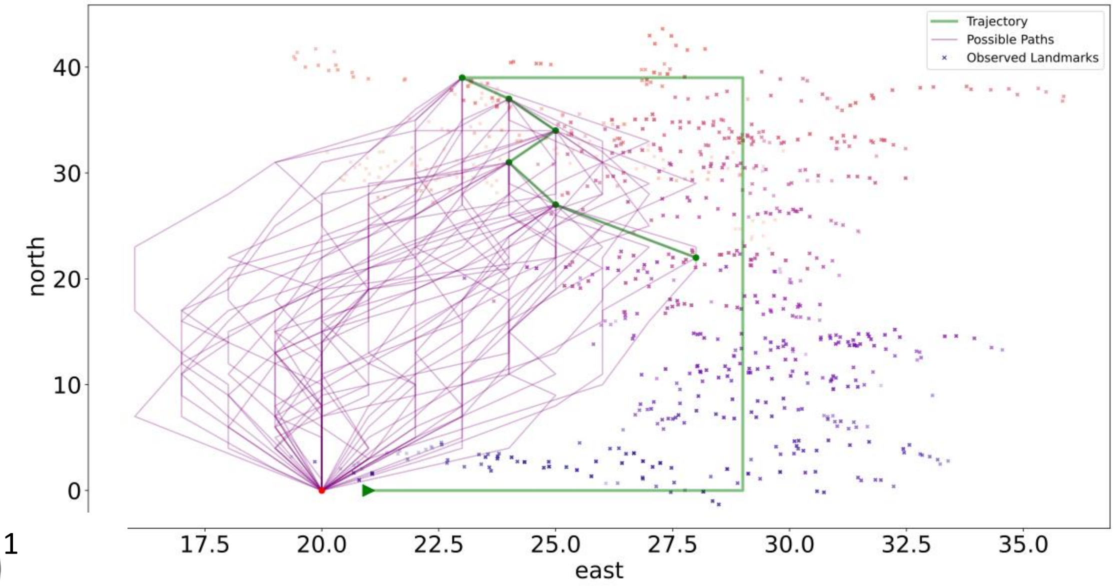

# Paths	# Factors	RP	rAMDL[2]	MP (ours)[1]
100	2956	No	11.521 ± 0.537	$\mathbf{6.888 \pm 0.155}$
100	2956	Yes	24.636 ± 1.381	$\mathbf{11.758 \pm 0.372}$
100	5904	Yes	84.376 ± 14.458	$\mathbf{32.069 \pm 4.913}$

Table: Total planning time in seconds (lower is better)

[1]T. Yotam and V. Indelman, “Measurement Simplification in ρ-POMDP with Performance Guarantees,” IEEE T-RO’24.
[2]D. Kopitkov and V. Indelman, “No Belief Propagation Required: Belief Space Planning in High-Dimensional State Spaces via Factor Graphs, Matrix Determinant Lemma and Re-use of Calculation,” IJRR’17.
[2]D. Kopitkov and V. Indelman, “General-purpose incremental covariance update and efficient belief space planning via a factor-graph propagation action tree”, IJRR’19.

Real Worl Experiment - Visual Active SLAM

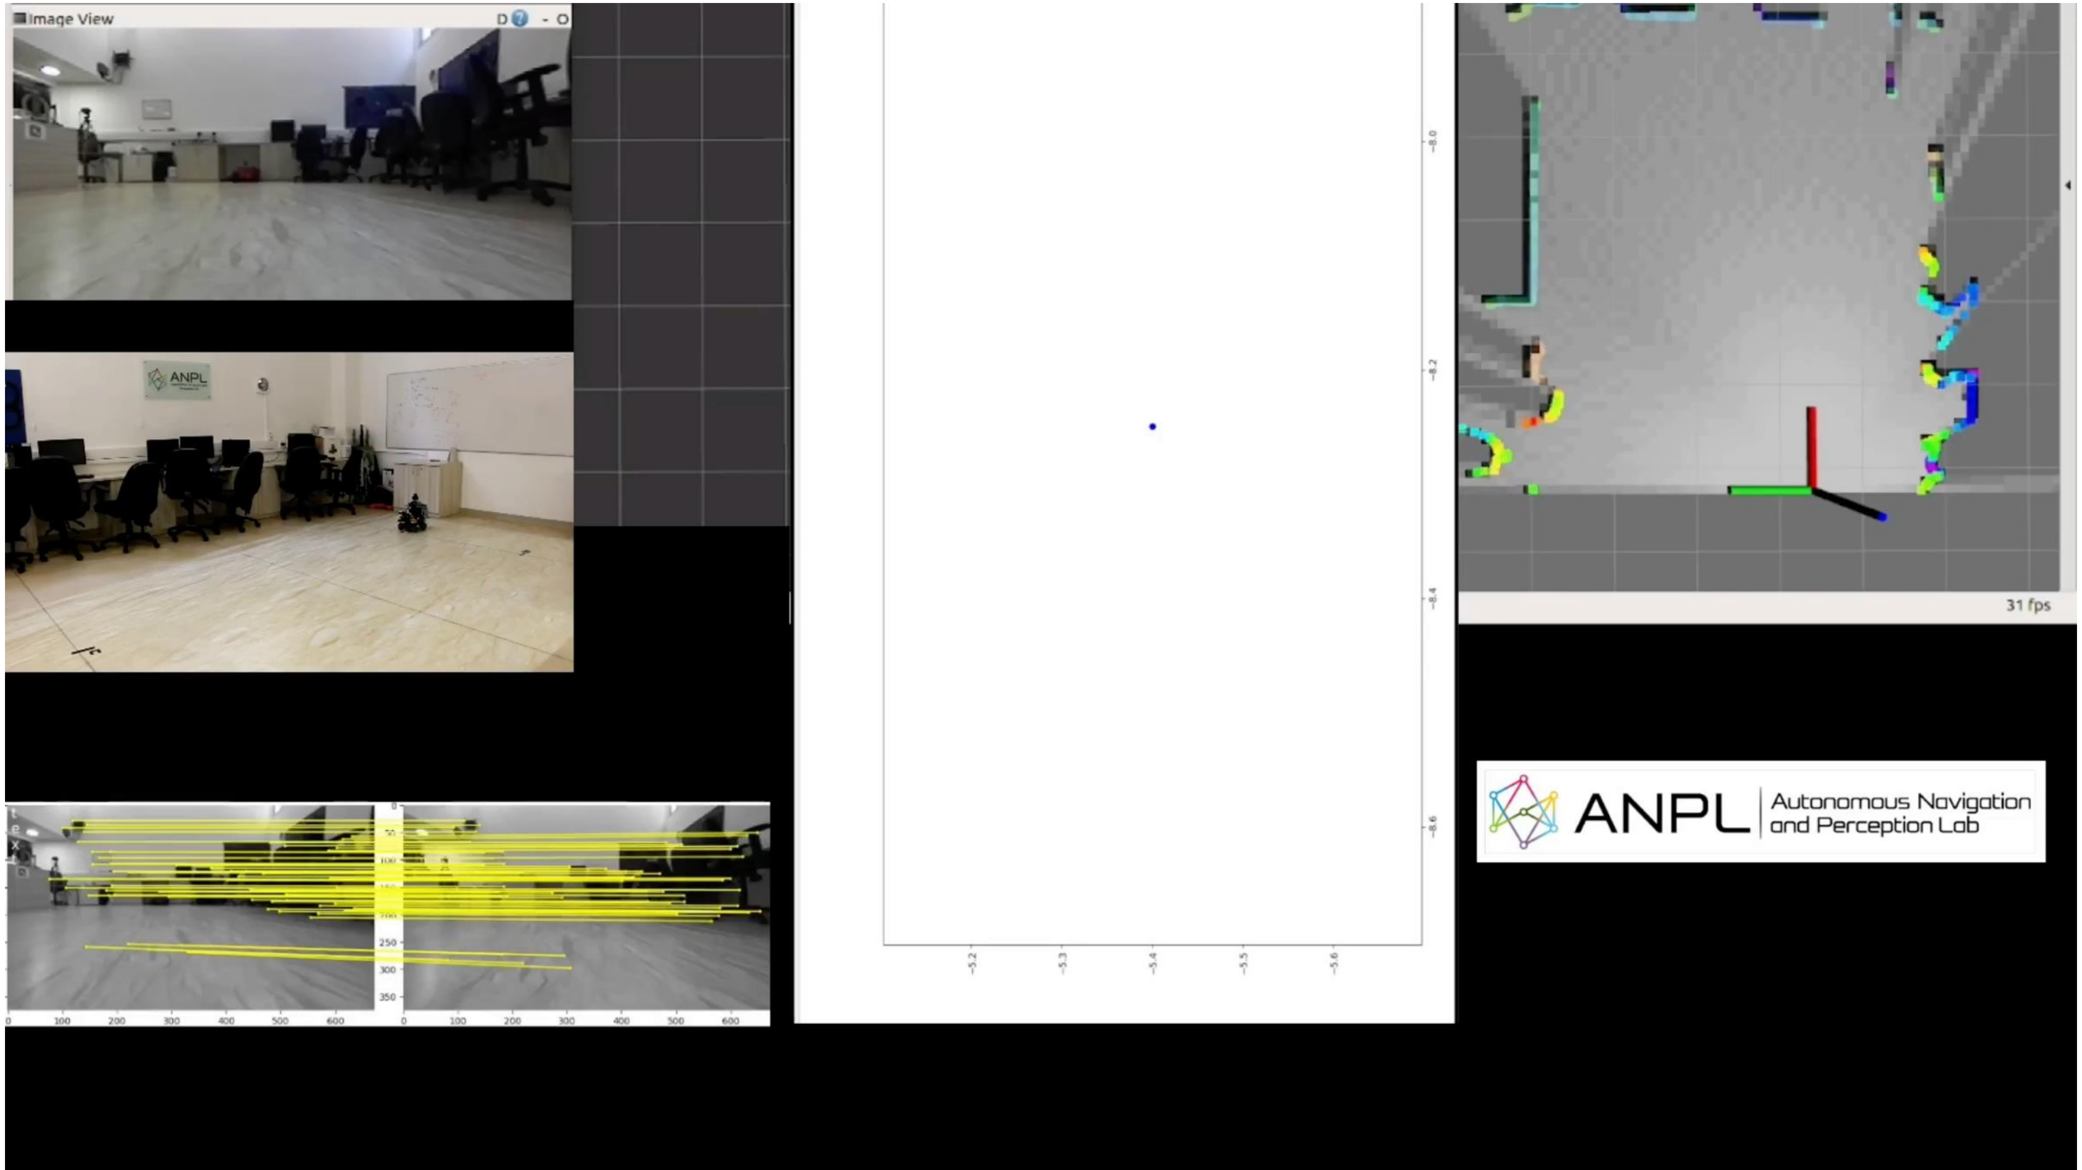

Partitioning of a Multivariate Observation Space

Application to Active SLAM

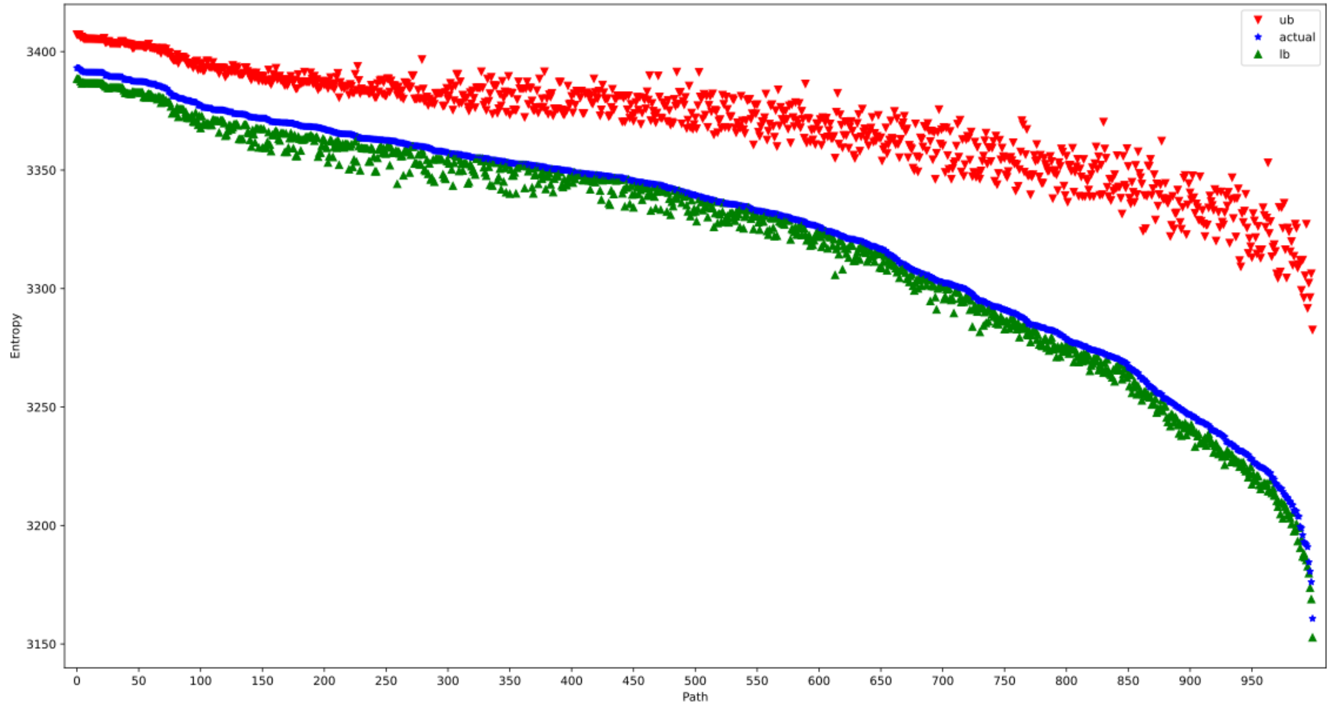

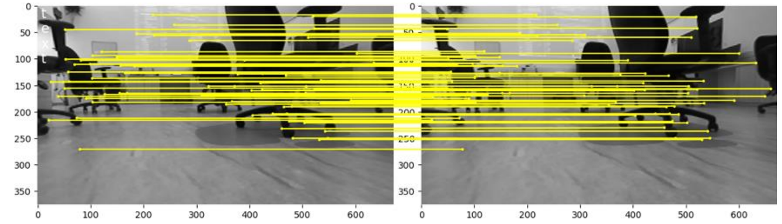

Method	time [sec]
MP (ours)[1]	$\mathbf{585.507 \pm 27.153}$
rAMDL[2]	802.545 ± 25.651
iSAM2[3]	1764.835 ± 26.521

Table: Total planning time in seconds (lower is better)

[1]T. Yotam and V. Indelman, "Measurement Simplification in ρ-POMDP with Performance Guarantees," IEEE T-RO'24.
[2]D. Kopitkov and V. Indelman, "No Belief Propagation Required: Belief Space Planning in High-Dimensional State Spaces via Factor Graphs, Matrix Determinant Lemma and Re-use of Calculation," IJRR'17.
[2]D. Kopitkov and V. Indelman, "General-purpose incremental covariance update and efficient belief space planning via a factor-graph propagation action tree", IJRR'19.
[3]M. Kaess, et al., "iSAM2: Incremental smoothing and mapping using the Bayes tree," IJRR'12.

Simplification of Decision-Making Problems

Concept:

- Identify and solve a **simplified (computationally) easier** decision-making problem
- Provide performance guarantees

Specific simplifications include:

- Sparsification of Gaussian beliefs (high dim. state)
- Topological metric for Gaussian beliefs (high dim. state)
- Utilize a subset of samples (nonparametric beliefs)
- Utilize a subset of hypotheses (hybrid beliefs)
- Simplified models and spaces
- Simplification of policy space
- Simplification of Risk-Averse & Robust Planning
- Simplification in a multi-agent setting

[Indelman RAL'16; Elimelech & Indelman IJRR'22; Sztyglic & Indelman IROS'22, Zhitnikov & Indelman AIJ'22, TRO'24; Shienman & Indelman ICRA'22; Kitanov & Indelman IJRR'24; Zhitnikov et al. IJRR'24, Barenboim & Indelman NIPS'23; Lev-Yehudi, Barenboim & Indelman AAAI'24, Kong & Indelman ISRR 2024]

Simplification of Risk Averse POMDP Planning

- **Distribution** over returns/rewards

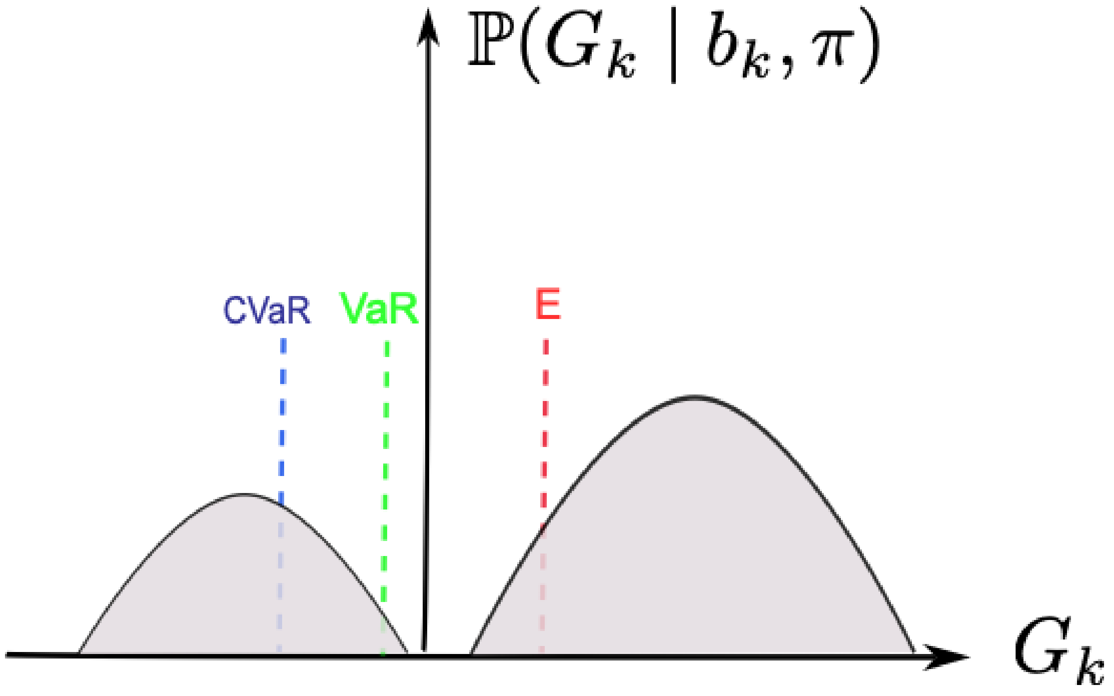

$$G_k = \sum_{t=k}^{L} r(b_t, a_t)$$

return

$$V^{\pi}(b_k) = \varphi\left(\mathbb{P}(G_k \mid b_k, \pi)\right)$$

Risk measure (e.g. CVaR)

A. Zhitnikov and V. Indelman, “Simplified Risk Aware Decision Making with Belief Dependent Rewards in Partially Observable Domains,” Artificial Intelligence, 2022.
Y. Pariente and V. Indelman, “Simplification of Risk Averse POMDPs with Performance Guarantees,” arXiv’24.
I. Nutov and V. Indelman, “Simplified Risk Aware CVaR-based POMDP With Performance Guarantees: a Risk Envelope Perspective”, TR’24.
Y. Pariente and V. Indelman, “Bounding Conditional Value-at-Risk via Auxiliary Distributions with Bounded Discrepancies,” arXiv’25.

Simplification of Risk Averse POMDP Planning

- Impact of simplification on **distribution** over returns/rewards
- Simplified **risk-aware** decision-making with formal guarantees

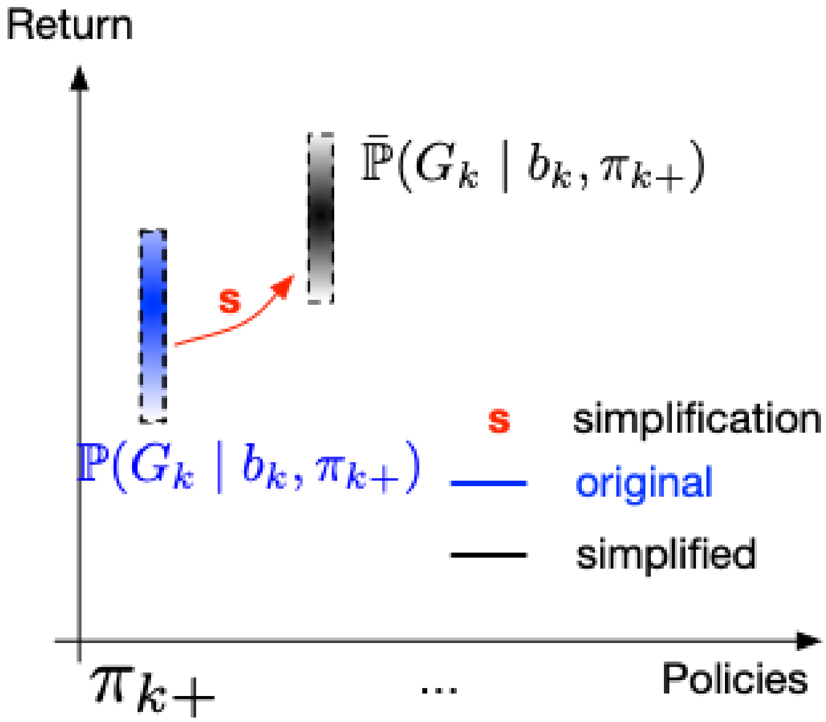

$$V^\pi(b_k) = \varphi\left(\mathbb{P}(G_k \mid b_k, \pi)\right)$$

Risk measure (e.g. CVaR)

$$\underline{V}^\pi(b_k) \leq V^\pi(b_k) \leq \bar{V}^\pi(b_k)$$

A. Zhitnikov and V. Indelman, “Simplified Risk Aware Decision Making with Belief Dependent Rewards in Partially Observable Domains,” Artificial Intelligence, 2022.
Y. Pariente and V. Indelman, “Simplification of Risk Averse POMDPs with Performance Guarantees,” arXiv’24.
I. Nutov and V. Indelman, “Simplified Risk Aware CVaR-based POMDP With Performance Guarantees: a Risk Envelope Perspective”, TR’24.
Y. Pariente and V. Indelman, “Bounding Conditional Value-at-Risk via Auxiliary Distributions with Bounded Discrepancies,” arXiv’25.

Probabilistically Constrained Belief Space Planning

$$\max_{\pi_{k+}} \mathbb{E}\left[\sum_{\ell=k}^{k+L-1} \rho_{\ell+1} \middle| b_k, \pi_{k+} \right]$$

$$\text{subject to } \mathrm{P}(c(b_{k:k+L}; \phi, \delta) = 1 | b_k, \pi_{k+}) \geq 1 - \epsilon$$

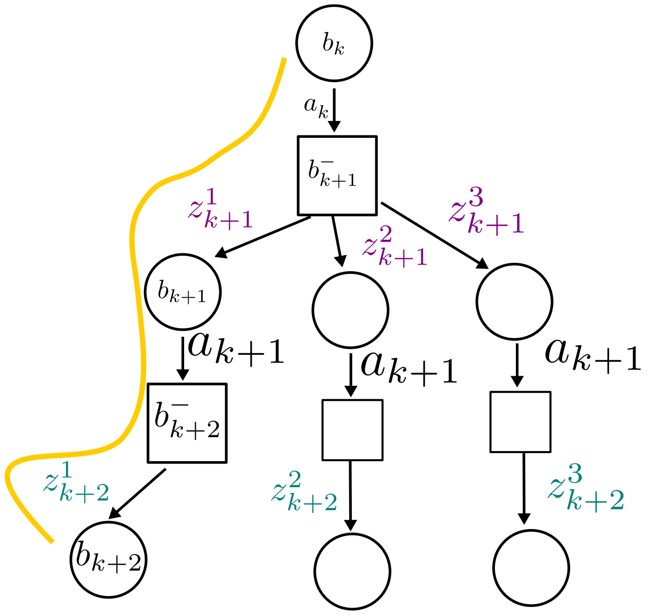

Information gain[1]:

$$c(b_{k:k+L}; \phi, \delta) \triangleq \mathbf{1}_{\{\left(\sum_{\ell=k}^{k+L-1} \phi(b_t, b_{t+1})\right) \geq \delta\}}(b_{k:k+L})$$

Information gain

Safety[2]:

$$c(b_{k:k+L}; \phi, \delta) \triangleq \prod_{\ell=k}^{k+L} \mathbf{1}_{\{b_\ell : \phi(b_\ell) \geq \delta\}}(b_\ell)$$

Safety

[1]A. Zhitnikov and V. Indelman, “Simplified Continuous High Dimensional Belief Space Planning with Adaptive Probabilistic Belief-dependent Constraints,” T-RO’24.
[2]A. Zhitnikov and V. Indelman, “Anytime Probabilistically Constrained Provably Convergent Online Belief Space Planning,” T-RO’25.

Robust Online Planning Under Uncertainty

- So far, models were assumed to be given and perfect
- In practice, models are learned from data
- What happens when the models are **uncertain**?

How to do **online robust** planning?

Uncertainty set:

$$P_t(S_{t+1} \mid S_t = s, A_t = a) \in \mathcal{P}_t^{s,a}$$

Robust value function:

$$V^{\pi}(s) = \min_{P \in \mathcal{P}} V^{\pi,P}(s)$$

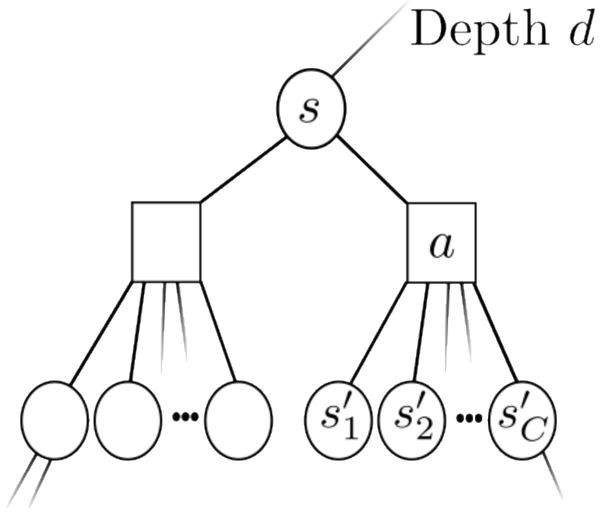

Robust Sparse Sampling (RSS) Algorithm:

- A sample-based online robust planner
- Applicable to infinite or continuous state spaces
- Finite-sample performance guarantees

T. Shazman, I. Lev-Yehudi, R. Benchetrit, and V. Indelman, "Online Robust Planning under Model Uncertainty: A Sample-Based Approach," *AAAI'26*.

Robust Online Planning Under Uncertainty

- So far, models were assumed to be given and perfect
- In practice, models are learned from data
- What happens when the models are **uncertain**?

How to do **online robust** planning?

Uncertainty set:

$$P_t(S_{t+1} \mid S_t = s, A_t = a) \in \mathcal{P}_t^{s,a}$$

Robust value function:

$$V^{\pi}(s) = \min_{P \in \mathcal{P}} V^{\pi,P}(s)$$

Robust Sparse Sampling (RSS) Algorithm:

- A sample-based online robust planner
- Applicable to infinite or continuous state spaces
- Finite-sample performance guarantees

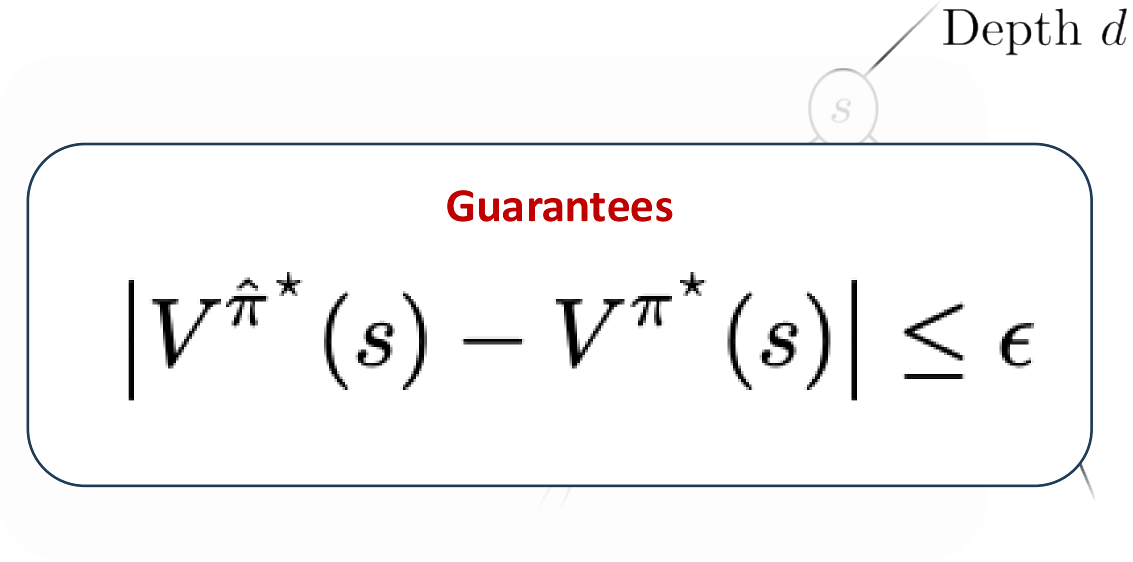

Guarantees

$$\left|V^{\hat{\pi}^\star}(s) - V^{\pi^\star}(s)\right| \leq \epsilon$$

T. Shazman, I. Lev-Yehudi, R. Benchetrit, and V. Indelman, "Online Robust Planning under Model Uncertainty: A Sample-Based Approach," *AAAI'26*.

Agenda

Experience Reuse in POMDP Planning

POMDP Planning with Hybrid Beliefs

Simplification of POMDP with Formal Guarantees

Multi-agent POMDP Planning with Inconsistent Beliefs

Simplification of Decision-Making Problems

Concept:

- Identify and solve a **simplified (computationally) easier** decision-making problem
- Provide performance guarantees

Specific simplifications include:

- Sparsification of Gaussian beliefs (high dim. state)
- Topological metric for Gaussian beliefs (high dim. state)
- Utilize a subset of samples (nonparametric beliefs)
- Utilize a subset of hypotheses (hybrid beliefs)
- Simplified models and spaces
- Simplification of policy space
- Simplification of Risk-Averse & Robust Planning
- Simplification in a multi-agent setting

[Indelman RAL'16; Elimelech & Indelman IJRR'22; Sztyglic & Indelman IROS'22, Zhitnikov & Indelman AIJ'22, TRO'24; Shienman & Indelman ICRA'22; Kitanov & Indelman IJRR'24; Zhitnikov et al. IJRR'24, Barenboim & Indelman NIPS'23; Lev-Yehudi, Barenboim & Indelman AAAI'24, Kong & Indelman ISRR 2024]

Multi-Robot Belief Space Planning

- **A common assumption**: Beliefs of different robots are consistent at planning time
- Requires prohibitively frequent data-sharing capabilities!

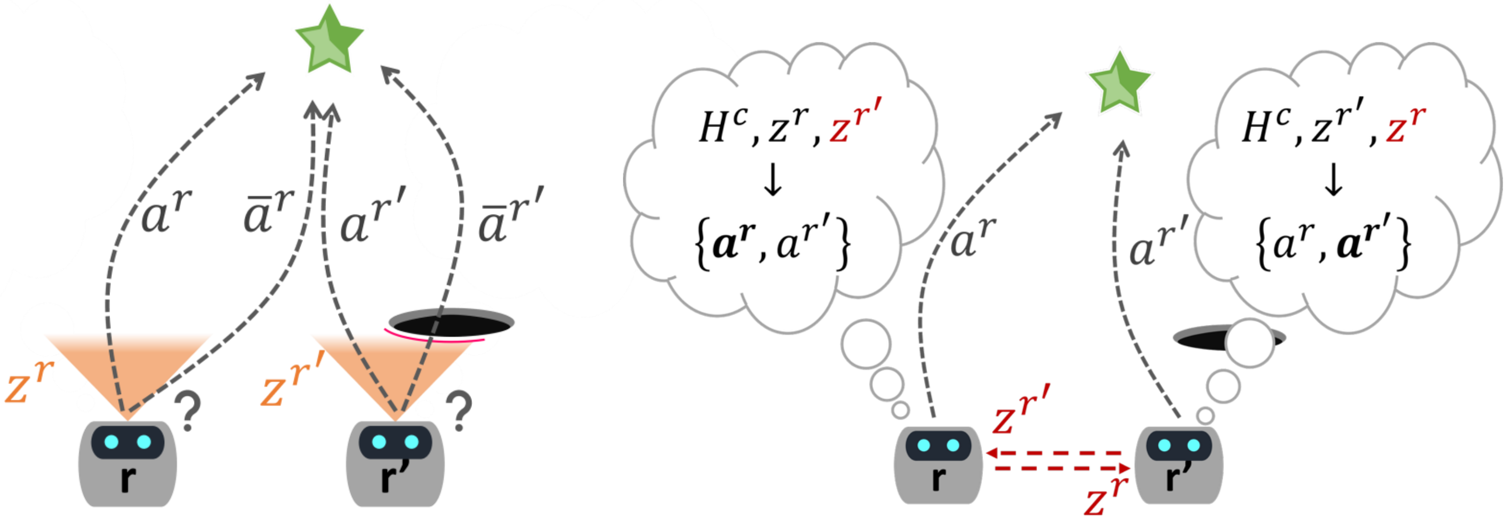

Multi-Robot Cooperative BSP with Inconsistent Beliefs

What happens when data-sharing capabilities between the robots are limited?

- Histories & beliefs of the robots may **differ** due to limited data-sharing capabilities

$$b_k^r = \mathbb{P}(x_k \mid \mathcal{H}_k^r) \qquad b_k^{r'} = \mathbb{P}(x_k \mid \mathcal{H}_k^{r'}) \qquad \mathcal{H}_k^r \neq \mathcal{H}_k^{r'}$$

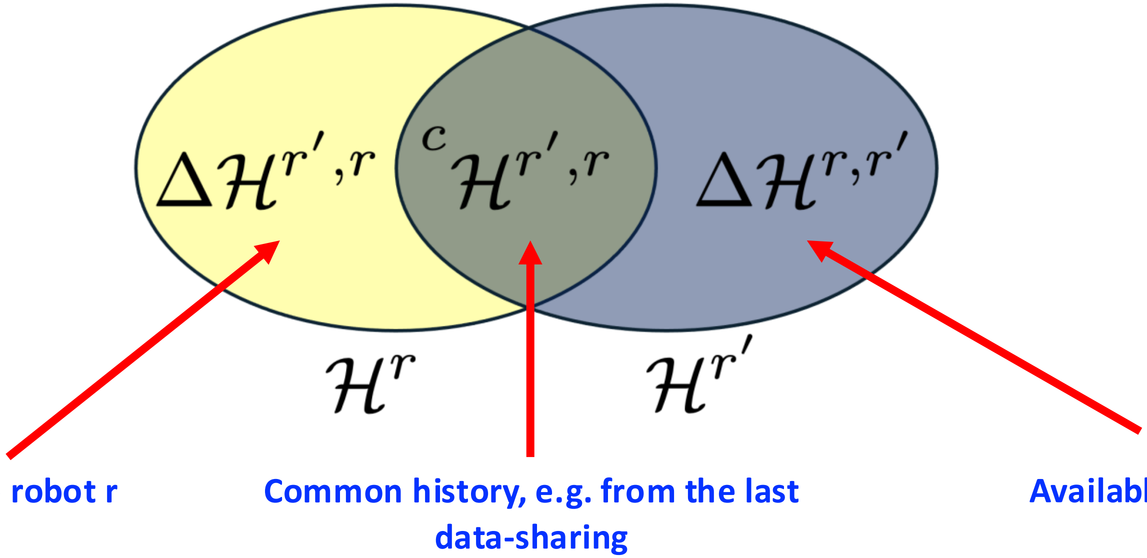

T. Kundu, M. Rafaeli, and V. Indelman, "Multi-Robot Communication-Aware Cooperative Belief Space Planning with Inconsistent Beliefs: An Action-Consistent Approach," IROS'24.
T. Kundu, M. Rafaeli, A. Gulyaev, and V. Indelman, "Action-Consistent Decentralized Belief Space Planning with Inconsistent Beliefs and Limited Data Sharing: Framework and Simplification Algorithms with Formal Guarantees," Submitted, 2025.
M. Rafaeli, and V. Indelman, "Towards Optimal Performance and Action Consistency Guarantees in Dec-POMDPs with Inconsistent Beliefs and Limited Communication," Submitted, 2025.

Multi-Robot Cooperative BSP with Inconsistent Beliefs

What happens when data-sharing capabilities between the robots are limited?

- Histories & beliefs of the robots may **differ** due to limited data-sharing capabilities

$$b_k^r = \mathbb{P}(x_k \mid \mathcal{H}_k^r) \qquad b_k^{r'} = \mathbb{P}(x_k \mid \mathcal{H}_k^{r'}) \qquad \mathcal{H}_k^r \neq \mathcal{H}_k^{r'}$$

- **Can lead to a lack of coordination and unsafe and sub-optimal actions**

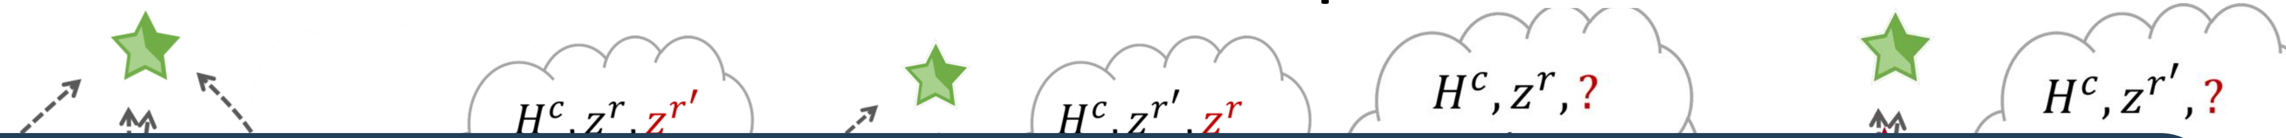

Challenge:

- **Guarantee** a consistent joint action selection by individual robots **despite** inconsistent histories
- Otherwise, self-trigger communication

T. Kundu, M. Rafaeli, and V. Indelman, "Multi-Robot Communication-Aware Cooperative Belief Space Planning with Inconsistent Beliefs: An Action-Consistent Approach," IROS'24.
T. Kundu, M. Rafaeli, A. Gulyaev, and V. Indelman, "Action-Consistent Decentralized Belief Space Planning with Inconsistent Beliefs and Limited Data Sharing: Framework and Simplification Algorithms with Formal Guarantees," Submitted, 2025.
M. Rafaeli, and V. Indelman, "Towards Optimal Performance and Action Consistency Guarantees in Dec-POMDPs with Inconsistent Beliefs and Limited Communication," Submitted, 2025.

Action Consistency

- If two decision-making problems have the **same action preference**, this implies both have the **same best action regardless of the actual objective/value function values**

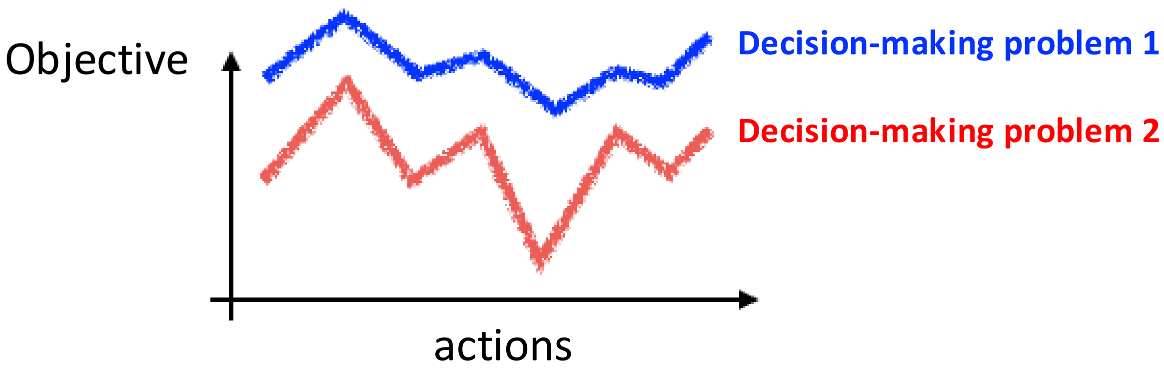

- **Key idea**: to guarantee consistent multi-robot decision-making, each robot
  - reasons about its own and other robots' action preferences while accounting for the **missing information** between the robots
  - checks if for all these realizations, we get the same best joint action

V. Indelman, "No Correlations Involved: Decision Making Under Uncertainty in a Conservative Sparse Information Space," IEEE RA-L'16.
K. Elimelech and V. Indelman, "Simplified decision making in the belief space using belief sparsification," IJRR'22.
A. Kitanov and V. Indelman, "Topological Belief Space Planning for Active SLAM with Pairwise Gaussian Potentials and Performance Guarantees," IJRR'24.

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

- From the perspective of robot r, MR-AC holds if the selected joint actions are the same based on:
  1. Its local information
  2. What it perceives about the reasoning of the other robot r'
  3. What it perceives about the reasoning of itself perceived by the other robot r'

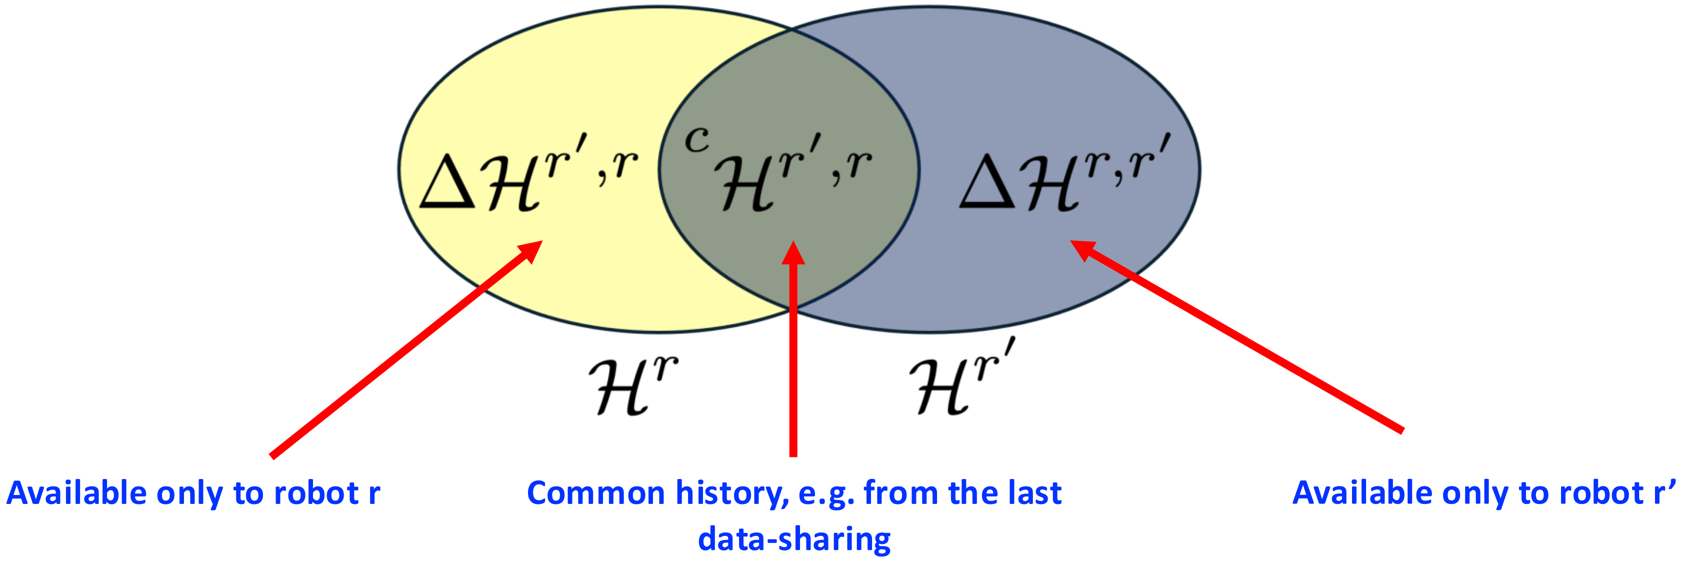

T. Kundu, M. Rafaeli, and V. Indelman, "Multi-Robot Communication-Aware Cooperative Belief Space Planning with Inconsistent Beliefs: An Action-Consistent Approach," IROS'24.
T. Kundu, M. Rafaeli, A. Gulyaev, and V. Indelman, "Action-Consistent Decentralized Belief Space Planning with Inconsistent Beliefs and Limited Data Sharing: Framework and Simplification Algorithms with Formal Guarantees," Submitted, 2025.
M. Rafaeli, and V. Indelman, "Towards Optimal Performance and Action Consistency Guarantees in Dec-POMDPs with Inconsistent Beliefs and Limited Communication," Submitted, 2025.

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

- From the perspective of robot r, MR-AC holds if the selected joint actions are the same based on:

  1. Its local information
  2. What it perceives about the reasoning of the other robot r'
  3. What it perceives about the reasoning of itself perceived by the other robot r'

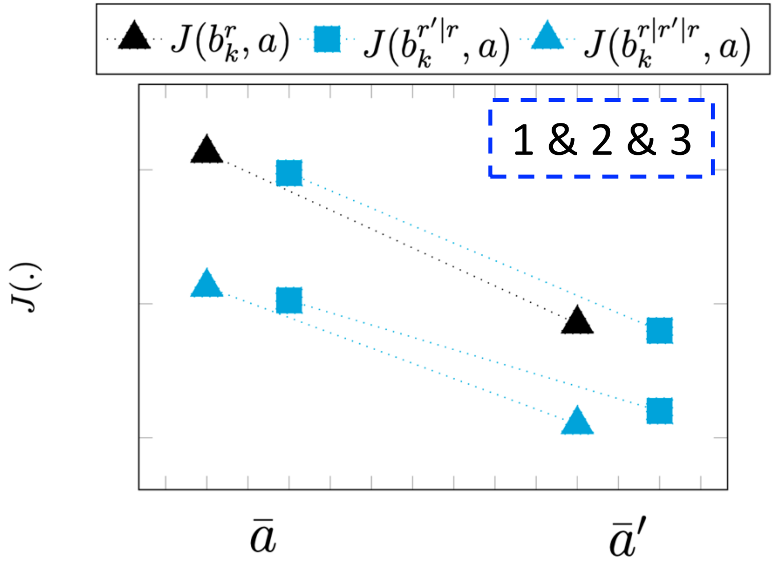

- Same best action in all cases?
  - **Yes**: MR-AC is **guaranteed** to be satisfied
    - Robots are **guaranteed** to choose the same joint action
    - **No further data sharing** is needed!
  - **No**: self-trigger communication, share some data, repeat Steps 1-3

T. Kundu, M. Rafaeli, and V. Indelman, "Multi-Robot Communication-Aware Cooperative Belief Space Planning with Inconsistent Beliefs: An Action-Consistent Approach," IROS'24.
T. Kundu, M. Rafaeli, A. Gulyaev, and V. Indelman, "Action-Consistent Decentralized Belief Space Planning with Inconsistent Beliefs and Limited Data Sharing: Framework and Simplification Algorithms with Formal Guarantees," Submitted 2025.
M. Rafaeli, and V. Indelman, "Towards Optimal Performance and Action Consistency Guarantees in Dec-POMDPs with Inconsistent Beliefs and Limited Communication," Submitted, 2025.

Agenda

- Experience Reuse in POMDP Planning
- POMDP Planning with Hybrid Beliefs
- Simplification of POMDP with Formal Guarantees
- Multi-agent POMDP Planning with Inconsistent Beliefs

Semantic Risk Awareness

Ambiguous Environments

Thank You

TECHNION Israel Institute of Technology

ANPL Autonomous Navigation and Perception Lab

ISRAEL SCIENCE FOUNDATION

NSF

BSF United States – Israel Binational Science Foundation

inistry of Science and Technology

Meta

intel

THE ISRAELI SMART TRANSPORTATION RESEARCH CENTER

10 Year ANPL Anniversary Social Event, November 2024

TECHNION Israel Institute of Technology

ANPL Autonomous Navigation and Perception Lab

ISRAEL SCIENCE FOUNDATION

NSF

BSF United States – Israel Binational Science Foundation

inistry of Science and Technology

Meta

intel

THE ISRAELI SMART TRANSPORTATION RESEARCH CENTER