AI WEEK
December 9th-11th, 2025
Tel Aviv University

Simplified Online Planning Under Uncertainty with Performance Guarantees

Vadim Indelman

Yuval Ne'eman Workshop for Science, Technology and Security Tel Aviv University

ICRC Blavatnik Interdisciplinary Cyber Research Center

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

TAD The Center for AI & Data Science Tel Aviv University

IDSAI Israel Data Science and AI Initiative

In Cooperation with

ISRAEL State of Israel Ministry of Foreign Affairs

רשות החדשנות Israel Innovation Authority

INCD Israel National Cyber Directorate

TEL AVIV YAFO תל אביב יפו TEL AVIV TECH

Advanced Autonomy

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






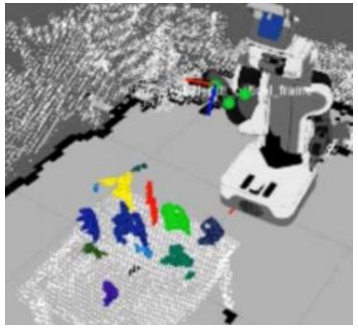

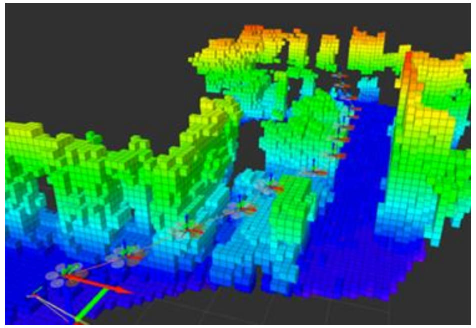


Advanced Autonomy

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**

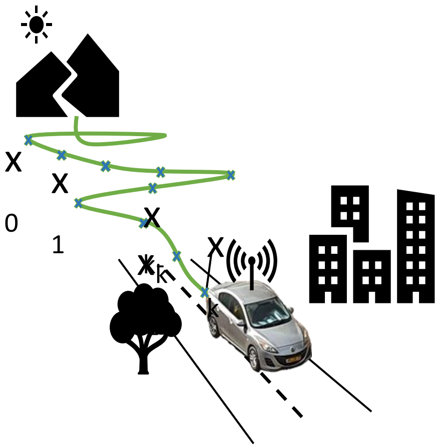

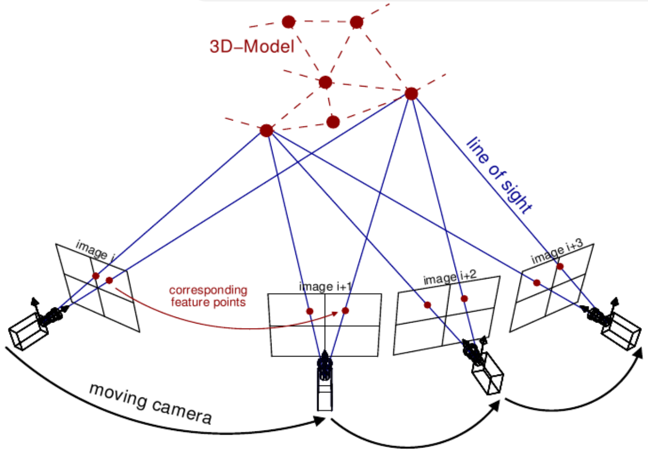

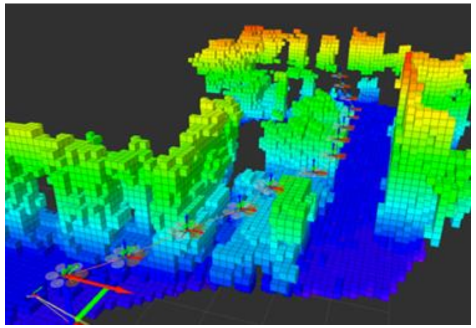

Challenge

Probabilistic Inference

Maintain a distribution over the state given data

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

Decision-making under uncertainty

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

Computationally intractable

More so, in high dimensional settings

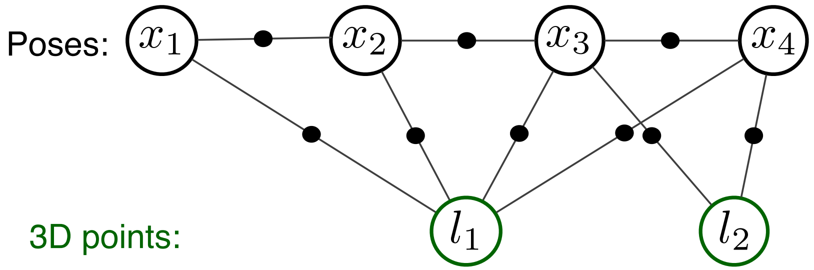

Example – grid world

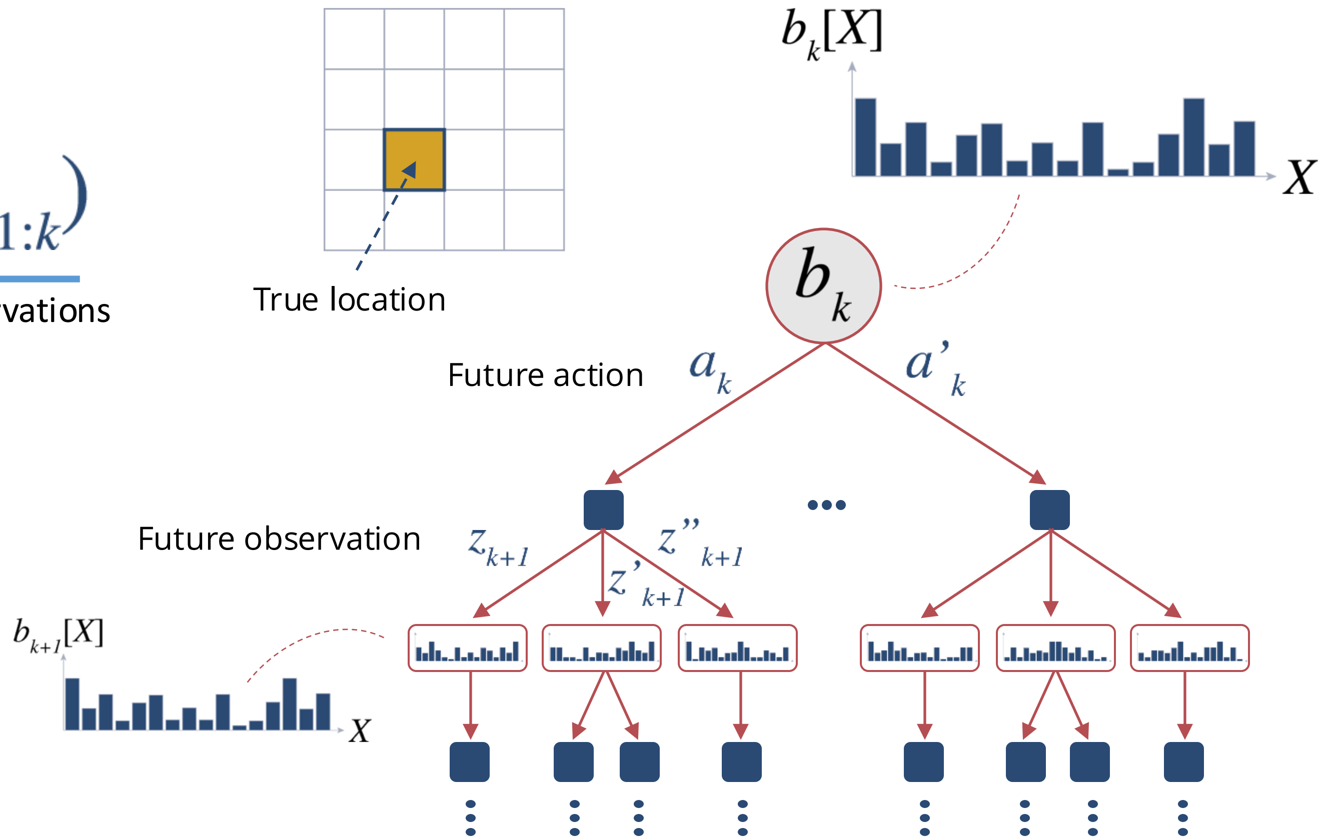

ct autonomously online and efficiently tasks in a safe and reliable fashion??

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

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

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 Risk-Averse POMDP 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

$$V^\pi(b_0) \triangleq \mathbb{E}[\sum_t \gamma^t r_t(b_t, a_t) \mid a_t = \pi_t(b_t)]$$

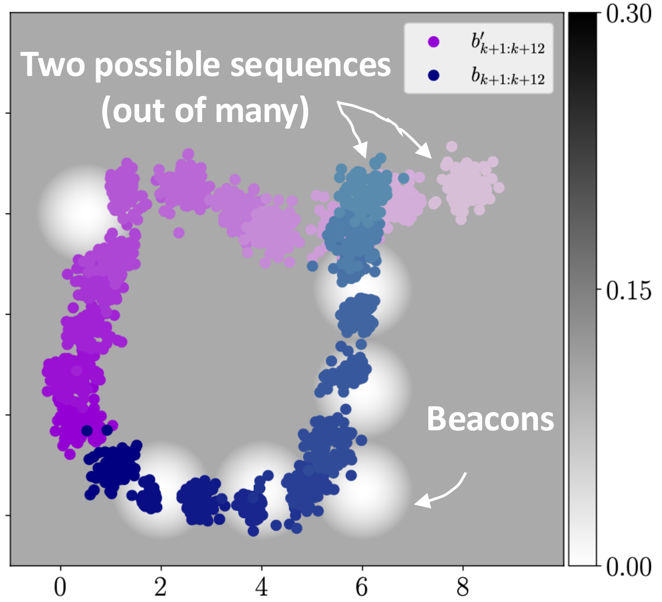

Simplification:

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

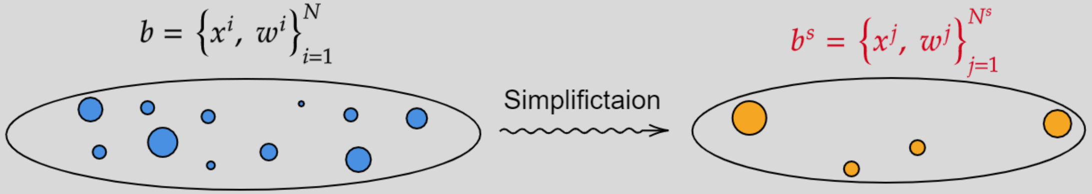

$$lb(b, b^s, a) \leq r(b, a) \leq ub(b, b^s, a)$$

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:

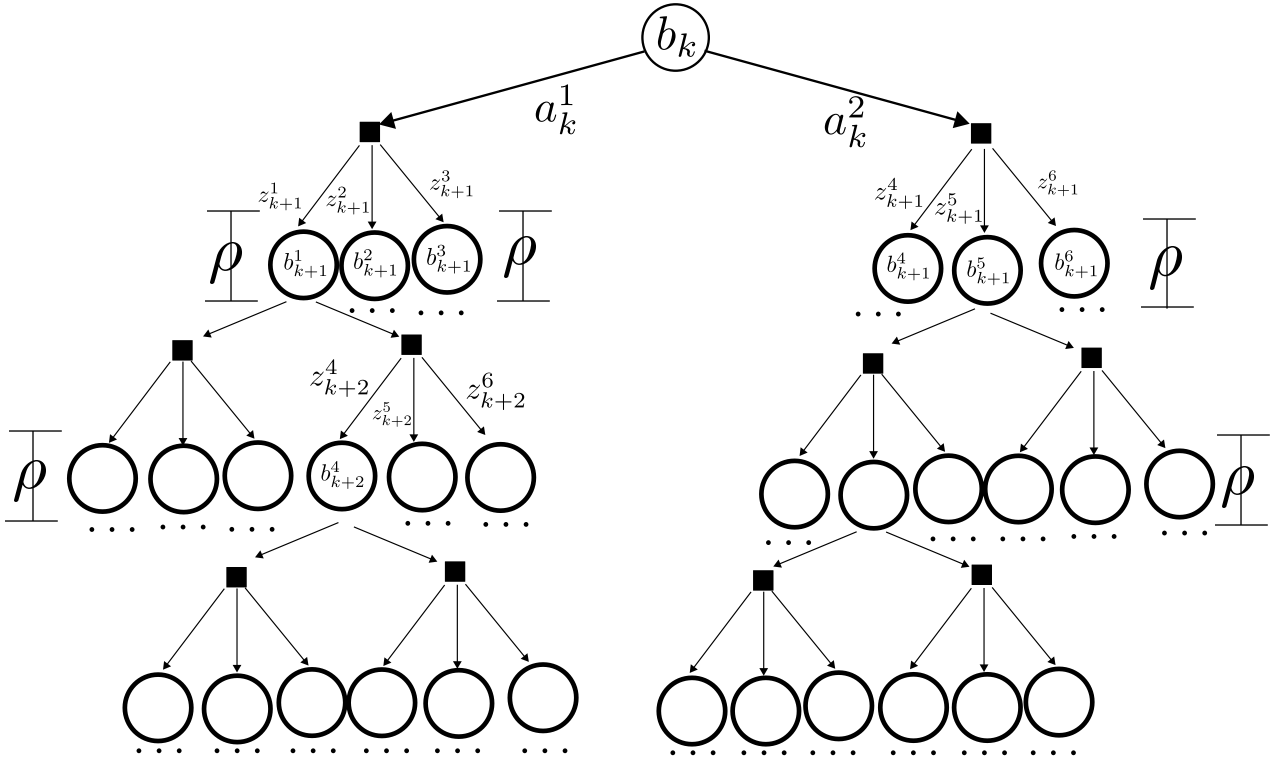

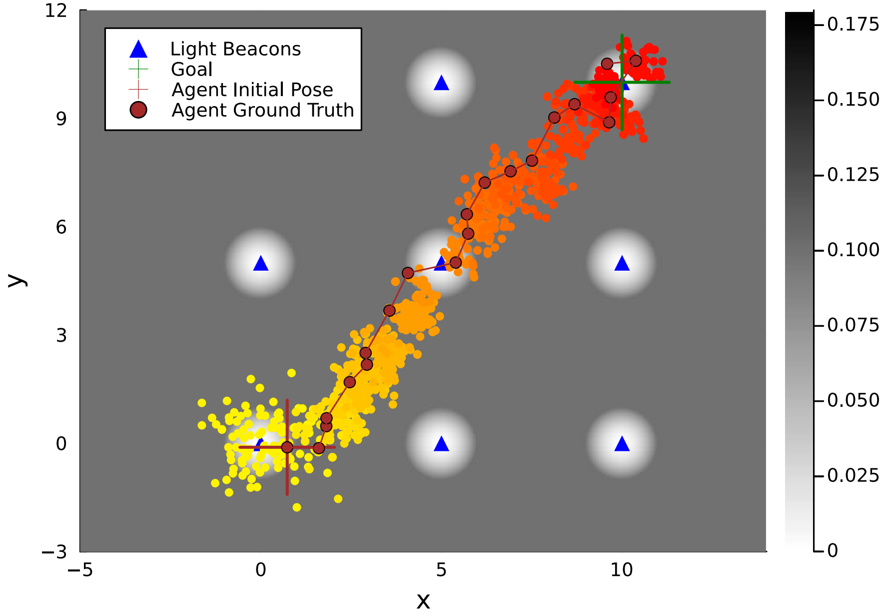

Typical speedup of 20% - 50%,
Same performance!

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:

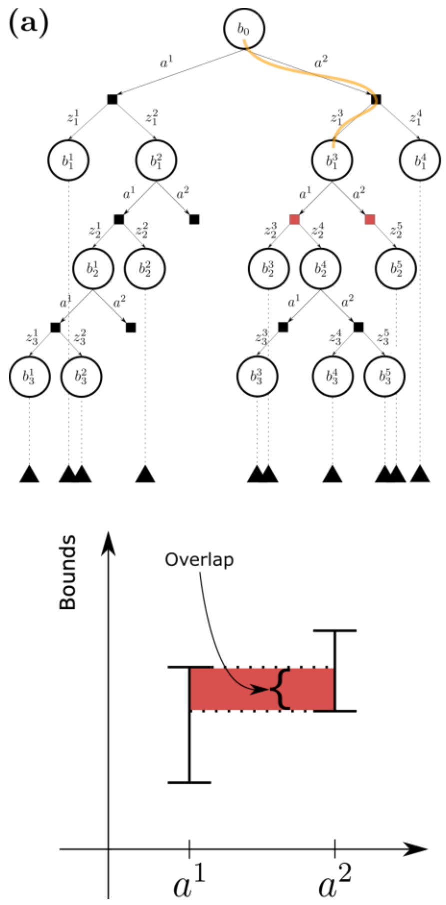

O. Sztyglic and V. Indelman, “Speeding up POMDP Planning via Simplification”, IROS’22.
A. Zhitnikov, O. Sztyglic, and V. Indelman, “No Compromise in Solution Quality: Speeding Up Belief-dependent Continuous POMDPs via Adaptive Multilevel Simplification”, IJRR’24.

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 Risk-Averse POMDP 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

…

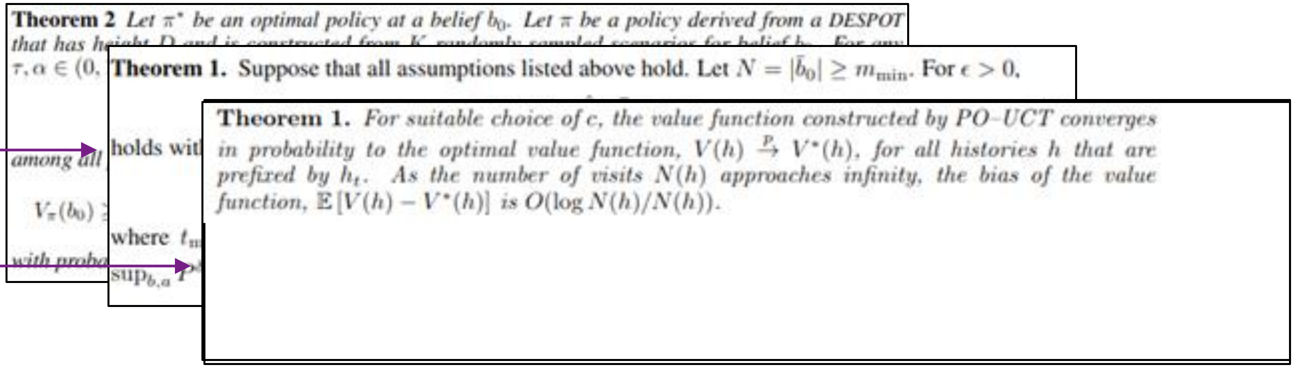

Can we get deterministic guarantees?

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 2025.

Online POMDP Planning with Anytime Deterministic Guarantees

Concept:

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

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

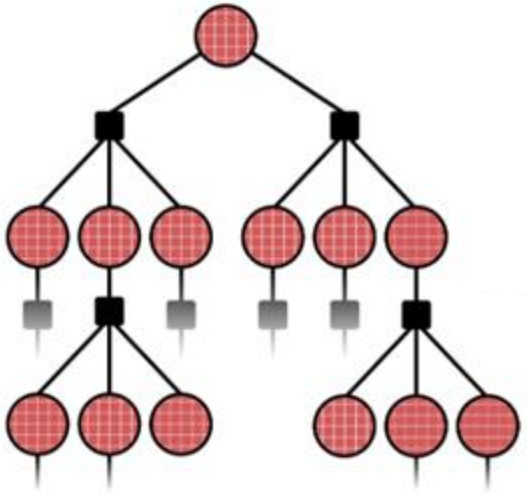

Derive a mathematical relationship between the solution of the simplified, and the theoretical POMDP.

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

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 2025.

Online POMDP Planning with Anytime Deterministic Guarantees

- Deterministic guarantees (assuming discrete spaces)

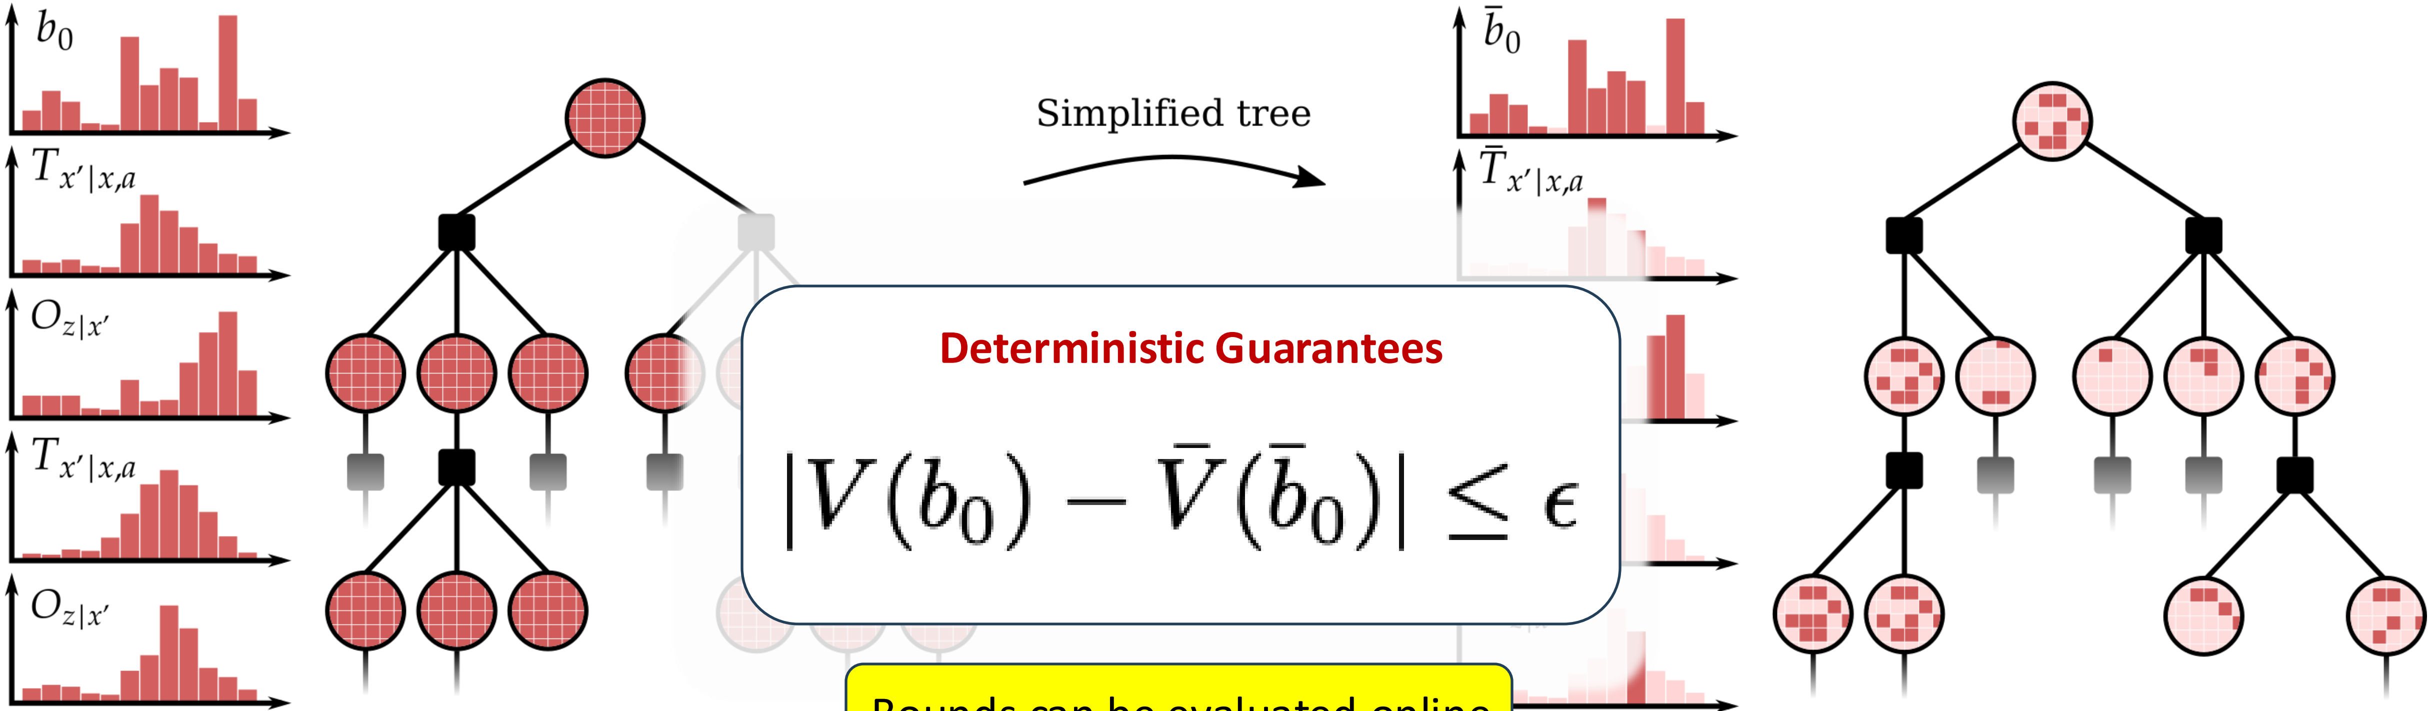

Deterministic Guarantees

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

Bounds can be evaluated online

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 2025.

Online POMDP Planning with Anytime Deterministic Guarantees

Importantly, the bounds can be calculated during planning.

Deterministic Guarantees

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

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

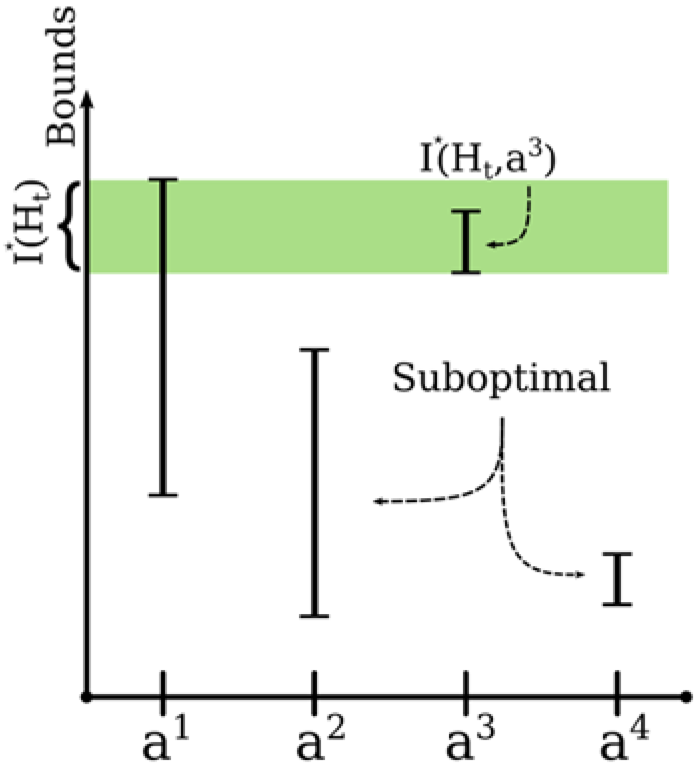

M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Guarantees," NeurIPS'23.
M. Barenboim and V. Indelman, "Online POMDP Planning with Anytime Deterministic Optimality Guarantees," Artificial Intelligence 2025.

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:

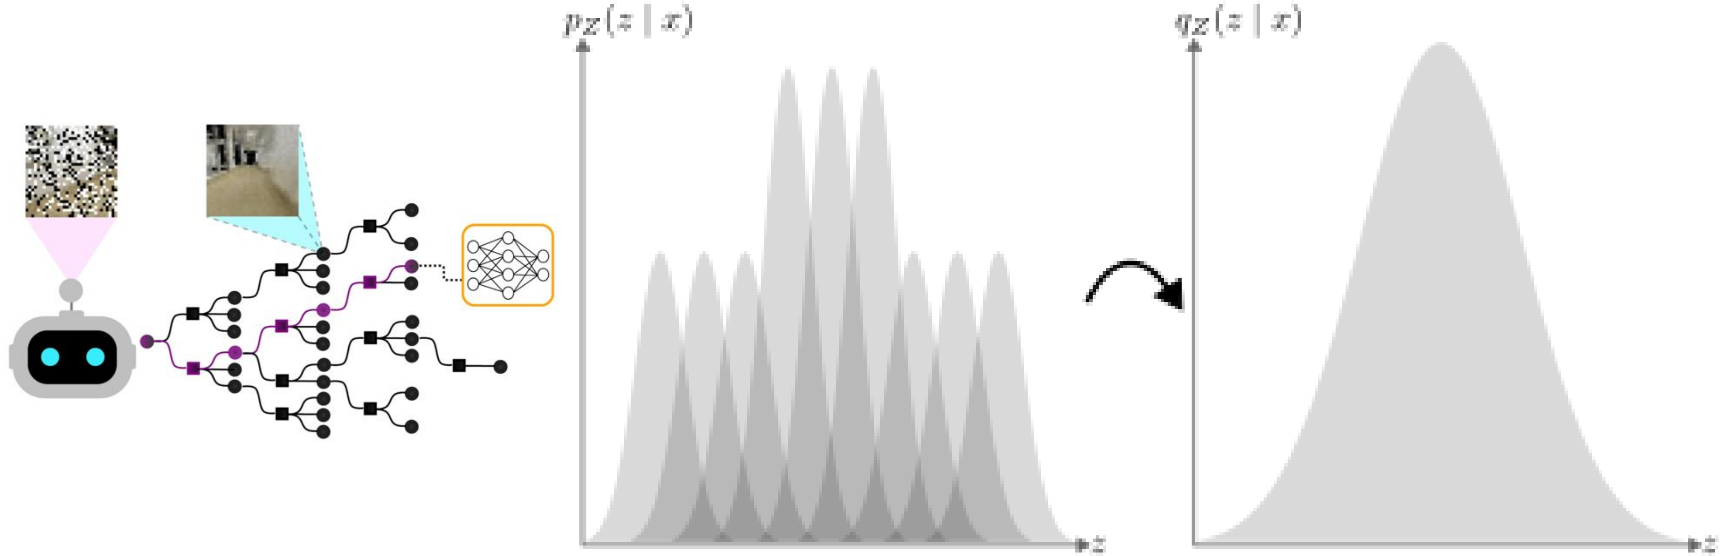

Simplified models

$p_\theta(z \mid x)$

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.

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

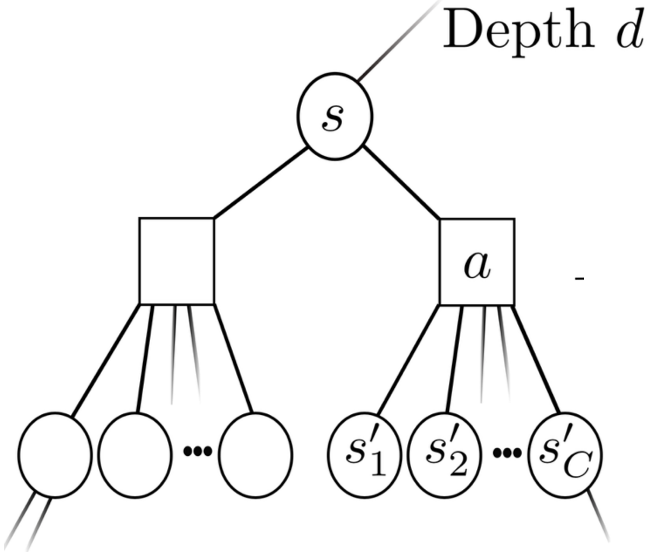

T. Shazman, I. Lev-Yehudi, R. Benchetrit, and V. Indelman, "Online Robust Planning under Model Uncertainty: A Sample-Based Approach," *AAAI'26*.

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:

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

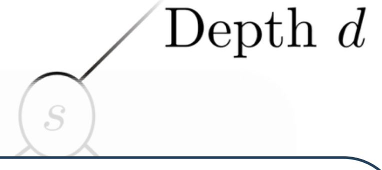

Robust Sparse Sampling (RSS) Algorithm:

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

Prob. 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*.

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 Risk-Averse POMDP 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!

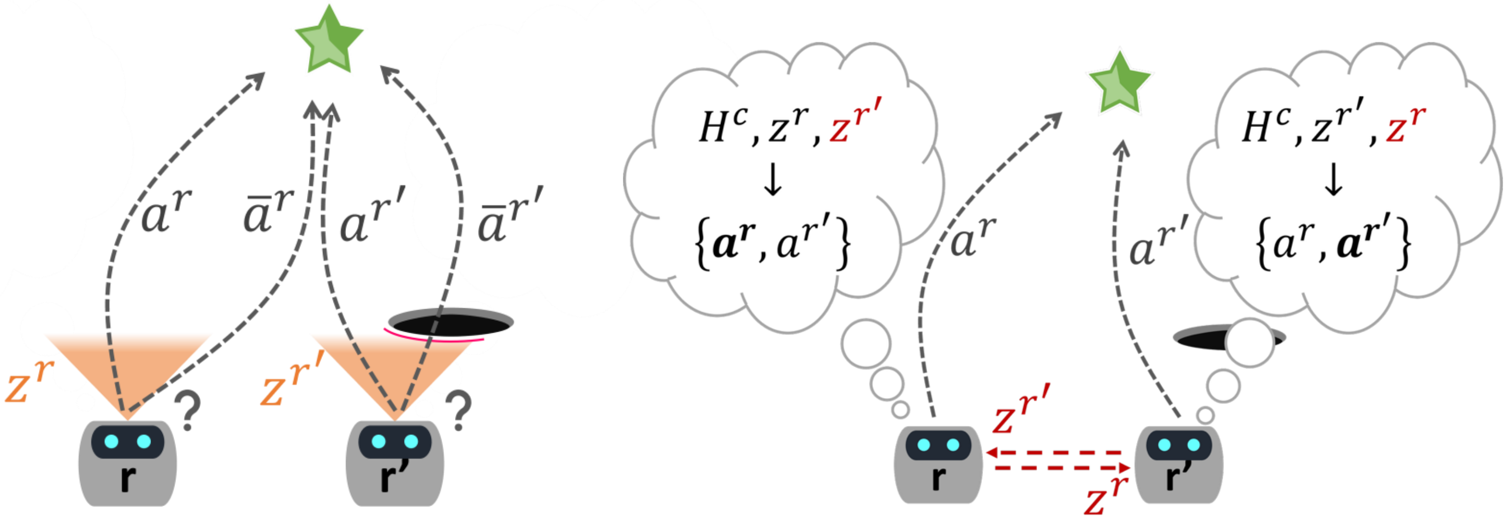

Multi-Robot Cooperative BSP with Inconsistent Beliefs

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

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

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

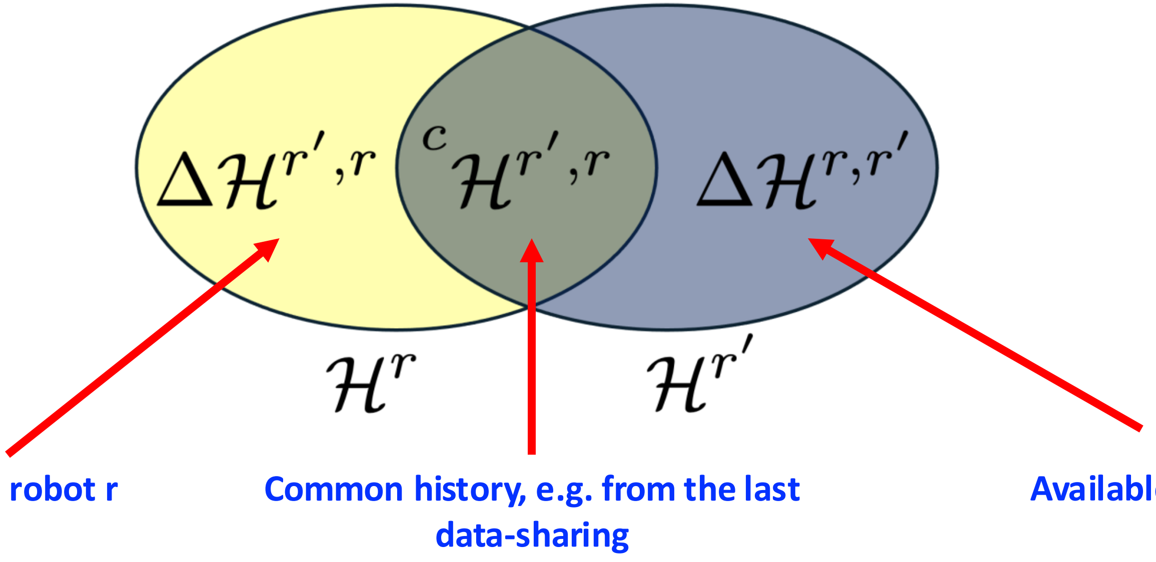

T. Kundu, M. Rafaeli, and V. Indelman, "Multi-Robot Communication-Aware Cooperative Belief Space Planning with Inconsistent Beliefs: An Action-Consistent Approach," IROS'24.

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**

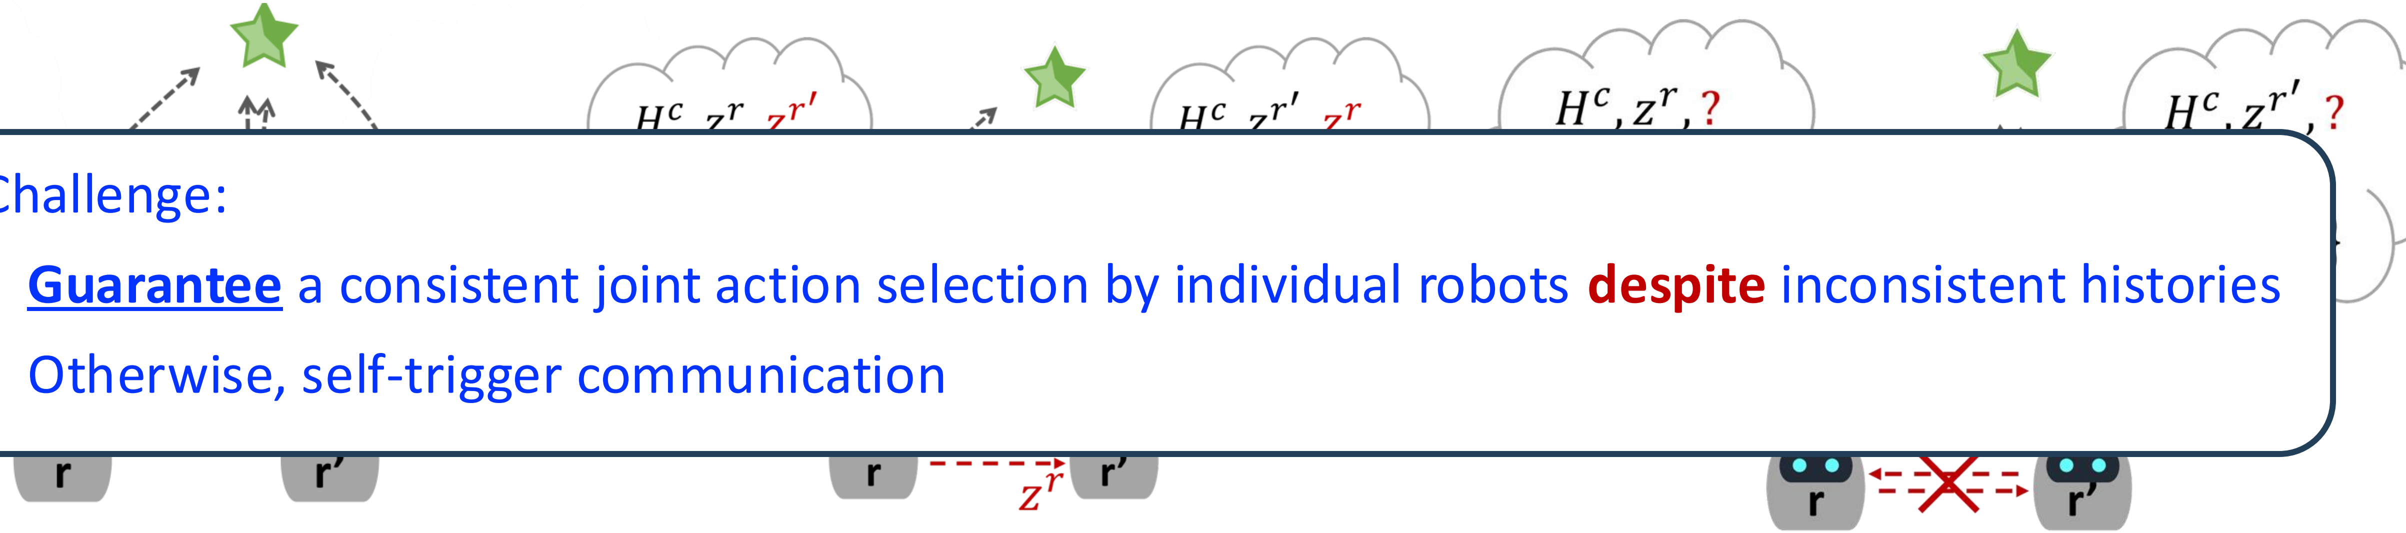

Challenge:

- **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.

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 Risk-Averse POMDP Planning
- Simplification in a multi-agent setting

See additional research directions on our website!

Feel free to reach out to explore **research opportunities**!