Multi-robot Communication-Aware Cooperative Belief Space Planning with Inconsistent Beliefs: An Action-Consistent Approach

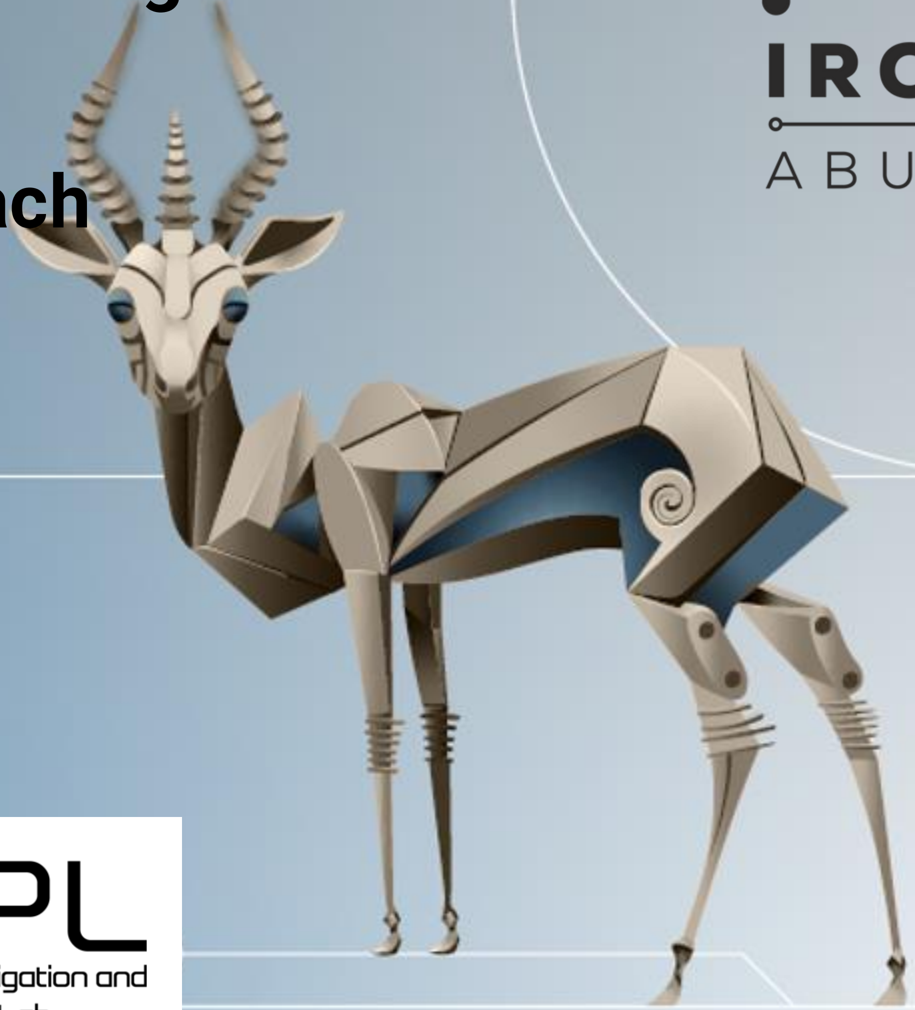

IROS '24
ABU DHABI

Tanmoy Kundu,
Moshe Rafaeli,
Vadim Indelman.

TECHNION Israel Institute of Technology

ANPL Autonomous Navigation and Perception Lab

Multi-Robot Belief Space Planning (MRBSP)

- Handling uncertainty in partially observable environments is fundamental in multi-robot decision making

- Approach: Belief Space Planning (BSP)
  - Computing a globally optimal solution in BSP is computationally intractable

- Previous research
  - Decentralized POMDP, BSP in cooperative and non-cooperative setting

Multi-Robot Belief Space Planning (MRBSP)

- Consider a group of N robots
- Cooperative setting, i.e. same task (reward function) for all robots
- Decentralized POMDP tuple from the perspective of robot r: $\langle \mathcal{X}, \mathcal{Z}, \mathcal{A}, T, O, \rho, b_k \rangle$

Joint state, observation, and action spaces

Joint transition and observation models

Belief-dependent reward function of robot r

Belief of robot r at planning time instant k

- Objective function

$$J^r(b_k, a_{k+}) = \mathbb{E}_{z_{k+1:k+L}}\left[\sum_{l=0}^{L-1} \rho(b_{k+l}, a_{k+l}) + \rho(b_{k+L})\right]$$

Multi-Robot Belief Space Planning (MRBSP)

- **A common assumption**: Beliefs of different robots are consistent at planning time
  - Data of each robot is available to all other robots

    ➢ **Requires prohibitively high number of communication capabilities**

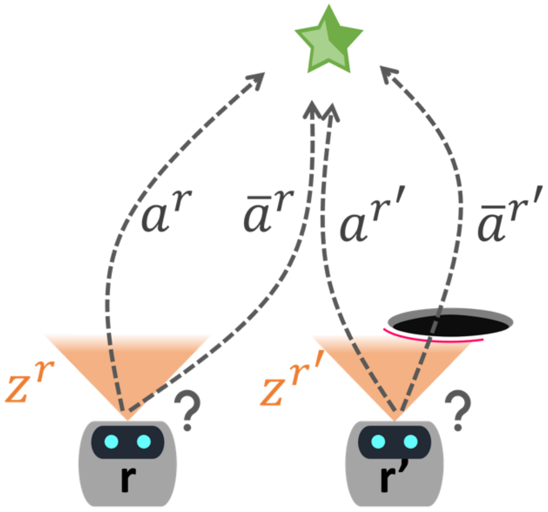

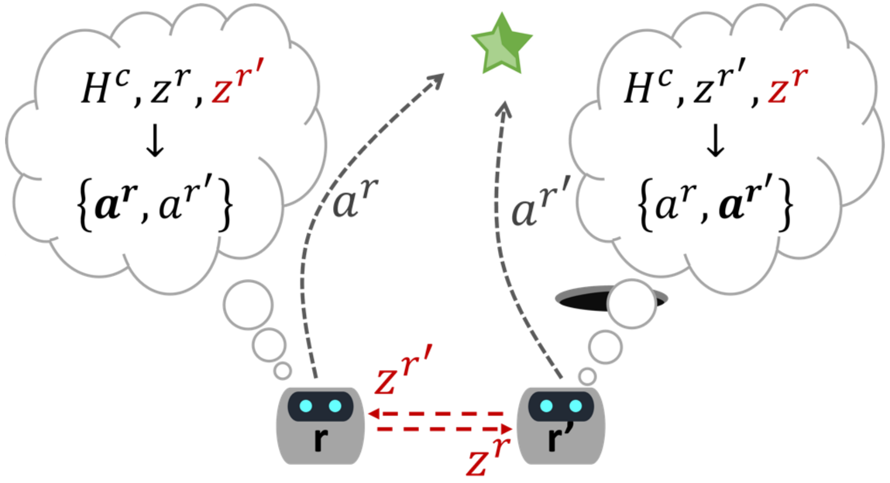

Our work relaxes previous assumption

- In reality
  - Frequent communications among the robots may not be possible
    - Sparse communications → inconsistent beliefs

➢ **Our work: Multi-robot coordination with inconsistent beliefs of the robots**

- Multi-robot cooperative BSP with inconsistent beliefs

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

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

- Decentralized POMDP tuple from the perspective of robot r:

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

- Objective function:

$$J(b_k^r, a_{k+}) = \underset{z_{k+1:k+L}}{\mathbb{E}} [\sum_{l=0}^{L-1} \rho(b_{k+l}^r, a_{k+l}) + \rho(b_{k+L}^r)]$$

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**

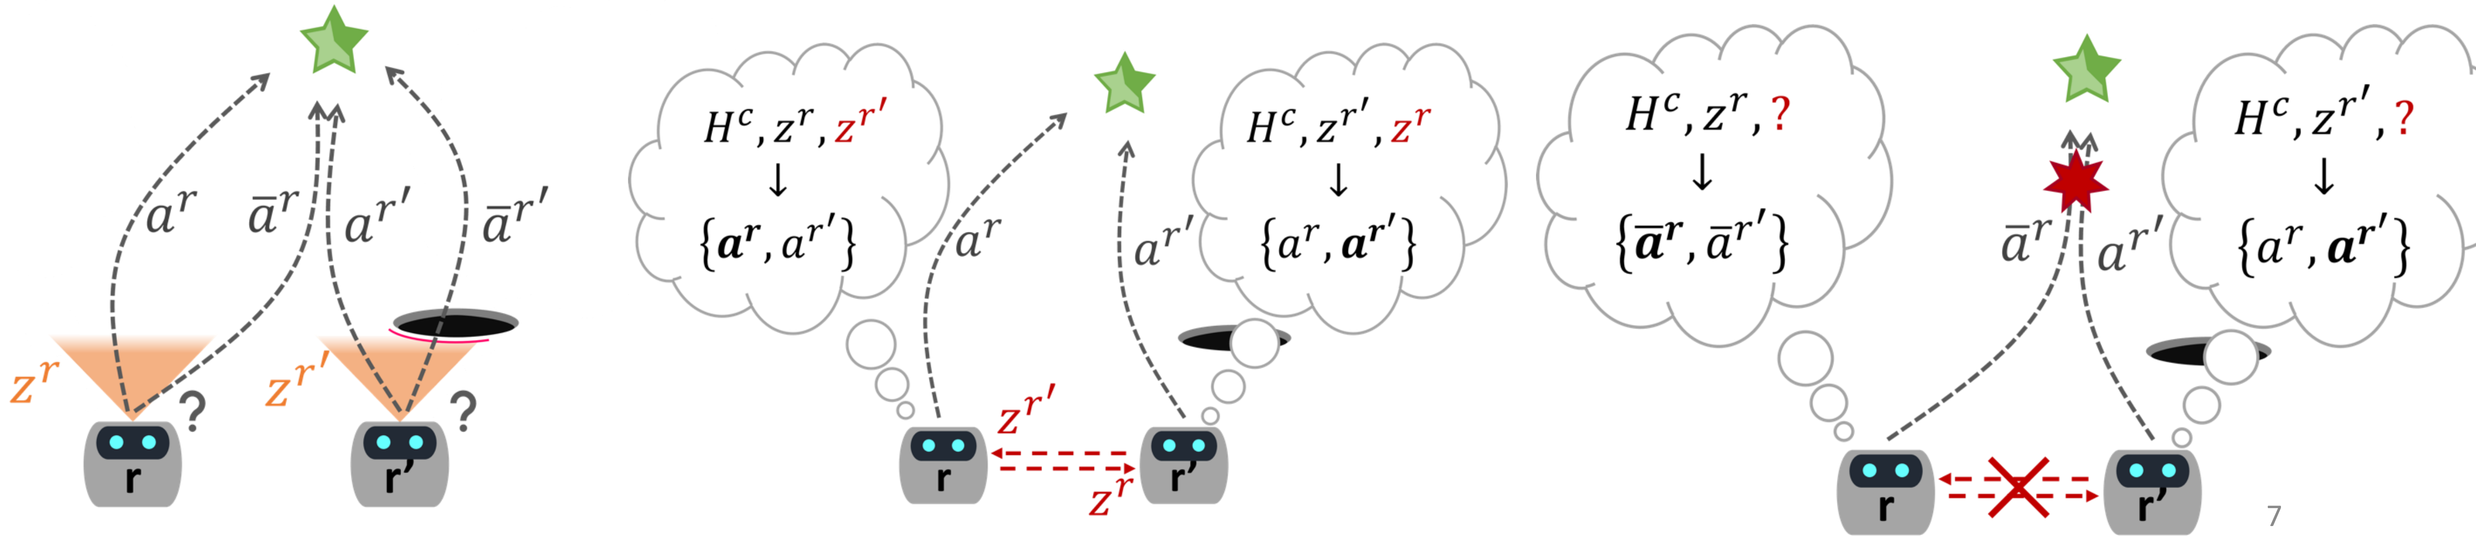

Multi-Robot Cooperative BSP with Inconsistent Beliefs

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 the robots having inconsistent beliefs; otherwise, self-trigger communication

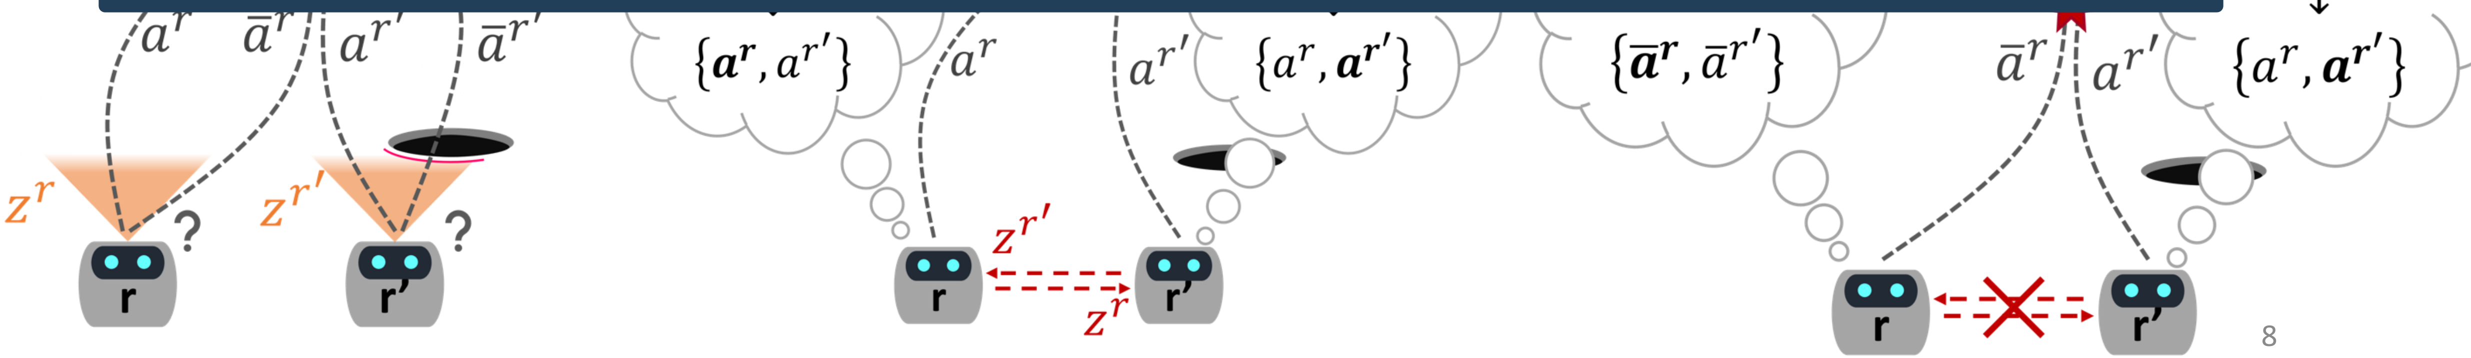

Action Consistency

[Indelman RA-L'16][Elimelech and Indelman, IJRR'22] [Kitanov and Indelman, IJRR'24]

- 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**

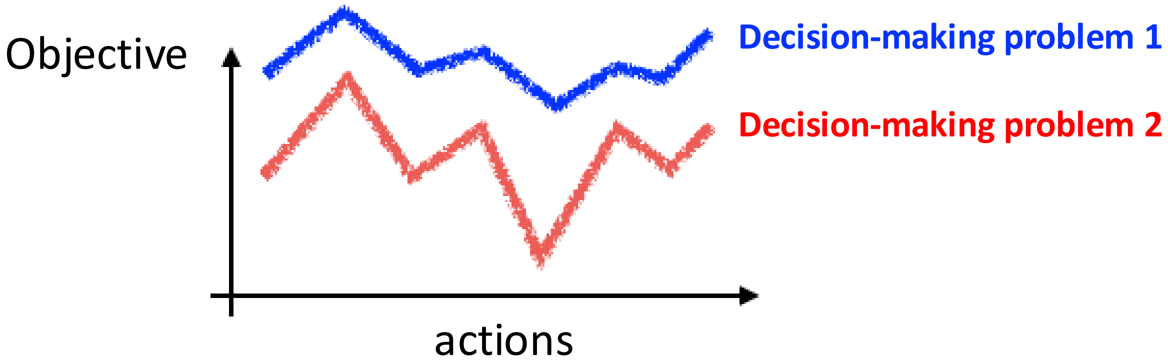

- **Key idea**: to guarantee consistent multi-robot decision-making, each robot
  - reasons about its own and other robots' action preferences while accounting for the **missing information** between the robots
  - checks if for all these realizations, we get the same best joint action

Inconsistent and Common Histories

$$b_k^r = \mathbb{P}(x_k \mid \mathcal{H}_k^r) \quad \Longleftrightarrow \quad b_k^r = \mathbb{P}(x_k \mid {}^c\mathcal{H}_k^{r,r'}, \Delta\mathcal{H}_k^{r',r})$$

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

$$\mathcal{H}_k^r \neq \mathcal{H}_k^{r'}$$

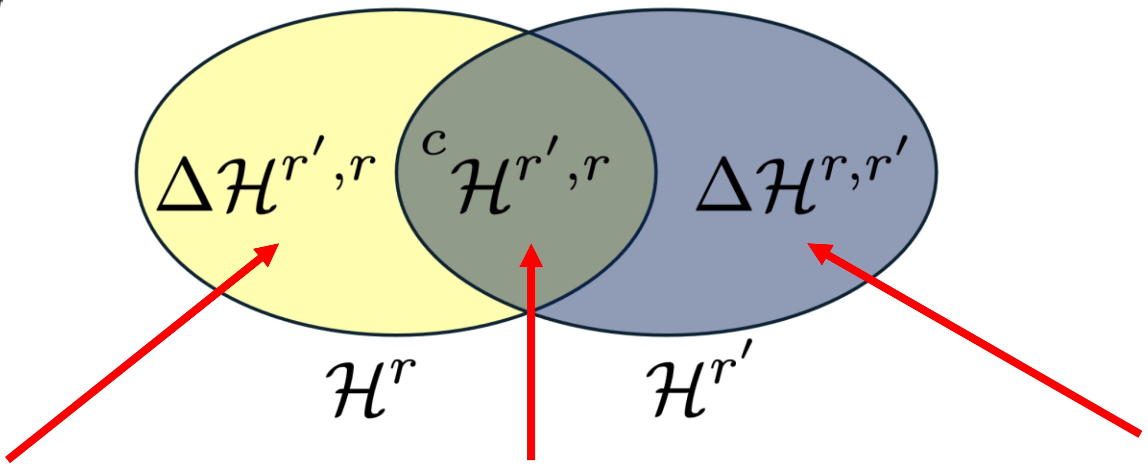

Corresponding observation **spaces**

$$\Delta\mathcal{H}_k^{r',r} \in \Delta\mathcal{Z}_k^{r',r}$$

$$\Delta\mathcal{H}_k^{r,r'} \in \Delta\mathcal{Z}_k^{r,r'}$$

TECHNION Israel Institute of Technology
ANPL Autonomous Navigation and Perception Lab
IROS '24 ABU DHABI

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'

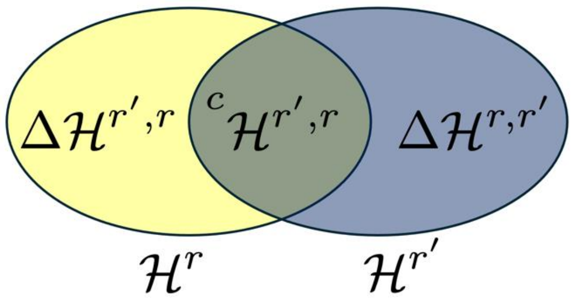

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

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

    1. Its local information

select $\bar{a} \in \mathcal{A}$ s.t. $J(b_k^r, \bar{a}) > J(b_k^r, \bar{a}') \;\; \forall \bar{a}' \in \mathcal{A}$

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

Toy example for $|\mathcal{A}| = |\mathcal{Z}| = 2$ **:**

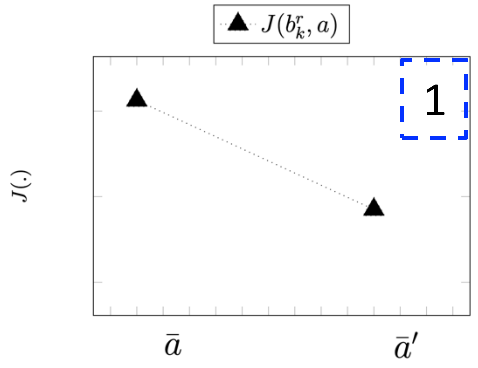

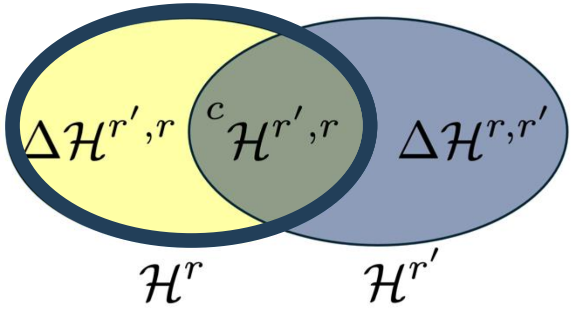

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

- 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’

     For each possible observation of r’, $\tilde{z}^{r'} \in \Delta\mathcal{Z}_k^{r,r'}$, robot r

     constructs a plausible belief of robot r’: $b_k^{r'|r}(\tilde{z}^{r'}) \triangleq \mathbb{P}(x_k \mid {}^{c}\mathcal{H}_k^{r,r'}, \tilde{z}^{r'})$

     evaluates $J(b_k^{r'|r}(\tilde{z}^{r'}), a) \quad \forall a \in \mathcal{A}$

     Checks if $\bar{a}$ is selected

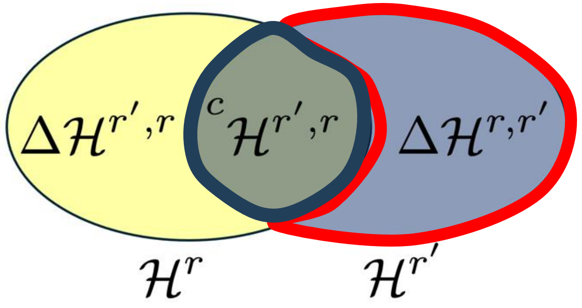

Toy example for $|\mathcal{A}| = |\mathcal{Z}| = 2$ **:**

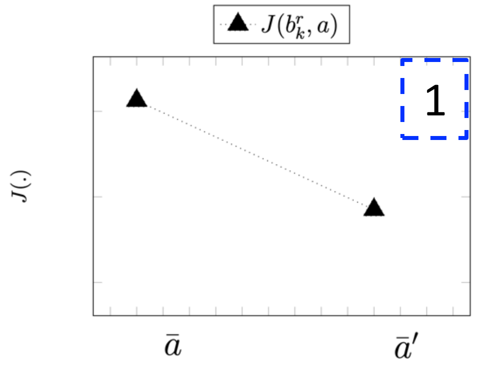

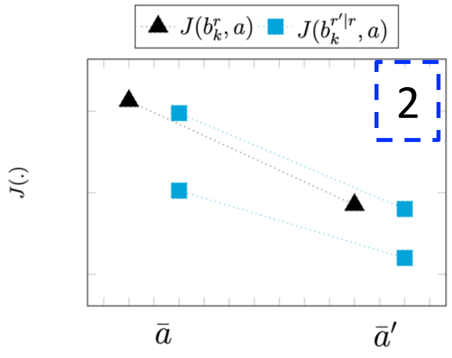

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

- 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'

For each possible observation of itself, $\tilde{z}^r \in \Delta \mathcal{Z}_k^{r',r}$, robot r

constructs a plausible belief of itself perceived by robot r': $b_k^{r|r'|r}(\tilde{z}^r) \triangleq \mathbb{P}(x_k \mid {}^c\mathcal{H}_k^{r',r}, \tilde{z}^r)$

evaluates $J(b_k^{r|r'|r}(\tilde{z}^r), a) \quad \forall a \in \mathcal{A}$

Checks if $\bar{a}$ is selected

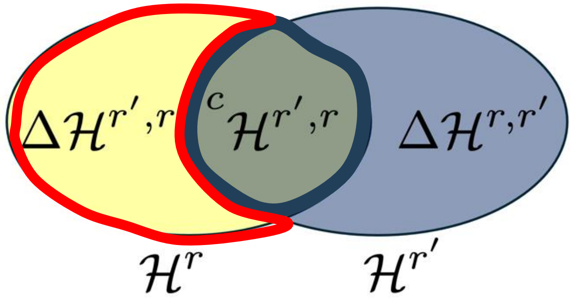

Toy example for $|\mathcal{A}| = |\mathcal{Z}| = 2$ **:**

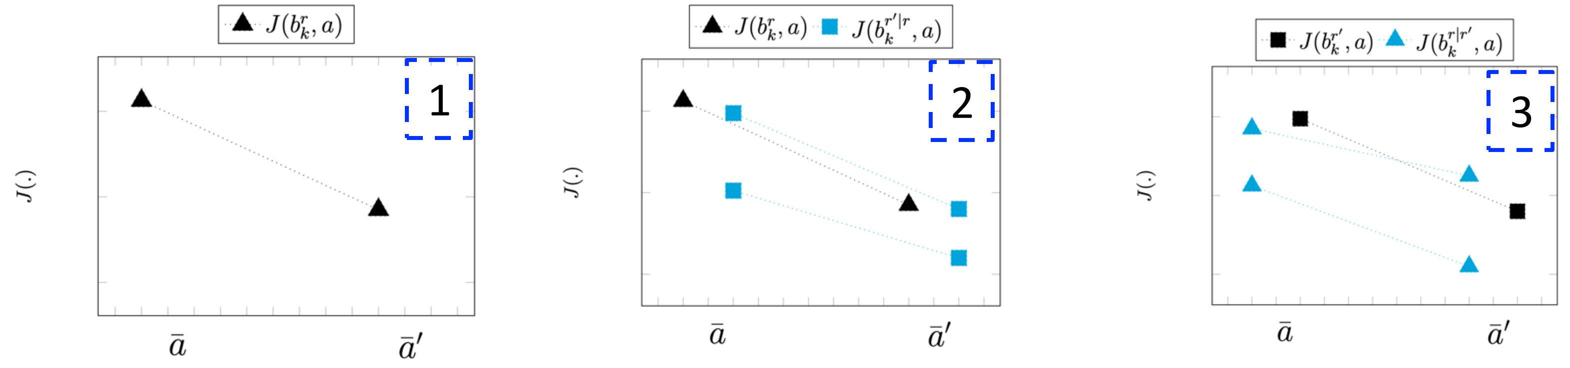

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

- 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’

     For each possible observation of r’, $\tilde{z}^r \in \Delta\mathcal{Z}_k^{r',r}$, robot r

     constructs a plausible belief of robot r’: $b_k^{r|r'|r}(\tilde{z}^r) \triangleq \mathbb{P}(x_k \mid {}^c\mathcal{H}_k^{r',r}, \tilde{z}^r)$

     evaluates $J(b_k^{r|r'|r}(\tilde{z}^r), a) \quad \forall a \in \mathcal{A}$

     Checks if $\bar{a}$ is selected

Toy example for $|\mathcal{A}| = |\mathcal{Z}| = 2$ **:**

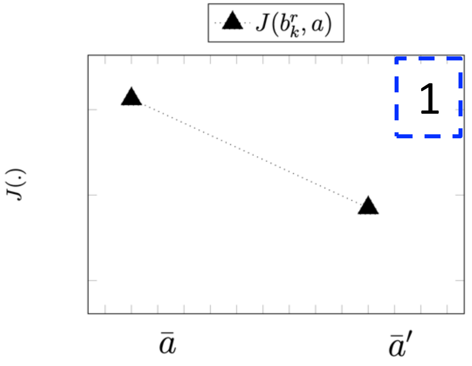

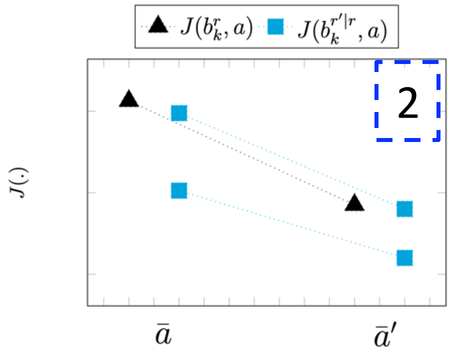

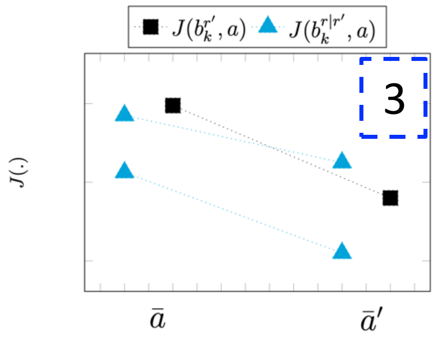

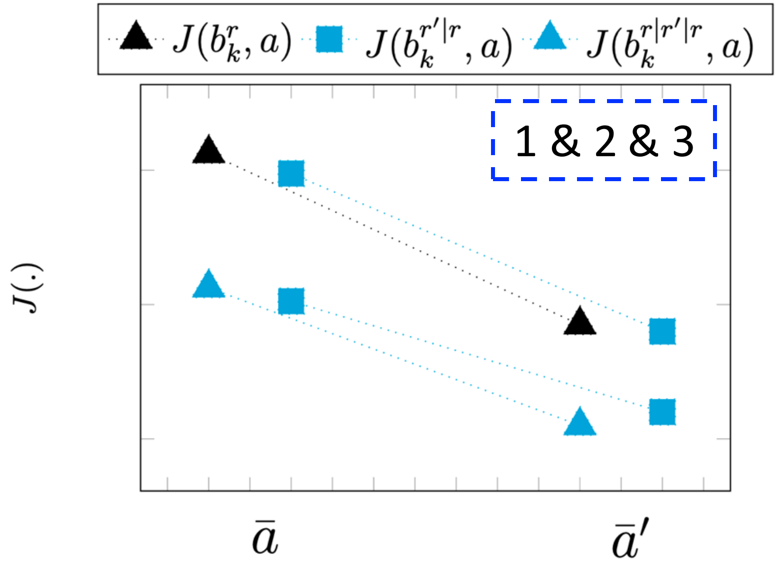

Decentralized Verification of Multi-Robot Action Consistency (MR-AC)

- 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’

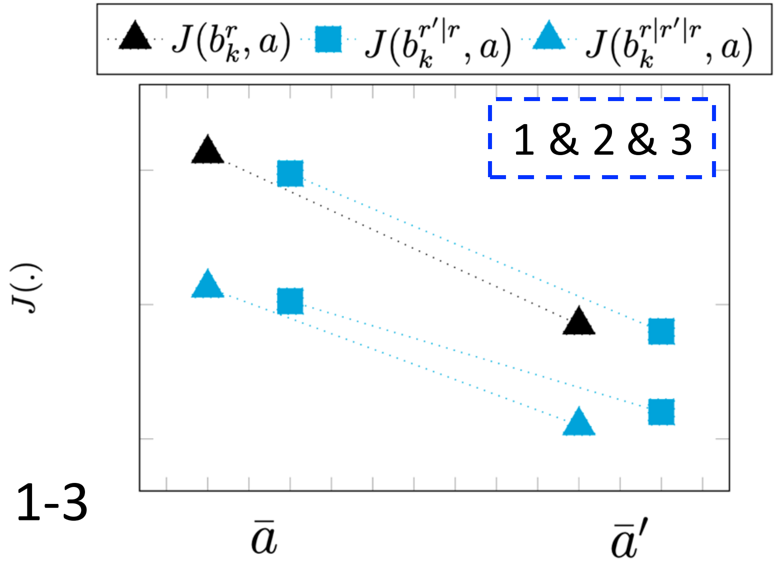

- Same best action in all cases?
  - **Yes**: MR-AC is **guaranteed** to be satisfied

    i.e. robots are guaranteed to choose the same joint action
  - **No**: self-trigger communication, share some data, repeat Steps 1-3

Simulation Results

- Search and Rescue operation in a disaster-hit area

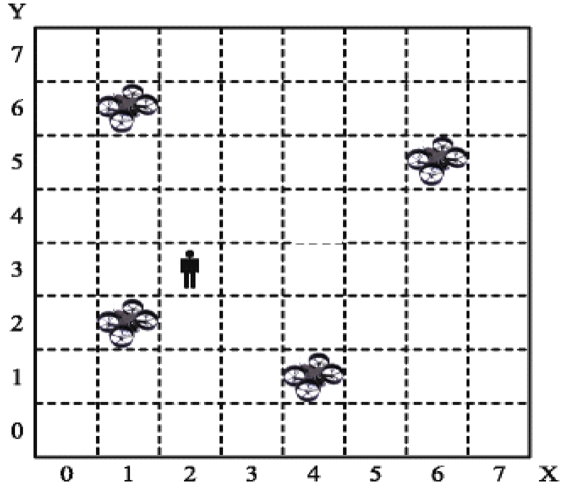

- Probability distribution over the joint state comprising cells $x \triangleq \{x_i\}$
- For simplicity assume cells are stat. independent and robots' poses are known:

$$b_k = p(x \mid \mathcal{H}_k, \xi^r_{0:k}, \xi^{r'}_{0:k}) = \prod_i p(x_i \mid \mathcal{H}_k, \xi^r_{0:k}, \xi^{r'}_{0:k}) \triangleq \prod_i b_k[x_i]$$

- Reward function: entropy (reduce uncertainty) $\rho(b_k) \triangleq -H[x] = \sum_i \sum_{j \in \{0,1\}} b_k[x_i = j] \log b_k[x_i = j]$

Simulation Results

- EnforceAC: our approach
- Baseline I: always communicate all data
- Baseline II: never communicate

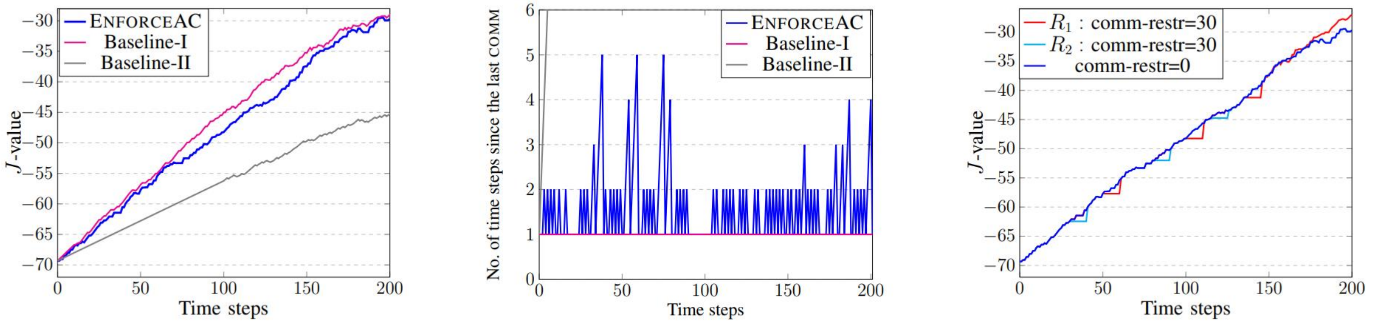

(a) *comm-restr* = 0 (b) *comm-restr* = 0 (c) *comm-restr* = 30

Not-AC (action inconsistency), comms and time for $E = 200$.

Input	Algorithm	Not-AC	COMM	Time
comm-restr = 0	Baseline-II	181	0	1.3s
Motion prim. = 4	Baseline-I	0	400	1.3s
MaxEntropy-Init	EnforceAC	0	238	12.4s
comm-restr = 0	Baseline-II	185	0	1.3s
Motion prim. = 4	Baseline-I	0	400	1.4s
Entropy-Init	EnforceAC	0	268	8.7s
comm-restr = 0	Baseline-II	194	0	3.6s
Motion prim. = 8	Baseline-I	0	400	3.5s
MaxEntropy-Init	EnforceAC	0	248	36.4s
comm-restr = 0	Baseline-II	188	0	3.6s
Motion prim. = 8	Baseline-I	0	400	3.6s
Entropy-Init	EnforceAC	0	278	31.1s
comm-restr = 20	Baseline-II	194	0	3.3s
Motion prim. = 8	Baseline-I	14	360	4.3s
MaxEntropy-Init	EnforceAC	13	224	94.9s
comm-restr = 20	Baseline-II	188	0	3.2s
Motion prim. = 8	Baseline-I	14	360	3.6s
Entropy-Init	EnforceAC	10	251	31.2s
comm-restr = 30	Baseline-II	188	0	3.4s
Motion prim. = 8	Baseline-I	22	340	4.0s
MaxEntropy-Init	EnforceAC	20	238	46.9s

Conclusions

- Formulation of a new problem: MRBSP with inconsistent beliefs
- A novel approach to address cooperative MR-BSP with inconsistent beliefs
- A self-triggering mechanism of communication between robots
- Our approach reduces number communications considerably compared to full-communication approaches.

Thank you!

IROS '24
ABU DHABI