I Would If I Could: Reasoning about Dynamics of Actions in Multi-Agent Systems

Rustam Galimullin [1] Hermine Grosinger [2] **Munyque Mittelmann** [3]

[1]University of Bergen (Norway)

[2]Örebro University (Sweden)

[3]CNRS, LIPN, Sorbonne Paris North University (France)

rustam.galimullin@uib.no, *hermine.grosinger@oru.se*, *mittelmann@lipn.univ-paris13.fr*

Context

- Multi-Agent Systems (MAS) have autonomous, strategic agents
- Real agents should operate in dynamic environments
- Actions set may not be fixed: new actions may need to be learned, old ones may be prohibited

Problem

Consequences of changing the available actions in the strategic abilities and system's properties

Context

- The robot can clean windows
- We also need help taking care of the plants...

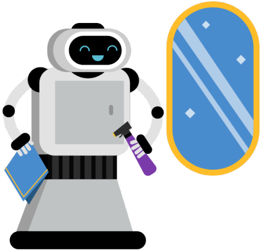

Context

- If we *give* it the ability to water the plants, will we be able to keep the plants healthy?
- Does the robot *know* that to keep the plants healthy he should learn how to water them?

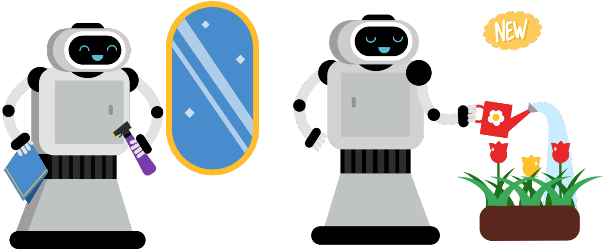

Context

- The cat likes to sleep by the window
- **Cat**astrophic failure: the robot may disturb the cat!

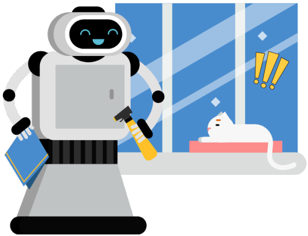

Context

- Whenever the cat is by the window, the action should not be available
- The problem is solved by *revoking* it conditionally, but allowing it otherwise

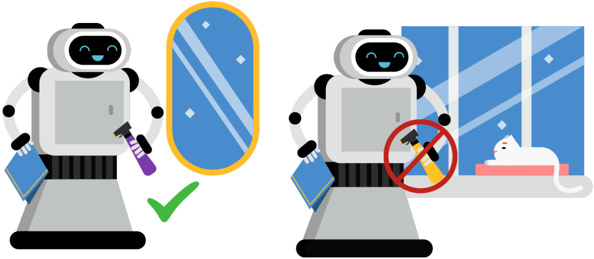

Alternating-time Temporal Logic (ATL)

- ATL and its epistemic extension ATEL: what agents can achieve (and their knowledge)
- **Static models:** cannot handle dynamic updates to agents' available actions

Alternating-time Temporal Logic (ATL)

- ATL and its epistemic extension ATEL: what agents can achieve (and their knowledge)
- **Static models:** cannot handle dynamic updates to agents' available actions
- Not ideal to reason about missing actions needed to achieve a goal & actions that should be forbidden

Our goal

Logic-based approach for strategic reasoning with dynamic action updates

Alternating-time Temporal Logic (ATL)

- ATL and its epistemic extension ATEL: what agents can achieve (and their knowledge)
- **Static models:** cannot handle dynamic updates to agents' available actions
- Not ideal to reason about missing actions needed to achieve a goal & actions that should be forbidden

Our goal

Logic-based approach for strategic reasoning with dynamic action updates

Blue Sky ideas

Proactive learning (Grosinger 2025) $\cap$ dynamic strategy logics (Galimullin, Gladyshev, et al. 2026)

Contribution

- Extension of ATL and ATEL with dynamic actions:
    - Update modalities for granting and revoking actions (conditionally)
    - In the imperfect information setting, updates affect the agents' knowledge

Results

- Complexity for the model checking and bounded update existence problems
- Expressivity results and relation with normative reasoning

Concurrent Game Structure (CGS)

- Ag — *agents*
- Ac — *actions*
- AP — *atomic propositions*
- V — *states*
- $\ell : V \to 2^{\mathsf{AP}}$ — *labeling function*
- $d : \mathsf{Ag} \times V \to 2^{\mathsf{Ac}} \setminus \emptyset$ — *availability function*
- $o : V \times \mathsf{Ac}^{\mathsf{Ag}} \to V$ — *transition function*

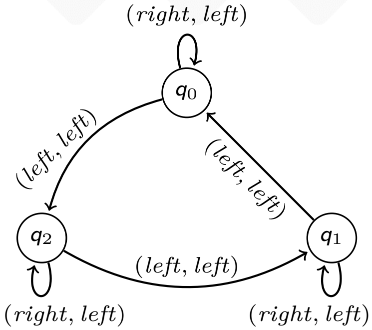

Concurrent Game Structure (CGS)

- Ag — *agents*
- Ac — *actions*
- AP — *atomic propositions*
- V — *states*
- $\ell : V \to 2^{\mathsf{AP}}$ — *labeling function*
- $d : \mathsf{Ag} \times V \to 2^{\mathsf{Ac}} \setminus \emptyset$ — *availability function*
- $o : V \times \mathsf{Ac}^{\mathsf{Ag}} \to V$ — *transition function*

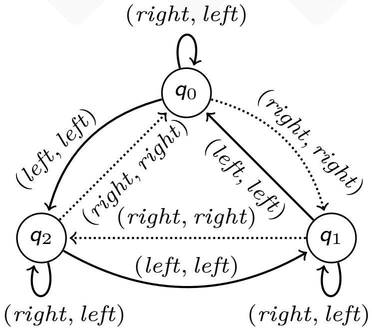

Alternating-time Temporal Logic with Dynamic actions (ATL-D)

Syntax of ATL-D

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\langle A \rangle\rangle \mathbf{X}\varphi \mid \langle\langle A \rangle\rangle \varphi \mathbf{U} \varphi \mid \langle\langle A \rangle\rangle \varphi \mathbf{R} \varphi \mid$$
$$[\pi]^{+}\varphi \mid [\pi]^{-}\varphi \mid has(a, \alpha)$$
$$\pi ::= \varphi : \alpha \to A \mid \pi, \pi$$

where p is a proposition, α is an action, a is an agent, and A is a coalition

Semantics

- $\langle\langle A \rangle\rangle \psi$: the coalition A can force the temporal goal ψ
- $has(a, \alpha)$: in the current state, the action α is available to agent a.
- $[\pi]^{+}\varphi$: after applying the granting actions update π, φ holds
- $[\pi]^{-}\varphi$: after applying the removing actions update π, φ holds

Alternating-time Temporal Logic with Dynamic actions (ATL-D)

Syntax of ATL-D

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\langle A\rangle\rangle \mathbf{X}\varphi \mid \langle\langle A\rangle\rangle \varphi \mathbf{U} \varphi \mid \langle\langle A\rangle\rangle \varphi \mathbf{R} \varphi \mid$$
$$[\pi]^{+}\varphi \mid [\pi]^{-}\varphi \mid has(a, \alpha)$$
$$\pi ::= \varphi : \alpha \rightarrow A \mid \pi, \pi$$

where p is a proposition, α is an action, a is an agent, and A is a coalition

Semantics

- $\varphi : \alpha \rightarrow A$: the action α is granted (or removed) to agents in A in *every state* where φ is true
- Updates of the same type can be applied simultaneously; update operators can be chained

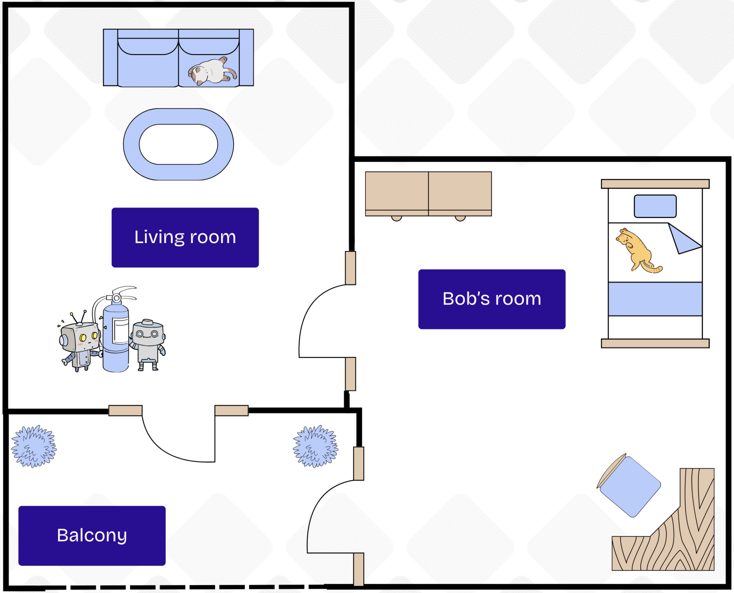

Goals:

- Bring the oxygen tank to Bob's room (quickly)
- Avoid bringing it to the balcony because it's warm outside

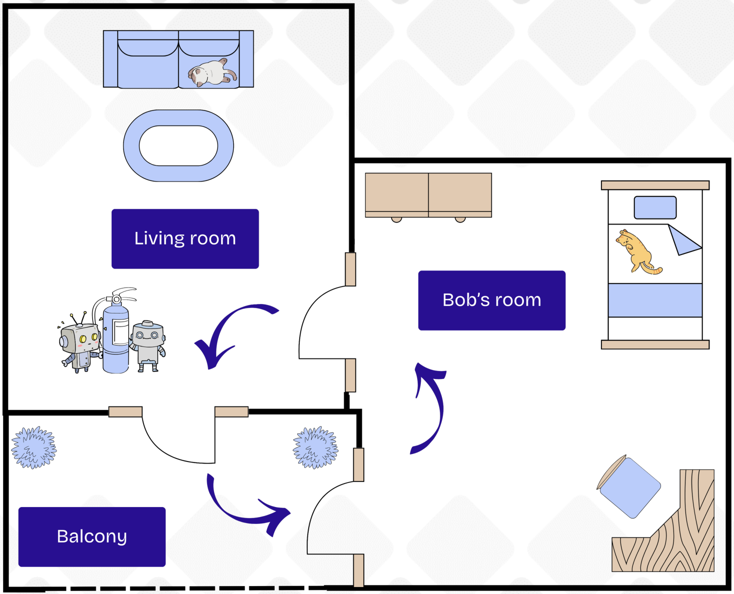

- Robots move when they push in the same direction
- Robot 2 doesn’t know how to turn to the right

Example

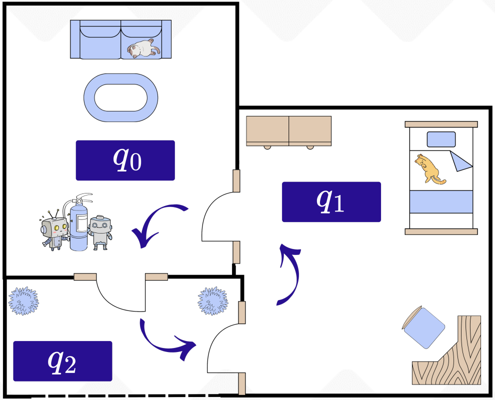

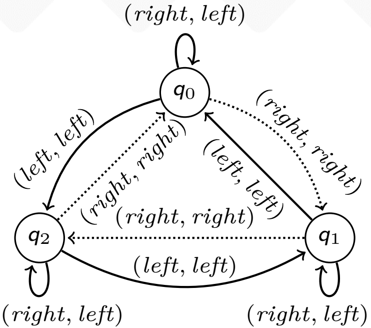

CGS $\mathcal{G}$

Example

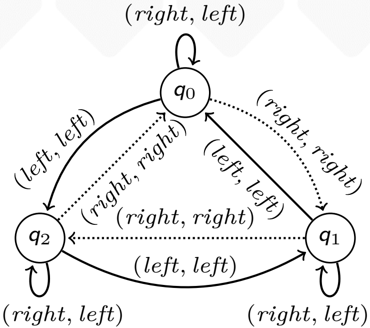

q_0 = *living room* $\quad$ q_1 = *Bob's room*

q_2 = *balcony* $\quad$ $q \in \{q_0, q_1, q_2\}$

$$\mathcal{G}, q_0 \not\models \langle\langle r_1, r_2 \rangle\rangle \mathbf{X}\, atBob$$

Example

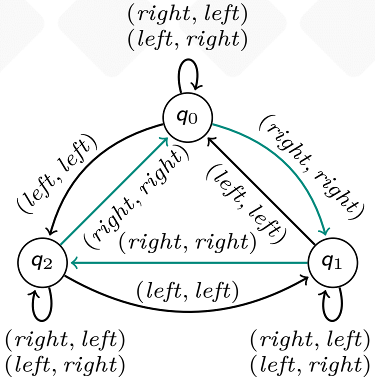

$$\mathcal{G}, q_0 \not\models \langle\langle r_1, r_2 \rangle\rangle \mathbf{X}\, atBob$$

$$\mathcal{G},\ q \models [\top : right \to r_2]^+ \langle\langle r_1, r_2 \rangle\rangle \mathbf{X}\, atBob$$

q_0 = *living room* $\quad$ q_1 = *Bob's room*

q_2 = *balcony* $\quad$ $q \in \{q_0, q_1, q_2\}$

Example

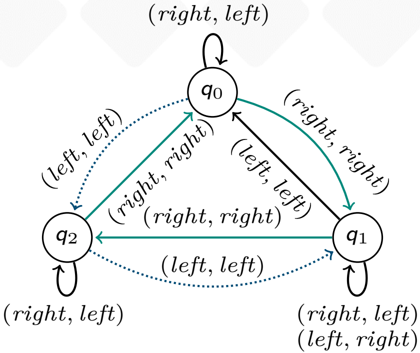

$$\mathcal{G}, q_0 \not\models \langle\!\langle r_1, r_2 \rangle\!\rangle \mathbf{X}\, atBob$$

$$\mathcal{G},\ q \models [\top : right \to r_2]^+ \langle\!\langle r_1, r_2 \rangle\!\rangle \mathbf{X}\, atBob$$

$$\mathcal{G},\ q_0 \models [\top : right \to r_2]^+\, [\neg atBob : left \to r_1]^- \langle\!\langle \emptyset \rangle\!\rangle\, atBob\, \mathbf{R}\, \neg warm$$

q_0 = *living room* $\quad$ q_1 = *Bob's room*

q_2 = *balcony* $\quad$ $q \in \{q_0, q_1, q_2\}$

Expressivity

The fragments without $[\pi]^{+}\varphi$ or $[\pi]^{-}\varphi$ are denoted ATL-D^{-} and ATL-D^{+}, resp.

Expressivity results

$$
\begin{tikzcd}
 & \text{ATL-D}^{+} \arrow[dr] & \\
\text{ATL} \arrow[ur] \arrow[dr] & & \text{ATL-D} \\
 & \text{ATL-D}^{-} \arrow[ur] &
\end{tikzcd}
$$

(An arrow from L_1 to L_2 means $L_1 \prec L_2$)

ATL-Based Normative Reasoning

- Hoek, Roberts, and Wooldridge (2007)'s social laws are pairs of goal formulas and *behavioural constraints* (defining prohibited actions at a given state)
- Galimullin and Kuijer (2024)'s atomic social laws are similar to behavioural constraints, but uses logical conditions to specify the states

ATL-D$^-$ captures both notions

Model Checking

Global model checking problem

Given a CGS $\mathcal{G}$ and a formula $\varphi \in$ ATL-D, compute $\{v \in V \mid \mathcal{G}, v \models \varphi\}$

Theorem

Model checking ATL-D is P-complete.

We need to compute the set of states from which the coalition can achieve the outcome not only for the original model but for the updated models (bounded by the formula size)

Update Existence

Bounded Update Existence Problem

Given a CGS $\mathcal{G}$, a formula $\varphi \in$ ATL-D, and a natural number n, determine whether *there is a sequence* $[\overrightarrow{\pi}]$ of updates $[\pi]^+$ and $[\pi]^-$ with $|\overrightarrow{\pi}| \leqslant n$ such that

$$\mathcal{G}, v \models [\overrightarrow{\pi}]\varphi$$

Theorem

The bounded update existence problem is NP-complete.

Reduction from the 3SAT problem (represent truth-values for atoms as states, and define a formula representing the choice of the truth values to make the 3SAT formula true)

Imperfect Information

- We can reason about granting and removing actions to affect the agents' abilities
- We would also like to reason about their knowledge:
  - Do I know which action I should acquire in order to achieve my goal?

Imperfect Information

- We can reason about granting and removing actions to affect the agents' abilities
- We would also like to reason about their knowledge:
  - Do I know which action I should acquire in order to achieve my goal?
  - We need to consider the epistemic dimension

CGS with Imperfect Information

- Ag — *agents*
- Ac — *actions*
- AP — *atomic propositions*
- V — *states*
- $\ell : V \to 2^{\mathrm{AP}}$ — *labeling function*
- $d : \mathrm{Ag} \times V \to 2^{\mathrm{Ac}} \setminus \emptyset$ — *availability function*
- $o : V \times \mathrm{Ac}^{\mathrm{Ag}} \to V$ — *transition function*
- $\sim_a \subseteq V \times V$ — *observation relation*

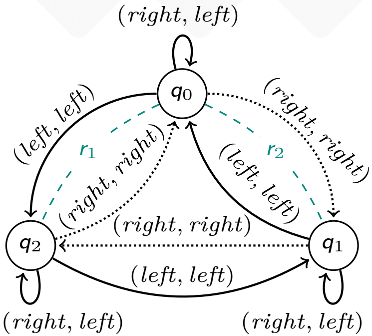

CGS with Imperfect Information

- Ag — *agents*
- Ac — *actions*
- AP — *atomic propositions*
- V — *states*
- $\ell : V \to 2^{\mathrm{AP}}$ — *labeling function*
- $d : \mathrm{Ag} \times V \to 2^{\mathrm{Ac}} \setminus \emptyset$ — *availability function*
- $o : V \times \mathrm{Ac}^{\mathrm{Ag}} \to V$ — *transition function*
- $\sim_a \subseteq V \times V$ — *observation relation*

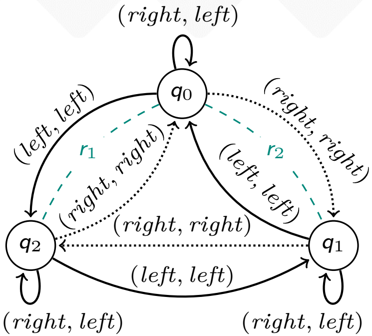

We consider uniform memoryless strategies

Epistemic ATL-D (ATEL-D)

Syntax of ATEL-D

$$\varphi ::= p \mid \varphi \vee \varphi \mid \neg\varphi \mid \langle\langle A \rangle\rangle \mathbf{X}\varphi \mid \langle\langle A \rangle\rangle \varphi \mathbf{U} \varphi \mid \langle\langle A \rangle\rangle \varphi \mathbf{R} \varphi \mid$$
$$[\pi]^{+}\varphi \mid [\pi]^{-}\varphi \mid has(a, \alpha) \mid K_a\varphi$$
$$\pi ::= \varphi : \alpha \to A, B \mid \pi, \pi$$

where p is a proposition, α is an action, a is an agent, A and B are coalitions.

Semantics

- $K_a\varphi$: agent a knows that φ

Epistemic ATL-D (ATEL-D)

Syntax of ATEL-D

$$\varphi ::= p \mid \varphi \vee \varphi \mid \neg\varphi \mid \langle\langle A \rangle\rangle \mathbf{X}\varphi \mid \langle\langle A \rangle\rangle \varphi \mathbf{U} \varphi \mid \langle\langle A \rangle\rangle \varphi \mathbf{R} \varphi \mid$$
$$[\pi]^{+}\varphi \mid [\pi]^{-}\varphi \mid has(a, \alpha) \mid K_a\varphi$$
$$\pi ::= \varphi : \alpha \to A, B \mid \pi, \pi$$

where p is a proposition, α is an action, a is an agent, A and B are coalitions.

Semantics

- $K_a\varphi$: agent a knows that φ
- $[\varphi : \alpha \to A, B]^{+}$: in all states satisfying φ, agents in A are granted action α and agents in A and in B are informed of φ (the precondition)
- $[\varphi : \alpha \to A, B]^{-}$: similar to above, but the actions are removed

Informing agents

- To inform an agent a of the precondition φ, we partition the states depending on the satisfaction of φ, then we 'cut' the relations $\sim_a$ that connect the partitions
- $[\varphi : \alpha \rightarrow A]^+$ is captured by the ATEL-D update $[\varphi : \alpha \rightarrow A, \emptyset]^+$
- $[\varphi : \alpha \rightarrow A, \mathrm{Ag}]^+$ is a fully public update
- $[\varphi : \alpha \rightarrow \emptyset, \mathrm{Ag}]^+$ is a public announcement

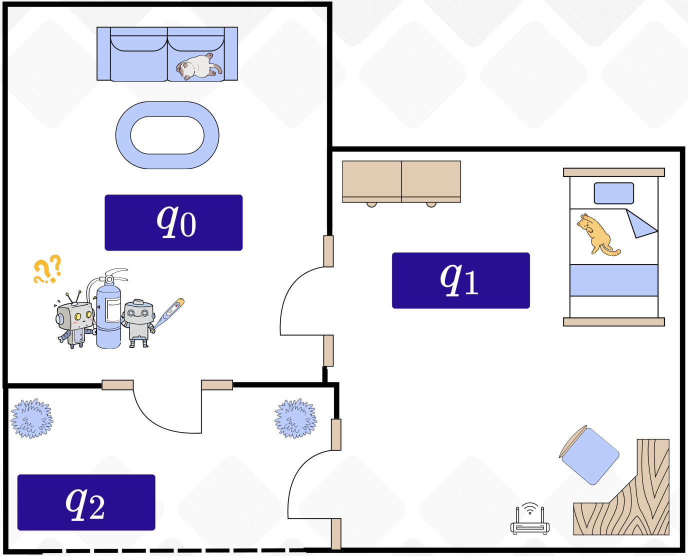

- The robots don't know where they are!
- r_2 has a temperature sensor and can distinguish the inside areas from the balcony (q_2)
- r_1 detects the Wi-Fi signal strength, which is only strong in Bob's room (q_1)

Example

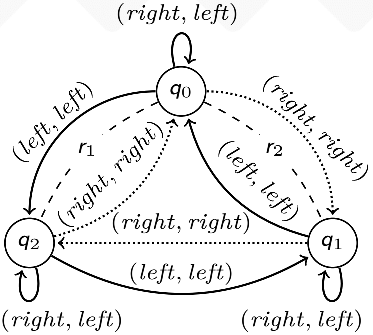

CGS with Imperfect Information $\mathcal{G}^{ii}$

Let us apply the same update as before, and inform the robots:

$$[\top : right \to r_2, r_1]$$

Example

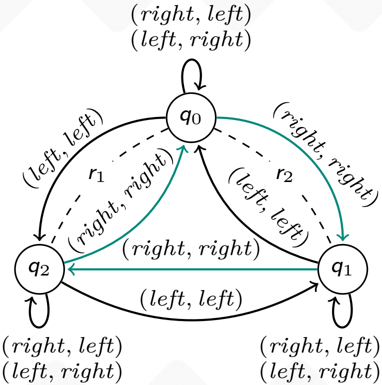

$$\mathcal{G}^{ii}, q \not\models [\top : right \to r_2, r_1]^{+} \langle\langle r_1, r_2 \rangle\rangle \mathbf{X}\, atBob$$

Because they don't know whether to go to the right or left

Example

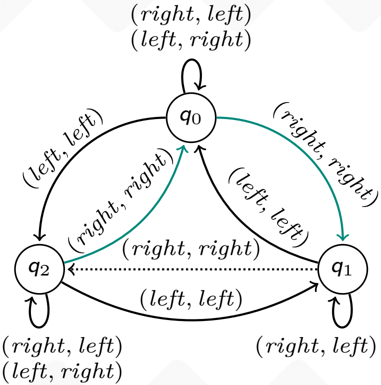

$$\mathcal{G}^{ii}, q \models [\neg atBob \land \neg warm : right \to r_2, r_1]^+ \langle\langle r_1, r_2 \rangle\rangle \mathbf{X}\, atBob$$

Model Checking

Theorem

Model checking ATEL-D is Δ_2^{P}-complete

The algorithm builds on the one for ATEL, with key new aspects:

- Label subformulas with the sequence of updates in their scope and order them so updates are computed before their effects are required
- Track available actions by labeling states, actions, transitions, and observation relations with update sequences

Conclusion

- Reasoning about updates of CGS: we focus on dynamic available actions
- Deciding which actions should be ‘learned’ or prohibited

Conclusion

- Reasoning about updates of CGS: we focus on dynamic available actions
- Deciding which actions should be 'learned' or prohibited
- Future work
  - Open problems: satisfiability, update existence problem for ATEL
  - Efficient search-based methods to synthesize updates
  - Agents that 'teach' other agents
  - Granting brand new actions, similar to (Galimullin, Gladyshev, et al. 2025)
  - Proactive agents reasoning about the interplay between knowledge and ability acquisition

Conclusion

- Reasoning about updates of CGS: we focus on dynamic available actions
- Deciding which actions should be 'learned' or prohibited
- Future work
  - Open problems: satisfiability, update existence problem for ATEL
  - Efficient search-based methods to synthesize updates
  - Agents that 'teach' other agents
  - Granting brand new actions, similar to (Galimullin, Gladyshev, et al. 2025)
  - Proactive agents reasoning about the interplay between knowledge and ability acquisition

Thank you for listening!

References

Galimullin, Rustam, Maksim Gladyshev, Munyque Mittelmann, and Nima Motamed (2025). "Changing the Rules of the Game: Reasoning About Dynamic Phenomena in Multi-Agent Systems". In: *Proceedings of the 24th AAMAS*. Ed. by Sanmay Das, Ann Nowé, and Yevgeniy Vorobeychik, pp. 829–838.

Galimullin, Rustam, Maksim Gladyshev, Munyque Mittelmann, and Nima Motamed (2026). "The Dynamic Turn in Strategy Logics". In: *Proceedings of the 24th AAMAS*.

Galimullin, Rustam and Louwe B. Kuijer (2024). "Synthesizing Social Laws with ATL Conditions". In: *Proceedings of the 23rd AAMAS*.

Grosinger, Hermine J. (2025). "The Next Level of Long-Term Agent Autonomy - Proactively Acquiring Knowledge and Abilities". In: *Proceedings of the 24th AAMAS*.

Hoek, Wiebe van der, Mark Roberts, and Michael J. Wooldridge (2007). "Social laws in alternating time: effectiveness, feasibility, and synthesis". In: *Synthese* 156.1, pp. 1–19.