Influencing Human-Robot Teams

Minae Kwon, Mengxi Li, Alex Bucquet, Sanjay Lall, Dorsa Sadigh

iliad intelligent and interactive autonomous systems

Stanford University

Motivation

Our goal is to enable a robot to model latent, dynamic leadership structures in human teams and use that information to influence a human team to reach some desired goal.

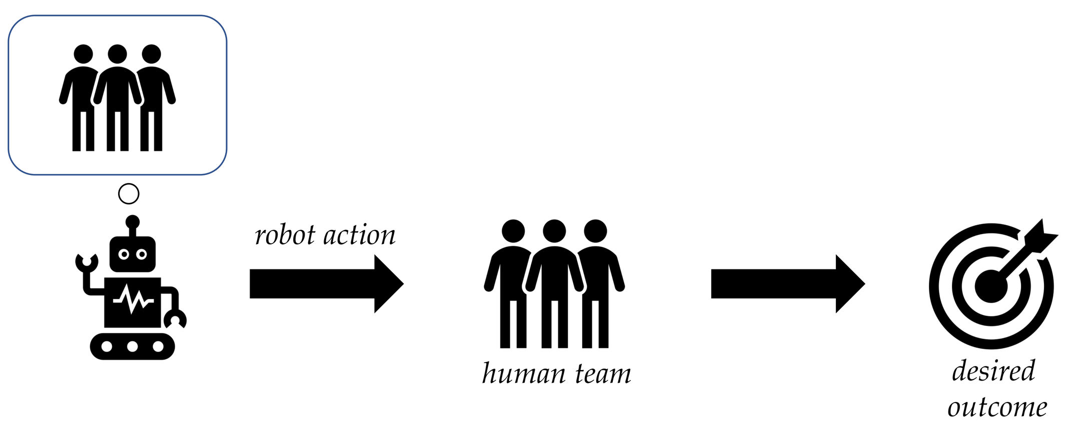

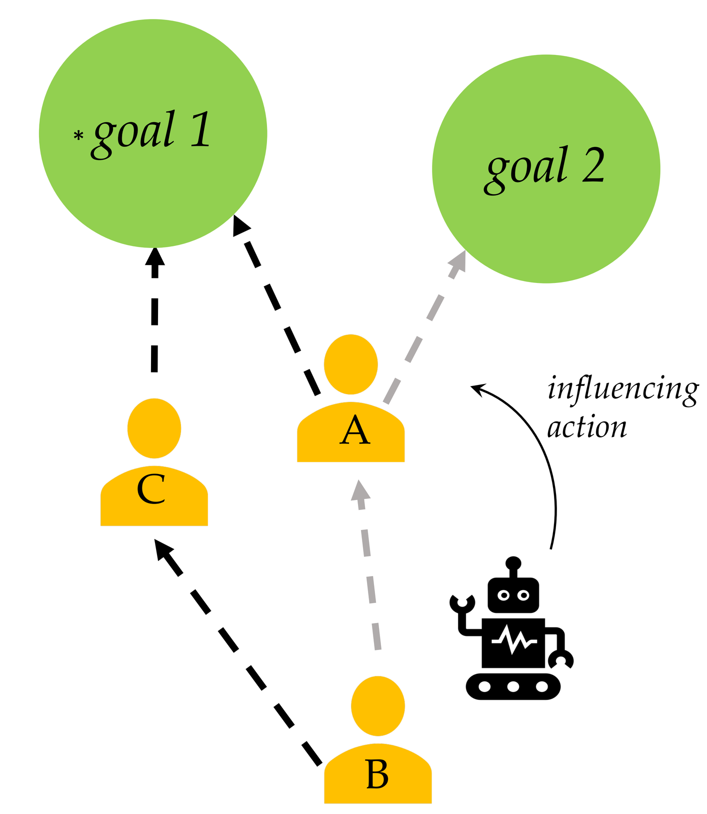

Leader-Follower Graph

Each agent *follows* a fixed goal or another agent, which we call a *leader*. An agent can change its leader at any time.

Leading and following behavior can be represented with a *leader-follower graph*. $w_{i,j}$ represents likelihood that j is i's leader, (i.e., a *leadership score*).

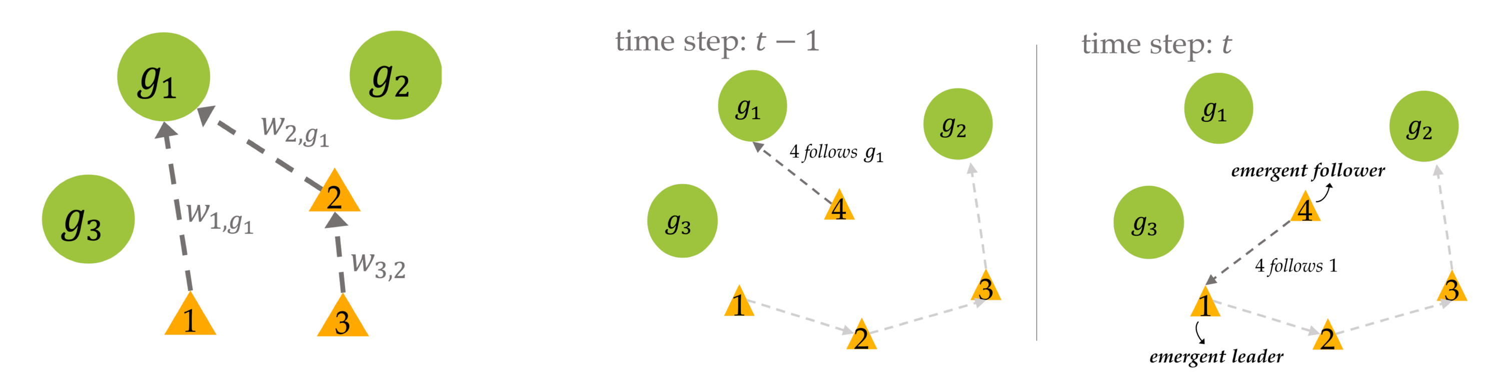

leader-follower graph

example of emergent leading and following behavior

Calculating Pairwise Leadership Scores

Determining agent 2's leader ($w_{2,j}$):

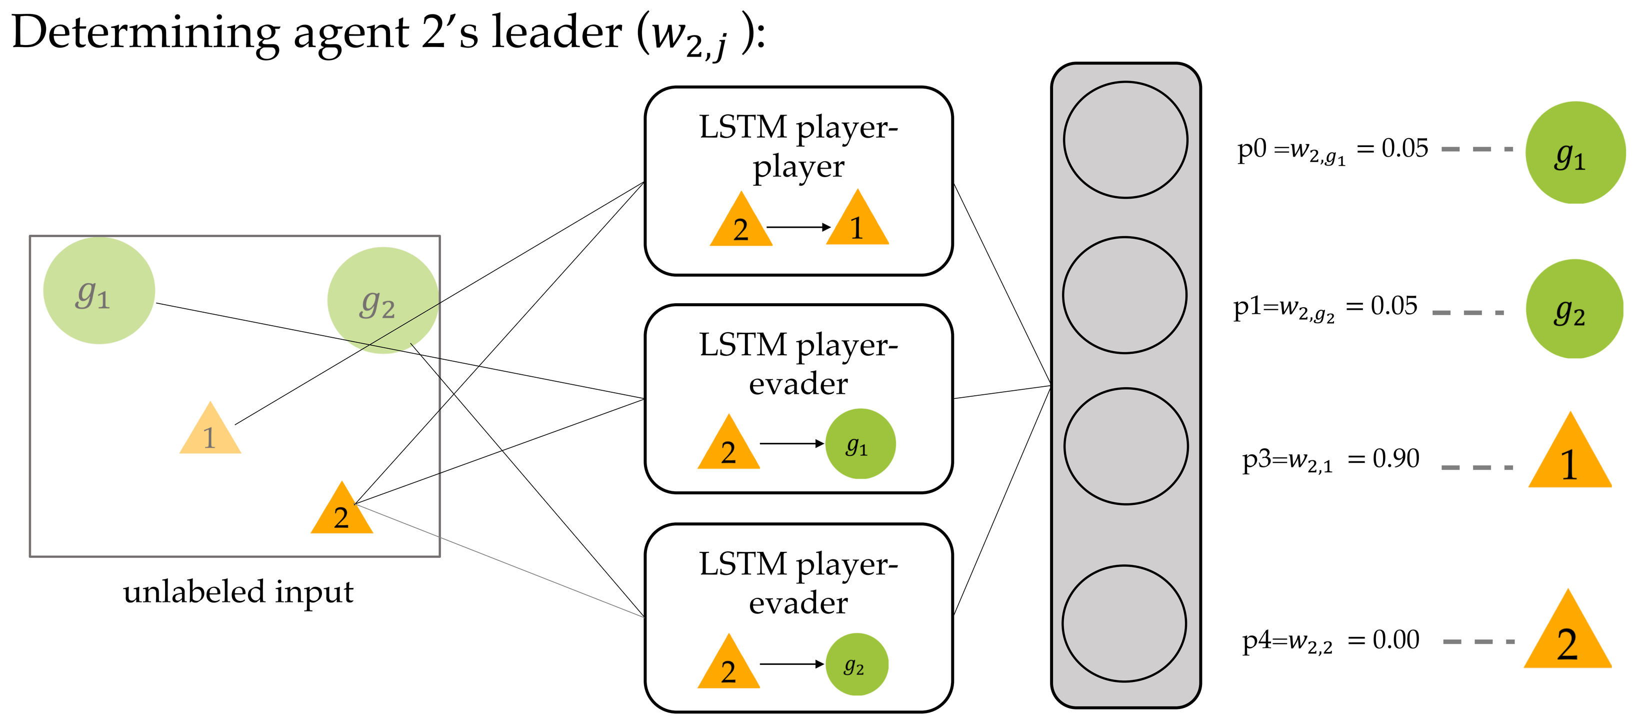

LSTM sub-units trained with supervised 2 player games

Softmax layer

Maximum Likelihood Leader-Follower Graph

Calculate pairwise leadership scores for all pairs of leaders and followers to create $\mathcal{G}$

- edge weights are $w_{i,j}$
- add a root node to extract maximum-weight arborescence

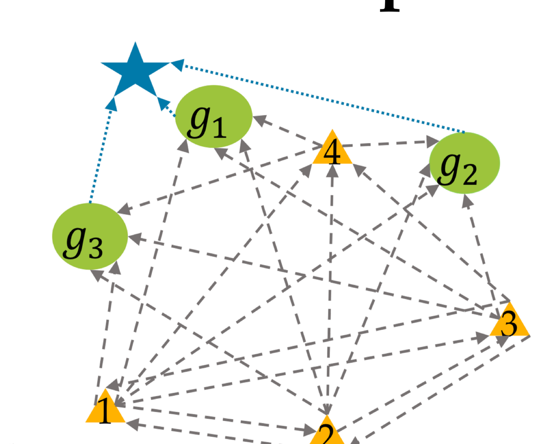

root node

Extract maximum-weight arborescence [1], equivalent to the maximum-likelihood leader-follower graph $\mathcal{G}^*$

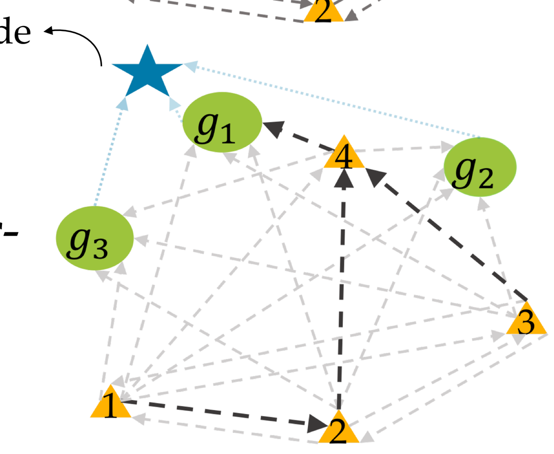

Planning to Encourage Role-Adaptation

Leveraging $\mathcal{G}^*$, robot identifies the *most influential leader*.

- a leader with the most number of followers
- influencing the most influential leader enables robot to indirectly influence its followers
- robot chooses a *leading* or *greedy* policy depending on who the most influential leader is

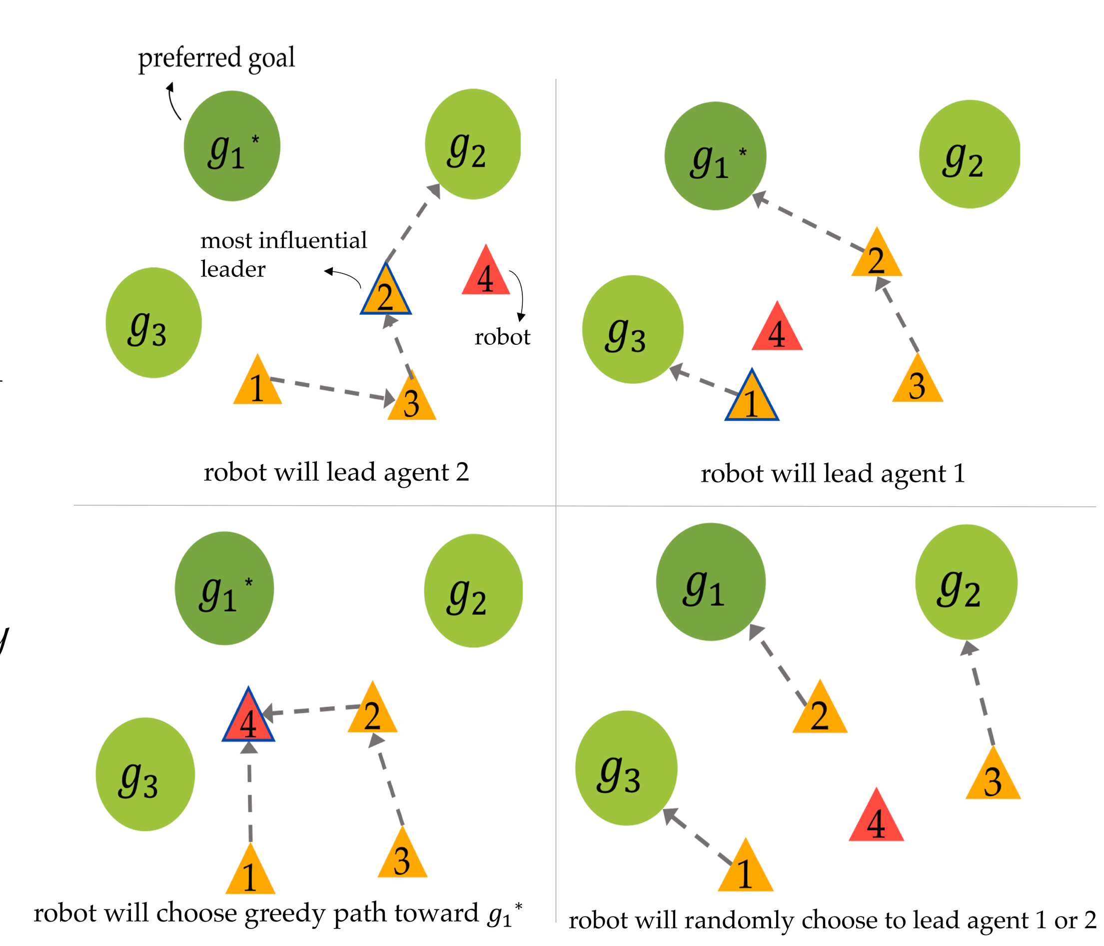

selecting best robot action:

$$a^*_{\text{robot}} = \begin{cases} \arg\max_{a_{\text{robot}}} w_{\text{most influential leader,robot}} & \text{if leading} \\ \arg\min_{a_{\text{robot}}} \|(x_{\text{robot}}, y_{\text{robot}}), (x_{g^*}, y_{g^*})\|_2 & \text{if greedy} \end{cases}$$

leading actions produces a "hovering" behavior around the most influential leader:

robot changes hovering distance when there are multiple potential influential leaders

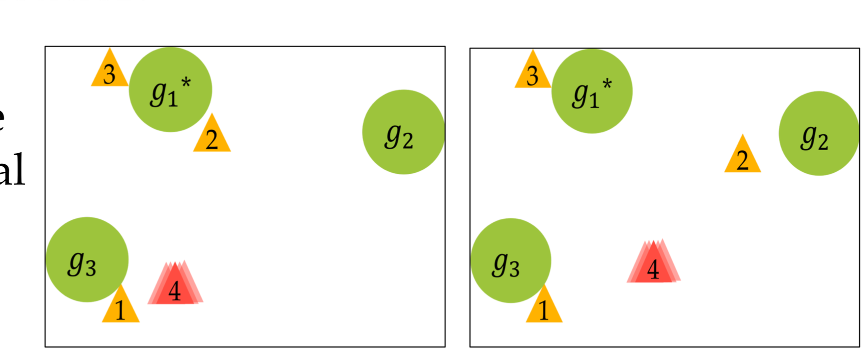

Evaluating the Robot as a Leader

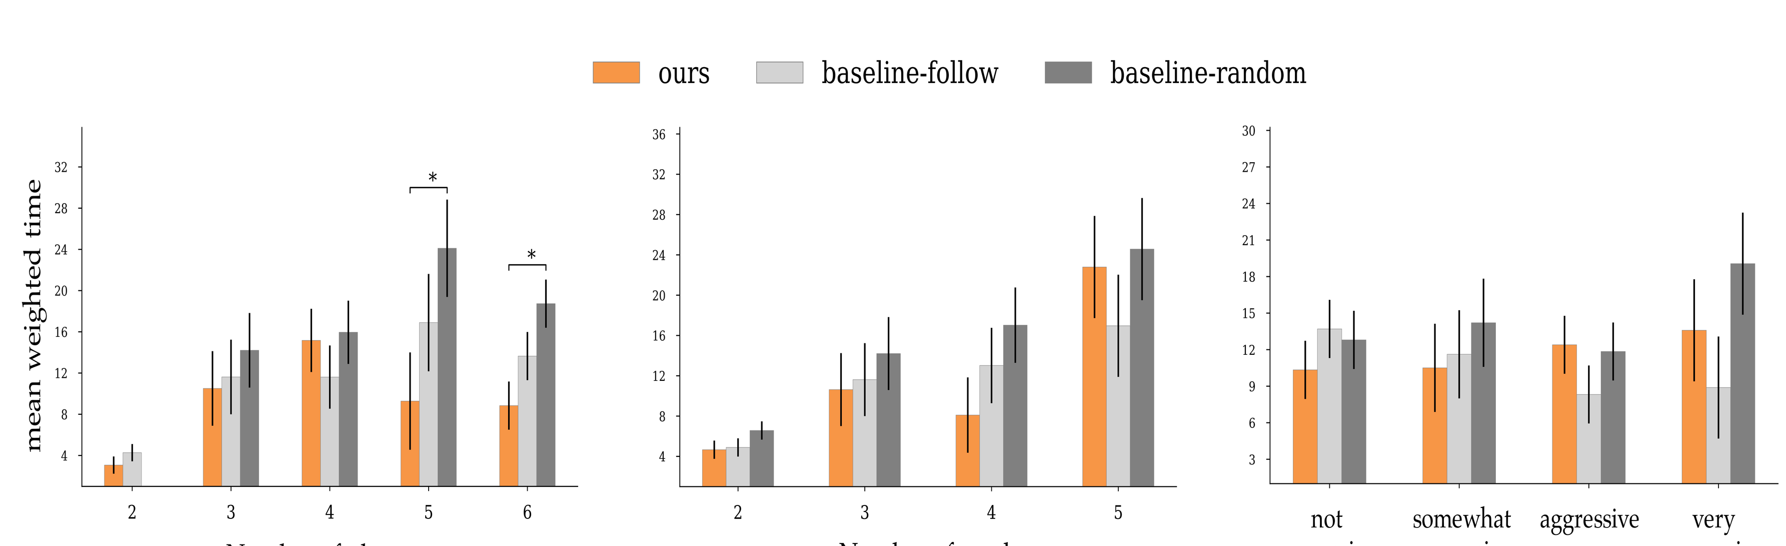

The robot's job was to lead its teammates toward the preferred goal. Only the robot knew where this goal was.

Our framework decreases the average time it takes for teams to reach the preferred goal across varying numbers of players, evaders, and robot aggressiveness.

Evaluating Emergent Adaptive Behaviors

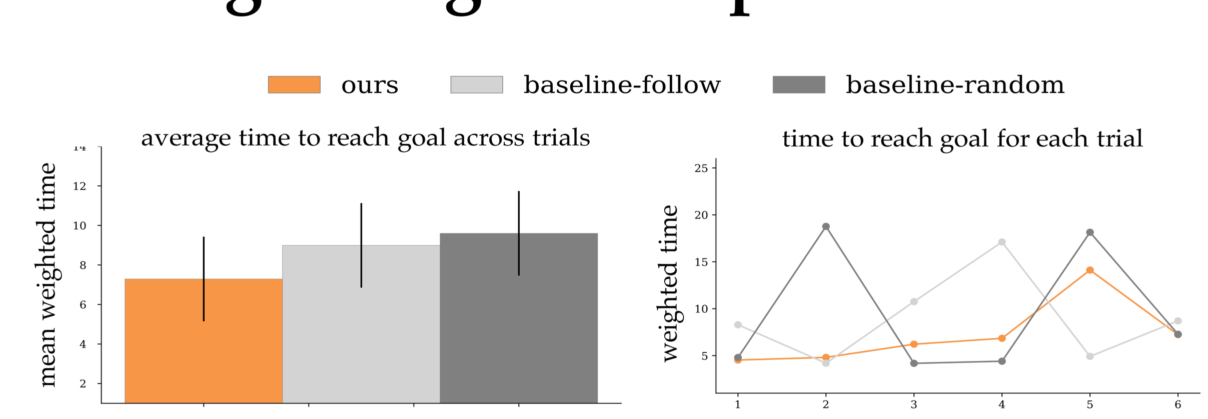

Knowledge of the preferred goal was randomly assigned to a human or robot teammate. This encourages teammates to adapt leader and follower roles each round.

Our framework decreases the average time it takes to capture the preferred goal with mixed leaders, suggesting teammates more fluidly adapted leader-follower roles compared to the baseline policies.

[1] J. Edmonds. Optimum branchings. Mathematics and the Decision Sciences, Part , 1(335-345): 412 25, 1968.