

Encouraging Human Interaction with Robot Teams: Legible and Fair Subtask Allocations

Soheil Habibian, *Student Member, IEEE*, and Dylan P. Losey

Abstract—Recent works explore collaboration between humans and teams of robots. These approaches make sense if the human is already working with the robot team; but how should robots encourage nearby humans to join their teams in the first place? Inspired by economics, we recognize that humans care about more than just team efficiency — humans also have biases and expectations for team dynamics. Our hypothesis is that the way inclusive robots divide the task (i.e., how the robots split a larger task into subtask allocations) should be both *legible* and *fair* to the human partner. In this paper we introduce a bilevel optimization approach that enables robot teams to identify high-level subtask allocations and low-level trajectories that optimize for legibility, fairness, or a combination of both objectives. We then test our resulting algorithm across studies where humans *watch* or *play* with robot teams. We find that our approach to generating legible teams makes the human’s role clear, and that humans typically prefer to join and collaborate with legible teams instead of teams that only optimize for efficiency. Incorporating fairness alongside legibility further encourages participation: when humans play with robots, we find that they prefer (potentially inefficient) teams where the subtasks or effort are evenly divided. See videos of our studies here: <https://youtu.be/cfN7O5na3mg>

Index Terms—Human-Robot Teaming, Intention Recognition, Acceptability and Trust

I. INTRODUCTION

IMAGINE sitting next to the team of robot arms in Figure 1. These robots are working together to clear a cluttered table: there are multiple tennis balls that the team needs to remove, and each robot arm is reaching to grab a different ball. You know the robots’ high-level task, but you do not know how the team will complete that task — or what you can do to help. How should the robot team *divide* and perform the task to *encourage* you to join in and collaborate?

Recent work on multi-agent reinforcement learning [1]–[3] and decentralized control [4] explores teams composed entirely of autonomous agents. Other current research brings a human into these teams: this includes learning to play with humans [5], influencing teams of humans [6], inferring the teammates’ roles [7], and forming new teams without pre-coordination [8]. Viewed together, these prior works take a step towards enabling teams of robots to collaborate with a human that is *already* participating in the team [9].

But how should robots encourage humans to join their teams in the first place? We hypothesize that the way the robots divide the task among agents — and the way each robot performs its subtask — will affect the human’s willingness




Figure 1. Robot team clearing a cluttered table. A human is thinking about joining this team, but the human does not know *a priori* which ball they should grab. We find that robot teams can encourage humans to collaborate by purposefully dividing the task into *legible* and *fair* subtask allocations.

to collaborate with the robot team. Prior work in behavioral economics [10]–[12] suggests that task efficiency is not the only factor that encourages humans to form teams. Instead of treating human teammates like robot partners, robots must account for human biases and expectations:

*Humans are inclined to participate in teams
where their roles are legible and fair.*

Returning to our motivating example from Figure 1, there are many equally efficient ways for the robots to grab the balls and clear the table. But robots which apply our insight reach for balls 1 and 2: by moving for the other end of the table they make the human’s role clear, and evenly divide the subtasks among the three teammates. Accordingly, in this paper we study two axes for subtask allocation: *legibility* and *fairness*. We experimentally test how both of these factors encourage participation when humans are watching robot teams and when humans join in and play with those teams. Our work is motivated by mixed-autonomy settings (such as factory floors) where we want to facilitate human-robot collaboration.

Overall, we make the following contributions:

Formalizing Legible and Fair Subtask Allocations. We first introduce a bilevel optimization approach that enables centralized robot teams to identify high-level subtask allocations and low-level trajectories. We then formulate legible and fair allocations when the human is watching and playing, and incorporate both into our bilevel optimization.

Encouraging Humans to Join Robot Teams. We conduct two online user studies where humans watch robot teams. In the first user study we find that humans prefer legible teams over teams that only optimize for task performance. In the

second user study we also incorporate fairness, and compare legible teams to teams that are both legible *and* fair.

Encouraging Humans to Keep Collaborating. We next evaluate our approach in an in-person user study where humans collaborate with two robot arms. We find that humans prefer to keep working with teams of robots that optimize for legibility instead of efficiency. We also find that fairness has a statistically significant impact: legible and fair teams better encourage collaboration than teams which are only legible.

II. RELATED WORK

Multi-Agent Teams. Recent works enable teams composed purely of autonomous agents to perform collaborative tasks [1]–[4]. However, it is still challenging for these autonomous teams to incorporate humans [5], [8], [13]. One common method for facilitating human-robot teams is introducing subtasks (or roles), and then assigning these subtasks to the robots and humans [6], [14], [15]. Gombolay *et al.* [16] find that humans prefer to work in teams where robots assign subtasks, and Wang *et al.* [7] suggest that humans can infer the subtasks of autonomous agents by observing their motion. We build on these approaches by similarly using subtasks. But unlike prior works — where often there is only one robot, and the human and robot have already formed a team [9] — here we focus on *bringing* one human into a team with *multiple* robots.

Legible Interaction. Robots can leverage their behavior to implicitly communicate goals, objectives, or uncertainty to human partners [17]–[19]. Most relevant is research by Roncone *et al.* [20], which indicates that transparency is key when humans and robots are deciding subtask allocations. But while prior work explores how one robot should convey its intent to the human [21], how should *teams* of robots communicate their overall allocation? We will extend legibility from dyads to teams, and optimize the team’s behavior so that human onlookers can infer their intended subtask(s).

Fairness in Human-Robot Teams. Research from psychology and economics indicates that humans have expectations for the teams they join: in particular, humans expect those teams to be fair [22], [23]. While prior work largely focuses on human-human teams, state-of-the-art studies extend those same principles to human-robot teams [24]–[26]. For instance, Claire *et al.* [24] impose a fairness constraint on resource distribution in multi-armed bandits, and find that this fairness impacts the the human’s trust in the system. Other works explore human perceptions of fairness in terms of assigned workload, member capability, and task type [25], [26]. In this paper we leverage definitions of fairness that are consistent with prior works, but we now focus on how the fairness of subtask allocations affects the human’s willingness to *join* robot teams.

III. PROBLEM SETTING

We explore scenarios where a robot team is collaborating to perform a task, and these robots want to encourage nearby humans to join in and participate. As our *running example*, consider the two robot arms that are clearing a cluttered table in Figure 1. We introduce additional structure by dividing

the overall task into *subtasks*. These subtasks could be steps towards a larger goal — e.g., removing one tennis ball from the cluttered table — or they could be roles within the task — e.g., leader and follower. We assume that the team of robots are centralized: they communicate with one another in real time and share a common controller.

MDP with Subtasks. From the robots’ perspective this is an instance of a Markov decision process (MDP) with subtasks: $\langle \mathcal{S}, \mathcal{A}, T, R, \gamma, \mathcal{T} \rangle$. Let $s \in \mathcal{S}$ denote the system state, let $a \in \mathcal{A}$ denote the system action, and let $s^{t+1} = T(s^t, a^t)$ be the system dynamics. We emphasize that the state s and action a contain the combined states and actions of every robot teammate: returning to our running example, the action a is the joint velocities of both robot arms. At each timestep t the robots take actions to interact with the environment. We write the team’s sequence of states and actions up to time t_K as a *trajectory* $\xi = \{(s^0, a^0), \dots, (s^{t_K}, a^{t_K})\}$.

The team of robots is collaborating to perform a task. We capture this objective through the sparse reward function $R(s) = \mathbb{1}\{\text{task solved in } s\}$, which indicates whether or not the task is complete at state s with discount factor $\gamma \in [0, 1)$. But we also break the overall task into subtasks: let $\tau \in \mathcal{T}$ be a subtask, and let \mathcal{T} be the finite set of required subtasks. By completing all of these subtasks the team reaches the goal state and receives reward $R = 1$. In our running example the task is to clear the table, and the three subtasks τ_1, τ_2 , and τ_3 are removing the tennis balls marked 1, 2, and 3.

Allocations. Introducing subtasks brings with it a challenge: how should the team of robots divide these subtasks among themselves and the nearby human? Let $\theta \in \Theta$ denote an *allocation*. This allocation θ determines which subtask(s) each robot will perform and which subtask(s) the robots want the human to complete. Our running example has three subtasks $\mathcal{T} = \{\tau_1, \tau_2, \tau_3\}$. Here one allocation could be $\theta = \{(human : \tau_2), (robot_1 : \tau_3), (robot_2 : \tau_1)\}$, indicating that the human is assigned subtask τ_2 and the robots will perform subtasks τ_1 and τ_3 . It is also possible for an agent to be assigned either *no* subtasks or *multiple* subtasks: for instance, $\theta = \{(human : \emptyset), (robot_1 : \tau_1, \tau_2), (robot_2 : \tau_3)\}$. Moving forward we will use $\mathcal{T}_i(\theta)$ to refer to set of subtask(s) assigned to the i -th agent under allocation θ .

Fixed Robots. The robots assume that the human will follow their chosen allocation $\theta \in \Theta$. Put another way, the robots are *fixed*: they select and execute a single allocation during each interaction, and do not switch their allocation in response to the human [17]. We make this assumption in order to isolate how the human responds to the robots, and avoid entangling this with how the robots respond to the human.

IV. OPTIMIZING FOR LEGIBLE AND FAIR ALLOCATIONS

Given the formulation from Section III, we search for a subtask allocation θ that encourages nearby humans to join in and participate with the robot team. We hypothesize that robots encourage human participation by optimizing along two axes: *legibility* and *fairness*. Here legibility refers to how clearly the robots convey their allocation θ to the human: e.g., which tennis balls are the robots clearing, and which ball should the

human reach for? Fairness captures how the evenly subtasks are divided among the teammates: e.g., is the human expected to remove the same number of balls as each robot teammate? Below we formally define legibility and fairness, and introduce a bilevel optimization approach for identifying legible and fair allocations and trajectories.

Bayesian Inference. At the start of the interaction the human is not sure what subtasks each agent should complete; however, the human can infer $\theta \in \Theta$ based on the robots' behavior ξ . Applying Bayes' theorem:

$$P(\theta | \xi) = \frac{P(\xi | \theta) \cdot P(\theta)}{\sum_{\theta' \in \Theta} P(\xi | \theta') \cdot P(\theta')} \quad (1)$$

where $P(\theta)$ is the prior over allocations and $P(\xi | \theta)$ is the likelihood of allocation θ given team trajectory ξ . Since the trajectory ξ is composed of conditionally independent state-action pairs (s, a) , we rewrite the likelihood function:

$$P(\xi | \theta) = \prod_{(s,a) \in \xi} P((s,a) | \theta) \quad (2)$$

We assume that the human views the team as a Boltzmann-rational agent. This model — commonly used in robotics [27], [28] and economics [29] — assigns higher likelihood to actions that lead to increased long-term reward:

$$P(\xi | \theta) \propto \exp \sum_{(s,a) \in \xi} Q_\theta(s, a) \quad (3)$$

Here $Q_\theta(s, a)$ is the cumulative reward of taking action a in state s and optimally completing allocation θ thereafter. Within our running example the allocation θ tells each agent which ball to reach for, and $Q_\theta(s, a) = -d(T(s, a), g_\theta)$ is the negative distance between the next state and the goal state g_θ . Combining Equations (1)–(3) provides a model that the robots can evaluate for $P(\theta | \xi)$.

A. Efficient Allocations

Before considering legibility or fairness, we start with an *efficient* baseline. This team of robots optimizes purely for task performance: put another way, the robots select an allocation $\theta \in \Theta$ that maximizes their long-term reward and completes the task as quickly as possible. In our experiments we select allocations that are noisily-optimal: $P(\theta_{\text{Eff}} | s) \propto V_\theta(s)$, where $V_\theta(s)$ is cumulative reward for starting at s and completing the task using allocation θ . Within our running example this produces allocations θ_{Eff} where each agent removes one tennis ball, and the robots reach directly towards their allocated ball.

B. Legible Allocations

We next leverage $P(\theta | \xi)$ to optimize for legible — but potentially inefficient — allocations. We emphasize that *legible allocations* are different from *legible motions* [21]. Within legible motions the team of robots is selecting a trajectory ξ that communicates one specific allocation θ , i.e., the robots are maximizing $P(\theta | \xi)$. Legible allocations require another level of optimization: now the team of robots must not only find the best way to convey a given θ , but they must also

determine which allocation $\theta \in \Theta$ they are able to convey most clearly. This results in a bilevel optimization problem across continuous trajectories and discrete allocations.

Watching. When the human is watching a team of robots — and is not participating in the task themselves — we identify legible allocations and trajectories by optimizing:

$$(\theta_{\text{Leg}}, \xi_{\text{Leg}}) = \underset{\theta \in \Theta, \xi \in \Xi}{\operatorname{argmax}} P(\theta | \xi) \quad (4)$$

Here the lower-level optimization problem is finding the trajectory ξ that maximizes the likelihood of θ , and the upper-level optimization problem iterates through each choice of $\theta \in \Theta$ to find the allocation that maximizes $P(\theta | \xi)$.

Playing. Legible allocations change when the human is joining in to collaborate with the robots. Now the human is not concerned with the overall allocation for every teammate — instead, the human only needs to know their own subtask(s). Returning to our running example, it does not matter which robot is picking up ball 1 or ball 3; the human just needs to know that their job is picking up ball 2. Hence, we introduce $\mathcal{H}(\theta)$, the set of all allocations $\theta^\theta \in \Theta$ where the human has the same subtasks. More formally, $\mathcal{H}(\theta) = \{\theta^\theta | \mathcal{T}_H(\theta^\theta) = \mathcal{T}_H(\theta)\}$. We then sum across this set:

$$(\theta_{\text{Leg}}, \xi_{\text{Leg}}) = \underset{\theta \in \Theta, \xi \in \Xi}{\operatorname{argmax}} \sum_{\theta^\theta \in \mathcal{H}(\theta)} P(\theta^\theta | \xi) \quad (5)$$

Intuitively, Equation (5) marginalizes out the robots' specific roles, leading to an allocation θ_{Leg} where the human can best infer what subtask(s) they are meant to do.

C. Fair Allocations

Optimizing for legibility enables the human to understand their role within the team. But just because their subtask(s) are clear does not mean that the human will want to complete these subtask(s). We therefore introduce *fairness* as a second axis for encouraging collaboration. Our approach is agnostic to the specific function used to quantify fairness, but in our experiments we tested *two definitions* from related works on economics [10], [11] and robotics [24], [25]. Let $f_i(\theta, \xi)$ be the fairness for agent i given allocation θ and trajectory ξ . Our first approach to fairness is *equality of allocation*:

$$f_i(\theta, \xi) = - \left| \frac{|\mathcal{T}_i|}{N} - |\mathcal{T}_i(\theta)| \right| \quad (6)$$

Here N is the total number of teammates, and allocation θ is fair if it assigns an equal number of subtasks to each agent. Our second approach is *equality of effort*:

$$f_i(\theta, \xi) = - \left| \frac{d(\xi)}{N} - d_i(\xi) \right| \quad (7)$$

where $d(\xi)$ is the overall distance the team must travel under trajectory ξ , $d_i(\xi)$ is the distance the i -th agent travels, and an allocation is fair if each team member travels the same distance. Regardless of our chosen definition for f , the team of robots again solves a bilevel optimization problem to identify allocations that are now *both fair and legible*.

Watching. When the human is watching a team they are not compelled to take the perspective of any specific agent. Put

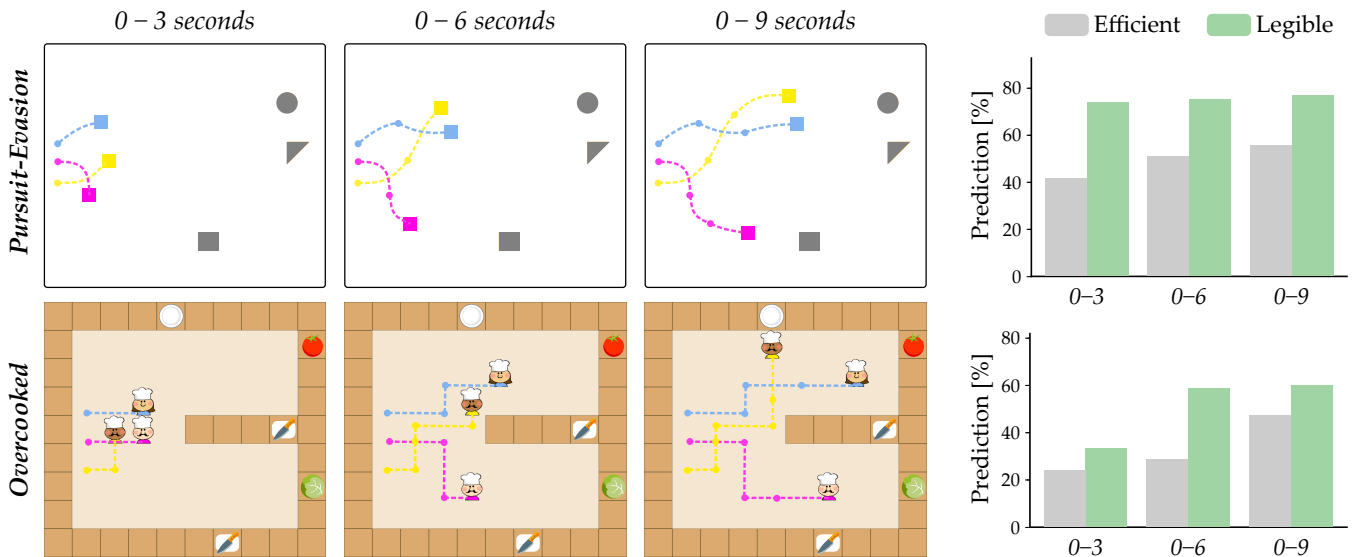


Figure 2. Environments and results from our online user study. (Left) Participants watched videos of simulated robot teams in *Pursuit-Evasion* and *Overcooked* environments. We showed videos of the team’s behavior during the first three seconds, six seconds, and nine seconds: here dotted lines depict an example of the trajectories the agents traveled in each video snippet. (Right) Based on these videos participants predicted the team’s subtask allocation. Users more accurately predicted the subtasks of teams that optimized for **Legible** allocations as compared to **Efficient** teams that optimized for task performance.

another way, the allocation should be fair for *every* agent. We therefore solve for fair and legible allocations by optimizing:

$$(\theta_{Fair}, \xi_{Fair}) = \operatorname{argmax}_{\theta \in \Theta, \xi \in \Xi} \sum_{i=1}^N f_i(\theta, \xi) + P(\theta | \xi) \quad (8)$$

We note that this builds on Equation (4), and now encourages the robots to leverage fair allocations which the human can correctly interpret.

Playing. The human’s perspective changes when they collaborate and actively participate with the team of robots. Here we hypothesize that the human focuses on the how fair the allocation is for *themselves* — e.g., is the human being asked to pick up and move more tennis balls than both of the robots? Consistent with Equation (5), we optimize:

$$(\theta_{Fair}, \xi_{Fair}) = \operatorname{argmax}_{\theta \in \Theta, \xi \in \Xi} f_H(\theta, \xi) + \sum_{\theta' \in H(\theta)} P(\theta' | \xi) \quad (9)$$

where f_H is the fairness for the human teammate.

V. ENCOURAGING PARTICIPATION WHEN WATCHING

Before describing this experiment we first want to outline our user studies in Section V and VI. In each section we will compare robot teams that optimize for efficiency, robot teams that optimize for legibility, and robot teams that optimize for legibility and fairness. The key difference between these two sections is whether the human is *watching* (Section V) or *playing* (Section VI) with the robot team. Both aspects are important: we want to determine which types of allocations draw the human into the robot team (*watching*) and encourage the human to continue collaborating with that team (*playing*). Moreover, recall that we have different equations for legibility and fairness in each context: when the human is watching they can consider the task from the perspective of any agent, but once the human joins they must focus on the legibility and fairness of their own subtask.

Here we start with *watching*. We perform two separate online user studies where humans observe multi-agent teams in two simulated environments (see Figure 2). In the first study we compare robot teams that only consider efficiency to robot teams that optimize for legibility. We test whether these legible allocations actually make it easier for humans to predict the subtasks, and whether legible allocations encourage onlookers to join robot teams. In the second study we compare legible teams to teams that are both legible *and* fair.

Environments. We used two simulated environments from prior work on multi-agent teams and legible motion. Both environments had an overall task that was divided into subtasks \mathcal{T} for the agents to complete. In *Pursuit-Evasion* [6], [21] the state-action space is continuous and the multicolored agents are trying to reach the gray targets. Each agent has at least one subtask (i.e., $\mathcal{T}_i(\theta) \neq \emptyset$), and it is possible for multiple agents to share the same target. In *Overcooked* [5], [7] agents operate in a discrete state-action space and their targets include ingredients (lettuce and tomato) or kitchen utensils (cutting board and plate). Here agents never shared the same target and an agent could be assigned no subtasks (i.e., $\mathcal{T}_i(\theta) = \emptyset$).

Participants. We recruited 100 total participants on Amazon Mechanical Turk (MTurk). Our experiment was only available to English speaking participants who had completed at least 100 Human Intelligence Tasks (HITs) with a 99% approval rating. Users had to correctly answer qualifying questions to ensure that they had read and understood our instructions before they could participate. We then divided these participants into two groups: 50 users completed the study in Section V-A that compares efficiency and legibility, and the other 50 users completed the study in Section V-B that compares legibility without fairness to legibility with fairness.




Figure 3. User preferences when watching **Efficient** and **Legible** robot teams. Fifty participants responded to three forced-choice comparisons by indicating which team they would prefer to work with in the future. For example, in *Pursuit-Evasion* users preferred **Legible** teams 111 times and **Efficient** teams 39 times. Here * denotes statistical significance ($p < .001$).

A. Legibility

Our baseline is a team of robots that optimize for efficiency: these robots choose allocations to complete the task as quickly as possible. Here we compare that baseline to a *legible* robot team. Participants watch teams of agents complete tasks in our simulated environments. We test whether legible teams better convey their allocation to the human, and whether humans prefer to join these teams.

Independent Variables. We compared two types of subtask allocations: **Efficient** and **Legible**. In **Efficient** the robots selected noisily-rational allocations θ_{Eff} . Under **Legible** the robots selected the allocation θ_{Leg} that optimizes Equation (4) and attempts to reveal every agent’s subtask. To better isolate legibility, we tested **Efficient** teams that *were not* legible.

Procedure. Our 50 participants first watched videos of **Efficient** and **Legible** allocations (see Figure 2). These videos showed the team’s behavior after three, six, and nine seconds had passed. While watching these videos participants indicated which subtask they thought each agent was completing; e.g., based on the motion from 0 – 9 seconds, a user might guess that the blue chef is reaching for the red tomato. Participants had to indicate their prediction at the current timestep before watching the next increment. Finally, participants were shown three side-by-side comparisons of **Efficient** and **Legible** teams. After watching both robot teams complete the task participants were asked to choose one team to join. We randomized the order of **Efficient** and **Legible** teams; participants were never told which allocations they were watching.

Dependent Measures. We counted the total number of times the participants correctly guessed the subtasks of all three agents when watching an **Efficient** allocation and when watching a **Legible** allocation. We also recorded the total number of times participants preferred **Efficient** allocations, and the total number of times participants preferred **Legible** allocations.

Hypotheses. We had two hypotheses in this user study:

H1. *Human observers will more accurately predict the subtasks of teams that optimize for legible allocations.*

H2. *Humans will prefer to join teams that optimize for legible subtask allocations.*




Figure 4. User preferences when watching **Legible** and **Legible+Fair** robot teams. Fifty participants watched eight pairs of teams and selected the teams they would prefer to join. We tested two definitions of fairness: in *Pursuit-Evasion* the robots maintained equality of effort, and in *Overcooked* the robots optimized for equality of allocation. Participants preferred **Legible+Fair** teams in the *Overcooked* environment ($p < .001$), but the differences were not statistically significant in *Pursuit-Evasion* ($p = 0.271$).

Results. The results of our first *watching* survey are shown in Figures 2 and 3. Across both *Pursuit-Evasion* and *Overcooked* environments, more participants correctly predicted the roles of each team member when watching **Legible** teams (Figure 2). This prediction accuracy increases as the interaction unfolds: when the agents get closer to completing the task their subtasks become increasingly clear to the human. Overall, the participants’ responses support **H1**.

Importantly, the *legibility* of a team’s allocation affected people’s willingness to join that team. Figure 3 displays the participant’s preference when asked: “if you had to join team A or team B, which would you join?” In both environments participants chose **Legible** teams more frequently. Two Wilcoxon signed-rank tests showed that these differences were statistically significant in *Pursuit-Evasion* ($Z = -5.879$, $p < .001$) or in *Overcooked* ($Z = -4.736$, $p < .001$). These results across 50 participants are in line with **H2** and suggest that legible multi-robot teams encourage human participation.

B. Fairness

The results of the first half of our *watching* study indicate that legible robot teams convey allocations to the human observer, and that humans prefer to join these legible teams. But is the transparency of subtask allocations the only parameter that encourages human participation? Here we test the effects of *fairness* when humans are watching robot teams.

Independent Variables. Remember that our fair teams optimize for legibility in addition to fairness — accordingly, we refer to this new condition as **Legible+Fair**. **Legible+Fair** teams used Equation (8) to identify allocations that were fair for all agents. As a baseline, we compared these teams to the purely **Legible** allocations from the previous part. To isolate fairness, we selected **Legible** teams that *were not* fair.

We also varied the definition of fairness used in Equation (8). Within *Pursuit-Evasion* we defined fairness as equality of effort from Equation (7), and within *Overcooked* we defined fairness as equality of allocation from Equation (6). In practice, **Legible+Fair** teams maintained an equal travel

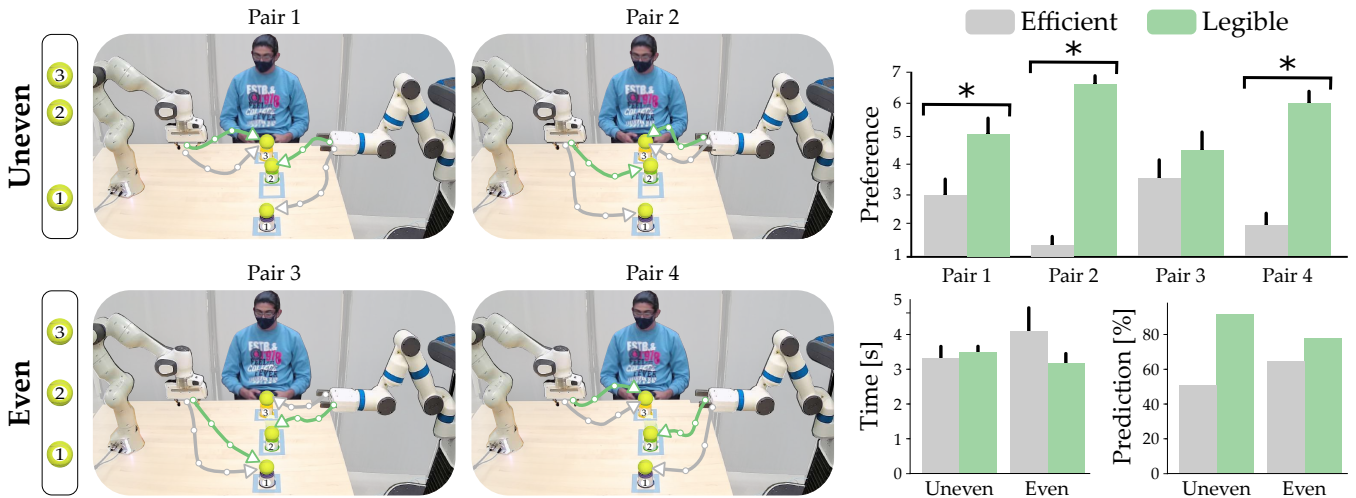


Figure 5. Comparing **Efficient** and **Legible** teams during our in-person user study. (Left) Participants collaborated with two 7-DoF robot arms to clear tennis balls off the table. Users compared pairs of teams: in each pair one team optimized for completing the task efficiently, and the other optimized for making the human’s role legible. Note that the **Legible** allocations were not necessarily fair, and often the human had to reach across the table to complete their subtask. (Right, Bottom) Participants more accurately predicted their role with **Legible** teams, but differences in prediction time were not statistically significant. (Right, Top) Participants preferred working with **Legible** teams. Here * denotes statistical significance ($p < .05$).

distance for each agent in *Pursuit-Evasion*, and gave all agents an equal number of subtasks in *Overcooked*.

Procedure. Participants watched sixteen pairs of robot teams with different subtask allocations (eight pairs in *Pursuit-Evasion* and eight pairs in *Overcooked*). Each pair contained a **Legible+Fair** team and a **Legible** team. We randomized the order of the robot teams, and did not tell participants what objective each team was optimizing for. After watching each pair of teams participants selected the one they would prefer to collaborate with. Our hypothesis was that users would prefer to join teams that were both legible and fair:

H3. *Humans will prefer to join legible and fair robot teams where all members contribute equally to the task.*

Results. The results of our second user survey are displayed in Figure 4. As a reminder, here we are tallying the total number of pairs where participants selected **Legible** teams, and the total number of pairs where participants selected **Legible+Fair** teams. Interestingly, the results were not consistent across environments. Within *Pursuit-Evasion* 47% of the participants favored **Legible+Fair** teams, while in *Overcooked* 77% of the users preferred **Legible+Fair**. Applying Wilcoxon signed-rank tests, the differences in participant preferences were not statistically significant in *Pursuit-Evasion* ($Z = -1.100$, $p = .271$). By contrast, participants *did* prefer to join **Legible+Fair** teams in *Overcooked* ($Z = -10.7$, $p < .001$).

To explain these results, recall that in the *Pursuit-Evasion* environment we defined fairness as equality of effort. This means that **Legible** teams could cause members to reach for goals that were farther away. However, this distance did not seem to affect the human’s perception: participants were just as willing to join teams where members had to travel unequal distances as teams that maintained equality of effort. By contrast, we focused on equality of allocation in *Overcooked*. Here participants preferred **Legible+Fair** robot teams where all agents have an equal role — i.e., people avoided teams where one or two chefs had to complete all the subtasks.

Our results partially support **H3**. People favored teams that optimized equality of allocation, but did not show a preference for teams that maintained equality of effort. This may have been because participants were only *watching* teams and not actively *playing* with those teams.

VI. ENCOURAGING PARTICIPATION WHEN PLAYING

In Section V we compared **Efficient**, **Legible**, and **Legible+Fair** allocations when humans were watching robot teams. Here we compare those same three approaches, but now with users that are *actively participating* with the robot team. We repeat these studies because *watching* is different from *playing*: when humans watch teams they can consider the perspective of any agent; but when humans join in and collaborate with robot teams they must focus on their own allocation, and respond in real-time to the behavior of their robot teammates.

Experimental Setup. Participants worked with two 7-DoF robot arms (Fetch, Fetch Robotics and Panda, Franka Emika). The two robots were centralized and shared a common controller. We placed three tennis balls within the workspace of the robot team: participants had to join in and help the robot team clear these tennis balls off the table (see Figure 5).

Participants. We recruited 11 participants (3 female, 1 non-binary, ages 26 ± 3.3 years) from the Virginia Tech community. All participants provided informed written consent consistent with university guidelines (IRB #20-755). We used a within-subjects design: every participant interacted with **Efficient**, **Legible**, and **Legible+Fair** robots, and performed both parts of the user study described in Sections VI-A and VI-B.

A. Legibility

In the first half of this user study we tested whether *legibility* encourages humans to keep playing with robot teams. We compared efficient allocations (that minimize interaction time) to legible allocations (that communicate the human’s role). We

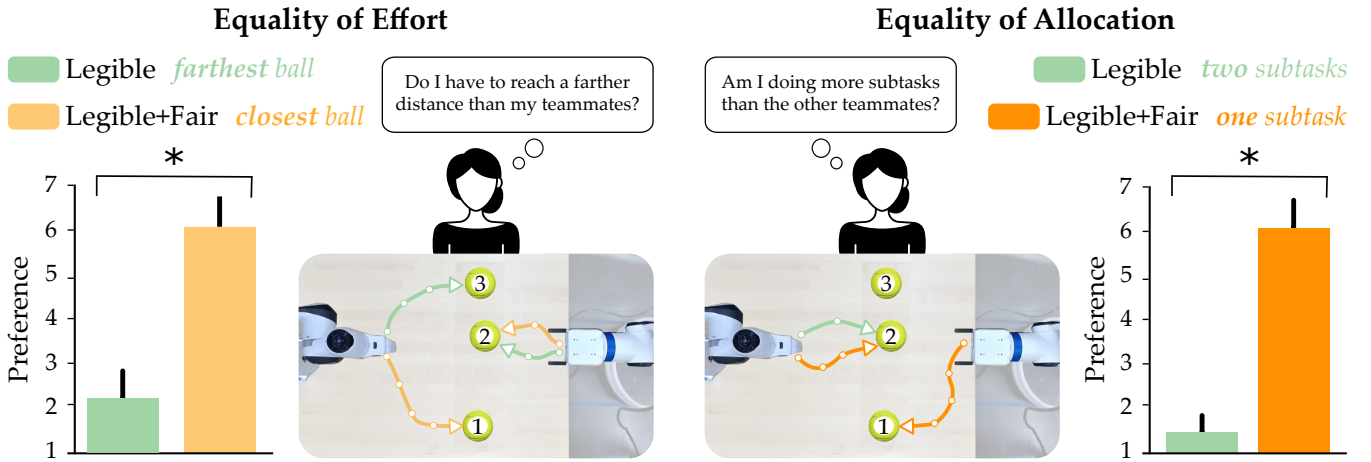


Figure 6. Comparing **Legible** and **Legible+Fair** teams during our in-person user study. We tested two definitions of fairness. (Left) Under equality of effort the **Legible+Fair** robot chose allocations where each agent travelled the same distance. (Right) Under equality of allocation the **Legible+Fair** robot chose allocations where the human was allocated one subtask. For both definitions users preferred **Legible+Fair** teams. Here * denotes $p < .001$.

measured how accurately each participant was able to identify their subtask, as well as the participant’s preferred team.

Independent Variables. The robots leveraged **Efficient** and **Legible** allocations. For **Legible** the robots optimized Equation (5) to select the allocation θ_{Leg} that best communicated the human’s role. This is different from legibility when watching: instead of trying to make every agent’s subtask clear, now the robots are *only* trying to convey the participant’s subtask (i.e., which tennis ball the human should remove from the table). To better isolate **Efficient** and **Legible**, we selected **Efficient** allocations that *were not* legible.

Procedure. Participants completed the cleaning task with four pairs of robot teams (see Figure 5). Each pair contained one **Efficient** team and one **Legible** team. We never told the participants what type of team they were interacting with. To increase the number of data points, and to ensure that the position of the tennis balls did not affect our results, we placed the tennis balls in two configurations: *Uneven* and *Even*. In *Uneven* balls 2 and 3 were clustered on one side of the table, while in *Even* all of the tennis balls were equally spaced. During interaction the user needed to remove a ball from the table, but the user did not know *a priori* which ball they should remove — participants had to infer their assigned subtask from the actions of other agents. Participants sat next to the table and observed both robots’ trajectories ξ . Once the participant was confident they knew which tennis ball to pick up, they pressed a button to temporarily pause the robots. After the user reached in and grabbed their tennis ball, the robots completed the rest of the task autonomously.

Dependent Measures. When the human was interacting with a robot team we recorded the amount of *time* that user waited before intervening, and whether or not the user *predicted* their subtask correctly. After the participant finished interacting with a pair of teams (e.g., Team A and Team B) participants indicated which team they would prefer to keep collaborating with on a scale of 1-7. Here a score of 1 denotes that the human had a strong preference for Team A, while a score of 7 denotes that the human had a strong preference for Team B.

Hypotheses. We had two hypotheses:

H4. *Humans that play with legible robot teams will more quickly and accurately recognize their role within the team.*

H5. *When given the choice, humans will prefer to keep playing with legible robot teams.*

Results. Our results are summarized in Figure 5. First, we conducted a repeated measures ANOVA to confirm that the arrangement of the tennis balls *did not* have an effect on the human’s preferences ($F(1, 21) = .000, p = 1.00$). Next, we counted the number of times the participants correctly predicted their subtasks with **Efficient** and **Legible** teams: similar to our online study in Figure 2, participants made the correct prediction more frequently when working with **Legible** teams. Interestingly, the type of robot team did not have an effect on the time it took for humans to make these predictions ($F(1, 21) = 1.843, p = .189$). The combination of these results partially supports **H4** — legible task allocations led to more accurate human predictions, but not faster predictions.

Finally, we analyzed whether participants preferred to keep playing with **Efficient** or **Legible** teams. We conducted a one-way repeated measures ANOVA with a Sphericity Assumed correction and determined that the type of allocation (**Efficient** or **Legible**) had a significant main effect on the users’ preference: $F(1, 10) = 77.13, p < .001$. In all four pairs we found higher average scores for **Legible**, and in three out of the four pairs our post hoc t-tests showed that this difference was statistically significant ($p < .05$). Although users scored **Legible** higher in Pair 3, here the difference between **Efficient** and **Legible** was not statistically significant ($t(20) = -1.056, p = .303$). These results are consistent with **H5**, and suggest that people prefer to continue collaborating with teams that make their roles clear.

B. Fairness

Our results so far suggest that humans prefer to interact with robot teams that optimize for legible (but potentially inefficient) allocations. Next, we test how adding *fairness* into these allocations changes the human’s preferences.

Independent Variables. We compared **Legible+Fair** teams that optimize Equation (9) to **Legible** teams which follow the

same approach as in Section VI-A. To better separate these conditions we purposely selected **Legible** teams that were not fair. Recall that Equation (9) optimizes for the fairness f_H from the human’s perspective. This incentivizes the robots to select allocations where the human has an equal share of the work, and does not penalize the robots for dividing the remaining effort unevenly among themselves (e.g., a single robot may be asked to pick up multiple tennis balls under **Legible+Fair**).

Similar to our *watching* user study in Section V, we studied two different definitions of fairness in teams: *equality of allocation*, Equation (6), and *equality of effort*, Equation (7). Examples of the trajectories generated by these robot teams are shown in Figure 6. Under equality of effort the robots assigned the human to closest ball, and under equality of allocation the robots always performed two of the three subtasks.

Procedure. Participants collaborated with two pairs of teams to clear the table. In one pair the **Legible** team asked users to pick up the farthest ball while the **Legible+Fair** team asked participants to get the closest ball. In the other pair the **Legible** team asked participants to pick up two balls while the **Legible+Fair** team asked users to get one ball. After interacting with a pair of teams users indicated their preference on a 1-7 scale (just as in Section VI-A). We hypothesized that: **H6.** *Participants will prefer to join robot teams that legibly and fairly distribute the allocations or effort.*

Results. Our results in Figure 6 indicate that humans have a preference for fairness when they are playing with robot teams. Across both definitions of fairness, participants rated **Legible+Fair** teams significantly higher than **Legible** teams (one-way repeated measures ANOVA: $F(1, 10) = 40.06, p < .001$). We contrast these results to Figure 4, where the watching humans marginally preferred **Legible** teams in Pursuit-Evasion. Comparing these results, we suggest that fairness may be less of a factor when users are watching the team, but more decisive when humans are actually playing with the team.

VII. CONCLUSIONS

We developed an optimization framework that enables teams of robots to encourage human participation. Under our approach centralized robot teams treat humans as humans, and actively search for legible and fair ways to allocate subtasks among agents. As compared to a baseline that purely optimizes for efficiency, robots that leverage legible and fair allocations better encourage *watching* humans to join the team and *playing* humans to keep collaborating with the team.

REFERENCES

- [1] J. K. Gupta, M. Egorov, and M. Kochenderfer, “Cooperative multi-agent control using deep reinforcement learning,” in *Int. Conf. Autonomous Agents and Multiagent Systems*, 2017, pp. 66–83.
- [2] J. Foerster, F. Song, E. Hughes, N. Burch, I. Dunning, S. Whiteson, M. Botvinick, and M. Bowling, “Bayesian action decoder for deep multi-agent reinforcement learning,” in *Int. Conf. Machine Learning*, 2019.
- [3] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, “Multi-agent actor-critic for mixed cooperative-competitive environments,” in *Advances in Neural Information Processing Systems*, 2017.
- [4] P. Culbertson, J.-J. Slotine, and M. Schwager, “Decentralized adaptive control for collaborative manipulation of rigid bodies,” *IEEE Trans. Robotics*, 2021.
- [5] M. Carroll, R. Shah, M. K. Ho, T. Griffiths, S. Seshia, P. Abbeel, and A. Dragan, “On the utility of learning about humans for human-AI coordination,” *Advances in Neural Information Processing Systems*, pp. 5174–5185, 2019.
- [6] M. Li, M. Kwon, and D. Sadigh, “Influencing leading and following in human–robot teams,” *Autonomous Robots*, vol. 45, pp. 959–978, 2021.
- [7] S. A. Wu, R. E. Wang, J. A. Evans, J. B. Tenenbaum, D. C. Parkes, and M. Kleiman-Weiner, “Too many cooks: Bayesian inference for coordinating multi-agent collaboration,” *Topics in Cognitive Science*, vol. 13, no. 2, pp. 414–432, 2021.
- [8] P. Stone, G. A. Kaminka, S. Kraus, and J. S. Rosenschein, “Ad hoc autonomous agent teams: Collaboration without pre-coordination,” in *AAAI*, 2010.
- [9] S. Sebo, B. Stoll, B. Scassellati, and M. F. Jung, “Robots in groups and teams: A literature review,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 4, pp. 1–36, 2020.
- [10] A. E. Clark and A. J. Oswald, “Comparison-concave utility and following behaviour in social and economic settings,” *Journal of Public Economics*, vol. 70, no. 1, pp. 133–155, 1998.
- [11] E. Fehr and K. M. Schmidt, “A theory of fairness, competition, and cooperation,” *The Quarterly Journal of Economics*, vol. 114, no. 3, pp. 817–868, 1999.
- [12] S. Eshghi, G.-R. Williams, G. B. Colombo, L. D. Turner, D. G. Rand, R. M. Whitaker, and L. Tassioulas, “Mathematical models for social group behavior,” in *IEEE SmartWorld*, 2017, pp. 1–6.
- [13] S. Parekh, S. Habibian, and D. P. Losey, “RILI: Robustly influencing latent intent,” *arXiv preprint arXiv:2203.12705*, 2022.
- [14] L. Johannsmeier and S. Haddadin, “A hierarchical human-robot interaction-planning framework for task allocation in collaborative industrial assembly processes,” *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 41–48, 2016.
- [15] S. M. Rahman and Y. Wang, “Mutual trust-based subtask allocation for human–robot collaboration in flexible lightweight assembly in manufacturing,” *Mechatronics*, vol. 54, pp. 94–109, 2018.
- [16] M. C. Gombolay, R. A. Gutierrez, S. G. Clarke, G. F. Sturla, and J. A. Shah, “Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams,” *Autonomous Robots*, vol. 39, no. 3, pp. 293–312, 2015.
- [17] C. Liu, J. B. Hamrick, J. F. Fisac, A. D. Dragan, J. K. Hedrick, S. S. Sastry, and T. L. Griffiths, “Goal inference improves objective and perceived performance in human-robot collaboration,” in *Int. Conf. Autonomous Agents and Multiagent Systems*, 2016.
- [18] S. Habibian, A. Jonnavittula, and D. P. Losey, “Here’s what I’ve learned: Asking questions that reveal reward learning,” *ACM Transactions on Human-Robot Interaction*, 2022.
- [19] T. Hellström and S. Bensch, “Understandable robots-what, why, and how,” *Journal of Behavioral Robotics*, vol. 9, pp. 110–123, 2018.
- [20] A. Roncone, O. Mangin, and B. Scassellati, “Transparent role assignment and task allocation in human robot collaboration,” in *IEEE Int. Conf. Robotics and Automation*, 2017, pp. 1014–1021.
- [21] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, “Legibility and predictability of robot motion,” in *ACM/IEEE Int. Conf. Human-Robot Interaction*, 2013, pp. 301–308.
- [22] M. R. Barrick, G. L. Stewart, M. J. Neubert, and M. K. Mount, “Relating member ability and personality to work-team processes and team effectiveness,” *Journal of Applied Psychology*, p. 377, 1998.
- [23] M. Rabin, “Incorporating fairness into game theory and economics,” *The American Economic Review*, pp. 1281–1302, 1993.
- [24] H. Claire, Y. Chen, J. Modi, M. Jung, and S. Nikolaidis, “Multi-armed bandits with fairness constraints for distributing resources to human teammates,” in *ACM/IEEE Int. Conf. Human-Robot Interaction*, 2020.
- [25] M. L. Chang, Z. Pope, E. S. Short, and A. L. Thomaz, “Defining fairness in human-robot teams,” in *IEEE Int. Conf. Robot & Human Interactive Communication*, 2020, pp. 1251–1258.
- [26] M. L. Chang, G. Trafton, J. M. McCurry, and A. L. Thomaz, “Unfair! Perceptions of fairness in human-robot teams,” in *IEEE Int. Conf. Robot & Human Interactive Communication*, 2021, pp. 905–912.
- [27] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, “Maximum entropy inverse reinforcement learning,” in *AAAI*, 2008.
- [28] H. J. Jeon, S. Milli, and A. D. Dragan, “Reward-rational (implicit) choice: A unifying formalism for reward learning,” in *Advances in Neural Information Processing Systems*, 2020.
- [29] R. D. Luce, *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.