

Generating Pattern-Based Conventions for Predictable Planning in Human-Robot Collaboration

CLARE LOHRMANN, University of Colorado Boulder, USA

MARIA STULL, University of Colorado Boulder, USA

ALESSANDRO RONCONE, University of Colorado Boulder, USA

BRADLEY HAYES, University of Colorado Boulder, USA

For humans to effectively work with robots, they must be able to predict the actions and behaviors of their robot teammates rather than merely react to them. While there are existing techniques enabling robots to adapt to human behavior, there is a demonstrated need for methods that explicitly improve humans' ability to understand and predict robot behavior at multi-task timescales. In this work, we propose a method leveraging the innate human propensity for pattern recognition in order to improve team dynamics in human-robot teams and to make robots more predictable to the humans that work with them. Patterns are a cognitive tool that humans use and rely on often, and the human brain is in many ways primed for pattern recognition and usage. We propose Pattern-Aware Convention-setting for Teaming (PACT), an entropy-based algorithm that identifies and imposes appropriate patterns over a robot's planner or policy over long time horizons. These patterns are autonomously generated and chosen via an algorithmic process that considers human-perceptible features and characteristics derived from the tasks to be completed, and as such, produces behavior that is easier for humans to identify and predict. Our evaluation shows that PACT contributes to significant improvements in team dynamics and teammate perceptions of the robot, as compared to robots that utilize traditionally 'optimal' plans and robots utilizing unoptimized patterns.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing**; **Empirical studies in collaborative and social computing**;

Additional Key Words and Phrases: collaboration, conventions, patterns, predictability

ACM Reference Format:

Clare Lohrmann, Maria Stull, Alessandro Roncone, and Bradley Hayes. 2023. Generating Pattern-Based Conventions for Predictable Planning in Human-Robot Collaboration. *ACM Trans. Hum.-Robot Interact.* 0, 0, Article 0 (2023), 23 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Human difficulty with accurately modeling and predicting robot behaviors prevents the integration of robots into human-populated environments. Prior work indicates that the more effectively a human can model their robot teammate, the better the team will be able to perform [32]. However, humans struggle to build accurate and effective models of robots [5, 24] and often find them

Authors' addresses: [Clare Lohrmann](mailto:clare.lohrmann@colorado.edu), clare.lohrmann@colorado.edu, University of Colorado Boulder, 1111 Engineering Dr, Boulder, Colorado, USA, 80309; [Maria Stull](mailto:maria.stull@colorado.edu), maria.stull@colorado.edu, University of Colorado Boulder, 1111 Engineering Dr, Boulder, Colorado, USA, 80309; [Alessandro Roncone](mailto:alessandro.roncone@colorado.edu), alessandro.roncone@colorado.edu, University of Colorado Boulder, 1111 Engineering Dr, Boulder, Colorado, USA, 80309; [Bradley Hayes](mailto:bradley.hayes@colorado.edu), bradley.hayes@colorado.edu, University of Colorado Boulder, 1111 Engineering Dr, Boulder, Colorado, USA, 80309.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

2573-9522/2023/0-ART0 \$15.00

<https://doi.org/XXXXXXXX.XXXXXXX>




Fig. 1. A participant plays a collaborative block-selection game with a robot. By using PACT to augment its planner, the robot's actions are more predictable to the participant over multiple episodes and multi-task time horizons, and the robot is viewed as a better teammate.

unpredictable even in very simple environments [4]. This limits team performance, as humans prefer to work with agents they find predictable and trust unpredictable agents less [9].

As humans struggle to predict agent behavior, agents are simultaneously attempting to predict and adapt to humans. Prior work has been done to improve an agent's ability to predict human actions [11, 14, 27, 37] as well as adapt to human behaviors [7, 17, 40]. However, for collaboration to succeed, both human and agent need to be mutually predictable, and there are significant technical gaps in improving humans' ability to predict agents' actions.

In contrast to human-robot teams, human-human teams are extremely skilled at collaborative tasks where synchronization, coordination, and prediction of each other's behavior is necessary—such as assembling a bookshelf or making a football pass. Part of this gap in performance can be explained by the distinct sets of tools that humans and robots each use to accomplish tasks. Humans do not often rely on optimization as a cognitive tool, and instead use heuristics and pattern recognition [12]. Notably, humans have difficulty in predicting what the robot will do next as the robot is not using cognitive tools the human is familiar with.

The cognitive processes humans rely on, while simpler than optimization methods, are more adaptable than and often outperform such techniques, especially in complex environments where optimization is computationally intractable [3, 12, 15, 31]. One of the cognitive tools that humans employ is the ability to identify and process patterns, such as recognizing rhythm in a song or the color order of changing stoplights. Pattern recognition has developed via our evolution as a species and is facilitated by specific structures in our brains [30]. Even preschoolers are capable of duplicating, extending, and abstracting patterns to new environments [35]. These are deeply ingrained cognitive processes that humans are adept at using.

Within the context of teaming, humans extensively rely on conventions in order to effectively coordinate behavior. Conventions are a form of shared knowledge [39] that teammates can use in collaborative tasks to synchronize their actions. Patterns—a predictable sequence [35] of actions—can make conventions easier to learn and follow. For a pattern-based convention to be meaningful to a human, the pattern must be human-perceptible, i.e., based on features that a human can observe. Using pattern-based conventions leverages innate human pattern-processing abilities, making the

99 conventions intuitive for people to identify [30] and predict. In order to facilitate human-robot
100 collaboration, in this work we propose a filtering algorithm which enables an embodied agent to
101 set conventions for the team by using a human-perceptible pattern to restrict its actions in a given
102 situation to a more predictable set. We refer to this algorithm as Pattern-Aware Convention-setting
103 for Teaming (PACT).

104 PACT can select patterns of varying complexity depending on the context, and is robust to
105 changes in the environment. A pattern selected by PACT can continue to be used without alteration
106 even if the task space changes over time. **Our results show that by using PACT, not only does
107 the robot become more predictable to its human teammates, but team performance as
108 well as perceptions of the robot improve.** These findings are supportive of the hypothesis
109 that, by leaning into familiar cognitive processes, humans can more readily identify the robot's
110 intentions, understand how the robot is making decisions, and abstract the robot's behavior into
111 new environments accurately. Not only is PACT effective, it demonstrates the benefit in adapting
112 human cognitive strengths for robots to use during collaboration.

113 2 BACKGROUND AND RELATED WORK

114 When humans collaborate, we build mental models of our teammates [44]. Mental models are
115 knowledge structures that help people to describe, explain, and predict events in our environment
116 [41]. Human teams are so astute, they can even create shared mental models for the team, creating
117 shared knowledge and expectations that lead to greater success [29]. Evidence shows that humans
118 also build mental models of robots [41]. In line with human factors research, work within human-
119 robot interaction indicates that when humans and robots can build accurate mental models of each
120 other, human-robot collaboration is more likely to be successful [21, 29, 32, 45]. Humans also trust
121 agents more when we find them predictable [9]. With traditional controllers, however, there are
122 no guarantees of pattern or regularity, so humans' mental models of robot teammates are often
123 incorrect or incomplete [5, 24].

124 Humans reason in a fundamentally different, and often contradictory way to our robot teammates.
125 Artificially intelligent agents, embodied or otherwise, are built to optimize, but humans do not
126 optimize when we plan or make decisions [12]. We satisfice—meaning we find a “good enough”
127 solution to the problem [3, 12, 28]. People employ a variety of cognitive tools to do this, from using
128 heuristics to pattern recognition [1, 31, 43]. Satisficing is not a weakness of human cognition; to
129 the contrary, heuristic usage approaches rationality over the long term, and our brains developed
130 it to navigate our environment, where optimization is computationally intractable [2, 3, 15, 43].
131 In human-robot teams, robot teammates are working on identifying and achieving the optimal
132 solution for a specific set of parameters, whereas human teammates are agreeing upon a “good
133 enough” solution, and these solutions are rarely the same.

134 Much of the recent work in human-robot collaboration focuses exclusively on improving the
135 performance of the robotic agent. Works that attempt to predict human actions or their path
136 directly have seen success within the environments they tested in [11, 14, 37]. There has also
137 been a significant effort to adapt successful methods in competitive environments to collaborative
138 environments [7, 20, 22, 42], though this is very difficult. Approaches that are highly effective in
139 competitive environments are challenging to adapt to collaborative environments [19]. What makes
140 many self-play approaches successful—a policy that is convoluted and difficult for opponents to
141 counter—is a drawback in collaborative settings. What results is a large drop in performance when
142 trained agents are tested with humans rather than other agents [22, 40]. These approaches also do
143 not capture the full scope of human collaboration within their environments [19].

144 There is strong evidence in the literature that using human cognitive tools within human-robot
145 and human-agent collaboration can be highly effective [23]. Work has shown that human partners
146
147

can learn conventions developed by artificially intelligent agents [39]. Significant work has also been done to integrate social conventions into collaborative agents [25, 40]. When robots explicitly adhere to human navigation conventions, humans find them more predictable and likeable, and the robot’s ability to navigate is not compromised [34, 36]. Further, having people rely on conventions that they create themselves [8] or are already familiar with, such as “pinch” and “pull” motions that people use on their smartphones [13], leads to improvement in their ability to collaborate with a robot.

Failing to account for human cognitive tendencies may obscure results, and limit future work [5]. One such cognitive tendency is pattern recognition. The human brain floods with dopamine upon recognizing a pattern, thus, humans are strongly incentivized to find them [26]. Some scientists even consider pattern recognition and reasoning to be a cornerstone of higher intelligence [30]. As the human brain is wired for pattern recognition, actions that are pattern-based are more likely to be recognizable to human participants. By using PACT, we can select patterns that are as recognizable as possible.

3 A FRAMEWORK FOR PATTERNS-BASED CONVENTIONS

In this section we detail PACT, an entropy-based algorithm to select the most appropriate pattern to use in a given environment. The central cognitive science concept that underpins this approach is the human tendency toward pattern recognition and usage. By playing into known strengths of human cognition, the robot’s behavior becomes more recognizable, predictable, and understandable to human teammates.

3.1 Definitions

PACT takes the tuple $\{T, F, r\}$ as input to determine the ideal pattern for a particular task space, where:

- $T = \{t_1, t_2, \dots, t_n\}$ is the finite set of subtasks an agent must complete. T is unordered, but subtasks within T may have ordering constraints imposed by prerequisites.
- $F = \{f_1, f_2, \dots, f_m\}$ is a set of functions that map from a subtask to a feature of that subtask. (e.g., $f_1(t) \rightarrow \text{“circle”}$, $f_2(t) \rightarrow \text{“red”}$)
 - $f_i = [v_1, v_2, \dots, v_k]$ is a feature vector representing a characteristic (e.g., for a feature “color” there may be categorical values {“red”, “green”, “blue”} encoded as a one-hot vector. A color feature could also be represented as a continuous three-dimensional vector of RGB values).
- A **Rule** is a function that sorts subtasks in T using a comparator function over output from one or more features in F .
- r is the maximum number of Rules that PACT is allowed to combine to form a *Pattern*, a hyperparameter selected by the user prior to Pattern formation.
- A **Pattern** is an ordered sequence of between 1 and r Rules that augments available subtasks in T for a planner to select from in a given state. Rules are applied sequentially to filter out or augment the cost of elements in T to inform plan generation.

3.2 Rule Formation and Application

A Rule is a data structure that contains a sorting function and a set of features to apply it to. Given a set of subtasks, a Rule filters it down to a subset of subtasks that the agent can perform (while still being consistent with the Rule). For example Figure 2 shows an environment in which an autonomous drone must check critical infrastructure after a natural disaster. Communications are down, so the drone is unable to communicate reliably with a human ground crew, making



Fig. 2. In this illustration of the PACT algorithm, we use a scenario in which a natural disaster has occurred in a coastal town. Critical infrastructure must be checked for damage, and an autonomous drone as well as a human team on the ground are tasked with damage assessment. In this time-sensitive task, communication between the drone and humans is limited. Each location indicated on the map has features used by PACT: whether the location contains humans that sheltered in place (red), the type of infrastructure (blue), and the likelihood the location is flooded (green). This scenario does not consider the distance traveled by the drone to be a constraint, but such constraints can easily be added to guide PACT's pattern-selection.

the predictability of the robot critical. Here, the set of subtasks T is the set of locations the drone must check and document. Each location has three features: the estimated flooding risk (which we discretize into low, medium, and high risk), the type of critical infrastructure (police station, power substation, water treatment plant, and hospital), and whether or not human staff sheltered in place there. A Rule based on the presence of humans could be ["no humans", "humans"], such that the robot would visit all places without humans sheltering, followed by those locations with humans. Another Rule based on the flood risk could be ["high", "medium", "low"]. Applying the flood risk Rule ["high", "medium", "low"] to the locations in T would result in filtering the locations down to the subset of locations with "high" flood risk. Each location in this subset would be visited by the drone. Then, with no more "high" flood risk locations, the drone would visit all "medium" flood risk locations, and so on. For this Feature, because each location has one of three possible values, there are $3!$ possible orderings, meaning this Feature (flood risk) has $3!$ possible Rules that could leverage it. Thus, given a set of n categorical Features, where each Feature i has k_i possible values, there are at most $\prod_{i=1}^n k_i!$ single-Feature Rules that can be generated. For Rules over features with non-categorical values, the space of orderings is technically infinite and depends entirely on how complex the comparator function encoded in the Rule is, but by imposing a restriction to sort values in either an ascending or descending order we may assume two Rules per continuous feature. In the drone example given by Figures 2 and 3 there are 32 possible Rules that can be used. Without loss of generality, in this work, each Rule is generated from a single feature.

3.3 Pattern Formation and Application

The hyperparameter r is set by the user prior to the generation of Patterns in order to determine the maximum allowable complexity for the Patterns. r can be at most equal to the number of Features

Features	Rules	Patterns
	R1: then	P1: [R1]
	R2: then	...
	R3: then then then	P33: [R1, R3] P34: [R3, R1]
...	R26: then then then	...
...	...	P462: [R32, R2, R15]
	R32: then then	...

Fig. 3. In this example, we describe the formation of Rules and Patterns from the scenario in Figure 2. The left column shows the three features of each location we are using (human presence, infrastructure type, and flood risk), and the possible values for each feature. The center column shows the ways Rules can be constructed from features by imposing an ordering on the possible values of a feature. The right column shows how a Pattern is constructed by applying one or more Rules. Note that Patterns may not have conflicting Rules; we choose at most one rule per feature.

and must be at least one. With a larger r , Patterns can be more complex and are thus more likely to be able to impose a fully deterministic ordering of tasks in a plan, but this increase in complexity may also make the Pattern too difficult for a human partner to identify and follow.

A Pattern is a data structure that contains a sequence of between 1 and r Rules. Given a set of subtasks, the Pattern determines the subset of next possible subtasks. The initial set of subtasks is passed to the first Rule in the sequence, which returns the subset of allowable subtasks. This subset is passed to the second Rule in the sequence, continuing through the full sequence of Rules to obtain the final subset of possible subtasks for the given Pattern. Figure 4 illustrates the Pattern [“low”, “medium”, “high”], [“no humans”, “humans”]. First, the set of locations the drone must visit is filtered down according to the first Rule (flood risk), leaving just the locations with “low” flood risk. This subset of locations is then passed on to the second Rule (human presence) to be filtered down to locations with “no humans”. The drone will have to visit all locations in this subset (B and G) in any order before moving on to locations with different values for these features. After these locations, pictured in the first box of Figure 4, are checked, the remaining locations are passed to the first Rule and then the second to obtain the subset of locations that are “low” risk with “humans” (location D). After this location, the Pattern filters down to an empty set, as there are no locations with “medium” flood risk and “no humans”. The Pattern will then identify locations with “medium” flood risk and “humans”, which will be visited before locations with “high” risk and “no humans” (F) and then locations with “high” risk and “humans” (E).

3.4 Pattern Trees

Calculating a score to evaluate the effectiveness of a given Pattern requires evaluating the possible orderings of subtasks that it imposes throughout the plan it generates (or a sampled subset if otherwise infeasible). To efficiently compute and organize this for each Pattern, we construct a Pattern Tree. An illustration of a portion of a Pattern Tree generated from the drone example can

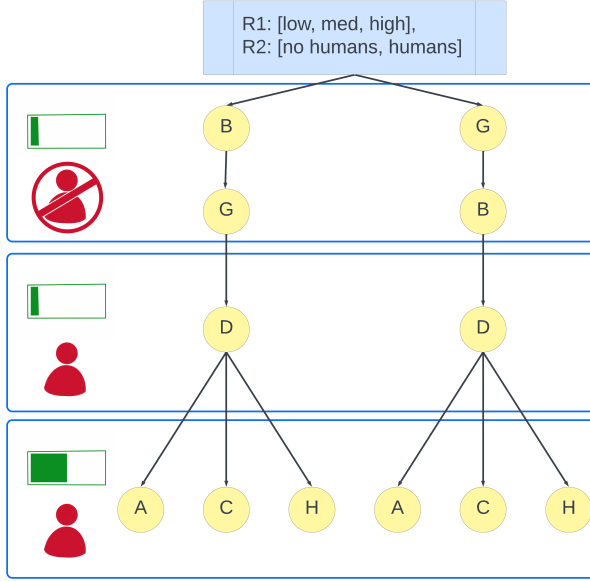


Fig. 4. For each Pattern, a Pattern Tree is constructed to easily identify all allowable orderings of subtasks. In this Pattern, the first subtasks are those locations that are low flood risk and have no humans sheltering in place (B,G). The second allowable subtasks are the remaining low risk, no human locations, which are appended to the tree. All allowable orderings of length 2 can thus be obtained by traversing the tree to depth 2. For the third subtask, there are no remaining low risk locations with no humans, so the low risk locations with humans are selected (D). There are no medium flood risk locations that do not have humans, so the fourth possible subtasks are those that are of medium risk with humans. The tree is constructed in this manner until it reaches a depth equal to the number of subtasks.

be seen in Figure 4. The first level of the tree is determined by the possible first subtasks given the Pattern and T . For each subsequent level, the children of a node are determined by assuming the path from root to parent node specifies the sequence being followed, and applying the Pattern to the remaining subtasks. The final tree will have $|T|$ levels, as the entire sequence will be generated. Thus, traversing to the i th level of the tree will reveal all possible subsequences of length i for a given Pattern. This simplifies Pattern evaluation calculations as matching subsequences of length $i - 1$ can be obtained for all Patterns quickly, and all possible i th subtasks can also be obtained by indexing the children of all nodes in the $(i - 1)$ th level of the tree. For tasks with prohibitively large amounts of subtasks, Monte Carlo methods can be applied to approximate the Pattern Tree.

3.5 Pattern Scoring Metric

To determine the most appropriate Pattern for a given T , we propose a scoring metric that can be applied to a set of possible Patterns (which we refer to as the *Pattern Bank*). Patterns with lower scores are more preferable. We define this score (λ) for a given Pattern (p) as:

$$\lambda_p = \prod_{i=1}^{|T|} \mathcal{H}(T_{i,p}) + \frac{|P_{i,shared}| - 1}{|P|} * \mathcal{H}(T_{i,shared}) \quad (1)$$

Where:

- p is the Pattern for which the score is being calculated.

- T is the set of subtasks the agent must perform.
- $\mathcal{H}(x)$ is an entropy calculation for the collection x .
- $T_{i,p}$ is the collection of all possible subtasks at the i th step of planning given the Pattern p .

This can be extended over sets of Patterns as follows:

- P is the set of possible Patterns given $\{F, r\}$.
- $P_{i,\text{shared}}$ is the set of Patterns that share at least one possible sequence of length $i - 1$ with the given Pattern p .
- $T_{i,\text{shared}}$ is the collection of all possible subtasks at the i th step for all Patterns in $P_{i,\text{shared}}$.
- All sequences of length $i - 1$ are allowable.

This scoring metric allows for selection of a pattern that is both as deterministic as possible (first term) as well as unique (second term). Favorable patterns are those that become unique in their possible sequences as soon as possible (easier to identify/ legible [10]), while also being as deterministic as possible (easier to follow).

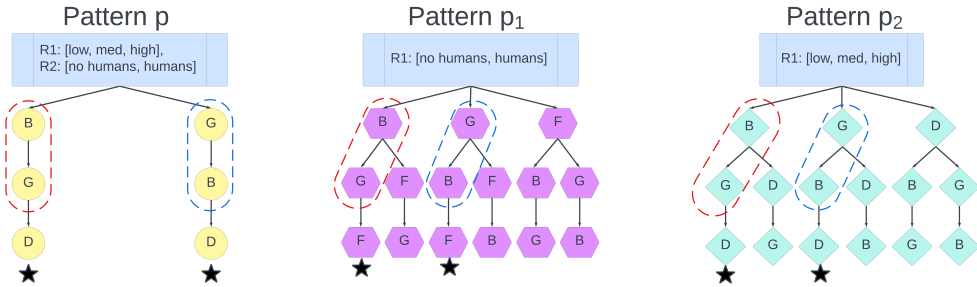


Fig. 5. This figure illustrates how the second term of the score in Eq. 1 is calculated for a given Pattern (the same Pattern shown in Figure 4) when $i = 3$. First, the possible orderings of length $i - 1$ are identified for the given Pattern, seen in the left tree. There are two possible subtask orderings of length $i - 1$, highlighted in red and blue. However, these orderings are not unique to this Pattern. There may be other Patterns in the Pattern Bank that share these orderings of length $i - 1$. Two such Patterns are shown here, with the matching orderings circled. If a human partner observes the robot going to B then G, they cannot distinguish between the Pattern the robot is following and these other Patterns. The (starred) children of these shared orderings are extracted from all Patterns in the Pattern Bank, and the entropy over this group is calculated. For this group of three trees, the group would be (D, D, F, F, D, D).

3.6 PACT

While the algorithm can be viewed in its entirety in pseudocode within Algorithm 1 in the appendix, we provide an intuitive walkthrough here for ease of understanding. Prior to applying PACT, we create a Pattern Tree for each Pattern in the Pattern Bank. We then initialize data structures that keep track of the best patterns and their scores. For each Pattern (p) in the Pattern Bank we calculate a pattern score, weighing how deterministic the pattern is and how much overlap in resulting plans there is with those generated by other Patterns in the Pattern Bank (i.e. how unique the pattern is compared to others). A pattern score is a summation of subscores calculated for the selection of each subtask in the sequence imposed by the Pattern. Scores start at zero and increase at each step. The subscores have terms related to entropy (i.e., how deterministic the plan imposed by the Pattern is) and uniqueness (to bias against Patterns that generate plans that can be explained by other Patterns).

393 The first term of Equation 1 is the entropy over the distribution of i th possible subtasks when a
394 sequence of subtasks is being constructed using the given Pattern p . When using p to order the
395 subtasks, the allowable sequences can be determined by traversing the tree. Thus, the subtasks
396 that could be i th in a sequence that conforms to p are all those nodes at a depth of i in the Pattern
397 Tree. In Figure 4, when $i = 3$, the nodes used for this calculation are in the middle box (depth of 3).
398 When $i = 4$, the nodes used for this calculation are those in the bottom box. The first term for i is
399 the calculated entropy for the set.

400 Figure 5 illustrates the calculation for the second term when $i = 3$. This calculates how unique p
401 is, i.e. how much overlap there is between p and other Patterns in the Pattern Bank. The term is
402 composed of an entropy value and a discount.

403 When determining how unique an ordering induced by p is, the possible orderings of subtasks
404 must be compared with those of other possible Patterns. In Figure 5, p is the Pattern on the far left,
405 with circular nodes. When $i = 3$, there are only two possible subtask orderings of length 2 that
406 follow the Pattern, circled in red and blue dashed lines. If the robot is using p to order its subtasks,
407 a human partner will observe one of the circled orderings. However, these orderings may also
408 comply with other Patterns within the Pattern Bank. The trees with hexagonal and diamond nodes
409 in Figure 5 are other Patterns in the Pattern Bank which have some orderings of length $i - 1$ (2) in
410 common with the target Pattern p —also circled with dashed lines. If the robot goes to location B
411 then on to G, this behavior can be explained by p , but also by these other Patterns, which may lead
412 to the human partner to believe the robot is following a Pattern other than p leading to confusion
413 or difficulty predicting the robot in the future, as their mental model of the robot is incorrect. The
414 second term identifies the children of these shared sequences, marked in the figure with stars,
415 and calculates the entropy over them. Thus, Patterns that produce sequences of subtasks that are
416 unique have lower scores, and Patterns that produce orderings of subtasks that are shared across
417 many Patterns have higher (worse) scores.

418 This entropy calculation is then discounted by the proportion of the Pattern Bank that has an
419 ordering of length $i - 1$ in common with p . This is done to penalize Patterns that could be mistaken
420 for a greater number of other Patterns. If p shares many orderings of length $i - 1$ with one other
421 Pattern, this will lead to less confusion on the part of a human partner than if p shares a few
422 orderings of length $i - 1$ with many Patterns in the Pattern Bank.

423 When all of the subscores have been calculated and summed, we compare the total score for
424 the Pattern to the minimum score, and store all minimum-scoring Patterns. When we have scored
425 all Patterns, we return every minimum-scoring Pattern for subsequent selection and use by the
426 planner.

427 Pattern scoring and selection is performed offline, done before the robot engages with an
428 environment. While the Pattern can be updated or changed, a Pattern deemed to be the most
429 suitable for a target set of environments can and should continue to be used in other environments
430 the robot acts in to maximize predictability, as long as the features used in the Pattern remain
431 present in these other deployment environments. Changes made to the Pattern during interaction
432 with humans may make the robot less predictable, and this work promotes the use of one Pattern
433 kept consistent even when the robot finds itself in previously unseen deployment environments.

434 4 EXPERIMENTAL EVALUATION

436 PACT can be applied to any planning problem for which the overarching task can be decomposed
437 into a predefined set of subtasks or goals (e.g., search a set of ten locations for survivors).

438 PACT can be applied to situations where the robot is working with one or more humans, such as
439 remote sample recovery, wherein PACT would make it easier to predict where the robot would be
440 retrieving samples from, allowing humans to parallelize efforts by focusing on areas that the robot
441

is not or to assist robots by traveling to their next destination without explicit communication requirements. PACT may also be used in scenarios where human and robot are simply sharing a workspace, where increased predictability of which object the robot will grab next allows humans to more easily navigate around or more safely work with the robot.

However, in these scenarios as well as in other more complex environments, there is a significant amount of extraneous side-channel information that people may use to predict the robot's behavior. People may wait for several moments to determine where the robot is headed next, take time to simply observe the robot, or even be provided with information from the robot itself. In order to effectively test PACT, and to show that the planner's order of subtasks alone is driving increased predictability, all of this information must be removed. Any effective testbed for such a system must be framed as a coordination problem, so that the human does not have the opportunity to observe the robot without taking any action themselves. The coordination task is structured such that only by accurately predicting the robot's actions can the team succeed, and there is no other information the human can rely on other than previous robot actions and their own mental model of the robot.

PACT can be applied to a broad range of planning problems that can be constructed from this maximally constrained coordination problem, by relaxing the testbed's requirements of forced simultaneous action selection or inability to wait and observe the robot. While this makes the task of coordinating with and predicting the behavior of the robot significantly more difficult, it allows for a stronger assessment of the effectiveness of PACT than would occur in a more realistic collaborative scenario with additional side-channel information available.

Thus, we evaluate the efficacy of PACT through a collaborative game involving a human and robot that rewards teams whose members' task selections are predictable.

4.1 Game Environment

The collaborative game is played on a five by five grid on a table in a shared workspace (Figure 6). At the beginning of a round, nine blocks are placed in unique locations on the grid. Every block is assigned a unique numerical value between one and nine, which is neither known nor observable by the human participant and is used to calculate the score for a successful move, representing the reward function that a traditional robot planner would attempt to maximize. At the start of each turn, the participant and Sawyer, a 7-degree-of-freedom robotic arm, both select a block without seeing the choice of their teammate. Participants make their selection on a tablet, and are allowed to select any block type (e.g., "blue triangle"). When both teammates have made their selection, Sawyer reveals its selection on its screen, and the participant receives an update on the tablet showing both players' selections as well as score information. If the team members each select blocks with fully matching visual features (e.g., both with yellow circles on them), Sawyer removes one block that matches those features from the grid. The team receives a positive reward based upon the numerical values of the blocks remaining and the number of blocks the team had to choose from; as the number of blocks on the board decreases (and it is more likely that teammates could coordinate by chance), the reward decreases. The game is scored as follows:

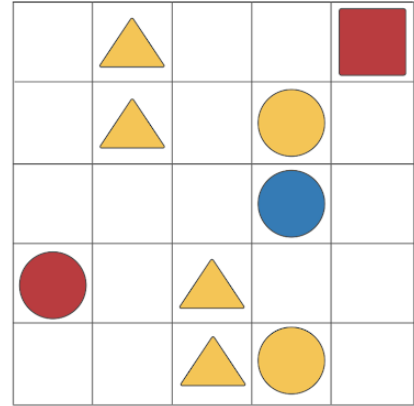


Fig. 6. The layout for the collaborative game. Nine blocks, each with a color, shape, and reward value are placed on the grid. Only the robot has knowledge of the rewards. Both players secretly select a block by color and shape, and if they coordinate, the robot removes a block from the grid.

$$S(t) = \begin{cases} -10t/(n+1) & n \text{ matching features} \\ B_{sum} + 5(|B_{rem}| + 1) & \text{all features match} \end{cases}$$

where:

- t = current turn number
- B_{rem} = set of blocks remaining on the board
- $B_{sum} = \sum_{i=0}^{|B_{rem}|-1} B_{rem}[i].numeric_value$

If the team does not agree on the same type of block, a penalty is assessed to the team. The size of the penalty increases as the game progresses, so an inability to coordinate early in the game is not penalized as harshly as failing to coordinate on the last few blocks. If the team is able to coordinate on a subset of features, i.e., both players select the same color or shape (but not both), the penalty assessed is reduced; teams that are able to coordinate along some axes are not penalized as much as teams that cannot coordinate on any features. Teams must coordinate to remove all nine of the blocks from the grid to complete a round.

It is important to note that the sequence of blocks that the robot will select is determined prior to the first turn: **this experimental setup is designed to test human understanding and prediction of robot behavior, not robot adaptation to human behavior**. Regardless of what the participant selects, the robot will always select the next block in the predetermined sequence.

This means that the robot will continue to select the same block until the participant matches the robot's selection. Upon the completion of each round, a new set of nine blocks is placed in the workspace in a new configuration. The numerical value of each block is also new, with no relationship between the value and any of the human-visible features. In other words, the reward function changes with each episode and is never shown or explained to the human teammate. This design decision illustrates the trade-offs between capability (reward maximization) and predictability (pattern adherence) when coordinating as a human-robot team.

4.2 Applying PACT to the Coordination Domain

The variables required to utilize PACT are defined as follows:

- T = A set of subtasks, one for each of the nine blocks in the workspace.
- F = A set of functions mapping each block subtask (by unique id) to values of features of the blocks— “color”: {“blue”, “red”, “yellow”}, “shape”: {“circle”, “square”, “triangle”}, “position”: {“row”:{1, 2, ..., 5}, “column”: {1, 2, ..., 5}}
- $r = 2$, such that only up to two Rules may be ordered to create a Pattern

The *position* Rules ordered the values in either ascending or descending order, operating over either only a single field (either “row” or “column”) or both (e.g., rows descending and columns ascending). In our environment—because the human and robot select blocks by color and shape on the interface (e.g., a “yellow triangle block”, without specifying location)—when calculating entropy over subsequent tasks to select, all tasks involving blocks of the same shape and color are treated equally regardless of position.

4.3 Experimental Design

Study participants ($n = 28$) were assigned randomly into one of three conditions in a between-subjects design:

- Reward-Maximizing: The participant works with a robot that selects blocks in the order that will maximize the team score in the event of perfect coordination, analogous to traditional reward optimization approaches.
- PACT Pattern: Participants are on a team with a robot that selects blocks following a pattern-based convention, generated and selected by PACT such that the pattern score is best for the set of tasks to be completed in the first round environment.
- Median Pattern: Participants work with a robot that selects blocks by following a pattern that achieved a median score when compared against all possible patterns in the first round environment.

Patterns selected for the PACT and Median groups are based on the first round environment and remain the same for all subsequent rounds of gameplay, despite environment changes. This allows us to evaluate team performance in both a ‘target’ environment that may be known and optimized against in advance (first round) and in new environments not explicitly optimized for (subsequent rounds).

4.4 Study Protocol

Consent was obtained from all participants, preceded by a brief check of participants’ ability to distinguish between the block colors. One participant self-identified as colorblind, though not a form of colorblindness that would prevent them from distinguishing between the colors used. Participants were then given a randomly generated six-digit identifier to link their survey responses, and were randomly assigned to an experimental group. Following this, participants filled out a pre-experiment survey about their experience with robots, attitudes about robots, and initial sentiments toward Sawyer, the robot used in the experiment. Experimenters then explained the collaborative game, answered questions, and participants began gameplay. After each round of the game, participants answered questions about their cognitive fatigue, ability to predict the robot’s behavior, and confidence in their team. After three rounds of the game, a third type of survey was administered. Participants were shown five novel game set ups, and were asked to identify which color and shape block the robot would select first and last in each given game. Participants were also given the option to mark that they were uncertain about either feature. Finally, a post-experiment survey was conducted, again surveying participants about their sentiments about the robot, their game comprehension, as well as questions about the team dynamics and performance of each team member. Following the completion of the survey, participants participated in a brief unstructured interview and debrief. The duration of the experiment was approximately sixty minutes.

4.5 Measurement

28 participants were recruited from the student community of our university for the IRB-approved human subjects study. Pre-experiment survey questions were taken from NARS, RoSAS, and previous HRI work [6, 33, 38]. Between rounds, participants answered selected questions from the NASA Task Load Index [16] to measure their cognitive fatigue and frustration, as well as several questions about their confidence in their choices. A “Round 4” survey consisting of five novel game setups was created specifically for this experiment in order to measure participants’ ability to abstract the robot’s behavior into a new environment. The post-experiment survey consisted of questions from RoSAS, identical to those asked in the pre-experiment survey, survey questions about the fluency of the team [18], as well as custom questions adapted from the team fluency questions.

4.6 Hypotheses

We conducted an ethics board approved human-subjects study to investigate the following hypotheses regarding the effectiveness of PACT within a human-robot collaborative coordination task:

- H_1 : Participants who work with the robot using PACT will have a more positive attitude about the dynamics of the team (i.e., coordination, mutual understanding, teamwork, etc) compared to all other groups.
- H_2 : Participants who engage with the robot using PACT will have a more positive perception of the robot than participants in the Reward-Maximizing and Median Pattern groups.
- H_3 : Constraining the robot's behavior to follow any patterns-based convention will result in better team performance on the task, as well as an improvement in participants' ability to predict the robot's actions.

5 RESULTS AND DISCUSSION

Of the 28 individuals who participated, the data of one participant was excluded due to noncompliance with instructions. We did not observe any multimodalities within the data.

We found a significant effect from the PACT Pattern condition on participant perceptions of the team's dynamics, **validating** H_1 . Post-hoc comparisons using Tukey's HSD test (Figure 7), indicate that participants felt that that robot picked the best block for the team during gameplay compared to the control condition of Reward-Maximizing ($p = 0.0353$) as well as the Median Pattern group ($p = 0.0493$). Additionally, PACT Pattern participants did not feel that swapping the robot out for a human teammate would result in better performance when compared to the Reward-Maximizing group ($p \checkmark 0.004$) as well as the Median Pattern group ($p \checkmark 0.009$), indicating that **PACT Pattern participants viewed the robot as performing at least as well as a human teammate would have**.

We also found a significant effect caused by the PACT Pattern condition on perceptions of team fluency, as indicated by Figure 8; the PACT Pattern condition resulted in a significantly higher perception of team fluency as compared to the Reward-Maximizing baseline ($p = 0.0178$), while there was no significant difference for participants in the Median Pattern group compared to the Reward-Maximizing condition. Additionally, there was a significant effect caused by the PACT Pattern condition on participants' perception of whether or not the participant and robot were good teammates to each other. PACT Pattern participants reported significantly more positive perceptions of themselves as a teammate to the robot when compared to the Reward-Maximizing ($p = 0.0129$) and the Median Pattern group ($p = 0.0336$), as well as whether the robot was a good teammate to the participant when compared to the Reward-Maximizing group ($p = 0.0121$). **The PACT Pattern group was also the only pattern-based**

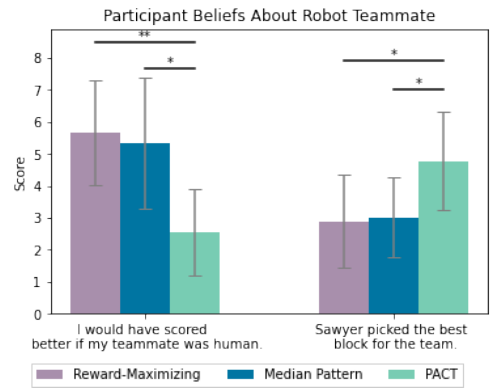


Fig. 7. There were significant improvements in PACT Pattern participant belief that the robot selected the right block for the team over the Reward-Maximizing ($p = 0.0353$) and Median Pattern ($p = 0.0493$) groups, as well as if the participants believed a human partner would have led to greater success. (Reward-Maximizing $p \checkmark 0.004$, Median $p \checkmark 0.009$).

Normalized Scores

Round	Group	Mean Score	p-value
1	Reward-Maximizing	6.81	—
1	Median Pattern	52.00	0.0007
1	PACT	71.27	0.0
2	Reward-Maximizing	27.12	—
2	Median Pattern	60.33	0.0246
2	PACT	78.68	0.0006
3	Reward-Maximizing	21.74	—
3	Median Pattern	79.49	0.0012
3	PACT	88.06	0.0003

Table 1. Normalized game scores and p-values obtained via Tukey’s HSD for each pattern-based group compared to the baseline Reward-Maximizing group. There were no significant differences between the PACT and Median groups for normalized scores across all rounds.

group that saw a significant difference over the Reward-Maximizing condition when asked if they would work with the robot again ($p = 0.0287$).

Post-hoc comparisons using Tukey’s HSD test indicate a **partial confirmation of H_2** . While there was a significant effect from the PACT Pattern treatment on the likeability of the robot when compared to the Median Pattern group ($p \checkmark 0.05$), there was no significant difference compared to the Reward-Maximizing group ($p = 0.1$), Figure 8.

We also found a significant effect from both pattern conditions on team performance, **validating H_3** . Post-hoc comparisons using Tukey’s HSD indicate significantly higher normalized scores for both the PACT Pattern and Median Pattern groups across all three rounds, as indicated by Table 1. As seen in Figure 9, the PACT Pattern group made significantly fewer errors when compared to the Reward-Maximizing group across all rounds of gameplay. The Median Pattern group made significantly fewer errors than the Reward-Maximizing group in rounds one and three, but there was no significant difference over the Reward-Maximizing in round two. Part of this may be due to the differences in the patterns seen by each group. Participants in the Median Pattern group saw much more ambiguous patterns than those in the PACT group, meaning that participants in the Median Pattern group could play at least half of the first round and obtain a perfect score by following a pattern other than the robot’s pattern. The Median Pattern group was the only group to show significance over the Reward-Maximizing after only one round of gameplay in their belief of understanding how the robot was choosing blocks ($p=0.0241$), but this effect was no longer significant after another round of gameplay.

Further validating H_3 , participants rated the predictability and understandability of the robot in a variety of questions in the Post-Experiment Survey. Using Tukey’s HSD, comparisons between

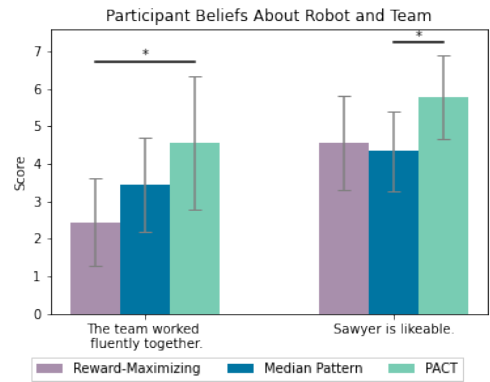


Fig. 8. Using PACT led to significant improvement in team fluency over baseline ($p = 0.0178$), as well as perceptions of robot likeability over the Median group ($p \checkmark 0.05$).

687 the Reward-Maximizing group and both patterns-based groups were significant (Figure 10). When
 688 compared to the Reward-Maximizing baseline, participants in both the PACT Pattern group ($p = 0.0001$)
 689 and the Median Pattern group ($p = 0.0001$) felt the robot was predictable. When asked
 690 about the understandability of the robot’s actions, the PACT Pattern group ($p = 0.0003$) and the
 691 Median Pattern group ($p = 0.0097$) both felt the robot was understandable compared to the Reward-
 692 Maximizing baseline. However, there is an important caveat to this finding. Participants were asked
 693 about the broader application of the system, and whether they believed most people would be
 694 able to understand the robot (Figure 10). **Only participants in the PACT group felt that most**
 695 **people would be able to understand the robot**, compared to both the Reward-Maximizing
 696 group ($p = 0.0009$) and Median Pattern group ($p = 0.0441$), confirming the premise of this work
 697 and validating the proposed contribution. While participants in the Median Pattern group believed
 698 at the end of gameplay that they understood the robot’s decisions, they did not believe that the
 699 system they saw would be broadly understandable.

700

701

701 5.1 Discussion

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

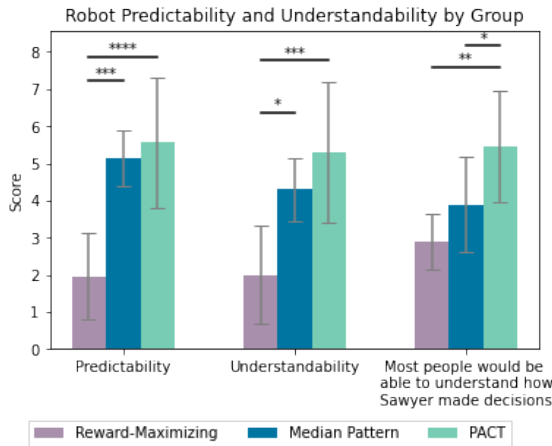


Fig. 10. Participants in the PACT Pattern ($p = 0.0001$) and the Median Pattern ($p = 0.0001$) both found the robot significantly more predictable than the baseline. Both groups also found the robot’s behavior more understandable than the baseline group. (PACT $p = 0.0003$, Median $p = 0.0097$) Only the participants who used PACT felt the robot would be broadly understandable to people when compared to the baseline ($p = 0.0009$) as well as the Median Pattern group ($p = 0.0441$).

of participants could fully explain the pattern they saw, despite many of them playing perfectly coordinated rounds with the robot.

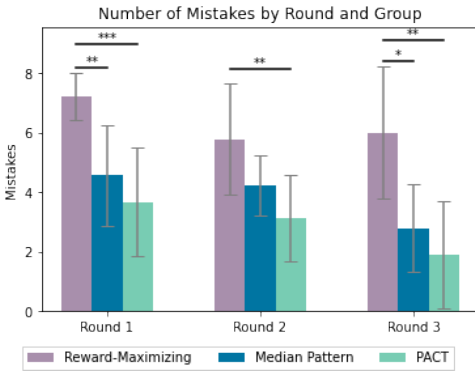


Fig. 9. Participants in the group that engaged with a robot using PACT made significantly fewer mistakes than the baseline group across all three rounds ($p = 0.0003, 0.0042, 0.0005$), whereas the Median group only made significantly fewer mistakes in two rounds ($p = 0.0047, 0.1138, 0.0053$).

robot. These study results reinforce the importance of leveraging convention in fluent human-robot collaboration, and confirm that PACT is an effective mechanism to do so.

This work demonstrates that intentionally leaning into human cognitive tendencies and de-emphasizing reward-maximizing behavior leads to substantially better outcomes along both objective and subjective metrics. Our proposed method does not preclude the usage of other planning tools, and can be used in tandem with other methods to make robots more predictable while remaining capable. Additionally, the tradeoff between optimal planning and predictability can be negotiated for any environment; PACT can create complex patterns similar to optimal planning, or simple ones to maximize predictability.

As robots are placed in environments where they will be trusted with a diversity of tasks, especially in cases where they will be in close contact with humans, it is critical to characterize and address the disparities between the way robots and humans reason. Robots that are exclusively optimizing for a given reward are reasoning about their environment and collaborations in a fundamentally different way from the humans that they work around and attempt to collaborate with. This leads to a lack of predictability, limiting collaboration. **PACT demonstrates that we can bridge this gap and make robots more predictable without limiting team performance.**

REFERENCES

- [1] Charles F. Abel. 2003. Heuristics and problem solving. *New Directions for Teaching and Learning* 2003, 95 (2003), 53–58. <https://doi.org/10.1002/tl.113>
- [2] Stéphane Airiau, Sandip Sen, and Daniel Villatoro. 2014. Emergence of conventions through social learning: Heterogeneous learners in complex networks. *Autonomous Agents and Multi-Agent Systems* 28, 5 (2014), 779–804. <https://doi.org/10.1007/s10458-013-9237-x>
- [3] Fatima M. Albar and Antonie J. Jetter. 2009. Heuristics in decision making. *PICMET: Portland International Center for Management of Engineering and Technology, Proceedings* (2009), 578–584. <https://doi.org/10.1109/PICMET.2009.5262123>

- 785 [4] Franziska Babel, Johannes Kraus, and Martin Baumann. 2022. Findings From A Qualitative Field Study with An
786 Autonomous Robot in Public : Exploration of User Reactions and Conflicts. *International Journal of Social Robotics* 14
787 (2022), 1625–1655. Issue 7. <https://doi.org/10.1007/s12369-022-00894-x>
- 788 [5] Serena Booth, Sanjana Sharma, Sarah Chung, Julie Shah, and Elena L Glassman. 2022. Revisiting Human-Robot
789 Teaching and Learning Through the Lens of Human Concept Learning. In *HRI '22: Proceedings of the 2022 ACM/IEEE*
790 *International Conference on Human-Robot Interaction*. 147–156. http://slbooth.com/HRI_Concept_Learning/index.html
- 791 [6] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. 2017. The robotic social attributes scale
792 (RoSAS) development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot*
793 *interaction*. 254–262.
- 794 [7] Micah Carroll, Rohin Shah, Mark K. Ho, Thomas L. Griffiths, Sanjit A. Seshia, Pieter Abbeel, and Anca Dragan. 2019.
795 *On the Utility of Learning about Humans for Human-AI Coordination*. Curran Associates Inc., Red Hook, NY, USA.
- 796 [8] Rui Chen, Alvin Shek, and Changliu Liu. 2021. Learn from Human Teams: a Probabilistic Solution to Real-Time
797 Collaborative Robot Handling with Dynamic Gesture Commands. *CoRR* abs/2112.06020 (2021). arXiv:2112.06020
798 <https://arxiv.org/abs/2112.06020>
- 799 [9] Sylvain Daronnat, Leif Azzopardi, Martin Halvey, and Mateusz Dubiel. 2021. Inferring Trust From Users' Behaviours;
800 Agents' Predictability Positively Affects Trust, Task Performance and Cognitive Load in Human-Agent Real-Time
801 Collaboration. *Frontiers in Robotics and AI* 8, July (2021), 1–14. <https://doi.org/10.3389/frobt.2021.642201>
- 802 [10] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In
803 *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (Tokyo, Japan) (HRI '13)*. IEEE
804 Press, 301–308.
- 805 [11] Ali Ghadirzadeh, Xi Chen, Wenjie Yin, Zhengrong Yi, Marten Bjorkman, and Danica Kragic. 2021. Human-Centered
806 Collaborative Robots with Deep Reinforcement Learning. *IEEE Robotics and Automation Letters* 6, 2 (2021), 566–571.
807 <https://doi.org/10.1109/LRA.2020.3047730> arXiv:2007.01009
- 808 [12] Gerd Gigerenzer. 2008. Why Heuristics Work. *Perspectives on Psychological Science* 3, 1 (2008), 20–29. <https://doi.org/10.1111/j.1745-6916.2008.00058.x> PMID: 26158666.
- 809 [13] Ji Han, Gopika Ajaykumar, Ze Li, and Chien Ming Huang. 2020. Structuring Human-Robot Interactions via Interaction
810 Conventions. *29th IEEE International Conference on Robot and Human Interactive Communication, RO-MAN 2020* (2020),
811 341–348. <https://doi.org/10.1109/RO-MAN47096.2020.9223468>
- 812 [14] Marc Hanheide, Annika Peters, and Nicola Bellotto. 2012. Analysis of human-robot spatial behaviour applying a quali-
813 tative trajectory calculus. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*
814 September (2012), 689–694. <https://doi.org/10.1109/ROMAN.2012.6343831>
- 815 [15] Sergiu Hart. 2005. Adaptive Heuristics. *Econometrica* 73, 5 (2005), 1401–1430. <http://www.jstor.org/stable/3598879>
- 816 [16] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical
817 and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in
818 Psychology, Vol. 52. North-Holland, 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- 819 [17] Bradley Hayes and Brian Scassellati. 2015. Effective robot teammate behaviors for supporting sequential manipulation
820 tasks. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6374–6380.
- 821 [18] Guy Hoffman. 2019. Evaluating Fluency in Human-Robot Collaboration. *IEEE Transactions on Human-Machine Systems*
822 49, 3 (2019), 209–218. <https://doi.org/10.1109/THMS.2019.2904558>
- 823 [19] Hengyuan Hu, Adam Lerer, Brandon Cui, Luis Pineda, Noam Brown, and Jakob Foerster. 2021. Off-belief learning. In
824 *International Conference on Machine Learning*. PMLR, 4369–4379.
- 825 [20] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. 2020. “Other-Play” for Zero-Shot Coordination.
826 In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research,*
827 *Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 4399–4410. <https://proceedings.mlr.press/v119/hu20a.html>
- 828 [21] Catholijn M. Jonker, M. Birna van Riemsdijk, and Bas Vermeulen. 2011. Shared Mental Models. In *Coordination,*
829 *Organizations, Institutions, and Norms in Agent Systems VI*, Marina De Vos, Nicoletta Fornara, Jeremy V. Pitt, and
830 George Vouras (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 132–151.
- 831 [22] Shauharda Khadka, Somdeb Majumdar, Tarek Nassar, Zach Dwiell, Evren Turner, Santiago Miret, Yinyin Liu, and
832 Kagan Turner. 2019. Collaborative evolutionary reinforcement learning. *36th International Conference on Machine*
833 *Learning, ICML 2019* 2019-June (2019), 5816–5827. arXiv:1905.00976
- 834 [23] Minae Kwon, Erdem Biyik, Aditi Talati, Karan Bhasin, Dylan P. Losey, and Dorsa Sadigh. 2020. When humans
835 aren't optimal: Robots that collaborate with risk-aware humans. *ACM/IEEE International Conference on Human-Robot*
836 *Interaction* (2020), 43–52. <https://doi.org/10.1145/3319502.3374832> arXiv:2001.04377
- 837 [24] Minae Kwon, Malte F. Jung, and Ross A. Knepper. 2016. Human expectations of social robots. *ACM/IEEE International*
838 *Conference on Human-Robot Interaction* 2016-April (2016), 463–464. <https://doi.org/10.1109/HRI.2016.7451807>
- 839 [25] Adam Lerer and Alexander Peysakhovich. 2019. Learning existing social conventions via observationally augmented
840 self-play. *AIES 2019 - Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (2019), 107–114.

- 834 <https://doi.org/10.1145/3306618.3314268> arXiv:1806.10071
- 835 [26] Fei Li, L. Phillip Wang, Xiaoming Shen, and Joe Z. Tsien. 2010. Balanced dopamine is critical for pattern completion
- 836 during associative memory recall. *PLoS ONE* 5, 10 (2010). <https://doi.org/10.1371/journal.pone.0015401>
- 837 [27] Olivier Mangin, Alessandro Roncone, and Brian Scassellati. 2022. How to be Helpful? Supportive Behaviors and
- 838 Personalization for Human-Robot Collaboration. *Frontiers in Robotics and AI* 8 (2022). <https://doi.org/10.3389/frobt.2021.725780>
- 839 [28] Barnaby Marsh. 2002. Heuristics as social tools. *New Ideas in Psychology* 20, 1 (2002), 49–57. [https://doi.org/10.1016/S0732-118X\(01\)00012-5](https://doi.org/10.1016/S0732-118X(01)00012-5)
- 840 [29] John E. Mathieu, Gerald F. Goodwin, Tonia S. Heffner, Eduardo Salas, and Janis A. Cannon-Bowers. 2000. The influence
- 841 of shared mental models on team process and performance. *Journal of Applied Psychology* 85, 2 (2000), 273–283.
- 842 <https://doi.org/10.1037/0021-9010.85.2.273>
- 843 [30] Mark P. Mattson. 2014. Superior pattern processing is the essence of the evolved human brain. *Frontiers in Neuroscience*
- 844 8, 8 AUG (2014), 1–17. <https://doi.org/10.3389/fnins.2014.00265>
- 845 [31] Shabnam Mousavi and Gerd Gigerenzer. 2014. Risk, uncertainty, and heuristics. *Journal of Business Research* 67, 8
- 846 (2014), 1671–1678. <https://doi.org/10.1016/j.jbusres.2014.02.013>
- 847 [32] Stefanos Nikolaidis and Julie A. Shah. 2012. Human-Robot Teaming using Shared Mental Models. *IEEE/ACM*
- 848 *International Conference on Human-Robot Interaction, Workshop on Human-Agent-Robot Teamwork (2012)* 17, 6 (2012),
- 849 1098–1106.
- 850 [33] Tatsuya Nomura, Tomohiro Suzuki, Takayuki Kanda, and Kensuke Kato. 2006. Measurement of negative attitudes
- 851 toward robots. *Interaction Studies* 7, 3 (2006), 437–454.
- 852 [34] Amit Kumar Pandey and Rachid Alami. 2009. A framework for adapting social conventions in a mobile robot motion
- 853 in human-centered environment. *2009 International Conference on Advanced Robotics, ICAR 2009* (2009).
- 854 [35] Bethany Rittle-Johnson, Emily R. Fyfe, Laura E. McLean, and Katherine L. McEldoon. 2013. Emerging Understanding
- 855 of Patterning in 4-Year-Olds. *Journal of Cognition and Development* 14, 3 (2013), 376–396. <https://doi.org/10.1080/15248372.2012.689897>
- 856 [36] Alessandra Rossi, Fernando Garcia, Arturo Cruz Maya, Kerstin Dautenhahn, Kheng Lee Koay, Michael L. Walters, and
- 857 Amit K. Pandey. 2019. Investigating the Effects of Social Interactive Behaviours of a Robot on People’s Trust During a
- 858 Navigation Task. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and*
- 859 *Lecture Notes in Bioinformatics)* 11649 LNAI, July (2019), 349–361. https://doi.org/10.1007/978-3-030-23807-0_29
- 860 [37] Basak Sakkak and Luca Bascetta. 2022. Safe Motion Planning for a Mobile Robot Navigating in Environments Shared
- 861 with Humans. (2022). arXiv:2206.07498 <http://arxiv.org/abs/2206.07498>
- 862 [38] Kristin E. Schaefer. 2016. *Measuring Trust in Human Robot Interactions: Development of the “Trust Perception Scale-HRI”*.
- 863 Springer US, Boston, MA, 191–218. https://doi.org/10.1007/978-1-4899-7668-0_10
- 864 [39] Andy Shih, Arjun Sawhney, Jovana Kondic, Stefano Ermon, and Dorsa Sadigh. 2021. On the Critical Role of Conventions
- 865 in Adaptive Human-AI Collaboration. In *Proceedings of the 9th International Conference on Learning Representations*
- 866 *(ICLR)*.
- 867 [40] DJ Strouse, Kevin McKee, Matt Botvinick, Edward Hughes, and Richard Everett. 2021. Collaborating with Humans
- 868 without Human Data. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin,
- 869 P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 14502–14515. <https://proceedings.neurips.cc/paper/2021/file/797134c3e42371bb4979a462eb2f042a-Paper.pdf>
- 870 [41] Aaquib Tabrez, Matthew B. Luebbbers, and Bradley Hayes. 2020. A Survey of Mental Modeling Techniques in Human-Robot Teaming. *Current Robotics Reports* 1, 4 (2020), 259–267. <https://doi.org/10.1007/s43154-020-00019-0>
- 871 [42] Mycal Tucker, Yilun Zhou, and Julie Shah. 2022. Latent Space Alignment Using Adversarially Guided Self-Play. *International Journal of Human-Computer Interaction* 0, 0 (2022), 1–19. <https://doi.org/10.1080/10447318.2022.2083463>
- 872 [43] Amos Tversky and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 4157
- 873 (1974), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- 874 [44] J. R. Wilson and A. Rutherford. 1989. Mental models: Theory and application in human factors. *Human Factors* 31, 6
- 875 (1989), 617–634. <https://doi.org/10.1177/001872088903100601>
- 876 [45] John Yen, Xiacong Fan, Shuang Sun, Rui Wang, Cong Chen, Kaivan Kamali, and Richard a. Volz. 2003. Implementing
- 877 Shared Mental Models for Collaborative Teamwork. *The Workshop on Collaboration Agents: Autonomous Agents for*
- 878 *Collaborative Environments in the IEEE/WIC Intelligent Agent Technology Conference, Halifax, Canada* (2003).

A ALGORITHM

Algorithm 1 Best Pattern Selection

Input: Set of tasks T , Set of Patterns P

Output: The pattern(s) best suited for T

```

888 1:  $minScore \leftarrow \infty$ 
889 2:  $bestPatterns \leftarrow \emptyset$ 
890 3: for  $p \in P$  do
891 4:    $score \leftarrow 0$ 
892 5:   for  $i \in 1 \leq i \leq |T|$  do
893 6:      $S_p \leftarrow$  every allowable sequence of length  $i - 1$  using  $p$ 
894 7:      $T_{i,p} \leftarrow []$ 
895 8:     for  $s \in S_p$  do
896 9:        $t_s \leftarrow$  all allowable next tasks after completing  $s$ , under pattern  $p$ 
897 10:       $T_{i,p}.extend(t_s)$ 
898 11:    end for
899 12:     $firstTerm = H(T_{i,p})$  // Calculate entropy
900 13:     $P_{i,shared} \leftarrow \{\}$  // Patterns sharing candidate seqs with  $p$ 
901 14:     $T_{i,shared} \leftarrow []$ 
902 15:    for  $q \in P$  do
903 16:       $S_q \leftarrow$  every allowable sequence of length  $i - 1$  using  $q$ 
904 17:       $S_q = S_q \cap S_p$  // Only sequences that also follow  $p$ 
905 18:      if  $|S_q| > 0$  then
906 19:         $P_{i,shared} \leftarrow P_{i,shared} \cup \{q\}$ 
907 20:        for  $s \in S_q$  do
908 21:           $t_s \leftarrow$  all allowable next tasks after completing  $s$ , under pattern  $q$ 
909 22:           $T_{i,shared}.extend(t_s)$ 
910 23:        end for
911 24:      end if
912 25:    end for
913 26:     $discount = \frac{|P_{i,shared}|-1}{|P|}$ 
914 27:     $secondTerm = discount * H(T_{i,shared})$ 
915 28:     $score \leftarrow score + firstTerm + secondTerm$ 
916 29:  end for
917 30:  if  $score = minScore$  then
918 31:     $bestPatterns \leftarrow bestPatterns \cup \{p\}$ 
919 32:  else if  $score < minScore$  then
920 33:     $minScore = score$ 
921 34:     $bestPatterns \leftarrow \{p\}$ 
922 35:  end if
923 36: end for
924 37: return  $bestPatterns$ 

```

B SURVEY QUESTIONS

Listed p-values are of the form (conventions/median, conventions/optimal, median/optimal).

B.1 Pre-Activity Survey

B.1.1 Experience with Robots. Questions in this section were either multiple choice, or select all that apply. Options for each question are listed below the question.

- Have you ever watched a movie or television show that includes robots? (0.86,0.28,0.55)
 - 0 shows/movies
 - 1-5 shows/movies
 - 6-10 shows/movies
 - 10+ shows/movies
- Have you ever interacted with a robot? (select all that apply) (0.22,0.22,0.22)
 - Museum or theme park animatronics
 - Toys such as Furby
 - Robot vacuum
 - Classroom robots or Battlebots
 - Sawyer (the robot in this experiment)
 - Everyday items such as cell phone, computer, ATM, or Xbox
 - Other
- Have you ever built a robot? (select all that apply) (0.11,0.22,0.11)
 - Classroom setting
 - Club setting
 - Other
- Have you ever controlled a robot? (select all that apply) (0.33,0.11,0.22)
 - Teleoperation or remote control
 - Speech, Gesture, Commands
 - Computer programmed
 - Other

B.1.2 Attitudes Towards Robots. The next set of questions detailed participants' attitudes towards robots in general. All questions were on a 7-point Likert scale, with 1 being Strongly Disagree and 7 being Strongly Agree. p-values in this section are based on the difference between pre- and post-activity surveys.

- I would feel uneasy if robots really had emotions. (0.27,0.14,0.92)
- Something bad might happen if robots developed into living beings. (0.12,0.95,0.21)
- I would feel relaxed talking with robots. (0.86,0.76,0.98)
- I would feel uneasy if I was given a job where I had to use robots. (0.003,0.06,0.45)
- If robots had emotions I would be able to make friends with them. (0.88,0.71,0.95)
- I would feel nervous operating a robot in front of other people. (0.02,0.84,0.06)
- I would hate the idea that robots were making judgements about things. (0.58,0.58,1.0)
- I would feel very nervous just standing in front of a robot. (0.26,1.0,0.26)
- I feel that if I depend on robots too much, something bad might happen. (0.71,0.99,0.78)
- I am good at working with robots. (0.39,1.0,0.39)
- I would feel paranoid talking with a robot. (0.98,0.58,0.68)
- I am concerned that robots would be a bad influence on children. (0.21,0.34,0.95)
- I feel that in the future society will be dominated by robots. (0.58,0.94,0.78)
- Most robots make poor teammates. (1.0,0.96,0.96)

- Most robots possess adequate decision making capabilities. (0.16,0.37,0.85)
- Most robots are easy to understand. (0.8,0.34,0.7)

B.1.3 Attitudes Towards Sawyer. This section of questions pertained to the participants' initial impression of the Sawyer robot. All questions are on a 7-point Likert scale. 1 was the adjective on the left, 7 was the adjective on the right. p-values in this section are based on the difference between pre- and post-activity surveys.

- I [blank] Sawyer. (Like/Dislike) (0.89, 0.97, 0.97)
- Sawyer is: (Unkind/Kind) (0.006, 1.0, 0.44)
- Sawyer is: (Ignorant/Knowledgeable) (0.07, 1.0, 0.07)
- Sawyer is: (Incompetent/Competent) (0.29, 0.92, 0.15)
- Sawyer is: (Unintelligent/Intelligent) (0.59, 0.98, 0.47)
- Sawyer is: (Foolish/Sensible) (0.31, 0.67, 0.07)
- Sawyer is a(n): (Individualist/Team Player) (0.66, 0.03, 0.15)
- Sawyer is: (Unlikeable/Likeable) (0.1, 0.9, 0.2)
- Sawyer is: (Unfriendly/Friendly) (0.53, 0.7, 0.16)
- Sawyer is: (Stubborn/Agreeable) (0.04, 0.52, 0.29)

B.2 Inter-Round Survey Questions

Other than the first question, which asked participants to select the round they had just completed, questions were on a 7-point Likert scale, and values for 1 and 7 are indicated in the form (adjective for 1 / adjective for 7) p-values in this section are written in the form (optimal r1/r2, optimal r1/r3, optimal r2/r3, median r1/r2, median r1/r3, median r2/r3, PACT r1/r2, PACT r1/r3, PACT r2/r3)

- Round
 - 1
 - 2
 - 3
- How mentally demanding was the task? (Very Low Mental Demand/Very High Mental Demand) (0.9, 0.9, 0.9, 0.83, 0.9, 0.9, 0.9, 0.9, 0.9)
- How successful were you in accomplishing what you were asked to do? (Perfect / Complete Failure) (0.9, 0.9, 0.9, 0.9, 0.9, 0.9, 0.9, 0.9, 0.9)
- How hard did you have to work to accomplish your level of performance? (Very Low Effort / Very High Effort) (0.72, 0.9, 0.8, 0.83, 0.9, 0.9, 0.75, 0.9, 0.9)
- How discouraged, irritated, stressed, and annoyed were you? (Very Low Frustration / Very High Frustration) (0.67, 0.53, 0.9, 0.82, 0.82, 0.9, 0.84, 0.9, 0.9)
- I was confident that Sawyer would choose the same block that I chose. (Very Low Confidence / Very High Confidence) (0.78, 0.56, 0.23, 0.9, 0.09, 0.17, 0.75, 0.16, 0.48)
- I understand how Sawyer was choosing blocks. (No Understanding / Complete Understanding) (0.79, 0.79, 0.44, 0.85, 0.65, 0.36, 0.82, 0.42, 0.75)

B.3 Post-Activity Survey

Listed p-values are of the form (conventions/median, conventions/optimal, median/optimal).

B.3.1 Game Comprehension. These questions concerned participants' understanding of the game. All questions are on a 7-point Likert scale. Value labels were Strongly Disagree (1) and Strongly Agree (7) unless otherwise stated.

- I understood the rules of the game. (0.9, 0.9, 0.9)
- I used the previous selections shown on the tablet to make my decisions. (0.28, 0.9, 0.44)

- 1030 • I knew things about the game that Sawyer didn't know. (0.9, 0.37, 0.32)
- 1031 • I understood the goal of the game. (0.9, 0.81, 0.86)
- 1032 • I kept track of our score at each turn. (0.9, 0.9, 0.9)
- 1033 • Sawyer knew things about the game that I didn't know. (0.79, 0.9, 0.79)
- 1034 • How much did your team's score influence the decisions you made? (No Influence / Score
- 1035 Was the Only Influence) (0.66, 0.54, 0.9)

1036 *B.3.2 Attitudes Towards Sawyer.* The questions in this section were identical to those asked in the
 1037 same section in the Pre-Activity Survey.
 1038

1039 *B.3.3 Team Fluency and Performance.* These questions concerned participants' perceptions of
 1040 their team. All questions are on a 7-point Likert scale. Value labels were Strongly Disagree (1) and
 1041 Strongly Agree (7) unless otherwise stated.

- 1042 • The robot and I contributed equally to the success of the team. (0.9, 0.6, 0.74)
- 1043 • Working with Sawyer was stressful or frustrating. (0.31, 0.67, 0.75)
- 1044 • I am responsible for the team's score. (0.9, 0.9, 0.9)
- 1045 • The team worked fluently together. (0.24, 0.07, 0.82)
- 1046 • I helped the robot accomplish the task. (0.9, 0.24, 0.46)
- 1047 • The team's coordination improved over time. (0.9, 0.02, 0.04)
- 1048 • The robot was cooperative. (0.26, 0.47, 0.87)
- 1049 • The robot is responsible for the team's score. (0.75, 0.41, 0.14)
- 1050 • If I were a robot, the team would have scored better. (0.9, 0.59, 0.61)
- 1051 • The robot perceived accurately what I was trying to do. (0.9, 0.72, 0.86)
- 1052 • I am good at working with robots. (0.53, 0.82, 0.24)
- 1053 • I contributed more to the success of the team. (0.83, 0.26, 0.59)
- 1054 • Working with Sawyer was difficult. (0.74, 0.25, 0.66)
- 1055 • The robot and I were working toward the same goal. (0.9, 0.9, 0.9)
- 1056 • The robot helped me accomplish the task. (0.9, 0.36, 0.58)
- 1057 • Sawyer is good at working with humans. (0.41, 0.56, 0.9)
- 1058 • I find what I am doing with the robot confusing. (0.9, 0.9, 0.9)
- 1059 • I was a good teammate to Sawyer. (0.075, 0.041, 0.9)
- 1060 • There was a team leader (True/False multiple choice) (0.77, 0.9, 0.9)
- 1061 • If there was a team leader, who was the team leader? (If there was no team leader, skip this
 1062 question) (Sawyer/Me) (0.56, 0.56, 0.56)
- 1063 • The robot contributed more to the success of the team. (0.86, 0.56, 0.29)
- 1064 • Over time, the way I selected blocks changed. (0.9, 0.37, 0.35)
- 1065 • Who is more responsible for the team's success or failure? (Sawyer / Me) (0.79, 0.9, 0.82)
- 1066 • Sawyer was a good teammate to me. (0.11, 0.04, 0.9)
- 1067 • I would have scored better if my teammate was human. (0.018, 0.004, 0.9)
- 1068 • I would work with Sawyer again. (0.35, 0.06, 0.66)

1069 *B.3.4 Robot Predictability and Understandability.* The questions in this section relate to the partici-
 1070 pant's understanding of the robot and how predictable they found the robot. All questions were on
 1071 a 7-point Likert scale from Strongly Disagree to Strongly Agree unless otherwise indicated.
 1072

- 1073 • Sawyer was unpredictable. (0.9, 0.014, 0.0395)
- 1074 • I understood why Sawyer made the decisions it did. (0.63, 0.0235, 0.18)
- 1075 • The way Sawyer selected blocks was unclear to me. (0.35, 0.0078, 0.2)
- 1076 • I could easily predict what block Sawyer would pick next. (0.31, 0.0078, 0.24)
- 1077 • The way Sawyer picked blocks made sense to me. (0.39, 0.0069, 0.16)

1078

- 1079 • As the game progressed, I was more easily able to predict which block Sawyer would pick
- 1080 next. (0.9, 0.001, 0.001)
- 1081 • Sawyer's decisions didn't make sense. (0.64, 0.07, 0.39)
- 1082 • Sawyer picked the best block for the team. (0.07, 0.16, 0.85)
- 1083 • Sawyer chose blocks randomly. (0.24, 0.001, 0.008)
- 1084 • Most people would be able to understand how Sawyer made decisions. (0.17, 0.01, 0.47)
- 1085 • I chose blocks (intuitively / analytically) (0.63, 0.24, 0.045)
- 1086 • Fill in the blank: By the end of Round [blank] I could easily predict which block Sawyer
- 1087 would pick next. (multiple choice)
- 1088 1
- 1089 2
- 1090 3
- 1091 None

1092 **B.4 Round 4 Survey**

1094 For this survey, participants were shown 5 novel game boards and were asked the same set of
 1095 multiple choice questions for each of them. Participants were instructed not to guess, and to select
 1096 "unsure" if they were not totally certain about their answer.

- 1097 • Which color is the block Sawyer will pick first?
- 1098 blue
- 1099 red
- 1100 yellow
- 1101 unsure
- 1102 • Which shape is the block Sawyer will pick first?
- 1103 circle
- 1104 triangle
- 1105 square
- 1106 unsure
- 1107 • Which color is the block Sawyer will pick last?
- 1108 blue
- 1109 red
- 1110 yellow
- 1111 unsure
- 1112 • Which shape is the block Sawyer will pick last?
- 1113 circle
- 1114 triangle
- 1115 square
- 1116 unsure

1117
 1118 Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127