



Multi-Armed Bandits with Fairness Constraints for Distributing Resources to Human Teammates

Houston Claire
Cornell University
Ithaca, New York
hbc35@cornell.edu

Yifang Chen
University of Southern California
Los Angeles, California
yifang@usc.edu

Jignesh Modi
University of Southern California
Los Angeles, California
jigneshm@usc.edu

Malte Jung
Cornell University
Ithaca, New York
mfj28@cornell.edu

Stefanos Nikolaidis
University of Southern California
Los Angeles, California
nikolaid@usc.edu

Abstract

How should a robot that collaborates with multiple people decide upon the distribution of resources (e.g. social attention, or parts needed for an assembly)? People are uniquely attuned to how resources are distributed. A decision to distribute more resources to one team member than another might be perceived as unfair with potentially detrimental effects for trust. We introduce a multi-armed bandit algorithm with fairness constraints, where a robot distributes resources to human teammates of different skill levels. In this problem, the robot does not know the skill level of each human teammate, but learns it by observing their performance over time. We define fairness as a constraint on the minimum rate that each human teammate is selected throughout the task. We provide theoretical guarantees on performance and perform a large-scale user study, where we adjust the level of fairness in our algorithm. Results show that fairness in resource distribution has a significant effect on users' trust in the system.

Keywords

Reinforcement Learning; Fairness; Multi-Armed Bandits; Trust

ACM Reference Format:

Houston Claire, Yifang Chen, Jignesh Modi, Malte Jung, and Stefanos Nikolaidis. 2020. Multi-Armed Bandits with Fairness Constraints for Distributing Resources to Human Teammates. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20), March 23–26, 2020, Cambridge, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3319502.3374806>

1 Introduction

For robots to function effectively in teams of multiple people, they have to be able to decide about the distribution of resources (e.g. assistance or social attention) [34, 54, 58]. Consider a factory robot that assists two workers by delivering parts needed for an engine

assembly. One worker is experienced and fast, the other inexperienced and slow. How should a robot take expertise into account when dividing its assistance among workers to achieve optimal outcomes?

The successful adoption of robots as parts teams requires not only the consideration of team performance (e.g. completion time, or cost), but also of team viability, that is the capability of team members to continue to work cooperatively over time [6]. Team viability requires trust. A robot distributing more resources to one worker than another might be perceived as unfair and consequently untrustworthy [43]. In fact, previous work has shown that ignoring human preferences in task allocation can negatively affect users' willingness to work with the system [26]. Groom and Nass [28] argue that our ability to build effective human-robot teams depends on a team's ability to build trust between all members of a team, and much work in human-robot interaction has focused on establishing perceived team fluency and trust in human-robot teams [5, 13, 19, 20, 25, 32, 35, 52, 55].

Here we focus on the notion of *fairness* in resource distribution. We formalize how a robot can take individual expertise into account to maximize team performance, while guaranteeing that each human teammate will be assigned a minimum rate of resources at any given time throughout the task. Our thesis is that, by accounting for fairness in resource allocation, we can significantly improve users' trust in the system.

To this end, we cast the problem as a multi-armed bandit, where each human teammate is represented as an arm with an unknown reward function corresponding to their skill level. We then propose a *multi-armed bandit algorithm with fairness constraints*, which builds upon the standard Upper Confidence Bound (UCB) algorithm [4]. We propose a stochastic version of the algorithm, where a minimum pulling rate for each arm is satisfied in expectation, and a deterministic version where the constraint is strictly satisfied anytime throughout the task. We provide a new definition of regret and theoretical guarantees of performance in the form of regret bounds for both algorithms.

To assess the effect of fairness on the users, we execute a large-scale user study on a Tetris game, where two players are sequentially assigned a batch of blocks by the algorithm. We selected the Tetris game, since it emulates collaborative tasks in human-robot interaction where a robot provides resources to human teammates [34, 52], it provides a clear and transparent way to assess the

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.

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6746-2/20/03...\$15.00

<https://doi.org/10.1145/3319502.3374806>

participants' performance and it can model a wide range of task characteristics [47].

We implement the algorithm with three levels of fairness, representing the required minimum allocation rate for each player: 25%, 33% and 50%. Results show that fairness significantly affects the trust of the players that performed worse than their teammates: those in the 33% condition trusted the system significantly more, compared to the 25% condition. Surprisingly, we did not observe a decrease in performance in the fairer distributions, even though the stronger player was selected less frequently. On the contrary, the median scores were higher when fairness increased. These results improve our understanding of the theory and implications of fairness in resource distribution in human-robot teams.

2 Background

2.1 Distributive Fairness in Resource Distribution

Fairness has been shown to be important for successful collaboration [17, 29]. While fairness can be construed in many ways, we adopt a distributive perspective on fairness [3] and operationalize it consistent with [39] as the degree of which resources are distributed equally to individuals within a group. While an equal distribution of resources across all members within a group seems ideal, researchers [22, 41] have shown that inequalities are deemed appropriate, particularly when they optimize the outcome of the group. Adam's model on equity suggests that allocation decisions are deemed appropriate if they are in proportion with the input of the individual [1, 57]. This model has been tested in various laboratory and real world scenarios suggesting that allocation decisions in groups follow such a model [7, 8, 27, 40, 59]. On the other hand, perceived inequalities have a strong impact on individuals' behavior, often motivating them to act contrary to their rational self-interest with the goal of eliminating the inequality [11, 42]. Previous work has shown that perceived lack of fairness affects job satisfaction [49] and can induce retaliation behavior from the affected party [56].

Interestingly, recent works have shown that individuals perceive fairness differently when decisions are made by an algorithm, compared to a human [43, 44]. As research in HRI advances, robots will be increasingly placed in complex environments where they will be making allocation decisions. From allocating time, resources, and attention, these robotic systems will require an understanding of the impact their allocation decisions can have on individual and organizational dynamics.

2.2 Stochastic Multi-Armed Bandits

The stochastic multi-armed bandits (MAB) framework without a minimum pulling rate requirement has been theoretically well studied. The gambler is tasked with choosing an arm, i , from K arms at each time step $t = 1, 2, 3, \dots, n$. At every time t , the gambler pulls an arm $i_t \in [K]$ while simultaneously the environment decides the reward vector $r_t \in [0, 1]^K$ from a fixed distribution with expectation $\mathbb{E}[r_t(i_t)] = \mu(i_t)$. The gambler, however, can only observe $r_t(i_t)$ but not the whole vector. Therefore, the gambler's goal is to pull the

sequence of arms, based on the past information, that can maximize the overall accumulated reward.

The best arm in hindsight is defined as $i^* = \operatorname{argmax}_{i \in [K]} \mu(i)$ and $\mu^* = \mu(i^*)$. We use regret to measure the performance of this algorithm, which is how worse our algorithm performs compared to the benchmark strategy – always pulling the best arm in each step.

$$\operatorname{Reg}_T = T\mu^* - \sum_{t=1}^T \mu(i_t)$$

An optimal solution to such a problem was proposed as the Upper Confidence Bound (UCB). It was originally introduced by Lai and Robbins [38] and expanded by Agrawal [2]. Building upon these works, Auer, Cesa-Bianchi & Fisher [4] introduced the Upper Confidence Bound Algorithm (UCB). In the most basic form of this algorithm, at each time t , we estimate the expected reward of each arm by using the mean of its empirical rewards in the past and the number of times it has been pulled, which gives us a confidence interval that the arm will lie in. Then the algorithm proceeds to pick the arm with the largest estimated expected reward.

This work has inspired a family of upper confidence bound variant algorithms for an array of different applications [21, 23, 37, 46, 48]. For a review of these algorithms we point readers to [10].

More recent work regarding multi-armed bandits has seen applications towards the improvement of human-robot interaction. Recent work has investigated using a MAB algorithm for the use of an assistive robotic system with the goal of exploring human preferences [12] and assisting human learning [50].

Of particular relevance is very recent work on sleeping bandits with fairness constraints [45], in a setting where multiple arms can be played simultaneously and some arms may be unavailable. Fairness is defined as a minimum rate satisfied in expectation and at the end of the task, whereas in our work we require the rate to be satisfied strictly and anytime throughout the task. Fairness in the context of MABs has also been studied in Joseph et al. [33]. The definition of fairness there is quite different, in that a worse arm should not be picked compared to a better arm, despite the uncertainty on payoffs. Their proposed algorithm chooses two arms with equal probability, until it has enough data to deduce the best of the two arms.

In addition, in parallel to our efforts, independent research [51] has very recently proposed similar definitions of fairness, where a fairness-satisfaction phase – that ensures that fairness is guaranteed – is succeeded by a regret minimization phase. We refer the readers to this coming interesting work as well.

3 Algorithm

We propose two new algorithms with optimal regret bound guarantees. The original unconstrained UCB algorithm fails in ensuring "fairness" because when time passes, a large set of "bad" arms will hardly be used again. Both of the algorithms we propose are based on the unconstrained UCB algorithm, where we adopt the idea of estimating the expected reward of each arm by using the mean of its empirical rewards in the past and the number of times it has

been pulled. We prove the following theorems in supplementary material.¹

3.1 Strict-rate-constrained UCB Algorithm

DEFINITION 1. Let S be any K -elements set whose elements are drawn from $[\frac{1}{v}]$ without replacement. Then define $g : S \rightarrow [K]$ as some one-to-one function.

Algorithm 1: Strictly-rate-constrained UCB

Input: time horizon T , arm set $[K]$, minimum pull rate v
Definition: Denote $UCB_t(i) = \frac{1}{t-1} \sum_{s=1}^{t-1} r_t(i) \mathbf{1}\{i_s = i\} + 2\sqrt{\frac{\ln T}{n_{t-1}(i)}}$, and τ_j be the starting time of block j .
Initialize: $t = 1, j = 1, \tau_1 = K + 1, \tau_j = \tau_1 + \frac{j-1}{v}$.
while $t \leq K$ **do**
 Pull arm $i_t = t$
 $t \leftarrow t + 1$.
for $j = 1, 2, 3, \dots$ **do** ▷ j indexes a block
 while $t < \tau_{j+1}$ **do**
 If $t - \tau_j + 1 \in S$, then pull the arm $i_t = g(t - \tau_j + 1)$,
 Otherwise, pull the arm $i_t = \operatorname{argmax}_{i \in [K]} UCB_t(i)$
 $t \leftarrow t + 1$

The algorithm divides T into blocks with length $\frac{1}{v}$. The algorithm is flexible in that there are multiple choices of S and g that satisfy the minimum rate constraint. For example, if $v = \frac{1}{4}$ and $K = 2$, we can choose $S = \{1, 3\}$ and $g(1) = 1, g(3) = 2$, which means we always pull arm 1 at τ_j and arm 2 at $\tau_j + 2$ for all j (see implementation in Section 4).

This algorithm guarantees that IN PRACTICE the pulling rate at any time for each arm is at least $v - \epsilon$, by fixing certain time slots where the algorithm will pull the prescheduled arms. Here $\epsilon = 1/t$.² In other time slots, the algorithm will behave just like the standard UCB.

Now the benchmark strategy for pulling an arm is always pulling the best arm in those non-prescheduled time slots. So the regret definition in this case becomes:

$$Reg_T = \mathbb{E}_{env} \left[\sum_{t \in \mathcal{I}} r_t(i^*) - r_t(i_t) \right]$$

where \mathcal{I} is all the non-prescheduled time slots among T .

THEOREM 1. By running Alg. 1, we obtain the regret bound that is close to the original unconstrained UCB,

$$Reg_T \leq \sum_{i: \Delta_i > 0} \left[\frac{16 \ln T}{\Delta_i} \left(\frac{1 - Kv}{1 - (K-1)v} \right) + 2(1 - Kv)^2 \Delta_i \right] + O(K)$$

If $\Delta_i \in [0, 1] \forall i$, we also get the worst case guarantee,

$$Reg_T \leq O(\sqrt{TK \ln T} + K \ln T)$$

¹<https://arxiv.org/pdf/1907.00313.pdf>

² We can prove this by observing that at time t , the arm i will be pulled at least $\lfloor tv \rfloor$ times according to the pre-schedule. So the pulling rate will be $\frac{\lfloor tv \rfloor}{t} \geq \frac{tv-1}{t} = v - \frac{1}{t}$.

3.2 Stochastic-rate-constrained UCB Algorithm

Algorithm 2: Stochastic-rate-constrained UCB

Input: time horizon T , arm set $[K]$, minimum pull rate v
Definition: Denote $UCB_t(i) = \frac{1}{t-1} \sum_{s=1}^{t-1} r_t(i) \mathbf{1}\{i_s = i\} + 2\sqrt{\frac{\ln T}{n_{t-1}(i)}}$.
Initialize: $t = 1, j = 1, \tau_1 = K + 1, \tau_j = \tau_1 + \frac{j-1}{v}$.
while $t \leq K$ **do**
 Pull arm $i_t = t$
 $t \leftarrow t + 1$.
for $t = K + 1, K + 2, K + 3, \dots$ **do**
 With probability $1 - Kv$, pull the arm
 $i_t = \operatorname{argmax}_{i \in [K]} UCB_t(i)$,
 Otherwise, uniformly pull an arm i_t from all K arms

This algorithm guarantees that the EXPECTED pulling rate at any time for each arm is at least v . Instead of rescheduling some arms as in the deterministic algorithm above, this algorithm introduces some randomness. At each time t , we ensure that each arm has at least v probability to be pulled; while with $1 - Kv$ probability, the algorithm will again pull the arm with the best UCB bound. We denote this distribution over arms as p_t where $p_t(\operatorname{argmax}_{i \in [K]} UCB_t(i)) = (1 - Kv) + v$ and $p_t(i) = v, \forall i \in [K] \setminus \operatorname{argmax}_{i \in [K]} UCB_t(i)$.

In this case, the benchmark strategy is pulling the best estimated arm with probability $(1 - Kv)$ at time t , otherwise uniformly drawing a random arm. We present this strategy with the distribution p^* over K arms where $p^*(i^*) = (1 - Kv) + v$ and $p^*(i) = v, \forall i \in [K] \setminus i^*$. So the regret definition in this case becomes:

$$Reg_T = \mathbb{E}_{env, learner} \left[\sum_{t=1}^T \mathbb{E}_{i_t \sim p^*} [r_t(i_t)] - \mathbb{E}_{a_t \sim p_t} [r_t(i_t)] \right]$$

THEOREM 2. By running Alg. 2, we obtain the regret bound that is close to the original unconstrained UCB,

$$Reg_T \leq \sum_{a: \Delta_i > 0} \left[\min \left\{ \frac{16 \ln T}{\Delta_i} + (1 - Kv) \Delta_i, (1 - Kv) \Delta_i T \right\} \right]$$

If $\Delta_i \in [0, 1] \forall i$, we also get the worst case guarantee,

$$Reg_T < O(\sqrt{TK \ln T} + K \ln T)$$

4 Evaluation

To evaluate our algorithm we conducted an online user study. We asked teams of two participants each to collaborate with a robot in completing a cooperative game (a modified version of Tetris). The robot's task is to decide which team-mate receives the next block to place. Similar to other human-robot collaboration tasks (e.g. [25, 26, 31, 34, 53]), the robot's role is that of a task or resource allocator as it provides a resource needed to complete a task (in this case Tetris blocks) to participants.

We designed a between-teams study with three conditions of increasing distributional fairness constraints maintained by the UCB

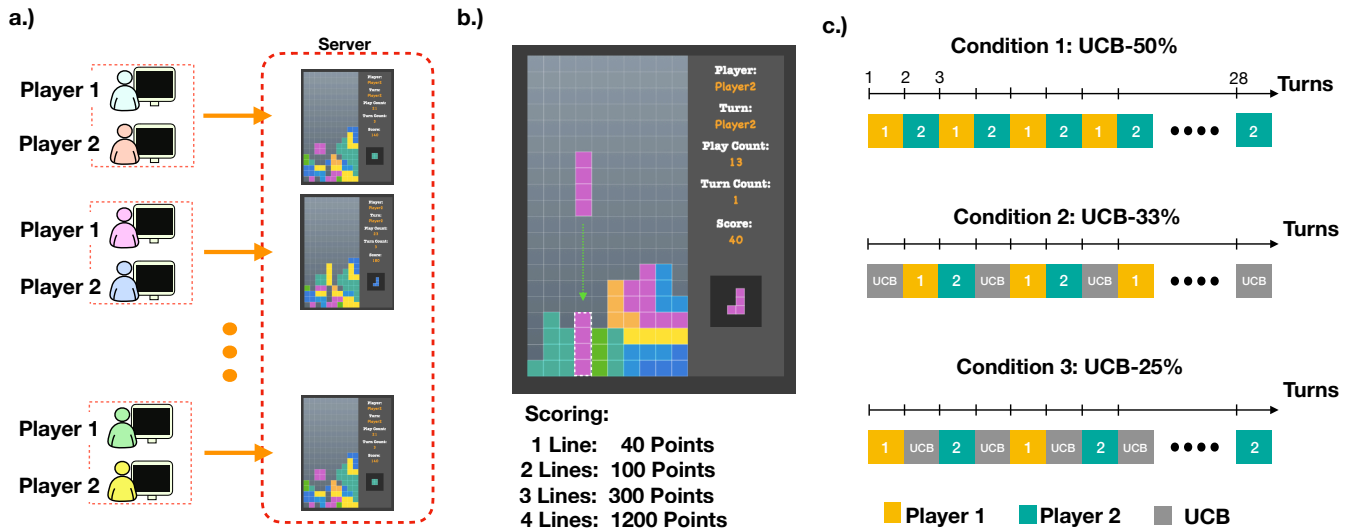


Figure 1: (a) Pairs of two remote human participants were connected to our cooperative Tetris game online. (b) The Tetris game followed the standard rules with two slight modifications. A different scoring metric as shown and only one participant had access to control the pieces per turn. (c) A visual representation of the three separate patterns that each condition offered.

Algorithm (Alg. 1): (UCB-25%,UCB-33%, UCB-50%).³ To reduce variance from sampling, we implemented the strict-rate-constrained UCB Algorithm.⁴

We characterize the player that has the best performance of the two, as observed at the end of the task, as *strong* and the other player as *weak*. The challenge of balancing between choosing the historically best player or a sub-optimal player allows us to investigate the impact of the system’s decision on team performance, perceived fairness and trust in the system.

Based on prior work which has shown that people react strongly to fairness in resource allocation [9, 43], with especially strong reactions occurring for the disadvantaged party (e.g. [56]), we expect that distribution rates (fairness) will have a significant effect on perceived fairness and trust in the system of the weak players. Specifically, as the distribution rates favor the stronger individuals at an increasing rate (UCB-50%,UCB-33%,UCB-25%, respectively) we expect fairness perceptions and trust in the system to decrease (**H1**). Furthermore, prior literature has shown that fairness in resource distribution has implications for team performance [16]. In our case, the fairer distributions favor the weak players, since they impose a constraint on the minimum number of pulls for both players. We expect that this will result in worse performance, compared to the less fair distributions that favor the strong player of the team (**H2**).

4.1 Methodology

4.1.1 Participants

³ We did not include a UCB-0% condition, since in our pilot studies in the team-Tetris task the UCB-0% and UCB-25% had similar allocation of pieces, because of the variance in the scores and the exploration inherent in the unconstrained UCB-0%.

⁴ https://github.com/anonym187/RL_Fairness_Constraints

We recruited 290 participants from Amazon Mechanical Turk (AMT) and paid \$1.00 for their participation in the task. We selected participants with previous ratings of 95% or higher. 8 data points were removed, since their AMT unique ID did not match the one given on Qualtrics. The final dataset contained $N = 94, 98, 90$ participants for UCB-50%, UCB-33%, and UCB-25% respectively (156 female, 124 male, 1 other, 1 did not disclose). The average age of participants was 36 years old ($SD = 11$). Participants were recruited if they could speak English, were from the United States, and had previous ratings of 95% or higher. Of the 282 participants, 6 of them reported to have never played Tetris before.

4.1.2 Task: Collaborative Tetris

Building on a task developed by [34] we developed a collaborative Tetris game that paired teams of two people to complete a game of Tetris together with our MAB algorithm. The goal for each team was to achieve the highest score possible. At each defined time step, the algorithm decides which team-mate should have control over the falling pieces—thus only one human player has control over the set of Tetris blocks at each time step, while both players observe the moves of the blocks at all times.

We chose Tetris as a collaborative task, since it has been shown to effectively model a broad range of common task characteristics [47], having been used as a testbed for several other studies (e.g. [30, 36]). It emulates previous settings where a robot assists users in a collaborative task [34, 53], providing a transparent and unambiguous way for the participants to observe each other’s performance.

We formally define our scenario as follows. The number of players in each game is set as $P = \{1, 2\}$ over a time horizon of $T = 30$. At each time step $t \leq T$, seven consecutive Tetris pieces are allotted to a player $p \in P$. In the turns where the UCB algorithm was run,




Figure 2: Responses to the subjective questions, grouped by player performance across each condition. Error bars represent the 95% confidence intervals.

we used as reward $r_{p,t}$:

$$r_{p,t} = \frac{S_{p,t}}{M * n_{p,t}}$$

where $S_{p,t}$ is the score achieved by player p up to turn t , $n_{p,t}$ is the number of turns of that player and M is a maximum value that we selected for normalization. After multiple pilot sessions we empirically set M to 300.

We defined each time step as a set of seven consecutive falling pieces that only the selected player could control. Our pilot sessions showed that allowing control of seven consecutive pieces together with limiting the width of the Tetris board prevented behaviors where one would place the blocks in such a way that prepares the groundwork for their partner to get the rewards. Observing the players' behavior in the pilot sessions, as well as their responses to questionnaires at the end of the study, confirmed that recorded scores matched observed performance.

4.1.3 Procedures

Upon providing informed consent and entering basic demographic information, AMT participants were instructed that they would be paired with a human partner and a robot that would decide who has control of the falling pieces and that the objective was to obtain the largest possible team score. Following standard Tetris rules, a player could rotate, speed up, or drop each falling piece. At the end of the time step the 50%, 33%, or 25% UCB, depending on condition, algorithm would run to select the next player.

Figure 1(c) shows the pattern of the distribution that was seen across each condition. This pattern was repeated for 30 time steps, with the exception of the first two time steps where each player played once. Each team was exposed to 210 pieces total. A code was given to participants upon the completion of the 30 rounds which enabled them to continue the Qualtrics survey.

4.1.4 Measures

Subjective: To measure levels of perceived *decision fairness* and *trust* we adapted survey questions from [43] (Table 1). Each response was measured on a seven-point Likert scale. Finally, we asked an open ended question, "In your own words, describe the strategy that you think the robot used to distribute the Tetris pieces."

Table 1: Subjective Metrics

Factor	Question Number	Question
Decision Fairness	Q1	How fair or unfair is it for your partner that the system gave them the designated number of Tetris pieces?
	Q2	How fair or unfair is it for you that the system gave you the designated number of Tetris pieces?
Trust	Q1	How much do you trust the system to make a good quality decision in the distribution of the Tetris pieces?

Objective: Information regarding an individual's performance was stored in a database during game play. We collected each player's individual score as well as the number of turns that was allocated to them. Additionally, we obtained the total score that each team accumulated at the end of the game play. Figure 1(b) shows the scoring convention that we used.

4.2 Results

Subjective: We grouped subjective responses of each pair of players based on their comparative performance in the game (Figure 2). We focus the analysis on the weak players, that is the players that performed worse than their teammate. We present the responses of the strong players as well for completeness.

A one-way ANOVA was performed for weak players across all conditions (UCB-50% vs. UCB-33% vs. UCB-25%) for each subjective metric. Analysis indicates a significant effect of the reported trust score of the weak players across the three conditions ($F(2, 138) = 3.172, p = 0.025$). A Tukey HSD with adjusted p-values demonstrated higher trust ($p = 0.037$) towards the system running the UCB-33% compared to the UCB-25%. While trust scores in UCB-50% were higher than in the UCB-25%, the difference was not significant ($p = 0.061$). Differences between all other factors were not significant.

Post-hoc Analysis. We observed a noticeable difference in the responses between the strong and the weak players for different fairness conditions (Fig. 2). Therefore, we conducted a post-hoc analysis to assess whether there were significant differences in the




Figure 3: (Top) Percentage of the number of pieces that each player received. Each bar represents a separate game. (Bottom) Plots of the total scores that each player individually and both players together achieved. Each bar represents a separate game.

responses between the weak and the strong players *within* each condition. Indeed, a 2 x 3 ANOVA with strength (weak vs. strong) and rate (UCB-50% vs. UCB-33% vs. UCB-25%) showed a main effect of players' strength for Decision Fairness Q2 ($F(1, 276) = 4.778, p = 0.0297$). There were no interaction effects. Post-hoc comparison with Bonferroni corrections looking at strength within the different fairness levels, showed that weak players ($M = 3.97, \sigma = 1.68$) reported significantly lower ratings on fairness (Q2) than their strong counterparts ($M = 4.82, \sigma = 1.49$) in the UCB 25% condition ($p = 0.035$), which was the least fair condition. We observed no significant difference in perceived fairness between strong and weak players in the other two conditions.

In summary, there was a significant difference between the weak and the strong players in their perception of fairness in the least fair condition (UCB-25%), and reducing the minimum rate from 33% to 25% had a negative effect on the trust of weak players. On the other hand, Figure 2 shows that trust scores between the UCB-50% and UCB-33% conditions were comparable.

To interpret these results, we observe the number of pieces received (arm pulls) for each condition in Figure 3(top). In the UCB-50% condition, all players received the same number of pieces regardless of their performance. In the UCB-33% condition, while the strong players received more pieces, the difference with the weak players was small. On the other hand, in the UCB-25% condition there were several games where the weak players received less than 30% of the pieces, resulting in lower reported trust in that condition.

We also examined participants' perceptions of resource distributions. Several participants in the UCB-33% and UCB-25% conditions noted how the system appeared to favor the "stronger" player during the gameplay:

"I felt the more competent player was given more turns. Which makes sense but was why it felt unfair."

Participants in the UCB-50% condition noted how the system gave each team member an equal number of turns:




Figure 4: Average Score per turn across the different conditions with 95% confidence intervals. Each point represents the average score of all participants at that time step split into weak and strong participants.




Figure 5: (a) Total scores for each condition. (b) Average score per turn for each condition.

"I think it was even, it made us take turns one after the other, enough that it made me feel I was making an equal contribution to the game."

Objective: A one-way ANOVA on the performance of the two-player teams across the three conditions indicated no statistical significance. In fact, Figure 5(a) shows that the medians of the total scores were higher for increasing levels of fairness. Plotting the individual scores of the players for each game in Figure 3(bottom) illustrates this tendency as well.

This result does not match our initial hypothesis. To interpret this result, we plot the average scores per turn for each condition in Figure 5(b). The average scores indicate how well the players performed on average every time they took a turn. Interestingly, we see that the distribution of the weak players' scores shifts towards lower scores as fairness decreases. While this result warrants further investigation, it indicates that assigning significantly less pieces to one of the players may negatively affect their performance, in addition to reducing their trust in the system. It showcases the importance of fairness when making resource distribution decisions.

We further tested for learning effects, since players may get better at the game over time. We fit a linear mixed effects model for the fixed effects of strength (weak vs. strong) and time as a continuous variable while including participants as random effects. We found a statistically significant increase in player scores over time in the UCB-50% ($F(1,1313)= 6.69, p= 0.010$) and UCB-25% condition ($F(1,1285)= 5.50, p= 0.019$). We found no statistically significant increase for participants in the UCB-33% condition ($F(1,1359)= 0.140, p= 0.708$). Indeed, Figure 4 shows that the average player scores tended to increase after the first few turns. These changes in performance occur in the first half of the game suggesting that learning effects did not drive our conclusions. Our goal was not to find the "true" stronger or weaker player within each game of Tetris, rather it was to assess how varying levels of distributions affect the team's performance within a given time period. The notions of "strong" and "weak" are with respect to their distribution of scores within the given 30 turns. While we have assumed fixed (albeit stochastic) reward distributions, Sliding-Window-UCB-based algorithms [14, 24, 60] have been proposed for evolving distributions. We suggest as future work extending these algorithms to account for fairness; providing theoretical guarantees for these algorithms would follow the same reasoning as in the proofs that we include in the supplemental material.

5 Discussion

We proposed a novel algorithm for a robot's resource distribution for human robot collaboration scenarios that include multiple human team members. Specifically, our MAB variant algorithm aimed to maintain a level of fairness by administering a minimum rate constraint limiting the number of times an individual may be assigned a resource.

An evaluation of the algorithm in a collaborative Tetris game showed that optimizing not only based on team member expertise, but also based on distributive fairness, lead to higher trust without a decline in team performance. Specifically, our study revealed a

statistically significant difference between the weak players in the UCB-33% and the UCB-25% conditions, partially supporting H1. We conjecture that this is due to the fact that weaker players in the UCB-25% were exposed to longer waiting periods and smaller number of turns than that of weaker players in the UCB-33%. This highlights the result of unfair distribution on trust, particularly that of weaker performing individuals.

Contrary to our expectations, our study did not find any differences in perceived fairness across the UCB-50% , UCB-33%, and UCB-25% conditions. Since the goal was stated as to maximize the team's overall score, some participants may have seen it as appropriate for the stronger player to receive more turns. For these participants, distributions that favoured one of the players might have been seen as procedurally fair. Procedural fairness "refers to the perception by the individual that a particular activity in which they are a participant is conducted fairly" [18]. In other words the algorithmic or procedural nature of the distribution may have contributed to a perception of fairness irrespective of the actual distribution. This interpretation is consistent with interview reports from some participants, for instance:

"I think it was judging that the other player was much better than me and thus they deserved to play more pieces than myself."

Additionally, we found no statistically significant differences in perceived fairness between the strong and weak participants within the UCB-50%, UCB-33% conditions. In these conditions the number of turns between the strong and weak participants did not differ much (Figure 3), suggesting that many participants did not perceive the distribution as unfair as they experienced almost equal participation in the game (and exactly equal participation in the UCB-50% condition). On the other hand, the UCB-25% condition saw a larger difference in the number of turns between weak and strong players and that lead to a statistically significant difference between the two groups when asked about the perceived fairness of the distribution (Decision Fairness Q2).

Finally, we see that the performance of the teams was not significantly impacted by the difference in resource distribution, which does not support H2. A possible reason is that, when weaker participants were limited in the number of turns and participation within the team, they might have put lower effort, affecting the overall performance, as shown by the low average scores of the weak players in UCB-25% in Figure 5(b).

These results are in line with previous work that highlights the importance of perceived allocation fairness on trust [15, 16]. As robots are increasingly placed in contexts where they are faced with allocation decisions, our work contributes important initial insights on a robot's impact on groups.

5.1 Limitations

In interpreting the findings from our evaluation, we need to address several limitations. First, the evaluation of our algorithm was conducted in the context of an online game which opens question about the generalizability of our findings to human-robot teamwork with physically embodied robots. While we cannot say how physical robot characteristics influence perceptions of trust and fairness, our overall approach to focus on resource distribution matches a scenario that has been proposed by Jung and colleagues [34] in

the context of physical human robot collaboration. Moreover, their work highlights how a robot's distribution of resources impacts team satisfaction yet does not influence the team's performance on a task, which is consistent with the results of our collaborative Tetris game. Our work extends this previous work [34], by examining reactions to an autonomous system rather than one based on Wizard of Oz control.

In our theoretical contribution, we have assumed the same minimum allocation rate for all arms. One can easily extend the proposed algorithms and theoretical results by defining a vector of different rates for each arm. By following the same reasoning as in the proofs provided in the supplemental material, the reader can verify that compared to an oracle that satisfies the same constraints, the regret will not be worse than unconstrained UCB.

A limitation of our paper is that Algorithm 1 allows for a set of possible schedules that satisfy the minimum rate constraint, based on our choice of S and g . For instance, in the UCB-25% condition we chose to play the arm with the highest UCB bound in the second and fourth timeslot, but we could also select the first and second timeslot. In fact, given a minimum rate v there are $\frac{(1/v)!}{(1/v-K)!}$ permutations, and we have not captured the effect of different schedules within a fairness condition.

6 Conclusion

This work explored the impact of resource allocation fairness by a system on human teams. We formulated the problem of distributing resources within a team as a multi armed bandit problem and developed two algorithms that constrains the number of resources an individual may acquire. Applying our algorithm with three distinct constraint rates (25%, 33%, 50%) as independent variables, we explored team member perception's of trust and fairness through a cooperative Tetris game. Results from our user study suggest that fairness in resource allocation can influence trust of weaker performing individuals but may not have an impact on the overall performance of the team.

Our work adds a new and novel framework for studying an important yet underexplored topic: robots in human teams. We demonstrated how fairness, as experienced from individuals within a team, did not vary across the different allocation rates. Results indicate a difference in perception of fairness between weak and strong performing players in the case where the allocation more heavily favored the stronger player, while we did not observe a significant difference in performance.

As robots are becoming more and more commonplace in everyday contexts at work and at home, they increasingly face situations that involve interactions with groups or teams of people. When interacting with multiple people robots have to make decisions about resource distributions. Our work demonstrates that such decisions can be made in ways that not only take task concerns into account but also human concerns of fairness.

References

- [1] J. Stacy Adams. 1965. Inequity In Social Exchange. *Advances in Experimental Social Psychology* (1965). [https://doi.org/10.1016/S0065-2601\(08\)60108-2](https://doi.org/10.1016/S0065-2601(08)60108-2)
- [2] Rajeev Agrawal. 1995. Sample mean based index policies by $O(\log n)$ regret for the multi-armed bandit problem. *Advances in Applied Probability* 27, 4 (12 1995), 1054–1078. <https://doi.org/10.2307/1427934>

- [3] Sheldon Alexander and Marian Ruderman. 1987. The role of procedural and distributive justice in organizational behavior. *Social justice research* 1, 2 (1987), 177–198.
- [4] Peter Auer and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem*. 47 (2002), 235–256.
- [5] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh Rao, and Minoru Asada. 2016. Initiative in robot assistance during collaborative task execution. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 67–74. <https://doi.org/10.1109/HRI.2016.7451735>
- [6] Murray R Barrick, Greg L Stewart, Mitchell J Neubert, and Michael K Mount. 1998. Relating member ability and personality to work-team processes and team effectiveness. *Journal of applied psychology* 83, 3 (1998), 377.
- [7] Alan A. Benton. 1971. Productivity, distributive justice, and bargaining among children. *Journal of Personality and Social Psychology* (1971). <https://doi.org/10.1037/h0030702>
- [8] Ellen Berscheid, David Boye, and Elaine Walster. 1968. Retaliation as a means of restoring equity. *Journal of Personality and Social Psychology* (1968). <https://doi.org/10.1037/h0026817>
- [9] Sarah F Brosnan and Frans BM de Waal. 2014. Evolution of responses to (un) fairness. *Science* 346, 6207 (2014), 1251776.
- [10] Giuseppe Burtini, Jason Loeppky, and Ramon Lawrence. 2015. A Survey of Online Experiment Design with the Stochastic Multi-Armed Bandit. (2015).
- [11] Colin Camerer. 2003. *Behavioral game theory : experiments in strategic interaction*. Russell Sage Foundation, 550 pages.
- [12] Lawrence Chan, Dylan Hadfield-Menell, Siddhartha Srinivasa, and Anca Dragan. 2019. The Assistive Multi-Armed Bandit. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 354–363. <https://doi.org/10.1109/HRI.2019.8673234>
- [13] Min Chen, Stefanos Nikolaidis, Harold Soh, David Hsu, and Siddhartha Srinivasa. 2018. Planning with Trust for Human-Robot Collaboration. In *ACM/IEEE International Conference on Human-Robot Interaction*. <https://doi.org/10.1145/3171221.3171264>
- [14] Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. 2018. Learning to optimize under non-stationarity. *arXiv preprint arXiv:1810.03024* (2018).
- [15] Yochi Cohen-Charash and Paul E. Spector. 2001. The role of justice in organizations: A meta-analysis. *Organizational Behavior and Human Decision Processes* (2001). <https://doi.org/10.1006/obhd.2001.2958>
- [16] Jason A. Colquitt, Michael J. Wesson, Christopher O.L.H. Porter, Donald E. Conlon, and K. Yee Ng. 2001. Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. <https://doi.org/10.1037/0021-9010.86.3.425>
- [17] Jacob W Crandall, Mayada Oudah, Fatimah Ishowo-Oloko, Sherief Abdallah, Jean-François Bonnefon, Manuel Cebrian, Azim Shariff, Michael A Goodrich, Iyad Rahwan, et al. 2018. Cooperating with machines. *Nature communications* 9, 1 (2018), 233.
- [18] Mary J Culnan and Pamela K Armstrong. 1999. Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization science* 10, 1 (1999), 104–115.
- [19] Munjal Desai, Poornima Kaniarasu, Mikhail Medvedev, Aaron Steinfeld, and Holly Yanco. 2013. Impact of robot failures and feedback on real-time trust. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*. IEEE Press, 251–258.
- [20] Munjal Desai, Mikhail Medvedev, Marynel Vázquez, Sean McSheehy, Sofia Gadea-Omelchenko, Christian Bruggeman, Aaron Steinfeld, and Holly Yanco. 2012. Effects of changing reliability on trust of robot systems. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 73–80.
- [21] Miroslav Dudik, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang. [n. d.]. *Efficient Optimal Learning for Contextual Bandits*. Technical Report.
- [22] Raymond Fisman, Shachar Kariv, and Daniel Markovits. 2007. Individual preferences for giving. *American Economic Review* 97, 5 (2007), 1858–1876.
- [23] Aurelien Garivier. 2008. *On Upper-Confidence Bound Policies for Non-Stationary Bandit Problems Eric Moulines*. Technical Report.
- [24] Aurelien Garivier and Eric Moulines. 2011. On upper-confidence bound policies for switching bandit problems. In *International Conference on Algorithmic Learning Theory*. Springer, 174–188.
- [25] Matthew C Gombolay, Reymundo A Gutierrez, Shanelle G Clarke, Giancarlo F Sturla, and Julie A Shah. 2015. Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Autonomous Robots* 39, 3 (2015), 293–312.
- [26] Matthew C Gombolay, Cindy Huang, and Julie A Shah. 2015. Coordination of Human-Robot Teaming with Human Task Preferences. *AAAI Fall Symposium Series on AI-HRI* (2015).
- [27] Richard G. Graf and Duane Green. 1971. The equity restoring components of retaliation. *Journal of Personality* (1971). <https://doi.org/10.1111/j.1467-6494.1971.tb00064.x>
- [28] Victoria Groom and Clifford Nass. 2007. Can robots be teammates?: Benchmarks in human-robot teams. *Interaction Studies* 8, 3 (2007), 483–500.
- [29] J Richard Hackman and Greg R Oldham. 1976. Motivation through the design of work: Test of a theory. *Organizational behavior and human performance* 16, 2 (1976), 250–279.
- [30] Richard J Haier, Benjamin V Siegel Jr, Andrew MacLachlan, Eric Soderling, Stephen Lottenberg, and Monte S Buchsbaum. 1992. Regional glucose metabolic changes after learning a complex visuospatial/motor task: a positron emission tomographic study. *Brain research* 570, 1-2 (1992), 134–143.
- [31] Bradley Hayes and Brian Scassellati. 2015. Effective robot teammate behaviors for supporting sequential manipulation tasks. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6374–6380.
- [32] Chien-Ming Huang, Maya Cakmak, and Bilge Mutlu. [n. d.]. *Adaptive Coordination Strategies for Human-Robot Handovers*. ([n. d.]).
- [33] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. 2016. Fairness in learning: Classic and contextual bandits. In *Advances in Neural Information Processing Systems*. 325–333.
- [34] Malte F Jung, Dominic DiFranzo, Brett Stoll, Solace Shen, Austin Lawrence, and Houston Claire. 2018. Robot Assisted Tower Construction-A Resource Distribution Task to Study Human-Robot Collaboration and Interaction with Groups of People. *arXiv preprint arXiv:1812.09548* (2018).
- [35] Poornima Kaniarasu, Aaron Steinfeld, Munjal Desai, and Holly Yanco. 2012. Potential measures for detecting trust changes. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 241–242.
- [36] David Kirsh and Paul Maglio. 1994. On distinguishing epistemic from pragmatic action. *Cognitive science* 18, 4 (1994), 513–549.
- [37] Robert Kleinberg, Aleksandr Slivkins, and Eli Upfal. 2008. Multi-armed bandits in metric spaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*. ACM, 681–690.
- [38] T L Lai Andherbertrobins. 1985. Asymptotically Efficient Adaptive Allocation Rules*. 6 (1985), 4–22.
- [39] Tian Lan, David Kao, Mung Chiang, and Ashutosh Sabharwal. 2010. *An axiomatic theory of fairness in network resource allocation*. IEEE.
- [40] Irving M. Lane and Lawrence A. Messe. 1971. Equity and the distribution of rewards. *Journal of Personality and Social Psychology* (1971). <https://doi.org/10.1037/h0031684>
- [41] Paul A M Van Lange, Daniel Batson, Ellen De Bruin, Sander Koole, and A M Van Lange. 1999. The Pursuit of Joint Outcomes and Equality in Outcomes: An Integrative Model of Social Value Orientation. 77, 2 (1999), 337–349.
- [42] Joseph Schumpeter Lecture, Ernst Fehr, and Armin Falk. 2002. Psychological foundations of incentives. 46 (2002), 687–724.
- [43] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (2018), 2053951718756684.
- [44] Min Kyung Lee and Su Baykal. [n. d.]. *Algorithmic Mediation in Group Decisions: Fairness Perceptions of Algorithmically Mediated vs. Discussion-Based Social Division*. ([n. d.]). <https://doi.org/10.1145/2998181.2998230>
- [45] Fengjiao Li, Jia Liu, and Bo Ji. 2019. Combinatorial Sleeping Bandits with Fairness Constraints. *arXiv preprint arXiv:1901.04891* (2019).
- [46] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2012. A Contextual-Bandit Approach to Personalized News Article Recommendation. (2012). <https://arxiv.org/pdf/1003.0146.pdf>
- [47] John Lindstedt and Wayne Gray. 2013. Extreme expertise: Exploring expert behavior in Tetris. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 35.
- [48] Odalric-Ambrym Maillard, Remi Munos, and Gilles Stoltz. 2011. *A Finite-Time Analysis of Multi-armed Bandits Problems with Kullback-Leibler Divergences*. Technical Report. 497–514 pages.
- [49] Dean B Mcfarlin and Paul D Sweeney. 1992. Distributive and Procedural Justice as Predictors of Satisfaction with Personal and Organizational Outcomes. 35, 3 (1992), 626–637.
- [50] Ravi Pandya, Sandy H Huang, Dylan Hadfield-Menell, and Anca D Dragan. [n. d.]. *Human-AI Learning Performance in Multi-Armed Bandits*. ([n. d.]).
- [51] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y Narahari. 2019. Achieving Fairness in the Stochastic Multi-armed Bandit Problem. *arXiv preprint arXiv:1907.10516* (2019).
- [52] Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. 2011. Improved human-robot team performance using chaski, a human-inspired plan execution system. In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11*. <https://doi.org/10.1145/1957656.1957668>
- [53] Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. 2011. Improved human-robot team performance using chaski, a human-inspired plan execution system. In *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 29–36.
- [54] Elaine Short and Maja J Mataric. 2017. Robot moderation of a collaborative game: Towards socially assistive robotics in group interactions. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 385–390.

- [55] Pan Shu, Chen Min, Indu Bodala, Stefanos Nikolaidis, David Hsu, and Harold Soh. 2018. Human Trust in Robot Capabilities across Tasks. In *ACM/IEEE International Conference on Human-Robot Interaction*. <https://doi.org/10.1145/3173386.3177034>
- [56] Daniel P Skarlicki and Robert Folger. 1997. Retaliation in the workplace: The roles of distributive, procedural, and interactional justice. *Journal of applied Psychology* 82, 3 (1997), 434.
- [57] Adams, Stacy. 1963. Toward and Understanding of Inequity. *Journal of Abnormal Psychology* (1963).
- [58] Marynel Vázquez, Aaron Steinfeld, and Scott E Hudson. 2016. Maintaining awareness of the focus of attention of a conversation: A robot-centric reinforcement learning approach. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 36–43.
- [59] Elaine Walster, Ellen Berscheid, and G. William Walster. 1973. New directions in equity research. *Journal of Personality and Social Psychology* (1973). <https://doi.org/10.1037/h0033967>
- [60] Lai Wei and Vaibhav Srivatsva. 2018. On abruptly-changing and slowly-varying multiarmed bandit problems. In *2018 Annual American Control Conference (ACC)*. IEEE, 6291–6296.