



Full length article

The social consequences of Machine Allocation Behavior: Fairness, interpersonal perceptions and performance

Houston Claire^{a,*}, Seyun Kim^b, René F. Kizilcec^c, Malte Jung^c^a Cornell University, Department of Mechanical Engineering, Ithaca, NY, USA^b Cornell University, Department of Computer Science, Ithaca, NY, USA^c Cornell University, Department of Information Science, Ithaca, NY, USA

ARTICLE INFO

Keywords:

Fairness in artificial intelligence
Group dynamics
Artificial intelligence
Perceived fairness

ABSTRACT

Machines increasingly decide over the allocation of resources or tasks among people resulting in what we call Machine Allocation Behavior. People respond strongly to how other people or machines allocate resources. However, the implications for human relationships of algorithmic allocations of, for example, tasks among crowd workers, annual bonuses among employees, or a robot's gaze among members of a group entering a store remains unclear. We leverage a novel research paradigm to study the impact of machine allocation behavior on fairness perceptions, interpersonal perceptions, and individual performance. In a 2×3 between-subject design that manipulates how the allocation agent is presented (human vs. artificial intelligent [AI] system) and the allocation type (receiving less vs. equal vs. more resources), we find that group members who receive more resources perceive their counterpart as less dominant when the allocation originates from an AI as opposed to a human. Our findings have implications on our understanding of the impact of machine allocation behavior on interpersonal dynamics and on the way in which we understand human responses towards this type of machine behavior.

1. Introduction

Algorithms shape court decisions (Campbell, 2020), whom we hire (Kodiyar, 2019), who gets approved for a loan (Mukerjee, Biswas, Deb, & Mathur, 2002), or who is admitted into college (Baig, 2018). Recognizing the importance and urgency to build understanding about the increasing influence that machines have over our lives as individuals and societies, Rahwan and colleagues proposed Machine Behavior as a field of study that “is concerned with the scientific study of intelligent machines, not as engineering artifacts, but as a class of actors with particular behavioral patterns and ecology. This field overlaps with, but is distinct from, computer science and robotics. It treats machine behavior empirically” (Rahwan et al., 2019).

A particularly ubiquitous type of machine behavior is what we call *Machine Allocation Behavior*. By machine allocation behavior, we refer to the overt behavior that results from a machine's decisions about the allocation of something of value (e.g., a resource or task) among people. This includes a humanoid robot turning its gaze towards one person in a group of people, an industrial robot arm handing over a tool to a team of maintenance workers, or the system message of a human-resources application that indicates which worker in an organization should get a raise. We are interested in this kind of behavior for

two reasons. First, allocation decisions are increasingly handed over to machines, making machine allocation behavior a phenomenon that more and more people encounter. For example, machines decide how donations are distributed among non-profit organizations (Lee, Kim, & Lizarondo, 2017), decide who receives support in a search and rescue scenario (Brandao, Jirotko, Webb, & Luff, 2020), determine how workload is distributed (Chang, Pope, Short, & Thomaz, 2020), or decide which ads to distribute to users (Singh, Nanavati, Kar, & Gupta, 2022). Second, while a growing body of research has studied the impact of this behavior on people and has developed novel approaches for allocations, our understanding about interpersonal consequences is limited. For example, prior work has examined how people respond to the algorithmic allocation of tasks among crowd workers (Yu et al., 2019), maintenance workers (Hassan, O'Riain, & Curry, 2013), and gig workers (Hodson, 2014). This work provides important insights into responses to algorithmic allocations, but it has predominantly focused on individual consequences such as an individual's perception of the fairness of an allocation behavior. In contrast, we know little about the social and interpersonal consequences of algorithmic allocation decisions. We specifically lack an understanding of how machine behaviors impact the way we interact with and perceive other people. Understanding the interpersonal consequences of allocation decisions made

* Correspondence to: Cornell University, 107 Hoy Rd, Ithaca, NY 14853, USA.
E-mail address: hbc35@cornell.edu (H. Claire).

by machines is important because the quality of our social interactions with others determines much of our lives including how we form and maintain relationships (Gottman & Levenson, 1992; Levenson & Gottman, 1985), and what we are able to achieve with others (Barrick, Stewart, Neubert, & Mount, 1998; Gottman, 2014; Jehn, 1995, 1997; Jung, 2016).

In this study, we explore the impact of machine allocation behavior on interpersonal perceptions, fairness perceptions, and task performance. We conducted a large online experiment (N = 232 participants) to understand how the degree of resources that a person receives relative to a collaborator (receiving less vs. receiving equal vs. receiving more resources) interacts with the type of allocation agent (human vs. AI system) to impact social judgments, fairness perceptions, and task outcomes. Our findings show that receiving less resources than a counterpart causes individuals to experience a larger interpersonal distance to their partner and allocator, have a lower performance, and have lower perception of fairness. Additionally, we show that the agency of the allocator changes how dominant individuals perceive their human counterparts.

This work makes three primary contributions to the literature: first, we introduce the concept of machine allocation behavior alongside a method to study its impact on people, and their interactions with each other. Second, while we find no differences in the impact of the allocator type (human vs. machine) on perceptions of the allocator we find evidence for the impact of the allocator type on interpersonal perceptions. Specifically, we find that individuals perceived their partner as more dominant when they received more resources from an AI system compared to a human. Finally, our work contributes a novel perspective for conceptualizing algorithmic fairness as a dynamic phenomenon with a temporal dimension.

2. Theoretical background and research questions

We situate our work in the context of machine behavior and propose the concept of machine allocation behavior. This type of behavior draws from prior research into the social consequences of distributing resources among humans.

2.1. Machine behavior

The study of machine behavior emerged as artificially intelligent machines increased in prevalence across workplaces and homes. Machines use engineered behaviors to perform complex tasks; for example, they determine who receives a bank loan (Yang, Li, Ji, & Xu, 2001), select the types of people shown on dating sites (Ma & Gajos, 2022), or even perform common household tasks such as folding clothes (Yang et al., 2016). This development has raised concerns about the potential societal impacts of these behaviors. Machine behavior as a discipline examines the relationship between machines and agents in an environment (Rahwan et al., 2019). Scholarship in this discipline seeks to draw conclusions about machine behaviors through observing different interactions in a similar way to how humans and animals are studied. To date, there are few studies highlighting concrete examples of the effects of specific machine behaviors. We are interested in a specific type of machine behavior that focuses on the social consequences of the allocation of resources by a machine, which we call machine allocation behavior.

2.2. Resource allocation behavior

Various social systems rely on the allocation of resources and other things of value to achieve their goals. Managers make decisions on who receives a raise, group leaders make decisions on how to distribute tasks, etc. The observable aspects of these allocation decisions can be called allocation behaviors. For example, focusing on resource allocation, Langholtz and colleagues (Langholtz, Marty, Ball, & Nolan,

2002) argued that "resource-allocation decisions are any decisions in which people make judgments about how they will allocate resources. Resource-allocation behavior is the outward, observable behavior in which people act upon their allocation decisions". Allocation behavior shapes task outcomes and relationships through determining the type of tasks group members get, reinforcing and strengthening behaviors that contribute to the solution of problems, or improving loyalty to a group (Leventhal, 1976).

Any discussion on the allocation of resources and its effect on human behavior is incomplete without a discussion of fairness. Sparking what would be years of research on fairness in the workplace, John Adams proposed equity theory in 1963, which describes individuals as motivated by fairness and willing to invest resources to address an inequity (Adams, 1965). It further motivated research into the different factors allocators account for when making their decisions. Allocation norms were introduced as the social rule that specifies criteria that define certain distributions of rewards and resources as fair and just (Leventhal, 1976). For example, an allocator may follow an equity norm of equality and ensure that every member in a group receives an equal level of resources. Or an allocator may follow an equity norm ensuring that every individual receives resources in proportion to their input. In an equitable allocation norm, research has shown that allocating high reward to good performers and low reward to poor performers facilitates productivity (Collins et al., 1964; Homans, 1961). Yet this type of norm can be counterproductive in situations where individuals' contributions are difficult to assess or when a task requires a high degree of cooperation (Lawler, 1971). Along a similar vein, studies exploring the equality norm showed evidence of increased solidarity and harmony amongst group members (Deutsch, 1953; Julian & Perry, 1967; Smith & Cook, 1973). These findings highlight the complexity of deciding on the appropriate allocation norm to apply as well as its context-dependent nature while hinting at the effects it can have on interpersonal behaviors.

In particular, "unfair" allocation decisions have vast social and behavioral consequences. Findings from economic game theory, such as experiments with the ultimatum game, show that individuals will reject unfair offers even if it has negative consequences to their own payoff (Stahl, 1977; Yamagishi et al., 2012). Similar types of responses are seen in the workplace where individuals will work against the interest of the team or group if they feel unfairly treated (Everton, Jolton, & Mastrangelo, 2007). For instance, workers have been found to sabotage customers even at the cost of their own payout when they feel that they have been unfairly treated (Skarlicki, Van Jaarsveld, & Walker, 2008). Other studies note a connection between workplace aggression and perceptions of unfairness, such as Folger et al.'s study of how unfair work layoffs can lead to aggression in the workplace (Folger & Baron, 1996). Many of the studies related to fairness in the workplace have relied on analyzing fairness through how fair an allocation outcome was perceived (distributive fairness) (Cook & Hegtvedt, 1983; Nozick, 1973) or how fair the allocation process was perceived (procedural fairness) (Konovsky, 2000).

The increased application of artificial intelligence systems as allocators has raised concerns about how their decisions will affect human responses and behavior. The work exploring differences in how people respond to unfair allocations from an AI system versus a human is limited and the results show to be contradicting. On one hand, several studies highlight the misconceptions that people have about the objectivity of AI and how this misconception can shape their perception of its decision (Dijkstra, Liebrand, & Timminga, 1998; Yeomans, Shah, Mullainathan, & Kleinberg, 2019). On the other hand, Lee et al. found differences in fairness perceptions depend on the context (Lee, 2018).

2.3. Machine allocation behavior

Analogous to defining resource allocation behavior as the observable behavior that results from resource allocation decisions that people

make (Langholtz et al., 2002), we use the term Machine Allocation Behavior to refer to the observable machine behavior (Rahwan et al., 2019) that results from the allocation decisions that machines make. Machine allocation behavior becomes increasingly prevalent in the allocation of rides (Lee, Kusbit, Metsky, & Dabbish, 2015), hospital supplies (Ljungblad, Kotrbova, Jacobsson, Cramer, & Niechwiadowicz, 2012), design suggestions (Zhang, Raina, Cagan, & McComb, 2021), tasks (Gombolay, Gutierrez, Clarke, Sturla, & Shah, 2015), etc.

The increased application of artificial intelligence systems as allocators has raised concerns about how their decisions will affect human responses and behavior. For example, a recent line of research has begun exploring the impact of machine allocation behavior on people. Studies by Lee and Rosenblatt (Lee et al., 2015; Rosenblatt & Stark, 2016) found that the way in which an algorithm dispatches the next ride among ride-share drivers influenced how they behave and approach their role. Additionally, remote gig workers can suffer from social isolation and overwork as a result of different forms of algorithmic decisions (i.e. rating and ranking features) within the platform (Wood, Graham, Lehdonvirta, & Hjorth, 2019). A study exploring how the inclusion of an AI system in a call center influenced employee's view of the workplace revealed feelings of psychological distress and even found that they were considering leaving the workplace altogether due to the perceived threat of the AI taking over their jobs (Presbitero & Teng-Calleja, 2022).

Fairness has received a considerable amount of attention in the discussion revolving how machines should allocate resources. People show to be attuned to the way resources are allocated by machines and make fairness judgments based on the level of resources they receive in comparison to others (Bartol & Srivastava, 2002; Dulebohn & Martocchio, 1998; Leventhal, 1976). On one hand, efforts to leverage human notions of fairness has led investigations into human responses to allocation decision from a machine (Christin, 2017; Hohenstein & Jung, 2020; Saxena, Badillo-Urquiola, Wisniewski, & Guha, 2020; Strohkorb Sebo, Traeger, Jung, & Scassellati, 2018; Završnik, 2021) while on the other hand, a separate stream has focused on formalizing definitions and creating metrics of fairness across different contexts (Barocas, Hardt, & Narayanan, 2018; Mehrabi, Morstatter, Saxena, Lerman, & Galstyan, 2021; Mitchell, Potash, Barocas, D'Amour, & Lum, 2021).

However, there is a lack of consensus on how the agency of the allocator shapes human response to allocation decisions. On the one hand, research on the computers as social actors theory suggests that there should be no difference between a human or machine allocator as people have a tendency to apply social heuristics for interactions with people mindlessly to machines (Reeves & Nass, 1996). On the other hand, there is growing evidence that humans have contrasting fairness judgments in cases where an allocation decision originates from a machine as opposed to a human. For example, Lee (2018) suggest that perceptions of human algorithmic decision-making depend on the nature of the task. For mechanical tasks, participants perceived human and algorithmic decisions as equally fair. Hence, human managers' decisions and perception of fairness was judged based on the "authority of the manager's position". However, for human tasks, participants considered human decision-makers fairer than algorithmic decisions. Algorithmic allocators were considered to be "less intuitive and subjective" for human tasks. Moreover, Lee and Rich (2021) suggests that cultural mistrust from certain demographics influence perception of fairness in human versus machine decision-making. Complicating current understanding are findings from several studies that highlight the misconceptions that people have about the objectivity of AI and how this misconception can shape their perception of its decision (Dijkstra et al., 1998; Yeomans et al., 2019), or how context can shape perceptions of machine behavior (Lee, 2018). Despite all this work there is surprisingly little understanding about how allocation behavior – whether it originates from people or machines – impacts not only individuals but their interactions and relationships with each other.

2.4. Research questions

The goal of our research is to compare human and machine allocation behavior in its impact on peoples' relationships and interpersonal perceptions (Rahwan et al., 2019). Prior work suggests that people respond to resource allocation decisions based on their fairness judgments (Skarlicki et al., 2008; Yamagishi et al., 2012). Yet there is conflicting evidence as to the degree to which this response will change if the allocation decision comes from a machine as opposed to a human (Lee, 2018; Lee & Rich, 2021). Furthermore, our understanding of the social consequences of machine allocation behavior is limited to observations of changes in perceptions towards a machine. We know little about how these allocation decisions actively shape the interpersonal relationships between the humans it interacts with. Finally, how machines affect performance is a topic of ongoing discussion. While many studies highlight improvements across different metrics of performance (Shirado & Christakis, 2017; Song et al., 2022; Strohkorb et al., 2016), others have demonstrated that the application of machines in teams can promote behaviors that hinder outcomes (Bansal et al., 2019; Zhang et al., 2021). We thus use the following research questions to guide our study of machine allocation behavior:

RQ1: *How does allocation behavior shape social judgments of different group members and how are those judgments shaped by the type of allocation agent?*

RQ2: *How does a machine's allocation behavior shape perceptions of fairness?*

RQ3: *How does machine allocation behavior shape task performance?*

3. Method

To examine our research questions, we conducted a between-participants study (N=232 participants) that manipulated the type of allocation agent (human vs. AI system) and the type of allocation behavior (receive more, vs. equal, vs. less turns) resulting in overall six conditions (see Fig. 2).

3.1. Participants

We recruited 232 US based MTurk participants that had overall HIT approval ratings of at least 80%. Participants were compensated \$4.00 dollars after completing the task. Of the 232 participants, 88 were female, 142 were male, and 2 selected other. Fig. 2b shows the distribution of individuals and gender across the different conditions. The average age of the participants was 40 years old with 27 of them having played less than one hour of Tetris in their life, 120 have played 1–10 hours, and 85 played more than 10 hours.

3.2. Materials: Co-tetris platform

To conduct our study, we developed a research platform that we call Co-Tetris. Co-Tetris builds on a platform developed by Claire and colleagues (Claire, Chen, Modi, Jung, & Nikolaidis), that allows multiple people to collaborate in playing a Tetris game.

Co-Tetris follows the rules and gameplay of a standard Tetris game where individuals have to manipulate falling geometric blocks in order to stack them without creating gaps. What sets Co-Tetris apart from a standard Tetris game is that it allows two people to work together to play the game instead of one. Only one player at a time has control over the current turn (i.e. the current block falling down that needs to be placed) and an "allocator" decides over the allocation of each turn among players.

We chose Tetris for our research for three specific reasons. First, Tetris has proven to be a game that allowed research to gain fundamental insights about human cognition and social behavior that

have implications far beyond the specific game context (e.g. Haier et al. (1992), Kirsh and Maglio (1994), Lindstedt and Gray (2013)). A prominent example is Kirsh and Maglio's (Kirsh & Maglio, 1994) discovery of the concept of epistemic action. Other studies have used Tetris to explore general learning phenomena such as expertise (Lindstedt & Gray, 2013), memory (Holmes, James, Coope-Bate, & Deeprose, 2009), and teaching (Knox, Glass, Love, Maddox, & Stone, 2012). Additionally, Claire and colleagues (Claire et al.) used a Tetris game to demonstrate the performance of a novel allocation algorithm. Second, the Tetris platform can model a wide range of task characteristics such as sequential allocation behaviors, different difficulty levels, learning opportunities, and objective performance metrics. It offers a clear and transparent method to capture various metrics, such as player behavior and performance, while completing a challenging task that is dependent on expertise (Lindstedt & Gray, 2013). Third, using a game platform like Tetris allows us to collect data at scale.

Fig. 3a on the left shows the Co-Tetris screen for the active player who has the current turn. The game screen provides information about the current group score, who has the active turn, the keyboard buttons that controls the blocks, and a visual of what the next block type will be. Fig. 3a on the right shows the screen for the inactive player who is waiting to be assigned a turn by the allocator. The screen shows the same elements but is grayed out to highlight that the player currently has no control over the game. After a player's turn is completed by placing a Tetris block, a pop up is displayed that shows the allocator's decision about which player receives the next turn. The allocator can either be human (Fig. 3b) or AI system (Fig. 3c).

Players complete several rounds of the game. A round ends when players are no longer able to keep blocks from piling up to the top of the game screen. After a set time a button to stop playing appears allowing participants to move on to the post task survey. However, the game leaves it to participants to decide if they want to play longer.

3.3. Procedure

After accepting the Human Intelligence Task (HIT) on MTurk and providing informed consent, participants were directed to a Qualtrics Survey that provided them with a set of instructions informing them about the overall game play and rules. Participants were instructed to complete a collaborative game of Tetris with another MTurk player and an allocation agent. Based on the randomly assigned allocation agent condition, participants received instructions that introduced the allocation agent either as another MTurk worker (allocation agent: human) or as an algorithm (allocation agent: machine). To test whether participants paid careful attention to the instructions, participants were asked to answer manipulation check questions on gameplay and the agency of the allocator. Participants had two opportunities to read the instructions and correctly answer the questions before being barred from attempting the experiment.

After completing the instructions portion of the survey, participants were given a link that directed them to the Co-Tetris game. After clicking on the link, participants were asked to wait to be matched with another player before starting the game. Players waited for 5 minutes to be matched before they were removed from the waiting room and compensated appropriately for their time. Based on the randomly assigned allocation type condition (more vs. less vs. equal), groups of two participants each were assigned to a group where one MTurk player would receive a cumulative 10 percent of the total blocks (allocation type: less) while the other would receive 90 percent of the total blocks (allocation type: more) or to a group where both MTurk players received an equal amount of the total cumulative blocks (allocation type: equal).

Once participants were matched with another MTurk player, they were asked to play the Co-Tetris game for a minimum of 5 minutes before a "Quit Game" button appeared allowing participants to leave the game if they wanted to. The platform waited for one player to

quit before automatically directing both players to the next phase. At the onset of the game and between each turn, participants saw one of two possible images of the allocator dependent on the study condition (an image symbolizing a human for the human allocator condition, Fig. 3b, or an image symbolizing an algorithm for the machine allocator condition, Fig. 3c). Unbeknownst to both players, the allocator is always an algorithm that distributes control of the turns in accordance with the study conditions.

The allocation algorithm used for players assigned to the more and less condition ensured that any player assigned to the more condition will receive a total of 90 percent of the Tetris blocks while the other player in the less condition receives a total of 10 percent of the blocks over the course of the game. To fulfill this, the randomly selected player in the less condition was set to receive a Tetris block once every ten turns. Across each batch of ten turns, the algorithm would randomly select one turn where the player in the less condition would receive a block. The algorithm for the equal condition ensures an equal distribution of blocks between players by randomly selecting a player every turn to obtain control of the falling blocks. To guarantee that an individual receives a total of 50% of the blocks over the course of the game, the algorithm checks that no player is selected more than three times before receiving a Tetris block (Fig. 1).

After completing the Co-Tetris game portion of the experiment, participants were directed back to the Qualtrics Survey and asked to complete the post-task survey measures (see Fig. 4).

3.4. Measures

To answer our research questions, we measured interpersonal perceptions, fairness perceptions, and task performance.

- *Interpersonal Perceptions*: To operationalize the interpersonal perceptions measure, we used the Inclusion of Others in the Self (IOS) scale (Aron, Aron, & Smollan, 1992) and the Revised Interpersonal Adjective Scale (IAS-R) (Wiggins, Trapnell, & Phillips, 1988). The IOS scale presents participants with seven images of increasingly overlapping circles that are labeled "self" and "other". The measures ask participants to rate which images best represent their relationship on a Likert scale from (1) very distant with no overlap between circles, to (7) very close with a high degree of overlap between circles. The IOS scale was used to measure perceptions between players as well as perceptions of the allocation agent. The Revised Interpersonal Adjective Scale (IAS-R) (Wiggins et al., 1988) asks participants to rate their group partner across 32 trait adjectives that sample the interpersonal dimensions of dominance and affiliation. Each adjective was rated on a 7-point Likert scale ranging from "Extremely inaccurate" (1) to "Extremely accurate" (7) ($\alpha = 0.88$). Ratings are then aggregated to an overall dominance and affiliation score.
- *Fairness*: We operationalized our fairness measure through survey scales for distributive and procedural fairness. We used four items from Colquitt's Distributive Justice Scale (Colquitt, 2001). This scale asks participants to indicate the extent to which they have received a fair distribution of Tetris blocks. Each item was rated on a 5-point Likert scale ranging from "To a small extent" (1) to "To a large extent" (5) ($\alpha = 0.90$). We used five items from Colquitt's Procedural Justice Scale (Colquitt, 2001). This scale measured the extent to which participant's found the procedures the allocator used to decide who receives a Tetris block per turn fair. Each item was rated on a 5-point Likert scale ranging from "To a small extent" (1) to "To a large extent" (5) ($\alpha = 0.83$).
- *Task Performance*: We measured task performance at the individual and group level. At the individual level, the measure was operationalized by taking each individual's score and dividing it by the number of turns a player had control over. We additionally captured the amount of time and the number of turns a player

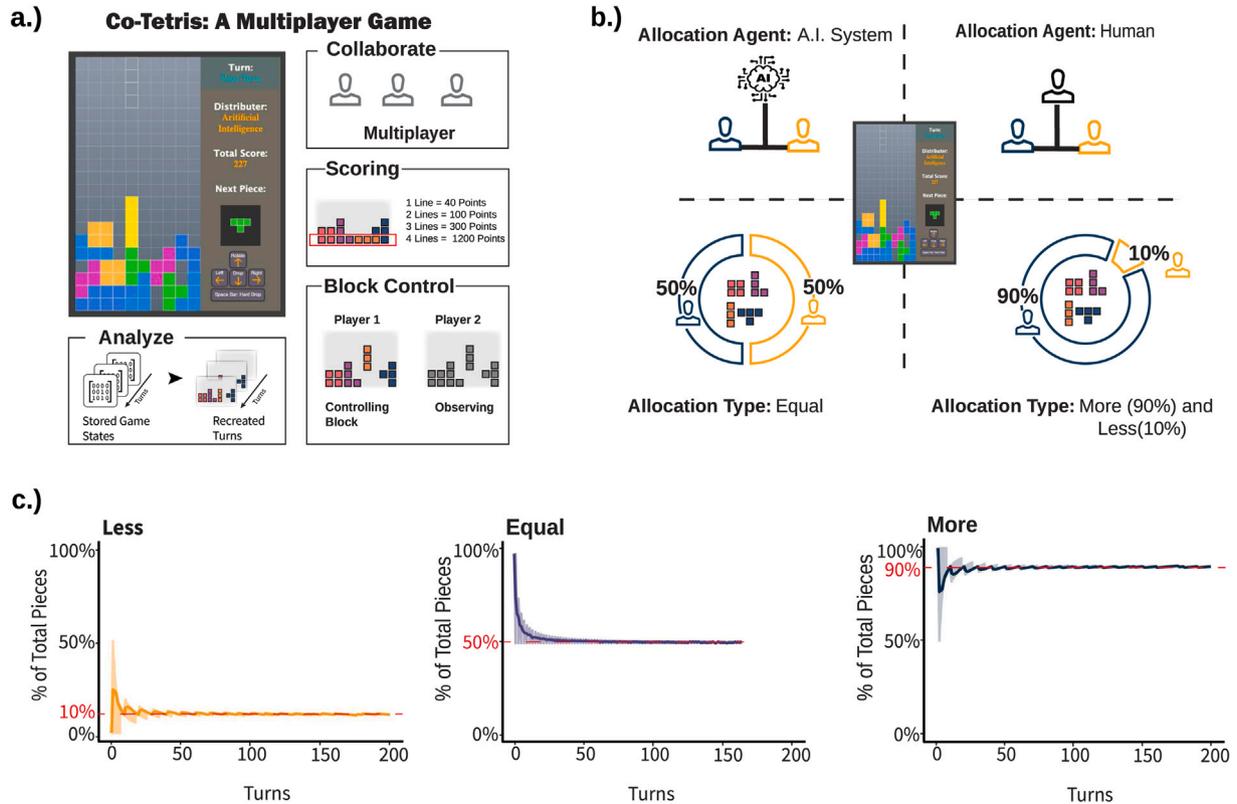


Fig. 1. (a) The Co-Tetris platform is a platform that allows two players to collaboratively play a game of tetris with an allocator deciding the allocation of the next falling block. All user behaviors are logged, allowing fine grained analysis of individual behavior. (b) A visual representation of the 2 Allocation Agent (AI system vs. Human) X 3 Allocation Type (Less vs. Equal vs. More) study design. (c) The algorithm used for each condition ensured participants in the more, equal, less condition received an accumulated sum of 90%, 50%, 10%, respectively, of all the Tetris blocks over the course of the game. Lighter lines represent the different games during this experiment while the bold line is the average across all the individuals in that condition.

had access to the falling blocks throughout the game. At the group level, the measure was operationalized by taking the overall score that both players achieved together at the end of the game.

4. Results

We report findings on interpersonal perceptions, fairness perceptions, and task performance. See Fig. 5 and Fig. 6 for visual plots of the results.

4.1. Manipulation checks

To confirm that participants were aware of the level of resources they received, we asked the question, “Around what percentage of the Tetris blocks did you control?”. Participants chose from a list of percentages that increased from 10 to 100 in increments of 10. A linear mixed effects model with the allocation type as the independent variable confirmed that a participant’s perception of how many resources they received was statistically different depending on which condition they were in. Results showed that participants who received less resources believed their share of blocks was significantly lower ($M = 23.5, SD = 20.1$) than individuals who received an equal share of blocks ($M = 55, SD = 17.8, b = 28.6, p < 0.001$) or more blocks ($M = 70.2, SD = 16.2, b = 47.0, p < 0.001$). Similarly, individuals who received an equal amount of blocks believed that they received a significantly lower number of blocks than individuals who received more blocks ($b = 18.4.0, p < 0.001$).

4.2. Interpersonal perceptions

To recall, our first research question addressed how allocation behavior (more vs. less vs. equal) shapes social judgments and how those judgments are shaped by the type of allocation agent (machine vs. human). To account for the grouping of participants in the study design, we report results from a linear mixed-effects regression (modeling the group identifier as a random intercept) that models IOS for the self and the allocator (see Table 1, model 1). Our results show a significant difference in ratings from participants who received less Tetris blocks and participants who received an equal number ($b = 0.64, p = 0.05$) or more Tetris blocks ($b = 1.16, p = 0.003$). Participants who received less blocks ($M = 2.56, SD = 1.96$) perceived a larger interpersonal distance to the allocator compared to those individuals who received an equal number of blocks ($M = 3.20, SD = 2.06$) or more blocks ($M = 3.73, SD = 1.91$). Additionally, results on IOS: self and human partner showed a significant difference in ratings between individuals who received less blocks and those who received an equal number of blocks ($b = 0.86, p = 0.009$) and those who received more blocks ($b = 0.75, p = 0.05$). Participants who received less blocks ($M = 2.73, SD = 1.83$) perceived a larger interpersonal distance to the other player compared to those individuals who received an equal number of blocks ($M = 3.57, SD = 2.02$) or more blocks ($M = 3.47, SD = 2.17$).

Results show an interaction effect (Table 1, model 3) between the allocation agent and receiving more blocks on perceptions of dominance ($b = -2.81, p = 0.05$). Individuals who received more blocks from a human ($M = 1.28, SD = 3.94$) saw their partner as less dominant than individuals who received more blocks from an AI system ($M = -1.38, SD = 3.92$). Of the covariates, the amount of turns an individual was an active player predicted perceptions of dominance ($b = 0.03, p = 0.04$).

Table 1
Results of Linear Mixed-effects Regressions with Groups Modeled as Random Intercept.

Outcome	(1) IOS1	(2) IOS2	(3) Dominance	(4) Affiliation	(5) Dist. Fair	(6) Proc. Fair	(7) Ind. Perf.
Allocation Type: Equal	0.64* (0.32)	0.86** (0.32)	-0.15 (0.59)	-0.15 (0.59)	0.85*** (0.17)	0.62*** (0.15)	228*** (101)
Allocation Type: More	1.16** (0.38)	0.75* (0.35)	-0.15 (0.70)	-0.15 (0.70)	1.02*** (0.18)	0.67*** (0.17)	754*** (102)
Allocation Agent: Human	0.08 (0.28)	0.12 (0.26)	-0.50 (0.51)	-0.50 (0.55)	-0.06 (0.15)	-0.15 (0.13)	-38.2 (93.3)
Agent: Human × Type: Equal	-0.12 (0.36)	-0.05 (0.36)	0.40 (0.66)	1.0 (0.71)	0.00 (0.19)	0.15 (0.17)	92.2 (120)
Agent: Human × Type: More	-0.47 (0.54)	-0.04 (0.54)	-2.81* (0.98)	-0.74 (0.95)	-0.09 (0.28)	-0.13 (0.25)	-53.4 (162)
Active Time (mean-centered)	0.02 (0.010)	0.01 (0.10)	-0.34 (0.19)	-0.05 (0.18)	0.02 (0.05)	0.01 (0.05)	82.3*** (31.1)
Active Turns (mean-centered)	0.00 (0.00)	0.00 (0.00)	0.03* (0.01)	0.01 (0.01)	0.00 (0.00)	0.00 (0.00)	-6.76*** (2.59)
(Intercept)	2.75 (0.38)	2.56 (0.38)	0.14 (0.69)	2.06 (0.67)	2.99 (0.19)	3.01 (0.18)	67.7 (113.5)
AIC	1008	1010	1275	1254	710	670	3554
R ²	0.03	0.01	0.06	0.02	0.13	0.09	0.20
Num. Groups	116	116	116	116	116	116	116
Num. Observations	232	232	232	232	232	232	232

Note: IOS1 stands for IOS for the Self and Allocator; IOS2 stands for the Self and Human Partner. The reference group is Allocation Type Less and Allocation Agent AI. Standard errors are in parenthesis. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

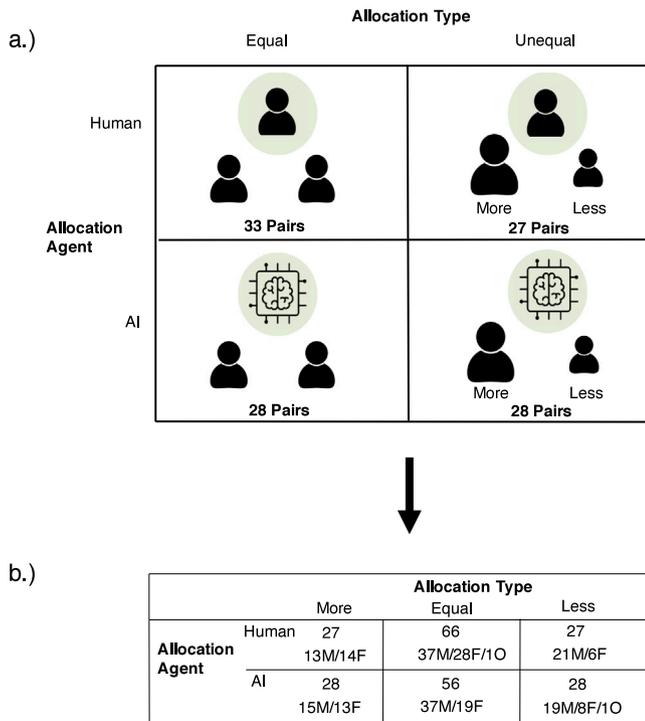


Fig. 2. (a) From a group level perspective, our study had a total of 116 pairs of MTurk workers that could be placed in a 2 Allocation Agent (AI vs Human) × 2 Allocation Type (Unequal vs. Equal) distribution. In the Unequal Allocation Type, one participant was randomly selected to receive a cumulative total of 90% of the total blocks by the end of the study. (b) We analyzed most our findings at the individual level which leveraged a 2 Allocation Agent (AI vs Human) × 3 Allocation Type (More vs. Equal vs. Less) study design.

4.3. Fairness perceptions

Our second research question sought to explore how machine allocation behavior can shape perceptions of fairness. Results based on linear mixed-effects regression shows models for distributive fairness (see Table 1, model 5). Our results shows a significant difference in

ratings from participants who received less Tetris blocks and participants who received an equal number of blocks ($b = 0.85, p < 0.001$) and those who received more blocks ($b = 1.02, p < 0.001$). Individuals who received less blocks ($M = 2.81, SD = 1.40$) felt that the distribution of blocks was less fair compared to individuals who received an equal number ($M = 3.66, SD = 0.89$) or more blocks ($M = 3.84, SD = 0.87$). Similarly, results on procedural fairness (Table 1, model 6) shows a significant difference in ratings from participants who received less Tetris blocks and participants who received an equal number of blocks ($b = 0.62, p < 0.001$) and those who received more blocks ($b = 0.67, p < 0.001$). Individuals who received less blocks ($M = 2.80, SD = 1.17$) felt that the process used to decide over the allocation of blocks was less fair compared to individuals who received an equal number ($M = 3.41, SD = 0.88$) or more blocks ($M = 3.47, SD = 0.81$).

4.4. Task performance

With our next research question, we sought to investigate how machine allocation can shape task performance. Results based on linear mixed-effects regression shows models for individual performance (see Table 1). Our results shows a significant difference in performance from participants who received less Tetris blocks and participants who received an equal number of blocks ($b = 228, p < 0.001$) and those who received more blocks ($b = 754, p < 0.001$). Participants who received less blocks ($M = 7.3, SD = 12.2$) performed worse than those who received an equal number of blocks ($M = 10.1, SD = 10.7$) and those who received more blocks ($M = 27.8, SD = 49.9$). Both of the covariates, the amount of time an individual was an active player ($b = 82.3, p < 0.001$) and the number of turns an individual was an active player ($b = -6.76, p < 0.001$), predicted individual performance.

Investigating performance at the group level (see Table 2), results revealed a significant difference between teams where Tetris blocks were distributed equally and teams where blocks were distributed unequally ($b = 193.9, p < 0.001$). Groups where blocks were distributed unequally ($M = 896, SD = 602$) performed better than groups where resources were distributed unequally ($M = 548, SD = 602$).

5. Discussion

We examined the social consequences of machine allocation behavior. We found that it impacts interpersonal relationships. In particular, we showed that the way in which a machine allocates resources among individuals affects the way they perceive one another.

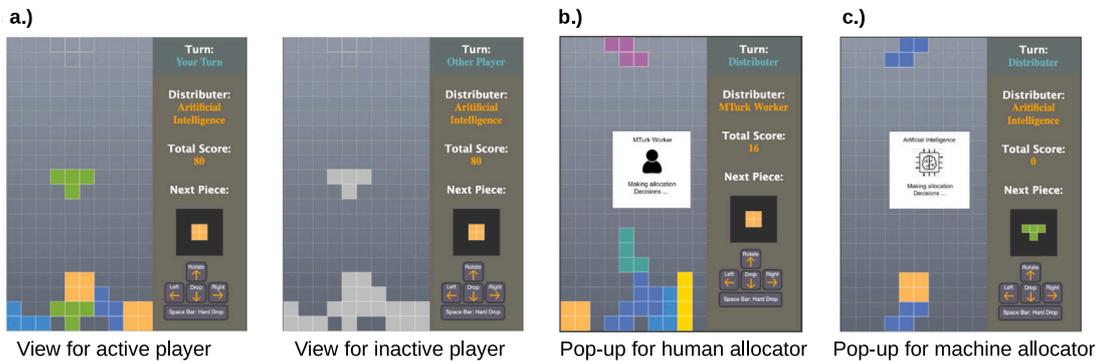


Fig. 3. (a) Left: Example of the Co-Tetris screen the active player who has control over the current turn. Right: Example of the Co-Tetris screen for the inactive player who does not have control over the current turn. The blocks are grayed out to make the transition in turns immediately visible, (b) pop-up image displayed in between turns indicating that the allocation of the next turn is decided by a human. The pop up is shown to both the active and the inactive player, (c) Pop-up image displayed in between turns indicating that the allocation of the next turn is decided by a machine.

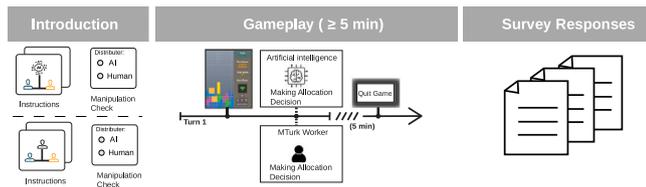


Fig. 4. An overview of the steps that each participant experienced during our experiment.

Table 2
Results of Linear Regression for Team Performance.

Outcome	Team Performance
Allocation Type: Unequal	193.9*** (174.8)
Allocation Agent: Human	-225.6 (168.1)
Agent: Human × Type: Unequal	291.8 (243.7)
(Intercept)	670.0
AIC	3833.4
R ²	0.03
Num. Observations	116

Note: Standard errors are in parenthesis. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.01$. Allocation Type = Equal and Allocation Agent = AI were used as reference levels.

5.1. Interpersonal perceptions

We found that the level of resources an individual receives relative to others affects interpersonal perceptions. This finding has implications for our growing understanding about the integration of machines in teams and organizations (e.g. Bailey, Faraj, Hinds, Leonardi, and von Krogh (2022), Jung and Hinds (2018)).

Research into groups and teams has linked the types of interpersonal relationships people form to important outcomes in the workplace. For instance, positive interpersonal relationships have been shown to reduce stress (Cohen & Wills, 1985), increase communication (Lee & Doran, 2017), and increase job satisfaction (Liden, Wayne, & Sparrowe, 2000). Furthermore, the inclusion of machines in groups has demonstrated that different machine behaviors, such as machine allocation behavior, can affect the social dynamics between group members (Jung et al., 2020; Mutlu, Shiwa, Kanda, Ishiguro, & Hagita, 2009; Sebo, Traeger, Jung, & Scassellati, 2018). We expand on these studies by showing that the degree of resources an individual receives affected the perceived interpersonal distance to their gaming partner as well as the allocator. This observed difference in interpersonal perceptions based on the level of resources received further expands on

the type of influence that machine allocation behavior has on group dynamics (Strohkorb Sebo et al., 2018; Tennent, Shen, & Jung, 2019; Terzioğlu, Mutlu, & Şahin, 2020; Traeger, Sebo, Jung, Scassellati, & Christakis, 2020). Understanding how machine allocation behavior can shape the dynamics of the collective will be important to the overall productivity and success of a group. Aside from improved work experiences, various studies highlight how strong interpersonal relationships affect group outcomes (Jehn & Shah, 1997; Liden et al., 2000; Skarlicki & Folger, 1997).

Additionally, we found that perceptions of dominance towards the group members who received less resources changed depending on the agency of the allocator. We saw that those who receive more resources viewed their game partner as less dominant when the allocation originates from an AI system as opposed to a human. Hohenstein et al. showed that machines can alter the way in which people evaluate one another by leveraging a study exploring the use of AI system smart replies in a conversation (Hohenstein et al., 2021). In contrast, Mieczkowski et al. showed no changes in how individuals perceived one another when an AI system generated smart reply was used during a collaborative game (Mieczkowski, Hancock, Naaman, Jung, & Hohenstein, 2021). Our finding aligns with the former and expands on these studies by showing differences in interpersonal perceptions depending on the allocator agency within a collaborative game. When individuals perceive that their partner is receiving less resources from a machine, they may believe that the machine is acting in a rational and objective manner (Dijkstra et al., 1998) and form judgments about the personal characteristics of their group partners. Believing that a machine is making decisions purely off objective metrics (i.e. performance) could raise suspicion that the partner receiving less resources is less capable of completing a task and change the way they are perceived.

5.2. Fairness

Our results show that individuals receiving less resources rated perceptions of fairness lower compared to those who received more or an equal level of resources. Our findings have implications for our understanding of algorithmic fairness.

Research on algorithmic fairness seeks to determine the relevant individual factors that shape individual judgments on a machine's allocation decision. Our results align with the results that illustrate the amount of resources a machine allots to an individual relative to others has a direct influence on perceptions of fairness (Christin, 2017; Saxena et al., 2020; Wang, Harper, & Zhu, 2020; Završnik, 2021).

We did not find differences in perceptions of fairness when it came from an AI system versus a human. Although this result contradicts the line of research that shows humans perceive decisions differently when it originates from a human versus a machine (Dietvorst, Simmons, & Massey, 2015; Dijkstra et al., 1998; Lee, 2018), it aligns and expands

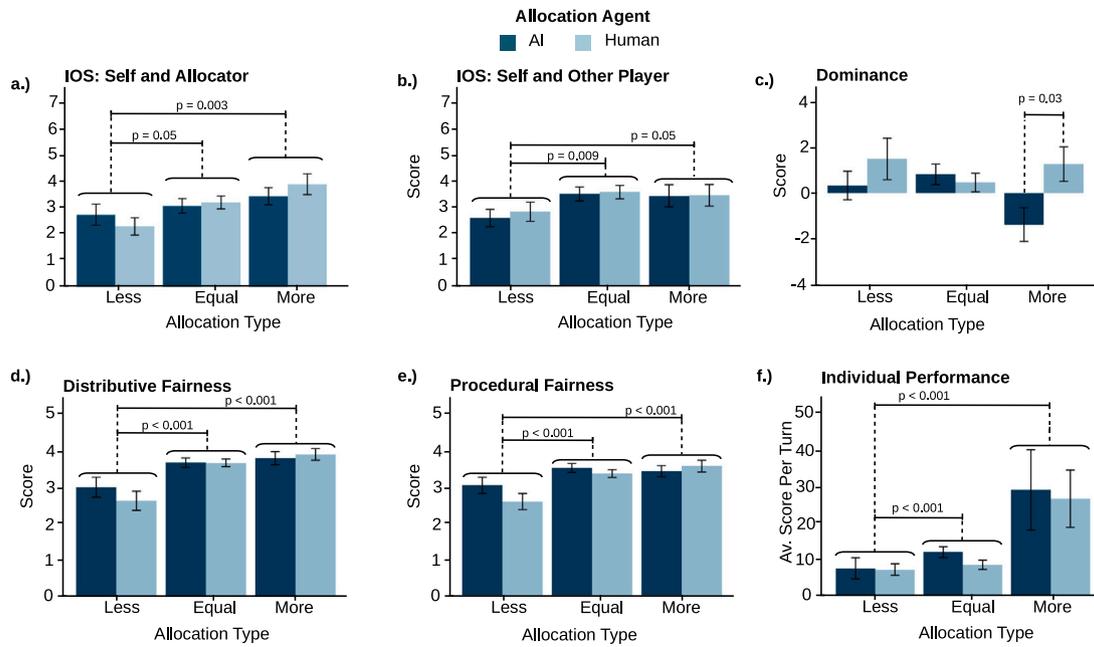


Fig. 5. Results of allocation agent and allocation type on: (a) IOS: self and the allocator, (b) IOS: self and the other player, (c) dominance, (d) distributive fairness, (e) procedural fairness, and (f) individual performance.

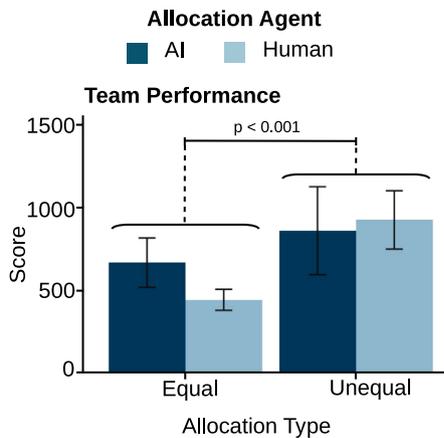


Fig. 6. Group level analysis of team performance.

on the work of (Lee, 2018). Their work identifies that people did not show differences in fairness perceptions when a human or algorithmic manager assigned mechanical tasks. They also found that individuals recognized that people differed in the way in which they reasoned about the underlying processes when the allocation decision came from a machine versus a human. Our results suggest that the agency of the allocator does not necessarily influence perceptions of fairness with respect to the process used (procedural fairness) or the level of resources allocated (distributive fairness). These results may be due to the cooperative context of the experiment which had all members working towards a similar goal. Future work should look into how these results translate to a competitive context where group members are rewarded based on their individual input.

5.3. Task performance

We found that machine allocation behavior has an impact on the performance of individuals and groups. Individuals, who received less resources performed worse than individuals who received an equal or more resources. However, surprisingly, groups (i.e. both players

together) in which both players controlled an equal amount of turns were outperformed by groups in which one player controlled more turns than the other.

The influence of machine allocation behavior on performance is critical to the adoption of machines in workplace groups and our work expands on literature exploring the role of AI on performance. While machines have shown to enhance different aspects of human abilities, work by Zhang et al. suggests negative effects to performance due to the social impact that machines can have on a group (Zhang et al., 2021). Our finding raises questions about what type of allocation strategy a machine should apply in certain allocation tasks. To avoid negative social consequences, an equal distribution may be the ideal method to pursue in certain. Humans can process an equal distribution of resources more easily and avoid putting any cognitive load trying to discern meaning behind such an action (Diekmann, Samuels, Ross, & Bazerman, 1997). Yet, our findings suggest that distributing resources equally across group members may come at a cost on group performance. Individuals within the group may have identified different approaches to the game that do not necessarily compliment one another. One individual may have wanted to achieve the highest score by clearing as many single lines of Tetris blocks possible while their partner aimed to stack and then clear multiple rows to obtain a higher group score. Further investigations into how individuals form different strategies in groups where resources are allocated equally versus unequally could give further insights into the effects of machine allocation behavior on performance.

5.4. Implications for theory

Most importantly, our research has implications for our understanding of algorithmic fairness. Existing work on algorithmic fairness has adopted a "snapshot" view that focuses on fairness assessments of single algorithmic decisions such as rankings in hiring decisions, or risk evaluations for court cases and loan approvals (e.g. Wang et al. (2020). This "snapshot" view, however, misses that in certain contexts fairness is better understood as a dynamic phenomenon that develops and varies over time. In situations such as in our Co-Tetris game, fairness has to be understood as the accumulation of a series of allocation decisions at a specific point in time. For example, the fair allocation of turns

among players in the Co-Tetris game is dependent on when in time an assessment is made. Additionally, algorithmic fairness of a sequence of decisions is not only dependent on the overall allocation of turns among players but also on the specific sequences of turn allocations over time. An allocation sequence in which one player receives most of the turns at the beginning of a game will likely be perceived differently from one in which each player receives alternating turns even when both players have received the same number of turn allocations overall.

Jones et al. proposes a model of organizational justice which posits that fairness is more malleable than previously believed (Jones & Skarlicki, 2013). The authors argue that individuals are constantly trying to set accurate perceptions of how fair or unfair certain events are and this can lead to adjustments in fairness judgments. This model can be extended to contexts such as ours in which machines make a sequence of allocation decisions. People's understanding of the underlying strategy that a machine is using will evolve with increasing interactions leading to changes in fairness judgments over time.

Our research also has implications for our understanding of human-agent teamwork and human-robot teamwork. Recent research on human-robot interaction has shown how robots can impact group dynamics and interpersonal behavior in groups and teams (e.g. Jung et al. (2020), Jung, Martelaro, and Hinds (2015), Sebo, Stoll, Scassellati, and Jung (2020), Strohkorb Sebo et al. (2018), Tennent et al. (2019), Traeger et al. (2020)). While existing work has focused predominantly on building understanding about the impact of specific verbal or non-verbal machine behaviors (e.g. expressions of a specific emotion, behavior strategy, or utterance), our research introduces machine allocation behavior as a key input factor to consider when building understanding about a machine's impact on groups and teams.

Finally, our research has implications for our understanding of machine behavior. The study of machine behavior involves a broader look at the social consequences of including machines into various contexts (Rahwan et al., 2019). Rahwan and colleagues argue that machine behavior should be investigated across four dimensions including the mechanisms that produce a specific machine behavior, the development of such behavior, its functional consequences, and the evolution of said behavior. They propose that an investigation of machine behavior across these four dimensions will aid in understanding of the societal effects a machine's behavior can have. With our focus on machine allocation behavior, we contribute to our understanding of machine behavior by exploring a specific type of machine behavior and show how it impacts the behavior of individuals and their social relations with each other.

5.5. Implications for design

The findings of our study reinforce the importance of designing machines that understand the interpersonal dynamics of groups and human notions of fairness judgments. Certain methods have explored using different verbal and nonverbal cues to inform machines about the current state of a group and take appropriate action. For instance, Seo et al. leveraged a surgical team's actions to make predictions about a team's mental model (Seo et al., 2021). Salam and authors used behavioral cues to infer levels of engagement using supervised learning (Salam & Chetouani, 2015). Based on our results, we argue that monitoring the level of resources each individual receives relative to others may be an appropriate measure to help infer or make predictions about a group's cohesion level. We show that individuals within a group have different judgments depending on whether resources were distributed in their favor (i.e. individuals who received more blocks found the distribution to be fair while those who received less found it less fair). Machines can leverage such cues to engage in actions to repair or support the social relationship between group members.

Machines should also have appropriate responses to scenarios where an equal distribution of resources cannot be obtained. In such scenarios, machines can maintain and promote feelings of inclusion within

the group to ensure that they do not feel isolated from their counterparts. This can be done through different verbal and nonverbal behaviors (Strohkorb et al., 2016; Tennent et al., 2019). For instance, Sebo et al. showed that a robot's utterance can influence the sense of inclusion that a group member may feel (Strohkorb Sebo et al., 2018). Another way of showing appropriate responses is communicating and being transparent about the machine's allocation decision. Explaining why and how the machine decided to allocate a resource to one person over another could improve perceptions of fairness and trust towards the machine (Kizilcec, 2016; Yu et al., 2020). As trust within groups influences group performance (Drescher, Korsgaard, Welp, Picot, & Wigand, 2014), establishing transparency in the machine's decision-making process is crucial. Moreover, providing transparency in why a decision was made has shown to resolve conflicts (Park, Karahalios, Salehi, & Eslami, 2022).

5.6. Limitations

Our findings are based on behavior observable from playing Tetris. We chose Tetris due to its popularity and recognition of the rules as well as its rich history in research across a broad range of research spaces including machine learning (Lu, Wei, Lin, Yan, & Li, 2018), psychology (Pilegard & Mayer, 2018), and biology (Iyadurai et al., 2018). While focusing on Tetris might raise questions about the ecological validity of our work, the broader research community has used Tetris as a reliable platform to understand human behavior across a wide range of topics and phenomena (Dabbish, Farzan, Kraut, & Postmes, 2012; Farzan, Dabbish, Kraut, & Postmes, 2011). Furthermore, although our study does not further elaborate how the specific platform influences the results, our Co-Tetris platform emulates a commonly observed setting where a machine has to allocate resources across group members to accomplish a task (e.g. Dabbish et al. (2012), Gombolay et al. (2015), Jung et al. (2018), Short and Mataric (2017)).

We constrained the MTurk participants to those who reside in the US; therefore, our participants may not be representative of a larger population (Bryant, Borenstein, & Howard, 2020; Keith, Tay, & Harms, 2017). However, MTurk participants are known to be more diverse than other modes of participant recruitment (e.g. college students) (Behrend, Sharek, Meade, & Wiebe, 2011; Keith et al., 2017). Moreover, concerns with MTurk include possible lack of attention and control of the participants (Keith et al., 2017; Kittur, Chi, & Suh, 2008). Despite such concerns, our study conducted a manipulation check to ensure that the participants were cognizant of the received resource.

Finally, our study is limited with respect to length of the task. Participants from MTurk were randomly placed in groups in order to complete a short one-time task. We do not know how the perceptions of individuals would change if they completed this task over longer periods of time or if they were given more information about their group members. Future work should explore how different group dynamics and perceptions form over longer periods of time.

6. Conclusion

People spend the majority of their waking hours in the presence of others (Kahneman, Krueger, Schkade, Schwarz, & Stone, 2004). We live and work with other people, and the interactions we have with them shape much of our well-being and what we can accomplish.

We introduced the notion of machine allocation behavior and showed how it can impact our relationships with others. The influence that machines can have on people goes far beyond the person directly interacting with it. We hope that our research lays the foundation for further research that deepens our understanding about the ways in which machines can impact not only individuals but groups of people and their interactions with each other.

CRediT authorship contribution statement

Houston Claire: Conceptualization, Methodology, Software, Data analysis. **Seyun Kim:** Data analysis, Software, Paper preparation. **René F. Kizilcec:** Conceptualization, Methodology, Data analysis, Paper preparation. **Malte Jung:** Conceptualization, Methodology, Paper preparation.

Data availability

The authors do not have permission to share data.

Acknowledgments

We want to thank Jignesh Modi and Hyun Kyo Jung for their help with the development of the Co-Tetris platform. The work was supported by the National Science Foundation, United States CPS #1646417, National Science Foundation, United States CAREER #1942085, and ONR #N00014-19-1-2299. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding sources.

References

- Adams, J. S. (1965). Inequity in social exchange. In *Advances in experimental social psychology*, vol. 2 (pp. 267–299). Elsevier.
- Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of other in the self scale and the structure of interpersonal closeness. *Journal of Personality and Social Psychology*, 63(4), 596.
- Baig, E. (2018). Who's going to review your college applications—A committee or a computer. *USA Today*, Dec, 2, 2018.
- Bailey, D. E., Faraj, S., Hinds, P. J., Leonardi, P. M., & von Krogh, G. (2022). We are all theorists of technology now: A relational perspective on emerging technology and organizing. *Organization Science*, 33(1), 1–18.
- Bansal, G., Nushi, B., Kamar, E., Weld, D. S., Lasecki, W. S., & Horvitz, E. (2019). Updates in human-ai teams: Understanding and addressing the performance/compatibility tradeoff. In *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01 (pp. 2429–2437).
- Barocas, S., Hardt, M., & Narayanan, A. (2018). *Fairness and machine learning*. fairmlbook.org, 2019.
- Barrick, M. R., Stewart, G. L., Neubert, M. J., & Mount, M. K. (1998). Relating member ability and personality to work-team processes and team effectiveness. *Journal of Applied Psychology*, 83(3), 377.
- Bartol, K. M., & Srivastava, A. (2002). Encouraging knowledge sharing: The role of organizational reward systems. *Journal of Leadership & Organizational Studies*, 9(1), 64–76.
- Behrend, T. S., Sharek, D. J., Meade, A. W., & Wiebe, E. N. (2011). The viability of crowdsourcing for survey research. *Behavior Research Methods*, 43(3), 800–813.
- Brandao, M., Jirotko, M., Webb, H., & Luff, P. (2020). Fair navigation planning: A resource for characterizing and designing fairness in mobile robots. *Artificial Intelligence*, 282, Article 103259.
- Bryant, D., Borenstein, J., & Howard, A. (2020). Why should we gender? The effect of robot gendering and occupational stereotypes on human trust and perceived competency. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction* (pp. 13–21).
- Campbell, R. W. (2020). Artificial intelligence in the courtroom: The delivery of justice in the age of machine learning. *Colorado Technology Law Journal*, 18, 323.
- Chang, M. L., Pope, Z., Short, E. S., & Thomaz, A. L. (2020). Defining fairness in human-robot teams. In *2020 29th IEEE international conference on robot and human interactive communication* (pp. 1251–1258). IEEE.
- Christin, A. (2017). Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society*, 4(2), Article 2053951717718855. <http://dx.doi.org/10.1177/2053951717718855>, Publisher: SAGE Publications Ltd.
- Claire, H., Chen, Y., Modi, J., Jung, M., & Nikolaidis, S. *Reinforcement learning with fairness constraints for resource distribution in human-robot teams: Technical report*.
- Cohen, S., & Wills, T. A. (1985). Stress, social support, and the buffering hypothesis. *Psychological Bulletin*, 98(2), 310.
- Collins, B. E., Guetzkow, H. S., et al. (1964). *A social psychology of group processes for decision-making*. New York: Wiley.
- Colquitt, J. A. (2001). On the dimensionality of organizational justice: A construct validation of a measure. *Journal of Applied Psychology*, 86(3), 386.
- Cook, K. S., & Hegtvædt, K. A. (1983). Distributive justice, equity, and equality. *Annual Review of Sociology*, 217–241.
- Dabbish, L., Farzan, R., Kraut, R., & Postmes, T. (2012). Fresh faces in the crowd: Turnover, identity, and commitment in online groups. In *Proceedings of the ACM 2012 conference on computer supported cooperative work* (pp. 245–248).
- Deutsch, J. (1953). A new type of behaviour theory. *British Journal of Psychology*, 44(4), 304.
- Diekmann, K. A., Samuels, S. M., Ross, L., & Bazerman, M. H. (1997). Self-interest and fairness in problems of resource allocation: Allocators versus recipients. *Journal of Personality and Social Psychology*, 72(5), 1061.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114.
- Dijkstra, J. J., Liebrand, W. B., & Timminga, E. (1998). Persuasiveness of expert systems. *Behaviour & Information Technology*, 17(3), 155–163.
- Drescher, M. A., Korsgaard, M. A., Welpe, I. M., Picot, A., & Wigand, R. T. (2014). The dynamics of shared leadership: Building trust and enhancing performance. *Journal of Applied Psychology*, 99(5), 771.
- Dulebohn, J. H., & Martocchio, J. J. (1998). Employee perceptions of the fairness of work group incentive pay plans. *Journal of Management*, 24(4), 469–488.
- Everton, W. J., Jolton, J. A., & Mastrangelo, P. M. (2007). Be nice and fair or else: Understanding reasons for employees' deviant behaviors. *Journal of Management Development*.
- Farzan, R., Dabbish, L. A., Kraut, R. E., & Postmes, T. (2011). Increasing commitment to online communities by designing for social presence. In *Proceedings of the ACM 2011 conference on computer supported cooperative work* (pp. 321–330).
- Folger, R., & Baron, R. A. (1996). *Violence and hostility at work: A model of reactions to perceived injustice*. American Psychological Association.
- Gombolay, M. C., Gutierrez, R. A., Clarke, S. G., Sturla, G. F., & Shah, J. A. (2015). Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Autonomous Robots*, 39(3), 293–312.
- Gottman, J. M. (2014). *What predicts divorce?: The relationship between marital processes and marital outcomes*. Psychology Press.
- Gottman, J. M., & Levenson, R. W. (1992). Marital processes predictive of later dissolution: Behavior, physiology, and health. *Journal of Personality and Social Psychology*, 63(2), 221.
- Haier, R. J., Siegel Jr, B. V., MacLachlan, A., Soderling, E., Lottenberg, S., & Buchsbaum, M. S. (1992). Regional glucose metabolic changes after learning a complex visuospatial/motor task: A positron emission tomographic study. *Brain Research*, 570(1–2), 134–143.
- Hassan, U. U., O'Riain, S., & Curry, E. (2013). Effects of expertise assessment on the quality of task routing in human computation. In *2nd international workshop on social media for crowdsourcing and human computation*, vol. 2 (pp. 1–10).
- Hodson, H. (2014). The AI boss that deploys Hong Kong's subway engineers. *New Scientist*.
- Hohenstein, J., DiFranzo, D., Kizilcec, R. F., Aghajari, Z., Mieczkowski, H., Levy, K., et al. (2021). Artificial intelligence in communication impacts language and social relationships. arXiv preprint arXiv:2102.05756.
- Hohenstein, J., & Jung, M. (2020). AI as a moral crumple zone: The effects of AI-mediated communication on attribution and trust. *Computers in Human Behavior*, 106, Article 106190.
- Holmes, E. A., James, E. L., Coode-Bate, T., & Deerpore, C. (2009). Can playing the computer game "Tetris" reduce the build-up of flashbacks for trauma? A proposal from cognitive science. *PLoS One*, 4(1), Article e4153.
- Homans, G. C. (1961). The humanities and the social sciences. *American Behavioral Scientist*, 4(8), 3–6.
- Iyadurai, L., Blackwell, S. E., Meiser-Stedman, R., Watson, P. C., Bonsall, M. B., Geddes, J. R., et al. (2018). Preventing intrusive memories after trauma via a brief intervention involving Tetris computer game play in the emergency department: A proof-of-concept randomized controlled trial. *Molecular Psychiatry*, 23(3), 674–682.
- Jehn, K. A. (1995). A multimethod examination of the benefits and detriments of intragroup conflict. *Administrative Science Quarterly*, 256–282.
- Jehn, K. A. (1997). A qualitative analysis of conflict types and dimensions in organizational groups. *Administrative Science Quarterly*, 530–557.
- Jehn, K. A., & Shah, P. P. (1997). Interpersonal relationships and task performance: An examination of mediation processes in friendship and acquaintance groups. *Journal of Personality and Social Psychology*, 72(4), 775.
- Jones, D. A., & Skarlicki, D. P. (2013). How perceptions of fairness can change: A dynamic model of organizational justice. *Organizational Psychology Review*, 3(2), 138–160.
- Julian, J. W., & Perry, F. A. (1967). Cooperation contrasted with intra-group and inter-group competition. *Sociometry*, 79–90.
- Jung, M. F. (2016). Coupling interactions and performance: Predicting team performance from thin slices of conflict. *ACM Transactions on Computer-Human Interaction*, 23(3), 1–32.
- Jung, M. F., DiFranzo, D., Shen, S., Stoll, B., Claire, H., & Lawrence, A. (2020). Robot-assisted tower construction—A method to study the impact of a robot's allocation behavior on interpersonal dynamics and collaboration in groups. *ACM Transactions on Human-Robot Interaction (THRI)*, 10(1), 1–23.
- Jung, M. F., DiFranzo, D., Stoll, B., Shen, S., Lawrence, A., & Claire, H. (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.

- Jung, M., & Hinds, P. (2018). Robots in the wild: A time for more robust theories of human-robot interaction. *ACM Transactions on Human-Robot Interaction (THRI)*, 7(1), 1–5.
- Jung, M. F., Martelaro, N., & Hinds, P. J. (2015). Using robots to moderate team conflict: The case of repairing violations. In *Proceedings of the Tenth Annual ACM/IEEE international conference on human-robot interaction* (pp. 229–236).
- Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306(5702), 1776–1780.
- Keith, M. G., Tay, L., & Harms, P. D. (2017). Systems perspective of Amazon Mechanical Turk for organizational research: Review and recommendations. *Frontiers in Psychology*, 8, 1359.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4), 513–549.
- Kittur, A., Chi, E. H., & Suh, B. (2008). Crowdsourcing user studies with Mechanical Turk. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 453–456).
- Kizilcec, R. F. (2016). How much information? Effects of transparency on trust in an algorithmic interface. In *Proceedings of the 2016 CHI conference on human factors in computing systems* (pp. 2390–2395).
- Knox, W. B., Glass, B. D., Love, B. C., Maddox, W. T., & Stone, P. (2012). How humans teach agents. *International Journal of Social Robotics*, 4(4), 409–421.
- Kodiyan, A. A. (2019). An overview of ethical issues in using AI systems in hiring with a case study of Amazon's AI based hiring tool. *Researchgate Preprint*, 1–19.
- Konovsky, M. A. (2000). Understanding procedural justice and its impact on business organizations. *Journal of Management*, 26(3), 489–511.
- Langholtz, H. J., Marty, A. T., Ball, C. T., & Nolan, E. C. (2002). *Resource-allocation behavior*. Springer Science & Business Media.
- Lawler, E. (1971). *Pay and organization effectiveness: A psychological view*. NY: McGraw Hill.
- Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, 5(1), Article 2053951718756684.
- Lee, C. T.-S., & Doran, D. M. (2017). The role of interpersonal relations in healthcare team communication and patient safety: A proposed model of interpersonal process in teamwork. *Canadian Journal of Nursing Research*, 49(2), 75–93.
- Lee, M. K., Kim, J. T., & Lizarondo, L. (2017). A human-centered approach to algorithmic services: Considerations for fair and motivating smart community service management that allocates donations to non-profit organizations. In *Proceedings of the 2017 CHI conference on human factors in computing systems* (pp. 3365–3376).
- Lee, M. K., Kusbit, D., Metsky, E., & Dabbish, L. (2015). Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (pp. 1603–1612).
- Lee, M. K., & Rich, K. (2021). Who is included in human perceptions of AI?: Trust and perceived fairness around healthcare AI and cultural mistrust. In *Proceedings of the 2021 CHI conference on human factors in computing systems* (pp. 1–14).
- Levenson, R. W., & Gottman, J. M. (1985). Physiological and affective predictors of change in relationship satisfaction. *Journal of Personality and Social Psychology*, 49(1), 85.
- Leventhal, G. S. (1976). The distribution of rewards and resources in groups and organizations. In *Advances in experimental social psychology*, vol. 9 (pp. 91–131). Elsevier.
- Liden, R. C., Wayne, S. J., & Sparrowe, R. T. (2000). An examination of the mediating role of psychological empowerment on the relations between the job, interpersonal relationships, and work outcomes. *Journal of Applied Psychology*, 85(3), 407.
- Lindstedt, J., & Gray, W. (2013). Extreme expertise: Exploring expert behavior in Tetris. In *Proceedings of the annual meeting of the cognitive science society*, vol. 35, no. 35.
- Ljungblad, S., Kotrbova, J., Jacobsson, M., Cramer, H., & Niechwiadowicz, K. (2012). Hospital robot at work: Something alien or an intelligent colleague? In *Proceedings of the ACM 2012 conference on computer supported cooperative work* (pp. 177–186).
- Lu, H., Wei, X., Lin, N., Yan, G., & Li, X. (2018). Tetris: Re-architecting convolutional neural network computation for machine learning accelerators. In *2018 IEEE/ACM international conference on computer-aided design* (pp. 1–8). IEEE.
- Ma, Z., & Gajos, K. Z. (2022). Not just a preference: Reducing biased decision-making on dating websites. In *CHI conference on human factors in computing systems* (pp. 1–14).
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
- Mieczkowski, H., Hancock, J. T., Naaman, M., Jung, M., & Hohenstein, J. (2021). AI-mediated communication: Language use and interpersonal effects in a referential communication task. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), 1–14.
- Mitchell, S., Potash, E., Barocas, S., D'Amour, A., & Lum, K. (2021). Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and Its Application*, 8, 141–163.
- Mukerjee, A., Biswas, R., Deb, K., & Mathur, A. P. (2002). Multi-objective evolutionary algorithms for the risk-return trade-off in bank loan management. *International Transactions in Operational Research*, 9(5), 583–597.
- Mutlu, B., Shiwa, T., Kanda, T., Ishiguro, H., & Hagita, N. (2009). Footing in human-robot conversations: How robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE international conference on human robot interaction* (pp. 61–68).
- Nozick, R. (1973). Distributive justice. *Philosophy & Public Affairs*, 45–126.
- Park, J. S., Karahalios, K., Salehi, N., & Eslami, M. (2022). Power dynamics and value conflicts in designing and maintaining socio-technical algorithmic processes. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1), 1–21.
- Pilegard, C., & Mayer, R. E. (2018). Game over for Tetris as a platform for cognitive skill training. *Contemporary Educational Psychology*, 54, 29–41.
- Presbitero, A., & Teng-Calleja, M. (2022). Job attitudes and career behaviors relating to employees' perceived incorporation of artificial intelligence in the workplace: A career self-management perspective. *Personnel Review*, (ahead-of-print).
- Rahwan, I., Cebrían, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., et al. (2019). Machine behaviour. *Nature*, 568(7753), 477–486.
- Reeves, B., & Nass, C. (1996). *The media equation: how people treat computers, television, and new media like real people*. Cambridge university press Cambridge, United Kingdom.
- Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication*, 10, 27.
- Salam, H., & Chetouani, M. (2015). Engagement detection based on multi-party cues for human robot interaction. In *2015 international conference on affective computing and intelligent interaction* (pp. 341–347). IEEE.
- Saxena, D., Badillo-Urquiola, K., Wisniewski, P. J., & Guha, S. (2020). A human-centered review of algorithms used within the us child welfare system. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1–15).
- Sebo, S., Stoll, B., Scassellati, B., & Jung, M. F. (2020). Robots in groups and teams: A literature review. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), 1–36.
- Sebo, S. S., Traeger, M., Jung, M., & Scassellati, B. (2018). The ripple effects of vulnerability: The effects of a robot's vulnerable behavior on trust in human-robot teams. <http://dx.doi.org/10.1145/3171221.3171275>.
- Seo, S., Kennedy-Metz, L. R., Zenati, M. A., Shah, J. A., Dias, R. D., & Unhelkar, V. V. (2021). Towards an AI coach to infer team mental model alignment in healthcare. In *2021 IEEE conference on cognitive and computational aspects of situation management* (pp. 39–44). IEEE.
- Shirado, H., & Christakis, N. A. (2017). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, 545(7654), 370–374.
- Short, E., & Mataric, M. J. (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* (pp. 385–390). IEEE.
- Singh, V., Nanavati, B., Kar, A. K., & Gupta, A. (2022). How to maximize clicks for display advertisement in digital marketing? A reinforcement learning approach. *Information Systems Frontiers*, 1–18.
- Skarlicki, D. P., & Folger, R. (1997). Retaliation in the workplace: The roles of distributive, procedural, and interactional justice. *Journal of Applied Psychology*, 82(3), 434.
- Skarlicki, D. P., Van Jaarsveld, D. D., & Walker, D. D. (2008). Getting even for customer mistreatment: The role of moral identity in the relationship between customer interpersonal injustice and employee sabotage. *Journal of Applied Psychology*, 93(6), 1335.
- Smith, R. J., & Cook, P. E. (1973). Leadership in dyadic groups as a function of dominance and incentives. *Sociometry*, 561–568.
- Song, B., Zurita, N. F. S., Gyory, J. T., Zhang, G., McComb, C., Cagan, J., et al. (2022). Decoding the agility of artificial intelligence-assisted human design teams. *Design Studies*, 79, Article 101094.
- Ståhl, I. (1977). An n-person bargaining game in the extensive form. In *Mathematical economics and game theory* (pp. 156–172). Springer.
- Strohkorb, S., Fukuto, E., Warren, N., Taylor, C., Berry, B., & Scassellati, B. (2016). Improving human-human collaboration between children with a social robot. In *2016 25th IEEE international symposium on robot and human interactive communication* (pp. 551–556). IEEE.
- Strohkorb Sebo, S., Traeger, M., Jung, M., & Scassellati, B. (2018). The ripple effects of vulnerability: The effects of a robot's vulnerable behavior on trust in human-robot teams. In *Proceedings of the 2018 ACM/IEEE international conference on human-robot interaction* (pp. 178–186).
- Tennent, H., Shen, S., & Jung, M. (2019). Micbot: A peripheral robotic object to shape conversational dynamics and team performance. In *2019 14th ACM/IEEE international conference on human-robot interaction* (pp. 133–142). IEEE.
- Terzioğlu, Y., Mutlu, B., & Şahin, E. (2020). Designing social cues for collaborative robots: The Role of gaze and breathing in human-robot collaboration. In *2020 15th ACM/IEEE international conference on human-robot interaction* (pp. 343–357). IEEE.
- Traeger, M. L., Sebo, S. S., Jung, M., Scassellati, B., & Christakis, N. A. (2020). Vulnerable robots positively shape human conversational dynamics in a human-robot team. *Proceedings of the National Academy of Sciences*, 117(12), 6370–6375.
- Wang, R., Harper, F. M., & Zhu, H. (2020). Factors influencing perceived fairness in algorithmic decision-making: Algorithm outcomes, development procedures, and individual differences. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1–14).

- Wiggins, J. S., Trapnell, P., & Phillips, N. (1988). Psychometric and geometric characteristics of the revised interpersonal adjective scales (IAS-R). *Multivariate Behavioral Research*, 23(4), 517–530.
- Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56–75.
- Yamagishi, T., Horita, Y., Mifune, N., Hashimoto, H., Li, Y., Shinada, M., et al. (2012). Rejection of unfair offers in the ultimatum game is no evidence of strong reciprocity. *Proceedings of the National Academy of Sciences*, 109(50), 20364–20368.
- Yang, B., Li, L. X., Ji, H., & Xu, J. (2001). An early warning system for loan risk assessment using artificial neural networks. *Knowledge-Based Systems*, 14(5–6), 303–306.
- Yang, P.-C., Sasaki, K., Suzuki, K., Kase, K., Sugano, S., & Ogata, T. (2016). Repeatable folding task by humanoid robot worker using deep learning. *IEEE Robotics and Automation Letters*, 2(2), 397–403.
- Yeomans, M., Shah, A., Mullainathan, S., & Kleinberg, J. (2019). Making sense of recommendations. *Journal of Behavioral Decision Making*, 32(4), 403–414.
- Yu, H., Liu, Y., Wei, X., Zheng, C., Chen, T., Yang, Q., et al. (2019). Fair and explainable dynamic engagement of crowd workers. In *IJCAI* (pp. 6575–6577).
- Yu, B., Yuan, Y., Terveen, L., Wu, Z. S., Forlizzi, J., & Zhu, H. (2020). Keeping designers in the loop: Communicating inherent algorithmic trade-offs across multiple objectives. In *Proceedings of the 2020 ACM designing interactive systems conference* (pp. 1245–1257).
- Završnik, A. (2021). Algorithmic justice: Algorithms and big data in criminal justice settings. *European Journal of Criminology*, 18(5), 623–642. <http://dx.doi.org/10.1177/1477370819876762>, Publisher: SAGE Publications.
- Zhang, G., Raina, A., Cagan, J., & McComb, C. (2021). A cautionary tale about the impact of AI on human design teams. *Design Studies*, 72, Article 100990.