

Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management

Big Data & Society
January–June 2018: 1–16
© The Author(s) 2018
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/2053951718756684
journals.sagepub.com/home/bds



Min Kyung Lee

Abstract

Algorithms increasingly make managerial decisions that people used to make. Perceptions of algorithms, regardless of the algorithms' actual performance, can significantly influence their adoption, yet we do not fully understand how people perceive decisions made by algorithms as compared with decisions made by humans. To explore perceptions of algorithmic management, we conducted an online experiment using four managerial decisions that required either mechanical or human skills. We manipulated the decision-maker (algorithmic or human), and measured perceived fairness, trust, and emotional response. With the mechanical tasks, algorithmic and human-made decisions were perceived as equally fair and trustworthy and evoked similar emotions; however, human managers' fairness and trustworthiness were attributed to the manager's authority, whereas algorithms' fairness and trustworthiness were attributed to their perceived efficiency and objectivity. Human decisions evoked some positive emotion due to the possibility of social recognition, whereas algorithmic decisions generated a more mixed response – algorithms were seen as helpful tools but also possible tracking mechanisms. With the human tasks, algorithmic decisions were perceived as less fair and trustworthy and evoked more negative emotion than human decisions. Algorithms' perceived lack of intuition and subjective judgment capabilities contributed to the lower fairness and trustworthiness judgments. Positive emotion from human decisions was attributed to social recognition, while negative emotion from algorithmic decisions was attributed to the dehumanizing experience of being evaluated by machines. This work reveals people's lay concepts of algorithmic versus human decisions in a management context and suggests that task characteristics matter in understanding people's experiences with algorithmic technologies.

Keywords

Algorithmic management, perception, folk theory, fairness, trust, emotion

This article is a part of special theme on Algorithms in Culture. To see a full list of all articles in this special theme, please click here: <http://journals.sagepub.com/page/bds/collections/algorithms-in-culture>.

Introduction

Advances in artificial intelligence, machine learning, and data infrastructure are transforming how people govern and manage citizens and organizations. Now more than ever, computational algorithms increasingly make decisions that human managers used to make, changing the practices of managers, policy makers, physicians, teachers, police, judges, on-demand labor platforms, online communities, and more. Algorithms match patients to therapists and doctors (Amino; Cloud 9 Psych; Idrees et al., 2013), passengers to

ridesharing drivers (Lee et al., 2015) and subway workers to maintenance tasks (Hodson, 2014). Predictive analytics are used on hiring platforms (Carey and Smith, 2016) like LinkedIn,¹ where algorithms sort through thousands of profiles to recommend promising

Carnegie Mellon University, Pittsburgh, PA, USA

Corresponding author:

Min Kyung Lee, Carnegie Mellon University, 4615 Forbes Ave, Pittsburgh, PA 15213, USA.

Email: mklee@cmu.edu



job candidates to company recruiters. In customer call centers, algorithms examine employees' calls with customers to evaluate their performance (Azadeh et al., 2013; Petrushin, 1999). Algorithms are also used by companies to determine which employees are at risk of quitting, allowing the companies to take preventative action (Silverman and Waller, 2015); by courts, to predict which people are likely to commit crimes again in order to make bail decisions (Electronic Privacy Information Center); and by police, to decide which areas of a city are likely to see crime in order to determine where to patrol (National Institution of Justice). How do people feel about algorithms taking over managerial decisions that used to be made by humans? Do people think that algorithmic decisions are more fair and trustworthy, or less?

Algorithms may enable efficient, optimized, and data-driven decision-making, and in fact this vision is one of main drivers of increasing adoption of algorithms for managerial and organizational decisions. However, the fact that these decisions are made by algorithms, rather than by people, may influence perceptions of the decisions that are made, regardless of the qualities of the actual decision-outcomes (Sundar and Nass, 2001). These perceptions may in turn influence people's trust in and attitudes toward algorithmic decisions, which are critical aspects of workplaces, communities, and societies that allow people to thrive. For example, previous research shows that if organizational and managerial decisions are perceived as unfair, the affected workers experience resentment and anger and may engage in retaliation and acts against the organization (Skarlicki and Folger, 1997). As we are in a transition period in which algorithms are increasingly making more managerial and organizational decisions, it is an opportune and critical time to understand societal attitudes toward this change and build a knowledge base to advance people's theories and mental models of algorithmic technology.

Increasing scholarly attention has been given to perceptions of algorithms, especially for online media content (Bucher, 2016; Eslami et al., 2015; French and Hancock, 2017; Rader and Gray, 2015). This line of work suggests that people form diverse mental models and folk theories about how algorithms operate, regardless of how algorithms actually work. Even before this recent focus on algorithms as a socio-technical academic subject, a long line of research in human-computer interaction, communication, and human factors has investigated how people perceive computers. Notably, the Computers Are Social Actors (CASA) paradigm has demonstrated that people interact with computers as if they were not just tools but social agents. For example, people may respond positively to a computer's flattery even though

the complement was not given by a person and was not genuine (Fogg and Nass, 1997; Nass and Moon, 2000; Nass and Steuer, 1993; Reeves and Nass, 1996). However, while algorithms are one of the computational methods that operate computers, algorithms are rather more abstract than physical computers, and many managerial algorithms do not directly interact with people. Our research investigates social perceptions and attitudes toward decisions made by algorithms as compared to people.

The context that we investigate is also different from the contexts used in the studies mentioned above. Most previous studies used computers and automation technologies as decision aids or interactive partners. However, the recent trend of algorithms assuming managerial roles puts people into a different power structure than when they are "users" or "consumers" of algorithmic systems. For consumer applications, people can decide to use algorithmic decisions or not; when those decisions are incorporated into managerial and governance processes, however, it is much more difficult for people to reject or refute them.

Our research takes a step toward systematically understanding perceptions of algorithmic decisions in management contexts and how perceptions differ depending on whether the decision-maker is a person or an algorithm. Although there are many ways perceptions of decisions vary, we focused on judgments of fairness and trust and emotional responses—all of which contribute to positive collaboration and job satisfaction (Hackman and Oldham, 1976) and can be deterred by the introduction of algorithms: perceived unfairness of decisions has been associated with workers taking action against their organizations (Leventhal, 1980); trust in decision quality and reliability has been suggested as a major determinant for effective adoption of automation technology (Lee and See, 2004); and affective experiences play a key role in work motivation (Seo et al., 2004), and previous work suggests negative emotions are often associated with the introduction of new automation technology (Zuboff, 1988). We conducted an online experiment in which participants read descriptions of a managerial decision that either algorithms or people had made. The managerial decisions were based on real-world examples of workplaces where algorithms have begun to change organizational practices. We then examined the influence of the decision-maker (algorithmic or human) on participants' perceptions of the decisions.

Our work makes two contributions to research on the social psychology of computing technologies in general, and in particular to emerging theories around people's experiences with algorithmic technologies. First, we experimentally demonstrate how people's

knowledge of the type of decision-maker—algorithmic or human—can influence their perceptions of the decisions made. We also find that whether tasks require more “human” or “mechanical” skills can influence perceptions, identifying an important construct to consider in future studies on understanding people’s experiences of algorithmic technologies. Second, our results offer insights into people’s opinions of and reactions to the transition from human to algorithmic decision-makers, which can help us as a society to create trustworthy, fair, and positive workplaces with algorithms.

Perception of algorithms vs. people

We posit that how people perceive algorithmic and human decision-makers may influence their perceptions of the managerial decisions that are made.

What are algorithms?

One dictionary definition of “algorithm” is “a process or set of rules to be followed in calculations or other problem-solving operations, especially by a computer.”² This generic definition includes any rules that people and/or computers can follow. In our paper, we use the term “algorithm” to mean a computational formula that autonomously makes decisions based on statistical models or decision rules without explicit human intervention. This reflects the recent advancement of the autonomous decision-making capabilities of algorithms from artificial intelligence and machine learning, and current usage of the term in popular media.

Our focus of inquiry—everyday people’s perceptions of algorithms—is closely related to recent discourses around socio-technical definitions of algorithms (Gillespie, 2014; Kitchin, 2016), which go beyond the technical, entitative definition described above. Our line of inquiry focuses on how algorithms are enacted in real-world social contexts and influenced by multiple stakeholders, ranging from the media to developers to users. For example, recent research has highlighted the role of human choices in the algorithm development process (Barocas and Selbat, 2016; Sweeney, 2013) and potential gaps between the mathematic, computational definitions of fairness used by algorithms and more social definitions of fairness (Lee and Baykal, 2017), which can lead to unintended consequences or biases in algorithmic decisions. Another thread of work focuses on folk theories and mental models of algorithms; the work of Rader and Gray (2015) shows how people make sense of algorithms in social media and suggests different mental models. Bucher (2016) proposes the concept of the “algorithmic imaginary,” through a qualitative study with Facebook users, to

highlight the importance of understanding people’s affective experiences with algorithms. Our research extends this line of work with a focus on how everyday people perceive and feel about decision-outcomes from algorithmic decision-makers, and how those perceptions differ when the decision-makers are human.

Algorithmic vs. human decision-makers

We posit that people may attribute different qualities to algorithmic and human decision-makers, which may in turn influence their perceptions of the decisions. On one hand, the CASA literature demonstrates that people respond to computers according to socio-psychological principles similar to those that regulate human–human interaction (see Reeves and Nass, 1996 for a review). This literature suggests that people may judge algorithmic and human decision-makers in a similar manner, especially when they are engaged in social interaction. On the other hand, research on bots and robot suggests that people may perceive computational systems as having less agency and emotional capability than humans (Gray et al., 2007; Gray and Wegner, 2012; Waytz and Norton, 2014). These results suggest that people will perceive algorithmic decision-makers as more rational, and less intentional and emotional than people. As our focus is on perceptions of decision-outcomes that do not involve direct interaction between algorithms and people, we believe the different qualities that people attribute to the decision-makers, rather than interpersonal interaction principles, will play a role in the evaluation of decision-outcomes.

Algorithmic vs. human decisions

Perceptions of algorithmic decisions may be influenced by the different qualities that people attribute to the two classes of decision-makers. Previous research on source bias suggests that attitudes towards an information source can influence judgments of the credibility and quality of the information (Sundar and Nass, 2001). Research in computer-supported collaborative work and human-robot interaction suggests that source bias can also influence the ways people interact with computational systems. For example, people cooperated with a computer less when it took the form of a dog than of a human (Parise et al., 1996), and attributed less responsibility to a machine-like robot than to a human-like robot (Hinds et al., 2004). However, previous work has not investigated how perceptions of computational decision-makers (e.g. algorithms) influence perceptions of the decisions themselves. This motivates our overall research question: how do perceptions of algorithmic and human decisions differ?

Tasks that require human vs. mechanical skills

We posit that how people perceive algorithmic and human decisions will depend on people's expectations about different tasks, particularly whether people think the tasks require skills that humans can do better than machines or vice versa. Literature from cognitive psychology and human factors defines types of knowledge and skills that are generally hard to code in computer programming languages (Reber, 1989). Explicit and declarative knowledge can be easily codified and transferred to other people, but procedural and tacit knowledge are more difficult to codify or transfer to other people. Often, this kind of knowledge is acquired from hands-on experiences and practices. Thus it has been thought that such knowledge cannot be programmed into computing technologies. Making subjective and intuitive judgments, understanding and expressing emotion, and navigating social nuances are still regarded as difficult for computers and machines to carry out, despite active research efforts to equip computers with such capabilities. For example, Waytz and Norton's study suggests that people think that computers and robots have less emotional capability than humans (2014). For these reasons, we believe that people will distinguish between tasks that require more "human" skills (e.g. subjective judgment and emotional capability), and those that require more "mechanical" skills (e.g. processing quantitative data for objective measures). We explain how these task characteristics could influence people's perceptions of decisions below.

Perceived fairness

Fairness is defined as treating everyone equally or equitably based on people's performance or needs (Leventhal, 1980). In particular, the focus of our inquiry is perceived fairness. Previous research suggests that people judge the fairness of decisions by considering the procedures that regulate the decision process and by their interpersonal interactions with the decision-makers, in addition to the decisions outcomes themselves (Leventhal, 1980; Skarlicki and Folger, 1997). We posit that algorithmic decision-makers will be perceived to have higher procedural fairness because algorithms follow the same procedures every time, are not influenced by emotional factors, and have no agency, and thus are perceived less biased than human decision-makers. Previous research has shown that, when viewed as the source of information, computational sources have been perceived as higher in quality and objectivity (Sundar and Nass, 2001). In tasks that require mechanical skills and do not involve subjective judgment and emotion, people may feel there

is less room for human bias. Therefore, in these tasks, people may think that algorithmic and human decisions are equally fair. On the other hand, in tasks that require more "human" skills, there is more room for human bias or preference. Examples from economics and social psychology show that perceptions of bias decrease perceived fairness and objectivity, especially within managerial and business contexts (Babcock et al., 1995; Greenberg, 1986; Konovsky and Folger, 1991; Prendergast and Topel, 1993). Thus we believe algorithmic decisions will be perceived as fairer than human decisions in tasks that require human skills.

H1. Algorithmic decisions are perceived as fairer than human-made decisions in tasks that require human skills but not in tasks that require mechanical skills.

Trust regarding the reliability of future decisions

Trust can be defined as the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability. Trust plays a central role in intraorganizational cooperation, coordination, and control (Kramer, 1999). With automation technology such as algorithms, establishing the right level of trust, or how much people believe in the reliability and accuracy of the technology's performance, can be a challenge, which can deter the adoption and efficacy of the technology (Lee and See, 2004). Previous research on algorithmic decision-aids suggests mixed results on people's acceptance of algorithmic recommendations. Some research reports that people trust algorithmic decisions more than human-made decisions (Madhavan and Weigmann, 2007). Other research suggests that people trust their own judgment more, especially after the algorithm has erred, probably because they believe algorithms cannot learn from their mistakes (Dietvorst et al., 2015). The differences in these findings may stem from the different natures of the tasks tested, such as predicting risk in luggage screening or stock prices vs. predicting students' success in college; the former may be perceived as requiring more mechanical skills, and the latter, more human skills. We posit that in tasks that require mechanical skills, people will trust algorithmic decisions as well as human decisions. On the other hand, in tasks that require human skills, people will trust algorithmic decisions less as they do not believe that algorithms have the capacity to successfully execute the tasks.

H2. Algorithmic decisions are trusted as well as human-made decisions in tasks that require mechanical skills but are trusted less than human-made decisions in tasks that require human skills.

Emotional responses

Affective experiences on the job contribute to job satisfaction (Weiss, 2002) and work motivation (Seo et al., 2004). One of the core differences between human and algorithmic decision-makers that may influence people's emotional responses to decisions is the presence (or lack of) intentionality. We posit that people will react to algorithmic decisions less emotionally because people attribute less agency and intentionality to algorithms. Previous research suggests that intentionality plays an important role in how people interpret other people's behavior (Clark, 1996), and the interpretation of the intention influences people's emotional responses to others' behaviors. For example, people self-reported greater pain when they thought that other people intentionally chose to give them an electric shock (Gray and Wegner, 2008). We expect that people will not perceive intention from algorithmic decision-makers, which will weaken their emotional response. This effect will not depend on the task type, as the lack of intentionality is not influenced by whether the tasks require mechanical or human skill sets.

H3. Algorithmic decisions evoke less emotional response than human-made decisions.

Method

We conducted a between-subjects online experiment in September 2016. Participants examined a managerial scenario (Mintzberg, 1975) in which a decision was made by either humans or algorithms. We examined the influence of the type of decision-maker on people's perceptions of the decisions by collecting both quantitative ratings of the decisions and qualitative reasons behind those ratings. We used a scenario-based method, commonly used in social psychology and ethics research, which investigates people's opinions, beliefs and attitudes (e.g. Petrinovich et al., 1993); studies have suggested consistency between people's behaviors in scenario-based experiments and their behaviors in real life (Woods et al., 2006).

Participants

We recruited participants on Amazon Mechanical Turk (MTurk) to take an online survey that took 6.2 minutes on average to complete. Participants had to reside in the US, be at least 18 years old, have completed at least 100 Human Intelligence Task (HITs, MTurk's task unit) and have at least a 95% HIT approval rate. Participants were compensated \$1.00 for their time, a rate more than the minimum wage in the US (\$7.25/

hour). 321 people responded. We omitted participants who filled out the survey in less than 2.5 minutes ($N=10$), completed the survey multiple times ($N=4$), indicated that they did not reside in the US ($N=7$), were younger than 18 ($N=9$), did not pass the first attention check ($N=37$), did not pass the consecutive attention checks ($N=22$), or did not finish the survey ($N=4$). These procedures left 228 participants in the sample, about three quarters of the original number. Their ages ranged from 19 to 77, with an average age of 35.9 ($SD=11.63$). Forty-five percentage of the sample was female. They were fairly well educated: the mean education score was 4.07 ($SD=1.36$; 3="Some college or currently enrolled", 4="2-year degree"). The sample was 79% Caucasian, 8.77% Asian, 6.58% Hispanic, 3.81% African American, 2.2% Native American, and 3.5% other.

Materials

We presented participants with four scenarios in which either algorithmic or human managers made decisions that had significant impacts on human workers (Table 1). Our scenarios were based on real-world situations in which algorithms are being employed in middle-management roles connected to the main functions of management (Mintzberg, 1975). The work assignment and scheduling scenarios involved managerial decisions that required "mechanical" skills. The work assignment scenario concerned maintenance tasks on a factory floor, which used algorithms to predict which components of machinery were likely to malfunction in order to prevent a larger stoppage in the overall workflow (Jardine et al., 2006) (Table 1(a)). The scheduling scenario used the example of an employee shift scheduling algorithm, which determined when café baristas would be called into work based on the predicted number of customers in the café at a given time (Table 1(b)) (Kantor, 2014; Pinedo et al., 2015).

The hiring and work evaluation scenarios involved managerial decisions that require "human" skills. The hiring scenario was based on job search websites, such as LinkedIn, that use algorithms to analyze resumes and select top candidates for onsite interviews (Table 1(c)). Finally, the work evaluation scenario involved a customer service call center that used a natural language-based algorithm to evaluate the performance of its employees (Petrushin, 1999) (Table 1(d)).

We used the projective, third-person viewpoint when creating scenarios so that participants were reading scenarios that describe another person's experience (e.g. Chris works [. . .]) as supposed to ones that directly put the readers into the scenario (e.g. you work [. . .]). The projective viewpoint has been shown to minimize

Table 1. Managerial scenarios presented to participants.

Managerial decision type	General policy and scenario
Tasks that require mechanical skills	
(a) Work assignment	<p>General policy: In a manufacturing factory, an/a (algorithm/manager) assigns their employees to check and update certain components of the machinery to prevent any critical operation failures. The component assignment is based on data that show how often different components have worn out and broken down in the past.</p> <p>Specific scenario: Chris works in the manufacturing factory. The (algorithm/manager) assigns him to check a specific component of the machinery and he does the maintenance work on it.</p>
(b) Work scheduling	<p>General policy: In a coffee shop, an/a (algorithm/ manager) decides who should come into work when on an hourly basis based on the predicted number of customers at different times of the day.</p> <p>Specific scenario: Riley works in the coffee shop. One day the (algorithm/ manager) calls him in to work in the afternoon with a few hours notice based on the increasing customer numbers.</p>
Tasks that require human skills	
(c) Hiring	<p>General policy: In a high-tech company, an/a (algorithm/manager) reviews resumes and personal statements on a job search website for each open engineering position. The (algorithm/manager) selects the top applicants for final on-site interviews.</p> <p>Specific scenario: Alex applies for an engineering position on the job search website by submitting his resume and personal statement. The (algorithm/manager) reviews thousands of applicants including Alex.</p>
(d) Work evaluation	<p>General policy: In a customer service center, an/a (algorithm/manager) evaluates employees by analyzing the content and tone of their calls with customers.</p> <p>Specific scenario: Jayln works at the customer service center. Based on past call recordings, the (algorithm/manager) evaluates his performance.</p>

social desirability effects, or the desire to present socially desirable answers rather than honest opinions, and to have considerable external validity (Nisbett et al., 1973). The scenarios had two parts. The first part described the general policy of management, and the second part described a specific instance of the policy involving a worker. We manipulated the type of decision-maker (algorithmic or human) for all scenarios.

We conducted a pilot test to check which decision-makers people thought would perform each task better and whether those perceptions were in line with our assumptions; specifically, whether people thought algorithms would do the mechanical tasks equally well or better than human managers and human managers would do the human tasks better than algorithms. The within-subjects survey ($N=21$) presented the general policy portion of the four scenarios in a random order and asked which type of decision-maker would perform the task better (1 = “Algorithms will perform the task better than human managers,” 2 = “Both algorithms and human managers will perform the task equally well,” 3 = “Human managers will perform the task better than algorithms”). We ran a multi-level analysis controlling for individuals for repeated measures.

The results show that people thought that algorithms would perform the work assignment and scheduling tasks slightly better than the human managers ($M_{\text{assign}} = 1.4$ ($SE = .15$), $M_{\text{schedule}} = 1.65$ ($SE = .15$)), and that human managers would perform the hiring and evaluation tasks better than algorithms ($M_{\text{hire}} = 2.5$ ($SE = 1.5$), $M_{\text{evaluate}} = 2.9$ ($SE = 1.5$)).

Procedure

After consenting and affirming that they were over the age of 18 and a US resident, participants were given an attention check question adopted from Egelman and Peer (2015). Failure to correctly answer the attention check resulted in immediate disqualification from the survey. Those who passed the check were randomly assigned to a decision-maker (human/algorithmic) in one of the four task scenarios. Participants in the algorithmic condition were shown this definition of “algorithm”: “Algorithms are processes or sets of rules that a computer follows in calculations or other problem-solving operations. In the situation below, an algorithm makes a decision autonomously without human intervention.”³ We presented the definition to ensure that participants had a similar definition in mind, and

bolded the word “computer” to emphasize that the decision-makers were computers rather than people or people using computers. Participants were then presented with a scenario, followed by survey questions about their thoughts on the scenario they had just read. The manipulation check and demographic questions were asked at the end.

Measures

Perceptions of decisions. Except for a few open-response questions, all survey items used a 7-point Likert-type scale. Response options varied based on the questions (i.e., Strongly Disagree, Very Unfair).

Decision fairness. The question on the decision’s fairness was adopted from previous research (Brockner et al., 1994; Konovsky and Folger, 1991): “How fair or unfair is it for [scenario subject] that the [algorithm/manager] takes the action specified in the scenario?” For example: “How fair or unfair is it for Jayln that the manager evaluates his performance?” The scale ranged from “Very unfair” (1) to “Very fair” (7).

Trust. To ascertain subjects’ trust in the reliability and accuracy of the decision presented in each scenario, we asked, “How much do you trust that the [algorithm/manager] make good-quality [decision specified in the scenario]?” The scale ranged from “No trust at all” (1) to “Extreme trust” (7).

Emotional response. To understand how participants thought the decision affected the scenario’s subject, we asked how much they agreed or disagreed that the decision-maker’s decision would make the scenario’s subject feel happy, joyful, proud, disappointed, angry, and frustrated (1 = “Strongly disagree”, 7 = “Strongly agree”) (Larsson, 2011; Weiss et al., 1999). We constructed an emotional response scale by averaging answers to the three positive adjectives and the reversed answers to the negative adjectives, such that greater numbers meant more positive emotion. The scale was very reliable (Cronbach’s $\alpha = .9$).

Open-response question. After each of the three questions above, we asked participants to explain their reasons for their numeric ratings.

Manipulation checks, attention checks, and demographic questions. At the end of the survey in the algorithmic decision conditions, participants were asked an open-ended question: “In your own words, please briefly explain what you think algorithms are.” The answers confirmed that participants perceived algorithms as autonomous decision-makers. Another manipulation

check question asked all participants: “Which of the following made the decisions in the situations that you read?” and provided a choice between humans and algorithms. All participants correctly answered this question by condition.

We used one attention check (Egelman and Peer, 2015) in the beginning of the survey to immediately disqualify participants. Throughout the survey, we also asked participants how much they agreed with the statement, “I do not read the questions in this survey.” At the end of the survey, we also asked participants to indicate their knowledge of algorithms (1 = “No knowledge at all”, 5 = “Expert knowledge of algorithms”) followed by demographic questions.

Analysis

We conducted a one-way ANOVA for the main effect of the decision-maker on perception of the decision for each decision scenario, and a multi-level analysis on the main and interaction effects of the decision-maker type and task types. We qualitatively analyzed participants’ reasons for their answers to the three questions about fairness, trust, and emotional response (Strauss and Corbin, 1990). We open-coded data at the response level in conjunction with participants’ survey ratings to identify emerging themes. We grouped different themes to explain how participants responded to and judged the decisions depending on the type of decision-maker.

Results

In this section, we present participants’ perceptions of the fairness, trustworthiness, and emotional effects of human-made and algorithmic decisions (Table 2). For each of the perception measures, we discuss whether relevant hypotheses were supported or unsupported for scenarios, grouped by task type (mechanical or human). We then summarize relevant results across conditions and discuss emergent themes in participants’ answers.

Fairness

The results suggest that H1 was only partly supported. As predicted, participants thought that algorithms’ and human managers’ decisions were equally fair on the mechanical tasks (Table 2(a) and (b)). On the other hand, human managers’ decisions were deemed fairer than algorithmic decisions on the human tasks (Table 2(c) and (d)). The open-response answers allow us to understand why and how people made their fairness judgments differently depending on the decision-maker.

Table 2. People's perceptions of human versus algorithmic managerial decisions.

Managerial decision type	Decision-maker type			
	Measures	Human	Algorithm	
Tasks that require mechanical skills				
(a) Work assignment	Fairness	6.21 (SE = .22)	5.90 (SE = .20)	NS
	Trust	5.33 (SE = .18)	5.52 (SE = .17)	NS
	Emotion	4.7 (SE = .22)	4.61 (SE = .20)	NS
(b) Work scheduling	Fairness	3.96 (SE = .33)	3.61 (SE = .33)	NS
	Trust	3.66 (SE = .31)	3.29 (SE = .30)	NS
	Emotion	3.07 (SE = .19)	2.86 (SE = .18)	NS
Tasks that require human skills				
(c) Hiring	Fairness	5.78 (SE = .23)	4.2 (SE = .24)	$F(1,59) = 21.76, p < .0001$
	Trust	5.37 (SE = .24)	3.55 (SE = .25)	$F(1,59) = 26.96, p < .0001$
	Emotion	4.55 (SE = .25)	3.75 (SE = .27)	$F(1,59) = 4.71, p < .05$
(d) Work evaluation	Fairness	6.11 (SE = .22)	3.39 (SE = .20)	$F(1,57) = 81.6, p < .0001$
	Trust	5 (SE = .24)	2.48 (SE = .23)	$F(1, 57) = 59.16, p < .0001$
	Emotion	4.39 (SE = .18)	3.23 (SE = .17)	$F(1, 57) = 22.13, p < .0001$

Algorithmic and human decisions are equally fair for mechanical tasks. Participants judged both algorithmic and human decisions to be similarly fair for the mechanical tasks, but participants' reasons for their fairness judgments varied depending on the decision-maker. In the work allocation scenario, when a human manager assigned a task, participants tended to view fairness in terms of the manager's authority or the employee's duties.

"We have to assume that the manager is in that position because he is more qualified. Therefore he would be more qualified than Chris to know which job he should do. Chris took the job as a worker. Therefore it is neither unfair or fair that he is assigned a task. It is just part of the job." (P2)

"I trust the manager knows best and knows which employee will function at a high level with a specific task, in this case Chris and the maintenance task." (P13). On the other hand, when an algorithm assigned a task, many participants focused more on the characteristic of the algorithm itself as an unbiased and efficient decision-maker. "The algorithm is merely following a set of instructions. It has no bias." (P130). "A machine is assigning the task based on a set of rules so all employees should be treated the same way." (P136). "I obviously don't have an intimate knowledge of the algorithm used, but I'm sure it was more random and fair than a human's choice." (P140).

With the work scheduling task, participants seemed to have felt the situation itself was unfair in the sense that the scenario subject should not have had such short notice about his work shift. Therefore, most participants' answers, regardless of the decision-maker

type, focused on the unfairness of calling in workers on short notice. However, some of the answers regarding the algorithmic decision-maker made an interesting comparison between the human manager and the algorithm by pointing out that the algorithm seemed unable to consider nuanced information about each individual worker's context and situation: "Because the algorithm doesn't take into account (as far as I know) how far Riley might have to travel, what plans he might have made for the day already, how much he's already worked for the week (or anything else), it's unfair." (P141). "I feel like the algorithm might not be as precise in handling the schedule like a human being and might over or under schedule." (P153). "It's not unfair assuming there's a basis for the computer's algorithm, but it's not very fair or fair because there could be errors that only a human could consider." (P159). "The algorithm is probably more impartial than a person would be, but there may be some unfairness due to failure to consider human concerns (e.g. maybe Riley is under more stress than other employees lately)." (P167).

Human decisions are fairer for tasks that involve human skills. H1 predicted that, with human tasks, participants would judge algorithmic decisions to be more fair than human decisions because they would be less affected by subjective biases. The results suggest the opposite: participants thought that human decisions were fairer than algorithmic decisions in the hiring and evaluation tasks. In the hiring task, as in the mechanical tasks, the majority of participants attributed the fairness of the decision to the authority of the manager's position. Participants thought that the human manager's decision was fair

because they expected the manager would be able to identify top candidates based on “skills and experiences” and “merit,” and would have the qualifications and authority to do so. In addition, participants mentioned that every applicant had to go through the same process, which made the outcome fair.

Some participants thought the manager’s decision was neither fair nor unfair, because basing the decision on documents was not really fair (e.g. “I feel that some people read well on paper, but don’t live up to par in real life. And vice versa” (P52)) and the manager’s judgment of the top applicants would be subjective. However, they were a small subset of the responses, and these perceptions did not significantly influence the overall fairness rating. Some also mentioned that the manager may overlook certain candidates, but did not think that this made the process unfair: “It is possible that the manager may overlook Alex. However, he ultimately has [the] same chance to be selected as anyone else.” (P59).

When an algorithm made the hiring decision, on the other hand, about half of the participants thought that the decision was unfair. Most participants thought that the algorithm would not have the ability to discern good candidates because it would lack human intuition, make judgments based on keywords, or ignore qualities that are hard to quantify. “An algorithm or program cannot logically decide who is best for a position. I believe it takes human knowledge and intuition to make the best judgment.” (P169). “He could have a great resume but not hit the key phrases.” (P174).

“Alex may feel like he has some qualifications that are not measurable through a computer program. For example, he might have the perfect personality and demeanor for the position, but the program cannot judge this as it is not quantifiable.” (P177).

“It’s unfair because he shouldn’t be judged by science or mathematics in using the process. [I]t should be a human.” (P170). The other half of the participants thought that the algorithmic decision was somewhat fair or fair. Some participants accepted it because it was the method that the company had chosen; other participants thought that it was fair because an algorithm is more efficient than a human; and a couple of participants mentioned that the algorithm’s decision would not be based on “social skills” or “favoritism.”

“It would take forever for a human to go through all of the applications. In the interest of time it is good to have an algorithm to do this. Assuming the algorithm assess all resumes the same way, I believe it is a somewhat fair selection process.” (P186).

With the evaluation task as well, participants thought that human decisions were fairer than algorithmic decisions. Like in the other tasks, most participants thought that evaluating worker performance was a manager’s

responsibility and the manager would therefore have the skill and authority to carry out the task. On the other hand, the majority of the participants thought that the algorithm’s performance evaluation would be unfair. Most participants thought that algorithms are not capable of evaluating tones of voice or human interaction, and were worried that a few errors would unfairly lower performance evaluation. “I do not believe a computer can evaluate how a human interacts with other humans fairly.” (P200). “Every day is a new day. [T]his person could have had a couple off days and that would influence the algorithm negatively, even though this person might be awesome.” (P214). “Human interaction and performance can’t always be analyzed mathematically without taking into consideration context and other non-quantifiable variables.” (P219).

A small set of participants mentioned that algorithmic evaluation is somewhat fair or fair because it is based on rules and is unbiased, but all acknowledged the limitations of the algorithm. “As long as the employee knows they are being evaluated in this way, I don’t see a problem with it. However, I’m not sure that I believe a computer can evaluate this better than a human ear.” (P223).

“By measuring things like tone and content, the computer does have a slight room for error if the employee were to joke around with the customer. On the other hand, the computer would be completely unbiased in its evaluation.” (P228).

“It’s based on the comparison of data to a set of rules. There will be times when the rules don’t perfectly match the situation that Jayln’s job dictated. People program computers and people are imperfect.” (P225).

Trust

H2 about trust was supported. Participants thought that both algorithms’ and human managers’ decisions were equally trustworthy on the work assignment and scheduling tasks (Table 2(a) and (b)). On the other hand, human managers’ decisions were trusted more than algorithmic decisions on the hiring and evaluation tasks (Table 2(c) and (d)).

Algorithmic and human decisions are equally trustworthy for tasks that involve mechanical skills. Participants trusted algorithmic and human decisions equally with mechanical tasks. There were some noticeable differences in the open-response answers concerning the source of trust between human and algorithmic decisions. In the work assignment task, as with the fairness judgment, participants mostly cited the human manager’s authority and the algorithm’s reliability and lack of bias as their reasons for trusting the decisions. However, participants mentioned a potential error or bias as a reason for not

giving algorithms complete trust, whereas there was no mention of potential human mistakes. “I trust computers to do their job but I didn’t give it 100 percent trust because there are glitches in the system at times” (P116); “Overall I trust the algorithm, but sometimes accidents happen that the algorithm cannot account for” (P128); and “I don’t know the rules that the algorithm is using. It could still be biased if it is programmed that way” (P135).

For the scheduling decision, half of the participants deemed the decision unfair because it was made last-minute, and thus they did not trust the manager to make a good decision next time. Interestingly, the other half still showed trust in the human manager’s decisions. These participants seemed to trust that the manager was making decisions based on both the company’s and the employees’ interests, and believed that even though the decision was last-minute, it was something that had to be done.

“The manager is trying to do what is best for the business, which is a large part of the position. It is not always easy to make decisions you know people will not like, but it has to be done regardless.” (P40).

“I expect that they do so based on the best interests of their and employees’ long-term success.” (P50). For the algorithmic scheduling decision, participants’ trust was similarly compromised by the last-minute nature of the decisions: “The Algorithm is missing the human aspect of things people would want like early notice” (P142).

A small subset of participants somewhat trusted algorithms because algorithms are based on rules and therefore seemed capable of scheduling tasks: “It is probably more consistent than a manager, if it’s programme[d] without bias” (P166); “It’s a logical way to make a decision to call someone in. It’s unbiased and goes by a set of rules” (P163); “The algorithm is only as good as its data and programming. [Its] failure to provide a timely schedule in this instance leads me to believe it is deficient; however, even an average algorithm should be able to do a fair job of scheduling” (P168). But these perceptions did not improve the overall trust rating.

Human decisions are more trustworthy for tasks that involve human skills. On human tasks, participants trusted human decisions more than algorithmic decisions. In the hiring task, most participants trusted the human manager’s hiring decision because it was his job, and thus he seemed qualified and had the authority to do it, consistent with answers for the fairness judgment: “It is in his best interest to make a quality decision. Usually people in a company who are in a particular position are capable of performing the job” (P74). While a few mentioned that they did not trust the decision because the manager’s decisions could be

biased or influenced by fatigue, this group comprised only a small set of the responses.

For the algorithmic hiring decision, the majority did not trust the algorithm. “Algorithms cannot apply exceptions and such and can only reliably do crude sorting (Even supposed AI).” (P182). “Too many factors can cause a candidate to be discarded: their name could indicate ethnicity, work and graduation dates could indicate age and make discriminatory exclusions.” (P180).

Some participants trusted the algorithmic decision, however, as long as the algorithm had been designed carefully. But these answers represented a small set of responses, and did not increase the overall trust rating. “I feel that if the algorithm is very carefully designed it can help remove harmful prejudices from the hiring process such as [bias based on] race, gender, or sexuality.” (P189).

“I think the algorithm is more than capable of assessing the applicability of people [for] the job. I.e., it ought to be able to easily assess GPA, and the number and duration of previous work experience. Therefore, I moderately trust it, barring any glitches in the technology. (P186).

Emotion

Our results suggest that H3 on emotional responses was not supported. H3 predicted that participants would have stronger emotional responses to human decisions than to algorithmic decisions because algorithms lack intentionality. On the contrary, participants felt similar or more negative emotion toward algorithmic decisions as compared to human decisions. On tasks that involved mechanical skills, participants did not differ in their emotional responses to algorithmic versus human decisions (Table 2(a) and (b)). On tasks that involved human skills, participants reported that they felt more negative about the algorithmic decisions than the human decisions (Table 2(c) and (d)).

Algorithmic and human decisions evoke similar emotion for tasks that involve mechanical skills. In the work assignment task, when asked to explain why they thought the scenario subject would feel either positively or negatively about the decision, many participants said there would be no reason for the scenario subject to feel any strong emotion because he had merely been asked to do “part of his job” (P18) or “his regular job” (P16). This kind of explanation was dominant in the work assignment scenarios regardless of the type of decision-maker. “It’s just a job and probably does not evoke strong emotion either way.” (P17). “He probably doesn’t feel anything or think about it. It’s just part of his normal duties.” (P114).

However, there were some noteworthy explanations in both the human manager and the algorithm scenarios, which could help us understand potential differences in the way people might feel about human versus algorithmic decisions. A couple of participants mentioned social recognition in the human manager scenario as follows: “Chris would feel proud that he was selected to do maintenance for a vital task, and he would not be angry or frustrated [because] his manager has faith in him and his performance.” (P13). “I can’t imagine being excited about getting assignments but I imagine that Chris might take pride in his job and the trust that his manager is placing in him.” (P20).

In the algorithm scenario, participants seemed to have mixed perceptions of the algorithm. Some expressed concern about having an algorithm as a decision-maker in the workplace, as the algorithm might make employees feel that they lacked “agency” (P121) or were “being watched” (P113). Others perceived the algorithm as a mere tool and therefore appeared to expect it would “help” (P119, P132) or make the job “easier” (P117, P128) or “take some of the stress away” from the job (P125).

In the scheduling task, participants generally reported negative emotion because of the unjust nature of the situation; the reasons for this emotion did not differ by decision-maker.

Algorithmic decisions evoke more negative emotion for tasks that involve human skills. In the hiring task, participants on average reported neutral emotion about the human manager’s hiring decision. Some participants expressed somewhat negative emotion because of the selective process in which “Alex is having to compete against thousands of other applicants” (P79). Others expressed somewhat positive emotion because Alex was being considered as a candidate for the job: “He would be happy that he is being considered for this great chance” (P64).

On the other hand, with the algorithmic hiring decisions, most participants expressed negative emotion. Not being reviewed by a human was one of the major sources of the negative emotion. “Alex may feel that there was no human even looking at his resume and [it] was probably discarded as soon as received.” (P180). “I think that Alex would feel frustrated knowing that no human was going to be reviewing him.” (P175). “He would feel like a machine being chosen by a machine.” (P188).

With the human evaluation decisions, some participants felt somewhat negative, as being evaluated is generally unpleasant: “Most people don’t feel much joy or happiness when being critiqued” (P88). However, most participants thought that evaluation was part of the job, which reflected their trust in the process: “Jayln

I’m sure understands that it is her manager’s job to evaluate her performance” (P108); “Because he knows that he will be evaluated . . . he will put his best voice forward and be good. He knows whether he did a good job or not and should have nothing to be upset about” (P104).

With the algorithmic evaluation decisions, most participants expressed negative emotion. Some responses suggested that the fact that a machine evaluated a person was demeaning and disrespectful. “I doubt Jaylyn would enjoy being evaluated by a machine” (P201); “He would know it would be wrong plus that is disrespectful” (P198). Other responses suggested that participants felt negative because they did not trust nor find it fair that algorithms could make evaluation decisions: “An algorithm would miss the ‘person’ in customer service” (P202); “I’m sure it would irritate them to be rated by a program. It’s not a person and cannot evaluate someone well” (P206); “I don’t think Jayln would appreciate being reviewed by a computer that is possibly prone to errors that could cost him his job” (P217).

Discussion

The results suggest that task characteristics—in particular, perceptions of whether tasks require more “human” or more “mechanical” skills—significantly influence how people perceive algorithmic decisions compared to human-made ones. With tasks that mainly involve mechanical skills, participants trusted algorithmic and human decisions equally, found them fair, and felt similar emotion toward them, consistent with our hypotheses. While the degree of perceived trust, fairness, and emotion was the same between algorithmic and human decisions, the reasons behind people’s perceptions differed. With human-made decisions, participants attributed fairness and trust to managerial authority; with algorithmic decisions, to reliability and the lack of bias. For the human-made decisions, some participants mentioned the manager’s social recognition as a factor that could positively influence workers’ emotions. For algorithmic decisions, on the other hand, participants mentioned that algorithms could act as tools to help workers complete their tasks, which could positively influence workers’ emotions; or workers might feel negatively about algorithms, if they felt they were being watched and monitored.

With tasks that require human skills, participants’ perceptions differed between algorithmic and human-made decisions. As in mechanical tasks, participants attributed fairness, trust, and positive emotions for human decisions to the authority of the manager’s position and the social recognition implied by the manager’s choice. However, participants judged

algorithmic decisions as less fair, trusted algorithmic decisions less, and felt less positive toward algorithmic decisions than human decisions. Participants felt that algorithms are incapable of discerning good candidates for jobs or evaluating worker performance because they lack human intuition, only measure quantifiable metrics, and cannot evaluate social interaction or handle exceptions. Some thought it was demeaning and dehumanizing to use machines to judge a person. A few participants in the algorithmic decision condition mentioned they trusted the algorithmic decisions because the organization chose to use the algorithm and the algorithmic process prevents favoritism and human biases. These responses are consistent with Sundar and Nass (2001)'s finding about people perceiving computers as more objective than human news editors. However, these opinions remained the minority.

Limitations

The study has several limitations which future work should address. We used a survey experiment based on hypothetical situations. While this scenario-based method is commonly used in social psychology and ethics research to study perceptions of decisions (Petrinovich et al., 1993), the findings of this study need to be complemented with other studies that involve people's actual experiences. Lee and Baykal (2017) examined people's actual experiences with algorithmic versus human decisions. The authors compared how people perceived task division decisions made by algorithms versus humans using an actual task in a laboratory, and found that algorithmic decisions were perceived to be less fair than human decisions. More studies would need to be conducted in real-world settings with those who are affected by algorithmic management in order to confirm these findings and build systematic theories on when and how people perceive algorithmic and human decisions similarly or differently in management contexts. Because the decision tasks and situations were drawn from real-world practices, we could not exercise complete control over task characteristics. Future work should investigate different task types and their impact on perceptions of decisions and decision-makers. Only four managerial decisions were used, and all judgments were made at one point in time; our hypotheses should be tested with different decisions and outcomes over time. In this first study, we focused on human versus algorithmic decisions without additional contextual information (such as managers' education or algorithms' programmers) in order to compare differences in judgments that arise from the knowledge of who the decision-maker is alone. However, contextual information about algorithms

may influence people's perceptions, and future research needs to explore the roles of people in the creation and operation of algorithms. Along the same lines, the goal of the study was to understand the general public's perceptions of algorithmic decisions when the algorithm is presented as a "black box" (without specific details of the mechanics), as is currently done in most algorithmic workplace applications. For that reason, we used a dictionary definition of "algorithm" in the study for people who may not have known the term. This definition may have suggested to people that algorithms are neutral. In our future studies, we will explore how different descriptions of algorithms may influence people's perceptions.

Finally, the survey was conducted on Amazon Mechanical Turk. Future work should be tested using different sampling techniques. MTurk populations have been reported to be large and diverse (Paolacci and Chandler, 2014) and more representative of the US population than in-person convenience samples, but less representative than Internet-based panels or national probability samples. Some published experimental work has been replicated using MTurk samples (Berinsky et al., 2012). However, to our knowledge, the unique biases of the mTurk population with respect to organizational psychology and perceptions of fairness are not yet known. Studies on the MTurk population suggest that MTurk workers tend to be young and liberal, which means they may be more open to technological change and algorithmic decision-makers (Berinsky et al., 2012; Paolacci and Chandler, 2014). The study findings therefore need to be evaluated with other populations, especially with workers who are likely to be or are currently affected by algorithmic managerial decisions; with managers themselves; and with people with varying levels of knowledge about algorithms.

Implications and future research

Our study offers implications for theory and practice, and future research questions.

Implications for theory

Our research offers two main implications for theories of algorithms, automation and intelligent technologies. First, our research contributes to emerging studies that investigate people's mental models and folk theories of algorithms (Bucher, 2016; Eslami et al., 2015; Rader and Gray, 2015), and social studies on automation, artificial intelligence and machine learning more broadly. Our results suggest that, regardless of the actual performance of algorithms, what people think

algorithms are capable of and their comparison with human decision-makers play important roles in people's judgments of trustworthiness and fairness, as well as their emotional responses. Our work also raises an ethical and sociological question about the impact of introducing algorithmic management on people's perceptions of particular tasks and jobs. Even when the degree of perceived fairness and trustworthiness was similar, the source of authority in decisions varied depending on the decision-maker: the manager's authority was attributed to their position in the organization, whereas the algorithm's authority was attributed to its efficiency and lack of bias. This difference in perceived source of authority may result in different behaviors around decisions not measured in this work, which require further investigation. In addition, our results suggest people currently feel that using algorithms and machines to assess humans could be demeaning and dehumanizing. This feeling may remain unchanged regardless of the actual performance of algorithms, and might deter the adoption of such algorithms. Alternatively, as industries continue to introduce algorithms, people may start to see tasks and jobs like these as more "mechanical" and hold them in lower esteem. Further research needs to be done in order to unpack this dynamic.

Second, our results also suggest a new construct not investigated in previous work on social understandings of algorithmic systems: *task characteristic*, particularly whether the task is better suited to human or mechanical skills. Previous research on decision aids and algorithmic recommendations showed conflicting results on people's trust in algorithmic decisions. Our results suggest that the different characteristics of the tasks used in previous experiments may explain the conflicting results. There is ongoing work that explores how decision-makers' expertise, the potential to affect oneself or others, and the subjectivity and objectivity of decisions can influence people's reliance on algorithmic advices (Logg, 2017). Our study joins this line of work to contribute to a more nuanced, context-sensitive understanding of human perceptions of algorithmic decisions.

Further research needs to be done in order to understand what contributes to the perception that certain tasks can be done well uniquely by humans. People's attitudes toward and perceptions of technologies have changed throughout history; some technologies originally considered to be socially awkward, rude, or unacceptable were eventually adopted as perceptions changed, or designs were improved to better fit human conceptions. Our study assesses people's contemporary perceptions of what "human" tasks are, and the limits of algorithmic technologies. We acknowledge that the kinds of tasks that people think only humans can do

will change; for example, speech recognition was a task that could previously be performed only by humans, but is now reliably performed by computer algorithms. Evaluating a person's potential based on an application or understanding social interactions might soon be tasks that algorithms can perform as well as or better than people. In this case, whether the decision-maker is human or algorithmic might not affect perceptions of hiring and evaluation decisions. However, we believe that the distinction people perceive between capabilities that are and are not uniquely human will remain a factor in their perceptions of algorithmic decisions so long as there are areas in which humans outperform machines. This distinction can be a predictor of people's perceptions of algorithmic decisions, even as the specific tasks themselves might change.

Another avenue for future research is to investigate the role of interaction, and how the outcomes and inputs of algorithmic decisions affect people's perceptions. Our study investigated people's perceptions of algorithmic decisions when algorithms were embedded in organizational contexts; participants did not directly interact with the algorithm and the algorithmic decision-maker was not portrayed as an agent. Whether algorithmic decisions are delivered through human-like interactions such as an interactive chat, along with whether the decision outcomes are positive or negative, may cause people to perceive algorithmic decisions differently. For example, one study (Shank, 2012) suggests that people perceived behaviors of computer agents as more just than human agents when both agents acted coercively to them. This suggests that if a computer agent informs people of unwelcome hiring decisions, people may perceive the decision to be fairer than the same human decision. Another factor may be the kind of input used for algorithmic decisions, specifically whether the input data is in numeric forms or not (Jago, 2017), as this may influence whether people think the decision tasks require mechanical or human skills. Future studies should unpack these interactional and situational factors in human perceptions of algorithmic decisions.

Implications for practice

Our research offers implications for practice. A recent article written by Crawford and Calo (2016) offers a critical perspective on current trends in algorithms and artificial intelligence in industry. They argue that people fear that artificial intelligence is taking over human jobs, when in fact the problem is that industries often incorporate technology whose performance and effectiveness are not yet proven, without careful validation and reflection. Our results reinforce this argument

that the general public does not fully trust algorithms or find it fair to use algorithms for decisions that involve subjective judgments of human workers.

Our results also shed insights on the upsides and downsides that people perceive in using algorithms for managerial decisions in organizations. Many in our study believed that algorithms could remove favoritism or human biases from managerial processes. They also mentioned algorithms' inability to accommodate exceptions, measure human properties commonly believed to be non-quantifiable (such as social interactions and personalities), or consider human concerns such as empathy and personal commitments, all of which contributed to distrust of algorithms and feelings of unfairness. Addressing these concerns both in the actual implementation of algorithms and in communication about algorithms to users can help us create workplaces that are efficient but also that workers can trust, find fair, and feel good about.

Conclusion

Algorithms are increasingly being introduced into online and offline workplaces and are used to manage interactions among human workers, taking on tasks that human managers used to do. The work presented in this paper explored how algorithmic managers as compared to human managers influence workers' perceptions of decision fairness, trustworthiness, and emotional response in tasks that require human or mechanical skills. The results of our online scenario-based experiment suggest that people perceive algorithmic decisions as less fair, less trustworthy, and more likely to evoke negative emotion for tasks that people think require uniquely human skills. This study offers preliminary support for the claim that algorithmic decision-makers evoke different beliefs and associations for workers than human managers do. This study is just a first step toward understanding how we can design better workplaces, where people and intelligent machines can work together.

Acknowledgement

We thank the organizers and the participants of the Algorithms in Culture 2016 Conference and the anonymous reviewers for their thoughtful feedback that improved our paper. We also thank Laura Dabbish, Nathaniel Fruchter, Su Baykal, Daniel Kusbit, and Evan Metsky for their helpful discussion and comments and Jiyeon Kim for assistance with data collection.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received financial support for the research, authorship, and/or publication of this article: This research was partially supported by NSF grants CNS-1205539 and ACI-1322278.

Notes

1. <http://talent.linkedin.com/recruiter-homepage-tour/?src=s-bl>
2. Oxford Dictionaries <http://www.oxforddictionaries.com/>
3. Oxford Dictionaries <http://www.oxforddictionaries.com/>

References

- Amino. Available at: <https://amino.com> (accessed 13 February 2018).
- Azadeh A, Saberi M, Rouzbahman M, et al. (2013) An intelligent algorithm for performance evaluation of job stress and HSE factors in petrochemical plants with noise and uncertainty. *Journal of Loss Prevention in the Process Industries* 26(1): 140–152.
- Babcock L, Loewenstein G, Issacharoff S, et al. (1995) Biased judgments of fairness in bargaining. *The American Economic Review* 85(5): 1337–1343.
- Barocas S and Selbst AD (2016) Big data's disparate impact. *California Law Review* 104(3): 671–732.
- Berinsky AJ, Huber GA and Lenz GS (2012) Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis* 20(3): 351–368.
- Brockner J, Konovsky M, Cooper-Schneider R, et al. (1994) Interactive effects of procedural justice and outcome negativity on victims and survivors of job loss. *Academy of Management Journal* 37(2): 397–409.
- Bucher T (2016) The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society* 20(1): 1–15.
- Carey D and Smith M (2016) How companies are using simulations, competitions, and analytics to hire. *Harvard Business Review*. Available at: <https://hbr.org/2016/04/how-companies-are-using-simulations-competitions-and-analytics-to-hire> (accessed 13 February 2018).
- Clark H (1996) *Using Language*. Cambridge: Cambridge University Press.
- Cloud 9 Psych. Available at: <http://www.cloud9psych.com> (accessed 13 February 2018).
- Crawford K and Calo R (2016) There is a blind spot in AI research. *Nature* 538(7625): 311–313.
- Dietvorst BJ, Simmons JP and Massey C (2015) Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144(1): 114–126.
- Egelman S and Peer E (2015) Scaling the security wall: Developing a Security Behavior Intentions Scale (SeBIS). In: *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, Seoul, South Korea, pp. 2873–2882..
- Electronic Privacy Information Center. Algorithms in the criminal justice system. Available at: <https://epic.org/>

- algorithmic-transparency/crim-justice (accessed 13 February 2018).
- Eslami M, Rickman A, Vaccaro K, et al. (2015) I always assumed that I wasn't really that close to [her]: Reasoning about invisible algorithms in the news feed. In: *Proceedings of the 33rd annual SIGCHI conference on human factors in computing systems*, pp. 153–162.
- Fogg B and Nass C (1997) Silicon sycophants: The effects of computers that flatter. *International Journal of Human-Computer Studies* 46(5): 551–561.
- French M and Hancock J (2017) What's the folk theory? Reasoning about cyber-social systems. Available at: <https://ssrn.com/abstract=2910571> (accessed 13 February 2018).
- Gillespie T (2014) The relevance of algorithms. In: Gillespie T, Boczkowski P and Foot KA (eds.) *Media Technologies: Essays on Communication, Materiality, and Society*, Cambridge: The MIT Press.
- Gray HM, Gray K and Wegner DM (2007) Dimensions of mind perception. *Science* 315(5812): 619.
- Gray K and Wegner DM (2008) The sting of intentional pain. *Psychological Science* 19(12): 1260–1262.
- Gray K and Wegner DM (2012) Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition* 125(1): 125–130.
- Greenberg J (1986) Determinants of perceived fairness of performance evaluations. *Journal of Applied Psychology* 71(2): 340–342.
- Hackman RJ and Oldham GR (1976) Motivation through the design of work: Test of a theory. *Organizational Behavior and Human Performance* 16(2): 250–279.
- Hinds PJ, Roberts TL and Jones H (2004) Whose job is it anyway? A study of human-robot interaction in a collaborative task. *Human-Computer Interaction* 19(1-2): 151–181.
- Hodson H (2014) The AI Boss that deploys Hong Kong's subway engineers. *New Scientist*.
- Idrees M, Iqbal W and Bazaz S (2013) Real-time doctor-patient assignment in a telemedicine system. In: *Proceedings of multi topic conference (INMIC)*, pp. 55–60.
- Jago A (2017) Technology and (In)discrimination. *Data Blitz presentation at the Psychology of Technology conference*. Berkeley, CA USA, 3–4 November 2017.
- Jardine AK, Lin D and Banjevic D (2006) A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing* 20(7): 1483–1510.
- Kantor J (2014) Available at: <http://www.nytimes.com/interactive/2014/08/13/us/starbucks-workers-scheduling-hours.html> (accessed 13 February 2018).
- Kitchin R (2016) Thinking critically about and researching algorithms. *Information, Communication & Society* 20(1): 1–16.
- Konovsky MA and Folger R (1991) The effects of procedures, social accounts, and benefits level on victims' layoff reactions. *Journal of Applied Social Psychology* 21(8): 630–650.
- Kramer RM (1999) Trust and distrust in organizations: Emerging perspectives, enduring questions. *Annual Review of Psychology* 50(1): 569–598.
- Larsson G (1987) Quick measurement of individual stress reaction level: Development of the Emotional Stress Reaction Questionnaire (ESRQ). Report from Defense Technical Information Center.
- Lee MK and Baykal S (2017) Algorithmic mediation in group decisions: Fairness perceptions of algorithmically mediated vs. discussion-based social division. In: *Proceedings of the ACM conference on computer supported cooperative work and social computing*, Portland, USA, 25 February–1 March 2017, pp. 1035–1048.
- Lee MK, Kusbit D, Metsky E, et al. (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*, Seoul, South Korea, 18–23 April 2015, pp. 1603–1612.
- Lee JD and See KA (2004) Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 46(1): 50–80.
- Leventhal GS (1980) What should be done with equity theory? In: Gergen KJ, Greenberg MS and Willis RH (eds.) *Social Exchange*. Boston, MA: Springer, pp. 27–55.
- Logg J (2017) Theory of Machine: When Do People Rely on Algorithms? Harvard Business School NOM Unit Working Paper No. 17-086. Available at: <https://ssrn.com/abstract=2941774> (accessed 13 February 2018).
- Madhavan P and Wiegmann DA (2007) Effects of information source, pedigree, and reliability on operator interaction with decision support systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 49(5): 773–785.
- Mintzberg H (1975) *The nature of managerial work*. New York, NY: Harpercollins College Div.
- National Institute of Justice (2014, June 9) Predictive policing. Available at: <https://www.nij.gov/topics/law-enforcement/strategies/predictive-policing> (accessed 13 February 2018).
- Nass C and Moon Y (2000) Machines and mindlessness: Social responses to computers. *Journal of Social Issues* 56(1): 81–103.
- Nass C and Steuer J (1993) Voices, boxes, and sources of messages. *Human Communication Research* 19(4): 504–527.
- Nisbett RE, Caputo C, Legant P, et al. (1973) Behavior as seen by the actor and as seen by the observer. *Journal of Personality and Social Psychology* 27(2): 154–164.
- Paolacci G and Chandler J (2014) Inside the Turk: Understanding mechanical Turk as a participant pool. *Current Directions in Psychological Science* 23(3): 184–188.
- Parise S, Kiesler S, Sproull L, et al. (1996) My partner is a real dog: Cooperation with social agents. In: *Proceedings of the 1996 ACM conference on computer supported cooperative work*, Boston, USA, 16–20 November 1996, pp. 399–408.
- Petrinovich L, O'Neill P and Jorgensen M (1993) An empirical study of moral intuitions: Toward an evolutionary ethics. *Journal of Personality and Social Psychology* 64(3): 467–478.
- Petrushin V (1999) Emotion in speech: Recognition and application to call centers. In: *Proceedings of artificial neural networks in engineering*, p. 710.

- Pinedo M, Zacharias C and Zhu N (2015) Scheduling in the service industries: An overview. *Journal of Systems Science and Systems Engineering* 24(1): 1–48.
- Prendergast C and Topel R (1993) Discretion and bias in performance evaluation. *European Economic Review* 37(2): 355–365.
- Rader E and Gray R (2015) Understanding user beliefs about algorithmic curation in the Facebook news feed. In: *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, Seoul, South Korea, 18–23 April 2015, pp. 173–182.
- Reber AS (1989) Implicit learning and tacit knowledge. *Journal of Experimental Psychology* 118(3): 219–235.
- Reeves B and Nass C (1996) *How People Treat Computers, Television, and New Media Like Real People and Places*. Stanford, CA: CSLI Publications and Cambridge University Press.
- Seo MG, Barrett LF and Bartunek JM (2004) The role of affective experience in work motivation. *Academy of Management Review* 29(3): 423–439.
- Shank DB (2012) Perceived justice and reactions to coercive computers. *Sociological Forum* 27(2): 372–391.
- Silverman RE and Waller N (2015) Available at: <http://www.wsj.com/articles/the-algorithm-that-tells-the-boss-who-might-quit-1426287935> (accessed 13 February 2018).
- Skarlicki DP and Folger R (1997) Retaliation in the workplace: The roles of distributive, procedural, and interactional justice. *Journal of Applied Psychology* 82(3): 434–443.
- Strauss AL and Corbin JM (1990) *Basics of Qualitative Research*. Newbury Park, CA: Sage.
- Sundar S and Nass C (2001) Conceptualizing sources in online news. *Journal of Communication* 51(1): 52–72.
- Sweeney L (2013) Discrimination in online ad delivery. *Queue* 11(3): 1–19.
- Waytz A and Norton MI (2014) Botsourcing and outsourcing: Robot, British, Chinese, and German workers are for thinking—not feeling—jobs. *Emotion* 14(2): 434–444.
- Weiss HM (2002) Deconstructing job satisfaction: Separating evaluations, beliefs and affective experiences. *Human Resource Management Review* 12(2): 173–194.
- Weiss HM, Suckow K and Cropanzano R (1999) Effects of justice conditions on discrete emotions. *Journal of Applied Psychology* 84(5): 786–794.
- Woods S, Walters M, Koay KL, et al. (2006) Comparing human robot interaction scenarios using live and video based methods: Towards a novel methodological approach. In: *Proceedings of 9th IEEE international workshop on advanced motion control*, pp. 750–755.
- Zuboff S (1988) *In the Age of the Smart Machine: The Future of Work and Power*. New York, NY: Basic books.