A Human-Centered Approach to Algorithmic Services: Considerations for Fair and Motivating Smart Community Service Management that Allocates Donations to Non-Profit Organizations

Min Kyung Lee

Center for Machine Learning &
Health, Machine Learning Dept.
Carnegie Mellon University
Pittsburgh, PA USA
mkleee@cs.cmu.edu

Ji Tae Kim

School of Design
Carnegie Mellon University
Pittsburgh, PA USA
jitaekim@andrew.cmu.edu

Leah Lizarondo

412 Food Rescue
Pittsburgh, PA USA
leah@412foodrescue.org

ABSTRACT

Algorithms are increasingly being incorporated into diverse services that orchestrate multiple stakeholders' needs and interests. How can we design these algorithmic services to make decisions that are not only efficient, but also fair and motivating? We take a human-centered approach to identify and address challenges in building human-centered algorithmic services. We are in the process of building an allocation algorithm for 412 Food Rescue, an organization that matches food donations with non-profit organizations. As part of this ongoing project, we conducted interviews with multiple stakeholders in the service—organization staff, donors, volunteers, recipient non-profits and their clients, and everyday citizens—in order to understand how the allocation algorithm, interfaces, and surrounding work practices should be designed. The findings suggest that we need to understand and account for varying fairness notions held by stakeholders; consider people, contexts, and interfaces for algorithms to work fairly in the real world; and preserve meaningfulness and social interaction in automation in order to build fair and motivating algorithmic services.

Author Keywords

Algorithmic services, fairness, allocation, motivation, community service, donation, service design, automation.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

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

CHI 2017, May 06 - 11, 2017, Denver, CO, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4655-9/17/05...\$15.00

DOI: <http://dx.doi.org/10.1145/3025453.3025884>

INTRODUCTION

Algorithms are increasingly being incorporated into services in diverse industries. They connect supply and demand in real time, matching drivers with riders and independent workers with clients [19, 35, 37, 38, 39]. They also optimize how resources are allocated in public services and urban planning, determining the locations of shared bicycle racks and electronic car charging stations or the routes of self-driving vehicles [2, 4, 23]. In these services, algorithms take on the role of managing multiple stakeholders' needs and participation. How can we build an algorithmic service that is not only efficient but also fair and motivating to participants? In these services, algorithms are often invisible; how do we ensure certain stakeholders are not overlooked, intentionally or unintentionally?

Understanding human considerations and requirements for these algorithmic services is a critical first step in addressing these questions. In response to the expanding applications of algorithms and emerging evidence that they can be unfair or unethical [15, 29, 34], scholars have increasingly sought ways to make these systems more beneficial to society; this line of research highlights the importance of defining what is desirable and ethical in the first place to inform the design of the technology [24, 27]. For example, even though “fair” algorithms have been developed in mathematics and economics, such algorithms may not be perceived as fair in social contexts, because “fairness” in society has diverse ontological roots [20].

When algorithms are integrated into services, further complications arise. Services involve multiple stakeholders [30], who may have conflicting expectations about fairness. Successful service delivery requires all these stakeholders' participation. In addition, previous studies suggest that choosing the “right” interface design and algorithmic assumptions about human behaviors and incentives is important for motivating participation [19, 37].

In this paper, we take a human-centered approach in order to identify considerations for building fair and motivating algorithmic services. We work with 412 Food Rescue, a

local organization that takes food donations, matches them with non-profits, and coordinates volunteers to pick up and deliver the donations to the recipient organizations. A human manager currently makes the allocation decisions, but as the organization is struggling to meet increasing demand, we are building an algorithmic service system to take on this role. Most algorithmic services exist in the commercial sector, in which both the design process and inner workings of the algorithms are hidden from the public; this project therefore offers an important and rare opportunity to explore the process and challenges of designing human-centered algorithmic services.

The algorithmic service in this case will improve the efficiency of the food rescue organization's work with limited resources. It also necessitates formalizing the heuristics and practices that the community manager uses to allocate donations and coordinate multiple stakeholders. Our research questions are: What does it mean for an allocation algorithm to be "fair" in this context? How can we design the automated aspects so that the service motivates people and works as well as or even better than the existing service? To address these questions, we interviewed multiple stakeholders in the service—the food rescue organization, donors, volunteers, recipient non-profits and their clients, and everyday citizens. Drawing from these interviews, we argue algorithmic service designers should understand and accommodate stakeholders' multiple fairness notions; consider people, contexts, and interfaces for algorithms to work fairly in the field; and carefully decide what and how to automate to preserve meaningfulness and social interaction.

Our work makes contributions to the field of human-computer interaction and emerging academic scholarship in the social and ethical implications of algorithms and artificial intelligence. It offers an example of a human-centered approach to algorithmic services, and identifies challenges and considerations for designing fair and motivating algorithmic service systems.

RELATED WORK

Our work is situated in the literature on algorithmic technologies, fairness, value-sensitive design and service design, and technologies for underserved populations and non-profit organizations.

Social implications of algorithmic technologies

As the use of algorithms is becoming more prevalent in many sectors of society, both scholarly and industry efforts are increasingly being made to understand the social implications of algorithmic technologies [11, 17, 19, 31] and how these algorithmic technologies can be designed to make ethical and moral decisions that benefit people and societies [22, 24, 27]. Our work contributes to this growing body of work by focusing on fairness.

Recent research has shown that machine learning algorithms can make unfair decisions due to the unintended

consequences of choices made in the development process or biases in the training data set [15, 34]. Even in the case of artificial intelligence, a recent study suggests that mathematically-proven fair division algorithms may not be perceived as fair because they do not match social concepts of fairness [20]. In response, interdisciplinary scholars have begun to investigate methods to achieve fairness in algorithmic systems [44]; Crawford and Calo recently pointed out a gap in current research on artificial intelligence, arguing for the importance of conducting social-systems analysis at every stage of development [5]. Our work fills this gap by presenting a human-centered approach to understanding fairness in algorithmic services in context at the conception and design stages.

Fairness

The key value that our work examines is fairness. In the Cambridge Dictionary, fairness is defined as "the quality of treating people equally or in a way that is right or reasonable." Fairness forms the basis of human society, and has multiple ontological bases. We review two main operationalizations of fairness.

Equality

An equality model ignores differences in recipient's circumstances and distributes resources or rewards equally [21]. Theories of equality are usually characterized by a concern for the welfare of those in society who are the least advantaged. By equally satisfying basic human needs, they form a principle of justice. Some social psychologists propose that equality is the fundamental principle of a multi-criterion system favored in cooperative, as opposed to competitive, relationships [18]. Studies on the effects of different allocation methods also support the idea that equally distributing rewards produces a high level of satisfaction and harmony among group members [1]. Equal allocation reduces negative socio-emotional behaviors, such as dissatisfaction and antagonism [21]. From the standpoint of the allocator, the task of choosing a recipient can become harder to make in some contexts, which is why some authors argue that an equality model might be the only just way to resolve indeterminacies in decision-making [32].

Equity

An equity model embraces individual differences and distributes resources in accordance with each recipient's need or performance. Walster et al. [41] posit the existence of one simple justice principle to which all notions of fairness can be reduced—namely, equity. Their view is that individuals consider an allocation fair if the outcome-input ratio of all group members is equal [26]. How an allocator defines and weighs need or performance is, however, context-dependent. In a blood donor matching system, patients' information, such as blood type, tissue type, size of the organ, medical urgency of the patient, time on the waiting list, and distance, are weighted and ranked in order of degree of need. In healthcare, a person's need depends on some combination of the degree of gravity, urgency, and entrenchment [16]. The assessment of an individual's need

or performance can also vary significantly depending on the allocator's ethnic background, age, and gender. In a study examining the effects of culture on equitable allocation, Indian respondents favored need alone much more than the Americans who favored equality or merit. Given the differences in the general level of available resources between India and the U.S., the results support the idea that culture impacts allocation decision-making [26].

Service design and value-sensitive design

Our approach draws on service design [30] and value-sensitive design [12]. Service design considers multiple stakeholders' needs and work processes in the design process. Value-sensitive design shows how design choices, intentional or unintentional, can lead to biases in decision-making or compromises in important user values. For example, Batya and Nissenbaum [14] explain how the order of items in a list can promote the selection of top items. Miller et al. [25] surveyed "value tensions" surrounding groupware in a company and design choices that can balance stakeholders' values. This orientation toward multiple stakeholders' values inspires our work. We focus on fairness in allocation, which has not been explicitly investigated through value-sensitive design methods.

Technologies for underserved populations and non-profit organizations

Many human-computer interaction researchers have explored the use and design of technologies for underserved populations [3, 6, 7]. Topically, the thread of research that addresses "food insecurity" is most relevant to our work. Dombrowski et al. [9] explored opportunities for designing location-based systems to address local food needs, and found that the role of outreach workers is critical in expanding food assistance programs [10]. She also highlighted the values of food justice, expanding food options, utilizing local resources, and localizing control [8], which we observe in our work. Previous research has highlighted the use of technology in non-profits and discussed how existing media and other technologies might not fit their purposes [40]. On the other hand, the food rescue organization that we work with is a new, non-profit start-up which embraces social media and mobile phone technologies to coordinate and execute its work.

BACKGROUND: SMART COMMUNITY SERVICE MANAGEMENT

We have been working with 412 Food Rescue, a local food rescue organization in Pittsburgh, since January 2016 as part of an ongoing project on improving services for food-insecure and homeless populations.

Food rescue organization

412 Food Rescue is a non-profit food recovery service established in March 2015. The organization works with restaurants, caterers, food retailers, wholesalers, and universities in Pittsburgh, Pennsylvania in the United States and "saves" food that is nearing its expiration date, is leftover from events, or is damaged, bruised, or partially

moldy. Since its launch, the organization has redistributed about one million pounds of food that otherwise would have been wasted to about two hundred forty community organizations, including food pantries, housing authorities, refugee camps, retirement homes, homeless shelters, and soup kitchens, to alleviate food insecurity and reduce food waste around Pittsburgh with the help of their volunteers. The organization prides itself on helping every donation reach a community organization in need. The organization's ability to quickly match donations with recipient organizations and coordinate volunteers for on-demand deliveries has been highly praised, and played a big role in their fast growth as an organization.

The core function of the food rescue organization is to connect donors, volunteers, and recipient organizations. When a donating organization calls the food rescue organization to notify them of leftovers to be picked up from their location, the manager asks about the size and contents of the food donation, and evaluates how many households or individuals the food can serve. The manager then calls community organizations that may be in need of the donation item. Once the food is matched to a recipient, the volunteer opportunity is posted on the food rescue organization's social media page, emailed to previous volunteers, or occasionally circulated by calling volunteers who have taken similar delivery routes in the past. The food retailer hands the donated food to the volunteers once they arrive on site and the volunteers deliver the food to the community organization the operation manager chose.

Decision to introduce an algorithmic service system

The food rescue organization has been experiencing growing pains, being a small team. Currently one community manager and an intern are responsible for matching donations with an appropriate recipient organization, as well as recruiting and coordinating volunteer efforts simultaneously. The task of orchestrating multiple stakeholders and coordinating volunteer efforts can become overwhelming for such a small team to do. It is particularly difficult when the coordinator has to find recipient organizations with available staff members to take donations and volunteers who can transport these items all within a very short time window. Donations come at random times throughout the day and occasionally multiple donations arrive at the same time. No matter the coordinator's skill level, it is simply not feasible for an individual to handle multiple requests simultaneously at scale. The organization also acknowledges that because of the demands for fast execution, many matching decisions are made for convenience, i.e., food is sent to organizations that the coordinator happens to remember at the time or that are known to pick up their phone. There is therefore room to improve the matching process, to distribute donations across diverse recipient organizations and make the distribution more "fair." For these reasons, the organization strategically decided to build an algorithmic system that could take on some of these roles. This system will have a

database of donating organizations, recipient organizations, and volunteers; when there is a donation request, an algorithm will recommend recipient organizations and once the manager selects one or confirms, it will send out a call to volunteers through a smart phone app.

RESEARCH QUESTIONS

In this context, we explore the following research questions: 1) The human manager uses her discretion and heuristics to make fair allocation decisions. An algorithm will require more formalized rules. To that end, what is “fair” allocation from the perspectives of multiple stakeholders? 2) Using this system will automate some of the decision-making tasks and human-human interactions in the service workflow. How can these automated aspects be designed so that the service motivates people and works as well as or even better than the existing service?

METHOD

To address the research questions, we conducted semi-structured interviews with thirty-one participants who represented multiple stakeholders in the service (Table 1).

Community Manager
Community Manager at Food Rescue Organization
Donors
D1 Fresh Produce Distributor D2 Farm Alliance Manager D3 Director of Purchasing at Dining Service Provider D4 Supermarket Chain Marketing Team
Recipient Organizations
R1 Food Pantry Director R2 Homeless Medical Outreach Specialist R3 Food Pantry Public Relations Coordinator R4 Food Pantry Operation Coordinator R5 Food Pantry Manager
Volunteers
V1 African American Female 60s V2 Caucasian Female 30s V3 Caucasian Female 40s V4 Caucasian Male 30s V5 Caucasian Female 60s
Clients
C1 African American Male 30s C2 Caucasian Male 30s C3 African American Female 40s C4 African American Female 50s C5 African American Female 20s
Everyday Citizens
E1 Taiwanese American Male 20s E2 Chinese American Male 20s E3 Caucasian Male 20s E4 Middle Eastern Male 20s E5 Indian Female 20s E6 Caucasian Female 20s

Table 1. List of Interviewees and Affiliations

Participants

We were introduced to the food rescue organization’s community manager as we built our relationship with the organization. Other stakeholders, such as the donors, volunteers, and recipient organizations, were recruited through an email list compiled by the food rescue organization. Food pantry clients were solicited for interviews at the food pantry with the permission of the director. Everyday citizens were recruited for interviews in public places, such as cafés.

Interview procedure

The first half of the interview was open-ended and began with questions on participants’ current practices and their general thoughts on what would be a fair way of distributing donations. Our algorithm concepts (described below) were used as a probe during the second half of the interview. The concepts helped participants further articulate and compare their notion of fairness to ones with which they disagreed.

Algorithm concept probes

Drawing from the fairness literature review, we created handouts with informational illustrations of three allocation models: efficiency-, equality-, and equity-centric allocations. All models assumed that an allocation algorithm would first filter the organizations that were open and in need of the items, and then make the allocation decision among them. The efficient allocation picked the recipient organization closest to the donation location; the equal allocation gave all the recipient organizations an equal chance and picked one randomly; and the equity allocation gave a higher probability of being chosen to the recipient organizations with greater needs before picking one randomly. These handouts were meant to initiate the conversation; participants could pick one, mix the ideas, or come up with their own. In order to probe what they thought a fair equity model would look like, we also prepared different factors that the equity model could use; the list included organizational (e.g., budget and size), client (e.g., belonging to a minority), and geographical (e.g., food desert, average income) characteristics. We asked them to rank the factors in importance in assessing degrees of need. This was again to promote the participants’ own thinking, rather than to make a statistical claim.

Interview questions

We interviewed the community manager twice for a total of three hours. We first asked her to describe her latest allocation decisions, her typical decision process, and what “fair” meant to her in this context. We then asked her opinions on the three allocation models, what she thought was the fairest, and how she felt about automation and its potential impact on work practices and relationships with multiple stakeholders. For donors, volunteers, and recipient organizations, each interview took about thirty minutes. We adopted and used the questions used for the community manager to include their activities with the food rescue organization. For food pantry clients, each interview took about fifteen minutes, as we tried to fit the interview within the time they spent waiting for their turn at the pantry. For everyday citizens, each interview took about thirty minutes. For the pantry clients and everyday citizens, the same set of questions was used, except for the question on the organization’s work practices and relationships.

Analysis

One of the two interviewers took notes during the interviews. Interviews were audio recorded if permission was granted from the participants. We did not record the

pantry client interviews as they were conducted in a sensitive location. We qualitatively analyzed our interview notes and transcripts [28, 33]. The interviewers met regularly as a group and went over the data to organize commonalities and differences into concepts of fairness and derive themes regarding the reactions to automation in each stakeholder group. We then compared different stakeholder groups in order to find challenges and opportunities in the interactions between the stakeholders, and to devise design principles that could address them.

FINDINGS

We report our findings from the interviews for each of the stakeholders.

Community manager

The community manager is responsible for matching donations with an appropriate non-profit and coordinating volunteer efforts simultaneously. She is dedicated to the organization's mission to reduce food waste and end hunger. She worked in non-profit organizations for seven years prior to joining the food rescue organization.

Combination of efficiency and equity as fair allocation

The community manager believes that a mix of the efficiency and equity models is the fairest, and has herself been mixing the two to make fair allocations: *"First [...] I think about what the proximity is. [...] I have [...] a general radius in my head [...] about all the organizations that are within a given geographic area. Those are [...] the first choice for a donation. Unless I happen to know that there's an organization with a dire need [...] that's outside of that radius, then they might jump ahead. Then the equity allocation [...] comes in [...] as I try to make sure that I'm distributing equitably throughout the whole universe."* As partly described in the quote, when the manager receives a call about a donation, she first looks at the size and content of the donation, thinks about potential recipients who could use that donation, who would require less than thirty minutes of travel, and who would be open to taking the donation. For example, perishable donations are sent to organizations that can keep the food refrigerated. She then narrows down this list using several heuristics to decide which organization needs the donation most. The heuristics combine food type and quantity, the organization's size, the last time that they received donations, etc. For example, she sent a donation of twenty sandwiches to an organization with a small set of clients, so that all of the clients could enjoy the food. When she remembers a particular organization that expressed their preference, she tries to match it to that organization. She also tries to keep small donations within the county for efficiency.

With 238 recipient organizations that the food rescue organization donates to, the community manager admitted some limitations of her current approach to being fair. *"[The information that] I feel my brain doesn't automatically boot up is an assessment of who haven't I donated to in a while, especially because sometimes they*

tend to fall off my radar, like it's easy for me to donate to the same place twice. I would like to make sure that geographically we are diverse in our distribution." She also mentioned other factors that play a role in decisions: *"[I might be over-donating to this place] because they're easy to get ahold of on the phone or maybe because I really like the staff member or maybe because a kid smiled at me the last time I was there."* She felt that an algorithmic system would help her make decisions more objectively, by reminding her of organizations that had not received donations in a while or had been receiving too much.

Algorithms' inability to empathize with recipients' needs

The manager worried that an algorithmic allocation would lack intuition and empathy in decision-making: *"My concerns are that a computer will never have that sort of the intimate human understanding of the recipients and their needs."* The manager on multiple occasions referred to a neighborhood which she cared for that was located in a food desert. She worried that the neighborhood might be further put at a disadvantage: *"I have trouble getting stuff to [the neighborhood] and I have this very non-logical passionate desire to take care of those people. A computer won't have that. If I have trouble making it happen and I ache for the people down in [the neighborhood] how is a computer ever going to prioritize that?"*

Social relationships and automated interaction

According to the manager, personally interacting with the volunteers and donors plays a big role in motivating volunteers and maintaining relationships with donors. She was concerned about the recent push by the food rescue organization for automated volunteer coordination to handle the greater demand they were facing: *"We're going in that direction where it's sort of like the Uber where you accept and you're just not going to have the same relationships with them. Maybe at that point you don't need them because the volunteer's [...] like, 'Who cares who does it?' It's [...] less important about who takes the volunteer opportunity."* Personally interacting with donors secures a strong partnership which seems to be in the best interest of the organization, which relies on donations. As the community organizer said, *"You develop [...] that interpersonal relationship with the donors which makes them feel better about donating and it makes them feel better about donating to us as opposed to anybody else. So just from an organizational survival standpoint. [...] I'd be really hesitant to [automate that interaction]."*

Volunteers

Volunteers transport donations to recipient organizations using their own vehicles. On average, completing one delivery takes an hour or longer when making multiple stops for larger donations. Our participants had various occupations, including two stay-at-home parents, one full-time worker, one business owner, and one retired individual. Participants with flexible schedules volunteered ad hoc, with some fitting regular, long-term opportunities

into their schedules. All supported the organization's mission of reducing food waste and ending hunger.

Varying fairness concepts dependent on resources & beliefs
Volunteers had the most varying perspectives on what they perceived as the fairest allocation model. Three participants viewed the equality model as the fairest, one chose the equity model, and the other chose the efficiency model. Their concepts of fairness were closely related to their personal beliefs and resources. For example, proximity-based, efficient allocation was considered fair because it uses less fuel and time, allowing volunteers to fit the deliveries into their busy schedules. V3, who delivered to a nearby community organization on a weekly basis, preferred this model for its local impact: *"I want something that is close because of the amount of time I spend doing it [...] but also because then it feels like it has a greater impact for my immediate community and it's making a difference with the people who live closer to me."* On the other hand, other volunteers did not think efficient allocation was fair. Participant V1, who regularly delivered donations to neighborhoods that are not centrally located, was particularly concerned that the recipients farthest from the donation site would be underserved by efficient allocation: *"But how is that fair? Say this is where the black community is, so again we are pushed out because we don't live close to the donation area."*

Three participants (V2, V4, V5) found equal allocation to be fair, and mentioned that it eliminates any discrimination in decision-making. They found the equity model to be unfair and expressed their concerns about ranking one organization over another. Two participants wondered how anyone could adequately rank the needs of the recipient organizations. Participant V5 said, *"I don't know how they measure the need. That is difficult to measure; in a food pantry, at any given day, [many] are in need. As far as ranking I don't know how to rank that."* Participant V2 who previously worked in a grocery store said, *"If an organization is actually accepting this food [that's about to expire] and taking the time to cut out the rotten parts [...] you won't see privilege in these organizations. Degree of need might not even exist. Anyone who is taking this food is high in need."* Through their up-close interactions with and observations of recipient organizations, these volunteers might have gained more concrete, experiential knowledge about recipient organizations, which may have made them feel less favorably toward supposedly equitable decisions that were based on abstract factors.

One other participant found equitable allocation to be fair. Participant V1 said, *"When you are looking at the greatest need, that's really fulfilling the mission of the organization because [...] it shouldn't be about what's easy for me to get around, it's about who needs it the most"*

Reluctance to travel farther or go to certain locations

While most volunteers found the equality or equity model fairer than an efficiency model, they were also worried that

these models would make food rescue require more energy and resources. Participant V2 expressed, *"I probably mostly would keep the same radius. It's not often that I go any further than around here."* Another concern was the perceived safety of the area. V2, who volunteers often with her son, was wary of certain areas: *"There are parts of the city I don't feel like taking him to because I have to leave him in the car to help bring food out and stuff."* Even when the travel time was minimal, she was hesitant to volunteer for particular neighborhoods. These concerns highlight the need for an algorithm that is fair to all stakeholders, including the volunteers themselves, not just the clients.

However, some volunteers mentioned a few factors that might encourage them to take more distant trips. For example, for participant V5, the quantity of rescued food mattered most: *"Saving a lot of food matters if I am going far. [...] It's somebody who needs food, I don't care who gets it."* For participant V1, helping her own underserved communities mattered most: *"They weren't delivering into the black community and what they needed was drivers [...] who had no problem going into the community [...] making sure that the people got a fair distribution of the food. So the food rescue organization asked me, because I'm so involved in the community."*

Efficiency and reduced biases in algorithmic allocation

Most interviewed volunteers welcomed an allocation algorithm. An advantage that participants repeatedly mentioned was that an algorithm would make the decision-making process easier and quicker, allowing the food rescue coordinator more time to recruit new donors and recipients (V2, V3, V4). Participant V3 believed that the algorithm could eliminate potential bias when having to decide between choosing one recipient over the other.

Automation as an opportunity for information exchange

Most of the interviewed volunteers were not regularly socially interacting with the community manager, and a proposal to automate existing modes of communication did not seem to impact social dynamics. An exception was participant V1, who showed a preference for social interaction with the food rescue organization. She regularly communicated with the manager about recipient organizations' preferences, to improve allocation and better meet the needs of her recipient communities. On the other hand, other volunteers also gained valuable information about the preferences of certain recipient organizations through their personal interactions during deliveries, but did not communicate this information to the manager. For example, participant V2 said, *"I'm getting to know the types of organizations and like who they serve and what they need,"* but when asked if she shared any of this knowledge with the manager, she said, *"I just kind of do what they say they need, I [...] figured they are short staffed and really busy so I keep it short and sweet with them. [...] Usually I'll send them a picture and tell them how many pounds of food [...] and they can post it on their page, but that's*

usually all the communication that we have.” A centralized communication channel in a mobile app could help volunteers share such information with the manager.

Donors

Donors are the suppliers of the food rescue organization. They donate extra food resulting from failed predictions of demand or products that fail to meet quality standards. Donors were glad to work with a food rescue organization that would take food on demand, rather than having to donate to food banks or throw it away. The quantities varied, ranging from a large pallet to a few boxes of food. The types of donated food included dented cans, fruits with moldy spots, unsold bread, and fresh produce that was leftover or had defects.

Efficiency- and equity-based fairness concepts

Overall, most donors emphasized the importance of proximity-based, efficient allocation. Participant D4 also mentioned that the efficient model was well-aligned with her organization’s philosophy: *“For us as a company that focuses on being hyper-local, the more of a local impact the better.”* On the other hand, Participant D3, who donated in large quantities, raised a concern that continuously donating to the closest recipient could overwhelm them with donations they could not handle.

In contrast to volunteers, who mostly favored an equality model, none of the interviewed donors preferred the equality model. They did not find it to be efficient nor necessarily considerate of the organizations’ needs, despite the cost of travelling further for food rescues. Instead, all wished to maximize the impact of their donations, and the equity model fit this motive. Participant D4 said, *“We like to have a local impact [...], but I do recognize that there are [...] communities [...] that are further out who have a higher need. [...] I want to give where I know our company money is going to have the greatest impact.”*

Reduced bias but lost intuition in algorithmic allocations

The majority of the donors preferred automating allocation decisions for its efficiency and to streamline the process. Like other stakeholders, some donors said that an algorithm would be more objective in decision-making; others were worried that automating the process would lose the human aspects involved in the decision-making process of the food rescue organization. Participants D2 and D3 wished for the automated algorithm to have a manual override, hoping that it could be used as a tool for guidance for the community manager rather than as a final decision-maker. Most donors, thought highly of what the food rescue organization was doing and trusted the organization to have a good overview of need across all the recipients.

Managing relationships with the food rescue organization

Overall, donors did not mind an automated system replacing the social interactions they had with the food rescue manager. When determining allocations of food, the social relationship seemed less of a concern to them. As participant D2 said, *“I can see the emotion gets lost, less*

personal, but I don’t know if that is a bad thing. [...] Human connection would be nice but [the service] won’t lose anything.” However, donors seemed to value personal relationships when negotiating partnerships, and showed preferences for organizations they got along well with. As participant D4 expressed, *“If someone is difficult to work with I cut off the relationship [...] I won’t put in the effort, because there are 100 more organizations in line. You want to work with people you would like to work with.”*

Recipient organizations

Most recipient organizations were serving communities near their locations. They all received government funding to stock food for their clients, and most were buying inexpensive items with a longer shelf life to maximize their limited budget. The organizations did not rely solely on food rescue donations, but enjoyed receiving donations that were rare to have in stock, such as fresh fruit and vegetables, milk, artisan bread, personal care products, and household items. However, while the food rescue donations were helpful, some also mentioned the considerable effort and labor required to coordinate pick-ups. Recipients especially felt this way for donations that were unfit for their clients’ needs or received in excessive quantities.

Equity in theory, equality in practice

Many organizations worried that an efficiency-based model would allocate donations only to the closest recipients, which would be unfair to those who were not centrally located. Two recipients acknowledged their pantries were located near major grocery chains and that they were receiving frequent donations from these stores; yet they said that they would be okay with an equality- or equity-based model even if it might give them less of an advantage.

Most recipients initially thought the equity model to be the fairest; they agreed that organizations had varying degrees of need. However, when asked to rank factors that could be used to assess the needs of different organizations, three participants opted out of the activity entirely. They expressed concerns about the difficulty of ranking organizations who are all equally in need. One participant said it would be tricky to prioritize factors; for example, organizations that primarily serve minors already receive more funding. The participant also thought factors like the general income level of a neighborhood would not accurately portray the needs of the people who are forced to move to different neighborhoods due to gentrification. Another said it would be difficult to judge different levels of their clients’ needs: *“I don’t know whose hunger is more important to address—whether they were hungry for 10 years or 10 weeks.”* Rather than thinking about the varying needs of organizations, R3 suggested adding a measure of efficiency to the model, which accounted for how quickly a recipient could distribute donations.

Regularity as important as fairness

Many participants who thought the equality or equity model to be fair were cautious of randomly assigning donations to

recipients because of the unpredictability of when and what donation to expect. Participant R5 said, *“These are extras and add-ons, what comes from food rescue are organic and fresh, the stuff that I usually don’t have access to, it is kind of a cool bonus, stuff that people don’t normally get to try. [The equality model] is fair [...] but then [there are always] the other ones [that] don’t get any of it. If [the allocation cycle] is once a month [...] there is a whole another month they don’t get a shot at the food [if they were not picked].”*

Reduced bias but lost intuition in algorithmic allocations

Like other stakeholders, more than half of the recipients explicitly said an algorithm would be less biased than a human making allocation decisions (R1, R4, R5). R5 said, *“It helps if they know who you are [...] That can work well but that can work against people because [...] whoever is giving it out might not like that organization [...] and skip over that organization so the algorithm takes that out of the mix.”* Many also mentioned efficiency as an upside of an algorithmic allocation: *“Trying to get matched up for these donations can involve a lot of phone calls, [...] having a computer making these decisions would be a more efficient process overall.”* (R4) However, R4 worried an algorithm would lack human intuition: *“Recently the [manager] knew that there were a lot of mangoes available, and asked if we would want them because we worked with a lot of refugee families from Southeast Asia [...] that was a neat kind of a donation that wouldn’t have been made in this model [...] but overall for efficiency it makes a lot of sense.”*

Clients of the recipient organizations

Clients get donated foods from the recipient organizations, such as food pantries and housing authorities, to get through the month or stretch their fixed income. Our interviewees were clients of one food pantry. Items donated by the food rescue organization added to their monthly food quota from the food pantry and were often items rarely seen in the pantry, like fresh produce, milk, and healthy alternatives. The food pantry clients did not know where a particular donation came from unless the item was labeled with the donor’s organization. This particular pantry offered a supermarket-style experience to give clients their choice of food, and ran on a first-come-first-serve model.

Equity-based fairness concept

All client participants found the equity model to be the fairest and acknowledged that factors like living in a food desert or working low-income jobs should increase a population’s chances of receiving donations. Participant C1 mentioned that the allocation model could benefit from collective votes from people about which factors they thought should be used in determining need.

Reduced biases in algorithmic allocation

Pantry clients approved of an algorithm making allocation decisions instead of a human manager. Participants C3 and C4 regarded computers to be blind to ethnicity and human discrimination, unless the computer algorithm itself

portrayed bias. Participant C4 also commented, *“Computers are taking human jobs anyways.”*

Everyday citizens

Everyday citizens found the equity model to be the fairest. Participant E1 explained, *“As of the principle, I lean more toward whoever needs it more gets it, so ideals are good.”* Everyday citizens, who were not direct stakeholders, rarely mentioned efficiency and had varying but strong opinions on how to assess degree of need. The equity allocation model was thought by all everyday citizen participants to be a better version of equality, except for participant E1. Participant E1 referred to equity as simply a “more intense” version of equality due to the random selection of recipients: *“An organization in high need with a 50 percent chance of receiving a donation might still lose to an organization of lower need. And for rare donations, the organization in high need can be left with a sour taste in their mouth.”* This concern also resonates with the recipient organizations’ emphasis on the importance of regularity. One participant proposed a model in which recipient organizations could connect and coordinate with each other in real time to decide who was in the most need.

DISCUSSIONS AND IMPLICATIONS

Our research suggests that taking a human-centered approach to understand multiple stakeholders’ perspectives is an important first step in designing an algorithm. It allows designers to balance multiple stakeholders’ needs and interests without overlooking anyone. Participants also appreciated having a transparent process that aimed to incorporate their interests and feedback into the design of the allocation and coordination process. In the following section, we discuss challenges and considerations for designing fair and motivating algorithmic services.

Fair algorithms: What is fair and how to decide

Many services that currently rely on human workers or experts’ allocation decisions are beginning to use algorithmic tools to replace or assist the decision-makers. Embedding algorithms in service management necessitates formalized rules of fairness, in line with the recent attention on the case of triage [43]. Our research reveals the importance of and challenges in defining fairness rules for algorithms that appeal to multiple stakeholders.

Plural fairness concepts within and across stakeholders

Literature on fairness has shown that the concept of fairness is context-dependent, varying by cultures and tasks. Consistent with the literature, our findings show multiple concepts of fairness within and across stakeholders in the context of allocation. Within stakeholders, personal beliefs and organizational philosophies all influenced individuals’ fairness concepts. Some people believed *efficient* allocation to be fair, because it enabled them to directly impact their surrounding communities. Some believed *equal* allocation to be fair, as they felt that areas further away from donor organizations should have the same chance of receiving donations as closer ones; these participants generally had

difficulty judging the degree of recipient organizations' needs, as they felt that all the organizations deserved help equally regardless of size, budget, or the characteristics of the populations served. Others believed *equitable* allocation to be fair, such that organizations that needed donations most should get them first; however, participants' methods for discerning degree of need varied.

Differences across stakeholders also emerged depending on their roles and whether they directly interacted with the clients of recipient organizations, whom the allocation decisions immediately impacted. The community manager and most of the donors, pantry clients, and everyday citizens deemed the equity-centric allocation model to be fair. They valued maximizing the utility of donations. On the other hand, most of volunteers and recipient organizations, who directly interacted with recipient organizations and their clients, thought the equality-centric allocation model to be fair. They expressed difficulties in judging the degree of needs or importance of different causes. This difference might be explained by psychological distance from donation recipients. As front-line workers, volunteers and staff in recipient organizations have visceral experiences with donation recipients; in-person observation of the facilities and donated goods and interpersonal interactions might enable them to see nuances and contexts that are difficult to capture in the abstract factors that are used to prioritize needs in the equity model. This tension in stakeholders' fairness concepts is similar to the difficulties that hospital workers experienced when they operationalized triage rules to decide which patients need to be treated first in a disaster situation [43].

Designing fair algorithms

The plurality of fairness concepts held by multiple stakeholders poses challenges in defining which fairness principles algorithms should embody. Drawing from our findings, we propose two potential ways to design algorithmic services—collectively defining one global fairness rule, or enabling local operations of fairness principles by stakeholders. A democratic, open process can be used to collectively build fair allocation algorithms. Many participants suggested a poll, or at least wanted to know the rules and weightings of different factors in an algorithm, so that they could tailor their input to the allocation algorithm. Our findings also show the importance of creating close feedback loops between stakeholders who make allocation decisions and front-line stakeholders who execute decisions and are directly impacted by algorithmic decisions. Decision-makers can propose a fair rule based on holistic information about recipients, whereas volunteers and recipients can provide hands-on feedback on recipients' preferences and how the decisions actually fare in the field.

Alternatively, instead of determining one global concept of fairness, designers can embrace the plurality of fairness. Algorithms can process massive amounts of data, which means they can store and respond to each person's

individual input. Instead of imposing an external model of fairness on users, what if we could enable different stakeholders to act according to what they individually believe fair? For example, different donors might specify the ways in which they prefer their donations be distributed per donation, such that community organizations could emphasize local impact while organizations targeting food deserts could specify as such. The algorithmic donation matching system could also transparently display the reasoning behind each assignment, motivating volunteers by giving them more say in the kinds of donations they deliver and the organizations they help. Further research needs to be done on the processes of determining fair algorithms and to understand the pros and cons of different operationalizations of fairness in the long run.

Beyond fair algorithms: Considering people, contexts, and interfaces in algorithmic services in the real world

Our findings suggest that to design fair algorithmic services, we need to go beyond asking what is conceptually fair. We need to also think about how to make the algorithmic service system work fairly when implemented in the real world. Even if the algorithms themselves make “fair” decisions, the systems may still deliver unfair results if the people participating in the service are not motivated to follow the decisions.

Transparency of algorithmic decisions for trust and adoption

A previous ethnographic study on emergency dispatchers' adoption patterns of a new automated dispatch decision-making system is relevant to our work [42]. The study showed that even though the error rates in automated dispatching were as low as they had been prior to automation, the dispatchers did not trust the results and wanted to verify everything themselves using pens and paper, even six months after its introduction. In our case, similar reservations were displayed by the community manager, who makes the final allocation decisions. She has been using her own heuristics and logic to make allocation decisions for 1.5 years. If the algorithmic suggestions do not seem to make sense, she may start to override the system's recommendations on a regular basis, reintroducing her preferences and implicit biases to the process. Thus to maintain the fairness of the system, it is important that the interface explain the algorithm's recommendations and help the decision-maker make sense of them.

Accounting for unintentional biases and motivations

Our findings also suggest that we need to understand stakeholders' unintentional biases that might make algorithmic systems function unfairly, and design strategies to account for them. Interviews with the volunteers suggested that even when an algorithmic service makes a decision to allocate donations to lower-income areas or food deserts, volunteers may not sign up if the areas are perceived as dangerous or require longer trips. These attitudes are consistent with on-demand workers' behaviors, such as when TaskRabbit workers require higher pay-offs to go to the suburbs [37] and Uber and Lyft drivers turn off

“driver-mode” in areas perceived as dangerous [19]. In such cases, it is important for volunteers to understand why the decision was made. For example, one volunteer said she did not mind driving longer distances if she knew she was rescuing high-quality or larger quantities of food, or if it was for a cause that she cared about, such as feeding hungry children. Highlighting this information and explicitly saying why the donation matters may motivate volunteers to take the longer trips to food deserts, or go to areas that are not on their usual routes. The food rescue organization specifically recruited a volunteer from an underserved community to facilitate transportation to the area. Actively recruiting volunteers from distant or low-income areas could help the allocation service function more fairly.

Preserving meanings and social interactions in automation of algorithmic services

Initially we were concerned that people might have negative attitudes toward an algorithmic service taking on some of the roles of a human community manager. On the contrary, all participants positively responded to the algorithm’s efficiency and lack of bias. However, they also pointed out important aspects of human interactions that need to be preserved in algorithmic interactions.

Respecting empowerment in automation

Our interviews also suggested that we should consider the importance of emotion and meaningfulness when deciding how much of the workflow to automate. Zuboff’s ethnographic observations of newly-automated workplaces suggest that different types of automated tasks can make workers feel powerless or lose expertise on the one hand, or, on the other hand, feel empowered and free from meaningless work [45]. In our interviews with the community manager, we observed that the very act of making the allocation decisions gives her an opportunity to find meaning in her work—she thinks about the “dire needs” of the recipient organizations, remembers their unique preferences, and finds joy when she successfully satisfies those preferences. Automating this process—by making the decision and only asking for her final confirmation, or even by suggesting the top three options—may make this task too simple and less meaningful. Further research would be needed to find the balance between increasing efficiency and fairness on the one hand and making the task meaningful on the other.

Automating social interactions

Donors, volunteers, and recipient organizations are the stakeholders with whom the community manager directly interacts, through phone calls, text messages, or face-to-face meetings. Our findings suggest that we should not automate all these interactions equally. While they were okay with communicating through an app or texts, donors still wanted to have opportunities to socially interact with the manager for smoother coordination or to get to know her; this interpersonal relationship seemed to contribute to their trust of the food rescue organization. It will therefore be important to preserve opportunities for the manager to

interact personally and directly with these stakeholders. On the other hand, recipient organizations and volunteers were more open to automating their interactions with the community manager, except for a subset of volunteers who already interact regularly with the manager. Automation may also “expand” the positive potential of social interactions between the community manager and volunteers and recipient organizations by facilitating communication. Volunteers who more frequently interacted with the community manager mentioned discussing recipient organizations’ needs and logistics. However, the manager did not seem to keep up this same level of interaction with the majority of volunteers or recipient organizations, at least based on their phone calls, social media, and text messages. It would therefore be extremely useful for the manager to have a central, automatic method of communicating with the larger set of volunteers at the same time and collecting what they learn about each donor or organization, including information about recipients’ particular preferences and constraints.

LIMITATIONS AND FUTURE WORK

Our work has limitations. We explored what people think is fair based on the principles and logic of allocation in a non-profit domain in the United States, but people’s opinions on fairness and reactions to algorithms may differ depending on their cultures and work contexts. Our interviews allowed us to thoroughly understand multiple stakeholders’ rich, nuanced thoughts and feelings about fairness and algorithmic automation, but the sample size was small. In addition, people’s thoughts and behaviors might change once they have actually interacted with an algorithm. Our future research will draw from the findings reported in the paper to create algorithms, simulations, prototypes, and systems, and investigate people’s experiences and behaviors with a larger sample.

CONCLUSION

Many services in our society are using algorithms to improve their functionality. We know that algorithms can make services more efficient, but how can we make these algorithmic services fair and motivating for multiple stakeholders as well? We conducted interviews with stakeholders involved in a food rescue organization that plans to introduce an algorithmic allocation and coordination service. Our findings suggest that we need to understand and account for varying fairness notions held by different stakeholders; consider people, contexts, and interfaces for algorithms to work fairly in the real world; and preserve meaningfulness and social interaction in automation to build fair and motivating algorithmic services. Our work offers an example of a social-systems analysis of an algorithmic system at both the conception and design stages, which we hope will inspire the design of emerging algorithmic systems.

ACKNOWLEDGEMENTS

This work was supported by NSF1651566. We thank our participants who shared valuable insights with us.

REFERENCES

1. Robert F. Bales, "A set of categories for the analysis of small group interaction." *American Sociological Review* 15, no. 2 (1950): 257-263.
2. John E. Bell, and Patrick R. McMullen. "Ant colony optimization techniques for the vehicle routing problem." *Advanced Engineering Informatics* 18, no. 1 (2004): 41-48.
3. Deana Brown and Rebecca E. Grinter. "Designing for Transient Use: A Human-in-the-loop Translation Platform for Refugees." In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 321-330. ACM, 2016.
4. Donna Chen, Kara M. Kockelman, and Moby Khan. "The electric vehicle charging station location problem: a parking-based assignment method for Seattle." In *Transportation Research Board 92nd Annual Meeting*, vol. 340, pp. 13-1254. 2013.
5. Kate Crawford and Ryan Calo. "There is a blind spot in AI research." *Nature* 538, no. 7625 (2016): 311.
6. Christopher A. Le Dantec, Robert G. Farrell, Jim E. Christensen, Mark Bailey, Jason B. Ellis, Wendy A. Kellogg, and W. Keith Edwards. "Publics in practice: Ubiquitous computing at a shelter for homeless mothers." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1687-1696. ACM, 2011.
7. Tawanna R. Dillahunt, and Amelia R. Malone. "The promise of the sharing economy among disadvantaged communities." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 2285-2294. ACM, 2015.
8. Lynn Dombrowski. "Designing SocioTechnical Food Justice." *iConference 2015 Proceedings* (2015).
9. Lynn Dombrowski, Jed R. Brubaker, Sen H. Hirano, Melissa Mazmanian, and Gillian R. Hayes. "It takes a network to get dinner: designing location-based systems to address local food needs." In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pp. 519-528. ACM, 2013.
10. Lynn Dombrowski, Amy Volda, Gillian R. Hayes, and Melissa Mazmanian. "The labor practices of service mediation: a study of the work practices of food assistance outreach." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1977-1986. ACM, 2012.
11. Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 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*. 153-162.
12. Batya Friedman. *Human values and the design of computer technology*. No. 72. Cambridge University Press, 1997.
13. Batya Friedman, Alan Borning, Janet L. Davis, Brian T. Gill, Peter H. Kahn Jr, Travis Kriplean, and Peyina Lin. "Laying the foundations for public participation and value advocacy: Interaction design for a large scale urban simulation." In *Proceedings of the 2008 international conference on Digital government research*, pp. 305-314. Digital Government Society of North America, 2008.
14. Batya Friedman and Helen Nissenbaum. "Bias in computer systems." *ACM Transactions on Information Systems (TOIS)* 14, no. 3 (1996): 330-347.
15. Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. "Measuring price discrimination and steering on e-commerce web sites." In *Proceedings of the 2014 conference on internet measurement conference*, pp. 305-318. ACM, 2014.
16. Tony Hope, Lars Peter Østerdal, and Andreas Hasman. "An inquiry into the principles of needs-based allocation of health care." *Bioethics* 24, no. 9 (2010): 470-480.
17. Lucas D. Introna, and Helen Nissenbaum. "Shaping the Web: Why the politics of search engines matters." *The information society* 16, no. 3 (2000): 169-185.
18. James Konow. 2003. "Which is the fairest one of all? A positive analysis of justice theories." *Journal of economic literature* 41, no. 4, 1188-1239.
19. Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. "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. ACM, 2015.
20. Min Kyung Lee and Su Baykal. Algorithmic mediation in group decisions: Fairness perceptions of algorithmically mediated vs discussion based social division. To appear in CSCW 2017.
21. Gerald S. Leventhal, "The distribution of rewards and resources in groups and organizations." *Advances in experimental social psychology* 9 (1976): 91-131.
22. Jany Li, Xuan Zhao, Mu-Jung Cho, Wendy Ju, and Bertram F. Malle. *From Trolley to Autonomous Vehicle: Perceptions of Responsibility and Moral Norms in Traffic Accidents with Self-Driving Cars*. No. 2016-01-0164. SAE Technical Paper, 2016.
23. Jenn-Rong Lin and Ta-Hui Yang. "Strategic design of public bicycle sharing systems with service level constraints." *Transportation research part E: logistics and transportation review* 47, no. 2 (2011): 284-294.

24. Bertram F. Malle, Matthias Scheutz, Thomas Arnold, John Voiklis, and Corey Cusimano. "Sacrifice One For the Good of Many?: People Apply Different Moral Norms to Human and Robot Agents." In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, pp. 117-124. ACM, 2015.
25. Jessica K. Miller, Batya Friedman, Gavin Jancke, and Brian Gill. "Value tensions in design: the value sensitive design, development, and appropriation of a corporation's groupware system." In *Proceedings of the 2007 international ACM conference on Supporting group work*, pp. 281-290. ACM, 2007.
26. Viginia Murphy-Berman, John J. Berman, Purnima Singh, Anju Pachauri, and Pramod Kumar. "Factors affecting allocation to needy and meritorious recipients: A cross-cultural comparison." *Journal of Personality and Social Psychology* 46, no. 6 (1984): 1267.
27. One hundred year study on artificial intelligence (AI100) <https://ai100.stanford.edu>
28. Michael Quinn Patton. *Qualitative evaluation and research methods*. SAGE Publications, inc, 1990.
29. Eli Pariser. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.
30. Andy L Polaine. Løvlie, and Ben Reason. "Service design." *From Implementation to Practice*. New York: Reosenfeld Media (2013).
31. Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. "Auditing algorithms: Research methods for detecting discrimination on internet platforms." *Data and Discrimination: Converting Critical Concerns into Productive Inquiry* (2014).
32. Peter Stone. "Why lotteries are just." *Journal of Political Philosophy* 15, no. 3 (2007): 276-295.
33. Anselm Strauss and Juliet Corbin. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Sage Publications, Inc, 1998.
34. Latanya Sweeney. "Discrimination in online ad delivery." *Queue* 11, no. 3 (2013): 10.
35. TaskRabbit. www.taskrabbit.com
36. Alex S. Taylor, Siân Lindley, Tim Regan, David Sweeney, Vasillis Vlachokyriakos, Lillie Grainger, and Jessica Lingel. "Data-in-place: Thinking through the relations between data and community." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 2863-2872. ACM, 2015.
37. Jacob Thebault-Spieker, Loren G. Terveen, and Brent Hecht. "Avoiding the south side and the suburbs: The geography of mobile crowdsourcing markets." In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 265-275. ACM, 2015.
38. Uber. www.uber.com
39. Upwork. www.upwork.com
40. Amy Volda, Ellie Harmon, and Ban Al-Ani. "Bridging between organizations and the public: volunteer coordinators' uneasy relationship with social computing." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1967-1976. ACM, 2012.
41. Elaine Walster, Ellen Berscheid, and G. William Walster. "NEW DIRECTIONS IN EQUITY RESEARCH" 2." *Advances in experimental social psychology* 9 (1976): 1. Harvard
42. Jack Whalen. "Expert systems versus systems for experts: Computer-aided dispatch as a support system in real-world environments." *Cambridge Series on Human Computer Interaction* (1995): 161-183.
43. Whose Lives Should Be Saved? Researchers Ask the Public. New York Times, 2016. <http://www.nytimes.com/triage>
44. Workshop on Fairness, Accountability, and Transparency in Machine Learning. <http://www.fatml.org>
45. Shoshana Zuboff. *In the age of the smart machine: The future of work and power*. Basic books, 1988.