A Human-in-the-loop Framework to Construct Context-aware Mathematical Notions of Outcome Fairness

Mohammad Yaghini

University of Toronto

Toronto, Canada

mohammad.yaghini@mail.utoronto.ca

Andreas Krause

ETH Zürich

Zürich, Switzerland

krausea@ethz.ch

Hoda Heidari

Carnegie Mellon University

Pittsburgh, PA, USA

hheidari@cmu.edu

ABSTRACT

Existing mathematical notions of fairness fail to account for the **context** of decision-making. We argue that moral consideration of contextual factors is an inherently *human* task. So we present a framework to learn *context-aware* mathematical formulations of fairness by eliciting people's *situated fairness assessments*. Our family of fairness notions corresponds to a new interpretation of economic models of *Equality of Opportunity (EOP)*, and it includes most existing notions of fairness as special cases. Our *human-in-the-loop* approach is designed to learn the appropriate parameters of the EOP family by utilizing human responses to pair-wise questions about decision subjects' *circumstance* and *deservingness*, and the *harm/benefit* imposed on them. We illustrate our framework in a hypothetical criminal risk assessment scenario by conducting a series of human-subject experiments on Amazon Mechanical Turk. Our work takes an important initial step toward empowering stakeholders to have a voice in the formulation of fairness for Machine Learning.

CCS CONCEPTS

• **Applied computing** → **Economics**; • **Human-centered computing** → **Empirical studies in collaborative and social computing**.

KEYWORDS

Equality of Opportunity, Fairness, Machine Learning, Human Judgments

ACM Reference Format:

Mohammad Yaghini, Andreas Krause, and Hoda Heidari. 2021. A Human-in-the-loop Framework to Construct Context-aware Mathematical Notions of Outcome Fairness. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, May 19–21, 2021, Virtual Event, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3461702.3462583>

1 INTRODUCTION

Despite the surge of interest in designing and guaranteeing mathematical notions of fairness for Machine Learning (ML) (see, e.g., [10, 13, 30]), existing formulations fail to take the decision-making *context* into account and consider the social *backdrop* and *ramifications*

of automated decisions for their subjects and society. Such critical considerations vary across domains, cultures, and societies, and therefore, it is a futile attempt to look for one-size-fits-all definitions of fairness that are considered acceptable across a wide range of situations. Even if we assume certain facets of a highly complex concept like fairness can be reasonably approximated and described in mathematical terms—which itself is a significant assumption—the precise formulation must be highly sensitive to the context. Relying solely on the intuitions of ML experts and practitioners to capture the relevant nuances is likely inadequate and ineffective. We argue that instead, ML practitioners must have a framework in place that allows them to elicit relevant human¹ judgment and utilize it to construct mathematical formalizations of fairness *tailored* to the decision-making domain at hand. Such a framework empowers individuals and stakeholders to have a voice in the formulation of fairness and offers practitioners valuable tools to *audit* ML-based decision-making alternatives for *unfairness* in outcomes.

To further illustrate the importance of the *context* and human judgments, consider a popular notion of fairness, called “equality of odds” [10]² in a *hypothetical* recidivism prediction context, where predictions about future crimes directly translate into pre-trial decisions about defendants. As we discuss in our Equality of Opportunity (EOP) framework (Section 2.1), the equality of odds requirements can be interpreted as making several *implicit moral assumptions*. Namely, all defendants who would commit a crime if granted bail should receive the same (high-risk) prediction—which subsequently leads to jail time. Similarly, all defendants who will not reoffend should receive the same (low-risk) prediction, leading to bail. In other words, when we employ equality of odds as the notion of fairness in this context, it implies that we agree with the following value-laden statements: “The true label is a good indicator of a defendant's *deserved* pretrial outcome, and the predicted label is a good indicator of the *benefit/harm/utility* imposed on them as the result of the decision they receive”. We argue that ML-practitioners should not make such value judgments in a vacuum without carefully examining whether they reflect stakeholders' situated perceptions of desert and utility. Regarding pretrial decisions, for example, a closer inspection may reveal that in addition to re-offense, people favor taking a defendant's age and family responsibilities into account to reach a just decision about them.

In this work, we present a practical framework to *construct* context-aware mathematical formulations of fairness utilizing people's *situated judgments* about fairness in outcomes. We first offer a new interpretation of the theoretical model proposed by Heidari

This work is licensed under a Creative Commons Attribution International 4.0 License.

AIES '21, May 19–21, 2021, Virtual Event, USA.
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8473-5/21/05.
<https://doi.org/10.1145/3461702.3462583>

¹The term ‘human’ here refers to *stakeholders*, including impacted communities and individuals, domain experts, practitioners, policymakers, and beyond.

²Equality of odds requires false positive and false negative rates to be equal for all socially salient groups.

et al. [11]—which shows that most existing notions of algorithmic fairness are special cases of economic models of *Equality of Opportunity (EOP)*. This new interpretation allows us to design questionnaires that elicit human judgments regarding such concepts as decision-subjects’ *circumstance*, *desert*, and *utility*, and utilize the answers to learn a context-appropriate formulation within the EOP family. As a proof-of-concept, we illustrate our proposal by running human-subject experiments on Amazon Mechanical Turk (AMT) for a *hypothetical* criminal risk prediction scenario synthesized using the COMPAS dataset [14]. Before we proceed to the overview of our empirical setup and findings, two remarks are in order: First, we must emphasize that our use of the COMPAS dataset to generate our questions is solely to ground our questionnaires in a real-world use-case of ML in making high-stakes decisions about human beings. For the lack of better alternatives, many Fair-ML studies have used the COMPAS dataset for demonstrations in the past. By following suit, we do *not* mean to imply that our work tackles the numerous issues and considerations involved in assessing the fairness of ML in the highly-charged and complicated domain of criminal justice. Second, we utilized crowdsourcing to illustrate the feasibility of our human-in-the-loop proposal in a *cost- and time-efficient* manner. We do *not* advocate for crowdsourcing (or, more generally, relying on the average opinion of ordinary people) as a reliable and ethically sound approach to settle moral dilemmas. We will present a more thorough discussion of this point in Section 5.

With the above caveats in mind regarding their intent and scope, our experiments are conducted as follows: In order to learn the notion that is most compatible with people’s perception of fairness, we ask our participants to respond to a series of three questionnaires. Each participant is required to answer a total of 55 questions—all concerning a *hypothetical* criminal risk assessment scenario. A typical question in our experiment provides the participant with information about two hypothetical decision subjects (defendants) whose features are synthesized using instances in the COMPAS data set. The participant is then prompted to choose (a) which of the two defendants he/she believes is more *deserving* of receiving a lenient decision? (b) which of them he/she believes would *benefit* more from their algorithmic decision? Using the participant’s answer to these questions, we estimate the mathematical formulation of fairness that best captures his/her judgments in the decision-making context at hand. In addition to estimating the fairness notions of each participant, we propose and investigate two methods to aggregate individual participants’ notions into a single definition of fairness. In that sense, our work takes an initial step toward incorporating people’s collective judgments into the process of formulating fairness for algorithmic decision-making. Our small-scale empirical findings show that participants favor taking factors beyond the defendant’s true- and predicted labels (namely, race, gender, age, and criminal history) into account when formulating fairness and determining his/her deserved sentencing outcome and utility. Importantly, these considerations are not reflected in any of the existing mathematical formulations of fairness that have been extensively debated in the criminal justice context [13].

Summary of contributions. We offer a well-justified framework to learn a formal notion of (outcome) fairness tailored to

the specific decision-making context at hand. We incorporate contextual factors by eliciting human responses to questions about decision-subjects’ deservingness, circumstance, and harm/benefit. Our contributions include:

- Identifying a well-justified family of fairness notions (EOP) includes many common notions as special cases.
- Learning a new fairness notion for a specific context (criminal risk assessment) by eliciting human judgment and utilizing the data to optimize the EOP family’s parameters.
- Collecting and analyzing a novel dataset on people’s judgments regarding the above use-case of ML through AMT.
- Demonstrating that prior notions of fairness fail to adequately capture the participants’ situated fairness judgments.

In conclusion, our work draws attention to two significant gaps in the growing literature on fairness for ML. First, fairness is a highly context-dependent concept, and designing an ethically acceptable formulation of fairness requires carefully examining the domain and accounting for the social- and individual-level implications of automated decisions. Second, accounting for the social context of decisions is an inherently human task, and the machine on its own lacks the faculties to understand and process the social and ethical aspects of its decisions. Therefore, we need a systematic approach to utilize human-judgment. Our work presents a practical—yet by nature simplified—toolkit to construct fairness formulations that account for the context and background of decisions—a crucial consideration that existing mathematical formulations of fairness entirely ignore. While it will take many additional iterations and refinements to render this framework ready to employ in real-world applications, we believe it offers a concrete step toward empowering people’s voice in value-sensitive ML design, and we hope that the variety of critical questions it brings to the fore lead to impactful directions for future work in Fair-ML.

1.1 Related Work

Numerous mathematical definitions of fairness have been recently proposed and studied; examples include demographic parity [7], disparate impact [30], equality of odds [10], and calibration [13]. While each of these notions may seem appealing on their own, they are incompatible with one another and cannot hold simultaneously (see [3, 13]). Furthermore, at least in the context of criminal risk assessment, it is far from settled which one of these notions (if any) is the most appropriate measure of algorithmic fairness (see [1] and [5]).

Several recent papers empirically investigate AI ethics and algorithmic fairness utilizing human-subject experiments. MIT’s moral machine [2] provides a crowd-sourcing platform for aggregating human opinions on how self-driving cars should make moral decisions by eliciting their preferences through a series of pairwise preference ordering questions. For the same setting, Noothigattu et al. [18] propose learning a random utility model of individual preferences, then efficiently aggregating those individual preferences through a social choice function. Lee et al. [15] proposes a similar approach for ethical decision-making in the context of food distribution. Like the above papers, we obtain input from human-participants by asking them *pairwise* questions (i.e., comparing two

alternatives from a moral standpoint). The intent and design of our experiments, however, are entirely different from prior work.

On the topic of algorithmic fairness, Grgic-Hlaca et al. [9] study the criteria along which people perceive the use of certain features as unfair in making recidivism predictions about defendants (e.g., whether using the feature “family’s criminal background” violates defendants’ privacy.) Veale et al. [27] interview *public sector machine learning practitioners* regarding the challenges of incorporating public values into their work. Holstein et al. [12] investigate *commercial product teams’* needs in developing fairer ML through semi-structured interviews. Woodruff et al. [29] focus on how affected communities feel about algorithmic fairness. Unlike our work, in which the primary focus is on *quantifying* people’s perceptions of fairness, the above papers provide *qualitative* suggestions and best-practices.

Several recent articles investigate people’s attitudes towards *existing* mathematical formulations of fairness. Saxena et al. [23] investigate people’s attitude toward three notions of *individual fairness* in the context of *loan decisions*. Srivastava et al. [25] investigate which of the *existing* group-level formulations of fairness best captures people’s perception of fairness. Saha et al. [22] investigate non-experts’ comprehension of demographic parity and the impact of the application scenario on it. To our knowledge, no prior work has attempted to *construct context-aware* notions of fairness utilizing human perception of fairness for ML.

2 AN EQUALITY-OF-OPPORTUNITY FRAMEWORK

In this section, we formally define our learning setting and cast the problem of formulating fairness in *outcomes* as estimating the parameters of an EOP-based criterion. We must emphasize at the outset that our goal is to narrowly capture *distributive* aspects of (un)fairness through algorithms—namely, we are concerned with disparities in *outcomes* and how harms/benefits are allocated to decision-subjects as the result of the ML-based decisions they receive. Our family of fairness notions solely addresses *outcome/distributive* disparities and remains agnostic to the *causes* and *processes* by which such disparities come to be. Such *procedural* considerations, while extremely important, are outside the scope of the current work.

The learning environment. Throughout, we consider the following *simplified* automated decision-making setting: a predictive model (e.g., a hypothetical recidivism prediction model) is trained using historical data records (e.g., the attributes and labels of past defendants) through the standard supervised learning pipeline. The predictions made by this model then directly translate into important decisions for never-before-seen subjects (e.g., a future defendant is sent to jail if the model predicts that he/she would re-offend in the future). More precisely, let $T = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ denote a training data set consisting of n instances, where $\mathbf{x}_i \in \mathcal{X}$ specifies the feature vector for individual i and $y_i \in \mathcal{Y}$, the true label for him/her. Unless otherwise specified, we focus on the binary classification task (i.e., we assume $\mathcal{Y} = \{0, 1\}$) and assume without loss of generality that $\mathcal{X} = \mathbb{R}^k$. Individuals in T are assumed to be sampled i.i.d. from a population. A learning algorithm uses the training data T to fit a predictive *model* (or hypothesis) $h : \mathcal{X} \rightarrow \mathcal{Y}$ that predicts

the label for out-of-sample instances. More precisely, let $\mathcal{H}$ be the hypothesis class consisting of all the models the learning algorithm can choose from. The learning algorithm receives T as the input and selects a model $h \in \mathcal{H}$ that minimizes some notion of empirical loss on T . For the ease of notation, when the predictive model in reference is clear from the context, we denote the predicted label for an individual with feature vector $\mathbf{x}$ by $\hat{y}$ where $\hat{y} = h(\mathbf{x})$.

2.1 Distributive Fairness as Equality of Opportunity

Next, we present the Equality-of-Opportunity (EOP) framework for capturing (un)fairness in the outcomes produced by predictive models for decision-subjects. The core idea at the heart of EOP is to distinguish between factors that can morally justify inequality in outcomes among decision subjects and factors that are morally irrelevant and, ideally, should *not* impact outcomes. An equal opportunity policy is one through which an individual’s outcome only depends on the former—and not the latter—types of factors. Several models have been proposed to translate EOP into precise mathematical terms (see, e.g., [16]). At a high-level, these models break down an individual’s attributes into two categories: *circumstance* and *desert*³, where *circumstance* captures *all* factors that should not affect the individual’s outcome; *outcome* refers the harm/benefit or *utility* he/she receives due to the decision made about him/her; and *desert/deservingness* encapsulates *all* factors that can justify unequal utilities among individual decision subjects. More precisely, let c denote the *circumstance*, capturing factors that are not considered legitimate sources of outcome inequality (e.g., race in recidivism prediction). Let d denote a *scalar* summarizing factors that are viewed as legitimate sources of inequality (e.g., prior criminal convictions). Let ϕ denote the policy that governs the distribution of *utility*, u , among subjects. An individual subject’s utility is a consequence of his/her *desert* and *circumstance* and the implemented policy. Formally, let $F^\phi(\cdot | c, d)$ specify the cumulative distribution of utility under policy ϕ among people with *desert* level d and *circumstance* c . EOP requires that for individuals with similar *desert* d , the distribution of utility should be the same—regardless of their *circumstances*:

DEFINITION 1 (EQUALITY OF OPPORTUNITY (EOP)). A policy ϕ satisfies EOP if for all *circumstances* c, c' and all *desert* levels d , $F^\phi(\cdot | c, d) = F^\phi(\cdot | c', d)$.

With similar reasoning as Heidari et al. [11]’s, we can show that existing notions of fairness can be cast as special cases of our EOP interpretation. (For instance, by assuming a subject’s true label y captures their *desert*, their sensitive group membership is their *circumstance*, and their predicted label $\hat{y}$ reflects their utility, we obtain equality of odds.)

2.2 Parameter Estimation for EOP

We restrict our attention to the family of fairness models that can be cast as instances of EOP, and set out to estimate the parameters

³Existing formulations of EOP use the term “effort” or “effort-based utility” instead of “desert”, but we argue that this model fundamental should summarize all factors that are viewed as legitimate sources of inequality in outcomes—even if those factors have nothing to do with the person’s literal effort. We propose the use of the concept and terminology of “desert” instead to capture our broader interpretation.

c , d , and u that best captures the human perception and judgment of fairness—given a particular decision-making context. To formulate fairness as EOP, we need to specify which attributes belong to a subject’s circumstance (c), what constitutes their desert (or deserved outcome) (d), and how the decision they receives impacts their utility (u). Providing substantive answers to these questions requires human feedback, so we provide a framework that turns the input from a human participant (indexed by p) into a mathematical formulation of fairness.

Identifying circumstance c . From each participant p , we ask a series of k questions inquiring whether p believes feature $j \in \{1, \dots, k\}$ can justify inequality in outcomes. Let $\mathbf{z}_p$ denote the subset of features that p deems morally irrelevant. We treat each possible value that features in $\mathbf{z}_p$ can take as a distinct circumstance c . **Estimating desert d .** Consider a participant p and a decision subject with characteristics $\mathbf{x}$ and y . Let d_p specify the subject’s overall desert according to p . We assume d_p is not directly observable, but there exists a function $\delta_p : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$, such that $d = \delta_p(\mathbf{x}, y)$. That is, δ_p maps the information the participant observes about the decision subject (i.e., $\mathbf{x}$ and y) to his/her desert. For simplicity, we assume δ_p is a linear function of $\mathbf{x}, y$, that is, there exists a coefficient vector $\boldsymbol{\delta}_p$ such that $d = \boldsymbol{\delta}_p \cdot [\mathbf{x}, y]$ (where $[\mathbf{x}, y]$ is the vector concatenation of feature vector $\mathbf{x}$ and true label y). To estimate the $\boldsymbol{\delta}_p$, we ask p a series of Q pairwise comparison questions. In each question $q = 1, \dots, Q$, we present p with the information about two hypothetical decision subjects, $i_1^q = [\mathbf{x}_1^q, y_1^q]$ and $i_2^q = [\mathbf{x}_2^q, y_2^q]$, and ask her which one of them she considers to be more deserving of receiving the positive prediction.⁴ We assume p would pick i_1^q with probability

$$\Phi(\boldsymbol{\delta}_p \cdot [\mathbf{x}_1^q - \mathbf{x}_2^q, y_1^q - y_2^q]) \quad (1)$$

where $\boldsymbol{\delta}_p$ is the unknown parameter we wish to estimate and Φ represents the cumulative distribution function of the standard normal distribution. Note that our modeling choices here closely matches prior work (see, e.g., [18]).

With this model in place, we can readily find the Maximum Likelihood Estimator (MLE) of $\boldsymbol{\delta}_p$ using p ’s answers to the pairwise questions presented to her. More precisely, let a^q be p ’s answer to the q ’th pairwise question, in which she has expressed confidence level $c \in \{1, 2\}$. a^q is $|c|$ if p chooses subject 1 in response to question q , and it is $-|c|$ otherwise. We find the maximum likelihood estimator of $\boldsymbol{\delta}_p$ by solving the following optimization:

$$\arg \min_{\boldsymbol{\delta}} - \sum_{q=1}^Q \log \Phi(a^q \boldsymbol{\delta} \cdot [\mathbf{x}_1^q - \mathbf{x}_2^q, y_1^q - y_2^q]) \text{ s.t. } \|\boldsymbol{\delta}\|_2 \leq 1 \quad (2)$$

Estimating individual utility u . Let u_p denote the utility a decision subject $[\mathbf{x}, y]$ earns as a result of being subject to decision rule h according to participant p . For simplicity, we assume u_p is a linear function of $\mathbf{x}, y, \hat{y}$. That is, there exists $\mathbf{v}_p$, such that $u_p = \mathbf{v}_p \cdot [\mathbf{x}, y, \hat{y}]$. To estimate $\mathbf{v}_p$, we ask p a series of Q pairwise comparison questions. In each question $q = 1, \dots, Q$, we present p with the information about two hypothetical decision subjects, $i_1^q = [\mathbf{x}_1^q, y_1^q, \hat{y}_1^q]$ and $i_2^q = [\mathbf{x}_2^q, y_2^q, \hat{y}_2^q]$, and ask her which one of them she believes would benefit more from his/her decision. We

assume p would pick i_1^q with probability

$$\Phi(\mathbf{v}_p \cdot [\mathbf{x}_1^q - \mathbf{x}_2^q, y_1^q - y_2^q, \hat{y}_1^q - \hat{y}_2^q]) \quad (3)$$

where $\mathbf{v}_p$ is the underlying parameters we wish to estimate. With this model in place, we can compute the MLE of $\mathbf{v}_p$ using p ’s answers to the pairwise questions presented to her.

2.3 Methods of Aggregation

For each individual participant p , we estimate the circumstance, desert, and utility functions that best capture his/her judgments. We will denote these as $\mathbf{z}_p$, $\boldsymbol{\delta}_p$, and $\mathbf{v}_p$, respectively. Using these estimates for $\mathbf{z}_p$, $\boldsymbol{\delta}_p$, and $\mathbf{v}_p$, we can construct an EOP-based notion of fairness (Definition 1) according to participant p ’s judgments. To aggregate these individual notions into one representing society as a whole, we consider two natural aggregation methods.⁵ The first method is inspired by Borda Count:

- A feature is part of a subject’s circumstance if more than half of participants (or voters) deem it to be so.
- $\boldsymbol{\delta}^*(\mathbf{x}, y) = \frac{1}{N} \sum_{p=1}^N \boldsymbol{\delta}_p \cdot [\mathbf{x}, y] = \left(\frac{1}{N} \sum_{p=1}^N \boldsymbol{\delta}_p \right) \cdot [\mathbf{x}, y]$.
- $\mathbf{v}^*(\mathbf{x}, y, \hat{y}) = \frac{1}{N} \sum_{p=1}^N \mathbf{v}_p \cdot [\mathbf{x}, y, \hat{y}] = \left(\frac{1}{N} \sum_{p=1}^N \mathbf{v}_p \right) \cdot [\mathbf{x}, y, \hat{y}]$.

The second method is inspired by the hierarchical Bayesian modeling approach. We assume there exists a $\boldsymbol{\theta}$ that captures society’s perception of desert. A participant p in the society has his own $\boldsymbol{\delta}_p$ which is a noisy version of $\boldsymbol{\delta}$. To estimate $\boldsymbol{\delta}$ and $\boldsymbol{\delta}_p$ ’s jointly, we solve the convex optimization problem:

$$\arg \min_{\boldsymbol{\delta}, \boldsymbol{\theta}} - \sum_p \sum_q \log \Phi(a^q \boldsymbol{\delta}_p \cdot [\mathbf{x}_1^{p,q} - \mathbf{x}_2^{p,q}, y_1^{p,q} - y_2^{p,q}]) \quad (4)$$

s. t. $\|\boldsymbol{\delta}_p - \boldsymbol{\delta}\|_2 \leq \lambda$ and $\|\boldsymbol{\theta}_p\|_2 \leq 1, \|\boldsymbol{\delta}\|_2 \leq 1$.

We estimate the individual and aggregate utility coefficients (that is, $\mathbf{v}_p$ ’s and $\mathbf{v}$) similarly.

3 STUDY DESIGN

We conduct a series of questionnaires on AMT to construct a notion of fairness that best captures participants’ perception of fairness in a hypothetical crime prediction context.

A note on the use of ML in criminal justice. The use of ML for risk assessment in the criminal justice system have been heavily and validly criticized (see, e.g., [1, 17, 26]). Many scholars argue that such tools must be entirely banished to give way to sorely-needed structural changes in the criminal justice system (including but not limited to bail reform in the pretrial context). Among noteworthy objections to the COMPAS tool, in particular, is its lack of transparency [21], efficiency gains [6], and the numerous biases encoded in the data it is based off of [9] (for example, one source of bias is that the target label “rearrest within two-years” is an imperfect proxy for “future crimes” [28]). While we do not condone the use of ML in criminal justice processes (in particular, the COMPAS tool for criminal risk assessment), given the current wide-spread use and impact of such systems, we believe it is paramount for the ongoing debate around ML-fairness in this context to be grounded in stakeholders’ values and collective moral judgments.

⁴Note that different participants may respond to different sets of questions, hence to be precise, all the notation defined in this part must have a p superscript. For brevity, when the participant in reference is clear from the context, we drop the superscript p .

⁵Many alternative aggregation methods are conceivable. For brevity and as reasonable starting points, we only study these two.

3.1 Scenarios and Contexts

We first introduce the goals and scope of our study to participants as follows:⁶ If the participant needs further information to understand the task, he/she can access a longer version, which includes several examples illustrating the task. Figure 15 in the Appendix shows a longer version of the introduction to the HIT, which includes several examples.

Background and Task Description. Data-driven decision-making algorithms are increasingly employed to automate the process of making important decisions for humans, in areas such as credit lending, medicine, criminal justice, and beyond. **Organizations** in charge of decision-making can utilize massive datasets of historical records to learn a **decision-making** rule capable of making accurate **predictions** about never-before-seen individuals. Such predictions often serve as the basis for consequential **decisions** for these individuals. (From this point on we will refer to **individuals** subject to decisions as **decision subjects**.)

In recent years, several studies have shown that **automated/algorithmic decisions** made in the above fashion may disparately impact certain groups and individuals. For instance, in the context of credit lending, the decision-making rule may systematically disadvantage loan applicants belonging to a certain racial group and reject their loan applications more frequently. These observations have raised many questions and concerns about the fairness of automated decisions.

The goal of our study is to understand your moral reasoning and perception about what it means for automated decisions to be fair—considering the specifics of the decision-making context. We would like to know your ethical judgment through your answers to the following questions:

- Which **attributes** of a decision subject do you consider **morally acceptable** for the decision-making rule to base its decisions on?
- Comparing the attributes of two decision subjects, which one of them do you believe is more **deserving** of receiving a **better (more desirable) decision**?
- How do you believe automated decisions will **impact** these subjects? We'd like you to imagine how algorithmic decisions may contribute to the overall **happiness, satisfaction, & well-being** of a decision subject.

Throughout the experiment, we focus on a criminal risk prediction context, described below:

Decision-Making Context. In court-rooms across the United States, data-driven decision-making algorithms are employed to predict the likelihood of future crimes by defendants. These algorithmic predictions are utilized by judges to make sentencing decisions for defendants (e.g., setting the bail amount, or time to be spent in jail). Decision-making algorithms use historical data about past defendants to learn about factors that highly correlate with criminality. For instance, the algorithm may learn from past data that: 1) a defendant with a lengthy criminal history is more likely to reoffend if set free on bail—compared to a first-time offender, or 2) defendants belonging to certain groups (e.g., residents of neighborhoods with high crime rate) are more likely

to reoffend if set free. These automated predictions may directly translate into sentencing decisions. For instance, a defendant who is predicted to have a high risk of reoffending may be sentenced to jail, whereas a defendant who is predicted to have a low risk of reoffending may be set free on bail.

3.2 User Interfaces

Through our *conversational interface*⁷ the participant then responds to three questionnaires. The first questionnaire contains five questions, and the second and third parts, each consists of 25 questions. Once the participant completes the main three questionnaires, they can optionally provide us with their demographic information. (Participants had the option to skip this exit questionnaire).

Part 1: Identifying $\mathbf{z}$. To determine which attributes of a decision subject belong to the category of morally irrelevant, we first display the following introductory text:

In the first part, we would like to understand your moral reasoning about the following: Which attributes of a defendant do you consider morally acceptable for the decision-making rule to base its predictions on?

Example: one may believe it acceptable for the decision-making rule to take the subject's criminal history into account, but find it unacceptable for his/her parents' criminal history to impact whether he/she is predicted to have a high or low risk of reoffending.

(Note that this example is only meant to illustrate the task. You may have a very different opinion.)

The participant then responds to a simple questionnaire consisting of 5 questions of the following form:

To what extent do you agree with the following statement:
It is ethically acceptable for the attribute [...] (which can take one of the following values: [...]) to impact the decision a defendant receives.

For an example of this type of question, see Figure 1. The participant can choose their response from a 4-point Likert scale ("Disagree", "Somewhat Disagree", "Somewhat Agree", and "Agree"). The participant was encouraged (but not required) to provide a justification for their choice in a free-form text.

To what extent do you agree with the following statement:

It is ethically acceptable for the attribute "**race**" (which can take one of the following values: white, non-white) to impact the decision a defendant receives.

Figure 1: A typical question in part 1 of our experiments.

Part 2: Estimating δ . In the second part, our goal is to estimate the desert function (i.e., δ). We first display the following introductory text:

⁶Throughout, text boxes display the content exactly as it was shown to our study participants through the user interface.

⁷Conversational UIs are known to increase the user's attention and feel more natural to work with. See [19].

In the second part, we would like to understand your moral reasoning about the following: Comparing the attributes of two defendants, which one of them do you believe is more deserving of receiving a more lenient decision?

Example: Consider two defendants with similar attributes, except for their employment status—one unemployed, the other a local government employee. One may consider the employed subject more deserving of the “low risk to reoffend” prediction.

(Note that this example is only meant to illustrate the task. You may have a very different opinion.)

We then prompt the participant to respond to a series of pairwise comparison questions of the following form:

From an ethical standpoint, between the following two decision subjects, who do you believe deserves a more lenient decision?

For an example of this type of question, see Figure 2. The participant can choose either “Clearly subject 1”, “Possibly subject 1”, “Possibly subject 2”, and “Clearly subject 2”. The participant was encouraged (but not required) to provide a justification for their choice in a free-form text.

From an ethical standpoint, between the following two decision subjects, who do you believe deserves a more lenient decision?

ATTRIBUTE	SUBJECT #1	SUBJECT #2
Age Category	Younger than 25	Older than 25
Race	Non-white	Non-white
Gender	Male	Male
Charge Degree	Felony	Felony
Prior Counts	4	9
Actual Outcome	Will reoffend	Will reoffend

Note: The decision subject differences are marked in blue. If you are unsure about the meaning of any attribute, hold the cursor on it to see a definition.

Figure 2: A typical question in part 2 of our experiments.

Part 3: Estimating v . To estimate the utility function (i.e., v), we first display the following introductory text:

In the third part, we would like to understand your moral reasoning about the following: Given the attributes of two defendants, which one of them do you believe would benefit more from their respective algorithmic decision? In responding to this question, imagine yourself in the circumstances of these two defendants and think about how the sentencing decision they receive may affect their lives.

Example: Consider two defendants with similar attributes, except for their number of dependants (one with two children and another with no dependants.). One may believe that a “low risk to reoffend” prediction would contribute more to the overall satisfaction, happiness, and well-being of the subject who has kids.

(Note that this example is only meant to illustrate the task. You may have a very different opinion.)

We then prompt the participant to respond to a series of pairwise comparison questions of the following form:

From an ethical standpoint, between the two following decision subjects, who do you think will benefit more from their algorithmic decision?

For an example of this type of question, see Figure 3. The participant can choose either “Clearly subject 1”, “Possibly subject 1”, “Possibly subject 2”, and “Clearly subject 2”. The participant was encouraged (but not required) to provide a justification for their choice in a free-form text.

From an ethical standpoint, between the two following decision subjects, who do you think will benefit more from their algorithmic decision?

ATTRIBUTE	SUBJECT #1	SUBJECT #2
Age Category	Older than 25	Older than 25
Race	Non-white	White
Gender	Male	Male
Charge Degree	Misdemeanor	Felony
Prior Counts	4	4
Algorithmic Decision	Low risk to reoffend	Low risk to reoffend
Actual Outcome	Will not reoffend	Will not reoffend

Note: The decision subject differences are marked in blue. If you are unsure about the meaning of any attribute, hold the cursor on it to see a definition.

Figure 3: A typical question in part 3 of our experiments.

3.3 Experimental Design

We used a summarized and pre-processed version of the COMPAS data set [1] to generate our questions. We restricted attention to the following attributes for each defendant: **gender** (1 if the defendant is male, 0 otherwise), **age** (1 if the defendant is less than 25, 0 otherwise), **race** (1 if the defendant is not Caucasian, 0 otherwise), **charge degree** (1 if felony, 0 otherwise), **prior counts** (a positive integer). We set the **true label** to 1 if the defendant reoffends within 2 years, 0 otherwise). Note that in the original COMPAS dataset, the true label reflects whether a defendant is *rearrested* for an alleged crime. This is clearly different from our choice of the true label (whether the defendant actually reoffends), which we chose to reduce the layers of complication. The **predicted label** indicates if the COMPAS tool gives a score of 5 or higher⁸ to the defendant.

In part 1, we asked one question for each of the five features in the data set. In part 2, to generate a question q , we picked the first defendant (i_1^q) i.i.d. from the data set, then chose the second defendant (i_2^q) such that it does not differ with the first one in more than two attributes. The reason for this restriction was two-fold: (1) As shown in Figure 4 it allows us to recover the underlying weight vector confidently with a smaller number of questions; (2) it reduces the cognitive burden of comparing the two defendants with one another for the participants⁹. Questions in part 3 were

⁸Note that the COMPAS tool bases its risk scores on many features other than the ones considered here. Our focus is not on the COMPAS tool and its inner workings. We simply use the COMPAS risk scores (provided by Angwin et al. [1]) to generate binary recidivism predictions for defendants in our hypothetical scenario.

⁹Comparing two objects that differ across multiple dimensions is a type of *multitasking*—which notoriously reduces human attention and leads to erroneous responses. See https://en.wikipedia.org/wiki/Human_multitasking and the references therein.

Figure 4: Cosine similarity between the true weight vectors (according to which responses are simulated) and the estimated ones (MLE base on simulated responses).

generated similarly. Part 2 and 3 each consisted of 25 questions. We chose the number of questions by simulating our model of participants responses (Equation 1) and finding the number of questions necessary to learn their weight vectors with a satisfactory accuracy (i.e., 90%). See figure 4.

We validated our interface design through two rounds of *pilot studies*—one internal (among members of our research group) and one on AMT (among 20 crowd-workers). Based on the feedback we received from pilot participants, we made several (minor) changes to the text to improve readability and allowed participants to edit their responses. We also decided to randomize the order in which questionnaires in part 2 and 3 are shown to the participants, and included one attention-check question in each part.

4 EXPERIMENTS

To illustrate our framework, we conduct a series of questionnaires on AMT to construct a notion of fairness that best captures people’s perceptions in our hypothetical criminal risk prediction context. Through our experiments on AMT, we gathered a data set consisting of 90 participants’ responses to our questionnaires. As a final step, we asked the participants to provide us with their demographic information, such as their age, gender, race, education, and political affiliation. The purpose of asking these questions was to understand whether there are systematic variations in perceptions of fairness across different demographic groups.¹⁰ Table 1 summarizes the demographic information of our participants and contrasts it with the 2016 U.S. census data. In general, AMT workers are not a representative sample of the U.S. population (they have Internet access and are willing to complete online tasks). In particular, our participants were younger and more liberal compared to the average U.S. population. The majority of our participants were white, with Black and Hispanic minorities underrepresented compared to census numbers. (As we mentioned earlier, our experiment is meant as a proof-of-concept; the proper representation of disadvantaged minorities is one of the essential preconditions for the ethical use of such participatory frameworks in the real world.)

Part 1 results: Identifying morally irrelevant features. Figure 5 shows the distribution of responses to questions in part 1 of our experiment. Responses represent participants’ extent of agreement with the statement “It is ethically acceptable for the attribute [...] to impact the decision a defendant receives?” The majority of

¹⁰ Answering to this part was entirely optional and did not affect the participant’s eligibility for monetary compensation.

Table 1: Demographics of our AMT participants compared to the 2016 U.S. census data.

Demographic Attribute	AMT	Census
Male	61%	49%
Female	39%	51%
Caucasian	73%	61%
African-American	9%	13%
Asian	10%	6%
Hispanic	4%	18%
Liberal	65%	33%
Conservative	24%	29%
High school	34%	40%
College degree	58%	48%
Graduate degree	8%	11%
18–25	12%	10%
25–40	65%	20%
40–60	17%	26%

Figure 5: Responses to part 1 questions. Participants believe the number of prior counts, charge degree, & age can morally justify inequality in sentencing decisions, while race and gender are morally irrelevant.

participants strongly believe the number of prior counts and charge degree can morally justify inequality in sentencing decisions. The majority also strongly believe that race is morally irrelevant. For gender and age, opinions are mixed. On average, participants consider gender as part of a defendant’s circumstance, and age as a morally relevant feature, but their degree of confidence is generally lower compared to the other three features.

Table 2: The average coefficient of each attribute in our estimated desert and utility functions.

	Race	Gender	Age	Charge	Prior	y	$\hat{y}$
δ	0.11	-0.19	0.13	-0.46	-0.25	-0.41	N/A
v	0.06	-0.07	0.09	-0.09	-0.08	-0.18	-0.43

Part 2 results: Estimating the desert function (δ). Figure 6 illustrates the distribution of weights allocated to the true outcome y when estimating δ . As one may expect, the true outcome has a significant weight, but somewhat surprisingly, on average participants gave similar weights to charge degree (-0.46) and the true label (-0.41). Figure 12 (left panel) in Appendix A shows the distribution of weights given to the five features of a defendant (race, gender, age, charge degree, and prior counts) when estimating desert (δ).

(For ease of interpretation, a Gaussian probability density function with the same mean and variance as the data is plotted on the opposite axis with bolder typeface. We have used 20 bins for these illustrations.) According to our participants, younger defendants generally deserve a more lenient sentencing decision. The same holds if a defendant is female or non-white. Participants consider a defendant less deserving of a lenient outcome if they are charged with a felony, or they have a large number of prior convictions. Perhaps surprising, participants, on average, gave a larger weight (-0.25) to charge degrees as opposed to prior conviction counts (-0.46).

Part 3 results: Estimating the utility function (v). Figure 7 illustrates the distribution of weights allocated to the true outcome y when estimating v and Figure 8 shows the distribution of weight on the predicted outcome, $\hat{y}$. The contribution of $\hat{y}$ is significantly larger than other factors. Figure 12 (right panel) in Appendix A shows the distribution of weights given to theft features (race, gender, age, charge degree, and prior counts) when estimating utility (v). Trends are similar to δ weights, although the spreads are generally larger. Unlike the case of desert, charge degree and prior counts contribute similarly to utility.

Our notion of fairness. Table 2 summarizes the average weights obtained by estimating δ and v . With circumstance, desert, and utility specified, our EOP-based notion of fairness requires that those who are similarly *deserving*, should have the same prospect for *utility* regardless of their *circumstance* (see Definition 1). Notice that our notion of fairness is much more nuanced than the two definitions of fairness at the center of COMPAS controversy—in particular, many factors beyond the true and predicted labels impact the assessment of fairness in our formulation. Also observe that while the majority of participants consider race and gender as morally irrelevant in part 1, on average, they believe these factors should play a role in determining a defendant’s deserved sentencing decision and the harm/benefit they perceive as a result of their sentencing outcomes.

Variations across gender, race, age, education, & political views.

Tables 3, 4 show the average weights allocated to each attribute of a defendant when δ and v are computed for various demographic subgroups of participants. In terms of desert (δ), liberal participants give a relatively large positive weight to race, while conservative participants give it a weight close to 0 (indicating that they believe race should not affect whether a defendant deserves leniency). A similar trend holds for female vs. male, non-white vs. white, and young vs. old participants. Younger participants allocate a large positive weight to age (indicating that they believe younger defendants are generally more deserving of leniency), while our older participants allocate a small negative weight to age.

In terms of harm/benefit or utility to defendants, conservative participants believe female defendants benefit more from leniency, while liberal participants give a relatively small (but still negative) weight to gender. Similar to the case of desert, younger participants allocate a large positive weight to age, while older participants give a small negative weight to age.

Sensitivity to the aggregation method. Figure 9 contrasts the weights computed according the two aggregation methods described in Section 2.3. While the weights are generally comparable, we observe that the hierarchical model puts a significantly large weight on prior counts. (This is to be expected as this aggregation

Table 3: Average δ (desert coefficient) for participants belonging to various demographic segments. Significant differences are highlighted.

Participant	Race	Gender	Age	Charge	Prior	y
Liberal	0.16	-0.22	0.15	-0.46	-0.27	-0.43
Conservative	0.01	-0.15	0.10	-0.45	-0.22	-0.34
High-school	0.08	-0.18	0.11	-0.51	-0.31	-0.40
University	0.13	-0.20	0.14	-0.44	-0.22	-0.41
Male	0.06	-0.20	0.17	-0.42	-0.27	-0.41
Female	0.19	-0.18	0.07	-0.53	-0.21	-0.40
White	0.09	-0.19	0.11	-0.48	-0.28	-0.45
Non-white	0.19	-0.23	0.20	-0.41	-0.17	-0.29
Young (<40)	0.14	-0.20	0.18	-0.44	-0.24	-0.41
Old	0.01	-0.17	-0.03	-0.53	-0.30	-0.39

Table 4: Average v (utility coefficient) for participants belonging to various demographic segments. Significant differences are highlighted.

Participant	Race	Gender	Age	Charge	Prior	y	$\hat{y}$
Liberal	0.05	-0.06	0.08	-0.10	-0.08	-0.20	-0.47
Conservative	0.07	-0.17	0.10	-0.09	-0.08	-0.07	-0.29
High-school	0.13	-0.09	0.13	-0.10	-0.08	-0.23	-0.32
University	0.02	-0.06	0.06	-0.09	-0.08	-0.15	-0.48
Male	0.06	-0.04	0.07	-0.06	-0.08	-0.13	-0.46
Female	0.05	-0.12	0.11	-0.14	-0.09	-0.26	-0.38
White	0.06	-0.06	0.09	-0.05	-0.07	-0.16	-0.42
Non-white	0.05	-0.10	0.07	-0.20	-0.11	-0.23	-0.45
Young (<40)	0.04	-0.09	0.13	-0.12	-0.08	-0.18	-0.43
Old	0.10	0.00	-0.06	0.00	-0.09	-0.17	-0.40

method zeros in on the points of agreement among participants, and indeed, the majority of participants believe that a defendant with a large prior count is less deserving of leniency.)

Comparison with existing notions of fairness. Figure 10 illustrates the effect of existing fairness-enhancing methods on the utility distributions across (desert, circumstance) groups. Note that to simplify the illustration, we have made the desert space binary as follows: a defendant is considered deserving of leniency (or in short, deserving, on the plots) if and only if his/her desert score is above the median score in his/her own circumstance group. This ranked-based approach to desert is compatible with luck egalitarian interpretations of Equality of Opportunity (see e.g., [20]).

Figure 10 demonstrates the importance of formulating fairness as the comparison between *distributions*—as opposed to single aggregate measures, such as accuracy, false positive, or false negative rates. Our key finding here is that compared to other fairness-enhancing methods, our fairness notion performs better in terms of *equalizing the utility distributions* across groups. In particular, it pushes up the population to higher utility levels (for example, 75% of non-deserving and all of deserving subjects receive utilities scores greater than or equal to -1), whereas other methods lead to long tails for the non-deserving sub-populations and lower mean utilities overall. For a more detailed comparison, see Appendix A.

4.1 Qualitative Findings

At a high-level, most participants thought race and gender should not be a deciding factor in determining a defendant’s sentencing outcome, even if they are statistically correlated with certain types of crimes. Opinions about age were mixed. Some thought younger people are less in control of their decision, so they should be treated with leniency. Others thought a punishment at a younger age is more effective at preventing future crimes. (See Appendix A.1 for

Figure 6: The distribution of δ for the true label (y) when δ is estimated for each participant, separately.

Figure 7: The distribution of v for the true label (y) when v is estimated for each participant, separately.

Figure 8: The distribution of v for the predicted label ($\hat{y}$) when v is estimated for each participant, separately.

Figure 9: (Left) δ for the two aggregation methods proposed in Section 2.3. (Right) v for the two aggregation methods.

the quotations supporting these statements.) The majority of participants thought prior convictions should definitely affect a defendant’s sentencing outcome. Although, some noted that in a biased judicial system, certain races might have systematically higher numbers of prior convictions. In part 2, aside from other features, we provided the participants with the defendants’ true labels (i.e., whether the defendant will reoffend in the future). While they generally thought a participant who will not reoffend in the future is more deserving of leniency, some noted that such information is not available at the time of decision-making, so neither defendant should get preferential treatment. In part 3, we provided the participants with the defendants’ predicted labels (i.e., whether the algorithm predicts the defendant has a high risk of reoffending in the future). Participants generally agreed that a low-risk prediction is more beneficial for defendants, while a high-risk prediction is particularly harmful if the defendant does not have a long criminal history and won’t reoffend.

Overall, participants found our task thought-provoking and engaging. Our findings indeed showcases their complex reasoning about desert and utility: Our participants took race and gender into account when determining a defendant’s desert and utility. For instance, if a defendant is female, participants generally consider her more deserving of a lenient sentencing outcome (compared to an identical male defendant). As another example, participants generally believed that the utility of staying out of jail is higher for

a young defendant (compared to an older but otherwise identical defendant). These beliefs are aligned with the participants’ perception of the social and economic consequences of sentencing decisions for female and young defendants, respectively (as reflected in their justifications).

5 CONCLUDING DISCUSSION

Our work addresses a critical gap in the growing literature on Fair-ML: Currently, ML practitioners are the primary designers of mathematical notions of fairness, and the literature lacks a systematic way of bringing human judgment into the process of formulating fairness and translating it into computationally tractable measures. While our work takes an important step toward empowering people’s participation in this process, our simplified framework and small-scale case-study on AMT must be viewed with their limitations and caveats in mind. Importantly, any tractable mathematical notion of fairness by definition fails to capture all the nuances and complexities involved in people’s ethical reasoning and assessments. *Framing* and *messaging* always play a crucial role in human-subject studies, especially when we attempt to elicit people’s moral judgment about a *hypothetical* decision-making scenario. For a detailed discussion of learning-based approaches to moral dilemmas, see [4]. We dedicate the rest of this section to elaborate on some of the specific limitations of our work. A more comprehensive discussion (including that of our various simplifying assumptions and modeling choices) can be found in Appendix B.

Whose opinion should matter? Attempts to bring human judgment into the process of formulating fairness should not stop at taking input from *ordinary people* (e.g., participants on AMT), for a variety of ethical considerations—chief among them is the principle that a simple majority should not be allowed to determine the rights and freedoms of the minority. In practice, one must strive to involve all stakeholders—including impacted individuals, community members, domain experts, ethicists, policymakers, and system designers—in the process. These parties must collectively deliberate and determine the *weight* of the input from each side (akin to [8]). For example, under some conditions, it may be appropriate to give higher weight to people who have been previously disenfranchised or excluded from deliberations, or those with higher stakes [29].

Figure 10: The distributions of utility across (desert, circumstance) groups when a classifier is trained on the COMPAS dataset using different fairness interventions.

There must also exist pathways for overruling certain opinions (e.g., those that violate basic rights, legal requirements, or ethical norms; or those reflecting harmful biases and stereotypes). Another consequential choice is the mechanism for aggregating opinions. For example, what are the axioms and desiderata that such mechanisms should satisfy [24]? While addressing these vital considerations is out of the current work's scope, they mark critical prerequisites for the ethical deployment of our framework in practice.

Focus on the EOP family of notions. We constructed a context-dependent mathematical formulation of outcome fairness by obtaining people's responses to pairwise questions and finding the EOP-based fairness notion that best captured those responses. We took it as a given that the EOP family of fairness notions can represent the human judgments of outcome fairness reasonably well—at

least in the hypothetical scenario presented to our participants. We acknowledge that real-world scenarios are always highly complex and cannot be faithfully squeezed into any mathematical equation, but that does not justify picking an arbitrary formulation of fairness off the shelf without careful deliberations involving stakeholders.

Causal links among desert, circumstance, and utility. We did not directly address the causal relationship between various quantities of interest in our framework. We note, however, that the luck-egalitarian models of equality of opportunity do account for the possibility of circumstance causally impacting desert by defining desert in relevant (quantile-based) terms. For a detailed discussion, we refer the reader to [11, 16, 20]. We treated the problem of estimating desert, circumstance, and utility as decoupled problems. A natural direction for future work consists of a holistic treatment of these problems and examining the impact of a joint estimation on our findings.

Linear models. We assumed that both desert and utility functions have a simple linear form. Linear models simplified our experiments because (1) they are easy to train with a small sample size without the risk of overfitting, and (2) are interpretable. In particular and as shown in Figure 4, we can learn the coefficients of a linear preference vector (up to a constant) by asking only 25 questions on average from each participant. While the linearity restriction facilitated our illustration, we must emphasize that in practice, such simplifying choices have ethical implications beyond scalability and interpretability, so they must be thoroughly analyzed before being used in practice.

Framing Bias. As with any human-subject experiment, we cannot completely rule out the influence of framing on our empirical observations. We attempted to limit this by ensuring neutral language throughout and reminding the participant they may have a different opinion every time we presented them with a subjective example to demonstrate the task. We emphasize, though, that framing is not an issue specific to our study—no single human-subject study on its own can claim full-robustness to its experimental setup. Several replications and variation studies of the initial work are always required to establish the robustness of its findings to framing and other psychological effects.

In conclusion, we believe that the insights offered by our work can inform future studies of this kind, and ultimately, contribute to the proper formulation of fairness in real-world applications of ML. We close by noting that a more comprehensive case-study of our framework must be broader in scope and resources; it should carefully provide participants with additional relevant information about decision subjects and inner workings of the decision-making algorithm; and account for the effect of participants' stakes, life experiences, political views, and expertise in their judgments.

ACKNOWLEDGMENTS

This work was in part done while Heidari was a postdoctoral fellow at ETHZ, and Yaghini was a visiting student from EPFL carrying out his Master's dissertation work at ETHZ. Heidari acknowledges partial support from NSF IIS2040929. Any opinions, findings, and conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation and other funding agencies

REFERENCES

- [1] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. *ProPublica* (2016).
- [2] Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonneau, and Iyad Rahwan. 2018. The moral machine experiment. *Nature* 563, 7729 (2018), 59.
- [3] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *arXiv preprint arXiv:1703.00056* (2017).
- [4] Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. 2017. Moral decision making frameworks for artificial intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [5] William Dieterich, Christina Mendoza, and Tim Brennan. 2016. COMPAS risk scales: Demonstrating accuracy equity and predictive parity. *Northpoint Inc* (2016).
- [6] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science advances* 4, 1 (2018).
- [7] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the Innovations in Theoretical Computer Science Conference*. ACM, 214–226.
- [8] Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P. Dickerson, and Vincent Conitzer. 2020. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence* 283 (2020), 103261.
- [9] Nina Grgic-Hlaca, Elissa M Redmiles, Krishna P. Gummadi, and Adrian Weller. 2018. Human perceptions of fairness in algorithmic decision making: A case study of criminal risk prediction. In *Proceedings of the 2018 World Wide Web Conference*. 903–912.
- [10] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. In *Proceedings of Advances in Neural Information Processing Systems*. 3315–3323.
- [11] Hoda Heidari, Michele Loi, Krishna P. Gummadi, and Andreas Krause. 2019. A Moral Framework for Understanding of Fair ML through Economic Models of Equality of Opportunity. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*.
- [12] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudík, and Hanna Wallach. 2018. Improving fairness in machine learning systems: What do industry practitioners need?. In *Proceedings of the 2018 CHI conference on human factors in computing systems*.
- [13] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. Inherent trade-offs in the fair determination of risk scores. In *Proceedings of the Innovations in Theoretical Computer Science Conference*.
- [14] Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. 2016. Data and analysis for 'How we analyzed the COMPAS recidivism algorithm'. <https://github.com/propublica/compas-analysis>.
- [15] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–35.
- [16] Arnaud LeFranc, Nicolas Pistolesi, and Alain Trannoy. 2009. Equality of opportunity and luck: Definitions and testable conditions, with an application to income in France. *Journal of Public Economics* 93, 11–12 (2009), 1189–1207.
- [17] John Monahan and Jennifer L. Skeem. 2016. Risk assessment in criminal sentencing. *Annual review of clinical psychology* 12 (2016), 489–513.
- [18] Ritesh Noothigattu, Snehal Kumar 'Neil' S. Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D. Procaccia. 2018. A voting-based system for ethical decision making. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [19] Bill Rice. 2019. Conversational UI is Becoming the Preferred Consumer Interface: 12 Reasons Why. <https://clinc.com/conversational-ui-preferred-consumer-interface/>. *Clinic* (2019).
- [20] John E. Roemer. 2002. Equality of opportunity: A progress report. *Social Choice and Welfare* 19, 2 (2002), 455–471.
- [21] Cynthia Rudin, Caroline Wang, and Beau Coker. 2018. The age of secrecy and unfairness in recidivism prediction. *arXiv preprint arXiv:1811.00731* (2018).
- [22] Debjani Saha, Candice Schumann, Duncan C. McElfresh, John P. Dickerson, Michelle L Mazurek, and Michael Carl Tschantz. 2020. Human Comprehension of Fairness in Machine Learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 152–152.
- [23] Nripsuta Ani Saxena, Karen Huang, Evan DeFilippis, Goran Radanovic, David C Parkes, and Yang Liu. 2019. How do fairness definitions fare? Examining public attitudes towards algorithmic definitions of fairness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 99–106.
- [24] Amartya Sen. 1986. Social choice theory. *Handbook of mathematical economics* 3 (1986), 1073–1181.
- [25] Megha Srivastava, Hoda Heidari, and Andreas Krause. 2019. Mathematical notions vs. human perception of fairness: A descriptive approach to fairness for machine learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2459–2468.
- [26] Megan Stevenson. 2018. Assessing risk assessment in action. *Minn. L. Rev.* 103 (2018), 303.
- [27] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 CHI conference on human factors in computing systems*.
- [28] Michael Wick, Swetasudha Panda, and Jean-Baptiste Tristan. 2019. Unlocking fairness: a trade-off revisited. In *Proceedings of Advances in Neural Information Processing Systems*.
- [29] Allison Woodruff, Sarah E Fox, Steven Rousso-Schindler, and Jeffrey Warshaw. 2018. A qualitative exploration of perceptions of algorithmic fairness. In *Proceedings of the 2018 CHI conference on human factors in computing systems*.
- [30] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the World Wide Web Conference*. 1171–1180.