POTs: Protective Optimization Technologies

Rebekah Overdorf
imec-COSIC KU Leuven
rebekah@esat.kuleuven.be

Bogdan Kulynych
EPFL SPRING Lab
bogdan.kulynych@epfl.ch

Ero Balsa
imec-COSIC KU Leuven
ebalsa@esat.kuleuven.be

Carmela Troncoso
EPFL SPRING Lab
carmela.troncoso@epfl.ch

Seda Gürses
imec-COSIC KU Leuven
sguerses@esat.kuleuven.be

ABSTRACT

In spite of their many advantages, optimization systems often neglect the economic, ethical, moral, social, and political impact they have on populations and their environments. In this paper we argue that the frameworks through which the discontents of optimization systems have been approached so far cover a narrow subset of these problems by (i) assuming that the system provider has the incentives and means to mitigate the imbalances optimization causes, (ii) disregarding problems that go beyond discrimination due to disparate treatment or impact in algorithmic decision making, and (iii) developing solutions focused on removing algorithmic biases related to discrimination.

In response we introduce *Protective Optimization Technologies*: solutions that enable optimization subjects to defend from unwanted consequences. We provide a framework that formalizes the design space of POTs and show how it differs from other design paradigms in the literature. We show how the framework can capture strategies developed in the wild against real optimization systems, and how it can be used to design, implement, and evaluate a POT that enables individuals and collectives to protect themselves from unbalances in a credit scoring application related to loan allocation.

1 MOTIVATION

We are facing a new type of digital system whose organizing principle is *optimization*. These systems became the dominant paradigm, as software engineering shifted from packaged software and PCs to services and clouds, enabling distributed architectures that incorporate real-time feedback from users [30].

Through this process, digital systems became layers of technologies, metricized under the authority of objective functions. These functions drive, among others, the selection of software features, the orchestration of cloud usage, and the design of user interaction and growth planning [24]. In contrast to traditional information systems, which treat the world as a static place to be known and focus on storage, processing, transport, and organizing information, *optimization systems* consider the world as a place to sense and co-create. They seek maximum extraction of economic value by optimizing the capture and manipulation of people's activities and environments [1, 13].

Optimization systems apply a logic of operational control that focuses on outcomes rather than the process [46]. While this introduces efficiency and allows systems to scale, they also pose social risks and harms such as social sorting, mass manipulation, majority dominance, and minority erasure. In the vocabulary of optimization,

these systems create substantial *externalities* that arise due to the inadequacy of their *objective functions* to address the world.

Moreover, optimization systems hold great potential to shift power. The fast pace at which they manipulate users and environments obscures their effect, making it difficult to devise strategies to contest them. Optimization also often leads to asymmetrical concentration of resources in the hands of a few companies which can collect large scale data and muster the computational power to process these in the pursuit of financial gain [28, 46]. This centralizes governance and reconfigures market structures, creating an imbalance of power that benefits a select portion of society.

Fairness frameworks, we claim, have come to being as a response to the rise of optimization systems. They aim at solving associated problems, but often don't provide an in-depth characterization of these systems. To address this gap, we take a step back to gain a better understanding of the problem that fairness intends to respond to. We explore some fundamental shifts in the way digital systems are engineered to organize the world around us. We find that the problems that may arise are much greater than algorithmic unfairness, and that they cannot simply be solved by diligent service providers. Instead, they require new mental models and techniques to reason about strategies to counter them.

Specifically, we introduce Protective Optimization Technologies (POTs) which enable those affected by optimization systems to influence, alter, and contest these systems *from the outside*. We show how POTs are different from other protective technologies. We demonstrate the suitability of our framework by showing how it can encompass existing protection strategies, and how it can be used to design new POTs, using credit scoring as a use case. Finally, we discus the limitations and challenges involved in the design and deployment of POTs.

2 THE OPTIMIZATION PROBLEM

We call *optimization systems* those systems that capture and manipulate user behavior and environments under the logic of optimization. That is, systems whose operation relies on an *optimization algorithm*. For instance, ride sharing applications such as Uber, which rely on optimization to decide on the pricing of rides; navigation applications such as Waze, which rely on optimization to propose best routes; banks, which rely on optimization to decide whether to grant a loan; and advertising networks, which rely on optimization to decide what is the best advertisement to show to a user.

In this section, we start with an overview of those aspects and challenges of optimization system design that result in the common negative outcomes that usually surface during deployment

—that are typically (dis)regarded as *'externalities'*— and end with an evaluation of the ability of service providers to mitigate these.

2.1 Externalities of Optimization Systems

We first present an overview of *'externalities'* of optimization system design that result in common negative outcomes, risks and harms that usually surface during deployment. Externalities refer to situations when the actions of a group of agents, e.g., consumption, production and investment decisions, have "significant repercussions on agents outside of the group" [50]. The following are some of the common externalities intrinsic to optimization systems:

Disregard for non-users and environments. Optimizing the service for targeted users results in non-users and inhabitants of environments affected by the system being outside the optimization model. Traffic and navigation services only take into account their users and how to move them the fastest through the city, exposing *non-users*, i.e., people that do not use the service, to heavier traffic. Hence, residents of streets that were neither intended nor designed for heavy or non-local traffic experience externalities [33].

Disregard for certain users. Many optimization systems provide the most benefit to a subset of *"high-value"* users or to a particular population segment that does not match their complete user base. For instance, in the popular augmented reality mobile game Pokémon Go the placement of Pokémon and in-game resource stations rely on real world locations and maps, heavily benefiting players in urban areas and leaving players in rural areas and black neighborhoods starved of rarer Pokémon and resources [27, 54].

Externalization of exploration risks to users and environments. Optimization systems benefit from experimentation to reduce risks associated with environmental unknowns. Common practices in software engineering such as trialling new features through A/B testing involve experimentation on users. However, exploration often means that risks stemming from unknowns are pushed to users and their surroundings [5], a problem exacerbated by the trend of frequent system updates and real time optimization.

Distributional shift. Optimization systems built on data from a particular area or "domain" may underperform or downright flounder when deployed in a different environment [51], e.g., a voice recognition algorithm that is only trained on men's voices fails to recognize women's voices [35, 47].

Unfair distribution of errors. As with distributional shift, this results in disproportionate allocation of errors to a minority group [26]. Here the cause is that optimization algorithms learn to maximize success by favoring the most likely option, i.e., they can misclassify minorities while maintaining high accuracy. Therefore, minorities underrepresented in training do not perform well under deployment. For example, facial recognition algorithms are known to misclassify faces of black women because of this issue [7].

Promotion of unintended actions to fulfill intended outcomes. Systems may find shortcuts to their optimization goals, also known as "reward hacking" [2], e.g., an autonomous vehicle recklessly tailing an ambulance to decrease travel time, or electricity grid manager choosing to cause a blackout in order to save energy [49].

Mass data collection. Optimization systems need massive amounts of data to function. The concentration of resources and power in data holders enables more accurate inferences about populations and individuals using the data. However, it puts the privacy of the individuals whose data is input to the optimization at risk, as it can be leaked through interactions with the system [55].

2.2 Solutions by Design?

Typically the experts argue that risks and harms of optimization systems arise because system providers (OSPs) *"choose 'wrong' objective functions"* or *"lack sufficient good-quality data"*, i.e., flaws and mistakes that amount to *poor design* [2]. A common response is therefore to devise countermeasures that allow OSPs to prevent or minimize the occurrence of these *'flaws'* [2], with the underlying assumption that, given adequate tools and means, OSPs will strive to *'fix'* or *'correct'* their systems. While developing methods that can improve design of optimization systems is absolutely necessary, for a variety of reasons, they may not always work.

First, assuming that the risks and harms optimization systems cause are *accidents* derived from *poor design* choices dismisses the possibility that those design choices may in fact be intentional, i.e., that the objective functions underlying optimization systems may actively *aspire for asocial or negative environmental outcomes*. OSPs may lack incentives to maximize society's welfare as opposed to their own benefit. For example, Uber aggressively optimizes users' fares to maximize the company's profit at the expense of meager earnings for its drivers [37]. Similarly, Waze prioritizes, above all else, route optimization for its users, thus it lacks incentives to modify the system so as to avoid troubling inhabitants of residential neighborhoods that do not use the application [20].

Second, even when OSPs have incentives to address the problems optimization causes, they may not be in a position or even be able to do so, e.g., they may lack knowledge about the needs of those affected by optimization. These may result from cost minimization strategies that compel them to relinquish feedback from the population subject to optimization, or the affected community may simply be negligible for their bottom line. OSPs may strive to collect the necessary data to mend optimization outcomes, yet such data may simply not be available for capture, e.g., in the case of Waze, the service provider cannot possibly aim to have real-time access to all the people affected by its navigation optimization, let alone their wishes and constraints. Ultimately, in practice, data often represents *"what is easy to capture"* and thus provides a biased account of the people and environments it supposedly measures, leaving out key nuances required to better optimization [23].

These assumptions about the ability of the OSP to affect change thus limit their ability to address the harms and risks that stem from optimization systems. Ideally, under such conditions, service providers should internalize harms and risks, e.g., through stringent design practices, regulation and taxation, and provide democratic forms of governance. However, as long as this is not the case, we can consider OSPs as potentially unable, lacking incentives or unwilling to address the externalities of their optimization systems, rather than considering these as the result of poor design choices or accidents.

As a result, we need new mental models and techniques that enable designers to reason about strategies that not only counter the negative effects of optimization from *within* the system, but also

from *outside* of the system. Furthermore, we need to capture the fact that we may consider optimization functions that are different from those embedded in the original optimization algorithms, to the extent that it may consider variables and contour conditions not even present in the original optimization system.

3 PROTECTIVE OPTIMIZATION TECHNOLOGIES

We consider that optimization systems operate on *users*' inputs and interact with the *environment* in which they are deployed. The system's outputs thus affect both users and environments, at both individual and collective levels. We leave the definition of environment open so as to cover any object, human, individual or collective, e.g., *non-users* that do not directly interact with the optimization system.

In this context, we introduce *POTs—technological solutions that those outside of the optimization system deploy to protect users and environments from the negative effects of optimization.* POTs build on the idea that optimization systems infer, induce and shape events in the real world to fulfill objective functions. POTs analyze how events (or lack thereof) affect users and environments, then reconfigure these events to influence system outcomes, e.g., by altering the optimization constraints or poisoning the system inputs.

We specifically conceive POTs to address the negative externalities of optimization. To this end, POTs take a holistic perspective, considering the interaction of the algorithm with the rest of the optimization system and the environment.

For simplicity, in this paper we consider systems relying on a single optimization algorithm. However, our definition of optimization systems does not equate them to a single algorithm. For example, in addition to the concrete optimization task of providing navigation routes, Waze as a whole also provides services such as notifying users of gas prices or coordinating friends for pickups. We however only consider functionality subject to optimization when designing POTs.

3.1 A POTs design framework

We propose a framework based on partially observable stochastic games [25] to reason about the design of protective optimization technologies. In Section 4 we discuss how this framework differs from other protective technologies. For illustration purposes, throughout this section we use the Waze navigation service as a reference optimization system. In Section 5.1 we instantiate the framework to design a new POT, and in Section 5.2 we show how it accommodates existing POTs against other optimization systems.

Table 1 summarizes the framework notation.

Agents and interactions. We consider an *optimization system* that relies on an *optimization algorithm* OPT() to operate. The optimization system interacts with a set of *agents* $I = \{i\}$ (see Figure 1). The agents include both users and non-users, and other elements of the environment in which the optimization system is deployed In Waze, among others, the set of agents includes application users, non-users that experience changes in traffic as Waze changes routing, and the roads and neighborhoods themselves.

We consider that the optimization system and the agents co-exist in the *World*. At any point in discrete time t, the World is in *state*

Table 1: Notation

i	Agent
a	Action
r	Optimization algorithm output
s	System's state
$\theta(t)$	Agent acting at time t
$\pi(\mathbf{s})$	Agent's action policy
$\kappa(\mathbf{s}, a)$	Optimization algorithm's output policy
$\omega(\mathbf{s}, a)$	Observation function
$\tau(\mathbf{s}, a, r)$	State transition function
OPT$(\mathbf{s}, a)$	Optimization algorithm

$s \in \mathcal{S}$, with $\mathcal{S}$ the space of possible states. We consider that a state contains all information about the status of any entity in the World. Therefore it contains the optimization system status, which may include past interactions with users, and the status of the external agents: users, non-users, and the environment. In Waze, the state is composed of all the internals of the Waze systems, including models trained from users' information used to make routing decisions, as well as information about non-users and the actual state of roads and other environmental factors.

Agents perform individual actions $a \in \mathcal{A}_i$, where $\mathcal{A}_i$ denotes the set of possible actions that an agent i can perform. This set includes interactions of users with the system, as well as actions by non-users, or actions that change the system's environment. In Waze examples of actions could be: for users, looking for routes or reporting incidents; for non-users, counting cars passing by their houses; and for the environment, having works on a road.

Naturally, neither agents, nor the optimization system can see the full World's state. They can only view the part of the state that reflects information they have seen, and their internal state. We model this incomplete vision of the state through the function $\omega(\mathbf{s})$. This function provides an agent with the subset of the state s that it is able to see. For instance, $\omega(\mathbf{s})$ gives the Waze system the system's internal state, i.e., all interactions with users, and its vision of the environment according to the received reports. To users, $\omega(\mathbf{s})$ gives their own record of actions, and their vision of the environment.

To capture that optimization systems consider changes in their environment to maximize the extraction of value, we consider that the World's state evolves every time an agent i performs an action. We consider that at time t one and only one agent acts, and we denote her index by $\theta(t)$. We note, however, that other cases can be accommodated, e.g., concurrency of actions can be achieved by considering that agents are in fact collectives, and the concurrent action is a joint action of the collective.

We consider that the action an agent i issues given a sequence of system's states $\mathbf{s}_t = [s_t, s_{t-1}, \ldots, s_1]$ is governed by a probabilistic *policy* $\pi_i(\mathbf{s}_t)$. In the case of Waze, users requesting a route, or non-users having an accident, can make the World transition to a new state that accounts for this information. As an example, the policy $\pi_i(\mathbf{s}_t)$ describes the probability of a user's next action being 'report incident', or 'request route', depending on her state.

In our framework, state transitions are governed by the probabilistic function $\tau(\mathbf{s}_t, a_t, r_t)$. This function receives as input the sequence of system states up to current time t $\mathbf{s}_t$, the triggering
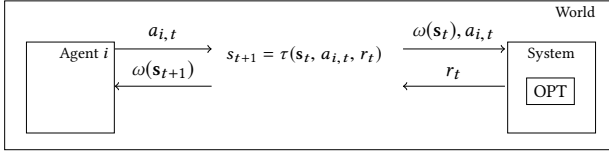
Figure 1: Optimization system interaction model

Figure 2: Optimization system and POTs design strategies. Both algorithms consider the state, the World's transition function, and their counterpart policies to find the best strategy to maximize their benefit. Note that for POT the benefit is defined over (a subset of) the agent's population.

action a_t, and $r_t \in \mathcal{R}$, the response of the optimization system to the agent's action. The space of outputs $\mathcal{R}$ depends on the application under consideration. In Waze, for instance, a report makes the World transition to a new state that records the incident.

The optimization algorithm OPT() provides a response according to a probabilistic policy $\kappa(\mathbf{s}_t, a_t)$ computed to meet an objective function as discussed below. In the Waze example, the optimization algorithm is called whenever a user requests navigation suggestions, and r is a proposed route, where $\mathcal{R}$ is formed by all possible routes.

Optimization systems' benefits. Optimization systems seek to maximize the economic value from their interactions with users and environment. We model the *immediate benefit* that an optimization system obtains when the World is in state s as a function $B_o(s)$. In a slight abuse of notation, we denote the expected benefit of issuing a reaction r_t to an agent's action $a_{i,t}$ given a sequence of state transitions $\mathbf{s}_t$, as $B_o(\mathbf{s}_t, a_{i,t}, r_t)$. It is computed as the expectation of the immediate benefit taken over all possible next states:

$$B_o(\mathbf{s}_t, a_{i,t}, r_t) = \mathbb{E}[B_o(s_{t+1}) \mid \mathbf{s}_t, a_{i,t}, r_t]$$
$$= \sum_{s' \in \mathcal{S}} B_o(s') \Pr[s' \mid \mathbf{s}_t, a_{i,t}, r_t]$$
$$= \sum_{s' \in \mathcal{S}} B_o(s') \Pr[\tau(\mathbf{s}_t, a_{i,t}, r_t) = s']$$

We can compute the optimization system's *expected discounted total benefit* $V_i^{\pi, \kappa}$, describing its benefit over n future actions as:

$$V_o^{\pi, \kappa}(\mathbf{s}_k) = \mathbb{E}\left[\sum_{t=k+1}^{n} \gamma^t B_o(\mathbf{s}_t, a_{\theta(t), t}, r_t) \mid \mathbf{s}_k \right]$$
$$= \sum_{t=k+1}^{n} \sum_{\substack{r' \in \mathcal{R} \\ a' \in \mathcal{A}}} \sum_{\mathbf{s}_t \in \mathcal{S}^t} B_o(\mathbf{s}_t, a', r') \cdot \Pr[\mathbf{s}_t, a', r' \mid \mathbf{s}_k],$$

where $\gamma \in [0, 1]$ is some discount factor. $\Pr[\tau(\mathbf{s}_t, a_{i,t}, r') = s']$ is the probability that r' leads the World into state s', and can be computed as follows:

$$\Pr[\mathbf{s}_t, a', r' \mid \mathbf{s}_k] = \Pr[r' \mid \mathbf{s}_t, a'] \Pr[a' \mid \mathbf{s}_t] \Pr[\mathbf{s}_t \mid \mathbf{s}_k]$$
$$= \Pr[\kappa(\mathbf{s}_t, a') = r'] \Pr[\pi_{\theta(t)}(\mathbf{s}_t) = a'] \Pr[\mathbf{s}_t \mid \mathbf{s}_k]$$

The total benefit enables us to define the optimization objective of the algorithm OPT (line 1 in Figure 2, left). OPT gets as input the current state of the World and the current action of an agent, as well as the models of World's transition function τ, agent order function θ, and the agents' policies π. With this information it finds the policy κ^* that maximizes its expected total benefit. Afterwards (line 2), the algorithm uses the found policy κ^* to select a reaction r to be outputted to the World.

Countering optimization systems. Similar to the optimization system, agents seek to increase their own benefit. We model the *immediate benefit* that an agent i obtains when the World is in state s, as a function $B_i(s)$. We denote as $B_i^\kappa(\mathbf{s}_t, a_{i,t})$ the expected benefit of performing an action $a_{i,t}$ at state $\mathbf{s}_t$ and compute it as follows:

$$B_i^\kappa(\mathbf{s}_t, a_{i,t}) = \mathbb{E}[B(s_{t+1}) \mid \mathbf{s}_t, a_{i,t}]$$
$$= \sum_{s' \in \mathcal{S}} \sum_{r' \in \mathcal{R}} \Pr[\tau(\mathbf{s}_t, a_t, r') = s']$$

Analogous to the optimization system, we compute an agent's *expected discounted total benefit* $V_i^{\pi, \kappa}$ describing her benefit over future n actions as:

$$V_i^{\pi, \kappa}(\mathbf{s}_k) = \mathbb{E}\left[\sum_{t=k+1}^{k+n} \gamma^t B_i^\kappa(\mathbf{s}_t, a_{\theta(t), t}) \mid \mathbf{s}_k \right]$$
$$= \sum_{t=k+1}^{n} \sum_{a' \in \mathcal{A}} \sum_{\mathbf{s}_t \in \mathcal{S}^t} \gamma^t B_i^\kappa(\mathbf{s}_t, a') \Pr[\mathbf{s}_t, a' \mid \mathbf{s}_k] \tag{1}$$

where $\gamma \in [0, 1]$ is some discount factor. $\Pr[\mathbf{s}_t, a' \mid \mathbf{s}_k]$, that models the probability of the agent acting at time t to perform action a' according to her policy $\pi_{\theta(t)}$, can be computed as follows:

$$\Pr[\mathbf{s}_t, a' \mid \mathbf{s}_k] = \Pr[\pi_{\theta(t)}(\mathbf{s}_t) = a'] \Pr[\mathbf{s}_t \mid \mathbf{s}_k]$$

Note that this computation requires that the agent has a model of other agents' actions to estimate the World's state evolution. Yet, $V_i^{\pi, \kappa}$ only considers i's own benefit, B_i^κ.

We also define a population's benefit function B_{pop} as a function over the immediate benefit of all agents and of the system itself:

$$B_{\text{pop}} = \sigma(B_o, \{B_i\}_{i \in \mathcal{I}})$$

Protectors is a group of agents (or a single agent) who wish to modify the benefit of the whole or a part of the population using a POT. To build the POT they first select the target users, and the sought effect on their benefits—by choosing the σ. This function can be defined to potentially override, cancel, or boost the benefit of any agent or the system. In order to affect the population benefit, each protector d solves an optimization problem to find a policy that maximizes the population's expected discounted total benefit $V_{\text{pop}}^{\pi, \kappa}$. It is computed analogously to an agent's expected total benefit (Eq. 1)—except it's based on the population benefit B_{pop}, not the individual benefit B_i.

The POT optimization problem is defined in Figure 2 (right). The first line illustrates how the POT is built, taking as input the state of the World, the state transition function τ, the agent order function θ, the optimization algorithm's reaction policy κ, and the policies of other agents' $\pi_{j \neq d}$. With this information, protectors find the best

policy π^* to select their actions (line 2) so as to obtain the desired result for the population.

In the Waze example a protector could build a POT to define actions on the optimization system that reduce traffic in front of her house. First, she sets herself as the 'beneficiary' ($B_{\text{pop}} = B_d$). The POT optimization problem takes the current state of the roads, a model of the World's state transitions, the Waze route optimization policy —based on users navigation requests and Waze's belief on the state of the map,— and a model of other agents' policies. The resulting POT optimization problem would find a policy π^* pointing where protector should report incidents to avoid cars being routed along her house [31].

We note that the OPT and POT problems are symmetric, and that they can consider each other's operation when computing their best strategy. This means that the proposed design method allows to capture the fact than the optimization system can deliberately counter the deployment of POTs. Analogously, since POTs consider the optimization system's underlying policy, the POTs can adapt to strategic changes of the system.

Acting under uncertainty. In reality, it is unlikely that the agents or the system have perfect knowledge about each other's policies or the World's state. To use the algorithms in Fig. 2, however, they need to establish models of all the parameters. These can be ad-hoc approximations, or data-driven models that get refined upon receiving observations.

Our framework accomodates for the possibility of inference of unknown parameters. For example, given a probabilistic model on the received observations $\omega(\mathbf{s})$, agents can use a maximum a posteriori estimate for the history of state transitions:

$$\hat{\mathbf{s}}_t = \arg \max_{\mathbf{s}_t \in \mathcal{S}^t} \Pr[\mathbf{s}_t \mid \mathbf{w}_{d, t-1}, \mathbf{a}_{d, t-1}],$$

where $\mathbf{w}_{d, t-1}$ is a sequence of observations $[\omega(\mathbf{s}_{t-1}), \omega(\mathbf{s}_{t-1}), \ldots, \omega(\mathbf{s}_1)]$. Similarly, given an initial guess of κ or π, they can refine their models from the received observations.

4 POTS AND OTHER PROTECTIVE TECHNOLOGIES

POTs are not the first technologies that aim to protect users from negative effects arising from the pervasive use of digital systems. Fairness technologies aim to mitigate unintended discrimination that arises in the use of machine learning and algorithmic decision making, privacy technologies aim to protect users' from surveillance, and security technologies aim to prevent the unauthorized use of information. Yet, these technologies fundamentally differ from POTs in their goals, their reasoning, and the inputs they take into account. We discuss these differences to position POTs among these protective technologies.

4.1 POTs and Fairness

Some of the negative externalities that POTs tackle relate to the problems that fairness solutions aim to mitigate. Even though the goals of each are comparable, we argue that the way in which fairness frameworks reason about the problems, and their underlying assumptions about incentives, limit their ability to address discontents that POTs attend to.

Research on fairness has eminently focused on decision-making systems [7, 18, 45]. Such systems typically take data from one individual as input and output a decision, based on an *algorithm*, that has a direct impact on that individual. For example, recidivism prediction systems take as input the data of a criminal defendant and predict how likely the individual is to reoffend [44]; fairness tools act on the algorithm to prevent or minimize inequitable outcomes for different groups [9, 11]. The use of algorithmic constraints to ensure a particular measure of fairness is what we refer to as *fairness frameworks* or, more simply, *fairness*. Over the last years, researchers have proposed various definitions of fairness [3, 11, 45], and there is a lack of consensus in what the paradigm encompasses. As such, we discuss fairness in broad terms and focus on the state-of-the-art.

Using the POTs framework to model the fairness approach, we observe the following. First, fairness focuses on the algorithm's bias, thus staying limited to the analysis of an algorithms' inputs and outputs. As a result, the system's state, s, tends to only contain information about the users of the system. Second, fairness generally considers that the benefit of an agent is a function of the output of the algorithm, i.e., $B_i(s_t) = B_i(r_t)$, determining that this benefit should be equitable across different groups. Finally, fairness typically assumes that the group of protectors is a single d that has access to the algorithm, i.e., the optimization system provider itself, which the POTs framework does not allow for[1].

The fairness approach hence adopts a narrower scope with respect to POTs. We now illustrate how that narrowness undermines fairness' ability to mitigate the harms optimization systems cause.

Static decision making. Fairness commonly considers decision-making as static, i.e., the algorithm and its environment are fixed.[2] In terms of the POTs framework, fairness assumes that the state of the world, s, that is fed to OPT does not evolve over time. Hence, fairness disregards the possibility of state transitions, τ, i.e., agents, i, feed actions, a, directly to OPT and directly receive the optimization output, r, that is input to the benefit function, B_o, as opposed to an updated s_{t+1}. However, under optimization systems, OPT does change, e.g., predictive policing results in increased crime reporting for patrolled neighborhoods if the algorithm is not updated to account for the fact that police visit certain neighborhoods more often than others [19]. Hence, fairness is ill-equipped to prevent the negative effects that arise when agents do evolve [32].

Post-deployment apathy. Fairness commonly focuses on the algorithm independent of the environment where it is deployed, thus, fairness stays limited to analyzing an algorithm's inputs and outputs.

This is reflected in that the state, s, is constrained to the data space of the inputs and outputs and that the benefit function, $B_i(r_t)$, is computed on the algorithm's output. However, in optimization systems many externalities occur only after the system is introduced into an environment. Because the state, s, that fairness considers does not reflect changes in the environment, fairness is made unaware of and cannot account for externalities, e.g., a navigation application may provide its users with *optimal* routes in a way

---

[1] With an exception of paper by McNamara et. al [36], which includes multiple authorities in the decision chain but not the users.

[2] A notable exception being a work on fairness in reinforcement learning [29], in which future changes of environments are considered.

that satisfies a given notion of fairness. Still, fairness cannot reason about the effect that the *fair* routes have on the environments that those routes traverse, i.e., traffic may increase in residential areas. This impact, especially when multiple routing applications are deployed, becomes evident only after deployment. POTs consider states, s_t, encompassing all agents and their environment as well as their evolution over time (according to τ). Thus POTs can address post-deployment externalities.

Disregard for non-users. By focusing solely on the algorithm's inputs and outputs, fairness limits itself to targeted users. While this may be due to post-deployment apathy, it may also be that non-users are simply ignored in spite of the OSP's awareness of potential post-deployment effects. The state, s, only includes information about the system's users, thus the solutions only consider benefit functions for them, i.e., B_i for agents i that directly interact with the algorithm. Optimization systems, however, have an impact on non-users and environments, or propagate externalities through environments to the people that populate them. For example, in Pokémon Go, the rural players are at a disadvantage, with fewer Pokémon and in-game resources. In parallel, non-users who live or work at popular Pokémon locations may find their environments flooded by gamers. Fairness solutions could ensure that rural and urban players' experience becomes more similar, but, as they are neither considering $B_i(s)$ for non-user agents i, nor include their data in the state s, they cannot be used to improve the situation for neighborhoods invaded by Pokémon players.

Blindness to allocation of resources. Fairness aims to ensure equitable outcomes across groups, one individual decision at a time, i.e., the algorithm produces each decision independently for each agent, although bound by the fairness constraints. While many optimization systems may indeed only be concerned with individual decisions, POTs consider the possibility where each individual decision is subordinate to a measure of optimality over a population. Since fairness only considers a state, s, as constituted solely by the algorithm's users data and, since this state does not evolve, it cannot consider interactions between individuals, or even the environment. Even though optimization systems may operate on individuals' inputs and provide individual outputs, they often focus on how to best distribute resources to a population or environment. For example, with Uber, each individual sees the outcome of the system as a single decision: match and price. However, the optimization system makes decisions about the entire ecosystem, optimizing the best matching of drivers to riders to reap the highest profit. The POTs framework captures interactions in the ecosystem and redistribution of resources by i) not constraining the agents, nor the state and ii) allowing the state to evolve to consider how actions from agents may affect future decisions for other agents. The framework allows both the *protectors* and the optimization algorithm OPT to include the effect of future actions, their own and from others, on the algorithm outputs. This is reflected on $V_i^{\pi,\kappa}(\mathbf{s}_t)$ taking into account the horizon of actions. Fairness, on the other hand, is commonly limited to objectives that can only depend on the pre-deployment state of the system.

Decontextualization. The decontextualization of fairness inherently implies that it accepts the adequacy of the algorithm under study and rarely questions its utility [17]. Thus, fairness solutions

mostly aim to ensure that different subgroups in the population are equally affected by the algorithm, only considering benefit functions, B_i, that are aligned with the goal of the algorithm.[3] In the context of optimization systems, this means that fairness does not question whether the objective function itself is just, only ensuring that people are equally subject to such effects. For instance, a predictive policing application that may be constrained for *fairness* but does not take into account the negative effects that predictive policing has on vulnerable populations —or a credit scoring algorithm that may be tuned to ensure that sub-prime loans are fairly distributed— are likely to lead to unjust outcomes. On the contrary, B_i in POTs design are not constrained. Hence, they enable the designer to consider alternative impacts arising from the goal of the optimization itself. Further decontextualization can occur when fairness work relies solely on legal conceptions of protected identities and fairness. One form of re-contextualization is to work with researchers, activists and community movements with intersectional aims in order to center affected communities, something that POTs are designed to do.

The intersection. Finally, we note that even though fairness and POTs design approaches differ in fundamental ways, as enumerated above, there exist fairness solutions that could be cast as POTs. These are de-biasing approaches that can be applied *after* the algorithm has been deployed, by agents *external* to the system. A paradigmatic example of such technology is Bolukbasi et al. proposal to de-bias word embeddings [6]. Their solution can be used to remove gender stereotypes, e.g., association of the words receptionist and female, after the embeddings have been computed, without destroying the utility of the associations they capture.

4.2 POTs and Security and Privacy

Security technologies are designed to protect systems' data *confidentiality* and *integrity*, and guarantee *availability* in the presence of adversaries. As such, they work within a system to ensure that it is operating as intended. Therefore, security technologies present similar limitations to fairness solutions when it comes to casting them as POTs.

Privacy enhancing technologies (PETs) are concerned with the collection and processing of data by systems (or individuals). They *may* be implemented by the system provider to limit the amount of data the optimization system collects from agents, or exposes to adversaries, e.g., using privacy-preserving cryptography to perform computation. In this case PETs are similar to fairness, i.e., they constrain the functionality of the system depending on privacy requirements much in the same way as fairness constrains decision making depending on a particular metric of fairness.

System users and non-users may also rely on PETs, e.g., obfuscation or encryption, to limit the amount of data an optimization system collects about them. In these scenarios PETs can overlap with POTs as they enable agents external to an optimization system to protect themselves against the system's negative outcomes.

For example, a PET may feed misinformation, or no information, to an optimization system to hinder its operation. Consider targeted advertising, where ad networks collect data about users such as

---

demographics, location, site visits, and ad clicks. Every time a user with a particular profile performs the action a = 'visit website' the ad network uses an optimization algorithm, OPT, to select the ad that has the highest click chance. A PET such as Tor [52] hides metadata that advertisers may use to identify individuals browsing the web, therefore preventing ad networks from profiling users. Another PET, AdNauseam [14] automatically clicks on every ad on every page a user visits so that ad networks obtain a polluted profile of the user that no longer represents their interests. Both PETs effectively undermine OPT's functionality; Tor exploits anonymity to thwart OPT, whereas AdNauseam exploits obfuscation.

Furthemore, PETs are often designed to protect individual agents and, as fairness, rarely consider non-users or the environment of the system. That means that they may not be suited to consider population benefits, i.e., they may only consider the case $B_{\text{pop}} = B_i$.

Finally, both security and privacy technologies consider an adversary that intentionally tries to harm the system and its users. POTs, on the other hand, target optimization systems whose negative externalities are not necessarily premeditated. Still, the POTs design framework we propose in Section 3.1 accommodates adversarial optimization algorithms that, e.g., consider counteracting POTs as part of their benefit function B_o, and strategic optimization systems that take into account the deployment of POTs by virtue of including the agents' policy π in OPT's operation (see Figure 2).

5 POTS INSTANCES

Next, we present an example of how to design a POT using the framework and discuss how the framework can model further strategies currently deployed against optimization systems.

5.1 Case study: credit scoring

We examine the viability of designing POTs to counteract a service provider that uses ML for credit scoring. By *credit scoring* we refer to the process that a bank undertakes to evaluate the risk of a particular loan application and decide as whether to accept or reject the application. Credit scoring systems are inherently designed to minimize the banks' risks and maximize their profit. The underlying algorithms that support these decisions can discriminate applicants on protected attributes like gender or ethnicity [10], or cause feedback loops for populations disadvantaged by the financial system [46]. These harms are often caused by capturing values that are a product of unjust realities as inputs, and then propagating these to the model's decisions. We aim to design POTs that help break out of such feedback loops and correct unjust decisions.

To design the POT, we rely on techniques from adversarial machine learning [34]. Unlike classical adversarial ML, however, the roles of the adversary and defender are reversed, i.e., the model deployer is the POTs' "adversary".

Setup. We simulate a bank (the OSP) and its clients (the agents), using the German credit risk dataset from the UCI Machine Learning Repository [16]. This dataset contains 1000 feature vectors representing loan applicants, where features include job type, housing type, amount of funds in the bank accounts, gender, etc., and the loan details: loan amount, duration, and purpose. Each example has a binary label encoding whether the loan was repaid (70% of examples) or not (30% of example).

We simulate the OSP as follows. First, following a common practice, we quantize all the continuous features into 5 quantiles. Then, the categorical and quantized features are one-hot encoded. We obtain 38-dimensional binary feature vectors. We split the dataset into 900 train examples and 100 test examples. We then train an SVM with an RBF kernel model to predict if applicants will default on the loan, choosing hyperparameters using 5-fold stratified cross-validation on the training dataset. If the model predicts 'default' for a given loan application, the OSP denies it, otherwise accepts it. This classifier achieves 77% accuracy on the test dataset (90% precision, 80% recall).

For this proof of concept we assume that protectors have perfect knowledge of the system workings: the ML model parameters and the actions of other agents. In reality they would need to estimate these parameters (see Section 3.1). Since this POT implements an adversarial machine learning attack, protectors can use model inversion [55] and take advantage of transferability of adversarial examples across different models [42].

The protectors aim to change the loan decisions to maximize different instantiations of population benefit functions, by changing loan applications. We assume that they cannot change any of the features that represent information about the individual, only the loan details: duration, amount, loan purpose. These features are arguably easy to modify in reality by changing the loan application form. For a given initial example x, we define its *transformations* as the feature vectors that change one of the following: 1) the loan amount bin, 2) the loan duration, 3) the loan purpose category.

Formalization. The OSP operates a supervised ML model. We denote by $\mathcal{D}$ a space of labeled examples $x, y \in \mathcal{D}$, where $x \in \mathcal{X}$ represents an individual and its loan details, and the label $y \in \{0, 1\}$ whether the individual defaulted on the loan or not. A model $\kappa : \mathcal{X} \rightarrow \{0, 1\}$ maps an example $x \in \mathcal{X}$ to an estimated binary label $\hat{y} \in \{0, 1\}$. The model takes as input a loan application and information about the applicant, and outputs a prediction that is interpreted as the the OSP's decision: deny if 'default' is predicted, accept otherwise.

An agent can perform two kinds of actions. First, apply for a loan. Second, after applying and *getting accepted*, she can repay the loan. Hence, the action space $\mathcal{A} = \{\text{'apply', 'repay'}\} \times \mathcal{X}$. The space $\mathcal{X}$ consists of all transformations of some initial loan application x.

The system's state at time t, s_t, consists of a training dataset $D_{\text{tr}} \subset \mathcal{D}$, and the previous loan decision. When the action a_t is 'apply', the reaction $r = \hat{y}$ of the optimization algorithm $\text{OPT}(s_t, a_t)$ is a decision. If the action is 'repay', however, the corresponding x is added to the training dataset D_{tr} with a 'repaid' label, effectively updating the model, and OPT returns the updated D'_{tr}. For simplicity, we assume that the model is retrained using the current dataset every time an 'apply' action is performed D_{tr}, even though in practice the model is retrained at each state update.

Protectors' goal—maximizing own benefit. In this scenario a protector has as goal to get her loan application to flip from deny to accept. She uses evasion attacks (also known as adversarial examples) [34, 53]. to modify her application form to obtain an 'accept' decision for the maximum loan amount. The $B_d^{\kappa}(s_t, a_t = (\text{'apply'}, x'))$ is hence equal to the loan amount if it gets accepted, and 0 otherwise. Denoting as $R(x')$ the loan amount encoded in a feature

vector x', the optimization problem that a protector d solves is:

$$\pi^* = \arg\max_{\pi \in \mathcal{X}} \sum_{x' \in \mathcal{X}} R(x') \Pr[\pi(\kappa) = x']$$

$$\text{s.t. } \kappa\left(\mathbf{s}_t, (\text{'apply'}, x')\right) = \text{'accept'}$$

which maximizes a degenerate case of $V_d^{\pi,\kappa}$, where only one step in the future is considered, and the protector only considers their own benefit as a component of the population benefit: $B_{\text{pop}} = B_d$.

Results. Our method is able to find loan applications that would be accepted for *all users* for whom the system denied the loan. We show in Figure 3 (left) a boxplot representing the distribution of the maximum increase in the loan amount for all users (the circles indicate how many individuals obtained that increase) depending on the number of manipulated features. The number of changed features reflects the *cost* of a transformation for the protector. Unsurprisingly, changing more features in the application enables larger increases in the loan amount. We note that a protector could change her benefit function to find a better trade-off between cost and profit for herself.

Protectors' goal—maximizing the benefit of others. In this scenario a group of protectors aim at ensuring that certain applicants, called target group G, have higher chances to obtain an 'accept' loan decision in the future. To this end they poison the training dataset of the model [4], i.e., apply and repay loans to trigger retraining operations. This strategy thus, imposes a significant burden on protectors, but we consider it realistic, e.g., the practice of taking out loans to improve one's own credit score is not uncommon in the United States [22, 40].

Let us assume that the benefit $B_i(s)$ of individuals in the target group G is equal to 1 if their loan is accepted, and 0 otherwise. Then, the population benefit that the protectors aim to maximize is the acceptance rate within the target group:

$$B_{\text{pop}}(s) = \frac{1}{|G|} \sum_{i \in G} B_i(s)$$

This is used by each protector to solve the POT optimization problem and obtain, a policy π^* defining actions that maximize $V_{\text{pop}}^{\pi,\kappa}$ over multiple steps, where θ is such that first all protectors act (i.e., poison the dataset), and afterwards, all agents from the target group act to get their loans.

For our experiments, we picked a target group of 22 loan applications made by individuals who have 'little' funds in both their checking and savings account, that would have repaid their loan application in reality, but were denied by the model. We *approximate* the policy π^* using a greedy probabilistic algorithm. We assemble the set of poisoning examples by randomly sampling an initial x from the dataset minus the target group, and adding its transformation x' to the set if 1) x' would get an 'accept' decision from the original model, and, 2) the acceptance rate for the target group by a model retrained on x' increases. We run the algorithm 10 times with different randomization seeds for sampling. To keep our simulation time reasonable we chose the best individuals to act as protectors and the best possible loan applications that maximize the objective. We note, however, that other groups could achieve such an objective, albeit at a higher computational cost to find the appropriate groups.
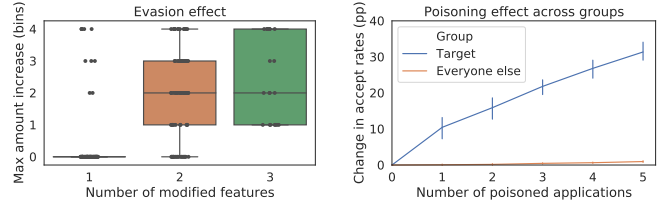

Figure 3: POTs simulation results

Results. Using the probabilistic algorithm, we are able to significantly improve the target population benefit $B_{\text{pop}}(s)$ in all simulations. Fig. 3 (right) shows the increase in the acceptance rate for the target group, and for others, depending on the number of protectors. A group of 5 protectors (about 0.5% of a poisoned dataset, initially containing 900 examples), is able to increase the target group's acceptance rate from 0% up to about 30%. As a side effect, the poisoning also slightly increases the acceptance rate for everyone else.

5.2 POTs in the Wild

In recent years, people have developed several strategies to counter the negative effects of optimization systems. Below we describe three such strategies and show how to formulate them as POTs.

Induced Uber surges. Uber manipulates prices in space and time, constituting geographies around supply and demand that both drivers and riders are unable to control, negatively impacting both with price falls and *surges*. In particular, Uber pushes prices up to respond to the scarcity of drivers, which drivers have collectively exploited by simultaneously turning off the Uber app on their devices to induce a price surge, then turning the app back on after a certain amount of time to take advantage of the surge [38].

In terms of a POT, drivers and riders are agents which perform the actions turn 'on'/'off' app and 'request ride', respectively. Responses r from Uber include the ride prices that OPT outputs as well as driver addition and removal. Protector drivers observe that prices depend on the number of drivers and riders recorded in the World's state, thus design a POT with a policy π^* to determine the number of absences that increases their group benefit, B_{pop}, by getting the OPT to output 'surge'. Then, following policy π^*, each protector performs action $a = \text{'off'}$, causing Uber's response to trigger a transition in the World's state whereby OPT outputs 'surge'.

Pokémon Go spawn encouragement. Pokémon Go users in rural areas enjoy fewer of the game's perks than those in more urban or even suburban areas: less Pokémon, less Pokéstops to collect resources, and less gyms to compete in. One obvious strategy to correct this imbalance is for users in rural areas to spoof their phone location to an urban area. A more elaborate strategy exploits the fact that Pokémon Go decides spawn points based on OpenStreetMaps[4] location information by reporting either false or previously unreported footpaths, swimming pools, or parks, all of which encourage new Pokémon to spawn [21].

---

[4]www.openstreetmap.org

In terms of POTs, players are agents whose actions are edits to OpenStreetMap and searches for Pokémon. The World's state depends, in part, on OpenStreetMap. The optimization system's reactions include r = 'add new Pokémon in a location'. A protector d designs a POT with a policy π^* that determines how Pokémon Go's OPT algorithm uses OpenStreetMap to spawn Pokémon, and then strategically chooses where and what to add on OpenStreetMap to increase her benefit B_d by triggering the desired spawns.

AdNauseam. Ad networks collect information on Internet users' browsing history and behavior to optimize ad placement, a practice that may lead to biased advertising [15]. AdNauseam[5] seeks to prevent this kind of negative optimization by blocking ads from view while clicking on each of the blocked ads, poisoning ad networks' user profiles to render them useless.

As a POT, AdNauseam users are agents that perform actions 'visit page' and 'click on ad'. The World's state encodes the agents' profiles the ad network builds. A protector d designs a POT with policy π^* to maximize the benefit B_d, defined as rendering the information the ad network collects useless. Then, according to π^*, AdNauseam clicks on every ad it sees.

6 DISCUSSION

Deployment challenges. That POTs work from outside the optimization system poses several challenges to their deployment. Firstly, the cost of POT design and deployment may be too high. The POT against credit scoring we study in Section 5.1 requires protectors to take loans, proving financially infeasible for many people, undermining collective action that requires coordination among a group of people. In fact, POTs that require collective action may be especially elusive, as they call for organization, communication, and mass participation of protectors (e.g., inducing Uber surges requires many drivers) which may not always be possible. OSPs may also try to actively inhibit or block POTs (e.g., Waze banning accounts reporting false traffic incidents), prompting protectors to engage in an arms race that further increases their costs. Such retaliation may further disadvantage those users who cannot afford the risk of being locked out of systems. Moreover, usability issues are important for POTs design, as POTs need to enable non-expert protectors to engage with a system without requiring them to master the underlying technical complexity.

The ethics of POTs. POTs may pollute input data to optimization systems, raising ethical issues analogous to those of obfuscation-based privacy technologies [39], which we discuss below.

Dishonesty. Since most optimization systems do not explicitly request users to provide information but rather sense or infer it —and often disregard non-users and environments,— the manipulation of inputs is not necessarily *dishonesty*, but a way to *introduce feedback into the cybernetic loop so that the system recognizes its externalities.*

Waste. POTs may require the investment of significant resources and thus considered wasteful, especially if they fail to yield meaningful results, or initiate an arms race with the OSP. However, POTs must be evaluated for their overall effect and the final allocation

of resources they achieve since they may, at the expense of investing additional resources, ensure far more equitable outcomes than those the system originally optimizes for.

Free riding. One may accuse protectors of free-riding whenever POTs maximize the benefit of one group at the expense of another. Free-riding is however inherent to optimization systems: Users of optimization systems may free-ride on non-users, whereas OSPs may free-ride on everyone. However, this free-riding may be tolerable taking into account the benefit that POTs bring by enabling those who bear the costs of optimization to compel OSPs to internalize the externalities they create.

On the other hand, if people deploy selfish POTs to maximize their benefit at any cost, the system may transition towards a *tragedy of the commons*, i.e., in the absence of a central policy that regulates allocation of resources, agents may turn to exhausting all available resources [41]. For instance, in the loan application example of Section 5.1, enabling agents to improve their chances of receiving money, and how much, may result in reckless decisions that leads individuals to acquire a debt that they cannot repay. While their misuse is a risk, POTs also enable agents to contest optimization systems that promote the very selfish strategies POTs may fall prey to [8].

Subversion and system damage. Since the goal of optimization systems is not to obtain knowledge per se but rather to use it to maximize a benefit function B_o, they selectively ignoring data deemed not relevant for the optimization implicitly distorting their vision of the populations and environments they impact. In this setting, POTs can be seen as tools to strategically codify information about agents that optimization systems choose to ignore. Moreover, charges of subversion and system damage must be measured up against the subversion and damage the optimization system causes itself. The blatant deployment of asocial or harmful optimization may justify agents' responses based on subversion and sabotage.

More worrying is the fact that by virtue of modifying, subverting, or sabotaging an optimization system POTs may elicit transitions in the system state that result on externalities, like those discussed in Section 2.1 for optimization systems) that a POT cannot possibly predict or account for. Note that this can be the case even when POTs are not engaging in the optimization game—e.g., opt to block or shut down optimization systems. They may still cause ripple effects in unpredictable and harmful ways.

If several POTs are deployed, and enter in an arms race, those agents with the most knowledge and resources are likely to deploy the most aggressive and effective POTs and have the most leverage. This in turn may undermine the ability of less powerful populations —who may need POTs the most— to have any impact effect on the system. This signals that well thought POTs must be built to provide less powerful actors with the means to respond to the potential abuse of power by those that have more information.

Finally, the existence of POTs that can address agents concerns after deployment may incentivize OSPs to deploy systems first, and address externalities only later, if and only if there is significant media or policy attention.

Accountability and Transparency. POTs can also overlap accountability and transparency frameworks.

Accountability. Rubin defines accountability as requiring *"the ability of one actor [. . .] to reward or punish [a] second actor on the basis of its performance or explanation"* [43, 48]. POTs promote accountability by enabling both disenfranchised agents to punish a system that negatively impacts them. When agents face an unaccountable OSP that offers no explanation or justification to the negative externalities of optimization, agents may rely on POTs to push back.

Pasquale further argues that *"an algorithmic accountability movement worthy of the name must challenge the [balance of power], rather than content itself to repair the wreckage left in their wake."* [43]. By enabling those outside the system to reclaim a share of the power OSPs wield, POTs have the potential to fit within such a movement.

Transparency. To tweak or intervene an optimization system using a POT one must understand what the system does, e.g., through reverse engineering, which in turn necessarily increases transparency. However, as a response to POTs, OSPs may render their systems more complex and opaque so that others (e.g., POT designers) cannot manipulate them. In addition, agents equipped with knowledge of the optimization system may choose to strategically refashion themselves to better fit the OSPs' optimization goals, providing the opposite outcome to empowerment, as POTs primarily seek to help users refashion the optimization system, but not themselves.

7 OUTLOOK

In this paper we identify negative externalities intrinsic to optimization systems that previous frameworks cannot address and introduce *protective optimization technologies* (POTs) as a solution. While this work represents a steady step forward in addressing the problems that arise from automation, many research questions remain open: What are further externalities of optimization? Can we design POTs to address them? Does our framework cover the design space of all possible POTs? What are non-technical approaches that can complement POTs (e.g., taxation)? Do POTs have other uses? E.g., can they support ongoing efforts to promote safety in automation or become useful not only post-deployment, but also to prevent the introduction of optimization systems in certain domains?

We believe that answering such questions requires an interdisciplinary approach that involves engagement with affected communities, policy and other academic disciplines. We hope that our work inspires the exploration of new avenues of research towards developing strong mechanisms to counter the negative effects of optimization in our society.

ACKNOWLEDGEMENTS

REFERENCES

[1] P. E. Agre. Surveillance and capture: Two models of privacy. *The Information Society*, 1994.

[2] D. Amodei, C. Olah, J. Steinhardt, P. F. Christiano, J. Schulman, and D. Mané. Concrete problems in AI safety. *CoRR*, abs/1606.06565, 2016.

[3] R. Berk, H. Heidari, S. Jabbari, M. Kearns, and A. Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 2018.

[4] B. Biggio, B. Nelson, and P. Laskov. Poisoning attacks against support vector machines. In *ICML*, 2012.

[5] S. Bird, S. Barocas, K. Crawford, and H. Wallach. Exploring or exploiting? social and ethical implications of autonomous experimentation in AI. In *FAT*, 2016.

[6] T. Bolukbasi, K. Chang, J. Y. Zou, V. Saligrama, and A. T. Kalai. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *NIPS*, 2016.

[7] J. Buolamwini and T. Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 2018.

[8] T. Cabannes, M. A. S. Vincentelli, A. Sundt, H. Signargout, E. Porter, V. Fighiera, J. Ugirumurera, and A. M. Bayen. The impact of GPS-enabled shortest path routing on mobility: a game theoretic approach. *Transportation Research Board 97th Annual Meeting*, 2018.

[9] A. Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 2017.

[10] D. K. Citron and F. Pasquale. The scored society: due process for automated predictions. *Wash. L. Rev.*, 2014.

[11] S. Corbett-Davies and S. Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*, 2018.

[12] S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq. Algorithmic decision making and the cost of fairness. In *ACM KDD '17*, 2017.

[13] M. R. Curry and D. Phillips. Surveillance as social sorting: Privacy, risk, and automated discrimination. In *Privacy and the phenetic urge: geodemographics and the changing spatiality of local practice*, 2003.

[14] C. H. Daniel and H. Nissenbaum. Engineering privacy and protest: a case study of adnauseam. In *IWPE*, 2017.

[15] A. Datta, M. C. Tschantz, and A. Datta. Automated experiments on ad privacy settings. *PoPETs*, 2015.

[16] D. Dheeru and E. Karra Taniskidou. UCI machine learning repository, 2017.

[17] J. Dressel and H. Farid. The accuracy, fairness, and limits of predicting recidivism. *Science advances*, 4(1), 2018.

[18] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 2012.

[19] D. Ensign, S. A. Friedler, S. Neville, C. E. Scheidegger, and S. Venkatasubramanian. Runaway feedback loops in predictive policing. *CoRR*, abs/1706.09847, 2017.

[20] L. W. Foderaro. Navigation apps are turning quiet neighborhoods into traffic nightmare. *The New York Times*, Dec 2017.

[21] A. Frank. Pokémon Go's maps now look a lot different. 2017.

[22] Z. Friedman. How personal loans can boost your credit score. *Forbes*, May 2017.

[23] A. Greenfield. *Radical technologies: The design of everyday life*. Verso Books, 2017.

[24] S. Gurses and J. van Hoboken. *Privacy after the Agile Turn*. Cambridge University Press, 2018.

[25] E. A. Hansen, D. S. Bernstein, and S. Zilberstein. Dynamic programming for partially observable stochastic games. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*, 2004.

[26] M. Hardt. How big data is unfair? Medium, September 2014.

[27] C. Huffaker. There are fewer Pokémon Go locations in black neighborhoods, but why? *The Miami Herald*, July 2016.

[28] T. Hwang. Computational power and the social impact of artificial intelligence. *arXiv preprint arXiv:1803.08971*, 2018.

[29] M. K. M. M. J. R. A. Jabbari, Shahin; Joseph. Fairness in reinforcement learning. *eprint arXiv:1611.03071*, 2016.

[30] I. Kaldrack and M. Leeker. There is no software, there are just services: Introduction. *There Is No Software, There Are Just Services*, 2015.

[31] I. Kuksov. Navigation wars. *Kaspersky Lab Daily*, October 2016.

[32] L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, and M. Hardt. Delayed impact of fair machine learning. In *ICML*, 2018.

[33] S. Lopez. On one of L.A.'s steepest streets, an app-driven frenzy of spinouts, confusion and crashes. *Los Angeles Times*, April 2018.

[34] D. Lowd and C. Meek. Adversarial learning. In *SIGKDD*, 2005.

[35] G. McMillan. It's not you, it's it: Voice recognition doesn't recognize women. *Time*, June 2011.

[36] D. McNamara, C. S. Ong, and R. C. Williamson. Provably fair representations. *CoRR*, abs/1710.04394, 2017.

[37] A. McNeilage. Uber's underpayment of drivers keeping it afloat, report finds. *The Guardian*, April 2018.

[38] M. Möhlmann and L. Zalmanson. Hands on the wheel: Navigating algorithmic management and uber drivers' autonomy. In *ICIS*, 2017.

[39] H. Nissenbaum and F. Brunton. Obfuscation: A user's guide for privacy and protest, 2015.

[40] B. O'Shea. What is a credit-builder loan? *Forbes*, April 2018.

[41] E. E. Ostrom, T. E. Dietz, N. E. Dolšak, P. C. Stern, S. E. Stonich, and E. U. Weber. *The drama of the commons.* 2002.

[42] N. Papernot, P. D. McDaniel, and I. J. Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *CoRR*, abs/1605.07277, 2016.

[43] F. Pasquale. Odd numbers. algorithms alone can't meaningfully hold other algorithms accountable. *Real Life*, August 2018.

[44] W. L. Perry. *Predictive policing: The role of crime forecasting in law enforcement operations.* Rand Corporation, 2013.

[45] G. Pleiss, M. Raghavan, F. Wu, J. M. Kleinberg, and K. Q. Weinberger. On fairness and calibration. In *NIPS*, 2017.

[46] M. Poon. Corporate capitalism and the growing power of big data: Review essay, 2016.

[47] J. A. Rodger and P. C. Pendharkar. A field study of the impact of gender and user's technical experience on the performance of voice-activated medical tracking application. *International Journal of Human-Computer Studies*, 2004.

[48] E. Rubin. The myth of non-bureaucratic accountability and the anti-administrative impulse. *Public Accountability: Designs, Dilemmas and Experiences*, 52, 2006.

[49] T. Simonite. When bots teach themselves to cheat. Wired, August 2018.

[50] D. A. Starrett. *Economic Externalities.* EOLSS, 2011.

[51] M. Sugiyama, N. D. Lawrence, A. Schwaighofer, et al. *Dataset shift in machine learning.* The MIT Press, 2017.

[52] P. Syverson, R. Dingledine, and N. Mathewson. Tor: The secondgeneration onion router. In *Usenix Security*, 2004.

[53] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013.

[54] P. Tassi. I am now a rural 'Pokémon GO' player and it's the worst. *Forbes*, August 2016.

[55] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart. Stealing machine learning models via prediction apis. In *USENIX Security Symposium*, 2016.