Bogdan Kulynych EPFL bogdan.kulynych@epfl.ch Rebekah Overdorf EPFL rebekah.overdorf@epfl.ch

# ABSTRACT

Algorithmic fairness aims to address the economic, moral, social, and political impact that digital systems have on populations through solutions that can be applied by service providers. Fairness frameworks do so, in part, by mapping these problems to a narrow definition and assuming the service providers can be trusted to deploy countermeasures. Not surprisingly, these decisions limit fairness frameworks' ability to capture a variety of harms caused by systems.

We characterize fairness limitations using concepts from requirements engineering and from social sciences. We show that the focus on algorithms' inputs and outputs misses harms that arise from systems interacting with the world; that the focus on bias and discrimination omits broader harms on populations and their environments; and that relying on service providers excludes scenarios where they are not cooperative or intentionally adversarial.

We propose *Protective Optimization Technologies (POTs)*. POTs, provide means for affected parties to address the negative impacts of systems in the environment, expanding avenues for political contestation. POTs intervene from outside the system, do not require service providers to cooperate, and can serve to correct, shift, or expose harms that systems impose on populations and their environments. We illustrate the potential and limitations of POTs in two case studies: countering road congestion caused by traffic-beating applications, and recalibrating credit scoring for loan applicants.

# **CCS CONCEPTS**

• Social and professional topics → Socio-technical systems;

# **KEYWORDS**

Fairness and Accountability; Protective Optimization Technologies

#### **ACM Reference Format:**

Bogdan Kulynych, Rebekah Overdorf, Carmela Troncoso, and Seda Gürses. 2020. POTs: Protective Optimization Technologies. In *Conference on Fairness, Accountability, and Transparency (FAT\* '20), January 27–30, 2020, Barcelona, Spain.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3351095. 3372853

# **1** INTRODUCTION

Advances in computational power, software engineering, and machine learning algorithms have been instrumental in the rise of

© 2020 Association for Computing Machinery.

Carmela Troncoso EPFL carmela.troncoso@epfl.ch Seda Gürses TU Delft / KU Leuven f.s.gurses@tudelft.nl

digital systems. Their ubiquity in our everyday activities raises concerns regarding the centralization of decisional power [1]. These concerns are amplified by the opaque and complex nature of these systems which results in hard-to-explain outputs [2] and unjust outcomes for historically marginalized populations [3–6].

Computer scientists counter these inequities through frameworks studied under the rubric of *fairness*, using a variety of formal fairness notions [7]. Their results have been instrumental in our understanding of the discriminatory effects of algorithmic decisions. The frameworks, however, rely on narrowing the inequity problem to primarily consider the discriminatory impact of algorithms, and assuming trustworthy providers. The narrow view has enabled valuable breakthroughs centered around the service providers' ability to address some inequities, but fails to capture broader harms or explore other ways to contest service-provider power [8].

In this work we investigate digital systems from a new perspective in order to understand how to address the broader class of harms that they cause. To achieve this, we characterize the type of systems in which algorithms are integrated and deployed. These systems typically build on distributed service architectures and incorporate real-time feedback from both users and third-party service providers [9, 10]. This feedback is leveraged for a variety of novel forms of optimization that are geared towards extraction of value through the system. Typically, optimization is used for technical performance and minimizing costs, e.g., by optimizing cloud usage orchestration or allocation of hardware resources. It has also become part and parcel of "continuous development" strategies based on experimentation that allow developers to define dynamic objective functions and build adaptive systems. Businesses can now design for "ideal" interactions and environments by optimizing feature selection, behavioral outcomes, and planning that is in line with a business growth strategy. We argue that optimization-based systems are developed to capture and manipulate behavior and environments for the extraction of value. As such, they introduce broader risks and harms for users and environments beyond the outcome of a single algorithm within that system. These impacts go beyond the bias and discrimination stemming from algorithmic outputs fairness frameworks consider.

Borrowing concepts and techniques from software and requirements engineering, and from economics and social sciences, we characterize the limitations of current fairness approaches in capturing and mitigating harms and risks arising from optimization. Among others, we show that focusing on the algorithms and their outputs overlooks the many 'externalities' caused by optimizing every aspect of a system; that discrimination is only one of the injustices that can arise when systems are designed to maximize gain; and that ignoring service providers' incentives and capabilities to enforce proposed solutions limits our understanding of their operation and further consolidates the power providers have regarding decisions and behaviors that have profound effects in society.

Bogdan Kulynych and Rebekah Overdorf contributed equally to this work.

FAT\* '20, January 27-30, 2020, Barcelona, Spain

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Conference on Fairness, Accountability, and Transparency (FAT\* '20), January 27–30, 2020, Barcelona, Spain,* https://doi.org/10.1145/3351095.3372853.

Finally, we propose *Protective Optimization Technologies (POTs)*, which aim at addressing risks and harms that cannot be captured from the fairness perspective and cannot be addressed *without* a cooperative service provider. The ultimate goal of POTs is to eliminate the harms, or at least mitigate them. When these are not possible, POTs can shift harms away from affected users or expose abusive or non-social practices by service providers. We illustrate the potential of POTs to address externalities of optimization-based systems in two case studies: traffic-beating routing applications and credit scoring. We also identify numerous techniques developed by researchers in the fields of security and privacy, by artists, and by others, that, though not necessarily designed to counter externalities, can be framed as POTs.

# 2 FAIRNESS FROM A SYSTEM'S PERSPECTIVE

We introduce Michael A. Jackson's theory of requirements engineering [11] to discuss the focus, goals, and assumptions behind fairness framework from a systems' perspective. This theory argues that computer scientists and engineers "are concerned both with the world, in which the *machine* serves a useful purpose, and with the machine itself [...] The purpose of the machine [...] is located in the world in which the machine is to be installed and used." In other words, our objective is to design systems that fulfill requirements *in the world* where the system is to be deployed.

More precisely, a portion of the world becomes the *environment*, or the *application domain*,<sup>\*</sup> of a *machine*. In this world there are *phenomena* in the application domain (e.g., events, behavior of people, activities) and in the machine (e.g., data, algorithms, state). The machine's inputs and outputs, i.e., the things the machine can sense or affect, are shared phenomena: they are observable both to the application domain and the machine.

We introduce machines into existing application domains to effect change in these domains. To achieve the desired change, we need descriptions of the phenomena before the machine is introduced, known as *domain assumptions* K; and statements about the desired situation once the machine is introduced to the domain, known as *requirements* R. A *specification* S is a set of requirements providing enough information for engineers to implement the machine. S typically describes phenomena shared between the machine and application domain. A *program* P derived from the specification is a description of a machine. If implemented correctly, programs satisfy the specification. If the specification is derived effects in the application domain, i.e., they fulfill the requirements.

Jackson's explicit treatment of the application domain and its interaction with the machine helps us to project known problems with algorithms to a systems view. It enables us to distinguish problems due to badly derived requirements (description of the problem) from those due to badly derived specifications (description of the solution) and those due to badly implemented programs (how solutions are implemented) [12].

# 2.1 The focus on algorithms is insufficient for addressing inequitable outcomes of systems

In the light of Jackson's theory, we claim that *the focus on algorithms leaves out the systems' view.* Algorithms are often a part of a larger technical system, which is deployed in an environment. Fairness proposals rarely evaluate the systems' environmental conditions, thereby leaving out the possibility that, *even* when a fairness metric is satisfied by the algorithm, the system (environment plus the machine) could still be unfair or have other negative side effects. Such focus on the specification of the machine also promotes the idea "that the difficulty in addressing [unfairness] lies in devising a solution" [11] and not necessarily in rethinking the world.

Most fairness frameworks focus on describing *S* independent of *K* and *R* and then guarantee that states of the machine, and its input and outputs, have certain properties. However, a machine that has fair inputs and outputs and fulfills a specification  $S_{fair}$  does not guarantee the fulfillment of requirements of fairness in the application domain. This would require evaluating *K*, establishing what phenomena the machine is expected to change in the application domain, and articulating requirements  $R_{fair}$  with fairness as a goal. Only then could we evaluate whether an  $S_{fair}$  fulfills  $R_{fair}$ .

Focusing on  $S_{fair}$  has a number of repercussions. First, it does not reflect how harms manifest themselves in the environment. Without an understanding of K and  $R_{fair}$ , a specification  $S_{fair}$  may simply not lead to a fair outcome. Imagine a hypothetical "fair predictive policing algorithm" that can fairly distribute police officers to different neighborhoods. If the algorithm does not consider that the policing institution is already configured to control minorities [13] and that interactions with police pose greater risk of harms for minorities, a "fair allocation" can still disparately impact those minorities [6].

Second, focusing on achieving fairness for users might *leave out the impact of the system on phenomena in the application domain that is not shared with the machine.* Unjust outcomes could arise due to the optimization of certain behavior in the application domain, and not because *S* was unfair. For instance, self-regulated housing markets such as Airbnb [14] may not actively discriminate against their users, but reports have shown that they can disrupt neighborhoods by changing rent dynamics and neighborhood composition [15].

Third, the focus on algorithms abstracts away potential harms of phenomena in the machine. Much of the fairness literature focuses on ways in which algorithms can be biased or on harms caused by decision-making algorithms. This overlooks that, when the system hosting the algorithms optimizes its operation, it may gather more inputs and outputs than those of the algorithm. Therefore, focusing on the algorithm may miss effects on the world that can go beyond those generated by the outputs of the algorithm actions.

When phenomena in the machine domain are subject to optimization, unfairness can arise from optimization programs  $P_{opt}$ fulfilling the specification  $S_{fair}$  but not the requirements  $R_{fair}$  in the application domain. For instance, prediction techniques to optimize targeted advertising can create discriminatory effects [16]; and exploration strategies to optimize  $S_{fair}$  may gather inputs from the application domain that put some users unfairly at risk [17].

<sup>\*</sup>Application domain does not refer to a class of applications, like health or banking domain, but to actors and things in the environment where the machine is introduced.

# 2.2 Discriminatory effects are not the only concern for building just systems

In Jackson's terms, considering only discriminatory effects constrains the requirements  $R_{fair}$  to a particular class of harms. This approach risks *missing other harms caused by the system* when evaluating the performance of the specification  $S_{fair}$  in the environment.

We assume the introduction of a machine in an environment aims to improve specific phenomena. The fact that this machine follows a specification  $S_{fair}$  that does not discriminate according to  $R_{fair}$  does not guarantee that this machine will not induce other harms to the environment in any applications domain. Take as examples unsafe housing or bad working conditions. If we use a machine to distribute these resources more efficiently, even if it does so fairly, *all* users are harmed. Bad housing conditions and lack of labor protection are problems in and of themselves so users will be badly served regardless of fairness conditions. When the system is by design unjust, or when the phenomena are structurally unjust or harmful, claims on  $S_{fair}$  are meaningless. In Jackson's terms: the requirements *R* are incomplete with respect to just outcomes.

Sometimes, however, injustices are tightly tied to the machine: S is not the solution but the source of problems. The way requirements R are optimized might lead to externalities for (a subpopulation of) users. When a specification S optimizes an asocial outcome, such as excessive user engagement in social networks [18], it can expose users to harms like addiction. Solving fairness in this system will not resolve the underlying problem: the system is harmful.

Striving to fulfill the requirements  $R_{fair}$  itself might bring new harms to the application domain. Consider a fairness solution that alleviates distributional shift based on increasing diversity in the training set. If its specification  $S_{fair}$  requires collecting data from more individuals or collects new attributes to implement the fairness measure, it will exacerbate privacy issues. These issues might result in many other harms in the application domain.

This ontology assumes that *S* is built to "solve problems" and "improve phenomena" in the world. Thus, it does not provide the conceptual tools to address adverse situations. This positive valence hinders the consideration of cases in which a machine amplifies existing injustices or introduce new ones. Neither this ontology nor fairness frameworks account for power imbalances or economic incentives, and how they impact how machine requirements are considered and prioritized. We consider these matters in the next section by augmenting our systems view to consider socio-economic aspects of the application domain.

#### **3 FAIRNESS, INCENTIVES, AND POWER**

In this section, we extend our analysis to problems related to the political economy of systems. To model harms, we borrow the term *negative externalities* from the economics literature. A system causes negative externalities when its consumption, production, and investment decisions cause significant repercussions to users, non-users, or the environment [19]. The introduction of a machine might cause externalities in the application domain, independent of the completeness or correctness of its requirements and specification. For example, the heavy use of traffic-beating apps such as Waze can worsen congestion for *all* drivers in the application *S* against

the requirements R is not enough. To build just systems one must consider externalities of the machine in the application domain.

Congruent with models in fair optimization and economics, to express externalities and incentives we introduce two utility functions that capture the machine's impact on the application domain: the *service provider's utility*, which measures how much value the provider extracts from introducing the machine, and the *social utility*, which measures the machine's utility for the environment and people. We define two versions of social utility: one with a "god's view" of the application domain and one from the specification's perspective. These utility functions enable us to capture *injustices* due to the introduction of the machine into the environment.

#### 3.1 Idealized Fair-by-Design Service Provider

We first consider an idealized *fair-by-design service provider* that is willing to address the externalities of the machine. That is, this provider aims to maximize both their own utility and the social utility. Using this setting, we show the ways in which fairness models fail to address a broad class of systems' externalities.

We consider that a system is parametrized by a vector of internal parameters  $\theta \in \Theta$  for some convex set of possible parameters  $\Theta$ . Let *P* be a *population*, a set of individuals or other environmental entities that might be affected by the system. Let  $U(\theta) : \Theta \to \mathbb{R}$ be the utility function of the provider when they use parameters  $\theta$ . Let  $B(\theta) : \Theta \to \mathbb{R}$  denote a hypothetical social-utility function, or *benefit*, defined in the requirements *R*. Let  $\hat{B}(\theta)$  be the social utility in the provider's specification. The provider optimizes its operation by solving the following multi-objective optimization problem:

$$\max_{\theta \in \Theta} \{ U(\theta), \hat{B}(\theta) \}$$
(1)

This problem is considered in fair learning literature in its scalarized form or constraint-form [21–25], with the social utility modeling a notion of fairness. We assume an ideal situation in which the chosen parameter vector  $\theta^*$  is a Pareto-optimal solution [26], that is, it cannot improve any of the objectives without hurting at least one of them. Pareto-optimality, however, is not sufficient to guarantee fairness or equity as different trade-offs between the objectives are possible [27]. We assume that out of the possible Pareto solutions, the provider chooses one that maximizes social utility:  $\theta^* \triangleq \arg \max_{\theta \in \Theta^*} \hat{B}(\theta)$ , where  $\Theta^*$  is the set of all Pareto-optimal solutions.

3.1.1 Limitations in the Face of Externalities. Consider  $\theta^{\circ}$ , the system obtained when using the "god's view" values of social utility:  $\theta^{\circ} \triangleq \arg \max_{\theta \in \Theta^{\circ}} B(\theta)$ , with  $\Theta^{\circ}$  being the set of the corresponding Pareto-optimal solutions when using *B* in Eq. 1.

Let  $\Delta B \triangleq B(\theta^*) - B(\theta^\circ)$ . We say that a system with parameters  $\theta$  induces *externalities* on its environment when the social utility of  $\theta$  is not equal to that of the system  $\theta^\circ: \Delta B \neq 0$ .

This definition is analogous to the neoclassical economic interpretation of externalities: due to an "inefficiency," a partial optimum  $(\theta^*)$  is achieved, that is different from the optimum if no inefficiencies were present  $(\theta^\circ)$  [28]. One can also parallel the divergence between social-utility values in  $\Delta B$  to the Pigouvian "divergence of social and private costs" [29]. Because we are analyzing harms, we focus on negative externalities:  $\Delta B < 0$ .



Figure 1: Pareto frontiers with real (U, B), and provider's version  $(U, \hat{B})$ . Misspecification of *B* results in externalities  $\Delta B$ , difference in values of the benefit function between provider's system  $B(\theta^*)$  and  $B(\theta^\circ)$ .

Clearly,  $\Delta B = 0$  when  $B = \hat{B}$ . If *B* does not precisely match  $\hat{B}$ , however, there will likely be externalities present, as illustrated in Fig. 1. In general, it is not known how exactly  $\Delta B$  is impacted by deviations of  $\hat{B}$  from *B*. In Appendix A, for a class of strictly concave utility functions we show that the *sensitivity* of  $\Delta B$  to infinitisemal error of  $\hat{B}$  is approximately quadratic in the magnitude of the error. Intuitively,  $\Delta B$  grows quadratically fast as  $\hat{B}$  diverges.

Incomplete Information of Social Utilities. We first study the case in which the provider's view of the social utility,  $\hat{B}$ , does not fully reflect the requirements and the context of the application domain.

Social benefit is often modeled as a function incorporating individual models of social utility. A common assumption in neoclassical economics is that social utility is the sum of individual utilities [30]:  $\hat{B}(\theta) = \sum_{i \in P} \hat{B}_i(\theta)$ , where  $\hat{B}_i : \Theta \to \mathbb{R}$  is a utility of individual  $i \in P$ . In practice, however, the fair-by-design provider has incomplete information: they are aware of the users, yet lack the full knowledge of their needs (similar to "imperfect knowledge" in economics and game theory [31]). That is, they could misspecify  $\hat{B}_i(\theta) \neq B_i(\theta)$  for some user  $i \in P$ , inadvertently ignoring the needs and well-being of this individual. Consider a hypothetical fair-by-design predictive policing [32] application that assumes that  $\hat{B}$  is maximized when there is equality of false positives in patrol dispatches across regions. The resulting dispatching rates might be fair for this definition, but the overall social benefit B might be unchanged as minorities could still be over-policed in terms of the number of dispatches.

Omitting Impact on Non-Users and Environmental Impact. Another case is when the fair-by-design provider has structural lack of knowledge of the application domain, e.g., knows the utility of their users, but not the utility of anything else (i.e., of phenomena not shared with the machine). Thus, a fair solution for the users could harm non-users as their benefits were never specified in the model  $\hat{B}(\theta)$ . This is exemplified in the case of self-regulating housing markets, which could hypothetically be optimized for fairness in acceptance rates for guests [33], but disrupt neighborhoods impacting housing conditions, especially for people with low incomes.

*3.1.2 Limitations Under Complete Knowledge.* Even in the presence of complete information regarding *U* and *B*, there can exist externalities that the provider cannot mitigate.

Bogdan Kulynych, Rebekah Overdorf, Carmela Troncoso, and Seda Gürses

First, the provider can maximize *B* and yet cause externalities if the system's goal *itself* is harmful [34]. For example, a facial recognition surveillance system might be completely fair with respect to skin color, but it still causes risks and harms to the population as a whole in terms of privacy loss.

Second, the fair-by-design model has as premise that there exists a solution  $\theta^{\circ}$  that, if not maximizes, then at least satisfies minimum standards of everyone's benefit. However, this may not be the case. Recall the examples about unsafe housing or bad jobs in Sec. 2.2. In those cases, the fair outcome might still be harmful for all users.

#### 3.2 Limitations on Ideals

Mitigating externalities becomes a greater issue when incentives, capacities, and power structures are not aligned. One way this manifests itself is in the modeling of social benefit as a function. Indeed, a function cannot encode all the nuance regarding human needs. Moreover, neither the provider nor the users' utility functions model the political context or the power asymmetries in the environment of the machine [35]. These can heavily influence or skew what we assume to be the ground truth *B*. Thus, power and politics may come to render the system specification unfair even when it was designed considering the perfect benefit function for users.

Furthermore, the fairness-by-design approach inherently assumes that the provider always has enough resources to implement the fair solution that maximizes social utility. However, it is unreasonable to believe that such an assumption will hold in practice. For instance, even if Waze was cognizant of the needs of all individuals and all the peculiarities in their streets, it is unlikely that they could afford an operations in which all of those constraints are taken into account or be capable of mitigating the impact of externalities due to the interaction of multiple traffic beating applications.

Finally, not all service providers have the incentives to implement fair-by-design solutions. The provider *might not be incentivized* to care about social utility or could benefit from externalizing certain costs. In such cases, the magnitude of the externalities  $\Delta B$  is likely to be more pronounced than in the case of incomplete knowledge.

# 4 PROTECTIVE OPTIMIZATION TECHNOLOGIES

The systems view on algorithms in previous sections enables us to systematically explore a problem space that had not been formalized before. A major source of problems is the use of optimization techniques that help to capture and manipulate phenomena in the application domains for the *extraction of value*. This practice causes intentional or unforeseen changes to the environment which result in (often neglected) harms to the environment.

We showed that existing fairness frameworks produce solutions that can only address a subset of these harms, at the discretion of a service provider. These frameworks have limited capability to mitigate harms arising from inequities inherent to the application domain and from harmful impacts introduced by the machine. By focusing solely on actions that can be taken by service providers, who have opposite incentives, fairness frameworks narrow both politics and contestation to the re-design of the algorithm which may not always be the site of either the problem or the solution.



Figure 2: Uber: a complex optimization system. Inputs to Uber optimization (black), effects on the environment (blue), and alternatives to deploy POTs (red).

In this section, we discuss means to address these issues. Although these means could be of socio-legal nature, we focus on technological approaches that we call *Protective Optimization Technologies (POTs)*. Their goal is to shape the application domain in order to reconfigure the phenomena shared with the machine to address harms in the environment. POTs are designed to be deployed by actors affected by the optimization system. As these actors directly experience the externalities, a) they have intimate knowledge of the system's negative effects, b) they are in position to have a better view of their social utility than a system provider can model—because it is their own utility [36]. Lastly, c) POTs do not rely on the incentives of the provider.

POTs seek to equip individuals and collectives with tools to counter or contest the externalities of the system. Their goal is not to maximize benefits for both users and service providers, nor to find the best strategy to enable optimization systems to extract value at minimum damage. POTs are intended to *eliminate* the harms induced by the optimization system, or at least *mitigate* them. In other cases, POTs may *shift* the harms to another part of the problem space where the harm can be less damaging for the environment or can be easier to deal with. Finally, when service providers react to reduce the effectiveness of POTs, this very action *exposes* the service providers' need to maintain the power relationship and their capability to manipulate the environment to their own benefit.

*POTs and Optimization Systems.* To extract value through optimization, service providers obtain inputs from their environments that help them make decisions. We consider three kinds of inputs: i) inputs that users generate when interacting with the system, ii) inputs about individuals and environments received from third parties [37], and iii) inputs from regulations and markets that define the political and economic context in which the system operates. Note that third parties can be public or private, and the data they provide can be gathered online or offline. We use Uber ride-hailing service [38] as an example to showcase this complexity (Fig. 2).

In order to maximize its profit, Uber optimizes the prices offered to riders and the wages offered to drivers. Uber uses the following inputs (black arrows). First, it uses the direct inputs it receives from both riders and drivers. Second, it uses data from other sources such as online service providers, e.g., Google for maps and data that they might collect using cookies or pixels on other sites their users visit [39]. Uber also receives offline data from parties like municipalities interested in promoting the use of Uber to reduce costs of public transport [40]. Lastly, Uber uses inputs from the market and regulators to evaluate the economic context in order to adjust wages and ride prices.

Uber ultimately uses these inputs and all the political and economic context in a combination of managerial and mathematical optimization to deliver outcomes to the environment: match riders to drivers and set ride prices. Reports and studies demonstrate that these outcomes cause externalities: Uber's activity increases congestion and misuse of bike lanes [41], increases pollution, and decreases public support for transit [42, 43].

In the previous sections' vocabulary, Uber and its optimization techniques are the machine; and the application domain comprises Uber drivers and riders, non-users, the online and offline environments, and the market and regulatory frameworks in which Uber operates. Roughly speaking, Uber's requirements R are to match drivers to riders with associated dynamic pricing. The specification S defines what Uber applications should do to fulfill R.

Uber is known to have unfairness problems. For example, Rosenblat et al. show that customer-based reviews are biased against minority drivers. As getting blocked from the system depends on these reviews, even if the Uber's algorithms do not discriminate drivers on their attributes per se, the rules of the system result in disparate impact on blocked drivers [44]. Further, even when Uber algorithms are fair, e.g., they reward all drivers the same irrespective of their protected attributes, the optimization processes underlying Uber's operation result in unjust outcomes: low wages [45]. The former is an externality stemming from structural biases in the application domain (thus a hypothetical  $S_{fair}$  does not result on  $R_{fair}$ ); whereas the latter is a problem of incentives misalignment.

In the Uber scenario, POTs can be deployed by users and nonusers with the goal of changing the phenomena captured by Uber. These can come in three forms (red lines in Fig. 2): by changing the inputs of the users to the system (e.g. the surge-induction POT [46] as we describe shortly in Sec. 4.1), by changing the online or offline signals gathered by Uber (e.g., mayors changing the city urban planning), or by affecting the market (e.g., by changing regulations or mandating salary increases [44]).

# 4.1 Examples of POTs

Our vision for POTs systematizes the use of technologies as tools to both explore the effects that algorithms and optimization systems have on our society, and the design of countermeasures to contest their negative effects in the absence of regulatory or other accountability measures.

We now revisit recent academic technologies, artistic interventions, and deployed tools, that can be reframed as POTs (Table 1 provides a summary). POTs formalize such technologies and interventions, enabling the systematic study and design of such solutions. We illustrate this in Sec. 5, where we design two POTs from scratch.

These technologies have different origins. First, we observe that there are technologies proposed in the academia that can be repurposed as POTs. For instance, in the field of computer security, Table 1: Repurposing technologies as POTs. We detail the origin of the technology; the externality it addresses, the optimization system causing the externality and the desired outcome of the POT; deployment requirements (individual or cooperative action); and the underlying design techniques (counter-optimization or heuristics to decide how to shape the environment).

Origin	Optimization System	Externality	РОТ	Desired Outcome	Deployment	Technique
Academic	Face Recognition	Privacy, discrimination	Wear printed eyeglasses [47]	Evade face detection	Individual	Optimization
Academic	Copyright Infringement Detection	Fair use takedowns	Adversarial examples [48]	Avoid a fair use takedown	Individual	Optimization
Academic	Psychometric Profiling	Privacy, manipulations	Text style transfer [49, 50]	Prevent from attribute inference	Individual	Optimization
Academic	YouTube Recommendations	Manipulation	Poisoning [51]	Breaking out of content bubbles	Individual	Optimization
Academic	Waze Routing	Local traffic congestion	Sybil devices simulate traffic [52]	Prevent routing into towns	Individual	Heuristic
Academic	GRE Scorer	Biased grading system	Generate essay to pass GRE [53]	Higher test score	Individual	Heuristic
Deployed	Ad Network	Privacy, manipulations	Click on all ads [54]	Ad Network destroyed	Collective	Heuristic
Deployed	Uber Pricing System	Low wages	Shut off app, turn it back on [46]	Induce surge	Collective	Heuristic
Deployed	Instacart Pricing	Low wages	Tip 22¢ in app, cash at door [55]	Fair pay for jobs	Collective	Heuristic
Deployed	Automated Hiring	Bias, discrimination	Edit resume [56]	Flip automated hiring decision	Individual	Optimization
Deployed	Pokemon Go Resource Spawn	Unfairness	Edit Open Street Maps [57]	Encourage resources to spawn	Individual	Heuristic
Deployed	FitBit for Insurance Premium	Privacy, surveillance	Spoof device location [58]	Get insurance benefits	Individual	Heuristic
Deployed	Pharma Optimizing Patents	End of humanity	Find potential drugs using ML [59]	Get drugs in the public domain	Individual	Mixed
Deployed	Insurance Coverage Optimization	High costs of treatment	Doctors changing claim codes [60]	Get higher reimbursements	Individual	Heuristic
Artistic	Face Recognition	Privacy, surveillance	Scarf that is classified as a face [61]	Evade face detection	Individual	Optimization
Artistic	Face Recognition	Privacy, surveillance	Camouflage to cover features [62]	Evade face detection	Individual	Heuristic
Artistic	Autonomous Cars	Exploration risks	Ground markings [63]	Trap autonomous cars	Individual	Heuristic

research that aims to protect against attackers gaming the YouTube algorithm [51] can be repurposed by users to fight against filter bubbles; in the field of adversarial machine learning, tools developed to evade copyright detection [48], originally developed to strengthen DRM, can be reframed as a way to prevent fair-use takedowns [64]. Second, we draw from works produced by artists looking into the impact of technology on society. For instance, counter-surveillance fashion that tricks facial recognition technologies [62] can be repurposed to evade discriminatory facial recognition algorithms or discriminatory uses of facial recognition. Finally, we look to deployed technologies that are already countering optimization, either intentionally or as a side effect. For instance, Jobscan [56] assists job applicants in getting past the automated sorting implemented by large companies and job-posting sites. This tool could be repurposed to reduce the gender or racial bias reported for these tools [65].

These examples can also be categorized based on the means used to design the POT: those based relying on adversarial machine learning [47, 49–51, 56, 61], and those that use heuristics to exploit the target optimization process. For instance, as a response to low wages, Uber drivers have developed heuristics for inducing surges [46].

Finally, some technologies can be deployed individually and others require collective action. AdNauseam [54], for example, reduces the utility of ad networks by clicking all ads served to the user, flooding the network with false information. Without a critical mass of users, however, AdNauseam's effect would not be observable.

# 5 DESIGNING PROTECTIVE OPTIMIZATION TECHNOLOGIES: CASE STUDIES

In this section, we show how, given an optimization system and a negative externality, one can design new POTs from scratch.

Like many of the technologies and interventions in Table 1, we make use of optimization techniques to design the operation of POTs. Starting from the model introduced in Sec. 3, we model the optimization system's objective as a function  $U(\mathbf{x}; \theta) : \mathcal{X}^m \times \Theta \rightarrow \mathbb{R}$ , where  $\mathcal{X}^m$  denotes the space of environment inputs to

the system (i.e., phenomena that are sensed by the machine), and the vector  $\mathbf{x} \in \mathcal{X}^m$  is a set of inputs coming to the system. Each concrete input  $x_i$  can come from a different source (users, non-users, or other actors in the environment). For simplicity, we model the time dimension through discrete time steps, and we assume that there are two possible time steps: 0 and 1.

To maximize its gain, the optimization system strives to solve a mathematical optimization problem  $\theta_t^* = \arg \max_{\theta} U(\mathbf{x}_t; \theta)$  for a set of inputs  $\mathbf{x}_t$  at each point in time t. Given this system, the goal of a POT is finding actionable—feasible and inexpensive—modifications to the inputs of the optimization system so as to maximize social utility that we denote as  $B_{pot}(\theta)$ . Note that this definition of social utility needs not to correspond to the social utility considered by the optimization system  $(\hat{B}(\theta) \text{ in Section 3})$ . We consider that each input has an associated modification cost  $C(\mathbf{x}_i \to \mathbf{x}'_i) : \mathcal{X} \times \mathcal{X} \to \mathbb{R}^+$ , that represents how hard it is to modify it. The cost of changing a set of inputs  $C(\mathbf{x} \to \mathbf{x}')$  can be any function of the individual costs.

In this model, we define the POT design as a bi-level multiobjective optimization problem:

$$\min_{\mathbf{x}'} \{ C(\mathbf{x} \to \mathbf{x}'), -B_{pot}(\theta_{t+1}^*) \}$$
  
s.t.  $\theta_{t+1}^* = \arg\max_{\theta_{t+1}} U(\mathbf{x}', \theta_{t+1})$  (2)

where **x** is the vector of inputs if no intervention happened, and **x'** are possible vectors of modified inputs.  $\theta_{t+1}$  represents the system's state in the next step, after the POT has been deployed through the modified inputs **x'**.

We now instantiate this problem for two different use cases, one where the POT can be deployed by an individual, and one where effective deployment requires a collective.

#### 5.1 Thwarting Traffic from Routing Apps

In our first case study, we look at Waze, a crowdsourced mobile routing application that optimizes the routes of its users based on information about traffic, road closures, etc. [66]. Waze causes negative externalities for residents of towns and neighbourhoods that are adjacent to busy routes. For example, take the town of Leonia, New Jersey, USA, which lies just outside one of the bridges into New York City. As Waze rose in popularity and directed an increasing amount of users through the town when the highway was busy, the town became crowded during rush hours. To prevent Waze traffic, the town was briefly closed off to non-local traffic, which was determined illegal [67]. In this section we propose a solution for discouraging Waze from selecting routes through the town while minimizing the impact on its inhabitants.

*Problem Setup.* We set up this problem as a planning problem in which the town's traffic network is modeled as a weighted directed graph, and the goal is to increase the cost of paths between the highway ramps. We define Waze's utility U as the capability to provide fastest routes for its users. Routing through town can increase this utility when it takes less time than traveling via the highway. The POT is designed to increase the minimum time through town so that Waze users are not routed through it. We define the social utility  $B_{pot}$  as a binary variable that takes value 1 when no vehicle is routed through the town (i.e., the cost in time of traversing the town is greater than traversing the highway) and 0 otherwise.

Let  $\mathcal{G} = (\mathbb{V}, \mathbb{E}, t)$  be a weighted directed graph representing the traffic network within the town. Each edge  $(x, y) \in \mathbb{E} \subset \mathbb{V} \times \mathbb{V}$  represents a road segment, and each vertex represents a junction of segments. The edges have associated time to pass the segment given by the function t(x, y). We define the *time cost* of traversing a path in the graph as the total time to pass its edges:  $t(\mathbf{e}) \triangleq \sum_{(x,y) \in \mathbf{e}} t(x, y)$ .

Let  $a, b \in \mathbb{V}$  be the source and sink vertices that represent the entry to the town from the highway (a), and the exit to return to the highway (b). Let the time to travel from point a to b via the highway be  $t^*$ . While we do not know the routing algorithm used by Waze, we assume that Waze will send users through the town when the path through town is quicker than the highway. That is, if there is a path **e** from a to b in  $\mathcal{G}$  with cost  $t(\mathbf{e}) < t^*$ , Waze routes users through the town.

5.1.1 Avoiding Routing via Planning. We aim to transform the graph G into a graph G' such that the time cost of any path e' from *a* to *b* in G' is  $t(e') \ge t^*$ . We focus on what the town can control: time to traverse a road segment. We express these changes as the increase in time that it takes to traverse a segment,  $\Delta t(x, y)$ . We abstract the exact method for adding time to the segment, which could be changes to speed limits, traffic lights, stop signs, speed bumps etc. We construct G' by modifying the time values in the original graph  $G: t'(x, y) = t(x, y) + \Delta t(x, y)$ , where t' is a function representing the edge time in G'.

We acknowledge that some roads are more costly to change than others by associating a modification cost to every road. In practice, this cost will be specified by the town. We use *length* of a road segment as such a cost, capturing that changing longer roads will likely have more impact on the town. Let  $\mu(x, y) \in$ {0, 1} be binary *interdiction variables* that represent whether we are modifying the edge (x, y) in graph G'. To express the cost *C* of modifying a graph G into G' we use the following modification-cost function:  $\sum_{(x,y) \in \mathbb{E}} c(x, y) \cdot \mu(x, y)$ , where  $c : \mathbb{E} \to \mathbb{R}^+$  represents the cost of modifying the edge (x, y).

We now formalize the POT through posing the multi-objective optimization problem in Eq. 2 in a constraint form: minimize *C* 

subject to the  $B_{pot}$  constraint, given that Waze is maximizing U:

$$\min_{\mu(\cdot) \in \{0,1\}} \sum_{(x,y) \in \mathbb{E}} c(x,y) \cdot \mu(x,y)$$
  
s.t.  $t(\mathbf{e}) \ge t^*$  for any path  $\mathbf{e}$  from  $a$  to  $b$  (3)

This is equivalent to the problem known as *shortest-path network interdiction* [68–71]. In the form equivalent to ours it can be solved as a mixed-integer linear program (MILP) [70, 71]. We refer to Appendix B for the exact formulation of the MILP in our context. We use ortools [72] for specifying this MILP in Python, and the CBC open-source solver [73] for solving.

*5.1.2 Empirical Evaluation.* We apply this POT to three towns, all of which reported issues with Waze: Leonia, NJ, USA; Lieusaint, France; and Fremont, CA, USA. We retrieve the map data for each via the Open Street Maps API [74].

To assign t(x, y) to each segment (x, y), we estimate the time it takes to traverse the segment using its length and type (e.g., we consider residential streets to have a speed of 25mph) and add extra time for traversing the intersection. We infer the intersection time using the travel time from Google Maps for accuracy. For each city, we run the solver for different  $t^*$  values, corresponding to different travel times on the highway. Larger values of  $t^*$  require more changes to roads in the town.



Figure 3: Solutions for Leonia, NJ (left) and Lieusaint, France (right), when the time of road segments is allowed to be increased by 75%. Town streets are marked in blue, highways in orange, and surroundings in grey. The red dots signify *a* and *b*, the closest points to the highway in the town. The roads marked by thick, black lines are the optimal set of segments in which the time should be increased.

For simplicity of implementation, we choose two points at the edge of the town (red dots in Fig. 3) and not the sink/source points, which would be on the highway, for the solver. We then add the time that it takes to travel from the actual source point on highway to the first point and from the second point back to the sink point in order the calculate the results. For Leonia, for instance, we approximated this as 30 seconds on each side of the town.

We consider scenarios in which the town is able to increase the time that it takes to traverse each road segment by 25%, 50%, and 75%. For each town we found the value of  $t^*$  for which Waze begins sending cars through the town and the value of  $t^*$  in which no further changes to the road can prevent Waze from sending its users through the town.

The graph for Leonia, NJ contains 251 vertices and 661 edges. Given the parameters we choose, without any changes to the town, Waze begins to send its users through Leonia when  $t^* = 4.0$ . That is, it normally takes 4 minutes to travel through Leonia, so when it takes longer than 4 minutes to traverse the highway, Waze routes users through town. This corresponds to traffic traveling at about 31mph on the highway. If we limit the amount of time that can be added to traverse a road segment to 25% of the original time, we can prevent cars from being routed through Leonia until the average highway speed has reached 26mph ( $t^* = 5.1$ ). That is, we can find a solution that prevents routing through town as long as the highway speed is greater than 26mph. For 75% time increase, we can change the town such that traffic will not be routed through Leonia until the average speed on the highway has fallen to 19mph. This solution for Leonia is shown in Figure 3.

Lieusaint is larger than Leonia, with 645 vertices and 1,248 edges. Given the parameters we choose, Waze routes its users through Lieusaint when  $t^* = 7.0$ , which corresponds to the speed on the highway dropping to 16mph. Allowing the road segments to be lowered by 75%, we can prevent traffic from being routed through Lieusaint until the highway speed has dropped below 10mph ( $t^* = 12.0$ ) (See Lieusaint in Figure 3). We report the solution for Fremont, CA, a significantly larger town, in Appendix B.

Finally, we measure the cost of implementing these solutions (Fig. 4). For each town, we consider the impact to the town to be how much longer, on average, it takes to travel between any two points in town. We compute the shortest path between every pair of points in the town and average these times before and after the POT solution. We then compute the percentage increase from the initial average time to the post-POT average time. The higher impact solutions are those which will tolerate a lower highway speed. That is, they will prevent cars from being routed through the town at lower highway speeds. We see that even though allowing road segments to take 75% longer to traverse can prevent cars from entering the town at a lower highway speed, the impact to the town is much higher. The inhabitants of Lieusaint also suffer more from the changes than the residents of the much smaller Leonia.

5.1.3 POT Impact and Limitations. The intervention we propose in this section focuses on alleviating the problems routing applications cause in *one* town. The POT would support this town's government to decide which changes should be implemented in the city layout and traffic rules so that the cost of traversing the city becomes undesirable for passthrough drivers. While this intervention indeed mitigates the effect of external traffic on the target town, it is likely that then vehicles are routed elsewhere, i.e., the POT *shifts* the harms from this particular town to other regions. Moreover, we acknowledge that our POT can only help in cases when vehicles can still circulate on the highway. The moment the congestion forces vehicles to halt, the town becomes the better option regardless of any modification on its road network.

#### 5.2 Reducing False Negatives in Credit Scoring

In this case study, we explore solutions for countering harmful effects of a machine-learning *credit scoring* system: A system that a bank uses to decide how to price a loan application according to the predicted risk. The underlying algorithms that support such



Figure 4: Effect of Changes on In-Town Travel.

decisions are designed to maximize banks' profits and minimize their risks. These algorithms can be discriminatory [75], or cause feedback loops for populations disadvantaged by the financial system [76]. These harms are often caused by inputs that represent unjust realities which are propagated to the model's decisions.

*Problem Setup.* We model the credit scoring system as a classifier  $f_{\theta}(x)$  that takes as input the information about the loan applicant and the loan details, x, and outputs a confidence score for whether the applicant would repay the loan or not. This function optimizes the bank's utility that we model here as the negative empirical loss over historical data:  $\theta^* \triangleq \arg \min_{\theta \in \Theta} \sum_{(x,y) \in X} L(x,y;\theta)$ , where  $\theta$  is the parameters of the classifier, L is the loss function, and X is the bank's dataset of historical loans (x) and their repayment or default outcomes (y). We assume that the classifier is retrained as new data points arrive.

In our case study, we model the harms of the system as the rate of false negatives (wrong default predictions) for economically disadvantaged populations. In this POT, we counter this problem using adversarial machine learning techniques deployed by a collective of individuals  $G_{pot}$  with the means to take and repay loans, as explained shortly. This POT aims to increase the social utility defined as the loss for a *target group*  $G: B_{pot}(\theta) \triangleq \sum_{(x,y) \in G} L(x, y; \theta)$ , where G is the disadvantaged subset of applicants (in this case study, we define disadvantaged as having little funds in the bank account) who were wrongfully denied a loan. This POT can be thought as promoting equality of false-negative rates between the target group and everyone else [77, 78], with the difference that we do not limit our view of externalities to a commonly protected subgroup, i.e., economical disadvantage is not commonly considered as protected.

FAT\* '20, January 27-30, 2020, Barcelona, Spain

5.2.1 Reducing False Negatives with Adversarial Machine Learning. We first identify what inputs (x in the abstract model) can be modified by the collective  $G_{pot}$  deploying the POT. First, the deployers can only add new inputs to the dataset by taking and repaying loans. Second, the demographic attributes of these added loan applications have to be similar to those of individuals in *G*. Thus, the POT must inform the collective  $G_{pot}$  about who and for which loans they should apply for and repay, in such a way that they reduce the false-negative rate on the target group after retraining. This POT is idealistic in that it assumes that this collective will include applicants of diverse backgrounds to be able to provide different inputs to the classifier. However, in the absence of other means of feedback, communication, and accountability, it represents the only means to influence an unjust system. It is also consistent with the existing practices people resort to in this setting [79, 80].

Finding which inputs to inject into a training dataset to modify its outputs is known as *poisoning* in adversarial machine learning [81]. Typically, poisoning attacks aim to increase the average error of the classifier or increase the error on specific inputs [82]. We note that our use of poisoning is different. First, we poison to *decrease the error* for a given target group. Second, our use of poisoning is not adversarial. On the contrary, we use it to protect users against harmful effects of the model.

With this in mind, we design the POT using the following bi-level optimization problem:

$$\min J(\theta^*) = \sum_{(x,y)\in G} L(x,y;\theta^*) + \lambda R(\theta^*)$$
  
s.t.  $\theta^* = \arg \min_{\theta \in \Theta} \sum_{(x,y)\in X \cup X_{pot}} L(x,y;\theta)$   
 $f_{\theta'}(x) = \text{`accept' for all } (x,y) \in X_{pot}$   
 $X_{pot} \subset X_{pool}, |X_{pot}| \le n$ 

where  $\theta'$  is the current parameters of the classifier,  $X_{pot}$  is the *set* of poisoned applications, *n* is the maximum number of poisoned applications, and  $X_{pool}$  is a *set of feasible loan applications*. That is, we minimize the classifier's loss for the target group, where the classifier  $\theta^*$  is trained using the poisoned examples  $X_{pot}$ . In our evaluation below, we additionally make use of a regularizer *R* to minimize the effect of poisoning on any other applications:  $R(\theta) \triangleq -\sum_{x,y \in X} L(x, f(x); \theta)$ .

This formulation makes two assumptions. First, we assume that when designing the POT we have access to the training dataset of the provider X, and we can obtain retrained loan-approval models. We consider this assumption reasonable as poisoning attacks tend to transfer even if the dataset and model do not match the real ones, especially when the models have low complexity [83]. Second, the added loan applications must be feasible: there has to exist a person in the  $G_{pot}$  with demographics required for this loan application.

We solve this problem by scoring each example in  $X_{pool}$  according to the value of our optimization objective  $J(\theta^*)$ , retraining  $\theta^*$  for each example, and then employing a greedy algorithm to assemble  $X_{pot}$  that satisfies the constraints. We refer to the Appendix C for the details of the algorithm.

5.2.2 *Empirical Evaluation.* We create a simulated loan approval system using the German credit risk dataset from the UCI Machine



Figure 5: Effects of the POT on the error rates of the classifier for the target group. Top: decrease in the number of false negatives (wrongly denied loans). Bottom: increase in the number of false positives (wrongly granted loans).

Learning Repository [84]. This dataset contains 1,000 feature vectors representing loan applications, including applicants' job type, bank account details, gender, etc., and loan details: amount, duration, and purpose. Each has a binary label encoding whether the loan was repaid (70%) or not (30%). We implement the loan approval system as a logistic regression classifier. For our simulation, we split the dataset into the bank's training dataset *X* (800 examples), and the test set (200). The classifier achieves 75.5% test accuracy. We obtained the best results with the use of  $\lambda = 0.5$  as the trade-off parameter for the regularization term (out of 0.25, 0.5, 0.1) in  $J(\theta^*)$ .

We simulate the set of feasible applications  $X_{pool}$  using a subset of *successfully repaid* applications from users that are not in the target group by generating all possible changes of modifiable attributes.

We evaluate our POT in two settings: a "clean" setting in which the only applications received by the bank are those from the collective  $(X \cup X_{pot})$ ; and a "noisy" setting in which other people take loans between the time when the POT is deployed and the time when the classifier is retrained. We report the results in the noisy setting for 10 and 50 additional loans (1.25% and 6.25% of the original dataset, respectively). We repeat the noisy experiments 10 times in which we draw the "noise" loan applications at random.

We present the effect of poisoning on the target group and everyone else in Fig. 5, top. Unsurprisingly, the more poisoning applications there are, the more significant the effect. With 10 poisoning applications (1.25% of the training dataset), our algorithm reduces the number of false negatives in the target group by 9; in the presence of noise this decreases to 7. The false negatives in the rest of the dataset are on average not impacted by our POT.

*5.2.3 POT Impact and Limitations.* Unfortunately, the POT increases the false positives (Fig. 5, bottom). That is, it *shifts* the harm to the bank, which would give more loans to people who cannot repay, increasing its risk. This effect sharply increases with the seventh

application (out of 10). This is due to the poisoning inputs starting to change the model parameter corresponding to the loan purpose. In turn, this increase in false positives could lead to adverse social impact over time if banks try to counter the POT effect [85]. The POT deployers could adjust the trade-off parameter  $\lambda$  to control for the side-effects. Yet, such a reaction would *expose* the incentives and motivations behind the bank's choice of parameters.

To be consistent with the fairness literature [77, 86], we assume that the target group is interested in getting higher credit ratings (through decrease in false-negatives). This model that postulates the access to credit as beneficial, however, is naïve. Even though the loan-making process that relies on risk-based pricing has economic backing, by definition it implies that the less finanically advantaged people will get more expensive loans. Hence, an intervention that aims at increasing inclusion of disadvantaged populations in the lending process can be seen as an instance of *predatory inclusion* [76, 87, 88]. Even if it results in lower loan prices in the short term, it can lead to dispossession in the long run. When harms are viewed through this lens, It is not clear if any technological intervention is capable of counteracting such systems.

# 6 DISCUSSION

Our inspiration to provide a more holistic view on optimization systems and their harms comes from works that point to the logic and potential impact of optimization systems. In particular, Poon has drawn attention to the ways in which optimization systems are driven by outcomes, as exemplified in our utility functions in Sec. 3. This allows for techniques like operational control and statistical management to be the primary mode with which machines interact with phenomena in the world [76]. As a result, these techniques function both as a means for engineering and as "a mathematical state that poses as a solution to political contention" [89]. Optimization is a technique long established in, for example, resource allocation. However, its increasing supremacy in systems that reconfigure every aspect of life (education, care, health, love, work etc.) for extraction of value is new and refashions all social, political, and governance questions into economic ones. This shift allows companies to commodify aspects of life in a way that conflates hard questions around resource allocation with maximization of profit and management of risk [90]. The impact is subtle but fundamental, as evident in the way even we start framing complex and historical questions of justice in terms of utility functions.

To ground ourselves in the world, we used the requirements engineering model of Michael A. Jackson, aligning it with calls for decentering technology and centering communities in the design of systems [91]. We extended the model to ensure that we are not suggesting that problems and solutions are neatly separable and void of power and political economy [92]. However, many elements that might be crucial for conceptualizing optimization systems are still missing. The ontology offers few concepts to capture matters around data, machine learning, or services, and it does not provide deeper insights into addressing issues like fairness or justice. It is, however, a nod to the importance of ontological work for systematizing reflections on our frameworks [93, 94].

To capture the political economy of optimization systems, we turned to utilitarian models and calculations of externalities. Such models are commonly used both in mathematical and managerial forms of optimization and are the cornerstone of neoclassical economics. However, utilitarian models have been thoroughly critiqued for, among others, perpetuating inequalities [35]. Most prominently, Sen has highlighted the limitations of assessing value through consequences, assessing value through subjective utility, maximizing welfare without regard for its distribution, and fetishizing resources over relations between resources and people [95]. Overall, utilitarian approaches are weak in capturing collective interests, social well-being, forms of power, and subjugation. Given these critiques, the central role that these models play in designing large-scale optimization systems is a problem in and of itself.

One possible way forward is to consider alternative economic models for the design and evaluation of systems, e.g., [35, 96–98]. POTs depend and build on the existence of such alternative economic models and the availability of collectivity, altruism and reciprocity. They assume there is something to be gained both individually and collectively, dismissing the selfish agents presumed in utilitarian approaches. In fact, we struggled to express POTs in the utilitarian logic: if we optimize for the utility of the service provider it is hard to justify any POT that may reduce the utility of the service provider.

Beyond economic gains, POTs strategically support people's agency. Optimization systems offer little agency to effectively contest their value proposition [75, 99–101] and offer more optimization as solutions for externalities [102]. POTs can be used to exercise some agency towards an unaccountable [103] and "authoritative system" [104]. Nevertheless, service providers may argue that POTs are gaming the system. Our focus is on social-justice contexts, in which POTs can be cast as "weapons of the geek" for the least equipped to deal with the consequences of optimization [105]. POTs can also serve to expose systems' injustices, achieving transparency and accountability goals. In that sense, they also can come to act like *rhetorical software* "that turn the logic of a system against itself [...] to demonstrate its pathology" [106]. This includes making apparent the damaging results of utilitarian forms of governance prominent in optimization systems.

Despite their positive potential, designing and deploying POTs is not trivial. By virtue of modifying, subverting, or sabotaging an optimization system, POTs may elicit transitions in the system state that result in externalities. 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 undermines the ability of less powerful populations, who may need POTs the most, to have any 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 capabilities.

Just as much, the multi-input form of most optimization systems poses a serious challenge: when optimization is based on continuous tracking across many channels, POTs cannot be built short of creating "optimized doubles" of entities in the environments [107]. The fact that a whole infrastructure for optimizing populations and environments is built in a way that minimizes their ability to object to or contest the use of their inputs for optimization is of great concern to the authors—and should be also to anyone who believes in choice, subjectivity, and democratic forms of governance.

# ACKNOWLEDGEMENTS

We are indebted to Martha Poon for her original framing of the optimization problem. We also thank Ero Balsa for his collaboration on the previous versions of this work, Jakub Tarnawski for his invaluable help with finding the right tool for the traffic-routing case study, and Muhammad Bilal Zafar for his helpful comments.

This work was supported in part by the Research Council KULeuven: C16/15/058, and the European Commission through KULeuven BOF OT/13/070, H2020-DS-2014-653497 PANORAMIX, and H2020-ICT-2015-688722 NEXTLEAP. Seda Gürses is supported by a Research Foundation–Flanders (FWO) Fellowship. Rebekah Overdorf is supported by the Open Technology Fund–Information Controls Fellowship.

### REFERENCES

- Karen Yeung. A study of the implications of advanced digital technologies (including AI systems) for the concept of responsibility within a human rights framework. MSI-AUT, 2018.
- [2] Brent Mittelstadt, Chris Russell, and Sandra Wachter. Explaining explanations in AI. In Conference on Fairness, Accountability, and Transparency (FAT\*), 2019.
- [3] Simone Browne. Dark matters: On the surveillance of blackness. Duke University Press, 2015.
- [4] Oscar H Gandy. Coming to terms with chance: Engaging rational discrimination and cumulative disadvantage. Routledge, 2016.
- [5] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias: There's software used across the country to predict future criminals. and it's biased against blacks. *ProPublica*, 2016. https://www.propublica.org/article/ machine-bias-risk-assessments-in-criminal-sentencing, visited 2019-12-05.
- [6] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. In Conference on Fairness, Accountability and Transparency, FAT\*, 2018.
- [7] Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness and Machine Learning. fairmlbook.org, 2018. http://www.fairmlbook.org, visited 2019-12-05.
- [8] Rebekah Overdorf, Bogdan Kulynych, Ero Balsa, Carmela Troncoso, and Seda F. Gürses. Questioning the assumptions behind fairness solutions. In *Critiquing and Correcting Trends in Machine Learning*, 2018.
- [9] Seda Gurses and Joris van Hoboken. Privacy after the agile turn. In Evan Selinger, Jules Polonetsky, and Omer Tene, editors, *The Cambridge Handbook of Consumer Privacy*, pages 579–601. Cambridge University Press, 2018.
- [10] Irina Kaldrack and Martina Leeker. There Is No Software, there Are Just Services, pages 9–20. meson press, 2015.
- Michael Jackson. The world and the machine. In International Conference on Software Engineering (ICSE), 1995.
- [12] Michael Jackson. Software Requirements and Specifications: A lexicon of practice, principles and prejudices. Addison-Wesley, 1995.
- [13] Sam Lavigne, Brian Clifton, and Francis Tseng. Predicting financial crime: Augmenting the predictive policing arsenal. CoRR, 2017.
- [14] AirBnB. https://airbnb.com, visited 2019-12-05.
- [15] Gaby Hinsliff. Airbnb and the so-called sharing economy is hollowing out our cities. *The Guardian*, 2018. https://www.theguardian.com/commentisfree/ 2018/aug/31/airbnb-sharing-economy-cities-barcelona-inequality-locals, visited 2019-12-05.
- [16] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes. *CoRR*, 2019.
- [17] Sarah Bird, Solon Barocas, Kate Crawford, and Hanna Wallach. Exploring or exploiting? social and ethical implications of autonomous experimentation in AI. In Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT-ML), 2016.
- [18] Zeynep Tufekci. Youtube, the great radicalizer. *The New York Times*, 10, 2018.
- [19] David A. Starrett. Economic externalities. In Fundamental Economics Vol I. EOLSS, 2011.
- [20] Theophile Cabannes, Frank Shyu, Emily Porter, Shuai Yao, Yexin Wang, Marco Antonio Sangiovanni Vincentelli, Stefanus Hinardi, Michael Zhao, and Alexandre M. Bayen. Measuring regret in routing: Assessing the impact of increased app usage. In *International Conference on Intelligent Transportation Systems*, *ITSC*, 2018.
- [21] Toshihiro Kamishima, Shotaro Akaho, and Jun Sakuma. Fairness-aware learning through regularization approach. In *IEEE International Conference on Data Mining Workshops*, 2011.
- [22] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. Fairness beyond disparate treatment & disparate impact: Learning

classification without disparate mistreatment. In International Conference on World Wide Web (WWW), 2017.

- [23] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness constraints: Mechanisms for fair classification. In International Conference on Artificial Intelligence and Statistics, AISTATS, 2017.
- [24] L. Elisa Celis, Damian Straszak, and Nisheeth K. Vishnoi. Ranking with fairness constraints. In ICALP, 2017.
- [25] Michele Donini, Luca Oneto, Shai Ben-David, John S Shawe-Taylor, and Massimiliano Pontil. Empirical risk minimization under fairness constraints. In Advances in Neural Information Processing Systems (NeurIPS), 2018.
- [26] Kaisa Miettinen. Nonlinear multiobjective optimization. Springer Science & Business Media, 2012.
- [27] Julian Le Grand. Equity versus efficiency: the elusive trade-off. Ethics, 1990.
- [28] Carl J Dahlman. The problem of externality. The journal of law and economics, 1979.
- [29] Arthur Pigou. The economics of welfare. London:McMillan and Co., 1920.
- [30] Ng Yew-Kwang. Welfare Economics: Introduction and Development of Basic Concepts. Springer, 1983.
- [31] Louis Phlips. The economics of imperfect information. Cambridge University Press, 1988.
- [32] Kristian Lum and William Isaac. To predict and serve? Significance, 2016.
- [33] Katie Benner. Airbnb vows to fight racism, but its users can't sue to prompt fairness. The New York Times, 2016. https://www.nytimes.com/2016/06/20/technology/ airbnb-vows-to-fight-racism-but-its-users-cant-sue-to-prompt-fairness. html, visited 2019-12-05.
- [34] Os Keyes, Jevan Hutson, and Meredith Durbin. A mulching proposal: Analysing and improving an algorithmic system for turning the elderly into high-nutrient slurry. In Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems, 2019.
- [35] Marilyn Power. Social provisioning as a starting point for feminist economics. Feminist economics, 10(3):3–19, 2004.
- [36] Paul Dolan. The measurement of individual utility and social welfare. *Journal of Health Economics*, 1998.
- [37] Carlos Barreneche and Rowan Wilken. Platform specificity and the politics of location data extraction. European Journal of Cultural Studies, 2015.
- [38] uber. https://uber.com, visited 2019-12-05.
- [39] uber. Uber: Cookie notice. https://www.uber.com/legal/privacy/cookies/en/, visited 2019-12-05.
- [40] Uber. The story of Innisfil. 2019. https://www.uber.com/en-CA/blog/ the-story-of-innisfil/, visited 2019-12-05.
- [41] Fredrick Kunkle. San Francisco police say most traffic tickets go to Uber and Lyft drivers. Washington Post, 2017.
- [42] Michael Graehler, Alex Mucci, and Gregory Erhardt. Understanding the recent transit ridership decline in major us cities: Service cuts or emerging modes? In *Transportation Research Board 98th Annual Meeting*, 2019.
- [43] Angie Schmitt. All the bad things about Uber and Lyft in one simple list. 2019. https://usa.streetsblog.org/2019/02/04/ all-the-bad-things-about-uber-and-lyft-in-one-simple-list/, visited 2019-12-05.
- [44] Alex Rosenblat, Karen EC Levy, Solon Barocas, and Tim Hwang. Discriminating tastes: Uber's customer ratings as vehicles for workplace discrimination. *Policy* & Internet, 2017.
- [45] Michael Sainato. I made \$3.75 an hour: Lyft and Uber drivers push to unionize for better pay. *The Guardian*, 2019. https://www.theguardian.com/us-news/2019/ mar/22/uber-lyft-ipo-drivers-unionize-low-pay-expenses, visited 2019-12-05.
- [46] Sam Sweeney. Uber, Lyft drivers manipulate fares at reagan national causing artificial price surges. WLJA, 2019. https://wjla.com/news/local/ uber-and-lyft-drivers-fares-at-reagan-national, visited 2019-12-05.
- [47] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In Conference on Computer and Communications Security (CCS), 2016.
- [48] Parsa Saadatpanah, Ali Shafahi, and Tom Goldstein. Adversarial attacks on copyright detection systems. CoRR, 2019.
- [49] Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. Style Transfer in Text: Exploration and Evaluation. In AAAI Conference on Artificial Intelligence, 2018.
- [50] Cícero Nogueira dos Santos, Igor Melnyk, and Inkit Padhi. Fighting offensive language on social media with unsupervised text style transfer. *CoRR*, 2018.
- [51] Minghong Fang, Guolei Yang, Neil Zhenqiang Gong, and Jia Liu. Poisoning attacks to graph-based recommender systems. 2018.
- [52] Gang Wang, Bolun Wang, Tianyi Wang, Ana Nika, Haitao Zheng, and Ben Y. Zhao. Ghost riders: Sybil attacks on crowdsourced mobile mapping services. IEEE/ACM Trans. Netw., 2018.
- [53] Les Perelman. Babel generator. http://lesperelman.com/ writing-assessment-robo-grading/babel-generator/, visited 2019-12-05.
- [54] Daniel C. Howe and Helen Nissenbaum. https://adnauseam.io, visited 2019-12-05.

- [55] Josh Eidelson. These instacart workers want you to leave them a 22-cent tip. Bloomberg, 2019. https://www.bloomberg.com/news/articles/2019-01-17/ these-instacart-workers-want-you-to-leave-them-a-22-cent-tip, visited 2019-12-05.
- [56] Jobscan. https://jobscan.co, visited 2019-12-05.
- [57] Allana Akhtar. Is Pokémon Go racist? how the app may be redlining communities of color. USA Today, 2016. https://usatoday.com/story/tech/news/2016/08/09/ pokemon-go-racist-app-redlining-communities-color-racist-pokestops-gyms/ 87732734/, visited 2019-12-05.
- [58] Tega Brain and Surya Mattu. https://www.unfitbits.com, visited 2019-12-05.
- [59] Francis Tseng and Sean Raspet. https://matter.farm/about/, visited 2019-12-05.
  [60] Matthew K Wynia, Deborah S Cummins, Jonathan B VanGeest, and Ira B Wilson. Physician manipulation of reimbursement rules for patients: between a rock and a hard place. Jama, 283(14):1858-1865, 2000.
- [61] Adam Harvey. https://ahprojects.com/hyperface/, visited 2019-12-05.
- [62] Adam Harvey. https://cvdazzle.com/, visited 2019-12-05.
- [63] James Bridle. https://jamesbridle.com/works/autonomous-trap-001, visited 2019-12-05.
- [64] Electronic Frontier Foundation. A guide to youtube removals. https://www.eff. org/issues/intellectual-property/guide-to-youtube-removals, visited 2019-12-05.
- [65] Jeffrey Dastin. Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*, 2018.
- [66] Waze. https://www.waze.com, visited 2019-12-05.
- [67] Svetlana Shkolnikova. Leonia's road closure ordinance ruled invalid. North Jersey Record, 2018. https://northjersey.com/story/news/bergen/leonia/2018/08/ 30/leonia-nj-road-closure-ordinance-ruled-invalid/1145418002/, visited 2019-12-05.
- [68] Bruce Golden. A problem in network interdiction. Naval Research Logistics Quarterly, 1978.
- [69] Delbert Ray Fulkerson and Gary C Harding. Maximizing the minimum sourcesink path subject to a budget constraint. *Mathematical Programming*, 1977.
- [70] Eitan Israeli and R Kevin Wood. Shortest-path network interdiction. Networks: An International Journal, 2002.
- [71] Xiangyu Wei, Cheng Zhu, Kaiming Xiao, Quanjun Yin, and Yabing Zha. Shortest path network interdiction with goal threshold. *IEEE Access*, 2018.
- [72] ortools. https://developers.google.com/optimization, visited 2019-12-05.
- [73] John Forrest, Stefan Vigerske, Ted Ralphs, Haroldo Gambini Santos, Lou Hafer, Bjarni Kristjansson, J P Fasano, Edwin Straver, Miles Lubin, rlougee, jpgoncal1, h-i gassmann, and Matthew Saltzman. coin-or/cbc: Version 2.10.3, 2019.
- [74] Open Street Map. https://openstreetmap.org, visited 2019-12-05.
- [75] Danielle Keats Citron and Frank Pasquale. The scored society: due process for automated predictions. Washington Law Review, 2014.
- [76] Martha Poon. Corporate capitalism and the growing power of big data: Review essay. Science, Technology, & Human Values, 41(6):1088–1108, 2016.
- [77] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In Advances in Neural Information Processing Systems (NeurIPS), 2016.
- [78] Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 2017.
- [79] Zack Friedman. How personal loans can boost your credit score. Forbes, May 2017. https://www.forbes.com/sites/zackfriedman/2017/05/25/ personal-loans-credit-score/#79bce94175ba, visited 2019-12-05.
- [80] Bev O'Shea. What is a credit-builder loan? Forbes, April 2018. https://www. nerdwallet.com/blog/finance/what-is-credit-builder-loan/, visited 2019-12-05.
- [81] Nicolas Papernot, Patrick D. McDaniel, Arunesh Sinha, and Michael P. Wellman. Towards the science of security and privacy in machine learning. CoRR, 2016.
- [82] Battista Biggio and Fabio Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 2018.
- [83] Ambra Demontis, Marco Melis, Maura Pintor, Matthew Jagielski, Battista Biggio, Alina Oprea, Cristina Nita-Rotaru, and Fabio Roli. Why do adversarial attacks transfer? Explaining transferability of evasion and poisoning attacks. In USENIX Security Symposium, 2019.
- [84] Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017. = http://archive.ics.uci.edu/ml, visited 2019-12-05.
- [85] Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. Delayed impact of fair machine learning. In *International Joint Conference on Artificial Intelligence, IJCAI*, 2019.
- [86] Ke Yang and Julia Stoyanovich. Measuring fairness in ranked outputs. In Proceedings of the 29th International Conference on Scientific and Statistical Database Management, 2017.
- [87] Louise Seamster and Raphaël Charron-Chénier. Predatory inclusion and education debt: Rethinking the racial wealth gap. Social Currents, 2017.
- [88] Keeanga-Yamahtta Taylor. Race for Profit: How Banks and the Real Estate Industry Undermined Black Homeownership. UNC Press Books, 2019.
- [89] Fenwick McKelvey. Internet daemons: Digital communications possessed. U of Minnesota Press, 2018.

- [90] Oscar H Gandy. Engaging rational discrimination: exploring reasons for placing regulatory constraints on decision support systems. *Ethics and Information Technology*, 12(1):29–42, 2010.
- [91] Seeta Peña Gangadharan and Jędrzej Niklas. Decentering technology in discourse on discrimination. *Information, Communication & Society*, 2019.
- [92] Julia Powles and Helen Nissenbaum. The seductive diversion of 'solving' bias in artificial intelligence, 2018. https://medium.com/s/story/ the-seductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53, visited 2019-12-05.
- [93] Sam Corbett-Davies and Sharad Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. CoRR, 2018.
- [94] Ben Green and Lily Hu. The myth in the methodology: Towards a recontextualization of fairness in machine learning. In the ICML 2018 Debates Workshop, 2018.
- [95] Amartya Sen. Utilitarianism and welfarism. The Journal of Philosophy, 1979.
- [96] Amartya Sen. Resources, values and development. Harvard University Press, 1984
- [97] Natalia Quiroga Diaz. Decolonial feminist economics: A necessary view for strengthening social and popular economy. *ViewPoint Magazine*, October 2015. https://doi.org/10.6084/m9.figshare.11336048.
- [98] Marc Fleurbaey, Maddalena Ferranna, Mark Budolfson, Francis Dennig, Kian Mintz-Woo, Robert Socolow, Dean Spears, and Stéphane Zuber. The social cost of carbon: Valuing inequality, risk, and population for climate policy. *The Monist*, 2018.
- [99] Elise Thomas. Facebook keeps failing in Myanmar. Foreign Policy, 21. June 2019. https://foreignpolicy.com/2019/06/21/ facebook-keeps-failing-in-myanmar-zuckerberg-arakan-army-rakhine/, visited 2019-12-05.
- [100] Myanmar Civil Society Groups (Phandeeyar, Mido, Burma Monitor, Center for Social Integrity, Equality Myanmar, Myanmar Human Rights Education Network). Open letter to Mark Zuckerberg. https://drive.google.com/file/d/ 1Rs02G96Y9w5dpX0Vf1LjWp6B9mp32VY-/view, visited 2019-12-05.
- [101] Veronica Miracle. Los Angeles councilman tries to work with map apps to alleviate traffic in neighborhoods. ABC7, April 2018. https://abc7.com/traffic/ la-councilman-tries-to-work-with-map-apps/3329737/, visited 2019-12-05.
- [102] James Vincent. Why AI isn't going to solve Facebook's fake news problem. The Verge, April 2018. https://www.theverge.com/2018/4/5/17202886/ facebook-fake-news-moderation-ai-challenges, visited 2019-12-05.
- [103] Frank Pasquale. Odd numbers. algorithms alone can't meaningfully hold other algorithms accountable. *Real Life*, August 2018. http://reallifemag.com/ odd-numbers/, visited 2019-12-05.
- [104] Malte Ziewitz. Rethinking gaming: The ethical work of optimization in web search engines. Social studies of science, 2019.
- [105] Gabriella Coleman. From internet farming to weapons of the geek. Current Anthropology, 2017.
- [106] Francis Tseng. Sabotage, heresy and traps. Conference Talk, June 2019.
- [107] Minna Ruckenstein. Visualized and interacted life: Personal analytics and engagements with data doubles. Societies, 2014.
- [108] Peter J Huber. Robust statistics. Springer, 2011.
- [109] Harold Kuhn and Albert W Tucker. Nonlinear programming. In Traces and emergence of nonlinear programming, pages 247–258. 2014.
- [110] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In International Conference on Machine Learning (ICML), 2017.
- [111] Stephen Boyd and Lieven Vandenberghe. Convex optimization. Cambridge university press, 2004.