



Optimization's Neglected Normative Commitments

Benjamin Laufer
bd156@cornell.edu
Cornell Tech
New York, NY, USA

Thomas Krendl Gilbert
tg299@cornell.edu
Cornell Tech
New York, NY, USA

Helen Nissenbaum
hn288@cornell.edu
Cornell Tech
New York, NY, USA

ABSTRACT

Optimization is offered as an objective approach to resolving complex, real-world decisions involving uncertainty and conflicting interests. It drives business strategies as well as public policies and, increasingly, lies at the heart of sophisticated machine learning systems. A paradigm used to approach potentially high-stakes decisions, optimization relies on abstracting the real world to a set of *decision(s)*, *objective(s)* and *constraint(s)*. Drawing from the modeling process and a range of actual cases, this paper describes the normative choices and assumptions that are necessarily part of using optimization. It then identifies six emergent problems that may be neglected: 1) Misspecified values can yield optimizations that omit certain imperatives altogether or incorporate them incorrectly as a constraint or as part of the objective, 2) Problematic decision boundaries can lead to faulty modularity assumptions and feedback loops, 3) Failing to account for multiple agents' divergent goals and decisions can lead to policies that serve only certain narrow interests, 4) Mislabeling and mismeasurement can introduce bias and imprecision, 5) Faulty use of relaxation and approximation methods, unaccompanied by formal characterizations and guarantees, can severely impede applicability, and 6) Treating optimization as a justification for action, without specifying the necessary contextual information, can lead to ethically dubious or faulty decisions. Suggestions are given to further understand and curb the harms that can arise when optimization is used wrongfully.

CCS CONCEPTS

• **Social and professional topics** → **Codes of ethics; Socio-technical systems; Computing / technology policy;** • **Applied computing** → *Law, social and behavioral sciences;* • **Theory of computation** → *Mathematical optimization.*

KEYWORDS

Optimization, ethics, modeling assumptions, values

ACM Reference Format:

Benjamin Laufer, Thomas Krendl Gilbert, and Helen Nissenbaum. 2023. Optimization's Neglected Normative Commitments. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*, June 12–15, 2023, Chicago, IL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3593013.3593976>



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

FAccT '23, June 12–15, 2023, Chicago, IL, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0192-4/23/06.
<https://doi.org/10.1145/3593013.3593976>

1 INTRODUCTION

Mathematical optimization¹ is a way to relate a set of potential decisions to a specific goal and use this relationship to identify the best-performing decision. It is a paradigm used in industry and government and aids in developing a variety of computing tools including machine learning (ML) models. Optimization's widespread use and applicability is enabled, in part, by its level of abstraction. By representing real decisions as a set of *decision variable(s)*, *objective(s)*, and *constraint(s)*, optimization methods identify solutions that maximize the objective while meeting the relevant constraints.

However, optimization does not necessarily imbue a decision with legitimacy or ethical justification. Because its parameters and functions are left to be specified, an optimization can be designed to serve subjective or parochial interests, out of step with social welfare. For example, when a social media advertiser uses optimization to target ads in a way that maximizes user engagement, it may not be the case that the so-called 'optimized' decision is one that is best for society.

This paper aims to highlight optimization's normative choices and assumptions. In Section 2, we describe the *life cycle* of an optimization, particularly as it is used in operational decision-making, machine learning, and research. Our treatment of the life cycle or pipeline associated with applied optimization is intended to be reminiscent of existing characterizations of the *machine learning life-cycle* which have proven useful for end-to-end analysis [20, 42, 55, 112]. We then focus on the various choices and assumptions inherent in the optimization life cycle, which we categorize into the following components: modeling, measurement, computing, and application. In Section 3, we identify six emergent issues that may be neglected: 1) *misspecified values*, 2) *problematic decision boundaries*, 3) *failing to account for multiple agents*, 4) *mislabeling and mismeasurement*, 5) *faulty use of relaxation and approximation*, and 6) *treating optimization as a justification for action*. Finally, Section 4 articulates suggestions and directions for research on optimization and its normative implications.

2 THE OPTIMIZATION LIFE CYCLE

Examining the assumptions and choices laden in an optimization requires thinking about optimization as a sociotechnical process. In this section, we describe the typical steps in the *life cycle* (or *pipeline*) associated with applying optimization in real-world settings. An equivalent life cycle for ML systems has been developed and widely used (see e.g. [42, 111]) including as a framework for evaluating normative assumptions and potential harms [5, 20, 69, 71, 112]. We propose an analogous pipeline that describes the use of optimization. Equipped with such a pipeline, we will systematically analyze each of the component choices and assumptions in order

¹For the formal definition used throughout this paper, see Appendix A.

to unearth normative commitments. A visual representation of the optimization life cycle is provided in Figure 1.

Although our goal is to provide a general characterization of the use of optimization from its inception (measurement and modeling) to its execution (solving and applying), it is important to note that optimization can look very different depending on the context. Thus, we'll start by discussing some domains where optimization is commonly applied. Then, we'll discuss the component steps involved in a typical optimization process.

Common Contexts

Here we discuss three common application domains for optimization: operational decision-making, machine learning, and scientific research.

Operational Decision-Making. Optimization is perhaps best understood as a method to approach complex logistical decisions. These decisions aim for operational efficiency for large-scale industrial or institutional endeavors with significant complexity. Canonical examples cited in introductory optimization courses include signal processing and networked communication [73, 87], portfolio optimization in financial mathematics [17, 29, 60, 96], and transportation [16, 74, 93].

For examples of high-impact optimization problems applied to decision-making, one can simply look to the history of linear programming.² LP problems gained prominence in both the US and USSR during the Second World War, where mathematicians and economists in both countries began considering how to devise optimal large-scale policies with limited budgets — these policies ranged from wood production to air force allocation [106].

Machine Learning. Optimization and machine learning are intimately related and machine learning often relies on optimization. The algorithms used to train a ML model may use a formal notion of loss—for example, prediction error—in the training process. Fitting parameters to *minimize* this loss can be written as an optimization problem.

For example, a univariate linear regression using ordinary least squares (OLS) finds a line to fit the data. Of all possible lines which could be used to describe the relationship in the data, the 'line of best fit' according to OLS is the line which minimizes the sum of squared errors between the predicted values and the observed data. This task can be written as an optimization problem, where the decision variables are the slope m and the intercept b of the line. The feasible set is all real numbers (the problem is unconstrained beyond $m, b \in \mathbb{R}^2$). And the objective is to minimize loss, namely the sum of squared errors.

For more complex models using high-dimensional datasets, the optimizations become more tedious to write out but the idea is the same: Some notion of loss is defined so that the model fit will resemble the observed training data, and defining the model involves solving an optimization to minimize loss.

²Linear programming is one particular type of optimization where the constraints are a set of linear inequalities and the objective is a linear equation. In the first half of the twentieth century, optimization was more commonly known as 'mathematical programming,' where the word *programming* referred to logistical decision-making rather than computer programs.

Although we predominantly draw from examples in logistical decision-making and management science, we believe that this paper's framework yields useful insights about optimization's role in the machine learning pipeline. For example, defining loss is a choice with particular priorities and commitments. Sometimes, this decision is made simply by choosing a loss metric which is computationally easy to work with. Another example of this paper's usefulness in scrutinizing ML is in assumptions about decisions and decision boundaries: Often, ML predictions are translated into decisions without considering the impacts, social implications, or feedback effects associated with the particular decision.

Scientific Research. Optimization crops up in a number of disciplines outside of operations research, math and statistics. Where optimization is applied to scientific research, it may guide design decisions for building scientific instruments [21, 90, 120] or in the discovery of new drugs to test [19, 26, 31, 94, 97]. These are examples of *operational* uses, similar to those described above.

Another use for optimization in research is as a descriptive tool. For example, researchers employ agent-based models to understand complex processes, and optimization often plays a role in simulating how modeled agents behave. The use of optimization in simulation or theoretical models is distinct because it is used to describe the behavior of an agent. The optimization is used to draw descriptive conclusions that do not necessarily guide real-world decisions. Optimization in this sense *describes* rather than *dictates*. For example, in game theory's applications to economics [2, 28, 102, 109] and biology [34, 48, 50, 70], optimization is used to model agent decisions and strategies. Pinpointing assumptions and normative imperatives within these optimizations may be useful—they may reveal values and priorities held by researchers, or assumptions made about economic subjects. However, because these optimizations are not used to justify a particular 'optimized' decision, we do not draw examples from these domains and consider them to be out of scope.

Choices, Abstractions and Assumptions

Formulating an optimization problem requires transforming complex information in the world into a simplified model. Models of any sort can (ideally) use abstraction and representation to make information communicable, actionable or useful [92, p 20]. Therefore, models that are useful often omit or collapse a significant amount of information.

The idea of abstraction is familiar to computer scientists, whose introductory computer science classes often involve creating a simple *class* or *object* such as a 'Bank Account' or an 'Employee' (see e.g. [110, p 316-319]). Doing so does not involve building a physical bank, of course; instead students simply define a set of variables (such as name and salary) and functions (such as "retire" or "promote"). By removing significant material and contextual aspects of a real-world system, a computer program can simulate a few key functions in order to make a particular point or perform a particular task. Abstraction is powerful and crucial to the field of computing; however, it also poses challenges: through making assumptions and collapsing information, models risk misrepresenting or omitting key information.

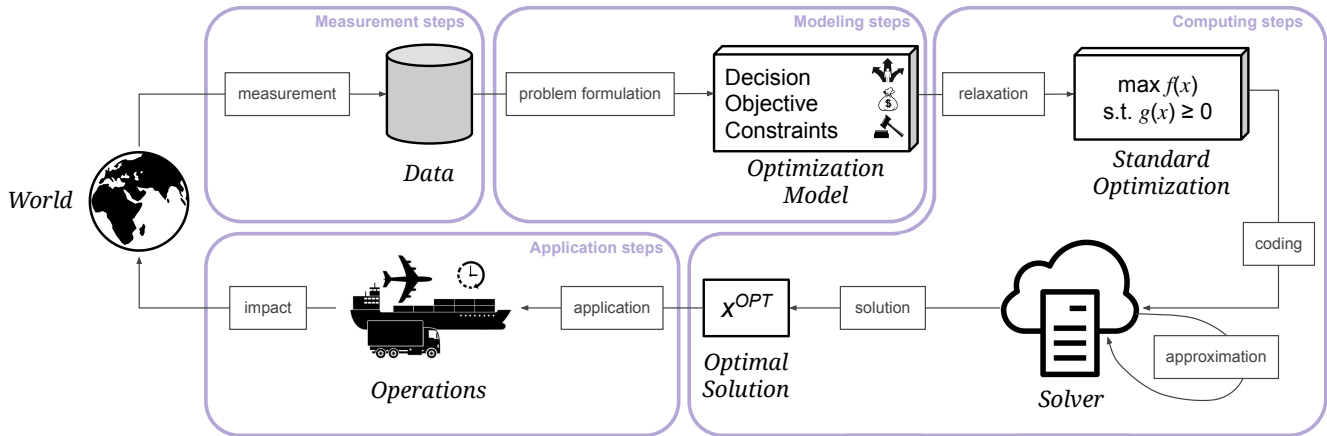


Figure 1: Optimization life cycle, beginning with measurement and data generation and ending with operational decision-making and real-world impact.

In the following sections, we trace the component steps that constitute the optimization process. The process spans modeling, measurement, solving, and deploying, as shown in Figure 1. Reasoning about these components requires assumptions—some axiomatic, others readily contestable. By systematically enumerating and formalizing these as constitutive parts of the optimization life cycle, the system’s risks, harms, indeterminacies,³ and concealed normative positions can be more easily identified and communicated.⁴

2.1 Modeling

A preliminary step in putting forward an optimization model is defining the set of variables and parameters that will be included in the model. These variables are then used to formalize the objective function, decision variable and constraints which together define the feasible set of potential decisions and their corresponding utility or performance. As the model is defined and worked with, modeling choices play a determining role in the larger optimization process and any resulting decisions.

In this section, we focus on the three components of a standard optimization and discuss the negotiable decisions and assumptions at play during the modeling process. We refer to Appendix A for the standard form of an optimization problem.

The Objective Function. Specifying an objective function, utility function or goal immediately invokes values judgements. Formalizing an objective function often means relying on quantifiable notions of utility or welfare that are relevant for a given decision. Formalization of this sort can also involve omitting values or desiderata from consideration. When multiple goals are valuable to consider, modeling assumptions are needed to incorporate these goals into a single objective. Often, this process involves defining a

welfare function, a term borrowed from economics (see e.g. [23, 62]). Instantiating an objective of this sort typically involves weighing the relative importance of different goals, and characterizing the overall objective function’s mathematical properties, like whether it is convex [18].

The Decision Variable. Specifying the decision variable(s) requires making a number of assumptions about the decision-making agent and their degree of control. The optimization paradigm presumes that a decision-maker has a certain set of options, and inquiry is narrowed to only comparing between the options. Such an approach avoids thorny questions about who is in control and whether they ought to be. Consider, for example, an optimization aimed at setting a regulatory fine. Defining the decision mathematically may necessarily involve stripping away particularities about how fines result from complex bureaucratic and deliberative processes. Instead, an optimization might model the decision simply as a single variable representing the dollar value of a fine. Abstractions of this sort can appear self-evident even as they conceal choices with decisive normative commitments.

The Constraints. Specifying constraints is also known as specifying the *feasible set* of decisions. This step involves making assumptions about the physical relationship between variables. For instance, an optimization problem may try to maximize production subject to a budgetary constraint. Budgets often play into optimizations and models must specify what kind of scarcity exists and how it is being measured: is the agent constrained monetarily? spatially? temporally? Though the feasible decision space is traditionally conceived of as the set of possible decisions, constraints are often used to guarantee that a decision meets some baseline ethical or social standard. We call these sorts of constraints *values-driven constraints*, and note that *fairness* is a particularly common values-driven desideratum that gets modeled as a constraint.⁵ Section 3.1.2

³See Dobbe et al. [32], which similarly uses assumptions, indeterminacies and *vague* aspects of socio-technical systems to explore their normative implications.

⁴We do not claim to provide an exhaustive method for anticipating future harms or *verifying* optimizations. Our project is influenced by Smith [108] on this matter: we seek to disambiguate and excavate the social and political aspects of optimization as a method, rather than make grand claims about its use.

⁵See Passi and Barocas [95] on problem formulation and modeling in the context of fairness.

includes a more in-depth discussion of the modeling assumptions involved in incorporating values into an optimization’s constraints.

2.2 Operationalization, Measurement and Quantification

Applying optimization requires a number of decisions aimed at *relating an abstract quantity to the external world*. Where modeling (discussed in Section 2.1) involves modes of knowledge aimed at making sense of the chaos of the world, measurement is how we empirically find new information or physically instrument the world. There is significant work on the elusive nature of measurement—here we highlight theoretical groundwork [8, 38, 57, 61, 84] as well as technical work on bias in measurement [15, 51, 79, 91]. In this section we categorize optimization’s measurement assumptions into two categories: 1) defining what is being measured and 2) deciding how to measure it.

What is being measured? Deciding what is relevant or worthwhile to measure is an impactful step in setting up an optimization model. When written out on paper, an optimization has finite number of variables, and the definitions of these variables are often formally listed. These variables constitute a highly abstracted and stylized version of the world—few of their characteristics may correspond to the real-world entity that they try to mimic or count. For example, in public transit optimization, a bus might get modeled as a capacity, route, and fuel level—all other characteristics are removed from consideration.

How is it being measured? Research on optimization often focuses on methods for solving and approximating certain types or groups of optimization problems. Accordingly, much of the attention devoted to optimization isn’t concerned with the measurement of particular quantities. Instead, it is assumed that quantities are well-defined and measurable. Therefore, measurement is notably far afield from academic research related to optimization. When optimization is used to guide real-world decision-making, however, measurement becomes a necessary component to understand the impacts and constraints of a decision. Measurement is never perfect; and any measurement technique inevitably introduces *error*. Designers of optimization models must decide what to do about error: for example, they may ignore it, leaving it out of the optimization, or they may use stochastic optimization [52, 100] to directly model uncertainty.

Example: Transportation. To better understand how measurement considerations play into an optimization model, consider the example of a logistics company that aims to optimize routes for its delivery vehicles. To perform this sort of optimization, the company may want to consider the way that refueling constrains its decision. The resulting optimization may assume that trucks have a certain unchanging fuel efficiency in miles per gallon (MPG) to make it easy to relate route distances to gasoline use. Doing so would rely on a certain kind of measurement (average MPG for a vehicle) to relate two modeling variables (route distance and gas usage), without requiring further measurement about how fuel efficiency varies depending on truck speed, number of traffic lights, fuel grade, price, and other factors. Assuming a static fuel efficiency greatly reduces what must be measured by making a simplifying assumption about how a quantity is measured.

2.3 Computing

A dominant strand of research on optimization is concerned with computing and computability. Research findings in the field of Operations Research include algorithms to solve optimizations (e.g. [10, 59, 88]) and methods to approximate optimal values (e.g. [63, 65, 89]). In some cases, computability is in tension with accurately and comprehensively modeling the state of the world. For example, an objective function may be multi-modal or non-differentiable in practice, but accurately modeling these features would make the optimization computationally difficult to solve. As such, optimization commonly relies on *relaxation* to transform a domain-specific, computationally-intensive optimization into a familiar, solvable type of optimization, ideally one that yields the same or a provably nearby solution. In the process of solving these optimizations, sometimes an optimal solution is still too computationally expensive to attain, so *approximation algorithms* [116] are similarly used to find a solution that is nearly-optimal. In this section, we describe these two steps — *relaxation* and *approximation* — and aim to excavate their normative assumptions.

Relaxations. Relaxations are the transformation of an optimization into another optimization that is easier to solve or work with. The formal definition of the sort of relaxation we refer to in this paper is as follows:

Definition 2.1 (Relaxation). A minimization problem Q is said to be a relaxation of a minimization problem P if the following conditions are met:

- (1) $F(P) \subseteq F(Q)$, where $F(x)$ is the feasible decision set defined by the decision and constraints of optimization x .
- (2) The objective function of optimization Q is less than or equal to that of P for all parameter values in $F(P)$. [44]

Approximations. Approximations are quantities or methods that are nearly correct. Relaxing an optimization and using the solution as an answer to the original problem (as described above) is one type of approximation. Here we focus specifically on *approximation algorithms* or ways of approximating certain types of optimization problems [114, 116].

Definition 2.2 (Approximation algorithm). An α -approximation algorithm for an optimization problem is a polynomial-time algorithm that for all instances of the problem produces a solution whose value is within a factor α of the value of an optimal solution [116, p 14].

Approximations can happen after an optimization has already been relaxed or fully defined; they do not necessarily rely on expanding the feasible set or changing the objective function as relaxations do. Instead, approximations are often algorithmic approaches that can be incorporated into a computational solver (see Figure 1). Using approximations to solve an optimization problem requires making a number of assumptions—for example, it is assumed that a sub-optimal solution with nearly-optimal utility performance is acceptable as a decision outcome. Without formally reporting and considering performance guarantees (α in Definition 2.2), those who apply optimization methods may land on approximate solutions that are far from optimal without knowing it. Such pitfalls are discussed further in Section 3.5.

Relaxations and approximations suggest a trade-off between at least two desiderata: reducing computational complexity and maintaining performance (e.g., α). Of course, modelers are juggling a number of other potential goals as well, including consistency [25, 36, 37], robustness [12–14, 41], and correctness [33, 40, 101]. Trade-offs between goals abound in optimization, and have been discussed especially in machine learning contexts [11, 24, 66, 119]. These trade-offs may or may not receive explicit consideration from model designers. However, whether or not the designer has explicitly taken these values into consideration or weighed between them, the underlying values and normative commitments are reflected in an optimized decision.

2.4 Application

Optimization research tends to categorize and solve problems based on particular groupings or structures of problems. Often, solution algorithms are found through a process of typifying the the objective function and constraints and demonstrating solutions that work on certain general categories of problems (e.g., convex problems [18]). Solutions, approximations and relaxations are then developed for these various already-abstract problems without considering domain-specific contexts or scenarios.

Our characterization of optimization research as abstracted from real-world scenarios is not intended as an indictment of the research but instead provides lessons for how optimization might be safely and ethically applied in real-world domains. When applying optimization, it is essential to consider whether the application is reasonable in the real, non-abstract context, or whether it is, for example, far-fetched, impossible, or unethical. Especially when objectives are contested, numerous, complex or conflicting, researchers and practitioners sometimes find that the right choice is refusal to adopt or use a technological tool [105, 118].

When optimization is applied in real-world settings, decision-makers necessarily make assumptions about *applicability* to conclude that an optimal solution is the right course of action. Assumptions are made about both the *impact* of the optimization and the *relation* between the decision and other decisions. Often, even though individual decisions are said to be optimized, the process of optimization requires tweaking, iteration, and repeated use over multiple decisions.⁶ The dynamics arising from applicability assumptions may generate forms of instability: system oscillations [121], runaway feedback effects [35, 72], amplification, or domain drift [113] over time.

3 EMERGENT PROBLEMS AND NEGLECT

The optimization paradigm requires a number of choices, simplifications, and assumptions, as demonstrated in Section 2. These choices can contain particular commitments and serve particular interests. A company's goals, priorities, values, norms, and conventions can influence the considerations that end up in an optimization model. As such, 'optimal' decision-making in one company or sector means something particular and context-dependent. A decision described as 'optimal' or 'optimized' is not, necessarily, a virtuous decision.

⁶Failing to model the decision as dynamic has been noted by Selbst et al. [103] as a potential vulnerability for decision-making in socio-technical systems. The authors refer to this issue as the *ripple effect trap*.

Our contention is that optimization is a powerful tool; its use is not ethically neutral simply because it is quantitative and scientific.⁷ Optimization can be used for a number of desirable or undesirable ends, and its application can lead to emergent issues and downstream consequences. We do not argue that optimization's constituent choices and assumptions are inevitably problematic. Instead, we use the process (or life cycle) of optimization to identify ethical pitfalls, harms, and wrongs that can emerge.

When are these problems 'neglected'? To neglect is to *fail to care for properly*. The worry with optimization is that it can appear to be objective and exhaustive. Optimization is a systematic paradigm used to approach decisions, constraints, and goals, however, its use does not guarantee that all relevant decisions, constraints, goals, interests, values, and other considerations have been adequately weighed. As we describe in this section, there are several cases where optimizations do not properly account for particular harms or values. When an optimization does not take proper consideration of ethically relevant interests or consequences, these interests or consequences may be described as *neglected*.

To be sure, some emergent problems and downstream consequences are unknowable and unforeseeable. For example, consequences that are far enough in the future, or that arise from apparent anomalies, may be exceedingly difficult to predict and account for in a decision calculation. These harms and wrongs are difficult to attribute to carelessness or malice. But, by observing and documenting known problems and failure modes, these cases can inform standards of care and accountability mechanisms.

This section puts forward six emergent issues that can be overlooked when optimization is used to justify or derive a decision. These issues arise at various steps throughout the optimization life cycle. Where relevant, we provide examples and offer suggestions.

3.1 Misspecifying Values

Optimization requires formally specifying at least one decision and at least one measure of utility. The values laden in optimization are made most apparent by the existence of an objective function, which formalizes the goal that the decision-maker is assumed to wish to maximize. Aligning optimization objectives with social values is a topic that has received significant interest [49, 78]. But values play into other components of a model besides the objective. For instance, models can include constraints motivated by ethical norms or values like fairness.

Here, we discuss the issues that arise when an optimization model is misspecified such that values are not given proper consideration. We identify two mechanisms that result in misspecified values: 1) leaving values out of an optimization completely and 2) mistreating or modeling values within an optimization. In discussing the latter type of misspecification, we consider instances where values are wrongfully incorporated as optimization constraints (instead of objectives) or as objectives (instead of constraints).

⁷Relatedly, applications of *mechanism design* have been referred to as 'technologies of depoliticization' [53]. See also Viljoen et al. [115], Hitzig et al. [54], and Finocchiaro et al. [39] on mechanism design's normative dimensions. Mechanism design can be described as a type of optimization and, similarly, can invoke objectivity or impartiality.

Table 1: Life cycle components, categories of neglected problems, and particular emergent issues in optimization.

Life Cycle Steps	Emergent Problems	Particular Issues
Modeling (2.1)	Misspecifying Values (3.1)	Omitting Values (3.1.1)
		Improperly Modeling a Value as a Constraint (3.1.2)
		Improperly Modeling a Value as an Objective (3.1.3)
	Problematic Decision Boundaries (3.2)	Faulty Modularity Assumptions (3.2.1)
		Feedback Loops (3.2.2)
Neglecting to Consider Multiple Agents (3.3)		Objective Reflects Narrow Interests (3.3.1)
		Decision Model Formalizes Notions of Control (3.3.2)
Measurement (2.2)	Mislabeling and Mismeasurement (3.4)	Imprecise Labels and Measurements (3.4.1)
		Biased Labels and Measurements (3.4.2)
Computing (2.3)	Faulty Solution Steps (3.5)	Faulty Use of Relaxations (3.5.1)
		Faulty Use of Approximation Algorithms (3.5.2)
Application (2.4)	Using “Optimal” as Justification (3.6)	Justificatory Language (3.6.1)
		Reverse Engineering/Retrofitting (3.6.2)

3.1.1 Omitting Values. Omitting values from an optimization problem can be either careless or deliberate. A careless omission of values arises when the humans designing and executing an optimization fail to properly consider the ramifications of their choices. A deliberate omission of values arises when the humans do consider the ramifications of their choices, and then purposefully choose to ignore or conceal them. For example, a company that makes custom shirts might decide where to source their textiles. An optimization problem that minimizes the cost of textiles might suggest changing which factory they contract with. However, it could be that the new factory has dangerous working conditions or otherwise harsh and unethical labor practices. If the optimization designer were unaware of this distinction, the omission of labor considerations from the decision would be *careless*; if they were aware, it would be *deliberate* and malicious.

Scrutinizing which values and implications are included or excluded from an optimization provides important information for attributing harms and establishing *accountability*. However, the optimization model alone does not contain all the relevant and necessary information. For example, one cannot differentiate negligent and deliberate omissions simply by looking at an optimization problem—further contextual evidence is needed to qualify the human decisions behind the model. If somebody omits a value from an optimization deliberately, they may be using the optimization as a way of *justifying* an otherwise faulty or problematic decision. We turn back to this idea in Section 3.6.

3.1.2 Improperly Modeling a Value as a Constraint. An optimization model may include a constraint that is intended to operationalize a moral or ethical imperative. For example, a new housing development may be required to include a certain number of low-income or rent-stabilized units, leading developers to treat this requirement as a constraint on building a profitable new tower. As another example, when companies boycott sales in regions to voice dissatisfaction with foreign states’ political decisions, their resulting business strategy is constrained by their stance; even while their goal is profit and expansion.

In the optimization paradigm, constraining a decision removes a set of choices from consideration. Constraints are categorical rules that determine the set of feasible choices. They are inflexible and not negotiable. Incorporating a values-driven constraint, therefore, lends itself to ethical considerations that are categorically satisfied over a range of (permissible) behaviors. Consider, as a toy example, a man who decides to go on a diet that forbids foods with added sugar. This diet *constrains* the set of possible decisions the man can make. As long as he is not eating food with added sugars, he is free to eat whatever he pleases. The man’s underlying goals are unknown to us: he might want to lose weight; he might want to snack less often; or he might want to protest the sugar industry. A constraint is categorically either satisfied or not satisfied—as long as it is satisfied, it does not guide a decision.

Though constraints play an important role in determining the feasible set of decisions available to a decision-maker, they can also be weak as ethics interventions. As long as two potential decisions are both allowed, constraints provide no reason to prefer one potential decision over the other. Put another way, within the feasible set, a constraint has no bearing on what option gets chosen. A constraint is described as *trivial* if instituting it does not change the optimal solution. Even when a constraint is non-trivial, it tends to yield an optimal solution that lies on the constraint boundary: *solutions tend to ‘bump up against’ constraints*. For example, if a for-profit housing development must contain a certain minimum number of rent-stabilized units, an ‘optimal’ business strategy might suggest including the minimum allowable rent-stabilized units, in order to maximize profit. In practice, the ‘optimal’ solution might do the *bare minimum* at meeting some values-driven imperative, if that imperative is modeled as a constraint.

In computing contexts, values may be operationalized as a constraint or as an objective. For example, the value of *fairness* might be operationalized as a constraint on an ML system, or alternatively, the ML system may be designed to advance *equity* as its objective. These choices can profoundly impact the ‘optimal’ decision.

Example: Access. Consider a city which has two potential ways of providing transportation to its residents. One option is to allow

private industry to provide transportation through contracts with the city. In this case, providers aim to *maximize profit*, subject to certain *contractual and legal constraints*. For instance, these companies could agree to provide a certain level of access to low-income neighborhoods. An alternative option is to treat transportation as a public service, where a single entity with a budget is tasked with providing transportation. In this public approach, an agency might aim to *maximize access* subject to *budgetary constraints*. These two approaches to optimization logic encode the same values (access and budgetary responsibility) but do so in ways that can lead to vastly different levels of access across communities.

3.1.3 Improperly Modeling a Value as an Objective. An optimization's objective is a way of defining utility as a quantifiable and measurable function of the decision variable. This way, for any two decision values, the utility of one decision can be quantitatively compared to the utility of the other decision. Defining the objective lends itself well to *ends-based* or *utilitarian* frameworks which treat value as quantifiable and to some extent predictable [107]. An optimization model presumes to model the relationship between a decision and its utility outcome—in non-stochastic optimizations, this relationship is presumed to be fully known and defined, whereas stochastic optimization models randomness explicitly [100]. Sometimes, values are treated as objectives when they ought to be modeled elsewhere in an optimization.

Example: Safety. Car companies communicate that they aim to build the safest cars. In reality, these companies have many goals including engine efficiency, speed, comfort, and profit. If safety were truly the utmost objective, a car would not move: staying parked is the safest state for a car. Of course, the functionality of a car requires that safety is not always modeled as a *sole* objective (even if such an objective is communicated in advertisements). More realistically, many car companies' safety considerations may be treated as *constraints* on business decisions.

3.2 Problematic Decision Boundaries

A first step in specifying an optimization model is defining and constraining the *decision variable*. Specifying which variables are freely tunable and which values are feasible lays the roadmap for solving the optimization. Doing so dictates the set of potential outcomes, how the problem may be solved, and, perhaps most importantly, which factors are within (versus outside) the control of the decision-maker.

Specifying the decision involves assumptions about *who makes the decision* as well as *what power they have over the state of the world*. Making such assumptions risks glossing over the mechanism through which a decision-maker enacts change, and a variety of requisite or related decisions that may be involved. The remainder of this section is devoted to 1) neglectful decision boundaries for synchronous decisions (which arise because of *faulty modularity assumptions*) and 2) neglectful decision boundaries for asynchronous decisions (which can lead to *feedback loops*).

3.2.1 Faulty Modularity Assumptions. It can be simpler or convenient to treat problems as separate optimizations when the number of decisions is too vast. The optimization paradigm makes it very

easy to divide and conquer by treating the solution of one optimization as an (unchanging) parameter in a second problem.

Imprecise decision boundaries, including incorrect assumptions about modularity, lead to narrow optimizations that may not serve the broader function of a system or a society. Divide-and-conquer optimization strategies, with incorrect modularity assumptions, can have unanticipated normative and political implications. Mulligan and Nissenbaum [85, p 245] provide the example of access control to illustrate this point: although tech companies can justify switching from passwords to fingerprints to face ID by saying each new technology more *optimally* performs the specified function of access control, such changes have notable political implications. For instance, each successive handoff might be said to curtail user control and diminish transparency.⁸ Improvements or optimizations directed at sub-functions of a larger system should consider broader societal implications and interests. Otherwise, optimization alone does not justify exchanging or updating a technological sub-component.

Example: Transit. Consider a metropolitan transit authority faced with the following decisions: given budget and staff constraints, how much to charge in tolls and whether (or which) lanes should be restricted to high-occupancy vehicles during peak travel times. In actuality, these decisions are not independent: how much a highway charges in tolls will impact the number of cars on the road by inducing demand, which in turn is important for evaluating to what extent high-occupancy vehicle lanes will improve overall traffic throughput. Similarly, how many lanes will be restricted—or the window of time during which the restriction is enforced, or how costly the fine—will constrain how many people commit to carpooling and significantly affect budget and revenue considerations for the toll decision. These choices could be treated in isolation: first, decide how many restricted lanes would serve the carpooling population, and then (conditional on a set number of lanes) decide what to charge in tolls. But to do so would ignore the fundamentally interdependent nature of the decision, and may generate spillover effects that the simplistic model does not account for.

3.2.2 Feedback Loops. Optimizations may arise in a repeated fashion where it is unreasonable to assume each sequential decision is independent. For the same reason numerous synchronous decisions may be haphazardly split into separate optimizations (discussed in Section 3.2.1), modelers may unreasonably assume sequential decisions are independent.

The phenomenon of feedback in optimization is represented in Figure 1's depiction of the entire optimization life cycle as a loop. Optimizations can be used repeatedly, even if the formal model does not account for one optimization's impact on the next. Similar feedback loops have been observed in systems that do not necessarily involve optimization — in domains ranging from mechanisms and acoustics to economic inequality, the concept of feedback and path-dependence explains patterns where compounding or overlapping impacts from multiple sequential events lead to unanticipated or chaotic outcomes. In the context of fairness for socio-technical systems, Selbst et al. [103, p 62] described this phenomenon (and corresponding neglect on behalf of designers) as the *ripple effect trap*. In the case of optimization, likewise, it may be that an optimal

⁸See also Goldenfein et al. [46].

decision in one instance is sub-optimal or actively harmful when repeated sequentially over time.

Example: Predictive Policing. Consider the example of PredPol (now re-branded as Geolítica), a commercial predictive policing software that used historical data to statistically train ‘optimal’⁹ policing recommendations [82]. The software was accused of formalizing and amplifying historical biases [56, 72, 104]. Subsequent studies used models of sequential predictive decisions to explain the reasons for these emergent biases: Decisions made as single-shot optimizations can introduce *runaway feedback effects* when applied sequentially over time, since predictive algorithms direct police to neighborhoods where crime was historically observed [35, 68].

3.3 Neglecting to Consider Multiple Agents

A notable feature of optimization is its strict delineation between decision variables and other (non-decision) variables. Certain parameters are considered to be under the control of the decision-making agent, and other parameters are exogenous and outside the direct control of the agent. This approach lends itself to treating decisions as belonging to a *singular agent*. Even though optimizations are commonly used by institutional bodies and passed between numerous stakeholders ranging from analysts to engineers to executives, an optimization model may not consider that the various constituents of an organization are distinct decision-makers. In this section, we consider how optimization problems may neglect to properly consider the agency and interests of multiple decision-makers, and how such neglects can be avoided or identified. When multiple agents are not taken into consideration, optimizations may only serve the interests of a few people, and additionally may codify or formalize particular structures of management and control. In this section, we consider each of these emergent issues.

3.3.1 Objective Represents Narrow Interests. Optimization is used to make logistical decisions that can impact many people, both inside and outside the decision-making apparatus. For a given optimization, the relative importance of different agents’ utility can be defined in the objective function. When an objective omits or de-values harmful impacts on certain populations, the resulting “optimal” decision can be ethically wrongful.

When an objective reflects only certain interests in an organization, the optimal decision might justify exploitative working conditions. Take the toy example of a board of directors deciding how to use a certain amount in earnings. Without union representation on the board, they might decide that their utmost priority is to keep investors happy and therefore pay dividends. With union representation, however, they might adjust their priorities and devote more earnings to employee compensation. The choice of a *goal* or *objective* is highly dependent on the organizational structure.¹⁰

Example: Technology Adoption. When financial mathematicians began to suggest that optimization could improve portfolio diversification schemes, firms were slow to adopt these schemes. In one academic paper devoted to the reasons why firms did not adopt these measures, a reason for the resistance was the social and

political structure of the firm: “Probably the single most important reason why many financial institutions don’t use portfolio optimizers is **political**. This is because the effective use of an optimizer mandates significant changes in the structure of the organization and the management of the investment process” [77, p 32]. Firm managers did not want to cede decision-making power to quantitative analysts, suggesting that organizational social structure influences the use of optimization.

The “narrow interests” described so far concern institutional decision-making involving more than one agent. But problems of narrowly defined goals can also have massive impacts on people external to a company or institution, who have no involvement or even awareness of the decision. The notion of *externalities* [6, 9, 27] (see also [22]) helps to describe how unaccounted-for impacts can affect people who might face barriers entering into contracts or fending off negative consequences. People downstream from a river might be negatively impacted by factory pollution and may not even know it. These effects can be the result of negligent or malicious tweaks to an objective function. They highlight the importance of critically evaluating optimizations and their stated goals.

3.3.2 Decision Model Formalizes Notions of Control. The *decision variables* in an optimization are an important indicator for what an agent has control over. In formalizing exactly what decision is assumed to be under the control of an agent or organization, an optimization can codify existing social structures.

Example: Workforce Optimization. Industries are fiercely competitive over employee talent, and retaining employees has become an important priority for companies. So-called ‘workforce optimization’ [3, 86, 99] or ‘staff planning problems’ [1, 4, 7] devise strategies to allocate, promote and pay workers in order to maximize retention and efficiency while minimizing costs. Such models make assumptions about a particular type of employee/employer relationship, for example that employees are contractors whose time can be allocated freely; or that layoff recommendations can come from performance metrics rather than human judgments. The clearest examples of decision models formalizing or implicating systems of control are warehouse management and logistics studies that create staffing tools for warehouse managers that balance worker discomfort, worker risk, and efficiency [67, 117].¹¹

3.4 Mislabeling and Mismeasurement

The choice of labels and measurements can lead to undesirable outcomes in the use of optimization. Measurements are erroneous when there is a disconnect between a measurement and an intended (or true) value. In this section, we discuss two types of mismeasurement: imprecise measures and biased measures. These issues can arise from a *conceptual error*, whereby proxies or incorrect labels are used instead of a true underlying value; or they can arise from a *measurement error*, whereby instruments or empirical methods do not work as expected. Accuracy is comprised of both precision and the absence of bias in measurement—below, we briefly discuss mislabeling and mismeasurement in terms of these two components.

3.4.1 Imprecise Labels and Measurements. Imprecise measurements are measurements with a high level of uncertainty or a low level

⁹The predictive algorithm developed by PredPol used expectation-maximization techniques developed in [81], as noted in [82] and [35].

¹⁰We refer readers to Kasy and Abebe [64], which puts forward a formal notion of *power* as it relates to designing the objective function.

¹¹These tools are further discussed in [30, 47].

of significance. A measurement can be precise and inaccurate if it tends to detect values to a certain level of precision, but those values are consistently biased. A measurement is imprecise when its measurement exhibits significant random error.

Example: Portfolio Optimization. After Harry Markowitz put forward a method to optimally diversify portfolios in such a way that maximizes expected return [75, 76], practitioners who tried applying the method found that it under-performed human decisions and even simple equal-weighting methods [58]. In an article aptly titled “The Markowitz Optimization Enigma: Is ‘Optimized’ Optimal?” [77] it was shown that, systematically, assets with the highest measurement error were also those that had the highest expected return. Jorion [60, p 68] described the issue: “A major drawback with the classical implementation of mean-variance analysis is that it completely ignores the effect of measurement error on optimal portfolio allocations ... optimization systematically overweights the assets with the highest estimation errors, hence overstates the true efficiency of the optimal portfolio.”

3.4.2 Biased Labels and Measurements. Biased labels and measurements are consistently inaccurate in a way that can systematically lead to different outcomes for different subjects of an optimization.

Example: Medical Risk Assessment. Obermeyer et al. [91] study a risk prediction algorithm that is used on 200 million Americans annually to optimally distribute health services. The objective function of the risk assessment algorithm trained the algorithm to predict future health *costs* rather than health *needs*, and due to a historical gap in access across race, the risk scores systematically recommended fewer Black patients receive medical resources. Notably, the paper invokes the algorithm’s *objective function* in framing its inquiry into the mechanism of bias: “An unusual aspect of our dataset is that we observe the algorithm’s inputs and outputs as well as its objective function, providing us a unique window into the mechanisms by which bias arises” [91, p 3].

3.5 Faulty Solution Steps

The steps involved in solving an optimization problem, from relaxations to iterative algorithmic solutions, involve assumptions and simplifications that ought to be dealt with carefully and notated precisely. In this section, we discuss two types of neglect in the steps involved with solving optimizations: *faulty use of relaxations* and *faulty use of approximation algorithms*. Our goal in this section is not to critique optimization research, which tends to precisely and clearly characterize the use of approximations, nor to critique all applications of optimization. More often, these areas of neglect can occur in private settings which emphasize speed and minimum viability [83, 98] in product development.

3.5.1 Faulty Use of Relaxations. When a relaxation is used for ease of computation or measurement, it can significantly change the optimization’s feasible set and objective function. As such, it is important to detail when relaxations are being used, and what assumptions are made in the process. When a problem is relaxed, its fundamental structure has changed. The solution set might be much bigger, or smaller, than the true optimization problem, for example. As such, calling a solution to a relaxed problem ‘optimal’ may be misleading without appropriately caveating. When relaxations are

purposefully concealed and a solution is described as *optimal*, such a practice is *faulty*. It postures as *legitimate* on false pretenses, even when its recommended choice could be misguided or non-optimal.

3.5.2 Faulty Use of Approximation Algorithms. It is common for operations researchers, mathematicians, and computer scientists to focus on a particular category of problem, come up with a way of approximating the solution, and prove that the approximate solution is within a certain error bound compared to the true solution. In optimization, maximization (minimization) problems are approximated with answers that perform below (above) the true optimal, in which case lower (upper) bounds are given. Approximations are applied carelessly when these guaranteed bounds are not given, and the performance guarantee or bound is not specified. Approximation algorithms are applied erroneously if they are derived for a certain type of problem (e.g., convex) and are applied on a different type of problem (e.g., non-convex). Such a behavior can be *faulty* when performance is purposefully over-stated, or uncertainty is concealed.

3.6 Using “Optimal” as Justification

The final problem identified in this paper arises when people assume that an *optimal decision* is an *ethically justified decision*. As optimization has entered common parlance, it is used to describe a variety of behaviors and strategies that aim to achieve a variety of ends. Optimization is sometimes invoked without specifying a particular decision or objective—an optimal workforce strategy or an optimal economic policy lack the kind of specificity that is required to understand if a particular strategy is morally or politically desirable. In this section, we discuss the justificatory use of optimization. Particular mechanisms by which optimization can be used to justify action are apparent in language (e.g., when a decision-makers calls their decision *optimal* without further context) and in cases where an optimization problem is retrofit to a decision that has already been made.

3.6.1 Implicit Subject-Object Relationship. The use of the verb “optimize,” when used in the sense we discuss in this paper, implies that there is a *decision* and a *goal*. Sometimes, when people use the word without specifying one of these constitutive components, the context can clarify: “optimize draft picks” might specify the decision and imply the goal, for example. When the decision and goal are indeterminate and a quantity is claimed to be *optimal*, however, we argue that this description is a category error and its use is a primary suspect for the sort of neglect we are describing. Especially in high-stakes settings, where the goals of a certain decision might be contested, such as admissions, sentencing, and urban policy, using the word *optimal* to describe a particular decision may be a faulty justification or optic maneuver.

3.6.2 Reverse Engineering. Even if an optimization model is defined and solved indicating that a particular decision is optimal, such a process does not morally justify making that decision. This point can be demonstrated by considering two optimizations, nearly identical, except one switches “minimize” for “maximize.” Both yield *optimal solutions* that have diametrically opposite utility performance. The best possible and worst possible strategies for a given

optimization can both be said to be optimal strategies, if the optimization is tweaked. A broader and perhaps deeper point is that *any decision can be framed as the solution to an optimization problem*. By limiting the feasible set enough, and framing the objective in just the right way, any number of optimization problems can yield any decision as its optimal solution. Using optimal as justification, then, does not translate to ethical justification.

4 DISCUSSION AND SUGGESTIONS

Optimization is a paradigm used to aid and support decision-making in potentially high-stakes situations. It offers a systematic way of balancing interests and considerations as objectives or constraints relevant to a particular decision. The required design choices — including which interests and considerations to model, and how to model them — are often guided by norms, commitments and priorities. For example, a social media company might use optimization to recommend content in a way that attracts attention and engagement from users. Or, a government might use optimization to set tariffs in a way that promotes growth while satisfying domestic workers and unions. Using optimization in a particular domain requires making choices not just about which factors to include, but also about which factors to omit. Much of the context around a decision is omitted in order to model the decision as an optimization. Accordingly, when an optimization is “solved” and a course of action is selected, this does not suggest that the action is morally or ethically or socially good, *all things considered*. Instead, the so-called ‘optimal’ solution satisfies a very specific set of considerations: namely, those that are specified and introduced through the process of modeling, measuring, computing and applying the paradigm.

In this paper, we put forward a framework for understanding the normative commitments that are contained, either implicitly or explicitly, in the use of mathematical optimization. To do so, we focused on the steps involving abstraction and assumption, many of which are choices that reflect practitioners’ priorities and values. We then attempted to excavate some of the assumptions, areas of neglect, and emergent issues that can arise along the various components of an optimization’s life cycle.

As researchers are reckoning with the potential for bias, unfairness, inequality, opacity, and other ethical issues in machine learning and dataset creation, a similar line of inquiry is needed to scrutinize the normative (and epistemic) elements of optimization. A number of research questions in this line of inquiry remain, including: which norms, commitments, and priorities might constitute appropriate or legitimate optimization?

These and other considerations around optimization and optimal behavior are open questions that, we hope, will motivate a program of future research. There are a number of directions for work, some of which stem directly from the present paper’s framework.

Empirical work. The use of optimization as a paradigm for decision-making requires constructing an abstract model and using it to make decisions in real-world contexts. Empirical studies about optimization’s use and impacts may find particular patterns, conventions, and mechanisms through which optimization causes social harms (or, conversely, benefits). Sometimes, simply observing and documenting the components of an optimization (e.g., the objectives) can

provide useful empirical insight about an individual or institution’s operations. For example, the fact that a social media company designs features to optimize for user engagement and attention might reveal pressing and legitimate issues about its products.

Available traces and artifacts pointing to the use of optimization are crucial for empirical study and, ultimately, conversations about appropriate use. To that end, *documentation* can play a role in making optimization use transparent and accountable. The work on documenting models [80], datasets [43], and reinforcement learning models [45] might hold lessons for optimization. These studies can ultimately enable audits, new standards, best practices, and other mechanisms for accountability.

Conceptual work. This paper provides evidence that optimization, as a paradigmatic technical tool, encodes certain normative commitments. Theorizing about whether and when these commitments are justified remains a promising direction for further inquiry. Normative analysis of this sort would benefit from careful consideration of a number of contextual parameters about the use of optimization: By whom? Of what? When? Where? For what end? It would also benefit from an aggregate and system-wide framing on optimization and its political implications. For instance, it seems evident that private corporations tend to optimize for profit, or a proxy thereof. In light of this, how can markets, institutions and processes improve the use of optimization and aim technological tools towards social ends?

ACKNOWLEDGMENTS

We thank the Digital Life Initiative (DLI), the AI, Policy and Practice Initiative (AIPP), the Cornell Tech Emma Pierson and Nikhil Garg (CTEN) Lab, and the Cornell University Information Science Colloquium for providing remarks and suggestions. We’d like to acknowledge Nikhil Garg, Yonatan Mintz, Eugene Bagdasaryan, Jon Kleinberg, Karen Levy, and David Williamson for providing invaluable feedback.

The authors are grateful for awards from the US National Science Foundation, NSF CNS-1704527 and the John D. and Catherine T. MacArthur Foundation, which generously supported this work.

REFERENCES

- [1] William J Abernathy, Nicholas Baloff, John C Hershey, and Sten Wandel. 1973. A three-stage manpower planning and scheduling model—a service-sector example. *Operations Research* 21, 3 (1973), 693–711.
- [2] Michele Aghassi and Dimitris Bertsimas. 2006. Robust game theory. *Mathematical programming* 107, 1 (2006), 231–273.
- [3] Douglas S Altner, Anthony C Rojas, and Leslie D Servi. 2018. A two-stage stochastic program for multi-shift, multi-analyst, workforce optimization with multiple on-call options. *Journal of Scheduling* 21, 5 (2018), 517–531.
- [4] Edward G Anderson Jr. 2001. The nonstationary staff-planning problem with business cycle and learning effects. *Management Science* 47, 6 (2001), 817–832.
- [5] Rob Ashmore, Radu Calinescu, and Colin Paterson. 2021. Assuring the machine learning lifecycle: Desiderata, methods, and challenges. *ACM Computing Surveys (CSUR)* 54, 5 (2021), 1–39.
- [6] Robert U Ayres and Allen V Kneese. 1969. Production, consumption, and externalities. *The American economic review* 59, 3 (1969), 282–297.
- [7] KR Balachandran and Ralph E Steuer. 1982. An interactive model for the CPA firm audit staff planning problem with multiple objectives. *Accounting Review* 57, 1 (1982), 125–140.
- [8] Karen Barad. 2003. Posthumanist performativity: Toward an understanding of how matter comes to matter. *Signs: Journal of women in culture and society* 28, 3 (2003), 801–831.
- [9] William J Baumol. 1972. On taxation and the control of externalities. *The American Economic Review* 62, 3 (1972), 307–322.

- [10] Amir Beck and Marc Teboulle. 2003. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters* 31, 3 (2003), 167–175.
- [11] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. 2019. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences* 116, 32 (2019), 15849–15854.
- [12] Aharon Ben-Tal and Arkadi Nemirovski. 2002. Robust optimization–methodology and applications. *Mathematical programming* 92, 3 (2002), 453–480.
- [13] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. 2011. Theory and applications of robust optimization. *SIAM review* 53, 3 (2011), 464–501.
- [14] Hans-Georg Beyer and Bernhard Sendhoff. 2007. Robust optimization—a comprehensive survey. *Computer methods in applied mechanics and engineering* 196, 33–34 (2007), 3190–3218.
- [15] Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, et al. 2021. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. AAAI/ACM, Virtual Event, USA, 401–413.
- [16] Maurizio Bielli, Massimiliano Caramia, and Pasquale Carotenuto. 2002. Genetic algorithms in bus network optimization. *Transportation Research Part C: Emerging Technologies* 10, 1 (2002), 19–34.
- [17] Fischer Black and Robert Litterman. 1992. Global portfolio optimization. *Financial analysts journal* 48, 5 (1992), 28–43.
- [18] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. 2004. *Convex optimization*. Cambridge university press, Cambridge.
- [19] Gary W Caldwell, David M Ritchie, John A Masucci, William Hageman, and Zhengyin Yan. 2001. The new pre-clinical paradigm: compound optimization in early and late phase drug discovery. *Current topics in medicinal chemistry* 1, 5 (2001), 353–366.
- [20] Yang Trista Cao and Hal Daumé III. 2021. Toward Gender-Inclusive Coreference Resolution: An Analysis of Gender and Bias Throughout the Machine Learning Lifecycle. *Computational Linguistics* 47, 3 (2021), 615–661.
- [21] Yeong-Jer Chen, AA Neuber, John Mankowski, JC Dickens, Magne Kristiansen, and Richard Gale. 2005. Design and optimization of a compact, repetitive, high-power microwave system. *Review of scientific instruments* 76, 10 (2005), 104703.
- [22] Ronald Harry Coase. 2013. The problem of social cost. *The journal of Law and Economics* 56, 4 (2013), 837–877.
- [23] James S Coleman. 1966. The possibility of a social welfare function. *The American Economic Review* 56, 5 (1966), 1105–1122.
- [24] A Feder Cooper and Ellen Abrams. 2021. Emergent Unfairness in Algorithmic Fairness-Accuracy Trade-Off Research. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. AAAI/ACM, Virtual Event, USA, 46–54.
- [25] A Feder Cooper, Solon Barocas, Christopher De Sa, and Siddhartha Sen. 2023. Variance, Self-Consistency, and Arbitrariness in Fair Classification. *arXiv preprint arXiv:2301.11562* 0, 0 (2023), 1–84.
- [26] Robert A Copeland. 2013. *Evaluation of enzyme inhibitors in drug discovery: a guide for medicinal chemists and pharmacologists*. John Wiley & Sons, Hoboken, New Jersey.
- [27] Richard Cornes and Todd Sandler. 1996. *The theory of externalities, public goods, and club goods*. Cambridge University Press, Cambridge.
- [28] Imma Curiel. 2013. *Cooperative game theory and applications: cooperative games arising from combinatorial optimization problems*. Vol. 16. Springer Science & Business Media, Maryland.
- [29] Jakša Cvitanic and Ioannis Karatzas. 1992. Convex duality in constrained portfolio optimization. *The Annals of Applied Probability* 4, 2 (1992), 767–818.
- [30] René BM De Koster, Andrew L Johnson, and Debjit Roy. 2017. Warehouse design and management. , 6327–6330 pages.
- [31] Lauro Ribeiro de Souza Neto, José Teófilo Moreira-Filho, Bruno Junior Neves, Rocio Lucia Beatriz Riveros Maidana, Ana Carolina Ramos Guimarães, Nicholas Furnham, Carolina Horta Andrade, and Floriano Paes Silva Jr. 2020. In silico strategies to support fragment-to-lead optimization in drug discovery. *Frontiers in chemistry* 8 (2020), 93.
- [32] Roel Dobbe, Thomas Krendl Gilbert, and Yonatan Mintz. 2021. Hard choices in artificial intelligence. *Artificial Intelligence* 300 (2021), 103555.
- [33] Vijay D'Silva, Mathias Payer, and Dawn Song. 2015. The correctness-security gap in compiler optimization. In *2015 IEEE Security and Privacy Workshops*. IEEE, USA, 73–87.
- [34] Lee Alan Dugatkin and Hudson Kern Reeve. 2000. *Game theory and animal behavior*. Oxford University Press on Demand, Oxford, United Kingdom.
- [35] Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway feedback loops in predictive policing. In *Conference on Fairness, Accountability and Transparency*. PMLR, New York City, USA, 160–171.
- [36] George I Evers and Mounir Ben Ghalia. 2009. Regrouping particle swarm optimization: A new global optimization algorithm with improved performance consistency across benchmarks. In *2009 IEEE International Conference on Systems, Man and Cybernetics*. IEEE, San Antonio, Texas, USA, 3901–3908.
- [37] Michele Fedrizzi and Silvio Giove. 2007. Incomplete pairwise comparison and consistency optimization. *European journal of operational research* 183, 1 (2007), 303–313.
- [38] Ludwik Finkelstein and MS Leaning. 1984. A review of the fundamental concepts of measurement. *Measurement* 2, 1 (1984), 25–34.
- [39] Jessie Finocchiaro, Roland Maio, Faidra Monachou, Gourab K Patro, Manish Raghavan, Ana-Andreea Stoica, and Stratis Tsirtsis. 2021. Bridging machine learning and mechanism design towards algorithmic fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Online, USA, 489–503.
- [40] Olivier Flückiger, Gabriel Scherer, Ming-Ho Yee, Aviral Goel, Amal Ahmed, and Jan Vitek. 2017. Correctness of speculative optimizations with dynamic deoptimization. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–28.
- [41] Virginie Gabrel, Cécile Murat, and Aurélie Thiele. 2014. Recent advances in robust optimization: An overview. *European journal of operational research* 235, 3 (2014), 471–483.
- [42] Rolando Garcia, Vikram Sreeakanti, Neeraja Yadwadkar, Daniel Crankshaw, Joseph E Gonzalez, and Joseph M Hellerstein. 2018. Context: The missing piece in the machine learning lifecycle. In *KDD CMI Workshop*, Vol. 114. ACM, London, UK, 1–4.
- [43] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92.
- [44] Arthur M Geoffrion. 1974. Lagrangean relaxation for integer programming. In *Approaches to integer programming*. Springer, Berlin, Germany, 82–114.
- [45] Thomas Krendl Gilbert, Sarah Dean, Nathan Lambert, Tom Zick, and Aaron Snoswell. 2022. Reward reports for reinforcement learning. *arXiv preprint arXiv:2204.10817* 0, 0 (2022), 1–47.
- [46] Jake Goldenfein, Deirdre K Mulligan, Helen Nissenbaum, and Wendy Ju. 2020. Through the Handoff Lens: Competing Visions of Autonomous Futures. *Berkeley Tech. LJ* 35 (2020), 835.
- [47] Yeming Gong and René de Koster. 2011. A review on stochastic models and analysis of warehouse operations. *Logistics Research* 3, 4 (2011), 191–205.
- [48] Niko A Grupen, Daniel D Lee, and Bart Selman. 2020. Low-bandwidth communication emerges naturally in multi-agent learning systems. *arXiv preprint arXiv:2011.14890* 0, 0 (2020), 1–10.
- [49] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. 2016. Cooperative inverse reinforcement learning. *Advances in neural information processing systems* 29 (2016), 1–9.
- [50] Peter Hammerstein and Reinhard Selten. 1994. Game theory and evolutionary biology. *Handbook of game theory with economic applications* 2 (1994), 929–993.
- [51] Miguel A Hernán and Stephen R Cole. 2009. Invited commentary: causal diagrams and measurement bias. *American journal of epidemiology* 170, 8 (2009), 959–962.
- [52] Daniel P Heyman and Matthew J Sobel. 2004. *Stochastic models in operations research: stochastic optimization*. Vol. 2. Courier Corporation, Mineola, New York.
- [53] Zoë Hitzig. 2020. The normative gap: mechanism design and ideal theories of justice. *Economics & Philosophy* 36, 3 (2020), 407–434.
- [54] Zoë Hitzig, Lily Hu, and Salomé Viljoen. 2019. The technological politics of mechanism design. *University of Chicago Law Review* 87, 1 (2019), 95–108.
- [55] Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Virtual, 560–575.
- [56] William S Isaac. 2017. Hope, hype, and fear: the promise and potential pitfalls of artificial intelligence in criminal justice. *Ohio St. J. Crim. L.* 15 (2017), 543.
- [57] Abigail Z Jacobs and Hanna Wallach. 2021. Measurement and fairness. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. ACM, Virtual, 375–385.
- [58] J. D. Jobson and Robert M Korkie. 1981. Putting Markowitz theory to work. *The Journal of Portfolio Management* 7, 4 (1981), 70–74. <https://doi.org/10.3905/jpm.1981.408816> arXiv:<https://arxiv.org/abs/https://jpm.pm-research.com/content/7/4/70.full.pdf>
- [59] Bjorn Johansson, Tamás Keviczky, Mikael Johansson, and Karl Henrik Johansson. 2008. Subgradient methods and consensus algorithms for solving convex optimization problems. In *2008 47th IEEE Conference on Decision and Control*. IEEE, Cancun, Mexico, 4185–4190.
- [60] Philippe Jorion. 1992. Portfolio optimization in practice. *Financial analysts journal* 48, 1 (1992), 68–74.
- [61] Cliff Joslyn. 2001. The semiotics of control and modeling relations in complex systems. *Biosystems* 60, 1-3 (2001), 131–148.
- [62] Mamoru Kaneko and Kenjiro Nakamura. 1979. The Nash social welfare function. *Econometrica: Journal of the Econometric Society* 47, 2 (1979), 423–435.
- [63] Adam Kasperski and Paweł Zieliński. 2006. An approximation algorithm for interval data minmax regret combinatorial optimization problems. *Inform.*

- Process. Lett.* 97, 5 (2006), 177–180.
- [64] Maximilian Kasy and Rediet Abebe. 2021. Fairness, equality, and power in algorithmic decision-making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Virtual, 576–586.
- [65] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, Washington DC, USA, 137–146.
- [66] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* 0, 0 (2016), 1–23.
- [67] José Antonio Larco, Rene De Koster, Kees Jan Roodbergen, and Jan Dul. 2017. Managing warehouse efficiency and worker discomfort through enhanced storage assignment decisions. *International Journal of Production Research* 55, 21 (2017), 6407–6422.
- [68] Benjamin Laufer. 2020. Compounding Injustice: History and Prediction in Carceral Decision-Making. *arXiv preprint arXiv:2005.13404* 0, 0 (2020), 1–89.
- [69] Benjamin Laufer, Emma Pierson, and Nikhil Garg. 2022. End-to-end Auditing of Decision Pipelines. In *ICML Workshop on Responsible Decision-Making in Dynamic Environments*. ACM, Baltimore, Maryland, USA, 1–7.
- [70] Cédric Leboucher, Hyo-Sang Shin, Rachid Chelouah, Stéphane Le Méneç, Patrick Siarry, Mathias Formoso, Antonios Tsourdos, and Alexandre Kotenkov. 2018. An enhanced particle swarm optimization method integrated with evolutionary game theory. *IEEE Transactions on Games* 10, 2 (2018), 221–230.
- [71] Michelle Seng Ah Lee and Jatinder Singh. 2021. Risk Identification Questionnaire for Detecting Unintended Bias in the Machine Learning Development Lifecycle. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) (AIES '21). Association for Computing Machinery, New York, NY, USA, 704–714. <https://doi.org/10.1145/3461702.3462572>
- [72] Kristian Lum and William Isaac. 2016. To predict and serve? *Significance* 13, 5 (2016), 14–19.
- [73] Zhi-Quan Luo. 2003. Applications of convex optimization in signal processing and digital communication. *Mathematical programming* 97, 1 (2003), 177–207.
- [74] Christoph E Mandl. 1980. Evaluation and optimization of urban public transportation networks. *European Journal of Operational Research* 5, 6 (1980), 396–404.
- [75] Harry M Markowitz. 1991. Foundations of portfolio theory. *The journal of finance* 46, 2 (1991), 469–477.
- [76] Harry M Markowitz and G Peter Todd. 2000. *Mean-variance analysis in portfolio choice and capital markets*. Vol. 66. John Wiley & Sons, Hoboken, New Jersey, USA.
- [77] Richard O Michaud. 1989. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial analysts journal* 45, 1 (1989), 31–42.
- [78] Smitha Milli, Luca Belli, and Moritz Hardt. 2021. From optimizing engagement to measuring value. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. FAccT, Virtual, 714–722.
- [79] Roger E Millsap and Howard T Everson. 1993. Methodology review: Statistical approaches for assessing measurement bias. *Applied psychological measurement* 17, 4 (1993), 297–334.
- [80] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. ACM, Atlanta, Georgia, 220–229.
- [81] George Mohler. 2014. Marked point process hotspot maps for homicide and gun crime prediction in Chicago. *International Journal of Forecasting* 30, 3 (2014), 491–497.
- [82] George O Mohler, Martin B Short, Sean Malinowski, Mark Johnson, George E Tita, Andrea L Bertozzi, and P Jeffrey Brantingham. 2015. Randomized controlled field trials of predictive policing. *Journal of the American statistical association* 110, 512 (2015), 1399–1411.
- [83] Dobrila Rancic Moogk. 2012. Minimum viable product and the importance of experimentation in technology startups. *Technology Innovation Management Review* 2, 3 (2012), 1–1.
- [84] Emanuel D Moss. 2021. The Objective Function: Science and Society in the Age of Machine Intelligence. *CUNY Thesis* 0, 0 (2021), 1–348. <https://jpm.pm-research.com/content/7/4/70>
- [85] Deirdre K Mulligan and Helen Nissenbaum. 2020. The Concept of Handoff as a Model for Ethical Analysis and Design. *The Oxford Handbook of Ethics of AI* 1, 1 (2020), 233.
- [86] Yehuda Naveh, Yossi Richter, Yaniv Altschuler, Donna L Gresh, and Daniel P Connors. 2007. Workforce optimization: Identification and assignment of professional workers using constraint programming. *IBM Journal of Research and Development* 51, 3.4 (2007), 263–279.
- [87] Michael J Neely. 2010. Stochastic network optimization with application to communication and queuing systems. *Synthesis Lectures on Communication Networks* 3, 1 (2010), 1–211.
- [88] John A Nelder and Roger Mead. 1965. A simplex method for function minimization. *The computer journal* 7, 4 (1965), 308–313.
- [89] Evdokia Nikolova. 2010. Approximation algorithms for reliable stochastic combinatorial optimization. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, Barcelona, Spain, 338–351.
- [90] Zhe Niu, Yulong Zhao, and Bian Tian. 2014. Design optimization of high pressure and high temperature piezoresistive pressure sensor for high sensitivity. *Review of Scientific Instruments* 85, 1 (2014), 015001.
- [91] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [92] Cathy O'neil. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, New York, NY, USA.
- [93] Carolina Osorio and Michel Bierlaire. 2013. A simulation-based optimization framework for urban transportation problems. *Operations Research* 61, 6 (2013), 1333–1345.
- [94] Si-sheng Ou-Yang, Jun-yan Lu, Xiang-qian Kong, Zhong-jie Liang, Cheng Luo, and Hualiang Jiang. 2012. Computational drug discovery. *Acta Pharmacologica Sinica* 33, 9 (2012), 1131–1140.
- [95] Samir Passi and Solon Barocas. 2019. Problem formulation and fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, Atlanta, Georgia, 39–48.
- [96] Andre F Perold. 1984. Large-scale portfolio optimization. *Management science* 30, 10 (1984), 1143–1160.
- [97] Diego Prada-Gracia, Sara Huerta-Yépez, and Liliana M Moreno-Vargas. 2016. Application of computational methods for anticancer drug discovery, design, and optimization. *Boletín Médico Del Hospital Infantil de México (English Edition)* 73, 6 (2016), 411–423.
- [98] Eric Ries. 2009. Minimum viable product: a guide. *Startup lessons learned* 3 (2009), 1.
- [99] Arindam Roy, Shamik Sural, Arun Kumar Majumdar, Jaideep Vaidya, and Vijayalakshmi Atluri. 2019. Enabling Workforce Optimization in Constrained Attribute-Based Access Control Systems. *IEEE transactions on emerging topics in computing* 9, 4 (2019), 1901–1913.
- [100] Nikolaos V Sahinidis. 2004. Optimization under uncertainty: state-of-the-art and opportunities. *Computers & Chemical Engineering* 28, 6-7 (2004), 971–983.
- [101] Hanan Samet. 1978. Proving the correctness of heuristically optimized code. *Commun. ACM* 21, 7 (1978), 570–582.
- [102] Gesualdo Scutari, Daniel P Palomar, Francisco Facchinei, and Jong-Shi Pang. 2010. Convex optimization, game theory, and variational inequality theory. *IEEE Signal Processing Magazine* 27, 3 (2010), 35–49.
- [103] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. ACM, Atlanta Georgia, 59–68.
- [104] Selena Silva and Martin Kenney. 2018. Algorithms, platforms, and ethnic bias: An integrative essay. *Phylon* (1960-) 55, 1 & 2 (2018), 9–37.
- [105] Audra Simpson. 2007. On ethnographic refusal: Indigeneity, voice and colonial citizenship. *Junctures: The Journal for Thematic Dialogue* 1, 9 (2007), 67–80.
- [106] Bismark Singh and Mark Eisner. na. A Brief History of Optimization and Mathematical Programming. <https://www.informs.org/Explore/History-of-O.R.-Excellence/O.R.-Methodologies/Optimization-Mathematical-Programming>. *Informs* na, na (na), na. [Online; accessed 17-April-2022].
- [107] Walter Sinnott-Armstrong. 2022. Consequentialism. In *The Stanford Encyclopedia of Philosophy* (Winter 2022 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Online.
- [108] Brian Cantwell Smith. 1985. The limits of correctness. *Acm Sigcas Computers and Society* 14, 1, 2, 3, 4 (1985), 18–26.
- [109] Mohammad Karim Sohrabi and Hossein Azgomi. 2020. A survey on the combined use of optimization methods and game theory. *Archives of Computational Methods in Engineering* 27, 1 (2020), 59–80.
- [110] Ian Sommerville. 2004. *Software Engineering* (7th edition ed.). Addison Wesley, Boston, MA, USA.
- [111] Renan Souza, Leonardo Azevedo, Vitor Lourenço, Elton Soares, Raphael Thiago, Rafael Brandão, Daniel Civitarese, Emilio Brazil, Marcio Moreno, Patrick Valduriez, et al. 2019. Provenance data in the machine learning lifecycle in computational science and engineering. In *2019 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS)*. IEEE, Denver, Colorado, USA, 1–10.
- [112] Harini Suresh and John Guttag. 2021. A framework for understanding sources of harm throughout the machine learning life cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, Virtual, 1–9.
- [113] Christian Tomani, Sebastian Gruber, Muhammed Ebrar Erdem, Daniel Cremers, and Florian Buettner. 2021. Post-hoc uncertainty calibration for domain drift scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Nashville, TN, USA, 10124–10132.
- [114] Vijay V Vazirani. 2001. *Approximation algorithms*. Vol. 1. Springer, Berlin, Germany.
- [115] Salomé Viljoen, Jake Goldenfein, and Lee McGuigan. 2021. Design choices: Mechanism design and platform capitalism. *Big data & society* 8, 2 (2021), 20539517211034312.

- [116] David P Williamson and David B Shmoys. 2011. *The design of approximation algorithms*. Cambridge university press, Cambridge.
- [117] Susanne Wruck, Iris FA Vis, and Jaap Boter. 2017. Risk control for staff planning in e-commerce warehouses. *International Journal of Production Research* 55, 21 (2017), 6453–6469.
- [118] Mike Zajko. 2022. Artificial intelligence, algorithms, and social inequality: Sociological contributions to contemporary debates. *Sociology Compass* 16, 3 (2022), e12962.
- [119] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. 2019. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*. PMLR, Long Beach, California, 7472–7482.
- [120] Xiaoyu Zheng, Joshua Deotte, Matthew P Alonso, George R Farquar, Todd H Weisgraber, Steven Gemberling, Howon Lee, Nicholas Fang, and Christopher M Spadaccini. 2012. Design and optimization of a light-emitting diode projection micro-stereolithography three-dimensional manufacturing system. *Review of Scientific Instruments* 83, 12 (2012), 125001.
- [121] Yongli Zhu, Chengxi Liu, Kai Sun, Di Shi, and Zhiwei Wang. 2018. Optimization of battery energy storage to improve power system oscillation damping. *IEEE Transactions on Sustainable Energy* 10, 3 (2018), 1015–1024.

A OPTIMIZATION DEFINITION

Optimization can be defined as a mathematical problem of the form specified in equation (1), where variables are defined on the (continuous) set of real numbers.

$$\begin{aligned} & \underset{x}{\text{maximize}} && f(x) \\ & \text{subject to} && g_i(x) \leq 0, \quad i = 1, 2, \dots, k \\ & && h_j(x) = 0, \quad j = 1, 2, \dots, l \end{aligned} \tag{1}$$

The **decision** is modeled as a vector of n variables $x \in \mathbb{R}^n$. x can only take values that satisfy the $k \geq 0$ inequality constraints and $l \geq 0$ equality **constraints**, $g_i(x) \leq 0$ and $h_j(x) = 0$, respectively. The **objective function** $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined on the *feasible set* dictated by the constraints. The solution is the decision value x^{OPT} which maximizes $f(x)$ out of all feasible values of x .