IEEE Control Systems Letters paper presented at
2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Flexible Optimization for Cyber-Physical and Human Systems

Andrea Simonetto[1]

***Abstract*— We study how to construct optimization problems whose outcome are sets of feasible, close-to-optimal decisions for human users to pick from, instead of a single, hardly explainable "optimal" directive. In particular, we explore two complementary ways to render convex optimization problems stemming from cyber-physical applications "flexible". In doing so, the optimization outcome is a trade off between engineering best and flexibility for the users to decide to do something slightly different. The first method is based on robust optimization and convex reformulations. The second one is stochastic and inspired from stochastic optimization with decision-dependent distributions.**

I. Introduction

Modern cyber-physical systems, such as the smart energy grid, are becoming tightly interlocked with the end users. Optimizing the operations of such systems is also being driven to the limits, by proposing personalized solutions to each user, e.g., to regulate their energy consumption. Sometimes we refer to these highly integrated systems as cyber-physical and human systems (CPHS) [1, Chapter 4D].

In this paper, we ask the question of whether we can optimize these systems by still allowing the end users to have a choice between different, yet reasonably similar, decisions. This becomes key in unlocking flexibility of the optimized decision to account for transparency and ease technology adoption. To fix the ideas on a concrete example, we could refer to optimizing a building heating control, where the set temperatures are determined by an algorithm. In this paper then, we study how to build algorithms that can deliver an allowed range of potentially good temperatures to the users to choose from *independently from each other*.

Human behavior and satisfaction modeling is a well-studied research area, and therefore optimizing a cyber-physical system with pre-trained or online-learned human models has received much attention (see, e.g., [1]–[7]). An active area of research in robotics and control is preference learning [8]–[10], whereby we try to capture users preferences among different actions. However, unlocking flexibility by delivering sets and *not single optimal solutions* to the users to choose from is not well explored, and mostly novel in optimization. In this paper, we were mainly inspired from the pioneering works [11], [12] which propose a set-delivering controller. Their analysis techniques stems from robust control and inverse optimization, which we will not use here since our setting is different.

In this paper, we propose the following main contributions,

[1] We propose a novel deterministic flexible optimization problem that can deliver to end users a set of feasible solutions to pick from. The problem is then solved by leveraging robust optimization reformulations;
[2] we propose a novel stochastic flexible optimization problem, which is less conservative and can fine tune the human-machine interaction. To solve this problem, we propose two primal-dual methods and prove their theoretical properties. To prove these results, we use tools from stochastic optimization with decision-dependent distributions [13]–[17].

The contributions yield two complementary views in flexible optimization and we finish by proposing a complete workflow, labeled Flex-O.

Numerical experiments showcase our theoretical development and their empirical performance.

II. Problem formulation

Let $f : \mathbb{R}^n \to \mathbb{R}$ be a convex cost, and let $X \subseteq \mathbb{R}^n$ be the convex feasible set defined as the intersection of m inequalities, $X := \{x \in \mathbb{R}^n | h(x) \leqslant 0\}$, for convex function $h : \mathbb{R}^n \to \mathbb{R}^m$. We model the problem we want to solve as a convex optimization problem,

$$\mathsf{P}_1 : \quad \min_{x \in \mathbb{R}^n} f(x), \qquad \text{subject to } h(x) \leqslant 0, \qquad (1)$$

where we partition the decision variable $x = [x_1 \in \mathbb{R}^{n_1}, \ldots, x_N \in \mathbb{R}^{n_N}] \in \mathbb{R}^n$ to highlight the presence of N users. For the sake of simplicity, we will let $n_i = 1$ without an over loss of generality.

Problem (1) has been used to model relevant cyber-physical systems, e.g., [18], and a variety of methods exist to find the optimal decisions. Here however, we wish to modify it to allow the users to have a choice: that is, we would like to assign to each user, not a single decision, but a set from which they can choose from independently from each other.

As described in [1, Chapter 4D], Problem (1) features a cyber-physical part: a system designer that wishes to optimize the decisions x, possibly subjected to physical constraints as $h(x) \leqslant 0$, and a human part: a set of users taking part in the decision making process and aiming at picking their best x_i, possibly in a selfish way. Our setting offers an alternative approach to learning preferences [8]–[10], since we do not want necessarily to learn users' wishes, just giving them a choice to pick from. Our setting offers also an alternative approach to multi-objective optimization, whereby one could compute the Pareto front and let the users decide a specific point. The latter workflow however has the disadvantage to have to compute the Pareto front, and to ensure that each user can decide their own solution independently.

*This work was partly supported by the Agence Nationale de la Recherche (ANR) with the project ANR-23-CE48-0011-01.

[1]A.S. is at the Unité de Mathématiques Appliquées, ENSTA Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France andrea.simonetto@ensta-paris.fr

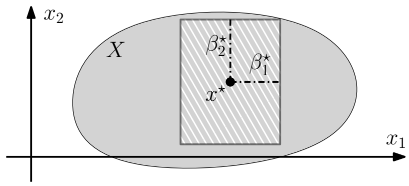

Fig. 1. Setting of the problem formulation in two dimensions with the best decision $x^\star$ and the optimal variations around it $\beta^\star$. We give to each user i the optimal set $[x_i^\star - \beta_i^\star, x_i^\star + \beta_i^\star]$.

We present now two ways that can be used to render Problem (1) flexible.

A. Deterministic approach

We start by looking at a deterministic approach. The intuition is to find the best decision $x^\star$ and an hyperbox of optimal size centered on it, so that all the points in the hyperbox are in the feasible set. This will allow us to assign to each user their component of $x^\star$ and the possible variations around it, determined by the size of the hyperbox.

To fix the ideas, Figure 2 depicts the intuition in a bi-dimensional setting. As we can see, depending on the nature of each user, we may allow for more or less flexibility.

Let us introduce scalar weights $w_i \geqslant 0$ and new scalar non-negative flexibility variables $\beta_i \geqslant 0$ for each user. We also introduce an uncertain variable $z_i \in [-1,1] \subset \mathbb{R}$ for each user. We collect β_i and z_i in column vectors as $\beta = (\beta_1, \ldots, \beta_n)^\top; z = (z_1, \ldots, z_n)^\top \in [-1,1]^n$, where $[-1,1]^n$ is the unitary hyperbox. We further let $y := [x^\top, \beta^\top]^\top$, $Y := \mathbb{R}^n \times \mathbb{R}^n_+$. We also define $[\mathbf{diag}(\beta)]$ as a $n \times n$ matrix with β as its diagonal. With this, we introduce the notation $\boldsymbol{A}(y) = [\mathbf{diag}(\beta)]$ and vector $\boldsymbol{b}(y) = x$, so that flexible hyperbox can be spanned with $\boldsymbol{A}(y)z + \boldsymbol{b}(y)$ for all $z \in [-1,1]^n$. Then, we render P_1 flexible by solving instead the following robust optimization problem,

$$\begin{aligned} \mathsf{P}_{2,\mathrm{d}}: \; y^\star \in \; & \operatorname*{argmin}_{y \in Y} \; g(y) := f(x) + \sum_{i=1}^n w_i \varphi_i(\beta_i), \\ & \text{subject to} \;\; h(\boldsymbol{A}(y)z + \boldsymbol{b}(y)) \leqslant 0, \; \forall z \in [-1,1]^n, \end{aligned} \tag{2}$$

for any convex function $\varphi_i : \mathbb{R} \to \mathbb{R}$. Typical examples would be $\varphi_i(\beta_i) = -\beta_i$, maximizing the flexibility for the user, and $\varphi_i(\beta_i) = -\beta_i + \frac{\epsilon}{2}\beta_i^2$ trading-off flexibility (small ϵ) and system performance (large ϵ).

Problem $\mathsf{P}_{2,\mathrm{d}}$ yields feasible solutions for all decisions in the hyperbox $X^\star := \prod_{i=1}^n X_i^\star$, with $X_i^\star := [x_i^\star - \beta_i^\star, x_i^\star + \beta_i^\star]$. The set $X_i^\star$ can be given to the user i for them to decide their optimal action, *independently of the other users*. Here, we are trading-off optimality of $x^\star$ while increasing the flexibility ensured by $\beta^\star$. That is, we are finding the optimal point $x^\star$ *and* the optimal variation around it $\beta^\star$, which still guarantees feasibility.

Problem (2) is a difficult optimization problem, for which however some approximation and reformulation procedures exist. We will discuss some of them in the following section.

Example 1: We consider the task of deciding the reference temperatures in different areas in an office building. Each group of users can set their thermostat in their office within an allowed range which we need to provide. Let $x \in \mathbb{R}^n$ be the temperature in n different areas, and $x_{\mathrm{ref}} \in \mathbb{R}^n$ be the engineering-best temperatures, which have been determined via an economic welfare trade-off. The problem we would like to solve can be the following one,

$$\begin{aligned} y^\star \in \; & \operatorname*{argmin}_{y \in Y} \quad \frac{1}{2}\|x\|^2 + \sum_{i=1}^n w_i \left(-\beta_i + \frac{\epsilon}{2}\beta_i^2\right), \\ & \text{subject to} \quad \left.\begin{array}{l} \|\boldsymbol{A}(y)z + \boldsymbol{b}(y) - x_{\mathrm{ref}}\|^2 \leqslant \gamma, \\ D\left(\boldsymbol{A}(y)z + \boldsymbol{b}(y)\right) \leqslant e, \end{array}\right\} \forall z \in [-1,1]^n. \end{aligned} \tag{3}$$

Here, the cost represents the wish to pick the smallest possible temperature, while the constraints impose a limited deviation with x_{ref} via a nonnegative scalar γ, and some additional affine constraints $D \in \mathbb{R}^{c \times n}, e \in \mathbb{R}^c$. The latter ones impose additional temperature bounds, and the fact that close-by areas cannot have very different temperatures. $\diamond$

B. Stochastic approach

In order to refine, and possibly render Problem (2) less conservative, we introduce a stochastic variant. Here, we assume that the users, *given a certain optimal decision $x_i^\star$ and allowed variation $\beta_i^\star$*, they pick a variable in the optimal set, say x_i, with a certain probability.

In this case, by introducing a nonnegative scalar $\delta \in [0,1]$, Problem (2) can be formulated as a chance-constrained decision-dependent non-convex problem as,

$$\begin{aligned} \mathsf{P}_{2,\mathrm{s}}: \; y^\star \in \; & \operatorname*{argmin}_{y \in Y} \quad g(y), \\ & \text{subject to} \quad \mathbb{P}_{z \sim \mathcal{D}(y)}\left[h(\boldsymbol{A}(y)z + \boldsymbol{b}(y)) > 0\right] \leqslant \delta, \end{aligned} \tag{4}$$

where $\mathbb{P}[\cdot]$ is the probability of a certain event, and $\mathcal{D}(y)$ is the decision-dependent distribution from which z is drawn from, whose support we assume is $[-1,1]^n$.

The chance constraint in Problem (4) can be conservatively rewritten in a convex-in-$h(y)$ form, by employing several bounds. Here, for reasons that will be clear later, we crucially need a smooth reformulation and we employ the Chernoff's bound. This allows one to write, for any scalar $u > 0$

$$\begin{aligned} & \mathbb{P}_{z \sim \mathcal{D}(y)}\left[h(\boldsymbol{A}(y)z + \boldsymbol{b}(y)) > 0\right] \leqslant \delta \Longleftarrow \\ & \mathop{\mathbb{E}}_{z \sim \mathcal{D}(y)} \left(\exp[u\, h_j(\boldsymbol{A}(y)z + \boldsymbol{b}(y))]\right) \leqslant \delta, \text{ for } j = 1, \ldots, m, \end{aligned}$$

see for instance [19]. With this in place, Problem (4) can be innerly approximated as,

$$\begin{aligned} \mathsf{P}_{3,\mathrm{s}}: \; y^\star \in \; & \operatorname*{argmin}_{y \in Y} \quad g(y), \\ & \text{subject to} \\ & \mathop{\mathbb{E}}_{z \sim \mathcal{D}(y)} \left(\exp[u\, h_j(\boldsymbol{A}(y)z + \boldsymbol{b}(y))]\right) \leqslant \delta, \text{ for } j = 1, \ldots, m. \end{aligned} \tag{5}$$

Note that, contrary to the case of decision-independent distribution, the constraints in (5) are still non-convex in y. Finally, we can write Problem (5) by its minimax formulation,

$$\begin{aligned} \mathsf{P}_{4,\mathrm{s}}: \; & \min_{y \in \mathbb{R}^{2n}} \max_{\lambda \in \mathbb{R}^m_+} \quad \varPhi(y, \lambda) := g(y) + \\ & \sum_j \lambda_j \left(\mathbb{E}_{z \sim \mathcal{D}(y)} \left(\exp[u\, h_j(\boldsymbol{A}(y)z + \boldsymbol{b}(y))]\right) - \delta\right) \end{aligned} \tag{6}$$

where the variables $\lambda \in \mathbb{R}^m_+$ are the Lagrangian multipliers associated to the constraints.

III. Solving The Flexible Problem

A. Robust optimization problem

Approximately (and conservatively) solving robust optimization problem (2) is possible, but the techniques depend on the type of constraints h one has. A recent framework, based on an extension of the Reformulation-Linearization-Technique, has been proposed by Bertsimas and coauthors in [20]. This framework is able to deal with any convex function $h(\boldsymbol{A}(y)z + \boldsymbol{b}(y)) \leqslant 0$ and any uncertainty set, and therefore it can be used here. The method is very general, so for specific problems, one would still use more standard techniques. For the sake of argument, we will show how one can transform Problem (3) into its convex worst-case reformulation, and let the reader explore different reformulations for their own settings. Problem (3) can be rewritten as,

$$(y^\star, s^\star) \in \underset{y\in Y, s\in\mathbb{R}^n}{\operatorname{argmin}} \quad \frac{1}{2}\|x\|^2 + \sum_{i=1}^n w_i \left(-\beta_i + \frac{\epsilon}{2}\beta_i^2\right), \tag{7}$$

$$\begin{aligned} \text{subject to} \quad & \|s\|^2 \leqslant \gamma, \\ & s_i \geqslant |\beta_i| + |x_i - x_{\mathrm{ref},i}|, \forall i \in \{1, n\} \\ & d_j x - e_j + \|[\mathbf{diag}(d_j)]\beta\|_1 \leqslant 0, \forall j \in \{1, c\} \end{aligned}$$

see for instance [21] and [22] for completeness, where d_j are the rows of D and e_j are the components of e.

The resolution of such convex reformulations, like (7), yields the optimal decision $x^\star$, as well as the optimal interval around it $\beta^\star$. As this may be conservative, we turn to the stochastic approach to refine it.

B. Stochastic optimization problem

We start our resolution strategy by rewriting (6) in the compact form,

$$\min_{y\in Y} \max_{\lambda\in\mathbb{R}^m_+} \quad \Phi(y,\lambda) := \mathbb{E}_{z\sim\mathcal{D}(y)}[\phi(y,\lambda,z)]. \tag{8}$$

Problem (8) is a stochastic saddle-point problem with decision-dependent distributions, which is in general non-convex and intractable in practice since one would need a full (local) characterization of $\mathcal{D}$. For this class of problems, since the optimizers are out of reach, one is content to find equilibrium points, as the points that are optimal w.r.t. the distribution they induce. In particular, one would start by assuming that the search space in y and λ is compact. In our case, this would be a reasonable approximation for y, since we can get an educated guess of a bounded search space by solving the deterministic problem (2) first. For λ, that would amount at clipping the multipliers, which is also a reasonable practice in convex and non-convex problems [23], [24]. With this in place, we let the search space for y be $\mathcal{Y} \subset Y$ and $\lambda \in \mathcal{M} \subset \mathbb{R}^m_+$. Then, one searches for equilibrium points, such that,

$$\bar{y} \quad \in \quad \arg\min_{y\in\mathcal{Y}} \left\{ \max_{\lambda\in\mathcal{M}} \mathbb{E}_{z\sim\mathcal{D}(\bar{y})}[\phi(y,\lambda,z)] \right\} \tag{9}$$

$$\bar{\lambda} \quad \in \quad \arg\max_{\lambda\in\mathcal{M}} \left\{ \min_{y\in\mathcal{Y}} \mathbb{E}_{z\sim\mathcal{D}(\bar{y})}[\phi(y,\lambda,z)] \right\}. \tag{10}$$

For continuous convex (in y) concave (in λ) uniformly in z functions ϕ, as in our case, assuming compactness of the sets $\mathcal{Y}$ and $\mathcal{M}$, as well as a continuous distributional map $\mathcal{D}$ under Wasserstein-1 distance W_1, then we know that the set of equilibrium points is nonempty and compact [15, Thm 2.5]. Consider further the following requirements.

Assumption 1: (a) Function ϕ is continuously differentiable over $Y \times \mathbb{R}^m_+$ uniformly in z, as well as μ-strongly-convex-strongly-concave (respectively in y and λ) for all z.

(b) The stochastic gradient map $\psi(y,\lambda,z) := (\nabla_y\phi(y,\lambda,z), -\nabla_\lambda\phi(y,\lambda,z))$ is jointly L-Lipschitz in (y,λ) and separately in z.

(c) The distribution map $\mathcal{D}$ is ε-Lipschitz with respect to the Wasserstein-1 distance W_1, i.e.,

$$W_1(\mathcal{D}(y), \mathcal{D}(y')) \leqslant \varepsilon\|y - y'\|, \quad \forall y \in Y.$$

$\diamond$

Strong convexity is ensured by properly defining the cost function, as well as φ_i, while differentiability holds thanks to the use of Chernoff's bound. If the engineering function is just convex, a regularization may be added. Strong concavity can be achieved by adding the dual regularization term $-\frac{\nu}{2}\|\lambda\|^2$, $\nu \geqslant \mu$, as in [25]. The various Lipschitz assumptions are mild (the ones on the gradient map hold trivially under compactness of the sets $\mathcal{Y}, \mathcal{M}$ and compact support for $\mathcal{D}$ as assumed).

With these assumptions in place, one can show that the distance between equilibrium points and optimal points of the original problem (8) is upper bounded by the constant of the problem and therefore solving for the former is a proxy for finding good approximate saddle points for the latter. Furthermore for $\varepsilon L/\mu < 1$, then the equilibrium point is unique [15, Thm 2.10]. To find such unique equilibrium point $(\bar{y}, \bar{\lambda})$ we employ a stochastic primal-dual method by generating a sequence of points $\{y_k, \lambda_k\}$, $k = 0, 1, \ldots,$ as,

$$z_k \quad \sim \quad \mathcal{D}(y_k), \tag{11a}$$

$$y_{k+1} \quad = \quad \mathsf{P}_{\mathcal{Y}}\left[y_k - \eta\nabla_y\phi(y_k,\lambda_k,z_k)\right], \tag{11b}$$

$$\lambda_{k+1} \quad = \quad \mathsf{P}_{\mathcal{M}}\left[\lambda_k + \eta\nabla_\lambda\phi(y_k,\lambda_k,z_k)\right], \tag{11c}$$

with step size $\eta > 0$ and projection operator $\mathsf{P}[\cdot]$. The stochastic gradient obtained by drawing z_k from the distribution generated at y_k is unbiased at k. Furthermore, we assume (as usual in the stochastic setting) that, for any given y, λ:

$$\mathop{\mathbb{E}}_{\omega\sim\mathcal{D}(y)}[\|\nabla_y\phi(y,\lambda,\omega) - \mathop{\mathbb{E}}_{z\sim\mathcal{D}(y)}[\nabla_y\phi(y,\lambda,z)]\| \leqslant \frac{\sigma}{\sqrt{2}}, \tag{12a}$$

$$\mathop{\mathbb{E}}_{\omega\sim\mathcal{D}(y)}[\|\nabla_\lambda\phi(y,\lambda,\omega) - \mathop{\mathbb{E}}_{z\sim\mathcal{D}(y)}[\nabla_\lambda\phi(y,\lambda,z)]\| \leqslant \frac{\sigma}{\sqrt{2}}, \tag{12b}$$

for a nonnegative constant σ.

Then we can derive the following result.

Theorem 1: Let $p = [y^\top, \lambda^\top]^\top$. Let Assumption 1 and the stochastic setting (12) hold. Assume $\frac{\varepsilon L}{\mu} < 1$ and pick the step size η as $\eta \in \left(0, \frac{2(\mu-\varepsilon L)}{L^2(1-\varepsilon^2)}\right)$. Then, the primal-dual method in Eq. (11) generates a sequence of points $\{y_k, \lambda_k\}$, such that in total expectation,

$$\limsup_{k\to\infty} \mathbb{E}[\|p_k - \bar{p}\|] = \frac{\eta\sigma}{1-\varrho},$$

with $\varrho := \sqrt{1-2\eta\mu+\eta^2L^2}+\eta\varepsilon L < 1$. ◇

Proof: Consider the deterministic method,

$$\check{y}_{k+1} \quad = \quad \mathsf{P}_{\mathcal{Y}}\left[\check{y}_k - \eta\mathbb{E}_{z\in\mathcal{D}(\check{y}_k)}\nabla_y\phi(\check{y}_k,\check{\lambda}_k,z)\right], \qquad (13a)$$

$$\check{\lambda}_{k+1} \quad = \quad \mathsf{P}_{\mathcal{M}}\left[\check{\lambda}_k + \eta\mathbb{E}_{z\in\mathcal{D}(\check{y}_k)}\nabla_\lambda\phi(\check{y}_k,\check{\lambda}_k,z)\right], \qquad (13b)$$

which is the deterministic version of (11), where we have substituted the stochastic gradients with their expectation. Under $\varepsilon L/\mu < 1$, for [15, Prop. 2.12], the fixed point of (13) is the unique equilibrium point $(\bar{y},\bar{\lambda})$. Compactly write (13) as, $\check{p}_{k+1} = \mathcal{G}(\check{p}_k;\check{p}_k)$, where we have indicated as $\check{p} = [\check{y}^\top, \check{\lambda}^\top]^\top$, and in the map $\mathcal{G}(\check{p}_k;\check{p}_k)$ the second argument represents the dependence of z on $\check{p}_k$. In the same way, we write (11) as, $p_{k+1} = \tilde{\mathcal{G}}(p_k;p_k)$. Then, for the Triangle inequality,

$$\|p_{k+1}-\bar{p}\| = \|\tilde{\mathcal{G}}(p_k;p_k)-\mathcal{G}(\bar{p};\bar{p})\| \leqslant \|\mathcal{G}(p_k;\bar{p})-\mathcal{G}(\bar{p};\bar{p})\|+ \\ + \|\mathcal{G}(p_k;p_k)-\mathcal{G}(p_k;\bar{p})\| + \|\tilde{\mathcal{G}}(p_k;p_k)-\mathcal{G}(p_k;p_k)\|. \quad (14)$$

For Assumption 1-(a) and (b), the gradient map ψ is μ-monotone and L-Lipschitz in (y,λ) [15]. As such, we can bound the first right-hand term as,

$$\|\mathcal{G}(p_k;\bar{p})-\mathcal{G}(\bar{p};\bar{p})\| \leqslant \sqrt{1-2\eta\mu+\eta^2L^2}\|p_k-\bar{p}\|. \quad (15)$$

For Assumption 1-(b-c) and [15, Lemma 2.9], the second right-hand term becomes,

$$\|\mathcal{G}(p_k;p_k)-\mathcal{G}(p_k;\bar{p})\| \leqslant \eta\varepsilon L\|p_k-\bar{p}\|. \qquad (16)$$

Passing now in total expectation, and defining $\varrho = \sqrt{1-2\eta\mu+\eta^2L^2}+\eta\varepsilon L$,

$$\mathbb{E}[\|p_{k+1}-\bar{p}\|] \leqslant \varrho\mathbb{E}[\|p_k-\bar{p}\|]+\mathbb{E}[\|\tilde{\mathcal{G}}(p_k;p_k)-\mathcal{G}(p_k;p_k)\|] \\ \underbrace{\leqslant}_{(12)} \varrho\mathbb{E}[\|p_k-\bar{p}\|]+\eta\sigma. \quad (17)$$

Finally, by iterating (17), the result is proven. ■

Theorem 1 tells us that if the step size is chosen sufficiently small, we can generate a sequence of points that approximates the unique equilibrium up to an error ball. The size of this ball depends on the variance of the stochastic gradient, as in many stochastic settings.

IV. A Human-Adapted Algorithm

Both solving the robust problem of Section (III-A) and the stochastic decision-dependent variant with (11) have their advantages and drawbacks. The robust program offers hard guarantee on feasibility but may be conservative. The primal-dual method can be closer to reality, however it achieves feasibility only asymptotically and it requires humans to give feedback at each iteration (since (11a) is achieved by asking humans to select their z_k), which can be unreasonable from an user-oriented perspective (e.g., if you are asked to adjust your thermostat multiple times).

A possible middle ground is to approximate the user's choice (11a) by a model, which we can run without asking the users to give feedback. Here, a pertinent notion of convergence will be provided in terms of the miss-match between the chosen model and the real distribution.

We approximate users as intelligent agents who respond to the requests by employing a best-response mechanism. In particular, we consider distributions as,

$$z_i \overset{d}{=} \left[\varPsi_i(x_i,\beta_i)+\xi_i\right]_{-1}^{+1}, \qquad \xi_i \sim \mathcal{D}_i, \qquad (18)$$

where $\mathcal{D}_i$ is a static distribution, $\overset{d}{=}$ indicates equality in distribution, and $[\cdot]_{-1}^{+1}$ indicates that the distribution is then truncated to have a $[-1,1]$ support. The reasoning behind (18) is that the users are selecting their best z_i depending on x_i,β_i, plus a static noise. The function $\varPsi_i(x_i,\beta_i):\mathbb{R}^2\to\mathbb{R}$ can be thought of as a function that encodes a user's optimization problem parametrized by (x_i,β_i). Models like (18) have been studied in economics and in game theory [16], [26] as well as in optimization [27] and control [12].

Let $\mathcal{D}_{\rm ms}(y)$ be the decision-dependent distribution induced by considering model (18). We assume that we are able to estimate the parameters of the approximate model (18) which is misspecified up to an error $B>0$ as follows.

Assumption 2: Cf. [16], [26]. The distribution $\mathcal{D}_{\rm ms}(y)$ is B-misspecified, in the sense that there exists a nonnegative scalar B, such that

$$W_1(\mathcal{D}_{\rm ms}(y),\mathcal{D}(y)) \leqslant B, \qquad \forall y\in\mathcal{Y}.$$

◇

Trivial bounds for B can be derived since $z\in[-1,1]^n$ for both true and estimated distributions, but better bounds can also be obtained with enough data on the users.

With this in place, we then use the misspecified model to run a deterministic model-based primal-dual method as,

$$y_{k+1} = \mathsf{P}_{\mathcal{Y}}\left[y_k - \eta\underset{z\sim\mathcal{D}_{\rm ms}(y_k)}{\mathbb{E}}\nabla_y\phi(y_k,\lambda_k,z)\right], \qquad (19a)$$

$$\lambda_{k+1} = \mathsf{P}_{\mathcal{M}}\left[\lambda_k + \eta\underset{z\sim\mathcal{D}_{\rm ms}(y_k)}{\mathbb{E}}\nabla_\lambda\phi(y_k,\lambda_k,z)\right]. \qquad (19b)$$

The advantage of Eq.s (19) is that they can be run *without* human intervention. For model (18), we need Assumption 1 to hold, which requires that the model is Lipschitz with respect to y as,

$$W_1(\mathcal{D}(y),\mathcal{D}(y')) = \|\varPsi(y)-\varPsi(y')\| \leqslant \varepsilon\|y-y'\|, \quad \forall y\in Y. \qquad (20)$$

For iterations (19), we have the following result.

Theorem 2: Let $p=[y^\top,\lambda^\top]^\top$. Let Assumption 1 and the misspecified setting of Assumption 2 hold. Assume $\frac{\varepsilon L}{\mu}<1$ and pick the step size η as $\eta\in\left(0,\frac{2(\mu-\varepsilon L)}{L^2(1-\varepsilon^2)}\right)$. Then, the model-based primal-dual method in Eq. (19) generates a sequence of points $\{y_k,\lambda_k\}$, such that deterministically,

$$\limsup_{k\to\infty}\|p_k-\bar{p}\| = \frac{\sqrt{2}\eta LB}{1-\varrho},$$

with $\varrho := \sqrt{1-2\eta\mu+\eta^2L^2}+\eta\varepsilon L < 1$. ◇

Proof: It follows by adapting the proof of Theorem 1 and by using Kantorovich and Rubinstein duality for the W_1 metric on the missmatch, and the Lipschitz assumption (Cf. Assumption 1-(b)). It is reported in [22]. ■

Theorem 2 says the iterations (19) converge up to an error bound, whose size is determined by the misspecification.

A. Warm-start and guarding step

Considering (19), we can now describe a final human-adapted algorithm. Consider the following workflow:

FleX-O: Flexible optimization algorithm

1) Solve the robust optimization problem (2) with the techniques of Section III-A;
2) Use the solution from 1) to warm-start the iterations (19) with an estimated model for T iterations;
3) (Optional: **Guarding step**) to make sure the solution is feasible $\forall z \in [-1,1]^n$, project the solution obtained from 2), say y_T, onto the robust feasible set, i.e., solve

$$\mathsf{P}_{3,\mathrm{d}}: \; y^\star \in \underset{y\in Y}{\operatorname{argmin}} \; \frac{1}{2}\left(\|x - x_T\|^2 + \varsigma\|\beta - \beta_T\|^2\right), \quad (21)$$
$$\text{subject to} \quad h(\boldsymbol{A}(y)z + \boldsymbol{b}(y)) \leqslant 0, \; \forall z \in [-1,1]^n.$$

with $\varsigma \geqslant 0$ and the techniques of Section III-A;
4) **Output:** an optimal set $[x_i^\star - \beta_i^\star, x_i^\star + \beta_i^\star]$ for each user i.

The Flex-O algorithm operates as follows. First, it solves the robust optimization problem to find a starting solution. Second, it uses this solution as a warm start from the primal-dual stochastic method (19). In T iterations, the stochastic method delivers a solution y_T, which comprises of the collection of sets $[x_{i,T}-\beta_{i,T}, x_{i,T}+\beta_{i,T}]$ to give to the users. Not all the points in the set will be feasible, but according to our probability model, the users will choose a solution in the set whose probability to be unfeasible is $\leqslant \delta$. If we want to be sure that we are feasible for any choice z, we can always do an optional guarding step, projecting y_T onto the robust feasible set. In this step, we can choose $\varsigma \geqslant 0$ to weigh more or less either x or flexibility β.

V. Numerical Results

Consider the setting of Example 1 with $n = 7$ for the cost:

$$g(y) = \frac{\epsilon_x}{2}\|x\|^2 + \sum_{i=1}^{7} w_i\left(-\beta_i + \frac{\epsilon_\beta}{2}\beta_i^2\right), \qquad (22)$$

with $\epsilon_x = 0.001$, $\epsilon_\beta = 0.01$, the weights w_i randomly drawn from the uniform distribution $\mathcal{U}(0.1, 1)$. We recall that we have set $y := [x^\top, \beta^\top]^\top$ and we express the temperatures, i.e., the units of y, in *degrees Celsius*. We also set x_{ref} randomly from a normal distribution $\mathcal{N}(19.5, 1.)$. We let $\gamma = 2n$. We consider a corridor with $n = 7$ offices, and therefore D is the matrix that represents the fact that two adjacent offices cannot have a very different temperature. In this case $c = 6$, and we let $e = [1, \ldots, 1]^\top$. For the primal-dual methods we let the dual regularization be $\nu = 0.01$, $u = 1.5$, and $\delta = 0.2$. By trial-and-error, we fix the step size at $\eta = 0.05$ for all the methods.

In Table I, we report the optimal solution of the problem obtained by the robust convex reformulation (7). We see that the variations around the optimal "imposed" temperature are zero for certain users. We then use this solution as a warm start for the primal-dual methods. In Figure 2, we report the evolution of the primal distance $\|y_k - \bar{y}\|$ for the baseline primal-dual (11) **[B-PD]**, for the misspecified-model-based primal-dual (19) **[MS-PD]**, and for Flex-O with a guarding

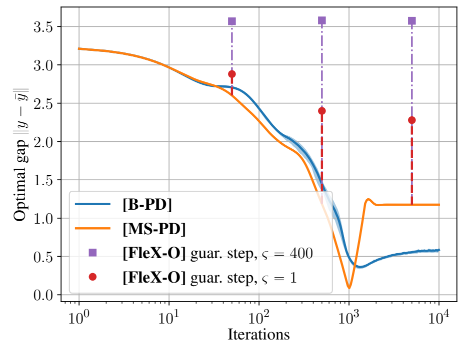

Fig. 2. Optimality gap vs. iterations for the considered algorithms. For **[B-PD]** we indicate mean and standard deviation over 50 realizations. Recall that **[B-PD]** cannot be implemented in practice, since it would require the users to select their best temperature at each iteration.

step (21) with $\varsigma = 1, 400$ at $T = 50, 500, 5000$, **[Flex-O]**. In all the cases, a nearly optimal value $\bar{y}$ is computed as follows. First we model the true *but unknown* distribution $\mathcal{D}(y)$ as,

$$z_i \overset{d}{=} \left[\xi_i + \begin{cases} \beta_i(19.0 - x_i) & \text{if } x_i \leqslant 19.0 \\ -\beta_i(x_i - 20.5) & \text{if } x_i \geqslant 20.5 \\ 0 & \text{otherwise} \end{cases}\right]_{-1}^{+1}, \quad (23)$$

with $\xi_i \sim \mathcal{N}(0, 0.1)$. Then $\bar{y}$ is computed by running the model-based primal-dual (19) on the true distribution (23). For the misspecified model, we instead take the deterministic,

$$z_i = \max\{-1, \min\{1, -\beta_i(x_i - 19.75)\}\}, \qquad (24)$$

which induces the distribution $\mathcal{D}_{\text{ms}}(y)$.

Model (23) indicates the natural propensity to accept the proposed optimal solution if it is within a suitable temperature range, while reacting if it falls outside. The strength of the reaction depends on the allowed variations.

In Figure 2, both **[B-PD]** and **[MS-PD]** reduce the optimal gap with respect to the stochastic equilibrium point $\bar{y}$, eventually reaching the error bound. For **[B-PD]** we have averaged the solution over 50 realizations. We also see the effect of the guarding step on **[Flex-O]**, which makes the solution less optimal w.r.t. $\bar{y}$, but feasible for all z's. More interestingly, if we observe the last iterate values in Table I, we see how the proposed primal-dual algorithms offer more flexibility (i.e., higher values of β). In the table $\langle \mathrm{CV}(z) \rangle$ indicate the average value over the last 100 iterations of the constraint violation,

$$\max_j\{\mathbb{E}_{z\sim\mathcal{D}(y)}\left[h_j(\boldsymbol{A}(y)z + \boldsymbol{b}(y))\right]\}.$$

From the results, we can appreciate the importance of having a good model for optimality. For **[Flex-O]**, we see how projecting onto the robust set trades-off flexibility with robustness. We note that even a few steps ~ 50 of the primal-dual (for both ς) can unlock more flexible solutions.

In summary: the robust problem (7) ensures feasibility, but it allows for flexibility to a limited number of users (the others have $\beta_i^\star = 0$). This may be considered not

TABLE I

SOLUTIONS OF DIFFERENT ALGORITHMS ON THE TEST PROBLEM.

Algorithm	Found x, found β
Robust (7)	$x^\star = [19.4, 19.4, 18.8, 18.3, 18.3, 18.3, 18.3]$ $\beta^\star = [1.0, 0.0, 0.5, 0.0, 1.0, 0.0, 1.0]$
True $\mathcal{D}(y)$ p.-d. (19)	$\bar{x} = [17.7, 17.7, 17.7, 17.7, 17.7, 17.7, 17.7]$ $\bar{\beta} = [1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0]$ $\langle \mathrm{CV}(z) \rangle = 0.023$
[MS-PD] (19)	$x = [17.3, 17.3, 17.3, 17.3, 17.3, 17.3, 17.3]$ $\beta = [1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0]$ $\langle \mathrm{CV}(z) \rangle = -0.450$
[B-PD] (11)	$x = [18.0, 17.9, 17.9, 17.9, 17.9, 17.9, 17.9]$ $\beta = [1.0, 0.9, 0.9, 0.9, 0.9, 0.9, 0.9]$ $\langle \mathrm{CV}(z) \rangle = -0.018$
[Flex-O], $\varsigma = 1$ at $T = 50$	$x = [19.1, 19.1, 18.6, 18.3, 18.3, 18.4, 18.3]$ $\beta = [0.7, 0.2, 0.3, 0.3, 0.6, 0.3, 0.5]$
at $T = 500$	$x = [18.8, 18.5, 18.3, 18.1, 18.1, 18.4, 18.1]$ $\beta = [0.3, 0.3, 0.3, 0.4, 0.5, 0.1, 0.5]$
at $T = 5000$	$x = [18.5, 18.3, 17.9, 17.7, 17.7, 18.2, 17.8]$ $\beta = [0.0, 0.0, 0.3, 0.4, 0.5, 0.0, 0.5]$
[Flex-O], $\varsigma = 400$ at $T = 50$	$x = [19.4, 19.3, 19.0, 18.7, 18.7, 18.7, 18.7]$ $\beta = [0.5, 0.3, 0.3, 0.3, 0.6, 0.3, 0.6]$
at $T = 500$	$x = [19.4, 19.3, 18.9, 18.7, 18.7, 18.7, 18.7]$ $\beta = [0.5, 0.3, 0.3, 0.4, 0.5, 0.4, 0.5]$
at $T = 5000$	$x = [19.4, 19.3, 18.9, 18.7, 18.7, 18.7, 18.7]$ $\beta = [0.5, 0.3, 0.3, 0.4, 0.5, 0.4, 0.5]$

acceptable. Then, **[MS-PD]** is doing better in allowing for greater flexibility (all β's are 1) but reducing the temperature with respect to $x^\star$. **[Flex-O]** with $\varsigma = 1$ or 400, with $T = 50$ seems to be the best compromise, allowing for shared flexibility, guaranteed feasibility, and a reasonable x.

VI. CONCLUSIONS

We have formulated flexible optimization problems that yield optimal decisions and per-user optimal variations around them. This allows the user to be given sets of possible decisions to take. The algorithms are based on robust optimization and stochastic decision-dependent distribution programs and they have been analyzed in theory and on a simple numerical example. Future research will look at possible links between this work and set-valued optimization [28] as well as decision-dependent distributionally robust optimization [29].

ACKNOWLEDGEMENT

The author thanks Killian Wood and Emiliano Dall'Anese for the careful read of an initial draft of the paper, and for the many useful suggestions for improvements.

REFERENCES

[1] A. M. Annaswamy, K. H. Johansson, G. J. Pappas, *et al.*, "Control for societal-scale challenges: Road map 2030," *IEEE Control Systems Society*, 2023.

[2] P. Chatupromwong and A. Yokoyama, "Optimization of charging sequence of plug-in electric vehicles in smart grid considering user's satisfaction," in *Proceedings of the IEEE International Conference on Power System Technology*, pp. 1 – 6, October 2012.

[3] R. Pinsler, R. Akrour, T. Osa, J. Peters, and G. Neumann, "Sample and Feedback Efficient Hierarchical Reinforcement Learning from Human Preferences," in *2018 IEEE ICRA*, pp. 596–601, 2018.

[4] D. D. Bourgin, J. C. Peterson, D. Reichman, T. L. Griffiths, and S. J. Russell, "Cognitive Model Priors for Predicting Human Decisions," in *Proceedings of the 36th International Conference on Machine Learning*, (Long Beach, California), 2019.

[5] A. Simonetto, E. Dall'Anese, J. Monteil, and A. Bernstein, "Personalized optimization with user's feedback," *Automatica*, vol. 131, 2021.

[6] Y. Zheng, B. Shyrokau, T. Keviczky, M. Al Sakka, and M. Dhaens, "Curve tilting with nonlinear model predictive control for enhancing motion comfort," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 4, pp. 1538–1549, 2021.

[7] A. D. Sadowska, J. M. Maestre, R. Kassing, P. J. van Overloop, and B. De Schutter, "Predictive control of a human–in–the–loop network system considering operator comfort requirements," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 53, no. 8, pp. 4610–4622, 2023.

[8] D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia, *Active preference-based learning of reward functions*. 2017.

[9] F. Bianchi, L. Piroddi, A. Bemporad, G. Halasz, M. Villani, and D. Piga, "Active preference-based optimization for human-in-the-loop feature selection," *European Journal of Control*, vol. 66, p. 100647, 2022.

[10] D. J. Hejna III and D. Sadigh, "Few-shot preference learning for human-in-the-loop RL," in *Conference on Robot Learning*, pp. 2014–2025, PMLR, 2023.

[11] M. Inoue and V. Gupta, ""Weak" Control for Human-in-the-Loop Systems," *IEEE Control Systems Letters*, vol. 3, no. 2, pp. 440–445, 2019.

[12] S. Shibasaki, M. Inoue, M. Arahata, and V. Gupta, "Weak control approach to consumer-preferred energy management," *IFAC-Papers online*, vol. 53, 2020.

[13] J. Perdomo, T. Zrnic, C. Mendler-Dünner, and M. Hardt, "Performative prediction," in *International Conference on Machine Learning*, pp. 7599–7609, PMLR, 2020.

[14] D. Drusvyatskiy and L. Xiao, "Stochastic optimization with decision-dependent distributions," *Mathematics of Operations Research*, vol. 48, no. 2, pp. 954–998, 2023.

[15] K. Wood and E. Dall'Anese, "Stochastic saddle point problems with decision-dependent distributions," *SIAM Journal on Optimization*, vol. 33, no. 3, pp. 1943–1967, 2023.

[16] L. Lin and T. Zrnic, "Plug-in performative optimization," *arXiv preprint arXiv:2305.18728*, 2023.

[17] Z. Wang, C. Liu, T. Parisini, M. M. Zavlanos, and K. H. Johansson, "Constrained optimization with decision-dependent distributions," *arXiv preprint arXiv:2310.02384*, 2023.

[18] J. A. Taylor, *Convex optimization of power systems*. Cambridge University Press, 2015.

[19] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal on Optimization*, vol. 17, no. 4, pp. 969–996, 2007.

[20] D. Bertsimas, D. den Hertog, J. Pauphilet, and J. Zhen, "Robust convex optimization: A new perspective that unifies and extends," *Mathematical Programming*, pp. 1–42, 2022.

[21] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*, vol. 28. Princeton university press, 2009.

[22] A. Simonetto, "Flexible optimization for cyber-physical and human systems (technical details)," *arXiv:2403.03847*, 2024.

[23] A. Nedić and A. Ozdaglar, "Approximate Primal Solutions and Rate Analysis for Dual Subgradient Methods," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1757 – 1780, 2009.

[24] T. Erseghe, "A Distributed and Maximum-Likelihood Sensor Network Localization Algorithm Based Upon a Nonconvex Problem Formulation," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 1, no. 4, pp. 247 – 258, 2015.

[25] J. Koshal, A. Nedić, and U. Y. Shanbhag, "Multiuser Optimization: Distributed Algorithms and Error Analysis," *SIAM Journal on Optimization*, vol. 21, no. 3, pp. 1046 – 1081, 2011.

[26] K. Wood, A. Zamzam, and E. Dall'Anese, "Solving decision-dependent games by learning from feedback," *arXiv:2312.17471*, 2023.

[27] P. Mohajerin Esfahani, S. Shafieezadeh-Abadeh, G. A. Hanasusanto, and D. Kuhn, "Data-driven inverse optimization with imperfect information," *Mathematical Programming*, vol. 167, pp. 191–234, 2018.

[28] A. A. Khan, C. Tammer, and C. Zalinescu, *Set-valued optimization*. Springer, 2016.

[29] F. Luo and S. Mehrotra, "Distributionally robust optimization with decision dependent ambiguity sets," *Optimization Letters*, vol. 14, pp. 2565–2594, 2020.