

Stochastic Programming Using Expected Value Bounds

Raphael Chinchilla and João P. Hespanha

Abstract—We address the problem of minimizing an expected value with stochastic constraints, known in the literature as stochastic programming. Our approach is based on computing and optimizing bounds for the expected value that are obtained by solving a deterministic optimization problem that uses the probability density function to penalize unlikely values for the random variables. The sub-optimal solution obtained through this approach has performances guarantees with respect to the optimal one, while satisfying stochastic and deterministic constraints.

We illustrate this approach in the context of three different classes of optimization problems: finite horizon optimal stochastic control, with state or output feedback; parameter estimation with latent variables; and nonlinear Bayesian experiment design.

By the means of several numerical examples, we show that our sub-optimal solution achieves results similar to those obtained with Monte Carlo methods with a fraction of the computational burden, highlighting the usefulness of this approach in real-time optimization problems.

I. INTRODUCTION

Optimization of an expected value, also called stochastic programming, appears in countless areas of applied probability and engineering. In optimal stochastic control, a dynamical system is subject to stochastic disturbances, and one wants to find the control that minimizes the expected value of the trajectory tracking error. In maximum likelihood estimation, unobserved variables may need to be integrated out through an expected value, to obtain the likelihood of the observed variables. In machine learning, training a neural network means finding the weights that best classify the expected value of a random variable.

Given a scalar function $V(\cdot)$ and a random vector D , the expected value of $V(D)$ can be lower and upper bounded, respectively, by the minimum and maximum values that $V(\cdot)$ takes over the support of D . Our first result in Section II shows how these very crude bounds can be improved by including information encoded in the probability density function (*pdf*) of D . In essence, we solve an optimization over the support of D that includes terms that penalize unlikely realizations for D . This means that we need to compute and solve optimality conditions — and therefore essentially compute derivatives and solve algebraic equations — rather than compute integrals.

The results in Section II actually define a family of bounds. Two instances of this family, which we call the additive and multiplicative bounds, are particularly useful. The first is more appropriate to problems where the cost function $V(\cdot)$ is polynomial, while the second one is more appropriate when the cost function is exponential. Both the additive and multiplicative bounds are parameterized by a scalar parameter

ϵ , which can itself be optimized. To guide the design of the bounds and select ϵ , we develop necessary and sufficient conditions with respect to ϵ that can be used to make sure that the additive and multiplicative bounds are finite.

Borrowing ideas from robust optimization, the bounds developed in Section II are used in Section III to compute approximate solutions to stochastic programming optimizations: Instead of minimizing an expected value subject to stochastic constraints, we minimize upper/lower bounds for the criterion subject to constraints on pessimistic/optimistic bounds for the stochastic constraints. For the lower bound, this leads to a minimization on an extended variable space; for the upper bound, it leads to a minmax problem.

In Section IV, we discuss three applications for our bounds. The first relates to finite-horizon stochastic optimal control, with either state feedback or output feedback. In the former case, the initial state is assumed known, but an expectation is needed over the realization of future disturbances. In the latter case, the initial state is unknown, and the expectation is taken with respect to a conditional distribution, given known realizations of past noisy measurements. Our approach can include stochastic constraints on the trajectory of the system, which we illustrate through a constraint on the final state.

The second application is related to Maximum Likelihood or Maximum a Posteriori estimation involving latent variables that cannot be measured [1]. These problems require the latent variables to be marginalized by an expectation that can be upper/lower bounded using the results from Section III.

The third application is in the area of Bayesian experiment design [2]–[4]. The goal is to optimize the values of experimental parameters to facilitate the estimation of unknown variables. Experiment design criteria typically involve taking expectations with respect to unknown variables, including the ones that need to be estimated. Also here, optimal experiment design can be performed by replacing expectations by bounds.

In the context of feedback control, all three applications discussed above typically need to be performed in real-time with limited computation, and benefit from the availability of bounds on how the approximate solution compares with the true optimum. It is in such scenarios that the approach proposed here is most attractive. In contrast, when computation is unlimited, Monte Carlo based methods can achieve arbitrarily accurate solutions to stochastic optimization problems as long as one uses a sufficiently large number of samples, and will thus eventually out-perform in accuracy the approach proposed here.

Related Work: Stochastic Programming has been an active area of research for the last 60 years, therefore a

complete overview of the literature is infeasible. Instead, we provide a brief overview of the most fundamental methods, some recent developments and how these relate to our work. We discuss separately four approaches: deterministic methods, stochastic methods, methods based on robust optimization and distributionally robust optimization .

Deterministic methods rely on computing the expected value using a numerical integration method such as Gauss-Kronrod [5], [6]. Since numerical integration is computationally infeasible for large problems, deterministic approximations of the expected value are often used. Common approximations include minimizing the truncated Taylor expansion of the expected value [7]–[9] and the Laplace and saddle-point approximations [10], [11]. A weakness of these methods is that they generally do not provide guarantees regarding how the solution found compares to the true optimum.

Stochastic optimization methods rely on some form of Monte Carlo sampling. These methods generally scale well and provide confidence intervals on the solutions. The most intuitive method is the Sample Average Approximation (SAA) (also known as Empirical Risk Minimization) [12]–[14], where the expected value is approximated by the empirical average obtained through sampling. Stochastic Gradient Descent (SGD) [2] is an easy to implementation and versatile alternative. The core idea of SGD is to directly draw samples of the gradient of the expected values, rather than using the gradient of the empirical mean to do gradient descent. Under appropriate assumptions, both SAA and SGD are guaranteed to converge to a (possibly local) minimum as the number of sample grows [14], but accurate results may require a very large number of samples, making these methods not suitable for real time applications.

The Scenario Approach approximately solves chance-constrained optimizations by sampling the constraints [15]–[17]. This approach guarantees constraint satisfaction with high probability. While the number of samples increases only logarithmically with the confidence parameter (usually denoted by β), it is also proportional to the dimensions of the optimization variables and inversely proportional to the risk parameter (usually denoted by ϵ). As a result, Scenario Approach may require many samples which can lead to high computational complexity. Moreover, while tight requirements on the confidence parameter β have a moderate impact in the number of samples, it typically also lead to more conservative results.

In robust optimization one minimizes for the worst possible perturbation, while guaranteeing some base level of performance [18], [19]. Robust optimization has been gaining popularity in recent years, for examples in fields such as Model Predictive Control [20] and Machine Learning [21]–[23]. Some new developments have also been made in numerical aspects, notably in [24], where the authors provide first and second order optimality conditions for minmax when the criteria is nonconvex on the minimization variable and nonconcave on the maximization variable. Robust optimization was traditionally not regarded as an approach to solve stochastic programming problems, but in the last decade some articles have connected the two areas, for instance [25]–[27].

At the intersection between robust and stochastic optimization lies distributionally robust optimization (DRO), where the objective is to minimize an expected value for the worst probability distribution within a set of admissible distributions. This set, called the ambiguity set, is often constructed from samples of the true distribution and its selection tries to balance between expressiveness (how rich is the information in the set) and tractability (how easy it is to solve the DRO). In [28], [29] the authors use the Wasserstein metric to construct the ambiguity set and show that for some classes of problem, the complexity of solving the associated DRO is similar to that of Sample Average Approximation. In [30]–[33] the ambiguity sets are constructed using sample statistics, such as mean, covariance and entropy. The DRO is reformulated into a minimization on a larger set of variables using tools from duality theory and convex optimization. When the ambiguity set is constructed to guarantee (with high probability) that it contains the true distribution, DRO can also be seen as a method to bound the true expected values, which can be used to solve stochastic programs. For a broader exposition on DRO we refer to [34] and the references within.

This paper expands in several ways the work reported in the conference paper [35]: We introduce a family of bounds, while the earlier paper considered only additive bound; we generalize the stochastic programming problem to include stochastic inequalities constraints; and we include applications to estimation and experiment design.

Notation: Given an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a random variable X and a scalar $x \in \mathbb{R}$, we denote by $\mathbb{P}(X \leq x)$ the probability measure of the set $\{\omega \in \Omega : X(\omega) \leq x\} \in \mathcal{F}$ and by $\mathbb{E}[X]$ the expected value of X . Given a measurable event $E \in \mathcal{F}$ with $\mathbb{P}(E) > 0$, we define conditional essential infimum and supremum by

$$\begin{aligned} \text{ess inf}[X \mid E] &= \sup\{x \in \mathbb{R} : \mathbb{P}(X \geq x \mid E) = 1\} \\ \text{ess sup}[X \mid E] &= \inf\{x \in \mathbb{R} : \mathbb{P}(X \leq x \mid E) = 1\}. \end{aligned}$$

Unconditional essential infimum and supremum are denoted simply by $\text{ess inf } X$ and $\text{ess sup } X$ and correspond to the case $E = \Omega$. The essential supremum and infimum relax the usual supremum and infimum by excluding sets of measure zero. One can informally think of them as $\sup_{x \in \mathcal{X}} x$ and $\inf_{x \in \mathcal{X}} x$ where \mathcal{X} is the support of X .

Given two random variables X, Y we use the notation $X \stackrel{\text{wpo}}{\leq} Y$ when $\mathbb{P}(X \leq Y) = 1$ and analogously for $\stackrel{\text{wpo}}{\geq}$, $<$, $>$.

II. BOUNDS ON AN EXPECTED VALUE

Given a random vector D taking values in $\mathcal{D} \subset \mathbb{R}^M$ and a scalar measurable function $V : \mathcal{D} \rightarrow \mathbb{R}$, the monotonicity of the expected value $\mathbb{E}[V(D)]$ provides the following basic bound

$$\text{ess inf } V(D) \leq \mathbb{E}[V(D)] \leq \text{ess sup } V(D). \quad (1)$$

The core idea of this section is to improve upon this crude bound by including information about D , for example, coming from its probability density function (*pdf*). To present our first result, we introduce the following terminology. Consider a

right-ordered group $G := (\mathcal{P}, \oplus)$ defined on a set $\mathcal{P} \subset \mathbb{R}$ for which the group operation \oplus satisfies the usual group properties of closure, associativity, existence of an identity element, and existence of inverse elements (which we denote using \ominus); as well as the right-ordered property

$$a \leq b \quad \Rightarrow \quad a \oplus c \leq b \oplus c, \quad \forall a, b, c \in \mathcal{P}$$

[36]. We say that $G := (\mathcal{P}, \oplus)$ is *distributive with respect to integration* (or *E-distributive for short*) if it is right-ordered and, for every random variable X taking values on \mathcal{P} , we have that

$$a \oplus \mathbb{E}[X] = \mathbb{E}[a \oplus X], \quad \forall a \in \mathcal{P}.$$

Theorem 1 (Bounds on an expected value): Consider an E-distributive group $G := (\mathcal{P}, \oplus)$, a random vector D taking values in $\mathcal{D} \subset \mathbb{R}^M$, and measurable functions $V, \alpha : \mathcal{D} \mapsto \mathcal{P}$. If $\mathbb{E}[V(D)]$ and $\mathbb{E}[-\alpha(D)]$ are finite, then

$$\text{ess inf } J(D) \leq \mathbb{E}[V(D)] \leq \text{ess sup } J(D) \quad (2)$$

where the function $J : \mathcal{D} \rightarrow \mathbb{R}$ is defined by

$$J(d) := V(d) \oplus \alpha(d) \oplus \mathbb{E}[-\alpha(D)]. \quad \square$$

Proof. We prove the upper bound, the proof for the lower bound can be obtained analogously. For every scalar $v \geq \text{ess sup } V(D) \oplus \alpha(D)$, we have that

$$\mathbb{P}(V(D) \oplus \alpha(D) \leq v) = 1,$$

by the definition of essential supremum. From the monotonicity of the expected value, we thus conclude that

$$\mathbb{E}[V(D) \oplus \alpha(D)] \leq v.$$

Since $\mathbb{E}[-\alpha(D)]$ is finite, we can use the right-ordered property of (\mathcal{P}, \oplus) to conclude that

$$\mathbb{E}[V(D) \oplus \alpha(D)] \oplus \mathbb{E}[-\alpha(D)] \leq v \oplus \mathbb{E}[-\alpha(D)]$$

and then the E-distributed property to obtain

$$\mathbb{E}[V(D) \oplus \alpha(D) \oplus -\alpha(D)] = \mathbb{E}[V(D)] \leq v \oplus \mathbb{E}[-\alpha(D)].$$

The upper bound then follows by taking an infimum on the right-hand side over the set of such scalars $v \geq \text{ess sup } V(D) \oplus \alpha(D)$. ■

The key idea of Theorem 1 is to improve upon (1) by including in $J(\cdot)$ terms that reduce the essential supremum and increase the essential infimum. To reduce the supremum, for example, one should select $\alpha(d)$ so that it is strongly negative (in the sense that $-\alpha(d)$ should be strongly positive) when $V(d)$ is large and while keeping $\mathbb{E}[-\alpha(D)]$ relatively small. In the remainder of the paper we mostly use two E-distributive groups G and associated functions α that achieve this for our applications of interest. Both bounds assume that D has a probability density function (*pdf*) that we denote by $p_D(\cdot)$.

Additive Bound: The E-distributive group $(\mathcal{P}, \oplus) = (\mathbb{R}, +)$ with the usual addition of reals, and $\alpha(d) = \epsilon \log p_D(d)$ with $\epsilon \in \mathbb{R}$, leads to

$$J(d, \epsilon) := V(d) + \epsilon \log p_D(d) + \epsilon \mathcal{H}_D, \quad (3)$$

where $\mathcal{H}_D := \mathbb{E}[-\log p_D(D)]$ is the differential entropy.

Multiplicative Bound: The E-distributive group $(\mathcal{P}, \oplus) = (\mathbb{R}_{>0}, \times)$ with the usual multiplication of positive reals, and $\alpha(d) = p_D(d)^\epsilon$ with $\epsilon \in \mathbb{R}$, leads to

$$J(d, \epsilon) := V(d) p_D(d)^\epsilon \mathcal{I}_D(\epsilon) \quad (4)$$

where $\mathcal{I}_D(\epsilon) := \mathbb{E}[p_D(D)^{-\epsilon}]$.

The functions J in (3) and (4) are not necessarily well defined on the measure zero set where $p_D(D) = 0$, but the value of J on such set is irrelevant, as it does not affect the value of the essential supremum or infimum in (2).

The key idea behind the additive bound is that unlikely values d for D will lead to a large negative value for $\log p_D(d)$ and reduce the value of $J(d)$. These unlikely values will contribute with a strong positive value in $-\log p_D(D)$, but precisely because they are unlikely, they will not increase $\mathcal{H}_D := \mathbb{E}[-\log p_D(D)]$ very much. Overall, this should thus decrease the supremum of $J(d)$ over \mathcal{D} to create a tighter bound. A similar reason can be used to justify the function α proposed for the multiplicative bound.

In Appendix A, we derive expressions for $\log p_D(d) + \mathcal{H}_D$ and $p_D(d)^\epsilon \mathcal{I}_D(\epsilon)$ for the Gaussian and for the uniform distributions.

Remark 1 (Bounds for conditional expectation): Theorem 1 can also be stated for conditional expectations, provided that the E-distributive property holds for the conditional expectation with probability one. In this case, the additive and multiplicative bounds should involve conditional *pdf*. ■

A. Selection of bound and ϵ

It is possible to establish necessary and sufficient conditions such that the additive and multiplicative bounds lead to non-trivial results, which are presented in Appendix C. Here, we present a corollary of those results that includes the sufficient conditions which, in practice, are the most useful in deciding which bounds to use. We require the following definition to present the corollary. Given a constant $\gamma > 0$ sufficiently small so that $\mathbb{P}(p_D(D) > \gamma) > 0$, we say that a measurable function $f(\cdot)$ is γ -essentially upper bounded if

$$\text{ess sup } [f(D) \mid p_D(D) > \gamma] < \infty,$$

γ -essentially lower bounded if

$$\text{ess inf } [f(D) \mid p_D(D) > \gamma] > -\infty,$$

and γ -essentially bounded if it is both γ -essentially upper and lower bounded. γ -essential boundedness is a much milder requirement than the usual notion of boundedness, as it allows functions to become very large (growing all the way to infinity) as long as the pdf becomes sufficiently small.

Corollary 1 (Sufficient conditions for finite bounds): Assume that $p_D(\cdot)$ is γ -essentially upper bounded and consider finite constants $\epsilon \in \mathbb{R}$ and $c \in (0, 1/\gamma)$ such that

$$\text{ess sup } [p_D(D) \mid p_D(D) > \gamma] \leq 1/c.$$

When $V(\cdot)$ is γ -essentially bounded, we have that

$$p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess inf} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] > \epsilon$$

$$\Rightarrow \text{ess inf} (V(D) + \epsilon \log p_D(D)) > -\infty.$$

and

$$p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess sup} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] < \epsilon$$

$$\Rightarrow \text{ess sup} (V(D) + \epsilon \log p_D(D)) < +\infty$$

Alternatively, when $\log V(\cdot)$ is γ -essentially bounded, we have that

$$p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess inf} \left[\frac{-\log V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] > \epsilon$$

$$\Rightarrow \text{ess inf} (\log V(D) + \epsilon \log p_D(D)) > -\infty.$$

and

$$p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess sup} \left[\frac{-\log V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] < \epsilon$$

$$\Rightarrow \text{ess sup} (\log V(D) + \epsilon \log p_D(D)) < +\infty. \quad \square$$

The first two implications in Corollary 1 involve $V(D)$ and are relevant for the additive bound, while the remaining ones involve $\log V(D)$ and the multiplicative bound.

Specifically, this result establishes that for the additive and multiplicative bounds to be non trivial (i.e., finite), it suffices to pick an ϵ such that $\log c p_D(D)$ dominates either $V(D)$ or $\log V(D)$, respectively. Therefore, which bound to use essentially depends on the rates of growth of $V(\cdot)$, $\log V(\cdot)$, and $\log p_D(\cdot)$. When both bounds have a finite value, we have observed that the approximations seems to be better when $V(\cdot)$ (or $\log V(\cdot)$) has roughly the same magnitude as $\log p_D(\cdot)$.

Among the values of ϵ that lead to a finite upper bound, the conservativeness of the bound can be minimized by selecting the value of $\epsilon \in \mathbb{R}$ that minimizes

$$\inf_{\epsilon \in \mathbb{R}} J^*(\epsilon), \quad J^*(\epsilon) := \text{ess sup} J(D, \epsilon) \quad (5)$$

with $J(d, \epsilon)$ as in (3) or (4). It turns out that such minimization over the scalar parameter ϵ is well-behaved as the function $J^*(\epsilon)$ in (5) has appropriate convexity properties, as noted in the following result proved in Appendix B:

Proposition 1 (Optimization over ϵ): The function $J^*(\epsilon)$ in (5) is convex for $J(d, \epsilon)$ in (3) and log-convex for $J(d, \epsilon)$ in (4). Moreover, $J^*(\epsilon)$ is finite on a convex set. \square

Remark 2 (Beyond the additive and multiplicative bounds): Most of the discussion in this section and the application examples discussed in Section IV make use of the additive and multiplicative bounds. However, these do not necessarily provide the tightest bounds. Consider for example a chi-square random variable D with 1 degree of freedom, whose pdf is given by $p_{\chi^2}(d) = \frac{e^{-\frac{d}{2}}}{\sqrt{2\pi d}}$, $\forall d > 0$ and is known to have an expected value $\mathbb{E}[D]$ equal to 1. The additive upper bound from (3) is not useful as it leads to $\forall \epsilon \in \mathbb{R}$

$$\sup_{d>0} \left(d + \epsilon \log \left(\frac{e^{-\frac{d}{2}}}{\sqrt{2\pi d}} \right) \right) + \epsilon \mathcal{H}_{\chi^2} = +\infty,$$

where \mathcal{H}_{χ^2} is the entropy of D . In contrast, the multiplicative upper bound from (4) leads to the following finite bound

$$\inf_{\epsilon \in \mathbb{R}} \sup_{d>0} d \left(\frac{e^{-\frac{d}{2}}}{\sqrt{2\pi d}} \right)^\epsilon \left(\frac{1}{2\pi} \right)^{\frac{1-\epsilon}{2}} \left(\frac{2}{1-\epsilon} \right)^{\frac{1+\epsilon}{2}} \Gamma \left(\frac{1+\epsilon}{2} \right)$$

$$\approx 1.478.$$

However, a tight bound can be obtained using the multiplicative group $(\mathcal{P}, \oplus) = (\mathbb{R}_{>0}, \times)$ together with the function $\alpha(d) = d^{-\epsilon}$, which leads to

$$\inf_{\epsilon \in \mathbb{R}} \sup_{d \geq 0} \left(d d^{-\epsilon} \frac{2^\epsilon \Gamma(0.5 + \epsilon)}{\Gamma(0.5)} \right) = 1.$$

While either the additive or the multiplicative bound typically lead to reasonable bounds, this example shows that it may be worth it to explore alternatives. \square

Remark 3 (Unknown pdf): When the pdf $p_D(\cdot)$ of D is not explicitly known, it is not easy to use the additive and multiplicative bounds in (3)–(4), because both include $p_D(d)$ in the criteria to be optimized over $d \in D$. In such cases, one can still use the bounds in Theorem 1, but with functions $\alpha(d)$ that do not explicitly include the pdf of D . We recall from the discussion right after Theorem 1, that the key to get a tight upper bound is to select for $\alpha(d)$ a function that is strongly “negative” when $V(d)$ is large, and yet $\mathbb{E}[-\alpha(D)]$ is relatively small. For the additive group, the function $\alpha(d) := \log p_D(d)$ typically has this property when large values for $V(d)$ have low probability. When the pdf is unknown, tight bounds can still be obtained as long as one selected for $\alpha(d)$ values that are strongly negative when $V(d)$ is large and yet $D = d$ is unlikely. \square

B. Connection to distributionally robust optimization

Distributionally robust optimization (DRO) can provide an alternative approach to compute bounds for an expected value by noting that

$$\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[V(D)] \leq \mathbb{E}_{\mathbb{P}}[V(D)] \leq \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[V(D)], \quad (6)$$

where the subscript in the expected value operator refers to the probability measure used for the computation of the expected value and \mathcal{P} denotes some class of probability measures that contains the actual measure \mathbb{P} . From a computational perspective, such bounds can be useful when the minimum and maximum over \mathcal{P} are achieved for measures $\bar{\mathbb{P}}$ for which the expectation $\mathbb{E}_{\bar{\mathbb{P}}}[V(D)]$ is easier to compute than the original $\mathbb{E}_{\mathbb{P}}[V(D)]$. For example, if we include in \mathcal{P} every distribution for which D is measurable, we essentially get the trivial bounds in (2).

It turns out that (6) can lead to bounds closely related to those obtained in Theorem 1: Suppose for simplicity that we focus our attention on a discrete random variable $D \in \{d_1, d_2, \dots, d_K\}$ and pick for \mathcal{P} the set of all distributions with entropy larger than or equal to the entropy $\mathcal{H}[\mathbb{P}]$ of the actual probability distribution \mathbb{P} . In this case, the upper bound in (6) is of the form

$$\mathbb{E}_P[V(D)] \leq \max_{\bar{p}_1, \dots, \bar{p}_K} \left\{ \sum_{k=1}^K V(d_k) \bar{p}_k : - \sum_{k=1}^K \log(\bar{p}_k) \bar{p}_k \geq \mathcal{H}[P] \right\}, \quad (7)$$

where the maximization is taken over the simplex of probability distributions. Because the entropy is a strictly concave function, as long as P is not the uniform distribution, $\bar{p}_k = 1/K, k \in \{1, \dots, K\}$ is a Slater point and strong duality holds, which allow us to replace the right-hand side of (7) by its dual problem:

$$\mathbb{E}_P[V(D)] \leq \inf_{\epsilon \leq 0} \max_{\bar{p}_1, \dots, \bar{p}_K} \sum_{k=1}^K V(d_k) \bar{p}_k + \epsilon \sum_{n=1}^K \log(\bar{p}_k) \bar{p}_k + \epsilon \mathcal{H}[P]. \quad (8)$$

For the same expected value, the additive upper bound provided by Theorem 1 is of the form

$$\mathbb{E}_P[V(D)] \leq \inf_{\epsilon \in \mathbb{R}} \max_k V(d_k) + \epsilon \log(p_k) + \epsilon \mathcal{H}[P], \quad (9)$$

where, as in (5), we pick the least conservative upper bound over the range of parameters $\epsilon \in \mathbb{R}$. It turns out that the maximum over k in (9) has the same numerical value as the following maximization over the simplex of distributions:

$$\mathbb{E}_P[V(D)] \leq \inf_{\epsilon \in \mathbb{R}} \max_{\bar{p}_1, \dots, \bar{p}_K} \sum_{n=1}^K V(d_k) \bar{p}_k + \epsilon \sum_{k=1}^K \log(p_k) \bar{p}_k + \epsilon \mathcal{H}[P], \quad (10)$$

leading to a bound strikingly similar to (8). However, the two bounds generally lead to different numerical values:

i) For distributions with large entropy, the DRO inequalities (7) and (8) lead to a tighter bound, because the family of distributions that satisfy the constraint in (7) becomes fairly small. In fact, for the uniform distribution $p_k = 1/K, \forall k$ with maximal entropy $\mathcal{H}[P] = \log K$, only the true distribution satisfies the constraint in (7) and the bound is exact.

ii) For distributions with small entropy, (10) leads to a tighter bound, which is exact for the extreme cases of minimum entropy $\mathcal{H}[P] = 0$. Note that when $\mathcal{H}[P] = 0$, all but one of the p_k is nonzero and a single value of k leads to a value of $V(d_k) + \epsilon \log(p_k)$ with $\epsilon > 0$ in (10) that is not $-\infty$.

Even though the DRO-based approach in (7)–(8) can often lead to a tighter bound than (9)–(10), an expected value $\sum_{k=1}^K V(u, d_k) \bar{p}_k$ still appears in (7)–(8) and therefore this bound is only helpful in simplifying computations if the optimal distribution $\bar{p}_1, \dots, \bar{p}_K$ has some particular structure that makes the computation of $\sum_{k=1}^K V(u, d_k) \bar{p}_k$ easier than the original computation $\mathbb{E}_P[V(D)] = \sum_{k=1}^K V(u, d_k) p_k$.

Remark 4: We focused this section on a discrete random variable D to avoid the technicalities that would arise from optimizations over general probability measures in (7) and (10), but all the key observations made in this section remain unchanged for a continuous random variable D . \square

III. STOCHASTIC PROGRAMMING

We define the following *stochastic programming problem* with a single scalar constraint, but the approach proposed can easily be extended to multiple constraints: Let D be a random vector taking values in $\mathcal{D} \subset \mathbb{R}^M$. Given measurable functions $V : \mathcal{U} \times \mathcal{D} \mapsto \mathbb{R}$ and $G : \mathcal{U} \times \mathcal{D} \mapsto \mathbb{R}$, with $\mathcal{U} \subset \mathbb{R}^N$ we want to solve

$$V^* := \inf_{u \in \mathcal{U}} \left\{ \mathbb{E}[V(u, D)] : \mathbb{E}[G(u, D)] \leq 0 \right\}. \quad (11)$$

The following results provides bounds on V^* , based on the bounds from Theorem 1.

Theorem 2 (Bounds to Stochastic Programming): Consider three E-distributive groups $(\mathcal{P}_V, \oplus_V)$, $(\mathcal{P}_G, \oplus_G)$, (\mathcal{P}, \oplus) ; functions $\alpha_V : \mathcal{D} \mapsto \mathcal{P}_V$, $\alpha_G : \mathcal{D} \mapsto \mathcal{P}_G$, $\alpha : \mathcal{D} \mapsto \mathcal{P}$ and define

$$J_V(u, d) := V(u, d) \oplus_V \alpha_V(d) \oplus_V \mathbb{E}[-\alpha_V(D)]$$

$$J_G(u, d) := G(u, d) \oplus_G \alpha_G(d) \oplus_G \mathbb{E}[-\alpha_G(D)]$$

$$J(u, d, \lambda) := (V(u, d) + \lambda G(u, d)) \oplus \alpha(d) \oplus \mathbb{E}[-\alpha(D)],$$

$\forall u \in \mathcal{U}, d \in \mathcal{D}, \lambda \geq 0$. If $\mathbb{E}[-\alpha_V(D)]$, $\mathbb{E}[-\alpha_G(D)]$, $\mathbb{E}[-\alpha(D)]$ are finite, then $V^\nabla \leq V^* \leq V^\Delta$ and $V^* \leq V^\ddagger$, with

$$V^\nabla := \inf_{u \in \mathcal{U}} \left\{ \left(\text{ess inf } J_V(u, D) \right) : \text{ess inf } J_G(u, D) \leq 0 \right\} \quad (12)$$

$$V^\Delta := \inf_{u \in \mathcal{U}} \left\{ \left(\text{ess sup } J_V(u, D) \right) : \text{ess sup } J_G(u, D) \leq 0 \right\} \quad (13)$$

$$V^\ddagger := \inf_{u \in \mathcal{U}} \sup_{\lambda \geq 0} \text{ess sup } J(u, D, \lambda). \quad (14)$$

Furthermore, if the infimum in the definition of V^Δ is achieved at some $u = u^\Delta$ that is feasible for (13), then u^Δ is also feasible for (11). Additionally, if the infimum in the definition of V^\ddagger is finite and achieved at some $u = u^\ddagger$ then u^\ddagger is also feasible for (11). \square

Theorem 2 guarantees that a solution u^Δ to the optimization (13) is feasible for the original stochastic program in (11) and provides performances guarantees for u^Δ , in the sense that the expected value $\mathbb{E}[V(u^\Delta, D)]$ obtained using u^Δ will be away from the optimal V^* by no more than $V^\Delta - V^\nabla$, which can be computed by solving the optimizations (12)–(13). Similarly, a solution u^\ddagger to the optimization (14) is also guaranteed to be feasible and the expected value $\mathbb{E}[V(u^\ddagger, D)]$ obtained using u^\ddagger will be away from the optimal V^* by no more than $V^\ddagger - V^\nabla$, which can be computed by solving the optimizations (12), (14).

It is important to note that $\mathbb{E}[-\alpha_V(D)]$ and $\mathbb{E}[-\alpha(D)]$ are constants that do not depend on either u nor d , and therefore their values do not affect the optimizations in (12)–(14). This means that, if one is not able to determine analytically $\mathbb{E}[-\alpha_V(D)]$ or $\mathbb{E}[-\alpha(D)]$, any errors in estimating these quantities will not introduce errors in determining u^Δ or u^\ddagger . This is specially relevant in large scale problems where obtaining accurate numerical estimates of $\mathbb{E}[-\alpha_V(D)]$ and $\mathbb{E}[-\alpha(D)]$ might be challenging.

As was the case for Theorem 1, the tightness of the bounds in Theorem (2) depends strongly on the choice of the groups,

the functions α_\cdot , the cost function, and the underlying random variable. Nevertheless, as we will see in the next section, the value of u that minimizes (11) and the value of u that minimizes (13) or (14) are often very close.

Proof of Theorem 2. In view of from Theorem 1, we have that $\text{ess inf } J_G(u, D) \leq \mathbb{E}[G(u, D)] \leq \text{ess sup } J_G(u, D)$, which guarantees that if $u = u^\Delta$ is feasible for (13), then u^Δ is also feasible for (11). Moreover,

$$\begin{aligned} & \inf_{u \in \mathcal{U}} \left\{ \mathbb{E}[V(u, D)] : \text{ess inf } J_G(u, D) \leq 0 \right\} \\ & \leq \inf_{u \in \mathcal{U}} \left\{ \mathbb{E}[V(u, D)] : \mathbb{E}[G(u, D)] \leq 0 \right\} \\ & \leq \inf_{u \in \mathcal{U}} \left\{ \mathbb{E}[V(u, D)] : \text{ess sup } J_G(u, D) \leq 0 \right\}. \end{aligned}$$

From Theorem 1, we can also conclude that $\text{ess inf } J_V(u, D) \leq \mathbb{E}[V(u, D)] \leq \text{ess sup } J_V(u, D)$, from which it follows that $V^\nabla \leq V^* \leq V^\Delta$.

To establish that V^\ddagger is also an upper bound on V^* , assume by contradiction that $V^\ddagger < V^*$, which means that there exists some $u \in \mathcal{U}$ such that $\text{ess sup } J(u, D, \lambda) < V^*$, $\forall \lambda \geq 0$. In view of Theorem 1, this would mean that $\mathbb{E}[V(u, D) + \lambda G(u, D)] < V^*$, $\forall \lambda \geq 0$, which is only possible if $\mathbb{E}[G(u, D)] \leq 0$ and consequently $\mathbb{E}[V(u, D)] < V^*$. The existence of such an u violates (11).

Finally note that if the infimum in the definition of V^\ddagger is finite and achieved at some $u = u^\ddagger$, then we must have $\text{ess sup } J(u, D, \lambda) \leq V^* < \infty$, $\forall \lambda \geq 0$. Reasoning as in the paragraph above, this allow us to conclude that $\mathbb{E}[G(u^\ddagger, D)] \leq 0$ and therefore u^\ddagger is feasible. ■

A. Combination with Monte Carlo methods

Any point u_{feasible} that is feasible for the optimization (11) can be used to construct an upper bound by using Monte Carlo averaging to compute

$$V^* \leq \mathbb{E}[V(u_{\text{feasible}}, D)] \approx \frac{1}{K} \sum_{k=1}^K V(u_{\text{feasible}}, d_k), \quad (15)$$

where the d_k are independent samples of D . Moreover, it is possible to control the error introduced by the Monte Carlo averaging by using a sufficiently large number of samples K . Essentially, to have an error smaller than δ with high probability we need

$$K \geq c \text{Var}[V(u_{\text{feasible}}, D)] / \delta^2, \quad (16)$$

where the constant c is typically small and depends on the desired confidence for the bound [10].

Any point u^Δ that achieves V^Δ and in feasible for (13) is also feasible for (11) and can be used in (15) to construct an upper bound that is typically tighter than V^Δ , provided that K is sufficiently large; the same reasoning is true for u^\ddagger and V^\ddagger . In fact, one can use Theorem 1 to compute other feasible points that may provide tighter upper bounds. For example, an alternative feasible point can be obtained by minimizing a

lower bound on the criterion constrained by an upper bound on the constraints, which leads to

$$V^\perp := \inf_{u \in \mathcal{U}} \left\{ \left(\text{ess inf } J_V(u, D) \right) : \left(\text{ess sup } J_G(u, D) \leq 0 \right) \right\}. \quad (17)$$

Unlike V^Δ and V^∇ in Theorem 2, V^\perp neither provides an upper nor a lower bound on V^* . However, any point that achieves the infimum and is feasible for (17) is also feasible for (11) and therefore can be used to construct the upper bound in (15). An alternative method to combine the results in Theorem 1 with Monte Carlos methods is obtained by replacing the optimization in (14) by

$$\inf_{u \in \mathcal{U}} \text{ess sup } J(u, D, \lambda)$$

for some fixed $\lambda \geq 0$. Rather than taking the supremum over $\lambda \geq 0$ that appears in (14), one could simply adjust λ and/or artificially tightening the constraint until a Monte Carlo estimate for $\mathbb{E}[G(u, D)]$ guarantees that the constraint is satisfied with a sufficiently large confidence.

Remark 5 (Contrast with Sample Average Approximation): It is important to emphasize the difference between using Monte Carlo averaging to estimate the value of the expected value for a given value of $u_{\text{feasible}} \in \mathcal{U}$, as in (15), and optimizing a Monte Carlo approximation of the criterion, as in

$$\min_{u \in \mathcal{U}} \frac{1}{K} \sum_{k=1}^K V(u, d_k), \quad (18)$$

which is typically referred to as the Sample Average Approximation (SAA). We can see in (16) that the number of samples required to achieve a desired error $\delta > 0$ depends mostly on the variance of $V(u, D)$ at the point u_{feasible} . However, the sample complexity required to obtain the same error in (18) is typically much larger as the numerator of (16) would be determined by the Vapnik-Chervonenkis (VC) dimension of the family of functions $u \mapsto \mathbb{E}[V(u, D)]$ [37], [38]. □

B. Numerically computing the bounds

We show next that under appropriate regularity assumptions, the essential infima and suprema in (12) and (13) are achieved at minima and maxima, respectively, and can be computed using numerical solvers. To formalize this observation we recall that a function $f : \mathcal{X} \rightarrow \mathbb{R}$, $\mathcal{X} \subset \mathbb{R}^n$ is said to have *compact sublevel sets* if its sublevel sets $\{x \in \mathcal{X} : f(x) \leq \lambda\}$ are compact for every finite $\lambda \in \mathbb{R}$ and it is said to have *compact suplevel sets* if $-f$ has compact sublevel sets. For the remainder of this section we assume that $\mathcal{U} \subset \mathbb{R}^N$ and that $\mathcal{D} \subset \mathbb{R}^M$ is the support of the random variable D , i.e., the smallest subset of \mathbb{R}^M for which $P(d \in \mathcal{D}) = 1$.

Theorem 3: Assume that the functions $J_V, J_G : \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}$ are continuous and that \mathcal{U} is compact. If J_V and J_G have compact sublevel sets and (12) is feasible, then V^∇ can be obtained by solving

$$V^\nabla = \min_{u \in \mathcal{U}, d, \tilde{d} \in \mathcal{D}} \left\{ J_V(u, d) : J_G(u, \tilde{d}) \leq 0 \right\}. \quad (19)$$

If $J_V(u, d)$ and $J_G(u, d)$ have compact suplevel sets and (13) is feasible, then $V^\Delta = \lim_{\mu \rightarrow \infty} V_\mu^\Delta$, with

$$V_\mu^\Delta = \min_{u \in \mathcal{U}} \max_{d, \tilde{d} \in \mathcal{D}} J_V(u, d) + \mu \left(\max \left\{ 0, J_G(u, \tilde{d}) \right\} \right)^2. \quad (20)$$

Whenever the minimum and maxima in (20) are achieved for values $u^\Delta \in \mathcal{U}$ and $d^\Delta, \tilde{d}^\Delta \in \mathcal{D}$, respectively, for which $J_G(u^\Delta, \tilde{d}^\Delta) \leq 0$, then u^Δ is feasible and V_μ^Δ is an upper bound for V^* . \square

The minimization in (19) is a regular constrained optimization and can be solved using commercial products like Knitro [39] or open-source solvers like IPOPT [40] and TensCalc [41]. For the sequence of minmax problem in (20), different algorithms are applicable depending on the convexity assumptions (convex-concave, nonconvex-concave, convex-nonconcave, nonconvex-nonconcave). These include methods based on robust counterpart [18], [42]–[44], cutting-set [45], and variations of gradient descent-ascent methods such as [46]–[56] among many others. For the examples in Section IV, we used TensCalc, which is based on a variation of interior point methods for gradient descent ascent.

Regardless of the solver used, for nonconvex problem like the ones typically arising in (19)–(20), it is generally hard to be certain that a local minimum or a local minmax [24] found by a numerical solver is actually a global optimum. An approach that can be used to obviate this problem is to replace the nonconvex optimization that arises from our bounds by pessimistic or optimistic convex relaxation, depending on whether we are interested in an upper or lower bound on the expected value, respectively. An alternative approach relies on analyzing the consequences of a solver getting stuck at a local optima and responding to the specific problems encountered: Theorem 3 essentially proposes to replace the stochastic optimization in (11) by the sequence of deterministic robust optimizations V_μ^Δ . A numerical solver for (20) can typically be “fooled” in three ways:

i) The solver could converge to a value d^Δ for d that is a local but not a global extremum to the inner maximization. This would mean that the value V_μ^Δ returned by the solver is actually not an upper bound on $\mathbb{E}[V(u^\Delta, D)]$. If it is important to obtain a high-confidence bound for this expected value and the inner maximization is not concave (or known to only have a unique local/global maximum), then one can use a Monte Carlo method to get an accurate estimate for $\mathbb{E}[V(u^\Delta, D)]$, which is typically computationally much easier than solving (11), as discussed in Section III-A.

ii) The solver may converge to a value u^Δ for u for which $J_G(u^\Delta, \tilde{d}^\Delta) \leq 0$ holds for a local maximum \tilde{d}^Δ that is not global and the expected value $\mathbb{E}[G(u^\Delta, D)]$ is actually positive. Again here, once the optimization finishes, we can use a Monte Carlo method to obtain an accurate estimate for $\mathbb{E}[G(u^\Delta, D)]$ and reject the solution u^Δ if the constraint is violated. Hopefully, different initialization for the solver would resolve this, but one could also tighten the constraint by asking $\max_{\tilde{d} \in \mathcal{D}} J_G(u, \tilde{d})$ to actually be negative.

iii) Finally the solver, may return a value u^Δ for u that satisfies the constraint but is a local (rather than a global) extremum of the outer minimization. In this case, it may

be possible to get a better solution, but the solver was unable to find it. In practice, for nonconvex problems there is little protection against this, rather than trying a different initialization for the solver.

It should be noted that any approach based on constructing (non-exact) convex relaxations to (20) will have very similar issues: pessimistic relaxations may overlook better solutions (as in iii), whereas optimistic relaxations may accept solutions that violate constraints (as in ii).

Lemma 1 (Equivalent compact subset): Consider a continuous function $J : \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}$ with \mathcal{U} compact. If J has compact sublevel sets, there exists a compact set $\mathcal{D}^\dagger \subset \mathcal{D}$ such that

$$J^{\text{inf}}(u) := \text{ess inf}_{d \in \mathcal{D}} J(u, d) = \min_{d \in \mathcal{D}^\dagger} J(u, d), \quad \forall u \in \mathcal{U}$$

and the function J^{inf} is continuous. Similarly, if J has compact suplevel sets, there exists a compact set $\mathcal{D}^\dagger \subset \mathcal{D}$ such that

$$J^{\text{sup}}(u) := \text{ess sup}_{d \in \mathcal{D}} J(u, d) = \max_{d \in \mathcal{D}^\dagger} J(u, d), \quad \forall u \in \mathcal{U}$$

and the function J^{sup} is continuous.

Proof of Lemma 1. First note that because \mathcal{D} is the support of the random variable D and J is continuous, the essential infimum of $J(u, D)$ is equal to the usual infimum of $J(u, d)$ over $d \in \mathcal{D}$. The same is true for the supremum,

We prove the result only for the minimization, as the proof for the maximization is analogous. Take an arbitrary point $d^\dagger \in \mathcal{D}$ and define

$$\lambda^\dagger := \max_{u \in \mathcal{U}} J(u, d^\dagger), \quad \mathcal{S}^\dagger := \{(u, d) \in \mathcal{U} \times \mathcal{D} : f(u, d) \leq \lambda^\dagger\}.$$

The constant λ^\dagger is finite because J is continuous and $\mathcal{U} \times \{d^\dagger\}$ is a compact set, and the set \mathcal{S}^\dagger is compact because J has compact sublevel sets. The desired set \mathcal{D}^\dagger is then given by the closure of

$$\mathcal{D}_o := \bigcup_{u \in \mathcal{U}} \left\{ d \in \mathcal{D} : (u, d) \in \mathcal{S}^\dagger \right\}.$$

Note that \mathcal{D}_o is bounded because \mathcal{S} is bounded and therefore its closure \mathcal{D}^\dagger is compact. To show that the infimum of $J(u, d)$ over \mathcal{D} is achieved at some point in \mathcal{D}^\dagger , assume by contradiction that there exists some $d^* \notin \mathcal{D}^\dagger$ such that $J(u, d^*) < J(u, d)$, $\forall d \in \mathcal{D}^\dagger$. Since $d^\dagger \in \mathcal{D}^\dagger$, we conclude that $J(u, d^*) < J(u, d^\dagger) \leq \lambda^\dagger$. This establishes a contradiction, because it would mean that $(u, d^*) \in \mathcal{S}^\dagger$ and therefore $d^* \in \mathcal{D}_o \subset \mathcal{D}^\dagger$. Continuity of J^{inf} then follows from Berge’s Maximum Theorem [57, Chapter E.3]. \blacksquare

Proof of Theorem 3. In view of Lemma 1, all the essential infima and suprema in (12)–(13) are achieved at some point inside a compact subset \mathcal{D}^\dagger of \mathcal{D} and

$$V^\nabla = \inf_{u \in \mathcal{U}} \left\{ J_V^{\text{inf}}(u) : J_G^{\text{inf}}(u) \leq 0 \right\} \quad (21)$$

for the continuous functions

$$J_V^{\text{inf}}(u) := \min_{d \in \mathcal{D}} J_V(u, d), \quad J_G^{\text{inf}}(u) := \min_{d \in \mathcal{D}} J_G(u, d). \quad (22)$$

Since $J_G^{\text{inf}}(u)$ is continuous and \mathcal{U} is compact, the feasible set $\{u \in \mathcal{U} : J_G^{\text{inf}}(u) \leq 0\}$ is compact and nonempty by

assumption. Weierstrass Theorem [58, Proposition A.8] then allow us to conclude that the inf is actually achieved at some point $u^\nabla \in \mathcal{U}$ of the feasible set. Denoting by d_V^∇ and d_G^∇ points in \mathcal{D} at which the minima in (22) are achieved for $u = u^\nabla$, we conclude that

$$V^\nabla = J_V(u^\nabla, d_V^\nabla), \quad J_G(u^\nabla, d_G^\nabla) \leq 0,$$

which shows that the right-hand side of (19) cannot be larger than V^∇ . By contradiction, assume that it is actually strictly smaller than V^∇ . This would mean there exist $u \in \mathcal{U}$ and $d, \tilde{d} \in \mathcal{D}$ such that

$$J_v(u, d) < V^\nabla, \quad J_G(u, \tilde{d}) \leq 0.$$

The right-hand side inequality shows that $J_G^{\text{inf}}(u) \leq 0$ and therefore such u is feasible for (21) and the left-hand side inequality shows that $J_V^{\text{inf}}(u) < V^\nabla$, which contradicts the fact that the infimum in (21) is equal to V^∇ .

Again using Lemma 1, we conclude that

$$V^\Delta = \inf_{u \in \mathcal{U}} \left\{ J_V^{\text{sup}}(u) : J_G^{\text{sup}}(u) \leq 0 \right\} \quad (23)$$

for the continuous functions

$$J_V^{\text{sup}}(u) := \max_{d \in \mathcal{D}} J_V(u, d), \quad J_G^{\text{sup}}(u) := \max_{d \in \mathcal{D}} J_G(u, d). \quad (24)$$

In view of [58, Proposition 4.2.1], $\lim_{\mu \rightarrow \infty} \bar{V}_\mu^\Delta = V^\Delta$, with

$$\bar{V}_\mu^\Delta := \min_{u \in \mathcal{U}} J_V^{\text{sup}}(u) + \mu \left(\max \{0, J_G^{\text{sup}}(u)\} \right)^2.$$

The result then follows by noting that

$$\max \{0, J_G^{\text{sup}}(u)\} = \max_{\tilde{d} \in \mathcal{D}} \max \{0, J_G(u, \tilde{d})\}.$$

and therefore $\bar{V}_\mu^\Delta = V_\mu^\Delta$ for positive μ . ■

IV. SELECTED APPLICATIONS

A. Stochastic control

Consider the dynamical system

$$x_{t+1} = f(x_t, \theta, u_t, d_t) \quad (25a)$$

$$y_t = h(x_t) + n_t, \quad (25b)$$

where x_t denotes the state of the system at time t , u_t the controlled input, d_t a random disturbance input, y_t the measured output, n_t measurement noise, and θ a random vector of parameters.

Our goal to select control inputs u_0, \dots, u_{T-1} to minimize a finite-horizon criterion of the form

$$\mathbb{E}[W(x_1, \dots, x_T, u_0, \dots, u_{T-1})], \quad (26)$$

subject to a constraint of the form

$$\mathbb{E}[U(x_1, \dots, x_T, u_0, \dots, u_{T-1})] \leq 0. \quad (27)$$

We consider two versions of this problem: First a *state-feedback* scenario in which the initial state x_0 is known and the expectation (26) is with regard to the random parameters θ and the disturbances d_0, \dots, d_{T-1} . We then consider an *output-feedback* scenario in which the initial state is not known, but one has available past measurements y_{-K}, \dots, y_0 . In this case, the expectation in (26) is conditioned to these past measurements and it regards the measurement noise n_{-K}, \dots, n_0 , the initial state x_{-K} , and the past disturbances d_{-K}, \dots, d_{-1} .

a) *State Feedback*: The state-feedback control problem can be viewed as an instance of (11), with the following associations

$$\begin{aligned} u &:= (u_0, \dots, u_{T-1}), \\ D &:= (\theta, d_0, \dots, d_{T-1}), \\ V(u, D) &:= W(x_1, \dots, x_T, u_0, \dots, u_{T-1}), \\ G(u, D) &:= U(x_1, \dots, x_T, u_0, \dots, u_{T-1}), \end{aligned}$$

with the understanding that the states x_1, \dots, x_T that appear in the definitions of $V(u, D)$ and $G(u, D)$ are obtained along solutions to (25a) for the control input in u and the parameters and input disturbances in D .

Assuming that the disturbances d_t are independent and identically distributed with pdf $p_d(\cdot)$ and differential entropy \mathcal{H}_d , and that the parameter θ has pdf $p_\theta(\cdot)$ and differential entropy \mathcal{H}_θ , we have that

$$\mathcal{H}_D = \mathcal{H}_\theta + T\mathcal{H}_d, \quad \log p_D(\theta, d) = \log p_\theta(\theta) + \sum_{t=0}^{T-1} \log p_d(d_t),$$

and the optimization in (13) with additive upper bounds for \oplus_G and \oplus_V takes the form

$$\begin{aligned} V^\Delta = \min_{u \in \mathcal{U}} \left\{ X(u) : U(\bar{x}_1, \dots, \bar{x}_T, u_0, \dots, u_{T-1}) + \bar{\epsilon}\mathcal{H}_D + \right. \\ \left. + \bar{\epsilon} \log p_D(\bar{\theta}, \bar{d}) \leq 0, \quad \forall \bar{\theta}, \bar{d} \right\} \end{aligned}$$

$$\begin{aligned} X(u) := \max_{\theta \in \Theta, d \in \mathcal{D}} W(x_1, \dots, x_T, u_0, \dots, u_{T-1}) + \\ \epsilon\mathcal{H}_D + \epsilon \log p_D(\theta, d), \end{aligned}$$

where \mathcal{U} denotes the set of admissible controls; Θ and \mathcal{D} the supports of the distributions for the random parameter and disturbance, respectively; $\bar{x}_1, \dots, \bar{x}_T$ the solution to (25a) for the control $u := (u_0, \dots, u_{T-1})$, parameter θ and disturbance $\bar{d} := (\bar{d}_0, \dots, \bar{d}_{T-1})$; x_1, \dots, x_T the solution to (25a) for the same control $u := (u_0, \dots, u_{T-1})$, but parameter θ and disturbance $d := (d_0, \dots, d_{T-1})$; and $\epsilon, \bar{\epsilon}$ the scalar parameters associated with additive upper bounds used for \oplus_G and \oplus_V , respectively. An equivalent formulation of the optimization in (12) gives V^∇ .

b) *Output Feedback*: The output-feedback problem can also be viewed as an instance of (11), but now with the following associations

$$\begin{aligned} u &:= (u_0, \dots, u_{T-1}), \\ D &:= (\theta, x_{-K}, d_{-K}, \dots, d_{T-1}), \\ V(u, D) &:= W(x_1, \dots, x_T, u_0, \dots, u_{T-1}), \\ G(u, D) &:= U(x_1, \dots, x_T, u_0, \dots, u_{T-1}), \end{aligned}$$

with the understanding that the states x_1, \dots, x_T that appear in the definition of $V(u, D)$ and $G(u, D)$ are obtained along solutions to (25a) for the control input in u and the parameters, initial state, and input disturbances in D . In addition, the expectation in (11) is now a conditional expectation, given measurements $Y = (y_{-K}, \dots, y_0)$ defined by (25b).

In this case, the optimization in (13) with additive upper bounds for \oplus_V and \oplus_G takes the form

$$V^\Delta = \min_{u \in \mathcal{U}} \left\{ X(u) : U(\bar{x}_1, \dots, \bar{x}_T, u_0, \dots, u_{T-1}) + \right.$$




Fig. 1: Linear system with unknown dynamics, comparison of the controls u^Δ obtained from Theorem 2 and u_{SAA} obtained using Sample Average Approximation. Using Monte Carlo integration, we obtain that $\mathbb{E}[V(u^\Delta, D)] = 1.79 \times 10^3$ and $\mathbb{E}[V(u_{SAA}, D)] = 1.62 \times 10^3$.

$$\begin{aligned} & \left\{ \bar{\epsilon} \log p_{D|Y}(\bar{\theta}, \bar{d}, \bar{x}_k) + \bar{\epsilon} \mathcal{H}_{D|Y}(y_{-K}, \dots, y_0) \leq 0, \forall \bar{\theta}, \bar{d}, \bar{x}_k \right\} \\ & X(u) := \max_{\theta \in \Theta, d \in \mathcal{D}, x_{-K} \in \mathcal{X}_{-K}} W(x_1, \dots, x_T, u_0, \dots, u_{T-1}) + \\ & \epsilon \log p_{D|Y}(\theta, d, x_{-K}) + \epsilon \mathcal{H}_{D|Y}(y_{-K}, \dots, y_0), \end{aligned} \quad (28)$$

where we use the version of the bounds for conditional expectation mentioned in Remark 1. The conditional *pdf* that appears in (28) can be computed using the following result.

Lemma 2 (Conditional pdf of a dynamical system): In addition to the assumptions made for the state feedback case, also assume that the observation noises n_t are independent and identically distributed with *pdf* $p_n(\cdot)$ and that the initial state x_{-K} has *pdf* $p_{x_{-K}}(\cdot)$. If $p_Y(y_{-K}, \dots, y_0) \neq 0$, the conditional probability density function $p_{D|Y}(\cdot)$ is given by

$$\frac{\prod_{t=-K}^0 p_n(y_t - h(x_t)) \prod_{t=-K}^{T-1} p_d(d_t) p_{x_{-K}}(x_{-K}) p_\theta(\theta)}{p_Y(y_{-K}, \dots, y_0)}$$

with the understanding that x_t is obtained along the solutions to (25a). \square

Proof of Lemma 2. Using the independence of n_t , one deduces that the observations y_t are conditionally independent:

$$p_{Y|D}(y_{-K}, \dots, y_0 | x_{-K}, \dots, x_0) = \prod_{t=-K}^0 p_{Y_t|D}(y_t | x_t).$$

As the noise n_t is additive in (25b), a change of variable gives $p_{Y_t|D}(y_t | x_t) = p_n(y_t - h(x_t))$. Using Bayes' theorem and the independence of d_t , θ , and x_{-K} finishes the proof. \blacksquare

The differential entropy $\mathcal{H}_{D|Y}(y_{-K}, \dots, y_0)$ that appears in (28) is typically difficult to compute (or even to estimate, e.g., through Monte Carlo integration); especially for a long sequence of *past* measurements y_{-K}, \dots, y_0 . However, this entropy is not affected by the optimization variable $u = (u_0, \dots, u_{T-1})$, which only includes *future* controls. This means that we can determine the optimal value for u in (28) without actually computing $\mathcal{H}_{D|Y}(y_{-K}, \dots, y_0)$.

Example 1 (Linear system with unknown dynamics): Consider a linear system, i.e., a system with dynamics

$$\begin{aligned} x_{t+1} &= A x_t + B u_t + d_t \\ y_t &= C x_t + n_t \end{aligned}$$

with d_t and n_t independent zero mean standard Gaussian processes. The system is time-invariant, C is an identity matrix, but the matrices A and B are unknown stochastic parameters of the form

$$A = \begin{bmatrix} A_{11} & A_{12} & 0 \\ 0 & A_{22} & A_{23} \\ 0 & 0 & A_{33} \end{bmatrix} \quad B = \begin{bmatrix} 0 \\ 0 \\ B_{31} \end{bmatrix},$$

where $A_{11}, A_{12}, A_{22}, A_{23}, A_{33}, B_{31}$ are independent Gaussian random variables with mean 1 and standard deviation 0.25. We chose a quadratic cost

$$\begin{aligned} W(u_0, \dots, u_{T-1}, x_0 \dots x_T) = \\ \sum_{t=0}^{T-1} 0.5 \|u_t\|_2^2 + 0.5 \|x_t\|_2^2 + 0.5 \|x_T\|_2^2 \end{aligned}$$

with a future horizon $T = 10$ and constraints on the control that $\|u\|_\infty \leq 1$. We suppose access to past measurements y_{-K}, \dots, y_0 with $K = 20$.

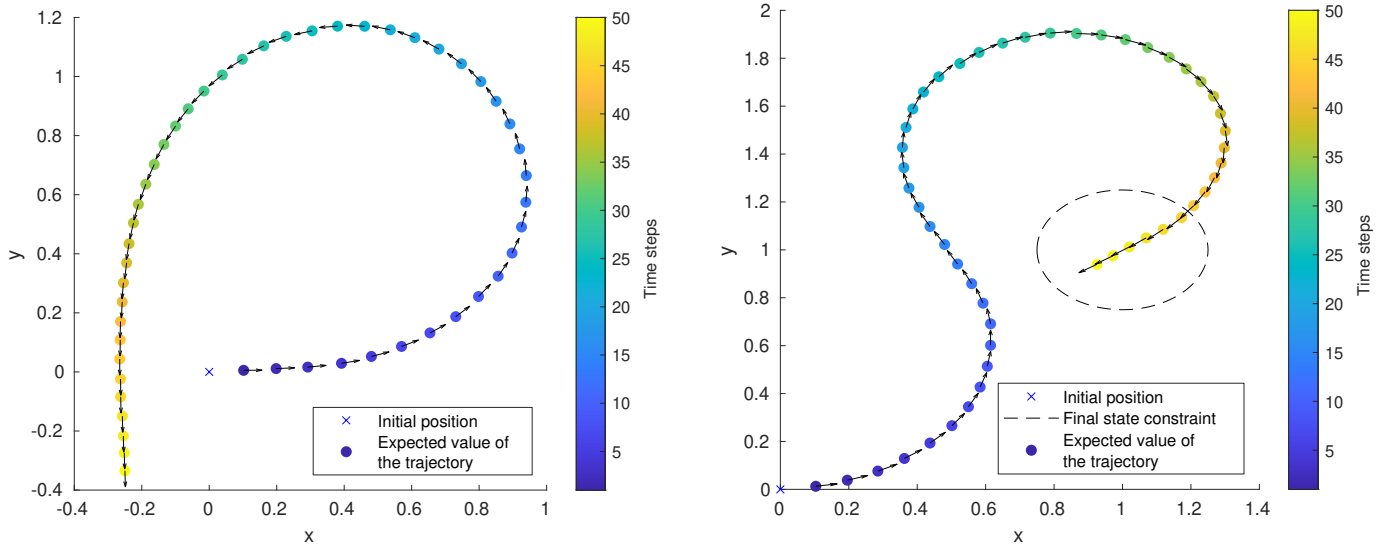
The value of the upper bound V^Δ is 5.04×10^5 and the value of the lower bound V^∇ is 28. We compare our results with an approximate solution obtained using Sample Average Approximation (SAA) (i.e., minimizing an empirical mean of the cost). Solving the upper bound and lower bound optimizations (Theorem 2) takes about 0.1 seconds, while solving the Sample Average Approximation takes about 5 minutes. In Figure 1 one can see that the controls match each other fairly closely until $t = 6$, when they start to slightly diverge. We also use Monte Carlo integration, as discussed in Section III-A, to estimate the expected value of the cost for the two controls, obtaining that they differ by about 10%.

Example 2 (Dubins vehicle): Consider a discrete time Dubins vehicle [59], [60] with dynamics

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \\ \omega_{t+1} \end{bmatrix} = \begin{bmatrix} x_t \\ y_t \\ \omega_t \end{bmatrix} + T_s \begin{bmatrix} v \cos(\omega_t) \\ v \sin(\omega_t) \\ u_t \end{bmatrix} + \frac{T_s^2}{2} \begin{bmatrix} -v \sin(\omega_t) u_t \\ v \cos(\omega_t) u_t \\ 0 \end{bmatrix} + d_t$$

where $T_s = 0.1$ is the sampling period, $v = 1$ is a constant forward speed. The initial state is known to be $[x_0, y_0, \omega_0]' = [0, 0, 0]$, and we want to optimize for a future horizon $T = 50$. The controls are constrained such that $\|u\|_\infty \leq \pi/2$. The disturbance $d_t = [d_t^{(x)}, d_t^{(y)}, d_t^{(\omega)}]'$ is such that $d_t^{(x)}, d_t^{(y)}$ are zero mean Gaussian random variables with variance T_s , and $d_t^{(\omega)}$ is a von Mises random variable, with probability density function $e^{\kappa \cos(x)} / (2\pi I_0(\kappa))$ with $\kappa = 5/T_s$ and where $I_0(\kappa)$ is the modified Bessel function of order 0. The cost function is

$$\begin{aligned} W(u_0, \dots, u_{T-1}, x_0 \dots x_T) = \\ \sum_{t=0}^{T-1} 0.5 \|u_t\|_2^2 + \sum_{t=0}^T 0.5 \|x_t\|_2^2 + 0.5 \|y_t\|_2^2 \end{aligned}$$



(a) Expected value of the trajectory without constraints on the final state; the expected value of the cost for this control is $\mathbb{E}(V(u^\Delta, D)) = 189$.

(b) Expected value of the trajectory with constraints on the final state; the expected value of the cost for this control is $\mathbb{E}(V(u^\ddagger, D)) = 250$.

Fig. 2: Expected value of the trajectory of the Dubins vehicle given two different controls estimated using Monte Carlo integration. Without constraints (a), the control brings the expected value of the trajectory back to near the origin. With the inclusion of constraints (b), the control drives the expected value of the final state towards the correct region.

We present two cases, one with no constraints on the states and one with a constraint on the final state. For both of them, we use the additive bounds.

The first case, without constraints, takes about 1 second to solve, the value of the upper bound V^Δ is 1.25×10^5 the lower bound V^∇ only provides the trivial value of 0. However, using a Monte Carlo integration we compute the expected value of the cost given the control and obtain 189. We use a Stochastic Gradient Descent to solve (26), which takes about 15 seconds, the optimal cost is 187 and the error between the solution obtained using the Stochastic Gradient Descent u_{SGD} and the solution obtained using the upper bound u^Δ is $\|u_{SGD} - u^\Delta\|_\infty = 0.029$, suggesting that u^Δ approximately finds the optimal solution to (26)

For the second case, we include the constraint

$$\mathbb{E} \left[\left\| \begin{bmatrix} x_T \\ y_T \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\|_2 \right] \leq 0.25 \quad (29)$$

i.e., we want to find a control such that the expectation of the final value of the trajectories of (x, y) be in neighborhood around the point $(1, 1)$ (look at Figure 2b for a visualization of the constraints). As the problem now has stochastic constraint, we have to choose between using the upper bound V^Δ from (13) which requires u to satisfy the constraint

$$\max_d \left\| \begin{bmatrix} x_T \\ y_T \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\|_2 + \epsilon \log p_D(d) + \epsilon \mathcal{H}_D \leq 0.25, \quad (30)$$

or the upper bound V^\ddagger from (14). Unfortunately the bound (30) of (29) is too conservative, and renders the problem infeasible. The upper bound V^\ddagger does not suffer from this problem. It takes about 30 seconds to solve the optimization for which we obtain a value for the upper bound V^\ddagger of

1.25×10^8 and the lower bound V^∇ provides only the trivial value of 0. However, using Monte Carlo integration, we obtain that the expected value of the cost is 250.

B. Maximum Likelihood and Maximum a Posteriori with latent variables

Consider an observation x of a random vector X taking values in \mathbb{R}^M whose distribution depends on an unknown parameter $\theta \in \mathbb{R}^P$ that one wants to estimate. The Maximum Likelihood Estimation (MLE) [1] of θ is a vector $\theta^* \in \mathbb{R}^P$ such that

$$\theta^* \in \arg \max_{\theta} p_X(x; \theta). \quad (31)$$

where the *pdf* of X is $p_X(x; \theta)$. The Maximum a Posteriori (MAP) is the analogous of the MLE in Bayesian estimation, *i.e.*, when one regards θ as a realization of a random variable Θ , called the prior, which has *pdf* $p_\Theta(\cdot)$. In this case, the MAP estimation of θ is a vector $\theta^* \in \mathbb{R}^P$ such that

$$\theta^* \in \arg \max_{\theta} p_{X|\Theta}(x | \theta) p_\Theta(\theta). \quad (32)$$

In many cases, constructing the model requires including latent variables that cannot be directly observed. This means that one does not know $p_X(x; \theta)$ but does know $p_{X|D}(x | d; \theta) p_D(d)$, where D is a "latent" random vector taking values in \mathbb{R}^N . In this case, the MLE θ^* is given by

$$\begin{aligned} \theta^* \in \arg \max_{\theta} p_X(x; \theta) &= \int_{\mathcal{D}} p_{X|D}(x | d; \theta) p_D(d) dd \\ &= \arg \max_{\theta} \mathbb{E}[p_{X|D}(x | D; \theta)]. \end{aligned} \quad (33)$$

For the MAP, the analogous deduction leads to

$$\theta^* \in \arg \max_{\theta} \mathbb{E}[p_{X|D}(x | D; \theta)] p_{\Theta}(\theta). \quad (34)$$

Computing the expected values in (33) or in (34) is normally intractable. The standard approach is to use the Expectation Maximization (EM) algorithm [61]. An issue with EM, in addition to a rate of convergence that might be very slow, is that it requires computing in closed form the expected value $\mathbb{E}_{D|X; \bar{\theta}}[\log p_{X,D}(X, D; \theta)]$, which is often not possible. In some cases, one can use Monte Carlo EM [61] to compute it, but with rates of convergence even slower.

The MLE optimization (33) can be viewed as an unconstrained form of (11), which using the multiplicative upper bound in (13) leads to

$$\theta^{\nabla} \in \arg \max_{\theta} \min_d p_{X|D}(x | d; \theta) p_D(d)^{\epsilon} \mathcal{I}_D(\epsilon) \quad (35)$$

or equivalently,

$$\theta^{\nabla} \in \arg \max_{\theta} \min_d \log(p_{X|D}(x | d; \theta) p_D(d)^{\epsilon} \mathcal{I}_D(\epsilon)), \quad (36)$$

which is numerically more stable. For the MAP, one would add $\log p_{\Theta}(\theta)$ to the right hand side of (36). The multiplicative bound is more amenable for the optimization than the additive as it allows to solve (35) in its logarithmic form (36).

Example 3 (Linear measurements with additive Gaussian noise): Let $D \sim \mathcal{N}(0, \sigma_D)$, $N \sim \mathcal{N}(0, \sigma_N)$. Consider T observations of the random variable $X_t = \theta + D_t + N_t$ where θ is the parameter to be estimated. This problem has a closed form solution, which is the empirical average of x_t . Applied to this problem, equation (36) reduces to ¹

$$\theta^{\nabla} \in \arg \max_{\theta} \min_{d_{1:T}} \sum_{t=1}^T -\|x_t - d_t - \theta\|_2^2 \sigma_N^{-1} - \epsilon \|d_t\|_2^2 \sigma_D^{-1} - T \log(1 - \epsilon) - T \log(2\pi\sigma_N). \quad (37)$$

If we take any ϵ such that $\epsilon < -\sigma_N^{-1}/\sigma_D^{-1}$, then the solution is $\frac{1}{T} \sum_{t=1}^T x_t$ which is the same as the exact solution.

Example 4 (Norm measurements with Gaussian disturbances and noise): We have T observations of the random variable $X_i = \|\theta + D_i\|_2 + N_i$ where θ is the parameter to be estimated, $D \sim \mathcal{N}(0, \Sigma_D)$ and $N \sim \mathcal{N}(0, \sigma_N)$. We also have a prior distribution $\Theta \sim \mathcal{N}(\bar{\theta}, \Sigma_{\Theta})$ on θ . Applied to this problem, equation (36) reduces to

$$\theta^{\nabla} \in \arg \max_{\theta} \min_{d_{1:T}} \sum_{t=1}^T -\left\|x_t - \|\theta + d_t\|_2\right\|_{\sigma_N^{-1}}^2 - \epsilon \|d_t\|_{\Sigma_D^{-1}}^2 - \|\theta - \bar{\theta}\|_{\Sigma_{\Theta}^{-1}}^2 - 2T \log(1 - \epsilon) - T \log(2\pi\sigma_N) \quad (38)$$

where we use the notation $\|v\|_Q^2 := v^T Q v$. We take the numerical values $T = 20$, $\Sigma_D = \begin{bmatrix} 2 & -1 \\ -1 & 1 \end{bmatrix}$, $\sigma_N = 1$ $\bar{\theta} = \begin{bmatrix} 1.8 \\ 1.8 \end{bmatrix}$, Σ_{Θ} the identity matrix.

The result of (38) is shown in Table I where we compare it with three other estimators. The first one is what we call naive

θ (actual value)	θ^{∇}	naive MAP	MC MMSE	MC MAP
$\begin{bmatrix} 2 \\ 2 \end{bmatrix}$	$\begin{bmatrix} 1.95 \\ 2.37 \end{bmatrix}$	$\begin{bmatrix} 2.8 \\ 0.21 \end{bmatrix}$	$\begin{bmatrix} 1.82 \\ 1.62 \end{bmatrix}$	$\begin{bmatrix} 2.13 \\ 1.70 \end{bmatrix}$

TABLE I: Comparison between the actual value of θ , of θ^{∇} obtained from (38) and three other estimators.

MAP, where one treats D_1, \dots, D_T not as a latent variable, but as a regular variable that one wants to estimate, *i.e.*,

$$\arg \max_{\theta, d_{1:T}} \sum_{t=1}^T -\left\|x_t - \|\theta + d_t\|_2\right\|_{\sigma_N^{-1}}^2 - \|d_t\|_{\Sigma_D^{-1}}^2 - \|\theta - \bar{\theta}\|_{\Sigma_{\Theta}^{-1}}^2.$$

The second and third are Monte Carlo methods, where we use a Markov Chain Monte Carlo to obtain 10^6 samples from $\Theta | X$, which takes about 30 minutes. Using these sample, the second estimator is the Monte Carlo estimate of the Minimum Mean Square Error (MC MMSE) estimator (*i.e.*, the empirical average of the samples). The third estimator, we use the sample based estimator of the mode described in [62] to compute a Monte Carlo estimate of the MAP (MC MAP).

Our estimator θ^{∇} is significantly closer to real θ and to the MCMC estimate of the MAP than the naive MAP. θ^{∇} is also approximately as distant to the true value of θ as the MMSE estimate. Although none of them is the real MAP, these results suggest that θ^{∇} accurately captures the estimation problem and provides a better result than naively trying to estimate $d_{1:T}$ as in the naive MAP.

C. Bayesian Optimal Experiment Design

The goal in experiment design is to find inputs for an estimation problem that will yield samples that provide "more information per sample". Consider a random vector X with *pdf* $p_X(x; u, \theta)$ where θ is a vector of unknown parameters and u a vector of control decision taking values in $\mathcal{U} \subset \mathbb{R}^N$. The Fisher Information Matrix is

$$\mathcal{FI}(u, \theta) = \mathbb{E} \left[\frac{d \log p_X(X; u, \theta)}{d\theta} \frac{d \log p_X(X; u, \theta)'}{d\theta} \right],$$

where the expected value is taken with respect to X and where we use the denominator-layout notation for the derivatives (producing column vectors). The Cramer-Rao lower bound states that, given any unbiased estimator $\hat{\theta}(u, X)$ of θ , its covariance $\mathbb{E}[(\hat{\theta}(u, X) - \theta)(\hat{\theta}(u, X) - \theta)']$ is lower bounded (in the positive definite matrix sense) by $\mathcal{FI}(u, \theta)^{-1}$. Therefore, if one minimizes (according to some criteria) $\mathcal{FI}(u, \theta)^{-1}$, one will decrease the covariance of any estimator achieving the Cramer-Rao bound.

In Bayesian optimal experiment design one assumes that θ is a realization of an underlying random vector Θ , with *pdf* $p_{\Theta}(\cdot)$, and select u^* to minimize the Bayesian D-optimality (the D stands for determinant), criteria:

$$u^* \in \arg \min_{u \in \mathcal{U}} \mathbb{E}[\log \det(\mathcal{FI}(u, \Theta)^{-1})], \quad (39)$$

where the expected value is taken with respect to Θ . It is shown in [4] that (39) optimizes the gain in the Shannon

¹We refer the reader to Appendix A for the deduction of the penalizing term.

information of the experiment when $\hat{\theta}(u, X)$ is a Gaussian distribution with mean θ and covariance $\mathcal{FI}(u, \theta)^{-1}$. In other words, it designs an experiment that brings more information on average. Alternative Bayesian criteria include A-optimality (the A stands for average), where one wants to find a u^* such that

$$u^* \in \arg \min_{u \in \mathcal{U}} \mathbb{E}[\text{tr}(\mathcal{FI}(u, \Theta)^{-1})]. \quad (40)$$

In this case, (40) minimizes the mean square error of any estimator $\hat{\theta}(u, X)$ that is unbiased and achieves the Cramer-Rao bound.

The experiment design in (39) and (40) is an unconstrained form of (11). Using the additive upper bound in (13) leads to

$$\begin{aligned} V^\Delta &= \min_{u \in \mathcal{U}, \epsilon} \max_{\theta \in \Omega} -\log \det(\mathcal{FI}(u, \theta)) + \epsilon \log p_\Theta(\theta) + \epsilon \mathcal{H}_\Theta \\ V^\nabla &= \max_{\epsilon} \min_{u \in \mathcal{U}, \theta \in \Omega} -\log \det(\mathcal{FI}(u, \theta)) + \epsilon \log p_\Theta(\theta) + \epsilon \mathcal{H}_\Theta. \end{aligned} \quad (41)$$

For Bayesian A-optimality (40), we obtain

$$\begin{aligned} V^\Delta &= \min_{u \in \mathcal{U}, \epsilon} \max_{\theta \in \Omega} \text{tr}(\mathcal{FI}(u, \theta)^{-1}) + \epsilon \log p_\Theta(\theta) + \epsilon \mathcal{H}_\Theta \\ V^\nabla &= \max_{\epsilon} \min_{u \in \mathcal{U}, \theta \in \Omega} \text{tr}(\mathcal{FI}(u, \theta)^{-1}) + \epsilon \log p_\Theta(\theta) + \epsilon \mathcal{H}_\Theta. \end{aligned} \quad (42)$$

Example 5 (Optimal trajectories for thermal air wind detection):

A glider is an air vehicle that flies without propellers, using only wind forces to change its altitude. In order to move up, a glider needs to estimate the location and intensity of the thermal vertical wind that would push it [63]–[65].

Given an air column, a common model for the intensity of the vertical wind speed at position $z = (x, y)$ is

$$w(\bar{w}, \gamma, \bar{z}, z) = \bar{w} e^{-\gamma \|z - \bar{z}\|^2}$$

where $\bar{z} = (\bar{x}, \bar{y})$ denotes the position of the thermal center, \bar{w} the wind speed at the thermal center and $1/\gamma$ the thermal radius. Our goal is to estimate the thermal parameters $\theta = (\bar{w}, \gamma, \bar{z})$ based on noisy measurements of the vertical air speed of the form

$$V_t = w(\bar{w}, \gamma, \bar{z}, z_t) + N_t$$

where z_t is the location where the measurement is taken and N_t are independent zero mean Gaussian distribution with variance σ^2 . The probability density function for T measurements $v = (v_1, \dots, v_T)$ is given by

$$p_V(v; \theta) = \frac{1}{(2\pi)^{T/2} \sigma^T} e^{-\frac{\sum_{t=1}^T (v_t - w_t)^2}{2\sigma^2}}$$

where $w_t = w(\bar{w}, \gamma, \bar{z}, z_t)$. The Fisher Information matrix associated to the estimation of θ is given by,

$$\begin{aligned} \mathcal{FI}(z_{1:T}, \theta) &= \mathbb{E} \left[\frac{d \log p_V(V; \theta)}{d\theta} \frac{d \log p_V(V; \theta)'}{d\theta} \right] \\ &= \frac{1}{\sigma^4} \mathbb{E} \left[\sum_{t=1}^T \sum_{l=1}^T (V_t - w_t)(V_l - w_l) \frac{dw_t}{d\theta} \frac{dw_l'}{d\theta} \right] \\ &= \frac{1}{\sigma^2} \sum_{t=1}^T \frac{dw_t}{d\theta} \frac{dw_t'}{d\theta} \end{aligned}$$

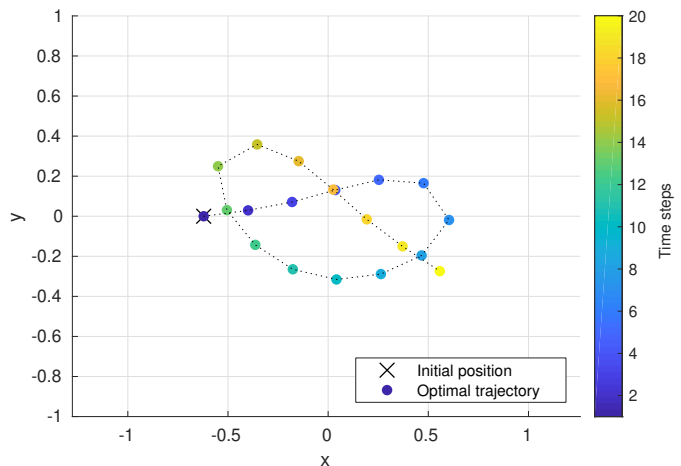


Fig. 3: Optimal trajectory for the Bayesian experiment design for detecting the parameters of a vertical thermal air flow

where

$$\begin{aligned} \frac{dw_t}{d\theta} &= \left[\frac{\partial w_t}{\partial \bar{w}}, \frac{\partial w_t}{\partial \gamma}, \frac{\partial w_t}{\partial \bar{z}} \right]' \\ &= e^{-\gamma \|z_t - \bar{z}\|^2} \left[1, -\bar{w} \|z_t - \bar{z}\|^2, -\bar{w} \gamma (z_t - \bar{z})' \right]'. \end{aligned}$$

Given prior distributions on \bar{w} , γ and \bar{z} , we want to find the measurement points z_1, z_2, \dots, z_T that minimize (39) subject to the constraint that the distance between two consecutive z_t should be no larger than Δz . As the problem is rotationally symmetric, we fix the y coordinate of the first point to be 0.

We assign the following prior distributions. Both \bar{w} and γ follow a Gamma distribution with parameters respectively $(\alpha_{\bar{w}}, \beta_{\bar{w}})$ and $(\alpha_\gamma, \beta_\gamma)$ and the thermal center \bar{z} follows a zero mean Gaussian distribution with covariance $\Sigma_{\bar{z}}$.

We take the following numerical values. The number of measurements is $T = 20$. The parameters of the priors are $\alpha_{\bar{w}} = \alpha_\gamma = 1.25$, $\beta_{\bar{w}} = \beta_\gamma = 0.25$ and $\Sigma_{\bar{z}} = 0.1I$. The maximum displacement between two sampling points is $\Delta z = 0.05$. The problem is highly nonconvex, requiring multiple initializations. For the lower bound, it takes about 6.56 seconds to run 100 optimizations with a random walk initialization, obtaining the lower bound $V^\nabla = -14.27$. For the upper bound it takes about 8.92 seconds to run 100 optimizations with random walk initialization, obtaining the upper bound $V^\Delta = 95.41$. Using Monte Carlo integration, as discussed in Section III-A, we obtain that the expected value of the log determinant of the Fisher Information Matrix given the trajectory is 1.203. The optimal trajectory can be seen in Figure 3.

V. CONCLUSIONS AND FUTURE WORK

We presented a general method to bound the expected value of any random variable with known probability density function. Stochastic programming is the main application of the bounds, where they can be used to determine an optimizer which has performance guarantees and satisfies inequality constraints. We illustrate the results with applications to finite-horizon stochastic control, estimation with latent variables and

experiment design. The numerical results in these applications show that optimizing the bound lead to solutions close to the optimal. They also suggest that even when the bounds are not tight, the argument that minimizes the upper bound is close to the one that minimizes the stochastic programming problem.

There are many future work directions to be considered. On the bounds themselves, most of the properties were determined for the additive and multiplicative bound, but other versions of the bounds could unlock other applications. The connection between the bounds we developed and distributionally robust optimization remains to be further understood, in particular for which kind of problem which approach is more suited. On obtaining solutions to the minmax optimization, an area for future research motivated by [15]–[17] arises from replacing the essential suprema used in the upper bound in (13) by maxima over independent samples of the random variable D and establishing sample complexity bounds to guarantee that the resulting optimization still provides an upper bound with high probability.

In terms of stochastic control, an evident extension would be stochastic model predictive control. In the estimation section, it would be interesting to study the asymptotic properties of the bound. As for new applications, machine learning is an area of significant potential. In particular, this method could either be used to accelerate the training of Neural Networks when there is a partial knowledge of the underlying model or in adversarial training.

ACKNOWLEDGMENT

The authors thank the anonymous reviewers for their several suggestions that added much value to the paper, especially regarding the content of Sections II-B, III-A and III-B.

APPENDIX

A. Penalization term for common distributions

a) *Gaussian distribution:* The probability density function of a Gaussian Distribution with mean μ and covariance matrix Σ is

$$p_D(d) = \det(2\pi\Sigma)^{-1/2} \exp\left(-\frac{1}{2}\|d - \mu\|_{\Sigma^{-1}}^2\right)$$

where we use the notation $\|v\|_Q^2 := v'Qv$.

For the additive bound, $\mathcal{H}_D := \mathbb{E}[-\log p_D(D)] = \frac{1}{2} \log \det(2\pi e\Sigma)$, therefore the penalization term simplifies to

$$\log p_D(d) + \mathcal{H}_D = -\frac{1}{2}\|d - \mu\|_{\Sigma^{-1}}^2 + \frac{1}{2}M$$

where M is the dimension of D .

For the multiplicative bound, $\mathcal{I}_D(\epsilon) := \mathbb{E}[p_D(D)^{-\epsilon}] = \det(2\pi\Sigma)^{\epsilon/2} (1-\epsilon)^{-M/2}$ if $\epsilon < 1$ and $+\infty$ otherwise, therefore for $\epsilon < 1$ the penalization terms simplifies to

$$p_D(d)^\epsilon \mathcal{I}_D(\epsilon) = \exp\left(-\frac{1}{2}\epsilon\|d - \mu\|_{\Sigma^{-1}}^2\right) (1-\epsilon)^{-M/2}$$

b) *Uniform distribution:* If D is a Uniform distribution over a bounded support \mathcal{D} , its pdf is $p_D(d) = \mathcal{V}_D^{-1} \mathbb{1}_{\mathcal{D}}(d)$ where $\mathbb{1}_{\mathcal{D}}(\cdot)$ is the indicator function of \mathcal{D} and $\mathcal{V}_D = \mathbb{E}[\mathbb{1}_{\mathcal{D}}(D)]$ is the volume of \mathcal{D} .

For the additive bound, $\mathcal{H}_D = \mathbb{E}[-\log(\mathbb{1}_{\mathcal{D}}(D))] + \log(\mathcal{V}_D) = \mathbb{E}[0] + \log(\mathcal{V}_D)$, therefore the penalization terms simplifies to $\log p_D(d) + \mathcal{H}_D = 0 \forall d \in \mathcal{D}$.

For the multiplicative bound, $\mathcal{I}_D(\epsilon) = \mathbb{E}[(\mathcal{V}_D)^\epsilon \mathbb{1}_{\mathcal{D}}(D)^{-\epsilon}] = (\mathcal{V}_D)^\epsilon$, therefore the penalization term simplifies to $p_D(d)^\epsilon \mathcal{I}_D(\epsilon) = 1 \forall d \in \mathcal{D}$.

B. Proofs of Section II

To prove the results that follow, we need the following properties of the essential supremum and infimum which we state without a proof.

Lemma 3: Given two random variables X and Y then

$$\begin{aligned} X \stackrel{\text{wpo}}{\geq} Y &\Rightarrow \text{ess inf } X \geq \text{ess inf } Y \\ X \stackrel{\text{wpo}}{\leq} Y &\Rightarrow \text{ess sup } X \leq \text{ess sup } Y \\ \text{ess inf}(X + Y) &\geq \text{ess inf } X + \text{ess inf } Y \\ \text{ess sup}(X + Y) &\leq \text{ess sup } X + \text{ess sup } Y \quad \square \end{aligned}$$

Proof of Proposition 1. Take $\epsilon_1, \epsilon_2 \in \mathbb{R}$ such that $J^*(\epsilon_1), J^*(\epsilon_2) < +\infty$ and $\lambda \in [0, 1]$

$$\begin{aligned} &J^*(\lambda\epsilon_1 + (1-\lambda)\epsilon_2) \\ &= \text{ess sup } V(D) + (\epsilon_1\lambda + \epsilon_2(1-\lambda)) \log p_D(D) \\ &= \text{ess sup}(\lambda + 1 - \lambda)V(D) + (\epsilon_1\lambda + \epsilon_2(1-\lambda)) \log p_D(D) \\ &\leq \lambda \text{ess sup } V(D) + \epsilon_1 \log p_D(D) \\ &\quad + (1-\lambda) \text{ess sup } V(D) + \epsilon_2 \log p_D(D) < \infty \end{aligned}$$

where the inequality follows from Lemma 3. This establishes that the additive upper bound is convex in ϵ and that $J^*(\epsilon)$ is finite on a convex set.

For the multiplicative bound, it remains to show that $\mathcal{I}_D(\epsilon)$ is log convex: take $\epsilon_1, \epsilon_2 \in \mathbb{R}$ such that $\mathcal{I}_D(\epsilon_1), \mathcal{I}_D(\epsilon_2)$ are finite and $\lambda \in [0, 1]$. By applying Hölder's inequality we obtain

$$\begin{aligned} &\mathbb{E}[p_D(D)^{-\lambda\epsilon_1} p_D(D)^{-(1-\lambda)\epsilon_2}] \\ &\leq \left(\mathbb{E}[p_D(D)^{-\lambda\epsilon_1/\lambda}]\right)^\lambda \left(\mathbb{E}[p_D(D)^{-(1-\lambda)\epsilon_2/(1-\lambda)}]\right)^{(1-\lambda)} \\ &= \left(\mathbb{E}[p_D(D)^{-\epsilon_1}]\right)^\lambda \left(\mathbb{E}[p_D(D)^{-\epsilon_2}]\right)^{(1-\lambda)} \end{aligned}$$

which establishes log convexity. \blacksquare

C. Necessary and sufficient conditions for finite bounds

Consider a constant $\gamma > 0$ sufficiently small so that $\mathbb{P}(p_D(D) > \gamma) > 0$. We say a measurable function $f(\cdot)$ is γ -essentially upper bounded if

$$\text{ess sup}[f(D) \mid p_D(D) > \gamma] < \infty,$$

γ -essentially lower bounded if

$$\text{ess inf}[f(D) \mid p_D(D) > \gamma] > -\infty,$$

and γ -essentially bounded if it is both γ -essentially upper and lower bounded.

Theorem 4 (Finite bounds): Suppose that $p_D(\cdot)$ is γ -essentially upper bounded and let $c \in (0, 1/\gamma)$ be any constant for which □

$$\text{ess sup} [p_D(D) \mid p_D(D) > \gamma] \leq 1/c, \quad (43)$$

and ϵ an arbitrary finite constant. Regarding the additive bound: Assuming that $V(\cdot)$ is γ -essentially lower bounded, then

$$\begin{aligned} p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess inf} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] > \epsilon \\ \Rightarrow \text{ess inf} (V(D) + \epsilon \log p_D(D)) > -\infty. \end{aligned} \quad (44)$$

Conversely,

$$\begin{aligned} \text{ess inf} (V(D) + \epsilon \log p_D(D)) > -\infty \\ \Rightarrow p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \exists L > 0 : \\ \text{ess inf} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] \geq \epsilon + \frac{L}{\log c \gamma}. \end{aligned} \quad (45)$$

Assuming that $V(\cdot)$ γ -essentially upper bounded, then

$$\begin{aligned} p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess sup} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] < \epsilon \\ \Rightarrow \text{ess sup} (V(D) + \epsilon \log p_D(D)) < +\infty \end{aligned} \quad (46)$$

Conversely,

$$\begin{aligned} \text{ess sup} (V(D) + \epsilon \log p_D(D)) < \infty \\ \Rightarrow p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \exists L > 0 : \\ \text{ess sup} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] \leq \epsilon - \frac{L}{\log c \gamma}. \end{aligned} \quad (47)$$

Regarding the multiplicative bound: Assuming that $\log V(\cdot)$ is γ -essentially lower bounded, then

$$\begin{aligned} p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess inf} \left[\frac{-\log V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] > \epsilon \\ \Rightarrow \text{ess inf} (\log V(D) + \epsilon \log p_D(D)) > -\infty. \end{aligned}$$

Conversely,

$$\begin{aligned} \text{ess inf} (\log V(D) + \epsilon \log p_D(D)) > -\infty \\ \Rightarrow p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \exists L > 0 : \\ \text{ess inf} \left[\frac{-\log V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] \geq \epsilon + \frac{L}{\log c \gamma}. \end{aligned} \quad (48)$$

Assuming that $\log V(\cdot)$ is γ -essentially upper bounded, then

$$\begin{aligned} p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \text{ess sup} \left[\frac{-\log V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] < \epsilon \\ \Rightarrow \text{ess sup} (\log V(D) + \epsilon \log p_D(D)) < +\infty \end{aligned}$$

Conversely,

$$\begin{aligned} \text{ess sup} (\log V(D) + \epsilon \log p_D(D)) < \infty \\ \Rightarrow p_D(D) \stackrel{\text{wpo}}{>} \gamma \text{ or } \exists L > 0 : \\ \text{ess sup} \left[\frac{-\log V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] \leq \epsilon - \frac{L}{\log c \gamma}. \end{aligned} \quad (49)$$

Proof. We will prove the theorem for the additive lower bound [i.e., (44) and (45)]. The proof for the other bounds can be obtained in an analogous way.

To prove (44), we note that since $V(\cdot)$ is γ -essentially lower bounded there exists a finite constant L such that

$$\text{ess inf} [V(D) \mid p_D(D) > \gamma] \geq L.$$

In view of this and (43), we have that

$$P(V(D) \geq L, p_D(D) \leq 1/c \mid p_D(D) > \gamma) = 1. \quad (50)$$

Since

$$\begin{aligned} V(D) \stackrel{\text{wpo}}{\geq} L, p_D(D) \stackrel{\text{wpo}}{\leq} 1/c, p_D(D) \stackrel{\text{wpo}}{>} \gamma \\ \Rightarrow V(D) + \epsilon \log c p_D(D) \stackrel{\text{wpo}}{\geq} L^* > -\infty, \end{aligned}$$

with $L^* := L - |\epsilon| |\log c \gamma|$, we conclude from (50) that

$$P(V(D) + \epsilon \log c p_D(D) \geq L^* \mid p_D(D) > \gamma) = 1. \quad (51)$$

In case $p_D(D) \stackrel{\text{wpo}}{>} \gamma$, we conclude that the corresponding unconditional probability satisfies the same bound and (44) follows. Otherwise,

$$\text{ess inf} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] > \epsilon$$

implies that

$$P\left(\frac{-V(D)}{\log c p_D(D)} \geq \epsilon \mid p_D(D) \leq \gamma\right) = 1. \quad (52)$$

Since

$$p_D(D) \leq \gamma \stackrel{\text{wpo}}{\implies} \log c p_D(D) \leq \log c \gamma < 0, \quad (53)$$

we also conclude from (52) that

$$P(V(D) + \epsilon \log c p_D(D) \geq 0 \mid p_D(D) \leq \gamma) = 1. \quad (54)$$

Combining (54) and (51), we conclude that the corresponding unconditional probability satisfies

$$P(V(D) + \epsilon \log c p_D(D) \geq \min\{0, L^*\}) = 1,$$

from which (44) follows.

To prove (45), we use the fact that $\text{ess inf} (V(D) + \epsilon \log p_D(D)) > -\infty$ implies that there exists some finite scalar $L > 0$, for which

$$P(V(D) + \epsilon \log c p_D(D) \geq -L) = 1. \quad (55)$$

When $p_D(D) \stackrel{\text{wpo}}{>} \gamma$ the implication in (45) is tautologically true, so we focus our attention on the case $P(p_D(D) \leq \gamma) > 0$, for which (55) implies that

$$P(V(D) + \epsilon \log c p_D(D)) \geq -L \mid p_D(D) \leq \gamma = 1. \quad (56)$$

Using (53), we conclude that

$$\begin{aligned} V(D) + \epsilon \log c p_D(D), p_D(D) \leq \gamma \\ \stackrel{\text{wpo}}{\implies} \frac{-V(D)}{\log c p_D(D)} \geq \epsilon + \frac{L}{\log c \gamma} \end{aligned}$$

and therefore (56) implies that

$$P\left(\frac{-V(D)}{\log c p_D(D)} \geq \epsilon + \frac{L}{\log c \gamma} \mid p_D(D) \leq \gamma\right) = 1. \quad (57)$$

This shows that

$$\text{ess inf} \left[\frac{-V(D)}{\log c p_D(D)} \mid p_D(D) \leq \gamma \right] \geq \epsilon + \frac{L}{\log c \gamma},$$

which completes the proof of the implication in (45). ■

REFERENCES

- [1] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory (v. 1)*. Prentice Hall, 1993.
- [2] J. C. Spall, *Introduction to stochastic search and optimization: estimation, simulation, and control*, vol. 65. John Wiley & Sons, 2005.
- [3] A. C. Atkinson, *Optimum Experimental Designs*. No. 8 in Oxford Statistical Science Series, Clarendon Press ; Oxford University Press, 1992.
- [4] K. Chaloner and I. Verdinelli, “Bayesian experimental design: A review,” *Statistical Science*, vol. 10, no. 3, pp. 273–304, 1995.
- [5] G. Monegato, “An overview of the computational aspects of kronrod quadrature rules,” *Numerical Algorithms*, vol. 26, no. 2, pp. 173–196, 2001.
- [6] M. Evans and T. Swartz, *Approximating integrals via Monte Carlo and deterministic methods*, vol. 20, ch. 5. OUP Oxford, 2000.
- [7] P. A. Samuelson, “The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments,” in *Stochastic Optimization Models in Finance*, pp. 215–220, Elsevier, 1975.
- [8] O. Loistl, “The erroneous approximation of expected utility by means of a Taylor’s series expansion: analytic and computational results,” *The American Economic Review*, vol. 66, no. 5, pp. 904–910, 1976.
- [9] W. Hlawitschka, “The empirical nature of Taylor-series approximations to expected utility,” *The American Economic Review*, vol. 84, no. 3, pp. 713–719, 1994.
- [10] C. Robert and G. Casella, *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- [11] R. W. Butler, *Saddlepoint Approximations with Applications*. Cambridge University Press, 1 edition ed., 2007.
- [12] A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello, “The sample average approximation method for stochastic discrete optimization,” *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2002.
- [13] S. Kim, R. Pasupathy, and S. G. Henderson, “A guide to sample average approximation,” in *Handbook of simulation optimization*, pp. 207–243, Springer, 2015.
- [14] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*. Society for Industrial and Applied Mathematics, second edition ed., 2014.
- [15] G. Calafiore and M. C. Campi, “Uncertain convex programs: randomized solutions and confidence levels,” *Mathematical Programming*, vol. 102, no. 1, pp. 25–46, 2005.
- [16] G. Calafiore and M. C. Campi, “The scenario approach to robust control design,” *IEEE Transactions on automatic control*, vol. 51, no. 5, pp. 742–753, 2006.
- [17] P. Mohajerin Esfahani, T. Sutter, and J. Lygeros, “Performance Bounds for the Scenario Approach and an Extension to a Class of Non-Convex Programs,” *IEEE Transactions on Automatic Control*, vol. 60, pp. 46–58, Jan. 2015. Conference Name: IEEE Transactions on Automatic Control.
- [18] A. Ben Tal and A. Nemirovski, “Robust optimization methodology and applications,” *Mathematical Programming*, may 2002.
- [19] V. Gabrel, C. Murat, and A. Thiele, “Recent advances in robust optimization: An overview,” *European Journal of Operational Research*, vol. 235, pp. 471–483, June 2014.
- [20] D. A. Copp and J. P. Hespanha, “Simultaneous nonlinear model predictive control and state estimation,” *Automatica*, vol. 77, pp. 143–154, 2017.
- [21] R. Huang, B. Xu, D. Schuurmans, and C. Szepesvari, “Learning with a Strong Adversary,” nov 2015.
- [22] U. Shaham, Y. Yamada, and S. Negahban, “Understanding adversarial training: Increasing local stability of supervised models through robust optimization,” *Neurocomputing*, vol. 307, pp. 195–204, sep 2018.
- [23] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards Deep Learning Models Resistant to Adversarial Attacks,” *arXiv:1706.06083 [cs, stat]*, Sept. 2019. arXiv: 1706.06083.
- [24] C. Jin, P. Netrapalli, and M. I. Jordan, “What is Local Optimality in Nonconvex-Nonconcave Minimax Optimization?,” feb 2019-02-01.
- [25] K. Margellos, P. Goulart, and J. Lygeros, “On the road between robust optimization and the scenario approach for chance constrained optimization problems,” *IEEE Transactions on Automatic Control*, vol. 59, no. 8, pp. 2258–2263, 2014.
- [26] X. Chen, M. Sim, and P. Sun, “A robust optimization perspective on stochastic programming,” *Operations Research*, vol. 55, no. 6, pp. 1058–1071, 2007.
- [27] C. Bandi and D. Bertsimas, “Tractable stochastic analysis in high dimensions via robust optimization,” *Mathematical programming*, vol. 134, no. 1, pp. 23–70, 2012.
- [28] P. M. Esfahani and D. Kuhn, “Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations,” *Mathematical Programming*, vol. 171, no. 1, pp. 115–166, 2018.
- [29] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, “Wasserstein distributionally robust optimization: Theory and applications in machine learning,” in *Operations Research & Management Science in the Age of Analytics*, pp. 130–166, INFORMS, 2019.
- [30] E. Delage and Y. Ye, “Distributionally robust optimization under moment uncertainty with application to data-driven problems,” *Operations research*, vol. 58, no. 3, pp. 595–612, 2010.
- [31] W. Wiesemann, D. Kuhn, and M. Sim, “Distributionally robust convex optimization,” *Operations Research*, vol. 62, no. 6, pp. 1358–1376, 2014.
- [32] A. Ben-Tal, D. den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen, “Robust Solutions of Optimization Problems Affected by Uncertain Probabilities,” *Management Science*, vol. 59, no. 2, pp. 341–357, 2013. Publisher: INFORMS.
- [33] J. Goh and M. Sim, “Distributionally robust optimization and its tractable approximations,” *Operations research*, vol. 58, no. 4-part-1, pp. 902–917, 2010.
- [34] H. Rahimian and S. Mehrotra, “Distributionally Robust Optimization: A Review,” p. 90, Aug. 2019.
- [35] R. Chinchilla and J. P. Hespanha, “Optimization-based estimation of expected values with application to stochastic programming,” in *Proc. of the 58th Conf. on Decision and Contr.*, Dec. 2019.
- [36] Springer Verlag and European Mathematical Society, “Right-ordered group,” *Encyclopedia of Mathematics*. URL: http://www.encyclopediaofmath.org/index.php?title=Right-ordered_group&oldid=43497. Accessed on 2020-04-08.
- [37] V. Vapnik, *The nature of statistical learning theory*. Statistics for Engineering and Information Science, New York: Springer, 2nd ed., 1999.
- [38] A. Shapiro, “Computational complexity of stochastic programming: Monte Carlo sampling approach,” in *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures*, pp. 2979–2995, World Scientific, 2010.
- [39] R. H. Byrd, J. C. Gilbert, and J. Nocedal, “A trust region method based on interior point techniques for nonlinear programming,” *Mathematical Programming*, vol. 89, pp. 149–185, Nov. 2000.
- [40] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical Programming*, vol. 106, pp. 25–57, Mar. 2006.
- [41] J. P. Hespanha, “TensCalc: A toolbox to generate fast code to solve nonlinear constrained minimizations and compute Nash equilibria.” tech. rep., Center for Control, Dynamical Systems and Computation; University of California, Santa Barbara, <https://github.com/hespanha/tenscalc>, 2017.
- [42] A. Ben-Tal and A. Nemirovski, “Robust optimization - methodology and applications,” *Mathematical Programming*, vol. 92, pp. 453–480, May 2002.
- [43] A. Ben-Tal, L. E. Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton University Press, Aug. 2009. Google-Books-ID: DttjR7lpjUEC.
- [44] D. Bertsimas, D. B. Brown, and C. Caramanis, “Theory and Applications of Robust Optimization,” *SIAM Review*, vol. 53, pp. 464–501, Jan. 2011. Publisher: Society for Industrial and Applied Mathematics.
- [45] A. Mutapcic and S. Boyd, “Cutting-set methods for robust convex optimization with pessimizing oracles,” *Optimization Methods and Software*, vol. 24, pp. 381–406, June 2009.
- [46] C. Daskalakis and I. Panageas, “The Limit Points of (Optimistic) Gradient Descent in Min-Max Optimization,” p. 11.
- [47] J. N. Foerster, R. Y. Chen, M. Al-Shedivat, S. Whiteson, P. Abbeel, and I. Mordatch, “Learning with Opponent-Learning Awareness,” *arXiv:1709.04326 [cs]*, Sept. 2018. arXiv: 1709.04326.

- [48] K. Ebrahimi, N. Elia, and U. Vaidya, "A continuous time dynamical system approach for solving robust optimization," in *2019 18th European Control Conference (ECC)*, pp. 1479–1485.
- [49] Y. Wang, G. Zhang, and J. Ba, "On solving minimax optimization locally: A follow-the-ridge approach,"
- [50] M. Razaviyayn, T. Huang, S. Lu, M. Nouiehed, M. Sanjabi, and M. Hong, "Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances," vol. 37, no. 5, pp. 55–66. Conference Name: IEEE Signal Processing Magazine.
- [51] H. Rafique, M. Liu, Q. Lin, and T. Yang, "Non-convex min-max optimization: Provable algorithms and applications in machine learning,"
- [52] G. Zhang, K. Wu, P. Poupard, and Y. Yu, "Newton-type methods for minimax optimization,"
- [53] T. Lin, C. Jin, and M. I. Jordan, "Near-optimal algorithms for minimax optimization,"
- [54] S. Lu, I. Tsaknakis, M. Hong, and Y. Chen, "Hybrid block successive approximation for one-sided non-convex min-max problems: Algorithms and applications," vol. 68, pp. 3676–3691. Conference Name: IEEE Transactions on Signal Processing.
- [55] E. V. Mazumdar, M. I. Jordan, and S. S. Sastry, "On finding local nash equilibria (and only local nash equilibria) in zero-sum games,"
- [56] M. Liu, H. Rafique, Q. Lin, and T. Yang, "First-order convergence theory for weakly-convex-weakly-concave min-max problems," version: 3.
- [57] E. A. Ok, *Real analysis with economic applications / Efe A. Ok*. University Press.
- [58] D. P. Bertsekas, *Nonlinear Programming*. Athena, 1999. 2nd Edition, 2nd edition ed., 1999.
- [59] S. Quintero, D. A. Copp, and J. P. Hespanha, "Robust coordination of small UAVs for vision-based target tracking using output-feedback MPC with MHE," in *Cooperative Control of Multi-agent Systems: Theory and Applications* (E. Garcia, Y. Wang, D. Casbeer, and F. Zhang, eds.), Hoboken, NJ: Wiley & Sons, 2017.
- [60] D. A. Copp, *Simultaneous Nonlinear Model Predictive Control and State Estimation: Theory and Applications*. PhD thesis, University of California, Santa Barbara, CA, USA, Nov. 2016.
- [61] G. J. McLachlan, *The EM algorithm and extensions*. Wiley series in probability and statistics, Wiley-Interscience, 2nd ed., 2008.
- [62] C. Abraham, G. Biau, and B. Cadre, "Simple Estimation of the Mode of a Multivariate Density," *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, vol. 31, no. 1, pp. 23–34, 2003. Publisher: [Statistical Society of Canada, Wiley].
- [63] K. Andersson, I. Kammer, V. Dobrokhodov, and V. Cichella, "Thermal Centering Control for Autonomous Soaring; Stability Analysis and Flight Test Results," *Journal of Guidance, Control, and Dynamics*, vol. 35, pp. 963–975, may 2012.
- [64] D. Edwards, "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, American Institute of Aeronautics and Astronautics, aug 2008.
- [65] V. N. Dobrokhodov, N. Camacho, and K. D. Jones, "Cooperative Autonomy of Multiple Solar-Powered Thermaling Gliders," *IFAC Proceedings Volumes*, vol. 47, pp. 1222–1227, jan 2014.



João P. Hespanha received his Ph.D. degree in electrical engineering and applied science from Yale University, New Haven, Connecticut in 1998. From 1999 to 2001, he was Assistant Professor at the University of Southern California, Los Angeles. He moved to the University of California, Santa Barbara in 2002, where he currently holds a Professor position with the Department of Electrical and Computer Engineering.

Dr. Hespanha is the recipient of the Yale University's Henry Prentiss Becton Graduate Prize for exceptional achievement in research in Engineering and Applied Science, a National Science Foundation CAREER Award, the 2005 best paper award at the 2nd Int. Conf. on Intelligent Sensing and Information Processing, the 2005 Automatica Theory/Methodology best paper prize, the 2006 George S. Axelby Outstanding Paper Award, and the 2009 Ruberti Young Researcher Prize. Dr. Hespanha is a Fellow of the International Federation of Automatic Control (IFAC) and of the IEEE. He was an IEEE distinguished lecturer from 2007 to 2013.

His current research interests include hybrid and switched systems; multi-agent control systems; game theory; optimization; distributed control over communication networks (also known as networked control systems); the use of vision in feedback control; stochastic modeling in biology; and network security.



Raphael Chinchilla received a B.Eng. from Universidade de São Paulo, Brazil, a M.Sc. from Université Paris-Saclay, France, and a Diplôme d'Ingénieur (M.Sc.Eng) from Télécom Paris, France, all in electrical engineering.

He currently is a Ph.D candidate at the Center for Control, Dynamical Systems and Computation at the University of California, Santa Barbara. He researches nonlinear robust optimization, developing new algorithms and studying how robust optimization can be used to model problems. Applications of

his research include control, estimation, stochastic optimization and machine learning.

Chinchilla has been an active member in his academic communities, having served as a student representative in multiple committees in all universities he attended. He was the first international student elected to the Graduate Student Association at UCSB, as Vice President of Academic Affairs, and was a Director of Academic Affairs at the Electrical Engineering Students Association (CEE Poli) at USP.