

A Kernel Mean Embedding Approach to Reducing Conservativeness in Stochastic Programming and Control

Jia-Jie Zhu

JZHU@TUEBINGEN.MPG.DE

Empirical Inference Department

Max Planck Institute for Intelligent Systems, Tübingen, Germany

Moritz Diehl

MORITZ.DIEHL@IMTEK.UNI-FREIBURG.DE

Department of Microsystems Engineering

University of Freiburg, Freiburg, Germany

Bernhard Schölkopf

BS@TUEBINGEN.MPG.DE

Empirical Inference Department

Max Planck Institute for Intelligent Systems, Tübingen, Germany

Editors: A. Bayen, A. Jadbabaie, G. J. Pappas, P. Parrilo, B. Recht, C. Tomlin, M. Zeilinger

Abstract

In this paper, we apply kernel mean embedding methods to sample-based stochastic optimization and control. Specifically, we use the reduced-set expansion method as a way to discard sampled scenarios. The effect of such constraint removal is improved optimality and decreased conservativeness. This is achieved by solving a distributional-distance-regularized optimization problem. We demonstrated this optimization formulation is well-motivated in theory, computationally tractable, and effective in numerical algorithms.

Keywords: Stochastic Control, Stochastic Programming, Kernel Methods, Robust Optimization, Data-Driven Optimization

1. Introduction

Robustification against uncertain events is at the core of modern optimization and control. From the classic S-lemma to recent advances in distributionally robust optimization, we have witnessed computational tools giving rise to new robustification designs. The classic “worst-case” approach sets out to robustify constraints against all realizations of disturbances in a mathematical model, resulting in the often over-conservativeness. Consider the illustrative example of a sample-based approach to solving constrained stochastic control problem in Figure 1. In this case, we sample different numbers of realizations of the uncertainty in (left) and (right), and then solve the control problem under those realized scenarios. Intuitively, one may expect the controller associated with more scenarios in (left) to be more robust against constraint violation than (right). However, this results in conservative designs which may be reflected in high cost.

The central idea of modern data-driven robust optimization, in non-technical terms, is to use data samples to form empirical understanding of the true distribution. Then, one only seeks to robustify against this empirical understanding instead of the whole distribution support. One concrete relevance to our discussion, for example in Figure 1, is that fewer scenarios translate to fewer constraints, which in turn lead to reduced cost. Therefore, this is a *trade-off between optimality and feasibility*.




Figure 1: Steering the state trajectories (blue) without crossing constraints (red). Two figures depict optimization over different number of scenarios ($N = 100$ vs. $N = 17$).

In this paper, we show that constraint removal can be formulated as an optimization problem aiming to form a new distribution close to the empirical data distribution. Our contributions are **(1)** We formulate the constraint removal in stochastic programming and control as a tractable convex optimization problem with reproducing kernel Hilbert space (RKHS)-distance regularization or constraint. This formulation is well motivated in theory and effective in numerical studies. **(2)** To our knowledge, this is the first use of RKHS-embedding reduced-set method in stochastic optimization and scenario approaches to control.

Notation. In this work, symbol \mathcal{H} denotes a reproducing kernel Hilbert space (RKHS). We write $\xi \sim P$ to denote that the random variable or vector (RV) ξ follows the distribution law P . By empirical distribution of the data, we mean the convex combination of Dirac-measures of the seen data $P_{\text{data}} := \frac{1}{N} \sum_i^N \delta(\xi_i)$ where $\{\xi_i\}_{i=1}^N$ is the data set.

2. Background & related work

2.1. Stochastic programming and scenario optimization for control

In this paper, the problem of interest is the (chance-constrained) stochastic programming (SP; also known as stochastic optimization) in the following canonical formulation.

$$\underset{x}{\text{minimize}} \mathbb{E}[F(x, \xi)] \quad \text{subject to} \quad \Pr\{C(x, \xi) \leq 0\} \geq 1 - \alpha. \quad (1)$$

As ξ is assumed to be an RV, program (1) may be intuitively understood as making decision x under uncertainty originated from ξ . We consider the following sample-based SP (a.k.a. scenario approach).

Suppose we have a set of i.i.d. realizations $\{\xi_i\}_{i=1}^N$ of ξ , we solve the sample-based program

$$\underset{x}{\text{minimize}} \mathbb{E}[F(x, \xi)] \quad \text{subject to} \quad C(x, \xi_i) \leq 0 \text{ for } i = 1, \dots, N. \quad (2)$$

If F and C are convex in x , measurable in ξ , it can be shown that this formulation is a convex approximation to the original SP (1). As $N \rightarrow \infty$, the solution recovers that of the SP with level

$\alpha = 0$. However, with a large N , the solution to (2) is overly conservative—it aims to satisfy the constraints *almost everywhere* in the distribution of ξ . Therefore, the size N trades off the conservativeness with constraint-satisfaction. Extensive research (e.g., [Calafiore and Campi \(2006\)](#); [Dentcheva et al. \(2000\)](#); [Luedtke et al. \(2010\)](#)) has focused on approaches to remove a subset of sampled constraints to reduce conservativeness of the solution.

Relevant to this paper, [Campi and Garatti \(2011\)](#) established guaranteed bounds for the constraint satisfaction probability and the number of removed constraints κ . Our method is built upon their sampling-and-discarding framework. [Campi and Caré \(2013\)](#) used $l1$ regularization to encourage sparsity in decision variables, which is different from our sparsity in RKHS expansion terms. For readers who are interested in sample-based stochastic programming, good text references are given by Ch.5 of [Shapiro et al. \(2009\)](#) and Ch.9 of [Birge and Louveaux \(2011\)](#).

2.2. Reproducing kernel Hilbert space (RKHS) embeddings

A positive definite kernel is a real-valued bivariate, symmetric function $k(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that $\sum_{i,j=1}^n \alpha_i \alpha_j k(x_i, x_j) \geq 0$ for any $n \in \mathbb{N}$, $(\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$, and $(x_1, \dots, x_n) \in \mathcal{X}^n$. One may intuitively think $k(x, x')$ as a generalized similarity measure (inner product) between x and x' after mapping them into the feature space \mathcal{H} , $k(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$. We refer to ϕ as *feature map* associated with the kernel k , and \mathcal{H} the associated RKHS. A canonical kernel is the Gaussian kernel $k(x, x') = \exp\left(-\frac{1}{2\sigma^2} \|x - x'\|_2^2\right)$ where $\sigma > 0$ is a bandwidth parameter.

RKHS embedding, or *kernel mean embedding* (KME) [[Smola et al. \(2007\)](#)] maps probability distributions to (deterministic) elements of a Hilbert space. Mathematically, the KME of a random variable X is given by the function $\mu_X(\cdot) = \int \langle \phi(x), \phi(\cdot) \rangle dP(x)$, which is a member of the RKHS. For example, the RKHS associated with the second-order polynomial kernel consists of quadratic functions whose coefficients preserve statistical mean and variance. Gaussian kernel embeddings, on the other hand, preserve richer information up to infinite order.

Reduced-set expansion method using RKHS embeddings. Given a data set $x_{i=1}^N$, the sample-based KME is given by $\hat{\mu}_x = \sum_{i=1}^N \alpha_i \phi(x_i)$, where one can simply choose $\alpha_i = \frac{1}{N}$. It has been shown that one may use fewer than the total N data sample to represent the distribution. This is the idea of reduced-set approximation. (cf. [Schölkopf et al. \(2002\)](#)) Mathematically, this method seeks to find an embedding with fewer expansion terms $\hat{\mu}_x^R = \sum_{i=1}^{N_R} \alpha^i \phi(x_i) \approx \hat{\mu}_x$, $N_R < N$, where the approximation is in the sense of RKHS distance measure. The reduced-set method forms the backbone of our approach. We also note that there are other related approximation methods such as those of [Chen et al. \(2012\)](#); [Bach et al. \(2012\)](#). Recently, [Zhu et al. \(2019\)](#) considered recursive applications of reduced-set method to uncertainty in stochastic systems.

3. Method

3.1. Stochastic programming with reduced-set expansion of RKHS embeddings

We consider the sample-based formulation of the stochastic programming problem (2). Our main idea is to perform constraint removal systematically using the aforementioned RKHS embedding reduced-set methods. Typically, constraint removal discards low-probability scenarios to reduce conservativeness of the resulting solution. Given a set of realized scenarios $\xi := \{\xi_1, \dots, \xi_n\}$ and

Algorithm 1 RKHS approximation to stochastic programming

- 1: Solve the sample-based stochastic programming problem (2).
- 2: Find the reduced-set RKHS embeddings

$$\hat{\mu}_x^R = \sum_{i \in \mathcal{R}} \alpha_i \phi(\xi_i).$$

by solving the convex optimization problem (3). \mathcal{R} is the reduced index set defined in (5).

- 3: Solve the stochastic programming problem according to the reduced set RKHS approximation (i.e., constraint removal by sparse optimization).

$$\underset{x}{\text{minimize}} L(x), \quad \text{subject to } C(x, \xi_j) \leq 0, \quad j \in \mathcal{R}. \quad (6)$$

- 4: **Output:** Solution of the above reduced stochastic program.
-

positive definite kernel k , we formulate optimization problem as

$$\underset{\alpha}{\text{minimize}} \|w^\top \alpha\|_1 \quad \text{subject to} \quad \left\| \sum_{i=1}^N \alpha_i \phi(\xi_i) - \hat{\mu}_\xi \right\|_{\mathcal{H}} \leq \epsilon, \quad (3)$$

where $w \in \mathbb{R}^N$ denotes scaling vector for the $l1$ -penalty. This can often be set to reflect specific concerns, such as the distance of states to the constraint. The KME expansion weights $\alpha := (\alpha_1, \dots, \alpha_N) \in \mathbb{R}^N$ need not sum to one. $\hat{\mu}_\xi$ denotes the empirical KME estimator of the distribution $\hat{\mu}_\xi := \frac{1}{N} \sum_{i=1}^N \phi(\xi_i)$. We further write down the equivalent Lagrangian form.

$$\underset{\alpha}{\text{minimize}} \quad \left\| \sum_{i=1}^N \alpha_i \phi(\xi_i) - \hat{\mu}_\xi \right\|_{\mathcal{H}}^2 + \lambda \|w^\top \alpha\|_1. \quad (4)$$

The resulting solution α^* is sparse due to the sparsity-inducing $l1$ term. We then discard the points $\xi_i, i \in \mathcal{I}$ with the index set $\mathcal{I} = \{i \mid \alpha_i = 0, i = 1, \dots, n\}$. Finally, we re-solve the stochastic programming problem with the reduced-set scenarios

$$\mathcal{R} := \{1, \dots, n\} \setminus \mathcal{I}. \quad (5)$$

The *intuition* of the optimization formulation (3) and (4) is to produce a subset of data whose distribution is close to the empirical data in the sense of RKHS-embedding distance $\|\hat{\mu}_x^R - \hat{\mu}_x\|_{\mathcal{H}}$. Meanwhile, the weighted $l1$ -penalty incentivizes the solution to become sparse. Therefore, the solution to (4) discards the ‘‘corner’’ cases while maintaining the statistical information. We outline the algorithmic procedure in Algorithm 1. The following lemma shows formulation (4) is computationally tractable.

If \mathcal{H} in problem (4) is the RKHS associated with a positive definite kernel, the objective of optimization problem is convex. To see this, we use the sample based estimator for the RKHS distance to rewrite the objective as

$$\underset{\alpha}{\text{minimize}} \quad \alpha^\top K \alpha + \alpha^\top K \beta + \lambda \cdot \|w^\top \alpha\|_1. \quad (7)$$

$\beta = (\frac{1}{N}, \dots, \frac{1}{N})$ is a constant vector. $K := \{k(x_i, x_j)\}_{i,j}$ is the gram matrix associated with the positive definite kernel, which implies $K \geq 0$. As $\|\cdot\|_1$ is convex, so is the optimization problem.

Remark 1 We can equivalently write the constraint of program (3) in the form of maximum mean discrepancy, $\|P^R - \hat{P}_{\text{data}}\|_{\text{MMD}} \leq \epsilon$. Optimization problems with constraints on the probability distributions are often referred to as moment problems.

3.2. Application to stochastic optimal control

Let us consider the following sample-based (scenario) formulation of stochastic optimal control problem (OCP).

$$\begin{aligned} & \underset{u}{\text{minimize}} && \frac{1}{N} \sum_{i=1}^N \int_0^T L(x^i(t), u(t)) dt \\ & \text{subject to} && \dot{x}^i = f(x^i, u, \xi^i), \\ & && h(x^i, u, \xi^i) \leq 0, \\ & && x^i(0) = \xi_0^i, i = 1, 2, \dots, N. \end{aligned} \tag{8}$$

where ξ, ξ_0 are uncertain variables and ξ^i, ξ_0^i their realizations. The uncertainty in the initial state ξ_0 is particularly relevant to MPC designs. After proper transcription and discretization, this OCP subsequently becomes the same form as the sample-based SP (2), solvable by Algorithm 1.

4. Numerical Experiments

4.1. Min-max robust regression

We first consider a synthetic stochastic programming problem given in the form of the following min-max robust regression.

$$\underset{x}{\text{minimize}} S, \quad \text{subject to } |A_i x - b_i| \leq S, \forall i. \tag{9}$$

For simplicity, we consider scalars A_i and b_i generated randomly according to the distributions. $n_1, n_2 \sim \text{N}(0, 1)$, $A_i = 3 + 3n_1$, $b_i = A_i x^* + 5n_2$. where x^* is the (unknown) true parameter drawn from $\text{Uniform}([2, 3])$.

Given the computed solution to the full program (9) \hat{x} , let us consider the quantity of interest $\hat{\xi}_i := A_i \hat{x} - b_i$, which is an RV due to the uncertainty in A_i and b_i . We now apply Algorithm 1 to find the reduced-set embedding of $\{\hat{\xi}_i\}_i^{N_R}$, $\hat{\mu}_{\hat{\xi}} = \sum_{i=1}^{N_R} \alpha_i^R \phi(\hat{\xi}_i)$. In step 3, in solving program (4), we used the scaling factor $w_i \propto \exp(T \hat{\xi}_i)$ to incentivize the removal of ‘‘corner’’ points (T may be thought of as the ‘‘softness’’ parameter of this softmax scaling factor). We then remove the constraints with identified index set \mathcal{I} , $|A_i x - b_i| \leq S, \forall i \in \mathcal{I}$, from the stochastic program and re-compute a solution. Following our discussion in the previous sections, this embedding captures the distribution information while discarding the rare scenarios. This is done by solving the sparse optimization problem (3). The results are illustrated in Figure 2. As we can see, scenarios associated with ‘‘corner’’ data points are not selected, causing the reduction in conservativeness. Figure 3 illustrates the effect of constraint removal on the optimal objective value and constraint violation. See the caption for detailed description.




Figure 2: Solutions of the min-max robust regression problem. Three figures correspond to three different regularization coefficients λ and number of discarded scenarios κ : (left) $\lambda = 0, \kappa = 0$. (center) $\lambda = 0.01, \kappa = 57$. (right) $\lambda = 0.05, \kappa = 144$. The shaded strip denotes the robust margin. Red points are the selected points by Step 2 in Algorithm 1. Dark points correspond to discarded scenarios and constraints. We used Gaussian kernel of bandwidth $\frac{1}{\sqrt{2}}$ to calculate the RKHS embedding in Algorithm 1.




Figure 3: Estimates produced by a large-sample ($N = 1000$) Monte Carlo simulation. (left). Constraint violation probability of the solution produced by Algorithm 1. This is estimated by the Monte Carlo estimation $\frac{1}{N} \sum_{i=1}^{N_R} \mathbb{I}(|A_i x - b_i| > S)$ where I is the indicator function of random events. (right). The new expected cost associated with the solution produced by Algorithm 1. This is estimated by the Monte Carlo estimation $\frac{1}{N} \sum_{i=1}^{N_R} F(x^*, \xi_i)$ where x^* is the solution by the proposed method.

4.2. Stochastic control

We now consider the Van der Pol oscillator model

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -0.1(1 - x_1^2)x_2 - x_1 + u \end{bmatrix}. \quad (10)$$

The goal of the control design is to steer the system state x_1 to a certain level. This is formulated as the following OCP.

$$\begin{aligned}
 & \underset{x(\cdot), u(\cdot)}{\text{minimize}} && \int_0^T \|x_1(t) - 3\|_2^2 dt \\
 & \text{subject to} && \dot{x}(t) = f(x(t), u(t)) \quad \forall t \in [0, T] \\
 & && -40 \leq u(t) \leq 40 \quad \forall t \in [0, T] \\
 & && -0.25 \leq x_1(t) \leq 2 + 0.1 \cos(10t) \quad \forall t \in [0, T] \\
 & && x(0) = s
 \end{aligned} \tag{11}$$

We sample the i.i.d. uncertainty realizations $\{s_1, \dots, s_n\} \sim N(m, \Sigma)$, where $m = [0.5 \ 0]^T$, $\Sigma = \begin{bmatrix} 0.01^2 & 0 \\ 0 & 0.1^2 \end{bmatrix}$. Because of the nonlinear dynamics, we cannot propagate the uncertainty in a tractable manner as in LQG without resorting to approximations. We use the sampled scenarios to form the OCP (8). The continuous-time dynamics is transcribed using multiple-shooting with CVodes (interfaced with CasADi) integrator. We then solve the discretized OCP with IPOPT to obtain the optimal control. An example of states associated with the solution is given in Figure 1 (left). The total time horizon is 1.0s and we consider 10 control steps in this experiment.

Let us consider the quantity of interest $\xi_i(t) := 2 + 0.1 \cos(10t) - x_1^i(t)$, the distance from state position to the upper bound constraint. This quantity reflects how close we are to be infeasible. It is random due to the states being a function of RV s . In Step 2 of Algorithm 1, the scaling factor is taken to be $w_i \propto \exp\left(\frac{C}{\epsilon + \min_t(\hat{\xi}_i(t))}\right)$ to encourage the removal of close-to-constraint trajectories.

We are now ready to apply Algorithm 1 to find the reduced-set embedding of $\{\hat{\xi}_i(t)\}_i^{N_R}, t = 1, \dots, T$, $\hat{\mu}_\xi = \sum_{i=1}^{N_R} \alpha_i \phi(\xi_i)$, where ξ_i is a vector comprising $\xi_i(t)$ at all time steps. Finally, we re-solve the subsequent reduced-set SP—OCP. Figure 1 (right) illustrates the reduced number of scenarios.

After we applied Algorithm 1, we obtain the “optimistic” controller u^* . To evaluate the performance of this controller, we use a large-sample Monte Carlo simulation to estimate the constraint violation probability, i.e., $\Pr\{C(x, \xi) \leq 0\}$ in the chance-constrained SP (1), as well as the expected cost over the large-sample simulation $\frac{1}{N} \sum_{i=1}^N \int_0^T L(x^i(t), u^*(t)) dt$. We plot the state trajectories with different number of removed constraints in Figure 4. The trade off between those is illustrated in Figure 5. The result makes intuitive sense that the more constraints we remove, the less the conservativeness, but with higher violation probability.

5. Discussion

This paper proposed a distributional-distance-regularized optimization formulation for stochastic programming under the framework of sampling-and-discarding. We demonstrated effective conservativeness reduction in data-driven optimization and control tasks. Although we did not study the guaranteed bounds, all analysis in Campi and Garatti (2011) applies to our case. Our on-going work is to apply our approach to distributionally robust optimization and control design.

Acknowledgments

We thank Joris Gillis for his help in the numerical experiment, Wittawat Jitkritum for his helpful feedback. This project has received funding from the European Union’s Horizon 2020 research

and innovation programme under the Marie Skłodowska-Curie grant agreement No 798321, the German Federal Ministry for Economic Affairs and Energy (BMWi) via eco4wind (0324125B) and DyConPV (0324166B), and by DFG via Research Unit FOR 2401.




Figure 4: (left) The optimistic controller evaluated in independent Monte Carlo simulations with $N = 100$ trajectories. The controller is produced with Algorithm 1 with $\kappa = 14$ removed scenarios ($l1$ regularization coefficient 10^{-5}). The estimated constraint-violation probability is 2%. (right) Controller produced with $\kappa = 83$ removed scenarios ($l1$ regularization 5×10^{-3}). The constraint violation probability is 5%.




Figure 5: Performance estimates produced by independent ($N = 100$) Monte Carlo simulations. (left). Constraint violation probability of the solution produced by Algorithm 1. This is estimated by the Monte Carlo estimation $\frac{1}{N} \sum_{i=1}^{N_R} \mathbb{I}(C(x, u) > 0)$ where I is the indicator function. $C(x, u) \leq 0$ denotes all the constraints in OCP (11). (right). The new expected cost of OCP associated with the solution produced by Algorithm 1. This is estimated by the Monte Carlo estimation $\frac{1}{N} \sum_{i=1}^N \int_0^T L(x^i(t), u^*(t)) dt$ where u^* is the optimistic controller.

References

- Francis Bach, Simon Lacoste-Julien, and Guillaume Obozinski. On the equivalence between herding and conditional gradient algorithms. *arXiv preprint arXiv:1203.4523*, 2012.
- John R Birge and Francois Louveaux. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.
- Giuseppe C Calafiore and Marco C Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.
- Marco C Campi and Algo Caré. Random convex programs with l_1 -regularization: sparsity and generalization. *SIAM Journal on Control and Optimization*, 51(5):3532–3557, 2013.
- Marco C Campi and Simone Garatti. A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality. *Journal of Optimization Theory and Applications*, 148(2):257–280, 2011.
- Yutian Chen, Max Welling, and Alex Smola. Super-samples from kernel herding. *arXiv preprint arXiv:1203.3472*, 2012.
- Darinka Dentcheva, András Prékopa, and Andrzej Ruszczyński. Concavity and efficient points of discrete distributions in probabilistic programming. *Mathematical Programming*, 89(1):55–77, 2000.
- James Luedtke, Shabbir Ahmed, and George L Nemhauser. An integer programming approach for linear programs with probabilistic constraints. *Mathematical programming*, 122(2):247–272, 2010.
- Bernhard Schölkopf, Alexander J Smola, Francis Bach, et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.
- Alexander J. Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A Hilbert space embedding for distributions. In *Proceedings of the 18th International Conference on Algorithmic Learning Theory (ALT)*, pages 13–31. Springer-Verlag, 2007.
- Jia-Jie Zhu, Krikamol Muandet, Moritz Diehl, and Bernhard Schölkopf. A new distribution-free concept for representing, comparing, and propagating uncertainty in dynamical systems with kernel probabilistic programming. *arXiv preprint arXiv:1911.11082*, 2019.