

Data-Driven Chance Constrained Control using Kernel Distribution Embeddings

Adam J. Thorpe¹ *

AJTHOR@UNM.EDU

Thomas Lew² *

THOMAS.LEW@STANFORD.EDU

Meeko M. K. Oishi¹

OISHI@UNM.EDU

Marco Pavone²

PAVONE@STANFORD.EDU

¹*Department of Electrical and Computer Engineering, University of New Mexico*

²*Department of Aeronautics and Astronautics, Stanford University*

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We present a data-driven algorithm for efficiently computing stochastic control policies for general joint chance constrained optimal control problems. Our approach leverages the theory of kernel distribution embeddings, which allows representing expectation operators as inner products in a reproducing kernel Hilbert space. This framework enables approximately reformulating the original problem using a dataset of observed trajectories from the system without imposing prior assumptions on the parameterization of the system dynamics or the structure of the uncertainty. By optimizing over a finite subset of stochastic open-loop control trajectories, we relax the original problem to a linear program over the control parameters that can be efficiently solved using standard convex optimization techniques. We demonstrate our proposed approach in simulation on a system with nonlinear non-Markovian dynamics navigating in a cluttered environment.

Keywords: kernel distribution embeddings, stochastic optimal control, joint chance constraints

1. Introduction

The deployment of reliable autonomous systems requires control algorithms that are robust to model misspecifications and to external disturbances. To enable safety-critical applications, these control algorithms should also explicitly enforce constraints. For instance, an autonomous car should always respect speed limits and avoid pedestrians at all times while accounting for uncertain road conditions and external disturbances. The presence of these two sources of aleatoric and epistemic uncertainty presents a significant challenge for traditional stochastic optimal control techniques, which typically rely upon an accurate model of the system and calibrated uncertainty quantification. Critically, the accuracy of the model and of its associated uncertainty estimates may deteriorate over time as the system is deployed in new environments. This motivates the use of data-driven techniques that leverage collected measurements of the system to design efficient adaptive control laws. However, existing data-driven techniques tend to be complex to implement (e.g., neural network controllers), may require system-specific assumptions restricting possible applications (e.g., dynamics that are linear in the uncertain parameters), or may not explicitly account for constraints that are crucial for the reliable deployment of these intelligent autonomous systems.

* These authors contributed equally to this work.




Figure 1: We propose a data-driven algorithm to efficiently compute control trajectories u for uncertain dynamical systems. The key consists of representing the state trajectory distribution $Q(\cdot|x_0, u)$ by a conditional distribution embedding $m(x_0, u)$ from a reproducing kernel Hilbert space $\mathcal{H}$. This approach allows approximating the true embedding $m(x_0, u)$ using a dataset $\mathcal{D}$ of observed trajectories of the uncertain system (see Assumption 1). The resulting embedding estimate $\hat{m}(x_0, u)$ defines an approximation $\hat{Q}(\cdot|x_0, u)$ of $Q(\cdot|x_0, u)$ that is used to efficiently compute a stochastic control policy to drive the system to a goal region $\mathcal{X}_{\text{goal}}$ while avoiding all unsafe sets $\mathcal{O}$ (e.g., representing obstacles) at all times.

Contributions: we present a data-driven control algorithm to efficiently compute stochastic control inputs for general joint chance constrained control problems given observed transitions from the system. The key consists of leveraging the theory of *kernel distribution embeddings*, which allows representing expectation operators as inner products in a reproducing kernel Hilbert space (RKHS). By applying this theory to the stochastic kernel that characterizes the uncertain system dynamics (Figure 1), we derive a tractable relaxation of the original joint chance constrained problem that is a linear program over the control parameters. Our approach allows computing randomized open-loop control strategies for non-Markovian nonlinear dynamical systems subject to nonconvex constraints.

Outline: we discuss related work in Section 2 and our joint chance constrained problem formulation in Section 3. Section 4 describes how stochastic kernel embeddings are used to relax the original intractable problem using a dataset of trajectories as a linear program over the stochastic open-loop control parameters. We validate our approach in Section 5 and conclude in Section 6.

2. Related work

A wide range of model-based stochastic control techniques have been developed to efficiently control uncertain systems. By leveraging a model of the system, these approaches explicitly account for uncertainty while enforcing constraints. For instance, stochastic model predictive controllers (Mesbah, 2016) explicitly enforce constraints along the state trajectory, dynamic programming approaches (Ono et al., 2015) use a Lagrangian relaxation and augment the control objective with a penalty on constraint violation, and control barrier functions (Clark, 2019) provide a condition that guarantees closed-loop forward invariance. Guaranteeing constraint satisfaction with high probability typically involves considering a chance constrained problem formulation. The most common formulation enforces *pointwise* chance constraints that ensure the independent satisfaction of each constraint at each time step with high probability (Castillo-Lopez et al., 2019; Lew et al., 2020; Hewing et al., 2020; Polymenakos et al., 2020; Khojasteh et al., 2020; Jasour et al., 2021). In contrast, joint chance constraints guarantee trajectory-wise constraints satisfaction with high probability (Blackmore et al., 2011; Frey et al., 2020; Schmerling and Pavone, 2017; Koller et al., 2018; Lew et al., 2022) which is of particular interest whenever all constraints should be satisfied at all times jointly. For instance, a drone transporting a package should always avoid obstacles and reach its destination with high probability over the distribution of possible payloads. Unfortunately, tackling

joint chance constrained formulations is particularly challenging due to the need to consider the full distribution of the state trajectory. Common approaches decompose such constraints using Boole’s inequality (Blackmore et al., 2011; Ono, 2016) which can be conservative (Schmerling and Pavone, 2017), or perform robust trajectory optimization given confidence sets for the model parameters (Koller et al., 2018; Lew et al., 2022) which requires assuming bounded external disturbances and relying on robust uncertainty propagation techniques that can be conservative.

The performance of model-based controllers depends on prior domain knowledge, which may become inaccurate over time as the system is deployed in new environments. For this reason, data-driven control techniques have been proposed to leverage measurements of the system to efficiently characterize its properties and adapt to changes in its operating conditions. Popular data-driven approaches leverage Gaussian processes models (Deisenroth et al., 2015; Ostafew et al., 2016; Khojasteh et al., 2020; Williams and Rasmussen, 2006; Berkenkamp et al., 2017; Koller et al., 2018; Lew et al., 2022) which make Gaussian-distributed predictions for the state transitions of the system, linearly-parameterized dynamics (Coulson et al., 2019; Berberich et al., 2021), Koopman operators (Abraham and Murphey, 2019), random Fourier features (Boffi, 2021), and neural networks (Chua et al., 2018). A key algorithmic feature of these approaches is to directly model single-step transitions of the system. While intuitive and sufficient from a statistical viewpoint, this approach makes uncertainty propagation challenging, requiring particle-based approaches to represent the distribution of the trajectory (Blackmore et al., 2010; Janson et al., 2015; Chua et al., 2018).

In this work, we take a drastically different approach and leverage the theory of *distributional kernel embeddings* to characterize the joint distribution of the state trajectory from data. This approach allows tackling general nonlinear, non-Markovian dynamical systems, representing joint chance constraints as a linear operation in an RKHS, and optimizing over open-loop stochastic control policies by solving a linear program over the control parameters. Kernel distribution embeddings have been thoroughly studied in the recent years (Song et al., 2009; Smola et al., 2007; Grünewälder et al., 2012; Park and Muandet, 2020), but are not yet popular within the control community. Controller synthesis applications have been explored in (Thorpe and Oishi, 2021), but the authors do not consider constraints and only consider one-step transition kernels. In contrast, we tackle a joint chance constrained problem formulation that explicitly accounts for constraints and represent the distribution of the entire state trajectory. As a byproduct, this approach allows handling non-Markovian dynamical systems that are resistant to traditional one-step modeling techniques.

3. Problem Formulation

First, we describe our notation. We denote a conjunction (logical AND) by $\wedge$ and a disjunction (logical OR) by $\vee$. $\mathbb{R}$ and $\mathbb{N}$ denote the sets of real and natural numbers, respectively. Given a space E and $N \in \mathbb{N}$, we denote the Cartesian product $E^N \triangleq E \times \dots \times E$ (N times). For a subset $A \subset E$, the map $\mathbf{1}_A : E \rightarrow \{0, 1\}$ denotes the indicator function of A , which satisfies $\mathbf{1}_A(x) = 1$ if $x \in A$ and $\mathbf{1}_A(x) = 0$ if $x \notin A$. We denote the Borel σ -algebra on a topological space E by $\mathcal{B}(E)$.

Definition 3.1 (Stochastic Kernel) *Let $(E, \mathcal{E})$ and $(F, \mathcal{F})$ be measurable spaces. A stochastic kernel from E to F is a map $\kappa : \mathcal{F} \times E \rightarrow [0, 1]$ such that $x \mapsto \kappa(A \mid x)$ is $\mathcal{E}$ -measurable for all $A \in \mathcal{F}$ and $B \mapsto \kappa(B \mid x)$ is a probability measure on $(F, \mathcal{F})$ for all $x \in E$.*

3.1. Chance Constrained Optimization

Let $(\Omega, \mathcal{G}, \mathbb{P})$ be a probability space and $f : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}^n$ be a continuous function. Given an initial state $x_0 \in \mathbb{R}^n$, our goal is to safely control a system with stochastic dynamics

$$x_{t+1} = f(x_t, u_t, w_t, \theta), \quad t \in \mathbb{N}, \quad (1)$$

where $x_t \in \mathcal{X} \subseteq \mathbb{R}^n$ denotes the state of the system at time t , $u_t \in \mathcal{U} \subset \mathbb{R}^m$ is the control input, $w = (w_t)_{t \in \mathbb{N}}$ is a stochastic process characterizing external disturbances, and θ is a random variable that describes uncertain model parameters (e.g., a drone transporting an uncertain payload of mass θ that is sampled at the beginning of the control task and held constant over time, see Section 5).

This formulation is particularly challenging due to the fact that the system described in (1) is non-Markovian. Specifically, at any time $t \in \mathbb{N}$, the stochastic state trajectory $x = (x_t)_{t \in \mathbb{N}}$ satisfies

$$x_t = f_t(x_0, u, w, \theta) \triangleq f(\cdot, u_{t-1}, w_{t-1}, \theta) \circ \cdots \circ f(x_0, u_0, w_0, \theta), \quad (2)$$

where $u = (u_s)_{s=0}^{t-1}$ denotes the control trajectory. Since the parameters θ are uncertain and the disturbances w_t are not necessarily independent, the increments $x_{t+1} - x_t$ are not independent, i.e. the state trajectory x is not a Markov process. Intuitively, the parameters are randomized only once and their uncertainty is propagated along the entire state trajectory, see (Lew et al., 2021).

We consider the problem of minimizing the sum of two (possibly non-convex) state and input cost functions $\ell^x : \mathcal{X}^N \rightarrow \mathbb{R}$ and $\ell^u : \mathcal{U}^N \rightarrow \mathbb{R}$. For instance, ℓ^x may be chosen as a quadratic cost that penalizes tracking error and ℓ^u may penalize control effort. In addition to this control objective, the system should reach a desired compact goal region $\mathcal{X}_{\text{goal}} \subset \mathcal{X}$ at a specified time $N \in \mathbb{N}$ while avoiding all unsafe sets of states $\mathcal{O}_t \subset \mathcal{X}$ (e.g., representing potentially nonconvex obstacles). Due to the stochasticity of the system, guaranteeing strict constraint satisfaction with probability one may be infeasible. Instead, given a tolerable failure probability threshold $\delta \in (0, 1)$, we require that all constraints are jointly satisfied with probability at least $1 - \delta$, a requirement often referred to as a joint chance constraint.

In general, the optimal solution of such chance constrained control problems are *stochastic* control policies (Altman, 1999; Ono, 2016). As such, we formulate a stochastic optimal control problem over the set Δ of all open-loop stochastic policies $\pi : \mathcal{B}(\mathcal{U}^N) \times \mathcal{X} \rightarrow [0, 1]$, where π is a stochastic kernel from $\mathcal{X}$ to $\mathcal{U}^N$. Optimizing over closed-loop stochastic policies is challenging; we leave such extensions for future work. The resulting control problem is expressed as follows:

$$\min_{\pi \in \Delta} \mathbb{E}[\ell^x(x) + \ell^u(u)] \quad (3a)$$

$$\text{s.t. } x_{t+1} = f(x_t, u_t, w_t, \theta), \quad u \sim \pi(\cdot \mid x_0), \quad t = 0, 1, \dots, N-1, \quad (3b)$$

$$\mathbb{P}\left(\left(\bigwedge_{t=1}^{N-1} x_t \notin \mathcal{O}_t\right) \wedge (x_N \in \mathcal{X}_{\text{goal}})\right) \geq 1 - \delta, \quad (3c)$$

where the state and control trajectories are denoted as $x = (x_t)_{t=1}^N$ and $u = (u_t)_{t=0}^{N-1}$. This problem is challenging due to the non-convexity of the cost function (3a), the nonlinearity of the uncertain dynamics (3b), the joint chance constraint (3c), and the non-Markovianity of the state-trajectory $(x_t)_{t=1}^N$ which requires reasoning about the joint probability distribution of the state trajectory.

Without imposing assumptions regarding the system parameterization and the distribution of the disturbances and parameters, the problem in (3) is intractable. Instead of relying on such assumptions, in this work, we assume access to a dataset of observed transitions from the system. Such a dataset may come from prior observations of the system evolution or from high-fidelity simulations.

Assumption 1 (Dataset) We have access to a dataset $\mathcal{D} = \{(x_0^i, u^i, x^i)\}_{i=1}^M$ of $M \in \mathbb{N}$ independent and identically distributed (i.i.d.) trajectories, where x_0^i and u^i are sampled i.i.d. from probability distributions on $\mathcal{X}$ and $\mathcal{U}^N$ and $x^i = (x_1^i, \dots, x_N^i)$ are i.i.d. trajectories satisfying (1) for each i .

In the next section, we propose an equivalent reformulation of the problem in (3) using stochastic kernels. This reformulation is the basis for our data-driven approach that leverages the dataset $\mathcal{D}$.

3.2. Reformulation using Stochastic Kernels

The state trajectory x that satisfies (1) can be characterized by a stochastic kernel $Q : \mathcal{B}(\mathcal{X}^N) \times \mathcal{X} \times \mathcal{U}^N \rightarrow [0, 1]$ that assigns a probability measure $Q(\cdot \mid x_0, u)$ to every initial condition $x_0 \in \mathcal{X}$ and control sequence $u \in \mathcal{U}^N$ on the measurable space $(\mathcal{X}^N, \mathcal{B}(\mathcal{X}^N))$. This kernel can be defined as $Q(A \mid x_0, u) \triangleq \mathbb{P}(\bigwedge_{t=1}^N x_t = f_t(x_0, u, w, \theta) \in A_t)$ for any $A = A_1 \times \dots \times A_N \in \mathcal{B}(\mathcal{X}^N)$ and $(x_0, u) \in \mathcal{X} \times \mathcal{U}^N$. As such, Assumption 1 states that one has access to M independent trajectory samples x^i that are distributed according to $Q(\cdot \mid x_0^i, u^i)$.

For any initial state $x_0 \in \mathcal{X}$, any control trajectory $u \in \mathcal{U}^N$, and any measurable function $g : \mathcal{X}^N \rightarrow \mathbb{R}$, we denote the expectation with respect to this probability measure as $\mathbb{E}_{Q(x \mid x_0, u)}[g(x)] = \int_{\mathcal{X}^N} g(x) Q(dx \mid x_0, u)$. Since $\mathbb{E}_{Q(x \mid x_0, u)}[g(x)]$ is a function of the initial condition x_0 and control input $u \in \mathcal{U}^N$, reformulating the original problem in (3) amounts to computing the expectation over the stochastic control policy. Since the distribution of the control inputs u is characterized by the stochastic kernel π , for any measurable function $h : \mathcal{U}^N \rightarrow \mathbb{R}$, we define the conditional expectation $\mathbb{E}_{\pi(u \mid x_0)}[h(u)] = \int_{\mathcal{U}^N} h(u) \pi(du \mid x_0)$.

The two stochastic kernels Q and π allow representing the joint chance constraint in (3c) as the expectation of an indicator function. Let $\mathcal{T} = (\mathcal{X} \setminus \mathcal{O}_1) \times \dots \times (\mathcal{X} \setminus \mathcal{O}_{N-1}) \times \mathcal{X}_{\text{goal}} \subset \mathcal{X}^N$ denote the set of all state trajectories which reach the goal set while avoiding the unsafe sets at all times, and let $\mathbf{1}_{\mathcal{T}}(x)$ denote the indicator of $\mathcal{T}$, which is one if $x_N \in \mathcal{X}_{\text{goal}}$ and $x_t \notin \mathcal{O}_t$ for all $t = 1, \dots, N-1$, and is zero otherwise. Then, the joint chance constraint in (3c) can be reformulated as

$$\mathbb{E}_{\pi(u \mid x_0)} [\mathbb{E}_{Q(x \mid x_0, u)} [\mathbf{1}_{\mathcal{T}}(x)]] = \mathbb{P} \left(\left(\bigwedge_{t=1}^{N-1} x_t \notin \mathcal{O}_t \right) \wedge (x_N \in \mathcal{X}_{\text{goal}}) \right) \geq 1 - \delta. \quad (4)$$

Using (4) and the stochastic kernel Q , we reformulate the original problem in (3) as

$$\min_{\pi \in \Delta} \quad \mathbb{E}_{\pi(u \mid x_0)} [\mathbb{E}_{Q(x \mid x_0, u)} [\ell^x(x) + \ell^u(u)]] \quad (5a)$$

$$\text{s.t.} \quad \mathbb{E}_{\pi(u \mid x_0)} [\mathbb{E}_{Q(x \mid x_0, u)} [\mathbf{1}_{\mathcal{T}}(x)]] \geq 1 - \delta. \quad (5b)$$

Since the stochastic kernel Q that characterizes the state trajectory x is unknown, the problem above is generally intractable. In the next section, we show how the dataset $\mathcal{D}$ allows relaxing the problem in (5) using a framework known as *kernel embeddings of distributions*. This approach allows representing the expectation operators in (5) as elements in a high-dimensional function space and subsequently reformulating the original problem as a linear program in this space.

4. Hilbert Space Embeddings of Distributions

The main challenge in solving (5) is evaluating the expectations with respect to the stochastic kernels Q and π that characterize dynamics uncertainty and the stochasticity of the control policy. Our key insight consists of embedding an integral operator as an element in an RKHS. With this approach,

evaluating the expectations in (5) amounts to performing a linear operation between two functions of an RKHS. We use the dataset from Assumption 1 to approximate the expectation operator associated to the unknown stochastic kernel Q . Then, by selecting a finite-dimensional representation for the stochastic policy, we propose a tractable finite-dimensional approximation of (5).

4.1. Reproducing Kernel Hilbert Space Embeddings

In this section, we endow our problem with an RKHS structure that enables representing the expectation operators in problem (5) as linear functions. First, we define the positive definite kernel function $k_{\mathcal{X}^N} : \mathcal{X}^N \times \mathcal{X}^N \rightarrow \mathbb{R}$ (Steinwart and Christmann, 2008, Definition 4.15). According to the Moore-Aronszajn theorem (Aronszajn, 1950), this kernel defines a unique corresponding RKHS $\mathcal{H}$ of functions from $\mathcal{X}^N$ to $\mathbb{R}$ with associated inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ such that: (i) $k_{\mathcal{X}^N}(x, \cdot) \in \mathcal{H}$ for all $x \in \mathcal{X}^N$, and (ii) $g(x) = \langle g, k_{\mathcal{X}^N}(x, \cdot) \rangle_{\mathcal{H}}$ for all $g \in \mathcal{H}$ and $x \in \mathcal{X}^N$. As a consequence of the reproducing property, $k_{\mathcal{X}^N}(x, x') = \langle k_{\mathcal{X}^N}(x, \cdot), k_{\mathcal{X}^N}(x', \cdot) \rangle_{\mathcal{H}}$, which is often referred to as the kernel trick (Steinwart and Christmann, 2008). We assume that the kernel $k_{\mathcal{X}^N}$ is $\mathcal{B}(\mathcal{X}^N)$ -measurable and bounded, such that $\sup_{x \in \mathcal{X}^N} \sqrt{k_{\mathcal{X}^N}(x, x)} < \infty$. With these conditions, according to (Song et al., 2009), for every $x_0 \in \mathcal{X}$ and $u \in \mathcal{U}^N$, there exists an element $m(x_0, u) \in \mathcal{H}$, called the *conditional distribution embedding*, which is a linear function from $\mathcal{X}^N$ to $\mathbb{R}$ defined as

$$m(x_0, u) \triangleq \mathbb{E}_{Q(x|x_0, u)}[k_{\mathcal{X}^N}(x, \cdot)]. \quad (6)$$

By the reproducing property and linearity of the expectation, the expectation of any function of the RKHS $\mathcal{H}$ with respect to the stochastic kernel Q can be evaluated as an inner product with the conditional distribution embedding. Specifically, for any $x_0 \in \mathcal{X}$ and any $u \in \mathcal{U}^N$,

$$\mathbb{E}_{Q(x|x_0, u)}[g(x)] = \langle g, m(x_0, u) \rangle_{\mathcal{H}} \quad \text{for any } g \in \mathcal{H}. \quad (7)$$

Assuming that the cost function ℓ^x and the indicator function $1_{\mathcal{T}}$ belong to $\mathcal{H}$, the inner expectations in (5) with respect to Q can be rewritten as a linear operation using the distribution embedding $m(x_0, u)$. As such, our approach does not rely on specific assumptions about the dynamics, the uncertain parameters, and the disturbances. The main challenge consists of computing an accurate approximation of the distribution embedding $m(x_0, u)$, which we pursue in the next section.

4.2. Empirical Embedding Estimate

Since the true stochastic kernel Q is unknown a priori, we do not have access to the conditional distribution embedding $m(x_0, u) \in \mathcal{H}$. To construct an empirical estimate $\hat{m}(x_0, u)$ of this operator, we leverage the dataset $\mathcal{D}$ from Assumption 1. Specifically, as in (Grünwälder et al., 2012; Caponetto and De Vito, 2007; Micchelli and Pontil, 2005), we search for a best-fit solution in the RKHS as the solution to the regularized least-squares problem $\hat{m} = \arg \min_{f \in \mathcal{H}} \frac{1}{M} \sum_{i=1}^M \|k_{\mathcal{X}^N}(x^i, \cdot) - f(x_0^i, u^i)\|_{\mathcal{H}}^2 + \lambda \|f\|_{\mathcal{H}}^2$, where $\lambda > 0$ is a regularization parameter and $\mathcal{Q}$ is a vector-valued RKHS of functions from $\mathcal{X} \times \mathcal{U}^N$ to $\mathcal{H}$ (see Micchelli and Pontil, 2005, for more information). The solution to this problem is unique:

$$\hat{m}(x_0, u) = \Phi^\top (G + \lambda MI)^{-1} K(x_0, u), \quad (8)$$

where $\Phi \in \mathcal{H}^M$ is a feature vector with elements $\Phi_i = k_{\mathcal{X}^N}(x^i, \cdot) \in \mathcal{H}$, $G \in \mathbb{R}^{M \times M}$ is a Gram matrix with elements $G_{ij} = k_{\mathcal{X}}(x_0^i, x_0^j) k_{\mathcal{U}^N}(u^i, u^j)$, where $k_{\mathcal{X}}$ and $k_{\mathcal{U}^N}$ are positive definite




Figure 2: We search for a stochastic kernel embedding $p(x_0)$ in the RKHS $\mathcal{U}$. This defines a stochastic kernel $\pi(u|x_0)$ that characterizes the control policy. Using this policy kernel and the approximate state trajectory stochastic kernel $\hat{Q}(x|x_0, u)$, we represent the state trajectory distribution and reformulate the original stochastic control problem.

kernels on $\mathcal{X}$ and $\mathcal{U}^N$, respectively, and $K(x_0, u) \in \mathbb{R}^M$ is a vector with elements $[K(x_0, u)]_i = k_{\mathcal{X}}(x_0^i, x_0) k_{\mathcal{U}^N}(u^i, u)$. This estimate $\hat{m}(x_0, u)$ of the conditional distribution embedding $m(x_0, u)$ enables an efficient approximation of the expectation in (7). Specifically, for any function $g \in \mathcal{H}$,

$$\mathbb{E}_{Q(x|x_0, u)}[g(x)] = \langle g, m(x_0, u) \rangle_{\mathcal{H}} \approx \langle g, \hat{m}(x_0, u) \rangle_{\mathcal{H}} = \mathbf{g}^\top (G + \lambda MI)^{-1} K(x_0, u), \quad (9)$$

where $\mathbf{g} = [g(x^1), \dots, g(x^M)]^\top \in \mathbb{R}^M$. With (9), the inner expectations in (5) with respect to the unknown stochastic kernel Q can be approximated using simple matrix operations.

The quality of this approximation generally depends on the dataset $\mathcal{D}$ from Assumption 1, on the kernel choice, and on the number of samples M . The convergence properties of conditional distribution embeddings have been studied in the literature, see (Song et al., 2009, 2010; Grünewälder et al., 2012; Park and Muandet, 2020). Notwithstanding minor variations in the underlying theoretical framework, existing results show that the empirical estimate of a kernel distribution embedding converges in probability to the true embedding as the sample size increases, and suggest an optimal rate of convergence of $\mathcal{O}_p(M^{-1/4})$ (Park and Muandet, 2020). Exactly quantifying the approximation error in (9) for a finite number of samples is beyond the scope of this work.

4.3. Stochastic Control Trajectory Representation

To reformulate the problem in (5), it remains to compute the expectation with respect to any chosen stochastic policy π . At this point, one could search for a deterministic control sequence $u \in \mathcal{U}^N$ and obtain a tractable approximate reformulation of the original problem. However, this reformulation would be non-convex if ℓ^u and $k_{\mathcal{U}^N}$ are non-convex.

Instead, optimizing over the larger set of stochastic policies allows for more efficient control strategies (Altman, 1999; Ono, 2016). Note that under the condition that $k_{\mathcal{U}^N}$ is $\mathcal{B}(\mathcal{U}^N)$ -measurable and bounded, any stochastic policy $\pi(\cdot | x_0)$ can be represented by a conditional distribution embedding $p(x_0)$ in the RKHS $\mathcal{U}$ associated with $k_{\mathcal{U}^N}$. Accordingly, the expectation of any function $h \in \mathcal{U}$ with respect to the stochastic policy can be expressed as an inner product in the RKHS $\mathcal{U}$ as $\mathbb{E}_{\pi(u|x_0)}[h(x_0, u)] = \langle h, p(x_0) \rangle_{\mathcal{U}}$. However, depending on the choice of kernel $k_{\mathcal{U}^N}$, the representation of the embedding $p(x_0)$ may be infinite-dimensional. Thus, optimizing over all possible stochastic embeddings is generally intractable. This observation motivates our finite-dimensional policy representation, which we characterize by the set of stochastic kernel embeddings

$$p(x_0) = \sum_{j=1}^P \gamma_j(x_0) k_{\mathcal{U}^N}(\tilde{u}^j, \cdot), \quad \text{where} \quad \sum_{j=1}^P \gamma_j(x_0) = 1, \quad 0 \preceq \gamma(x_0), \quad (10)$$

where $\gamma(x_0) \in \mathbb{R}^P$, the elements $\tilde{u}^j \in \mathcal{U}^N$ are user-specified control sequences, and $j = 1, \dots, P$. The finite set of control sequences $\mathcal{A} = \{\tilde{u}^j\}_{j=1}^P$ can be chosen strategically to uniformly cover the control space $\mathcal{U}^N$, can be sampled randomly (independently of Assumption 1), or can be pre-specified depending on the application (e.g., informed by a sampling-based planner with an approximate deterministic dynamics model, [Ichter et al., 2018](#)). The coefficients $\gamma_j(x_0)$ characterize the probability values that weight the admissible control sequences in $\mathcal{A}$. Correspondingly, the last two conditions in (10) constrain the coefficient vector $\gamma(x_0)$ to lie within a probability simplex $\mathcal{S} = \{\gamma \in \mathbb{R}^P \mid \sum_{i=1}^P \gamma_i = 1, 0 \preceq \gamma\}$. This guarantees that the kernel distribution embedding $p(x_0)$ defines a valid stochastic kernel π characterizing the control policy.

Restricting the search of stochastic kernel embeddings to a finite-dimensional subspace of $\mathcal{U}$ allows a tractable relaxation of the original problem. Indeed, for any function $h \in \mathcal{U}$, we have

$$\mathbb{E}_{\pi(u|x_0)}[h(u)] = \langle h, p(x_0) \rangle_{\mathcal{U}} = \left\langle h, \sum_{j=1}^P \gamma_j(x_0) k_{\mathcal{U}^N}(\tilde{u}^j, \cdot) \right\rangle_{\mathcal{U}} = \sum_{j=1}^P h(\tilde{u}^j) \gamma_j(x_0). \quad (11)$$

Assuming that $\ell^u \in \mathcal{U}$, this relaxation enables a tractable evaluation of the expectation operator with respect to the stochastic control policy $\pi(\cdot \mid x_0)$. Intuitively, $\pi(\cdot \mid x_0)$ is a sum of Dirac functions centered at each control trajectory $\tilde{u}^j \in \mathcal{U}^N$ weighted by the probability coefficient $\gamma_j(x_0)$ ([Berlinet and Thomas-Agnan, 2011](#), Chapter 4, §1.1). Importantly, (11) is linear in the unknown parameter $\gamma(x_0) \in \mathbb{R}^P$. This relaxation offers a tractable linear reformulation of the original problem, although it induces sub-optimality relative to the computationally intractable problem of optimizing over the infinite-dimensional space of stochastic policies. We leave the quantification of this sub-optimality gap (perhaps by first quantifying the error between the empirical approximation in (10) and any feasible stochastic embedding $p(x_0)$, see ([Kanagawa et al., 2018](#), §6.2)) to future work.

4.4. Approximate chance constrained Optimization Problem

Using the estimate $\hat{m}(x_0, u)$ and the policy estimate $p(x_0)$, the expectations in (5) can be approximated using simple matrix multiplications and inner products. Specifically, for any function $g \in \mathcal{H}$, combining (9) with (11), we obtain $\mathbb{E}_{\pi(u|x_0)}[\mathbb{E}_{Q(x|x_0,u)}[g(x)]] \approx \langle \langle g, \hat{m}(x_0, \cdot) \rangle_{\mathcal{H}}, \hat{p}(x_0) \rangle_{\mathcal{U}} = \sum_{j=1}^P \mathbf{g}^\top (G + \lambda MI)^{-1} K(x_0, \tilde{u}^j) \gamma_j(x_0)$. This motivates the following approximation of (5):

$$\min_{\gamma(x_0) \in \mathbb{R}^P} \quad \left(\ell^x{}^\top (G + \lambda MI)^{-1} R(x_0) + \ell^u{}^\top \right) \gamma(x_0) \quad (12a)$$

$$\text{s.t.} \quad \left(\bar{\mathbf{1}}_{\mathcal{T}}{}^\top (G + \lambda MI)^{-1} R(x_0) \right) \gamma(x_0) \geq 1 - \delta, \quad \mathbf{1}^\top \gamma(x_0) = 1, \quad 0 \preceq \gamma(x_0), \quad (12b)$$

where $R(x_0) \in \mathbb{R}^{M \times P}$ is a matrix such that each column $[R(x_0)]_j = K(x_0, \tilde{u}^j)$ such that $[R(x_0)]_{ij} = k_{\mathcal{X}}(x_0^i, x_0) k_{\mathcal{U}^N}(u^i, \tilde{u}^j)$, and $\ell^x, \ell^u, \bar{\mathbf{1}}_{\mathcal{T}} \in \mathbb{R}^M$ are vectors with elements $\ell_i^x = \ell^x(x^i)$, $\ell_i^u = \ell^u(\tilde{u}^j)$, and $[\bar{\mathbf{1}}_{\mathcal{T}}]_i = \mathbf{1}_{\mathcal{T}}(x^i)$, respectively. The approximate problem in (12) is a linear program over the coefficients $\gamma(x_0)$ (that represent probability weights over the admissible control sequences in $\mathcal{A}$) that can be efficiently solved via interior-point or simplex algorithms. Thus, the solution to (12) is a randomized control strategy, i.e., an open-loop stochastic control policy with values in $\mathcal{A}$. The quality of the resulting stochastic control strategy that solves the problem in (12) generally depends on the quality of the dataset $\mathcal{D}$ from Assumption 1 and on the choice of the control inputs in $\mathcal{A}$.




Figure 3: Solutions from our proposed data-driven algorithm for different values of the tolerable failure probability threshold δ . Blue trajectories from Monte-Carlo (MC) simulations denote feasible trajectories which reach the goal set $\mathcal{X}_{\text{goal}}$ and avoid obstacles $\mathcal{O}$. Red trajectories violate constraints. We use $\text{MC} = y\%$ to denote the percent of sampled trajectories that violate constraints.

5. Numerical Experiments

We consider the problem of controlling a quadrotor system operating in turbulent conditions, with nonlinear dynamics given by $x_{t+1} = Ax_t + B(\theta)u_t + d(x_t, \theta) + w_t$, where

$$A = \begin{bmatrix} 1 & \Delta t & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta t \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B(\theta) = \frac{1}{m} \begin{bmatrix} \Delta t^2/2 & 0 \\ \Delta t & 0 \\ 0 & \Delta t^2/2 \\ 0 & \Delta t \end{bmatrix}, \quad d(x_t, \theta) = -\alpha \begin{bmatrix} \Delta t^2 |v_x| v_x / 2 \\ \Delta t |v_x| v_x \\ \Delta t^2 |v_y| v_y / 2 \\ \Delta t |v_y| v_y \end{bmatrix}, \quad (13)$$

and $\Delta t = 1$. The state and control inputs of the system are denoted as $x_t = [p_{x,t}, v_{x,t}, p_{y,t}, v_{y,t}] \in \mathbb{R}^4$ and $u_t = [u_x, u_y] \in \mathbb{R}^2$. Starting from an initial condition $x_0 = [-0.5, 0, -0.5, 0]$, the goal consists of computing a stochastic policy such that the system reaches the set $\mathcal{X}_{\text{goal}} = \{x \in \mathbb{R}^4 \mid \|(p_x - 10, p_y - 10)\| \leq 2.5 \text{ at time } N = 15\}$ while avoiding two polytopic obstacles $\mathcal{O}$ shown in Figure 3. The disturbances w_t are i.i.d. zero-mean Gaussian-distributed, with diagonal covariance matrix $\Sigma = \text{diag}([10^{-2}, 0, 10^{-2}, 0])$. The dynamics are parameterized by uncertain parameters $\theta = (m, \alpha)$ representing the mass $m > 0$ of the system and an uncertain drag coefficient $\alpha > 0$. We assume that the parameters (m, α) of the system are uncorrelated random variables such that $(m - 0.75)/(0.5) \sim \text{Beta}(2, 2)$ and $(\alpha - 0.4)/(0.2) \sim \text{Beta}(2, 5)$, where $\text{Beta}(a, b)$ denotes a Beta distribution with shape parameters (a, b) . As discussed in Section 3, this system is not Markovian due to the temporal correlation between the state trajectory x and the uncertain parameter θ .

We generate a dataset $\mathcal{D} = \{(x_0^i, u^i, x^i)\}_{i=1}^M$ as in Assumption 1 with $M = 2500$ by uniformly sampling initial conditions $x_0^i \in [-0.5, -0.05, -0.5, -0.05] \times [0.5, 0.05, 0.5, 0.05]$. The control sequences u^i are generated by uniformly sampling $u_t^i \in [0, 1]^2$ for $t = 0, 1, 2$ to impose variability in the sample trajectories, and then selecting controls for $t \geq 3$ via a feedback controller, where $u_t^i = K(x_t - (10, 0, 10, 0))$ given closed-loop trajectories x simulated using (13) and a pre-specified linear feedback gain $K = \begin{bmatrix} 1 & 2.25 & 0 & 0 \\ 0 & 0 & 1 & 2.25 \end{bmatrix} \in \mathbb{R}^{2 \times 4}$. These control trajectories u^i are then used to simulate open-loop trajectories x^i again, which ensures that these trajectories correspond to open-loop control inputs. Finally, the control sequences $\mathcal{A} = \{\tilde{u}^j\}_{j=1}^P$ with $P = 2500$ are selected with the same approach as described above, except that we use a uniform grid to select the control inputs $\tilde{u}_t^j$ at the first three time steps, $t = 0, 1, 2$, and use the dynamics in (13) with constant parameters $(m, \alpha) = (1, 0.005)$. We then presumed no knowledge of the system dynamics in (13) or the stochastic kernel Q to compute a solution to the problem.

Using the datasets $\mathcal{D}$ and $\mathcal{A}$, we compute a kernel-based approximation of the original chance constrained problem and stochastic policy using (12). We set the regularization parameter to be $\lambda = 10^{-7}$ and use a Gaussian kernel $k(x, x') = \exp(-\sigma \|x - x'\|_2^2)$ with the kernel parameter σ chosen as the median Euclidean distance between sample points. We use the `linprog` function from

`scipy` to formulate and solve the linear programs in (12). Our Python implementation, as well as code to reproduce all figures and analysis is available at <https://github.com/ajthor/socks>.

Results are shown in Figure 3 for different values of δ . Figure 3 shows 1,000 Monte-Carlo simulations using the open-loop policy computed via our method. For all values of δ , the algorithm computes a randomized control policy that satisfies the desired success rate $1 - \delta$. The results in Figure 3 show that with larger values of δ , the algorithm is more likely to select control inputs that generate trajectories passing through the riskier middle corridor between the obstacles. This is expected behavior, as δ directly controls the maximum allowable probability of constraint violation.

We compare our approach with the method proposed in (Ono, 2016), which consists of decomposing the joint chance constraint using Boole’s inequality and combining a Lagrangian relaxation approach with a bisection technique to obtain a mixed control strategy. Since this approach considers linear dynamics, we implement this baseline assuming that $m = 1$ and $\alpha = 0$ using the dynamics in (13). We present results for $\delta = 0.05$ in Figure 4 and generate 1,000 Monte-Carlo simulations using the resulting mixed control strategy and the nonlinear dynamics in (13). We observe that the solution from this approach is similar to the output of our data-driven method, in that it mixes two control strategies that generate trajectories from two different control sequences. However, 10.2% of trajectories violate constraints using the method in (Ono, 2016). This comparison demonstrates the need to consider the nonlinear, non-Markovian dynamics of this system to reliably solve this problem.




Figure 4: Solution from the approach proposed in (Ono, 2016). Blue Monte-Carlo trajectories denote feasible trajectories that reach the goal set $\mathcal{X}_{goal}$ and avoid obstacles $\mathcal{O}$. Red trajectories indicate trajectories which violate constraints.

6. Conclusion

In this paper, we presented a novel data-driven technique leveraging stochastic kernel embeddings to compute stochastic control trajectories for general joint chance constrained control problems. By directly approximating the integral operator associated to the distribution of the state trajectory, this model-free approach is agnostic to the particular parameterization of the dynamics and of the distribution of the uncertain parameters and external disturbances. It is also computationally efficient, since the original stochastic problem is approximately reformulated as a linear program over the control parameters. Numerical experiments demonstrate the efficacy of our approach for controlling a challenging nonlinear non-Markovian dynamical system subject to non-convex constraints.

An important direction of future work is the accurate characterization of the approximation error in (9). Specifically, given further distributional and smoothness assumptions, finite-sample error bounds between the true and empirical embeddings in (9) could allow the strict satisfaction of the original joint chance constraint (3c). Future work will investigate different control input selection strategies, e.g., using non-convex optimization or cross-entropy policy optimization approaches. Of practical interest is investigating different choices of kernels to yield better accuracy and less conservative behavior (e.g., by meta-learning a kernel tailored to the true system). Finally, our data-driven approach could motivate reinforcement learning strategies to actively gather data and improve the stochastic kernel embedding approximation quality while balancing the exploration-exploitation trade-off that arises from this constrained Markov decision process formulation.

Acknowledgments

The authors thank Masahiro Ono for insightful discussions about using mixed strategies for joint chance-constrained control and Apoorva Sharma for his helpful feedback. This material is based upon work supported by the National Science Foundation under NSF Grant Numbers CNS-1836900 and CMMI-271742. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The NASA University Leadership initiative (Grant #80NSSC20M0163) provided funds to assist the authors with their research, but this article solely reflects the opinions and conclusions of its authors and not any NASA entity. This research was supported in part by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA-0003525. The views expressed in this article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

- O. Abraham and T. D. Murphey. Active learning of dynamics for data-driven control using koopman operators. *IEEE Transactions on Robotics*, 35(5):1071–1083, 2019.
- E. Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- N. Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950.
- J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer. Linear tracking MPC for nonlinear systems, 2021. Available at <https://arxiv.org/abs/2105.08560> and at <https://arxiv.org/abs/2105.08567>.
- F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause. Safe model-based reinforcement learning with stability guarantees. In *Conf. on Neural Information Processing Systems*, 2017.
- A. Berlinet and C. Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Transactions on Robotics*, 26(3):502–517, 2010.
- L. Blackmore, M. Ono, and B. C. Williams. Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics*, 27(6):1080–1094, 2011.
- N. M. Boffi. Nonparametric adaptive control and prediction: theory and randomized algorithms. In *Proc. IEEE Conf. on Decision and Control*, 2021.
- A. Caponnetto and E. De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.

- M. Castillo-Lopez, S. A. Ludivig, P. Sajadi-Alamdari, J. L. Sanchez-Lopez, M. A. Olivares-Mendez, and H. Voos. A real-time approach for chance-constrained motion planning with dynamic obstacles. *IEEE Robotics and Automation Letters*, 5(2):3620 – 3625, 2019.
- K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Conf. on Neural Information Processing Systems*, 2018.
- A. Clark. Control barrier functions for complete and incomplete information stochastic systems. In *American Control Conference*, 2019.
- J. Coulson, J. Lygeros, and F. Dörfler. Data-enabled predictive control: In the shallows of the DeePC. In *European Control Conference*, 2019.
- M. Deisenroth, D. Fox, and C. Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 37(2):408–423, 2015.
- K. M. Frey, T. J. Steiner, and J. P. How. Collision probabilities for continuous-time systems without sampling. In *Robotics: Science and Systems*, 2020.
- S. Grünewälder, G. Lever, L. Baldassarre, S. Patterson, A. Gretton, and M. Pontil. Conditional mean embeddings as regressors. In *Int. Conf. on Machine Learning*, 2012.
- L. Hewing, J. Kabzan, and M. N. Zeilinger. Cautious model predictive control using Gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6), 2020.
- B. Ichter, J. Harrison, and M. Pavone. Learning sampling distributions for robot motion planning. In *Proc. IEEE Conf. on Robotics and Automation*, 2018.
- L. Janson, E. Schmerling, and M. Pavone. Monte Carlo motion planning for robot trajectory optimization under uncertainty. In *Int. Symp. on Robotics Research*, 2015.
- A. Jasour, W. Han, and B. Brian Williams. Convex risk bounded continuous-time trajectory planning in uncertain nonconvex environments. In *Robotics: Science and Systems*, 2021.
- M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*, 2018.
- M. J. Khojasteh, V. Dhiman, M. Franceschetti, and N. Atanasov. Probabilistic safety constraints for learned high relative degree system dynamics. In *2nd Annual Conference on Learning for Dynamics & Control*, 2020.
- T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause. Learning-based model predictive control for safe exploration. In *Proc. IEEE Conf. on Decision and Control*, 2018.
- T. Lew, R. Bonalli, and M. Pavone. Chance-constrained sequential convex programming for robust trajectory optimization. In *European Control Conference*, 2020.
- T. Lew, A. Sharma, J. Harrison, E. Schmerling, and M. Pavone. On the problem of reformulating systems with uncertain dynamics as a stochastic differential equation, 2021. Available at <https://arxiv.org/abs/2111.06084>.

- T. Lew, A. Sharma, J. Harrison, A. Bylard, and M. Pavone. Safe active dynamics learning and control: A sequential exploration-exploitation framework. *IEEE Transactions on Robotics*, 2022. In Press.
- A. Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016.
- C. A Micchelli and M. Pontil. On learning vector-valued functions. *Neural computation*, 17(1): 177–204, 2005.
- M. Ono. Control by coin flips: Mixed strategy for finite-horizon stochastic optimal control. In *American Control Conference*, 2016.
- M. Ono, M. Pavone, Y. Kuwata, and J. Balaram. Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39(4):555–571, 2015.
- C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot. Robust Constrained Learning-based NMPC enabling reliable mobile robot path tracking. *Int. Journal of Robotics Research*, 35(13):1547–1563, 2016.
- J. Park and K. Muandet. A measure-theoretic approach to kernel conditional mean embeddings. In *Conf. on Neural Information Processing Systems*, 2020.
- K. Polymenakos, L. Laurenti, A. Patane, J. P. Calliess, L. Cardelli, M. Kwiatkowska, A. Abate, and S. Roberts. Safety guarantees for planning based on iterative Gaussian processes. In *Proc. IEEE Conf. on Decision and Control*, 2020.
- E. Schmerling and M. Pavone. Evaluating trajectory collision probability through adaptive importance sampling for safe motion planning. In *Robotics: Science and Systems*, 2017.
- A. Smola, A. Gretton, L. Song, and B. Schölkopf. A Hilbert space embedding for distributions. In *Proc. Computational Learning Theory*, 2007.
- L. Song, J. Huang, A. Smola, and K. Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Int. Conf. on Machine Learning*, 2009.
- S. Song, A. Gretton, and C. Guestrin. Nonparametric tree graphical models. In *Proc. of the Thirteenth Int. Conf. on Artificial Intelligence and Statistics*, 2010.
- I. Steinwart and A. Christmann. *Support Vector Machines*. Springer Publishing Company, Incorporated, 2008.
- A. J. Thorpe and M. M. K. Oishi. Stochastic optimal control via Hilbert space embeddings of distributions. In *Proc. IEEE Conf. on Decision and Control*, 2021.
- C. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. MIT press, 2006.