

Kernel-Based Learning of Safety Barriers

OLIVER SCHÖN*, ETH Zürich, Switzerland

ZHENGANG ZHONG, University of Warwick, United Kingdom

SADEGH SOUDJANI, Max Planck Institute for Software Systems, Germany and University of Birmingham, United Kingdom

The rapid integration of AI algorithms in safety-critical applications such as autonomous driving and healthcare is raising significant concerns about the ability to meet stringent safety standards. Traditional tools for formal safety verification struggle with the black-box nature of AI-driven systems and lack the flexibility needed to scale to the complexity of real-world applications. In this paper, we present a data-driven approach for safety verification and synthesis of black-box systems with discrete-time stochastic dynamics. We employ the concept of control barrier certificates, which can guarantee safety of the system, and learn the certificate directly from a set of system trajectories. We use conditional mean embeddings to embed data from the system into a reproducing kernel Hilbert space (RKHS) and construct an RKHS ambiguity set that can be inflated to robustify the result to out-of-distribution behavior. We provide the theoretical results on how to apply the approach to general classes of temporal logic specifications beyond safety. For the data-driven computation of safety barriers, we leverage a finite Fourier expansion to cast a typically intractable semi-infinite optimization problem as a linear program. The resulting spectral barrier allows us to leverage the fast Fourier transform to generate the relaxed problem efficiently, offering a scalable yet distributionally robust framework for verifying safety. Our work moves beyond restrictive assumptions on system dynamics and uncertainty, as demonstrated on two case studies including a black-box system with a neural network controller.

JAIR Track: Integration of Logical Constraints in Deep Learning

JAIR Associate Editor: Matteo Zavatteri

JAIR Reference Format:

Oliver Schön, Zhengang Zhong, and Sadegh Soudjani. 2026. Kernel-Based Learning of Safety Barriers. *Journal of Artificial Intelligence Research* 86, Article 25 (July 2026), 36 pages. DOI: [10.1613/jair.1.21679](https://doi.org/10.1613/jair.1.21679)

1 Introduction

We are living in a truly remarkable time of technological progress powered by AI. Autonomous cars, for instance, have recently been reported to have surpassed human level safety whilst driving under controlled conditions — an achievement supported largely by extensive and costly testing (Kalra and Paddock 2016). Yet, achieving a level of safety sufficient to earn societal trust in AI-powered safety-critical systems presents a significant challenge. With a rapid adoption of embodied AI across various domains, including smart grids and advanced medical technologies, we are facing an unprecedented demand for safe and trustworthy AI that testing cannot meet. In fact, in a world governed by non-determinism, exhaustively testing embodied AI against all possible edge cases is

*Corresponding Author.

Authors' Contact Information: Oliver Schön, ORCID: [0000-0002-0214-6455](https://orcid.org/0000-0002-0214-6455), oschoen@ethz.ch, ETH Zürich, Zürich, Switzerland; Zhengang Zhong, ORCID: [0009-0002-8262-5217](https://orcid.org/0009-0002-8262-5217), zhengang.zhong@warwick.ac.uk, University of Warwick, Coventry, United Kingdom; Sadegh Soudjani, ORCID: [0000-0003-1922-6678](https://orcid.org/0000-0003-1922-6678), sadegh@mpi-sws.org, Max Planck Institute for Software Systems, Kaiserslautern, Germany and University of Birmingham, Birmingham, United Kingdom.



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

© 2026 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.21679](https://doi.org/10.1613/jair.1.21679)

infeasible. To address these challenges, the ability to enforce logical constraints — such as safety and invariance — on AI systems is emerging as a critical capability.

Formal methods for the verification and synthesis of dynamical systems are being developed to bring forth tools for certifying whether an AI system can be trusted (Belta et al. 2017; Lavaei et al. 2022; Yin et al. 2024). More precisely, rich (stochastic) models of the target system’s behavior are subjected to temporal logic constraints to obtain rigorous probabilistic measures of their well-behavedness. As exact system representations are rarely available, models are often derived from data. To compensate for the arising epistemic uncertainty, formal methods are equipped to deal with model ambiguity. The formal nature of the results, however, comes at the cost of onerous assumptions and a substantial computational burden. In fact, formal reasoning amidst unknown dynamics is only possible if the data is supplemented with information on the regularity of the system dynamics. However, this often restricts the applicability of results to specific classes of systems, such as linear, polynomial, or control-affine systems. Extending these methods to encompass more general AI-driven systems, such as those governed by neural network (NN) controllers, could act as a catalyzer for the development of safe AI systems.

Existing work on formal methods is dominated by a bipartite narrative of approaches that are either fundamentally *abstraction-based* or *abstraction-free*. Whilst the former is attributed to suffer from the excruciating *curse of dimensionality*, the latter is expected to be more scalable. At a closer examination, however, this division becomes less clear. As the arguably most popular abstraction-free formal tool, *control barrier certificates*¹ (CBCs) allow one to reformulate the verification of safety by proving the existence of a function — a so-called barrier — satisfying a given set of mathematical inequalities. The resulting global optimization problem involving the latent system dynamics is in general a robust optimization problem featuring infinitely many constraints. Thus, many approaches rely on some form of spatial abstraction, i.e., partitioning of the system domain into discrete states and inputs, to solve the complex optimization problem (cf. Section 6.2). To address these shortcomings, we draw from a line of research that has thus far received only little attention from the formal methods community: (reproducing) kernel methods.

Recent advancements in machine learning (ML) have underscored the significance of representing models as functions within *reproducing kernel Hilbert spaces* (RKHS) (Berlinet and Thomas-Agnan 2004; Schölkopf and A. J. Smola 2002; Steinwart and Christmann 2008). An RKHS is a Hilbert space of functions where evaluation at any point is a continuous linear operation, making the RKHS formulation well-suited for analyzing and manipulating functions in high-dimensional spaces. The RKHS framework has sparked a deeper understanding of complex ML algorithms and led to the development of tools such as *conditional mean embeddings* (CMEs) (Klebanov et al. 2020; Song et al. 2009). CMEs provide a way to embed conditional probability measures into an RKHS, enabling the computation of conditional expectations of random variables as simple inner products. This approach is particularly powerful for analyzing the expected behavior of stochastic processes, as it captures dependencies and relationships within data without requiring explicit density estimation.

In control applications, kernel-based tools such as CMEs offer several benefits over traditional data-driven optimization techniques. They allow the reformulation of nonlinear problems as linear ones within the RKHS framework, thus simplifying computations (Berlinet and Thomas-Agnan 2004; Schölkopf and A. J. Smola 2002). Additionally, they bypass intermediate steps such as density estimation and numerical integration, facilitating the direct computation of conditional expectations on the observed data. Unlike many existing formal methods for verification and synthesis, which often rely on structural assumptions or Lipschitz continuity, CMEs can enable the approximate reformulation of the original problem using finite data based only on a standard assumption used in CME theory. Paired with statistical results based on concentration inequalities, probabilistic guarantees for the latent data-generating system can be obtained from small and noisy data sets (Z. Li et al. 2022). This versatility and computational efficiency make CMEs particularly appealing for formal verification and synthesis.

¹We refer to CBCs and discrete-time control barrier functions (CBFs) (Cosner et al. 2024) jointly as CBCs since their definitions are equivalent.

1.1 Contributions

In this work, we study utilizing the CME theory for verification and control synthesis of stochastic systems without explicit model knowledge via the concept of control barrier certificates (CBCs) (Prajna, Jadbabaie, et al. 2007). By reformulating the probabilistic CBC constraints into a data-driven optimization problem, we derive distributionally robust characterizations based on an ambiguity set of candidate transition kernels. We show how these characterizations can be extended to satisfy temporal logic specifications beyond safety using an automata representation of the specification and the concept of Streitt supermartingales (Abate, Giacobbe, et al. 2024). From a computational perspective, our central innovation lies in the application of a Fourier expansion to the barrier function, yielding a computationally efficient and expressive *Fourier control barrier certificate*. We show that for safety verification, this approach collapses the complex CME term into a tractable spectral representation, enabling constraint evaluation at equidistant sample lattices via the fast Fourier transform (FFT). The resulting spectral formulation reduces the semi-infinite program to a finitely constrained *linear program* (LP), leveraging bounding results for trigonometric polynomials (Pfister and Bresler 2018). Our method relaxes restrictive assumptions of existing approaches and scales favorably across benchmarks.

A subset of the results of this paper was published in the conference paper by Schön, Zhong, et al. (2024). This manuscript provides substantial extensions of previous results along the following directions.

- (a) We generalize the theoretical framework to handle non-autonomous dynamics, yielding data-driven CBC conditions based on a general standard assumption used in CME theory. This removes the need for restrictive assumptions prevalent in related work (cf. related work in Section 1.3).
- (b) In order to address temporal logic specifications beyond safety, we provide robust inequalities based on CME theory, an automata representation of the specification, and Streitt supermartingales, by raising appropriate assumptions on the system.
- (c) With focus on safety verification, we provide a detailed study of robust optimization and semi-infinite programming techniques applicable to the derived program. These include monolithic and divide-and-conquer strategies, which we juxtapose based on their theoretical complexity and our experimental findings.
- (d) For the general squared-exponential kernel and safety verification, we develop a truncated Fourier expansion that eliminates the reliance on large amounts of spatial support vectors, yielding a computationally efficient method. A scalable sampling-based scheme, leveraging the FFT, enables practical LP reformulations for complex stochastic systems.
- (e) The newly provided benchmarks include complex safety specifications and neural network controllers.

1.2 Paper Organization

The rest of the paper is organized as follows. After a brief review of related work, we present the preliminaries and problem statement in Section 2. The concepts of RKHS and kernel mean embeddings for (conditional) probability measures are introduced in Section 3. We derive inequality constraints for computing CBCs directly from data in Section 4. The application of kernel mean embeddings to general classes of temporal logic specifications beyond safety is discussed in Section 5. With a focus on safety specifications, we characterize the resulting semi-infinite problem and provide a thorough comparison of possible solutions in Section 6. In Section 7, we present a spectral approach based on a spectral Fourier expansion that admits a solution of the data-driven verification problem as an LP. We demonstrate the performance of the proposed approach on two benchmarks followed up by concluding remarks in Sections 8 and 9, respectively.

1.3 Related Work

There exists a substantial body of work on data-driven formal approaches for safety verification and control synthesis of stochastic systems. We will focus on “abstraction-free” approaches via CBCs/CBFs, and refer to the

complimentary publications by Banse et al. (2023), Chekan and Langbort (2023), Gracia et al. (2023), Kazemi, Majumdar, et al. (2024), Makdesi et al. (2024), Nazeri et al. (2025), Schön, Huijgevoort, et al. (2024), Schön, Naseer, et al. (2024), and Zhang et al. (2024) for a recent selection of abstraction-based work addressing uncertain systems.

For data-driven CBCs/CBFs, many existing approaches are limited to linear or control-affine dynamics (see, e.g., Cohen et al. 2022; Jagtap, Pappas, et al. 2020; Lopez and Slotine 2023). Furthermore, many approaches rely on known Lipschitz constants to provide formal guarantees and/or address only partially unknown dynamics. For instance, Gaussian processes are employed by L. Wang et al. (2018) and by Jagtap, Pappas, et al. (2020) to learn partially unknown dynamics of nonlinear systems whilst assuming the affine control-dependent part of the dynamics to be known. Systems with unknown additive disturbance are addressed by Chekan and Langbort (2023). Similarly, Mathiesen et al. (2024) assume the deterministic part of the dynamics to be accurately known, whilst the CBC-based safety control method proposed by Mazouz et al. (2024) focus on known noise distributions instead. Approaches to fully unknown dynamics are scarce. Salamati et al. (2024) study the computation of CBCs for fully unknown discrete-time systems relying on known Lipschitz constants. Uncertain continuous-time systems are studied by C. Wang et al. (2023), employing Bayesian inference and local Lipschitzness. Importantly, Lipschitz bounds directly constrain the variation of the unknown dynamics, whereas our RKHS-based assumption instead restricts the complexity of the function in a kernel-induced hypothesis space, which may be significantly more flexible depending on the choice of kernel.

In the pursuit of a more flexible method, NN-based approaches for synthesizing so-called neural barriers have gained popularity due to their functional expressiveness (Abate, Ahmed, et al. 2021; Abate, Bogomolov, et al. 2024; So et al. 2023). For instance, So et al. (2023) train neural CBFs based on finite sample sets, acknowledging that their verification can be performed using NN verification tools such as the bound propagation techniques leveraged by X. Wang et al. (2024). The field of NN verification is highly active, with numerous tools and methodologies being developed to ensure reliability and safety in AI systems (see, e.g., Henriksen and Lomuscio 2021; Katz et al. 2017; Kouvaros and Lomuscio 2021; Xu et al. 2020). Neural networks as compact representations have been utilized by Majumdar et al. (2023) for memory-efficient formal verification and synthesis. Noteworthy adjacent work by Kazemi and Soudjani (2020) and Kazemi, Perez, et al. (2024) studies temporal logic control via model-free reinforcement learning with convergence guarantees and assume guarantee contracts. Apart from our previous work (Schön, Zhong, et al. 2024), the only known publication on correct-by-design control via CMEs is due to Romao et al. (2023), which embeds the transition kernel for abstraction-based control. In contrast, we consider CMEs in an abstraction-free setting. In extension to our previous work, Chen, Y. Li, et al. (2025) propose a CBC-based approach for reach-avoid properties based on the Wasserstein distance.

2 Preliminaries and Problem Statement

Notation. We denote the sets of positive integers and non-negative reals as $\mathbb{N}_{>0}$ and $\mathbb{R}_{\geq 0}$, respectively. Consider a Polish sample space $\mathbb{X}$ (Bogachev 2007). Let $(\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{P})$ be the underlying probability space equipped with a Borel σ -algebra $\mathcal{B}(\mathbb{X})$ defined over $\mathbb{X}$, and a probability measure $\mathbb{P}$. For a random variable X , let p_X be the pushforward probability measure of $\mathbb{P}$ under X such that $X \sim p_X(\cdot)$. The expected value of a function $f(X)$ on $\mathbb{X}$ is written as $\mathbb{E}_{p_X}[f(X)]$. If it is clear from the context, we abbreviate and write $\mathbb{E}[f(X)]$. We denote the set of all probability measures for a given measurable space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$ as $\mathcal{P}(\mathbb{X})$. The n -dimensional Gaussian measure with mean $\mu \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$ is given by

$$\mathcal{N}(A \mid \mu, \Sigma) := \int_A \frac{d\lambda^n(x)}{\sqrt{(2\pi)^n |\Sigma|}} \exp \left[-\frac{1}{2} (x - \mu)^\top \Sigma^{-1} (x - \mu) \right],$$

where $\lambda^n: \mathcal{B}(\mathbb{X}) \rightarrow [0, \infty]$ denotes the usual n -dimensional Lebesgue measure and $|\Sigma|$ is the determinant of Σ . The Dirac delta measure $\delta_a: \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$ concentrated at a point $a \in \mathbb{X}$ is defined as $\delta_a(A) = 1$ if $a \in A$ and $\delta_a(A) = 0$ otherwise, for any measurable set $A \in \mathcal{B}(\mathbb{X})$. We denote the uniform distribution over $\mathbb{X}$ as

$\mathcal{U}_{\mathbb{X}}$ with realizations $x \sim \mathcal{U}_{\mathbb{X}}(\cdot)$. For two measurable spaces $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$ and $(\mathbb{Y}, \mathcal{B}(\mathbb{Y}))$, a *probability kernel* is a mapping $\mathbf{p}: \mathbb{X} \times \mathcal{B}(\mathbb{Y}) \rightarrow [0, 1]$ such that $\mathbf{p}(X = x, \cdot): \mathcal{B}(\mathbb{Y}) \rightarrow [0, 1]$ is a probability measure for all $x \in \mathbb{X}$, and $\mathbf{p}(\cdot, B): \mathbb{X} \rightarrow [0, 1]$ is measurable for all $B \in \mathcal{B}(\mathbb{Y})$. A probability kernel associates to each point $x \in \mathbb{X}$ a measure denoted by $\mathbf{p}(\cdot | X = x)$.

The transpose of a vector or matrix A is indicated as A^T . Let the $N \times N$ dimensional identity matrix be given by I_N . Let $X_N := [x_i]_{i=1}^N$ be a column vector with $x_i \in \mathbb{X}$. We denote the element-wise evaluation of a function $f: \mathbb{X} \rightarrow \mathbb{R}$ on X_N as $f(X_N) := [f(x_i)]_{i=1}^N$. Similarly, we may write $A = [a_{ij}]_{i,j=1}^N$ to denote a matrix with its elements. For a vector v , we denote by v^{-1} its element-wise inverse.

2.1 Discrete-Time Stochastic Systems

In this work, we consider systems expressible as Markov decision processes over continuous state and input spaces, formally defined as follows.

DEFINITION 1 (MARKOV DECISION PROCESS (MDP)). *An MDP is a tuple $\mathbf{M} = (\mathbb{X}, \mathbb{X}_0, \mathbb{U}, \mathbf{t})$, comprising a state space $\mathbb{X} \subset \mathbb{R}^n$ with states $x \in \mathbb{X}$; initial states $x_0 \in \mathbb{X}_0 \subset \mathbb{X}$; an input space $\mathbb{U}$ with inputs $u \in \mathbb{U}$; and a probability kernel $\mathbf{t}: \mathbb{X} \times \mathbb{U} \times \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$.*

In every execution, given a current state $x \in \mathbb{X}$ and control input $u \in \mathbb{U}$, the MDP evolves to a consecutive state $x_+ \in \mathbb{X}$, which is obtained as a realization $x_+ \sim \mathbf{t}(\cdot | X = x, U = u)$. As a notable class of systems that can be captured using MDPs, we may consider black-box systems with Markovian discrete-time stochastic dynamics

$$\mathbf{M}: \{ x_{t+1} = f(x_t, u_t, w_t), \quad w_t \sim p_w(\cdot), \quad (1)$$

where the system state and control input at the t^{th} time-step are denoted by x_t and u_t , respectively. The state evolution of the system, described by the function $f: \mathbb{X} \times \mathbb{U} \times \mathbb{W} \rightarrow \mathbb{X}$, is subject to independent, identically distributed (i.i.d.) noise $w_t \sim p_w(\cdot)$ supported on a set $\mathbb{W}$. The corresponding probability kernel is given by

$$\mathbf{t}(dx_+ | x_t, u_t) = \int_{\mathbb{W}} \delta_{f(x_t, u_t, w)}(dx_+) p_w(dw).$$

2.2 System Safety

In this paper, we aim to synthesize a control policy π such that the resulting closed-loop system $\pi \times \mathbf{M}$ is safe. We characterize *safety* associated with a time horizon $T \in \mathbb{N}_{>0}$ as a property of state trajectories: a trajectory is said to be safe if it does not enter a designated unsafe set $\mathbb{X}_u \subset \mathbb{X}$ for all $t \in [0, T]$. As stochastic systems do generally not admit a binary notion of safety, this induces a safety probability on the underlying closed-loop system, defined as the probability that the trajectory of $\pi \times \mathbf{M}$ starting from any initial state $x_0 \in \mathbb{X}_0$ remains in the safe set $\mathbb{X} \setminus \mathbb{X}_u$. In the following, we focus on a finite-horizon formulation. Thus, we quantify the *probability* of reaching $\mathbb{X}_u$ in a given finite horizon $T < \infty$ and compute its complement. A (finite-horizon) *safety specification* $\psi_{\text{safe}} := (\mathbb{X}_u, T)$ is hence fully characterized by an unsafe set $\mathbb{X}_u$ and a time horizon T . The system $\pi \times \mathbf{M}$ satisfies ψ_{safe} with probability at least $p \in (0, 1)$ if the probability of its trajectories starting in $\mathbb{X}_0$ and avoiding $\mathbb{X}_u$ within horizon T is at least p . This is denoted by $P_{\text{safe}}^{\pi}(\mathbf{M}) := \mathbb{P}(\pi \times \mathbf{M} \models \psi_{\text{safe}})$ (Baier and Katoen 2008), for which we require $P_{\text{safe}}^{\pi}(\mathbf{M}) \geq p$. To illustrate this, consider the following concrete example.

EXAMPLE 1. *Fig. 1 shows an autonomous vehicle (in blue) overtaking a leading vehicle. The dynamics of the ego vehicle are given by Dubin's car model with additive noise:*

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \\ \phi_{t+1} \end{bmatrix} = \begin{bmatrix} x_t \\ y_t \\ \phi_t \end{bmatrix} + \tau \begin{bmatrix} v \cos(\phi_t) \\ v \sin(\phi_t) \\ u_t \end{bmatrix} + \begin{bmatrix} w_t^1 \\ w_t^2 \\ w_t^3 \end{bmatrix}, \quad (2)$$




Fig. 1. Examples of safety-critical embodied AI systems in the transport sector.

with time discretization $\tau := 0.5$, where (x, y) denotes the relative position in longitudinal and lateral direction, and ϕ is the heading angle. The maneuver starts in $\mathbb{X}_0$ (blue region). Let the noise w_t^i , $i \in \{1, 2, 3\}$, be zero-mean Gaussian with standard deviation 0.01, 0.01, and 0.001, respectively. The steering angle $u_t = \pi(x_t)$ is supplied by a neural network (NN) controller $\pi: \mathbb{X} \rightarrow \mathbb{U}$, which is trained to steer clear of the sides of the road and overtake the leading vehicle without collision; that is, as illustrated in Fig. 1, avoiding the unsafe regions $\mathbb{X}_u$ and reaching a target set ahead of the leading vehicle. The goal of this paper is to determine (probabilistically) if the controlled system $\pi \times \mathbf{M}$ is safe.

We remark that for finite-horizon safety, maximizing $P_{\text{safe}}^\pi(\mathbf{M})$ with respect to the policy π would require policies that are Markov and time-variant (Puterman 1994). In contrast, the inequality $P_{\text{safe}}^\pi(\mathbf{M}) \geq p$ may be satisfiable by any type of policies. For simplicity, we consider stationary (i.e., time-invariant) policies of the form $\pi: \mathbb{X} \rightarrow \mathbb{U}$.

2.3 Control Barrier Certificates

Whilst certifying the safety of a continuous-space stochastic system is generally challenging, *control barrier certificates* (CBCs) and *control barrier functions* (CBFs) leverage the concept of set invariance to arrive at an abstraction-free formulation. This has made CBCs/CBFs popular tools for safety verification and synthesis (Prajna, Jadbabaie, et al. 2007). We briefly recall the definition of a CBC, based on the theory of Kushner (1967), dubbed *stochastic barrier functions* by Santoyo et al. (2021).

DEFINITION 2 (CONTROL BARRIER CERTIFICATE (CBC)). A function $B: \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is called a CBC of an MDP $\mathbf{M} = (\mathbb{X}, \mathbb{X}_0, \mathbb{U}, \mathbf{t})$ with reference to an unsafe set $\mathbb{X}_u$, if we have

- (a) $\forall x_0 \in \mathbb{X}_0: B(x_0) \leq \eta$;
- (b) $\forall x_u \in \mathbb{X}_u: B(x_u) \geq 1$; and
- (c) $\forall x \in \mathbb{X}, \exists u \in \mathbb{U}: \mathbb{E}_t[B(X^+) | X = x, U = u] - B(x) \leq c$;

for some constants $\eta \in [0, 1)$ and $c \geq 0$.

Intuitively, condition (c) of Definition 2 restricts the CBC B to elicit a relaxed *supermartingale* property for any non-zero constant c . Note that this implies the existence of a stationary policy $\pi: \mathbb{X} \rightarrow \mathbb{U}$ generating the corresponding control inputs $u \in \mathbb{U}$. If a CBC can be found for a system $\mathbf{M}$, then, a lower bound on the probability of $\mathbf{M}$ being safe is given by the following proposition due to Kushner (1967, Theorem 3).

PROPOSITION 1 (FINITE-HORIZON SAFETY). *Consider an MDP $\mathbf{M} = (\mathbb{X}, \mathbb{X}_0, \mathbb{U}, \mathbf{t})$ and a safety specification $\psi_{\text{safe}} = (\mathbb{X}_u, T)$. Suppose there exists a CBC B with respect to $\mathbb{X}_u$ (Definition 2) with constants η and c . Then, there exists a stationary policy π such that*

$$P_{\text{safe}}^{\pi}(\mathbf{M}) \geq 1 - (\eta + cT).$$

REMARK 1 (INFINITE-HORIZON SAFETY). *Finding a CBC (Definition 2) for a non-negative constant $c > 0$ is generally easier than for $c = 0$. If, however, a CBC can be obtained for $c = 0$, Proposition 1 provides safety guarantees for an unbounded time horizon $T \rightarrow \infty$. This recovers the results by Prajna, Jadbabaie, et al. (2007).*

2.4 Problem Statement

The aim is to synthesize a safety policy π for a black-box system $\mathbf{M}$, where the transition kernel $\mathbf{t}$ is unknown and only observed through a finite set of observations of the form $\{x^i, u^i, x_+^i\}_{i=1}^N$, $N \in \mathbb{N}_{>0}$, where $x_+^i \sim \mathbf{t}(\cdot | X = x^i, U = u^i)$ for uniformly drawn $x^i \sim \mathcal{U}_{\mathbb{X}}(\cdot)$ and $u^i \sim \mathcal{U}_{\mathbb{U}}(\cdot)$. This corresponds to drawing i.i.d. samples from the joint distribution

$$(X, U, X^+) \sim \mathbf{t}(X^+ | X, U) \mathcal{U}_{\mathbb{X}}(X) \mathcal{U}_{\mathbb{U}}(U). \quad (3)$$

Note that this setting admits a plethora of target systems. We formalize the problem statement as follows.

PROBLEM 1. *Consider a given safety specification ψ_{safe} and a confidence level $1 - \rho \in (0, 1)$. Without knowledge of the transition kernel $\mathbf{t}$ of $\mathbf{M}$ and based only on i.i.d. observations $\{x^i, u^i, x_+^i\}_{i=1}^N$, $N \in \mathbb{N}_{>0}$, find a policy $\pi: \mathbb{X} \rightarrow \mathbb{U}$ and threshold $p_N^{\pi} \in (0, 1)$ such that $P_{\text{safe}}^{\pi}(\mathbf{M}) \geq p_N^{\pi}$ with confidence at least $1 - \rho$.*

The confidence in Problem 1 is w.r.t. a random draw of N data samples used to obtain the safety bound p_N^{π} , i.e., we have $\mathbb{P}^N(P_{\text{safe}}^{\pi}(\mathbf{M}) \geq p_N^{\pi}) \geq 1 - \rho$, where $\mathbb{P}^N$ denotes the joint probability distribution over N data samples.

In this work, we address Problem 1 under the assumption of minimal complexity information (see Assumption 1) and use CBCs to certify probabilistic safety. The main challenge in establishing CBCs for unknown stochastic systems arises from the stochastic constraint in condition (c) (Definition 2). We show how CBCs can be generated from data by embedding the conditional probability measure via kernel methods, which cast the stochastic constraint in terms of an inner product. We construct an RKHS ambiguity set that is centered at the empirical mean embedding and can be inflated to robustify against out-of-distribution dynamics. As a result, the problem of synthesizing a robust policy is cast as a semi-infinite program. For the special case where a policy π is given, i.e., system verification, we present an efficient algorithmic solution. Based on a finite Fourier expansion of the kernel, we establish a relaxation of the semi-infinite program as a scalable linear program.

In the following section, we give a brief introduction to RKHS theory and provide the results for embedding (conditional) probability measures into RKHSs.

3 Kernel Mean Embeddings

RKHS basics. A symmetric function $k_{\mathbb{X}}: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ is called a (positive definite) *kernel* (note the distinction from *probability kernels*) if for all $N \in \mathbb{N}_{>0}$ we have $\sum_{i=1}^N \sum_{j=1}^N a_i a_j k_{\mathbb{X}}(x_i, x_j) \geq 0$ for $x_1, \dots, x_N \in \mathbb{X} \subset \mathbb{R}^n$ and $a_1, \dots, a_N \in \mathbb{R}$. A prominent example is the *squared exponential* (SQExp) kernel (Kanagawa et al. 2018; Rasmussen and Williams 2005):

$$k_{\mathbb{X}}(x, x') := \sigma_f^2 \exp \left(-\frac{1}{2} (x - x')^\top \Sigma^{-1} (x - x') \right), \quad \Sigma := \text{diag}(\sigma_l)^2, \quad (4)$$

with amplitude $\sigma_f^2 \geq 0$ and lengthscale coefficients $\sigma_l \in \mathbb{R}^n$. For this work, we assume that all kernels are bounded on their domain, i.e., $\mathbb{E}[k_{\mathbb{X}}(x, x)] < \infty$, $x \in \mathbb{X}$. Given a kernel $k_{\mathbb{X}}$ on a non-empty set $\mathbb{X}$, there exists a unique corresponding *reproducing kernel Hilbert space* (RKHS) $\mathcal{H}_{k_{\mathbb{X}}}$ of functions $f: \mathbb{X} \rightarrow \mathbb{R}$ equipped with an inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_{k_{\mathbb{X}}}}$ with the celebrated *reproducing property* such that for any function $f \in \mathcal{H}_{k_{\mathbb{X}}}$ and $x \in \mathbb{X}$ we have $f(x) = \langle f, k_{\mathbb{X}}(\cdot, x) \rangle_{\mathcal{H}_{k_{\mathbb{X}}}}$. Note that $\phi_{\mathbb{X}} := k_{\mathbb{X}}(\cdot, x): \mathbb{X} \rightarrow \mathcal{H}_{k_{\mathbb{X}}}$ is a real-valued function, which is also called an implicit *canonical embedding* or *feature map* such that $k_{\mathbb{X}}(x, x') = \langle \phi_{\mathbb{X}}(x), \phi_{\mathbb{X}}(x') \rangle_{\mathcal{H}_{k_{\mathbb{X}}}}$ for all $x, x' \in \mathbb{X}$. For an RKHS $\mathcal{H}_{k_{\mathbb{X}}}$, we use the associated feature map $\phi_{\mathbb{X}}$ and kernel $k_{\mathbb{X}}$ interchangeably for ease of notation and comprehensibility. The inner product induces the norm $\|f\|_{\mathcal{H}_{k_{\mathbb{X}}}} := \sqrt{\langle f, f \rangle_{\mathcal{H}_{k_{\mathbb{X}}}}}$ of the RKHS. A Hilbert space $\mathcal{H}$ is said to be separable if it admits a countable dense subset. Throughout this paper, we assume that all RKHSs are *separable*. Refer to the monograph by Berlinet and Thomas-Agnan (2004) for a comprehensive study on RKHSs. Given N i.i.d. samples $\hat{X}_N := [\hat{x}_i]_{i=1}^N$ with $\hat{x}_i \in \mathbb{X}$, the *Gram matrix* of $k_{\mathbb{X}}$ is given by $K_{\hat{X}}^N := [k_{\mathbb{X}}(\hat{x}_i, \hat{x}_j)]_{i,j=1}^N$. Furthermore, we define the vector-valued function $k_{\hat{X}}^N(x) := [k_{\mathbb{X}}(x, \hat{x}_i)]_{i=1}^N$.

Tensor product Hilbert spaces. For two kernels $k_{\mathbb{X}}: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ and $k_{\mathbb{Y}}: \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}$ with associated RKHSs $\mathcal{H}_{\mathbb{X}}$ and $\mathcal{H}_{\mathbb{Y}}$, $\mathcal{H}_{\mathbb{X} \otimes \mathbb{Y}} := \mathcal{H}_{\mathbb{X}} \otimes \mathcal{H}_{\mathbb{Y}}$ is the *tensor product of the Hilbert spaces* $\mathcal{H}_{\mathbb{X}}$ and $\mathcal{H}_{\mathbb{Y}}$ with reproducing kernel $k_{\mathbb{X} \otimes \mathbb{Y}}((x, y), (x', y')) = k_{\mathbb{X}}(x, x') k_{\mathbb{Y}}(y, y')$, for $x, x' \in \mathbb{X}$, $y, y' \in \mathbb{Y}$. $\mathcal{H}_{\mathbb{X}} \otimes \mathcal{H}_{\mathbb{Y}}$ is equipped with the inner product $\langle \phi_{\mathbb{X}} \otimes \phi_{\mathbb{Y}}, \phi'_{\mathbb{X}} \otimes \phi'_{\mathbb{Y}} \rangle_{\mathcal{H}_{\mathbb{X}} \otimes \mathcal{H}_{\mathbb{Y}}} = \langle \phi_{\mathbb{X}}, \phi'_{\mathbb{X}} \rangle_{\mathcal{H}_{\mathbb{X}}} \langle \phi_{\mathbb{Y}}, \phi'_{\mathbb{Y}} \rangle_{\mathcal{H}_{\mathbb{Y}}}$, for $\phi_{\mathbb{X}}, \phi'_{\mathbb{X}} \in \mathcal{H}_{\mathbb{X}}$ and $\phi_{\mathbb{Y}}, \phi'_{\mathbb{Y}} \in \mathcal{H}_{\mathbb{Y}}$, where its elements $\phi_{\mathbb{X}} \otimes \phi_{\mathbb{Y}}$ are called *tensors*. The associated norm is given by $\|\phi_{\mathbb{X}} \otimes \phi_{\mathbb{Y}}\|_{\mathcal{H}_{\mathbb{X}} \otimes \mathcal{H}_{\mathbb{Y}}} = \|\phi_{\mathbb{X}}\|_{\mathcal{H}_{\mathbb{X}}} \|\phi_{\mathbb{Y}}\|_{\mathcal{H}_{\mathbb{Y}}}$ (Arai 2018, Chapter 2). For feature vectors $\phi_{\mathbb{X}}(x)$ and $\phi_{\mathbb{Y}}(y)$, $\phi_{\mathbb{X}}(x) \otimes \phi_{\mathbb{Y}}(y)$ indicates their outer product.

3.1 Embedding Probability Measures

To reason about the expected value of a random variable, embedding the variable into a (higher dimensional) space is a well-established concept in ML (Schölkopf and A. J. Smola 2002; Steinwart and Christmann 2008). The (*kernel*) *mean embedding* (ME) follows the same reasoning and represents the projection of a probability measure into an RKHS (A. Smola et al. 2007).

DEFINITION 3 (MEAN EMBEDDING (ME)). Given an RKHS $\mathcal{H}_{k_{\mathbb{X}}}$ induced by a kernel $k_{\mathbb{X}}: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$, the mean embedding of a probability measure $p: \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$ is computed via the mean map $\mu_{k_{\mathbb{X}}}: \mathcal{P}(\mathbb{X}) \rightarrow \mathcal{H}_{k_{\mathbb{X}}}$ as

$$\mu_{k_{\mathbb{X}}}(p) := \mathbb{E}_p[\phi_{\mathbb{X}}(X)] = \int_{\mathbb{X}} \phi_{\mathbb{X}}(x) dp(x).$$

Note that the reproducing property of the kernel carries on to the ME, facilitating the computation of the expected value of a function $f \in \mathcal{H}_{k_{\mathbb{X}}}$ via the inner product (A. Smola et al. 2007), that is

$$\mathbb{E}_p[f(X)] = \langle f, \mu_{k_{\mathbb{X}}}(p) \rangle_{\mathcal{H}_{k_{\mathbb{X}}}}. \quad (5)$$

If $k_{\mathbb{X}}$ is *characteristic*, any probability measure $p \in \mathcal{P}(\mathbb{X})$ is injectively mapped to a unique ME $\mu_{k_{\mathbb{X}}}(p) \in \mathcal{H}_{k_{\mathbb{X}}}$ (Gretton et al. 2012). This gives rise to the definition of a distance metric between two probability measures in

terms of their embeddings in RKHS – the *maximum mean discrepancy* (MMD). For a characteristic kernel $k_{\mathbb{X}}$, the MMD between two prob. measures $p, p' \in \mathcal{P}(\mathbb{X})$ in $\mathcal{H}_{k_{\mathbb{X}}}$ is defined as $\|\mu_{k_{\mathbb{X}}}(p) - \mu_{k_{\mathbb{X}}}(p')\|_{\mathcal{H}_{k_{\mathbb{X}}}}$ (Gretton et al. 2012).

3.2 Embedding Conditional Probability Measures

Analogous to generic probability measures, there exists a similar notion for embedding *conditional* probability measures of the form $\mathbf{p}: \mathbb{X} \times \mathcal{B}(\mathbb{Y}) \rightarrow [0, 1]$ with realizations $Y \sim \mathbf{p}(\cdot | X = x)$ for a conditioning variable taking concrete values $x \in \mathbb{X}$. More specifically, we use the measure-theoretic *conditional mean embedding* (CME) introduced by Park and Muandet (2020). To this end, we equip the space of the conditioning random variable $X \in \mathbb{X}$ and the space of the target random variable $Y \in \mathbb{Y}$ with individual kernels $k_{\mathbb{X}}: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ and $k_{\mathbb{Y}}: \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}$, respectively.

DEFINITION 4 (CONDITIONAL MEAN EMBEDDING (CME)). *Given two RKHSs $\mathcal{H}_{k_{\mathbb{X}}}$ and $\mathcal{H}_{k_{\mathbb{Y}}}$ with the associated kernels $k_{\mathbb{X}}$ and $k_{\mathbb{Y}}$, the CME of a conditional probability measure $\mathbf{p}: \mathbb{X} \times \mathcal{B}(\mathbb{Y}) \rightarrow [0, 1]$ is an X -measurable random variable taking values in $\mathcal{H}_{k_{\mathbb{Y}}}$ given by*

$$\mu_{k_{\mathbb{Y}}|k_{\mathbb{X}}}(\mathbf{p})(\cdot) := \mathbb{E}_{\mathbf{p}}[\phi_{\mathbb{Y}}(Y) | X = \cdot].$$

Refer to Park and Muandet (2020) for a comprehensive mathematical dissemination of the CME. Analogous to the non-conditional case, we can compute the conditional expectation of a function $f \in \mathcal{H}_{k_{\mathbb{Y}}}$ via its inner product with the CME, i.e., almost surely

$$\mathbb{E}_{\mathbf{p}}[f(Y) | X = x] = \langle f, \mu_{k_{\mathbb{Y}}|k_{\mathbb{X}}}(\mathbf{p})(x) \rangle_{\mathcal{H}_{k_{\mathbb{Y}}}}. \quad (6)$$

Since the CME of a conditional measure $\mathbf{p}$ is generally unknown (and potentially infinite dimensional), we may obtain an empirical estimate from a finite set of training data.

PROPOSITION 2 (EMPIRICAL CME). *Let a finite dataset $(\hat{X}_N, \hat{Y}_N) := ([\hat{x}_i]_{i=1}^N, [\hat{y}_i]_{i=1}^N)$ of samples $(\hat{x}_i, \hat{y}_i) \in \mathbb{X} \times \mathbb{Y}$ from a conditional probability measure $\mathbf{p}: \mathbb{X} \times \mathcal{B}(\mathbb{Y}) \rightarrow [0, 1]$ be given such that $\hat{y}_i \sim \mathbf{p}(\cdot | X = \hat{x}_i)$. For kernels $k_{\mathbb{X}}$, $k_{\mathbb{Y}}$, the empirical CME of $\mathbf{p}$ given $(\hat{X}_N, \hat{Y}_N)$ is*

$$\hat{\mu}_{k_{\mathbb{Y}}|k_{\mathbb{X}}}^N(\cdot) := k_{\hat{X}}^N(\cdot)^\top \left[K_{\hat{X}}^N + N\lambda I_N \right]^{-1} \phi_{\mathbb{Y}}(\hat{Y}_N),$$

with a regularization constant $\lambda \geq 0$. For $N \rightarrow \infty$, the empirical CME converges in expectation, that is,

$$\mathbb{E}_{\mathcal{U}_{\mathbb{X}} \mathcal{U}_{\mathbb{Y}}} [\|\hat{\mu}_{k_{\mathbb{Y}}|k_{\mathbb{X}}}^N(X, U) - \mu_{k_{\mathbb{Y}}|k_{\mathbb{X}}}(\mathbf{p})(X, U)\|_{\mathcal{H}_{k_{\mathbb{Y}}}}] \rightarrow 0.$$

Here and in the following, we assume that $[K_{\hat{X}}^N + N\lambda I_N]$ (and equivalent terms) are invertible. Note that this is always true if the regularization constant λ is strictly positive. By virtue of the reproducing property we have for any function $f \in \mathcal{H}_{k_{\mathbb{Y}}}$ almost surely that

$$\mathbb{E}_{\mathbf{p}}[f(Y) | X = x] \approx \langle f, \hat{\mu}_{k_{\mathbb{Y}}|k_{\mathbb{X}}}^N(x) \rangle_{\mathcal{H}_{k_{\mathbb{Y}}}} = k_{\hat{X}}^N(x)^\top \left[K_{\hat{X}}^N + N\lambda I_N \right]^{-1} f(\hat{Y}_N). \quad (7)$$

Note that the empirical estimate improves in probability as more data becomes available.

4 Data-Driven Barrier Formulation

In this section, we establish the supermartingale-like property (c) in Definition 2 for a system with unknown transition kernel $\mathbf{t}$ using its CME based on i.i.d. training data generated from the unknown true system (3), that is, data of the form

$$(\hat{X}_N, \hat{U}_N, \hat{X}_N^+) := ([\hat{x}_i]_{i=1}^N, [\hat{u}_i]_{i=1}^N, [\hat{x}_i^+]_{i=1}^N), \quad \text{with} \quad \hat{x}_i^+ \sim \mathbf{t}(\cdot | X = \hat{x}_i, U = \hat{u}_i), \quad (8)$$

generated from the unknown black-box system as shown in (3).

Table 1. Overview of the kernels and RKHSs.

RKHS	Kernel	Feature Map
$\mathcal{H}_x$	$k_x: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$	$\phi_x: \mathbb{X} \rightarrow \mathcal{H}_x$
$\mathcal{H}_u$	$k_u: \mathbb{U} \times \mathbb{U} \rightarrow \mathbb{R}$	$\phi_u: \mathbb{U} \rightarrow \mathcal{H}_u$
$\mathcal{H}_+$	$k_+: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$	$\phi_+: \mathbb{X} \rightarrow \mathcal{H}_+$
$\mathcal{H}_{xu}$	$k_{xu}: (\mathbb{X} \times \mathbb{U}) \times (\mathbb{X} \times \mathbb{U}) \rightarrow \mathbb{R}$	not used
$\mathcal{G}$	$\Gamma: (\mathbb{X} \times \mathbb{U}) \times (\mathbb{X} \times \mathbb{U}) \rightarrow \mathcal{L}(\mathcal{H}_+)$	not used

Our theoretical analysis assumes uniform i.i.d. samples to simplify concentration arguments — a common assumption in formal settings. In particular, this assumption stems from established results (e.g., [Park and Muandet 2020](#)) on the consistency of the empirical CME (see Proposition 2). Whilst access to i.i.d. data of the form (8) is a strong assumption, the guarantees established in this manuscript can be preserved under substantially weaker conditions, such as ergodicity, geometric mixing, or risk-based notions, by replacing the nominal sample size N with an effective sample size N_{eff} that accounts for temporal dependence (see, e.g., [Massiani et al. 2024](#); [Steinwart and Christmann 2009](#); [Ziemann and Tu 2022](#)). Whilst the development of new concentration inequalities lies beyond the scope of this paper, it has been shown that kernel-based methods such as the CME can be applied even with dependent data — such as data sequences — in many settings (see, e.g., [Massiani et al. 2024](#); [Steinwart and Christmann 2009](#); [Ziemann and Tu 2022](#)).

In the following, we will use three kernels, one for each of the three components $\hat{X}_N$, $\hat{U}_N$, and $\hat{X}_N^+$, namely

$$k_x: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}, \quad k_u: \mathbb{U} \times \mathbb{U} \rightarrow \mathbb{R}, \quad k_+: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}, \quad (9)$$

with their associated RKHSs $\mathcal{H}_x$, $\mathcal{H}_u$, and $\mathcal{H}_+$, respectively. We will assume that k_+ is characteristic. Furthermore, we define the tensor product space $\mathcal{H}_{xu} := \mathcal{H}_x \otimes \mathcal{H}_u$ with the associated kernel $k_{xu} := k_x \cdot k_u$. Table 1 gives a brief summary of the notation used in the remainder of the paper.

To reason about the unknown dynamics in a mathematically rigorous way, we must restrict the system dynamics to a known function space. In particular, we will assume that the CME of the transition kernel $\mathbf{t}$ lives in a vector-valued RKHS of mappings $\mathbb{X} \times \mathbb{U} \rightarrow \mathcal{H}_+$, which we indicate by $\mathcal{G}$ ([Park and Muandet 2020](#), Section 2.3).² We will call $\mathcal{G}$ the latent *hypothesis (function) space* of the system. Under standard assumptions, $\mathcal{G}$ is dense in $L_2(\mathbb{X} \times \mathbb{U}, \mathcal{B}(\mathbb{X} \times \mathbb{U}), \mathbf{t}; \mathcal{H}_+)$, where $L_2(\mathbb{X} \times \mathbb{U}, \mathcal{B}(\mathbb{X} \times \mathbb{U}), \mathbf{t}; \mathcal{H}_+)$ is the space of strongly $(\mathcal{B}(\mathbb{X} \times \mathbb{U}) - \mathcal{B}(\mathcal{H}_+))$ -measurable and Bochner square-integrable functions $\mathbb{X} \times \mathbb{U} \rightarrow \mathcal{H}_+$ with respect to $\mathbf{t}$, with $\mathcal{B}(\mathbb{X} \times \mathbb{U})$ and $\mathcal{B}(\mathcal{H}_+)$ being the Borel σ -fields of $\mathbb{X} \times \mathbb{U}$ and $\mathcal{H}_+$, respectively ([Z. Li et al. 2022](#)). This assumption imposes minimal structural constraints on the underlying dynamics. In fact, in contrast to many existing approaches based on finite Taylor approximations of the unknown dynamics — which are limited to approximating analytic functions of bounded degree — the RKHS framework adopted here allows approximation of a far broader class, e.g., of all smooth functions via the squared exponential kernel. This property affords significantly greater modeling flexibility and expressiveness. Moreover, the choice of kernel provides a natural mechanism for incorporating prior knowledge, where available. With this, we raise the following standard assumption.

²For a fixed choice of $(x, u) \in \mathbb{X} \times \mathbb{U}$, the vector-valued RKHS $\mathcal{G}$ reduces to the regular RKHS $\mathcal{H}_+$.

ASSUMPTION 1. Let the CME $\mu_{k_+|k_{xu}}(\mathbf{t})$ be well-specified, i.e., $\mu_{k_+|k_{xu}}(\mathbf{t}) \in \mathcal{G}$. Furthermore, given confidence $1 - \rho \in (0, 1)$, let a bound $\varepsilon \geq 0$ be known such that

$$\mathbb{P} \left(\left\| \mu_{k_+|k_{xu}}(\mathbf{t}) - \hat{\mu}_{k_+|k_{xu}}^N \right\|_{\mathcal{G}} \leq \varepsilon \right) \geq 1 - \rho,$$

for the empirical CME based on $(\hat{X}_N, \hat{U}_N, \hat{X}_N^+)$, which is given by

$$\hat{\mu}_{k_+|k_{xu}}^N := k_{\hat{X}\hat{U}}^N(\cdot)^\top \left[K_{\hat{X}\hat{U}}^N + N\lambda I_N \right]^{-1} \phi_+(\hat{X}_N^+). \quad (10)$$

It is common practice to introduce a robustness radius ε (as in Assumption 1) to account for uncertainty in empirical estimates such as the estimated conditional expectation (7) by constructing an RKHS ambiguity set centered at the empirical CME. Although theoretical results exist that relate this radius to the probability of the true CME of the black-box dynamics $\mathbf{t}$ lying within the ambiguity set — achieving minimax rates of order $O(\log(N)/N)$ (Z. Li et al. 2022; Mollenhauer et al. 2022) — such concentration bounds tend to be excessively conservative in practice. Consequently, the radius is often chosen manually or calibrated using data-driven techniques such as MMD-based bootstrapping (see, e.g., Nemmour et al. 2022, for an approach using non-conditional MEs).

To learn CBCs from data, we start by reformulating the conditional expectation in the left-hand side of condition (c) in Definition 2 via the inner product with the CME $\mu_{k_+|k_{xu}}(\mathbf{t})$ as shown in (6), that is, almost surely

$$\mathbb{E}_{\mathbf{t}}[B(X^+)|X = x, U = u] = \langle B, \mu_{k_+|k_{xu}}(\mathbf{t})(x, u) \rangle_{\mathcal{H}_+}.$$

Since the CME $\mu_{k_+|k_{xu}}(\mathbf{t})$ is unknown, we construct an ambiguity set $C_\varepsilon^N \subset \mathcal{G}$ centered at the empirical CME $\hat{\mu}_{k_+|k_{xu}}^N$ constructed from the data in (8) (via Proposition 2) and choose an MMD radius $\varepsilon \geq 0$ such that the CME of $\mathbf{t}$ lies within the ambiguity set C_ε^N with a confidence of at least $1 - \rho \in (0, 1)$, i.e., such that $\mathbb{P}(\mu_{k_+|k_{xu}}(\mathbf{t}) \in C_\varepsilon^N) \geq 1 - \rho$, with the ambiguity set

$$C_\varepsilon^N := \left\{ \mu \in \mathcal{G} \mid \left\| \mu - \hat{\mu}_{k_+|k_{xu}}^N \right\|_{\mathcal{G}} \leq \varepsilon \right\}. \quad (11)$$

With this, we obtain the following result.

THEOREM 1. Let data $(\hat{X}_N, \hat{U}_N, \hat{X}_N^+)$ in (8) from an unknown MDP $\mathbf{M}$ and kernels k_{xu} and k_+ (characteristic) be given. Consider the resulting empirical CME $\hat{\mu}_{k_+|k_{xu}}^N$ in (10) and the ambiguity set C_ε^N in (11) with confidence bound $1 - \rho$. If there exists a function $B: \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, $B \in \mathcal{H}_+$, satisfying

$$\forall x \in \mathbb{X}, \exists u \in \mathbb{U}, \forall \mu \in C_\varepsilon^N: \langle B, \mu(x, u) \rangle_{\mathcal{H}_+} - B(x) \leq c, \quad (12)$$

for some constant $c \geq 0$, then, B satisfies CBC condition (c) of Def. 2 w.r.t. $\mathbf{M}$ with probability at least $1 - \rho$.

With the following theorem, we provide a way of establishing Theorem 1 by reformulating the left-hand side of (12) for the RKHS norm-ball ambiguity set in (11). For this, let $K_{\hat{X}^+}^N$ and $K_{\hat{X}\hat{U}}^N$ be the Gram matrices associated with $(k_+, \hat{X}^+)$ and $(k_{xu}, [(\hat{x}_i, \hat{u}_i)]_{i=1}^N)$, respectively.

THEOREM 2. Consider the setup of Theorem 1. If there exists a function $B: \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, $B \in \mathcal{H}_+$, with some $\bar{B} \geq \|B\|_{\mathcal{H}_+}$ such that

$$\forall x \in \mathbb{X}, \exists u \in \mathbb{U}: w(x, u)^\top B(\hat{X}_N^+) - B(x) \leq c - \varepsilon \bar{B} \kappa(x, u), \quad (13)$$

for some constant $c \geq 0$, $\kappa(x, u) := \sqrt{k_x(x, x)}\sqrt{k_u(u, u)}$, and weighting function

$$w(x, u)^\top := k_{\hat{X}\hat{U}}^N(x, u)^\top \left[K_{\hat{X}\hat{U}}^N + N\lambda I_N \right]^{-1}, \quad (14)$$

with constant $\lambda \geq 0$, then, B satisfies CBC condition (c) of Definition 2 with respect to $\mathbf{M}$ with probability at least $1 - \rho$. Note that this implies the existence of a policy $\pi: \mathbb{X} \rightarrow \mathbb{U}$ satisfying (13).

PROOF. We start by rewriting the inner product in (12) as

$$\langle B, \mu(x, u) \rangle_{\mathcal{H}_+} = \langle B, \mu(x, u) - \hat{\mu}_{k_+|k_{xu}}^N(x, u) \rangle_{\mathcal{H}_+} + \langle B, \hat{\mu}_{k_+|k_{xu}}^N(x, u) \rangle_{\mathcal{H}_+}.$$

Via (7), the latter term yields $w(x, u)^\top B(\hat{X}_N^+)$. For the prior term, let $\Gamma: (\mathbb{X} \times \mathbb{U}) \times (\mathbb{X} \times \mathbb{U}) \rightarrow \mathcal{L}(\mathcal{H}_+)$ be the operator-valued positive definite kernel of $\mathcal{G}$ (see, e.g., Z. Li et al. 2022, Definition 1) given by $\Gamma((x, u), (x', u')) := k_{xu}((x, u), (x', u')) \text{Id}_{\mathcal{H}_+}$, where $\mathcal{L}(\mathcal{H}_+)$ is the Banach space of bounded linear operators from $\mathcal{H}_+$ to $\mathcal{H}_+$ and $\text{Id}_{\mathcal{H}_+}$ is the identity operator on $\mathcal{H}_+$. Then, it follows from the reproducing property of Γ (cf. Z. Li et al. 2022, Equation (2)ff.) that

$$\begin{aligned} \langle B, \mu(x, u) - \hat{\mu}_{k_+|k_{xu}}^N(x, u) \rangle_{\mathcal{H}_+} &= \langle \Gamma(\cdot, (x, u))B, \mu - \hat{\mu}_{k_+|k_{xu}}^N \rangle_{\mathcal{G}}, \\ &\leq \varepsilon \|\Gamma(\cdot, (x, u))B\|_{\mathcal{G}}, \\ &= \varepsilon \sqrt{\langle \Gamma(\cdot, (x, u))B, \Gamma(\cdot, (x, u))B \rangle_{\mathcal{G}}}, \\ &= \varepsilon \sqrt{\langle B, \Gamma((x, u), (x, u))B \rangle_{\mathcal{H}_+}}, \\ &\leq \varepsilon \bar{B} \sqrt{\|\Gamma((x, u), (x, u))\|}, \\ &= \varepsilon \bar{B} \sqrt{k_{xu}((x, u), (x, u))}. \end{aligned}$$

Reordering yields (13), concluding the proof. $\square$

Intuitively, the conservatism introduced due to the worst-case approach w.r.t. the unknown transition kernel $\mathbf{t}$ is captured by the term $\varepsilon \bar{B} \kappa$, thus proportional to the radius ε . Ensuring that (13) holds for an adequate offset $\varepsilon \bar{B} \kappa$ gives the necessary headroom to provide guarantees for dynamics deviating from the empirical observations, captured in the CME. As we design the CBC B , an upper bound $\bar{B}$ can be estimated via analytical expressions or numerical estimation (Berliner and Thomas-Agnan 2004; Kanagawa et al. 2018). With the following proposition, we quantify the safety probability of a black-box system based on the results from Theorem 2.

THEOREM 3 (DATA-DRIVEN FINITE-HORIZON SAFETY). *Consider the data-based setup in Theorem 2 and a safety specification $\psi_{\text{safe}} = (\mathbb{X}_u, T)$. Suppose there exists a function B satisfying the conditions in Theorem 2 for a constant $c \geq 0$ and policy π . If there exists a constant $\eta \in [0, 1)$ such that*

- (a) $\forall x_0 \in \mathbb{X}_0: B(x_0) \leq \eta$; and
- (b) $\forall x_u \in \mathbb{X}_u: B(x_u) \geq 1$;

are satisfied, then, with probability at least $1 - \rho$ we have

$$P_{\text{safe}}^\pi(\mathbf{M}) \geq 1 - (\eta + cT).$$

Theorem 3 follows trivially from Theorem 2 and Proposition 1.

5 Extension to General Classes of Temporal Logic Specifications

The results presented in the previous section can be extended to develop a data-driven approach for computing a lower bound on the satisfaction of more general specifications beyond safety. Examples include developing

AI systems satisfying liveness, fairness, and complex temporal properties in scenarios involving cooperation or human interaction.

5.1 Linear Temporal Logic over Finite Traces

The work by Jagtap, Soudjani, et al. (2021) introduces a model-based procedure for using CBCs to compute a lower bound on the probability of satisfying *LTL specifications over finite traces* (LTL_F). The procedure comprises the following steps:

- (i) The specification is negated;
- (ii) The negated specification is translated into a DFA;
- (iii) The resulting DFA is decomposed into a sequence of reachability tasks;
- (iv) For each individual task, an upper bound on its satisfaction probability is computed using CBCs; and
- (v) The individual bounds are composed, based on the structure of the DFA, to obtain a lower bound on the probability of satisfying the original LTL_F specification.

This procedure relies primarily on knowledge of the specification and uses CBCs as a subroutine to establish bounds on the individual reachability probabilities. In scenarios where a model of the system is *not* available but an LTL_F specification is given, the identical procedure outlined above can be applied in conjunction with data-driven CBCs constructed using the results of Section 4. Further technical details are omitted here, as the overall approach remains unchanged from that of Jagtap, Soudjani, et al. (2021), with the only difference being the substitution of model-based CBCs by their data-driven counterparts developed in Section 4.

5.2 Specifications Modeled via Co-Büchi Automata

In an effort to address a broader class of specifications beyond LTL_F , the work by Anand et al. (2024) proposes a method for computing a lower bound on the probability of satisfying ω -regular specifications in networked stochastic systems. However, the approach is based on a fundamental simplification: the original specification is reduced to an automaton with a co-Büchi acceptance condition — that is, an acceptance criterion requiring certain states to be visited only finitely often. Consequently, the method is effectively restricted to a *strict subset* of ω -regular specifications.

The approach of Anand et al. (2024) comprises the following steps:

- (i) The ω -regular specification is modeled as a DFA with a Rabin acceptance condition;
- (ii) The specification is strengthened by replacing the Rabin condition with a co-Büchi acceptance condition;
- (iii) A subset of satisfying traces with repeating cycles is extracted from the resulting automaton;
- (iv) These traces are decomposed into sequences of safety specifications;
- (v) CBCs are employed to compute individual lower bounds on the probability of satisfying these safety specifications; and
- (vi) The individual bounds are combined to obtain a lower bound on the probability of satisfying the original specification.

As for the approach proposed by Jagtap, Soudjani, et al. (2021) reviewed previously, the above procedure relies primarily on knowledge of the specification and uses CBCs as a subroutine to obtain bounds on the individual safety probabilities. Thus, the data-driven results of Section 4 can be used to extend the procedure to a model-free setting. The underlying steps remain unchanged, and so the mathematical details are omitted.

5.3 ω -Regular Specifications

To address the general class of ω -regular specifications, we build on the concept of *Streett supermartingales*, proposed by Abate, Giacobbe, et al. (2024), and provide a data-driven version that can be established without model knowledge and based only on data from the system. For this, let a *Streett pair*, i.e., a pair of sets (F, I) with

$F, I \in \mathcal{B}(\mathbb{X})$, be given. The following result is an extension of Theorem 3 by [Abate, Giacobbe, et al. \(2024\)](#) to systems with control inputs.

PROPOSITION 3 (STREETT SUPERMARTINGALE). *For an MDP $\mathbf{M} = (\mathbb{X}, \mathbb{X}_0, \mathbb{U}, \mathbf{t})$ and a Streett pair (F, I) , suppose there exist two functions $V: \mathbb{X} \rightarrow \mathbb{R}$, $B: \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, and positive constants $\chi, \nu > 0$, such that for all $x \in \mathbb{X}$ there is a $u \in \mathbb{U}$ satisfying the following conditions:*

- (a) $V(x^+) \leq V(x)$ almost surely with $x^+ \sim \mathbf{t}(\cdot | X = x, U = u)$, if $x \in W$;
- (b) $\mathbb{E}_{\mathbf{t}}[B(X^+) | X = x, U = u] - B(x) \leq -\chi$, if $x \in (F \setminus I) \cap W$;
- (c) $\mathbb{E}_{\mathbf{t}}[B(X^+) | X = x, U = u] - B(x) \leq \nu$, if $x \in I \cap W$; and
- (d) $\mathbb{E}_{\mathbf{t}}[B(X^+) | X = x, U = u] - B(x) \leq 0$, if $x \in W \setminus (F \cup I)$;

where $W = \{x \in \mathbb{X} \mid V(x) \leq 1\}$. Then, we call B a Streett supermartingale and there exists a control policy $\pi: \mathbb{X} \rightarrow \mathbb{U}$ under which the MDP starting from any $x_0 \in W$ almost surely either visits F a finite number of times or visits I infinitely often.

From the four conditions in Proposition 3, the last three involve the computation of conditional expectations. These conditions can be replaced with their data-driven version using CMEs. The first condition, however, requires the inequality to hold almost surely. To satisfy the first condition, we raise the following assumption.

ASSUMPTION 2. *For the MDP $\mathbf{M} = (\mathbb{X}, \mathbb{X}_0, \mathbb{U}, \mathbf{t})$, the support of $\mathbf{t}$ is known, i.e., there is a set $S(x, u)$ such that $\mathbf{t}(S(x, u) | x, u) = 1$ for all $x \in \mathbb{X}$ and $u \in \mathbb{U}$.*

REMARK 2. *When the system is directly represented with dynamics in (1), the above assumption can be satisfied by requiring that the deterministic part of the dynamics f and the support $\mathbb{W}$ of w_t are known, but the measure $p_w \in \mathcal{P}(\mathbb{W})$ of w_t is unknown. In this case, $S(x, u) = \{f(x, u, w) \mid w \in \mathbb{W}\}$.*

The next theorem presents the conditions for the CME-based computation of Streett supermartingales under Assumption 2.

THEOREM 4. *Let data $(\hat{X}_N, \hat{U}_N, \hat{X}_N^+)$ in (8) from an unknown MDP $\mathbf{M}$ under Assumption 2 and kernels k_{xu} and k_+ (characteristic) be given. Consider the resulting empirical CME $\hat{\mu}_{k_+ | k_{xu}}^N$ in (10) and the ambiguity set C_ε^N in (11) with confidence bound $1 - \rho$. For a Streett pair (F, I) , suppose there exist two functions $V: \mathbb{X} \rightarrow \mathbb{R}$, $B: \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, $B \in \mathcal{H}_+$, and positive constants $\chi, \nu > 0$, such that for all $x \in \mathbb{X}$ there is a $u \in \mathbb{U}$ such that*

- (a) $V(x^+) \leq V(x)$, for all $x^+ \in S(x, u)$, if $x \in I$;
- (b) $w(x, u)^\top B(\hat{X}_N^+) - B(x) \leq -\chi - \varepsilon \bar{B} \kappa(x, u)$, if $x \in (F \setminus I) \cap W$;
- (c) $w(x, u)^\top B(\hat{X}_N^+) - B(x) \leq \nu - \varepsilon \bar{B} \kappa(x, u)$, if $x \in I \cap W$; and
- (d) $w(x, u)^\top B(\hat{X}_N^+) - B(x) \leq -\varepsilon \bar{B} \kappa(x, u)$, if $x \in W \setminus (F \cup I)$;

where $W = \{x \in \mathbb{X} \mid V(x) \leq 1\}$, $\bar{B} \geq \|B\|_{\mathcal{H}_+}$, $\kappa(x, u) := \sqrt{k_x(x, x)} \sqrt{k_u(u, u)}$, and $w(x, u)^\top$ given in (14). Note that this implies the existence of a policy $\pi: \mathbb{X} \rightarrow \mathbb{U}$ satisfying the conditions above. Then, we call B a Streett supermartingale and the MDP $\mathbf{M}$ starting from any $x_0 \in W$ and under the policy π almost surely either visits F a finite number of times or visits I infinitely often.

Any ω -regular specification can be modeled with a deterministic finite automaton with a Streett acceptance condition ([Baier and Katoen 2008](#); [Streett 1981](#)). The acceptance condition consists of a finite set of Streett pairs $\{(F_i, I_i), i = 1, 2, \dots, p\}$. A trajectory satisfies the acceptance condition if for each $i = 1, 2, \dots, p$, it either visits F_i a finite number of times or visits I_i infinitely often. Hence, almost-sure satisfaction of an ω -regular specification can be checked by finding Streett supermartingales for each pair, where each Streett supermartingale associates a

winning domain W_i . Then, a lower bound on the satisfaction probability can be computed by solving reachability to the intersection of the winning domains $\cap_i W_i$.

THEOREM 5. *Consider the data-driven setup in Theorem 4 and a Streett condition $\{(F_i, I_i) : i = 1, \dots, k\}$. If for a shared control policy $\pi: \mathbb{X} \rightarrow \mathbb{U}$ each pair (F_i, I_i) admits a Streett supermartingale B_i according to Theorem 4 with an associated winning domain W_i , then the trajectories under π satisfy the Streett acceptance condition almost surely from any initial state $x_0 \in \cap_i W_i$.*

6 Problem Characterization and Comparison of Solution Approaches

For a comprehensive investigation of computational techniques, we focus on *safety verification* in the rest of the paper. Therefore, the control policy is assumed to be given, and the controlled system is denoted by the same notation after eliminating the input from the mathematical expressions.

We start by summarizing and characterizing the optimization problem identified for computing CBCs from data, which implements the conditions from Theorems 2–3:

$$\min_{B, c, \eta} \quad \eta + cT, \quad (15a)$$

$$\text{s.t.} \quad \forall x_0 \in \mathbb{X}_0: B(x_0) \leq \eta, \quad (15b)$$

$$\forall x_u \in \mathbb{X}_u: B(x_u) \geq 1, \quad (15c)$$

$$\forall x \in \mathbb{X}: w(x)^\top B(\hat{X}_N^+) - B(x) \leq c - \varepsilon \bar{B} \kappa, \quad (15d)$$

$$\forall x \in \mathbb{X}: B(x) \geq 0, \quad (15e)$$

$$c \geq 0, \eta \in [0, 1), B \in \mathcal{H}_+,$$

with coefficients $\kappa \geq \sup_{x \in \mathbb{X}} \sqrt{k_x(x, x)}^3$, $\bar{B} \geq \|B\|_{\mathcal{H}_+}$, and weight function

$$w(x)^\top := k_{\hat{X}}^N(x)^\top \left[K_{\hat{X}}^N + N\lambda I_N \right]^{-1}.$$

Clearly, (15) constitutes an optimization problem with infinitely many constraints. Furthermore, recall that any barrier $B \in \mathcal{H}_+$ can be represented based on a linear combination of spatial support vectors $\bar{x}_1, \bar{x}_2, \dots \in \mathbb{X}$ and coefficients $b_i \in \mathbb{R}$, namely

$$B(x) := \sum_i b_i k_+(\bar{x}_i, x). \quad (16)$$

For popular choices of the kernel k_+ such as the Gaussian, SQExp, or Matérn kernels, the semi-infinite program (15) is generally not convex. We elaborate on possible (partial) dualizations in Subsection 6.1. The optimization problem (15) therefore belongs to the general category of *non-convex semi-infinite problems* (SIPs). There exists a collection of approaches from the robust optimization and semi-finite programming literature relevant to this problem category (Djelassi et al. 2021). In the following subsections, we provide an overview of the approaches we tested in implementation. More concretely, we juxtapose the frameworks based on their theoretical complexity and report our experimental findings. We also note that, as all the approaches discussed take (15) as a starting point, they are all based on the fundamental Assumption 1. That being said, we emphasize that this assumption can be set arbitrarily general and we do not raise/consider any strong assumptions such as linear, affine, or polynomial dynamics that could be exploited by special solvers.

³The reformulation of κ for translation invariant kernels k_x is exact.

6.1 Monolithic Solutions

Control barriers have been developed to reduce the computational burden associated with formal methods. In particular, they have been developed as a supposedly “abstraction-free” alternative to address the so-called *curse of dimensionality*, that arises from approaches based on the spatial partitioning of the system domain, aptly dubbed “abstraction-based”. Over the years, a list of customary assumptions has been established in the literature to obtain feasible formulations of the barrier problem that can be efficiently solved monolithically, i.e., without spatial or spectral subdivision. In the model-based case, the most typical assumptions include control-affine polynomial dynamics and polynomial barriers (both of known maximum degree). In the model-free/data-driven case, additional Lipschitz bounds on the dynamics and intermediate functions (Salamati et al. 2024) are inevitable, knowledge assumed given even though the dynamics are deemed unknown. Evidently, it is unclear how these assumptions can be rigorously justified.

6.1.1 Sum-of-Squares. The *sum-of-squares* (SOS) approach has gained notable popularity for establishing barriers in polynomial system models (Prajna, Papachristodoulou, et al. 2005), with Schön, Zhong, et al. (2024) being the first to cast *data-driven* barrier problems into SOS form. This method relies on several assumptions: polynomial barrier functions, sets $\mathbb{X}, \mathbb{X}_0, \mathbb{X}_u \subset \mathbb{R}^n$ given in semi-algebraic form, and, in the data-driven case, dynamics that are polynomially dependent on the state x (i.e., $\mathcal{H}_x$ limited to the RKHS of the polynomial kernel). Furthermore, the computational complexity hinges directly on the maximum polynomial degree appearing in the resulting program. There exist extensions to bases beyond polynomials, known as *kernel SOS* (KSOS) (Marteau-Ferey et al. 2020). For instance, Bagnell and Farahmand (2015) restrict a kernel machine to functions of SOS form from the product RKHS $\mathcal{H}' := \mathcal{H} \otimes \mathcal{H}$ based on a kernel

$$k'(x, x') := \langle \phi(x) \otimes \phi(x), \phi(x') \otimes \phi(x') \rangle_{\mathcal{H}'} = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}^2.$$

We provide further examples of such methods for *hard-constraining* kernel-machines in Subsection 6.2. Note, that in the multivariate case $n > 1$ SOS functions only constitute a subset of non-negative functions, hence introducing some conservatism (Parrilo 2003). Nevertheless, the SOS approach is very popular for obtaining well-behaved/convex relaxations of complex programs such as the SIP in (15) as it avoids domain discretization and nonlinear optimization.

Although nomenclature-wise abstraction-free, the SOS approach is not immune to the curse of dimensionality, as it covertly associates a semi-definite program (SDP) for which the number of variables scales exponentially with the system dimension n and the maximum polynomial degree d . The resulting complexity of the SOS is polynomial in the number of variables. Additionally, solving barrier problems such as the SIP (15) usually requires the expression of the sets $\mathbb{X}, \mathbb{X}_0, \mathbb{X}_u$, the constants η and $c - \varepsilon \bar{B} \kappa$, and additional Lagrange multipliers in the chosen kernel basis, which can be cumbersome (especially in an infinite basis) and increase the size of the SDP. There exist sampling-based extensions that can (partly) bypass this additional overhead and even exploit sparsity under the assumption of poisedness (Cifuentes and Parrilo 2017; Lofberg and Parrilo 2004). As developing efficient SOS solvers is an active field of research, there exist many relevant recent advancements (Coey et al. 2022; Cunis and Legat 2022). Refer to the paper by Parrilo (2003) for a more detailed dissemination of the SOS approach as well as Jagtap, Soudjani, et al. (2021, Sec. V.C.1) and Schön, Zhong, et al. (2024) for its application to barrier certificates.

6.1.2 Dual Reformulation. In an effort to find similar variants of the SIP (15) that can be efficiently computed monolithically but that require less restrictive assumptions, *dual reformulation* of the constraints may be considered. For example, for a polyhedral set $\mathbb{X}_u := \{x_u \in \mathbb{X} \mid H_u x_u \leq h_u\}$ with matrix H_u and vector h_u of appropriate dimension, the equivalent dual to (15c) is given by

$$\min_{\lambda_u \geq 0} B^*(H_u^\top \lambda_u) - \lambda_u^\top h_u \leq -1, \quad (17)$$

where B^* is the Fenchel conjugate⁴ of B . Closed-form solutions for (17) can be obtained for specific choices of the kernel k , as shown with the following example.

EXAMPLE 2 (EXPONENTIAL KERNEL). Consider $B \in \mathcal{H}$ with $\mathcal{H}$ the RKHS of the exponential kernel $k(x, x') := \alpha e^{\beta x^\top x'}$, $\alpha, \beta > 0$. Then, B can be expressed in the form $B(x) = \alpha \sum_i a_i e^{\beta x_i^\top x}$, where $\{x_i\}_{i \in I} \subset \mathbb{X}$ is an orthogonal basis of $\mathbb{X}$. Furthermore, the Fenchel conjugate B^* is given by

$$B^*(\lambda) = \frac{1}{\beta} \sum_i \frac{1}{\|x_i\|} x_i^\top H_u^\top \lambda. \quad (18)$$

For an orthonormal basis $\{x_i\}_{i \in I}$ we additionally have $\|x_i\| = 1, \forall i \in I$. The derivation is omitted for brevity.

Even though this shows that we can indeed obtain convex closed-form reformulations of individual constraints for specific choices of the kernel, it is *not* possible to convexify the entire optimization problem (15) in this way. In fact, non-linear convexity of both constraints (15b) and (15c) is contradictory. Furthermore, the Kushner constraint (15d) is generally neither convex nor concave. We will revisit dual reformulation in the context of local optimization problems as used in branch-and-bound optimizers in Subsection 6.2.

6.2 Divide-and-Conquer Solutions

The main take-away from the previous subsection is that solving the non-convex SIP (15) in a monolithic fashion is only feasible under carefully chosen assumptions. But even then, the solution approach suffers — similar to abstraction-based approaches — from the curse of dimensionality. For those that strive for maximal flexibility and minimal assumptions, in this subsection, we discuss approaches that (at least partially) leverage a spatial subdivision of the domain of the system/barrier. Although bearing similarities to traditional abstraction-based approaches, they leverage the barrier as a mathematical vehicle for storing information and guiding the spatial abstraction process. In comparison, abstraction-based approaches often suffer from significant memory requirements and adaptive abstraction regimes rely on ad-hoc principles such as the approximation of local Lipschitz constants (Kazemi, Majumdar, et al. 2024; Soudjani and Abate 2013).

One-stage vs. Two-stage solutions. Approaches to solving the SIP (15) can be applied in both one-stage and two-stage frameworks. One-stage methods determine valid barriers directly, while two-stage methods involve (i) *candidate barrier identification* and (ii) *candidate barrier verification*. For instance, the CEGIS approach is a two-stage method that iterates between learning candidate barriers and verifying them using a counterexample approach (Edwards et al. 2024). Typically, verifying a candidate barrier is less complex than generating it, rendering two-stage methods generally more computationally efficient. In Section 7, we present an efficient sampling-based algorithm for generating barriers that are valid *by design*, thus avoiding the computational challenges associated with CEGIS. For completeness, we will discuss and compare various verification approaches in the following.

6.2.1 SAT/SMT Solvers. *Satisfiability* (SAT) solvers and *satisfiability modulo theory* (SMT) solvers are useful algorithmic tools for determining if a given set of equations can be satisfied. Thus, SMT solvers such as Z3 (De Moura and Bjørner 2008) and dREAL (Gao, Kong, et al. 2013) have found widespread use in applications requiring formal correctness guarantees. For example, Huang et al. (2017) use Z3 to verify the safety of neural networks. Similarly, the recent barrier tool by Edwards et al. (2024) relies on SMT solvers to verify the soundness of generated candidate barriers. Although the aforementioned solvers can handle (some) nonlinear expressions, their numerical complexity is prohibitive. Whilst the exact complexity of solving an SIP such as (15) via SMT solvers is problem dependent, the general complexity of δ -SMT programs such as the ones solved by dREAL is NP-complete (Gao, Avigad, et al. 2012). The involved δ -decision procedures construct a symbolic representation

⁴The Fenchel conjugate of a function $f: \mathbb{X} \rightarrow \mathbb{R}$ is defined as $f^*(y) = \sup_{x \in \mathbb{X}} (y^\top x - f(x))$ (Boyd and Vandenberghe 2004).

of the problem by replacing nonlinear expressions with symbolic variables. The constraints are subsequently evaluated over the continuous domain by applying interval arithmetic. The resulting number of paths to be checked scales exponentially in the number of symbolic variables; a phenomenon known as *path explosion*.

6.2.2 Global Optimizers. Another approach to finding a feasible solution to the SIP (15) are global optimizers with deterministic guarantees for finding the global optimum (within tolerance) in finite time (Horst and Tuy 2013). Global optimizers themselves can be subcategorized depending on their operating principle. For instance, *branch-and-bound* (B&B) optimizers such as BARON (Sahinidis 1996), SCIP (Bestuzheva et al. 2021), and MAiNGO (Bongartz et al. 2018) solve the problem by iteratively subdividing the domain and performing local convex relaxations to obtain upper and lower bounds on possible function values. Similar to SMT solvers, they decompose the nonlinear program by assigning symbolic variables to nonlinear expressions. For kernel-based methods such as GPs, the number of nonlinear terms scales quadratically in the number of training samples. See Schweidtmann et al. (2021) for a dissemination of global optimization of GPs based on branch-and-bound algorithms. For the specific case of neural barrier verification, interval analysis tools such as IMMRAX (Harapanahalli et al. 2024) can be used. In contrast to SMT solvers, global optimizers can provide information beyond logical feasibility and determine the actual (approximate) value of the global optimum. However, SMT solvers typically leverage combinatorial techniques such as conflict-driven clause learning (CDCL) and are generally not less well equipped to solving SIPs over continuous domains than global optimizers leveraging convex relaxations. Notably, the tool MAiNGO provides dedicated relaxations for common kernel functions. The worst-case complexity of any B&B algorithm is exponential. However, the average complexity of the algorithm can be significantly lower on a case-by-case basis. In our experiments using the SQExp kernel, neither the SMT solver dREAL nor any of the previously mentioned B&B optimizers was able to produce results in a reasonable time frame for a minimal example of the SIP (15) involving all constraints.

6.2.3 Compact Coverings. Most approaches to hard shape constraints for kernel methods use a compact covering of the domain. For instance, the covering-based method outlined by Aubin-Frankowski and Szabo (2020) invokes inequality constraints on function derivatives by sampling equidistant points. Similar to global optimizers, compact coverings leverage regularity information of the optimization problem. In fact, a compact cover abandons the iterative B&B scheme and constitutes its limit case by forming branches according to a globally unified regularity property. Thus, compact coverings do not exclude regions of sure constraint satisfaction (that could be known from a convex relaxation) from the computations. This results in a inferior computational complexity of $O(r^n)$, where r is the resolution of the discretization and n the number of system dimensions (Lasserre 2001). To give an intuition, for the SQExp kernel (4), the resolution r required to obtain an error is highly dependent on the amplitude σ_f^2 of the kernel.

6.2.4 Other Approaches. There are a few noteworthy approaches that do not directly fit into the aforementioned categories or are currently impractical. For example, the non-convex scenario optimization approach recently presented by Garatti and Campi (2024) breaks down the solution of an SIP into the identification of the least complex finite (i.e., having the fewest possible finitely many constraints) subproblem. However, obviously, determining this subproblem is generally difficult. In contrast, B&B identifies a relaxed problem via a strategical iterative procedure.

Next, we present our approach to the SIP (15). Whilst all of the previously mentioned approaches operate on the spatial domain of the SIP, choosing a kernel inducing a Fourier basis allows us to obtain a spectral representation of the problem. More specifically, we choose a finite Fourier feature expansion of the SQExp kernel that is essentially a spectral abstraction of the system. We draw from the mathematical foundation of kernel-based methods to construct correct-by-design barriers based on an explicit mathematical characterization of the space of candidate barriers and the system’s hypothesis space $\mathcal{G}$.

7 Fourier Barrier Solution

To address the SIP in (15), a central observation is made: although modeling nonlinear state evolutions, at its core, a CME corresponds to a *linear* operator in RKHS. However, as characteristic RKHSs are typically infinite-dimensional, performing explicit computations within RKHS is generally not possible. To obtain a tractable approximation, a Fourier series expansion of the characteristic squared exponential kernel is employed. This yields a finite-dimensional spectral approximation of the associated RKHS, which remains expressive enough to capture both the system dynamics and candidate barrier functions.

In contrast to compact covering methods (see Subsubsection 6.2.3), where the number of spatial support vectors scales exponentially with the dimensionality of the system, the spectral density of the squared exponential kernel is concentrated near the origin (see Fig. 3), allowing for a compact and efficient basis requiring comparatively few spectral support vectors. The resulting formulation is a semi-infinite linear program (SILP), which can be solved efficiently as a standard linear program (LP) using trigonometric bounding techniques (Pfister and Bresler 2018). Notably, the ambiguity arises here in the frequency domain, introducing uncertainty over the frequency response of the system. While this paper focuses on establishing the overall framework, quantifying the induced abstraction error and identifying principled strategies for selecting an optimal spectral basis remain open challenges for future work. We provide an overview of the proposed algorithmic verification solution in Fig. 2.




Fig. 2. Sequence of steps to generate a Fourier barrier certificate.

7.1 Finite Fourier Feature Expansion

It is well known that for popular translation-invariant kernels Bochner's theorem provides an infinite series expansion (Rahimi and Recht 2007). In the remainder of this paper, we focus on the SQExp kernel (4), which admits a Fourier expansion of the form

$$k(x, x') \equiv \sigma_f^2 \int_{\mathbb{R}^n} \mathcal{N}(d\omega|0, \Sigma^{-1}) e^{i\omega^\top (P(x) - P(x'))}, \quad \Sigma := \text{diag}(\sigma_l)^2, \quad (19)$$

associated with an infinite-dimensional RKHS, where $x \mapsto P(x)$ is an affine mapping from $\mathbb{X}$ onto the unit hypercube $[0, 1]^n$. To this end, we assume $\mathbb{X}$ to be of hyperrectangular form. Recall that for the CME we have $\mu_{k_t|k_{xu}}(\mathbf{t})(x, u) \in \mathcal{H}_+$ for any $(x, u) \in \mathbb{X} \times \mathbb{U}$. We deliberately restrict ourselves to designing barriers from a finite space $\mathcal{H}_M$, that is spanned by a truncated Fourier basis

$$\phi_M(x) = \sigma_f \begin{bmatrix} \mathbf{w}_0 \\ \sqrt{2}\mathbf{w}_1 \cos(\omega_1^\top P(x)) \\ \sqrt{2}\mathbf{w}_1 \sin(\omega_1^\top P(x)) \\ \vdots \\ \sqrt{2}\mathbf{w}_M \cos(\omega_M^\top P(x)) \\ \sqrt{2}\mathbf{w}_M \sin(\omega_M^\top P(x)) \end{bmatrix},$$

characterized by $M \in \mathbb{N}_{>0}$ wavenumbers (spatial frequencies) $\omega_j \in \mathbb{R}^n$, $j = 1, \dots, M$, and weights $\mathbf{w}_0, \dots, \mathbf{w}_M \in \mathbb{R}_{\geq 0}$.

Frequency Selection. Here, we define equally spaced frequency bands on the hypercube $[-\vartheta, \vartheta] \subset \mathbb{R}^n$, with $\vartheta := 6\sigma_l^{-1}/(2M+1)$, capturing 99.73% of the kernel's spectral measure (see Fig. 3). In particular, we have $\omega_j := \text{diag}(\zeta_j) \cdot \vartheta$, with multi-indices $\zeta_j \in \mathbb{N}_{\geq 0}^n$. Importantly, note that the resulting basis ϕ_M is generally not periodic on $\mathbb{X}$. Moving forward, it will prove useful to define the domain

$$\tilde{\mathbb{X}} := P^{-1}([0, 2\pi] \otimes \vartheta), \quad (20)$$

which is the smallest domain on which ϕ_M is periodic. In particular, we have $\tilde{\mathbb{X}} \equiv \mathbb{X}$ for when $\sigma_{l,i} = 6/(2\pi(2M+1))$, for all $i = 1, \dots, n$, or equivalently $\vartheta = \mathbf{1}_n$, with unit vector $\mathbf{1}_n \in \mathbb{R}^n$. We denote the highest frequency $f_{\max} \in \mathbb{N}_{\geq 0}$ w.r.t. $\tilde{\mathbb{X}}$ as $f_{\max} := \max_{j=1, \dots, M; i=1, \dots, n} \zeta_{ji}$.

REMARK 3 (EFFECT OF THE LENGTHSCALE). Note that as the lengthscale σ_l increases — dilating the spatial extent of the SQExp kernel k in (4) on $\mathbb{X} \times \mathbb{X}$ — the associated Gaussian spectral measure depicted in Fig. 3 contracts. This corresponds to restricting the associated RKHS to progressively smoother functions.

Fourier CBC. As shown in Fig. 3, the corresponding weights of ϕ_M are efficiently computed via the n -dimensional cumulative distribution function (CDF):

$$\mathbf{w}_j^2 := \int_{\omega_j - \vartheta/2}^{\omega_j + \vartheta/2} \mathcal{N}(d\xi|0, \Sigma^{-1}), \quad j = 0, \dots, M.$$

Note that $\omega_0 := 0$. This allows us to expand and write the barrier B as a finite Fourier series with $2M+1$ coefficients:

$$B(x) = \phi_M(x)^\top b = \alpha_0 + \sum_{i=1}^M \alpha_i \cos(\omega_i^\top P(x)) + \beta_i \sin(\omega_i^\top P(x)), \quad (21)$$

$$\text{with } b = \left[\frac{\alpha_0}{\sigma_f^2 \mathbf{w}_0^2} \quad \frac{\alpha_1}{2\sigma_f^2 \mathbf{w}_1^2} \quad \frac{\beta_1}{2\sigma_f^2 \mathbf{w}_1^2} \quad \cdots \quad \frac{\alpha_M}{2\sigma_f^2 \mathbf{w}_M^2} \quad \frac{\beta_M}{2\sigma_f^2 \mathbf{w}_M^2} \right]^\top,$$

where the coefficients $b \in \mathbb{R}^{2M+1}$ act as spectral amplitudes. Naturally, we dub barriers of the spectral form (21) as *Fourier control barrier certificates*. We visualize several components of an example Fourier barrier in Fig. 4.




Fig. 3. Abstraction of a 1-dimensional Gaussian spectral measure of the SQExp kernel for $M = 3$.

Compared to the typical spatial form (16) based on spatial support vectors $\bar{x}_1, \bar{x}_2, \dots \in \mathbb{X}$, the spectral form (21) is a much more efficient representation for higher-dimensional systems. The number of spatial support vectors needed to construct adequately expressive barriers grows exponentially in the dimensionality of the system. We avoid this. In fact, the same principle is leveraged by memory-efficient file formats such as JPEG.

Filtered Dynamics. Apart from the memory efficiency of the spectral barrier form, we can apply the same idea to collapse the complex CME term in (15d) via the *fast Fourier transform* (FFT). That is, we obtain spectral coefficients $H \in \mathbb{R}^{(2M+1) \times (2M+1)}$ such that

$$k_{\hat{X}}^N(x)^\top \left[K_{\hat{X}}^N + N\lambda I_N \right]^{-1} (\Phi_{M, \hat{X}^+}^N)^\top b \approx \phi_M(x)^\top Hb, \quad (22)$$

where $\Phi_{M, \hat{X}^+}^N := [\phi_M(\hat{x}_1^+), \dots, \phi_M(\hat{x}_N^+)]$. In fact, this approximation corresponds to applying a low-pass filter to the dynamics with a cut-off frequency of $f_{\max}$. Whilst projection onto a truncated Fourier basis introduces an approximation error, the error decreases exponentially with the cardinality of wavenumbers in the basis (Rahimi and Recht 2007).

Linear SIP. Substituting the barrier $B \in \mathcal{H}_M$ of the form (21) and the approximation of the CME term in (22) into the SIP (15), we obtain the *semi-infinite linear program* (SILP)

$$\min_{b, c, \eta} \quad \eta + cT, \quad (23a)$$

$$\text{s.t.} \quad \forall x_0 \in \mathbb{X}_0: \phi_M(x_0)^\top b \leq \eta, \quad (23b)$$

$$\forall x_u \in \mathbb{X}_u: \phi_M(x_u)^\top b \geq 1, \quad (23c)$$

$$\forall x \in \mathbb{X}: \phi_M(x)^\top (Hb - b) \leq c - \varepsilon \bar{B} \kappa, \quad (23d)$$

$$\forall x \in \mathbb{X}: \phi_M(x)^\top b \geq 0, \quad (23e)$$

$$c \geq 0, \eta \in [0, 1), b \in \mathbb{R}^{2M+1},$$

with coefficients $\kappa \geq \sigma_f$ and $\bar{B} \geq \|b\|_2$.




Fig. 4. Components of the spectral basis expansion used for the barrier in Section 8.1 with a maximum frequency of $f_{\max} = 2$. Note that Fourier barriers are thus composed of a superposition of standing waves yielding a Fourier series expansion.

7.2 Finite-Constraint Relaxation

There exist results that establish deterministic sampling-based bounds for multivariate trigonometric polynomials (Cifuentes and Parrilo 2017; Pfister and Bresler 2018) that can be applied to the bandlimited spectral barrier in (21). This allows us to relax the SILP (23) to a *linear program* (LP) with finitely many constraints by sampling a finite set $\Theta_{\tilde{N}} \subset \mathbb{X}$ of cardinality $\tilde{N}$. By choosing $\tilde{N} = Q^n$ equidistant points that form a discrete lattice of points to cover the (state) space, the corresponding constraints can be efficiently evaluated in $O(Q \log(Q))$ operations using the inverse FFT. Note that this also substantially reduces the number of evaluations of the computationally expensive CME term in (15d), which requires $O(N^2)$ operations per evaluation. Refer to the Petersen-Middleton theorem (Petersen and Middleton 1962), which generalizes the Nyquist-Shannon sampling theorem to higher-dimensional Euclidean spaces, for further details.

We establish the following lemma based on the results of Pfister and Bresler (2018).

LEMMA 1 (TRIGONOMETRIC BOUNDS). *Consider the Fourier CBC $B: \mathbb{X} \rightarrow \mathbb{R}$ in (21) with maximum degree $f_{\max} \in \mathbb{N}_{\geq 0}$ and define the sampling lattice*

$$\Theta_{\tilde{N}} := \left\{ P^{-1} \left(\frac{2\pi l}{\tilde{Q}} \vartheta \right), l = 0, \dots, \tilde{Q} - 1 \right\} \subset \tilde{\mathbb{X}}, \quad \tilde{N} := \tilde{Q}^n, \quad \tilde{Q} \geq 2f_{\max} + 1,$$

with contraction/dilation coefficient $\vartheta := 3\sigma_l^{-1}/(2M+1) \in \mathbb{R}^n$ and $\tilde{\mathbb{X}}$ in (20). Let $\hat{B}_{\tilde{N}} := \max_{x \in \Theta_{\tilde{N}}} B(x)$ and $\check{B}_{\tilde{N}} := \min_{x \in \Theta_{\tilde{N}}} B(x)$. Then, we have for all $x \in \mathbb{X}$,

$$\frac{1}{2} \left(\hat{B}_{\tilde{N}} + \check{B}_{\tilde{N}} - C_{\tilde{N}} \left(\hat{B}_{\tilde{N}} - \check{B}_{\tilde{N}} \right) \right) \leq B(x) \leq \frac{1}{2} \left(\hat{B}_{\tilde{N}} + \check{B}_{\tilde{N}} + C_{\tilde{N}} \left(\hat{B}_{\tilde{N}} - \check{B}_{\tilde{N}} \right) \right),$$

with a constraint-tightening coefficient given by

$$C_{\tilde{N}} := \left(1 - \frac{2f_{\max}}{\tilde{Q}} \right)^{-\frac{n}{2}}. \quad (24)$$

PROOF. Note that the barrier in (21) is a real-valued trigonometric polynomial with maximum degree $f_{\max}$ that is periodic on the domain $\tilde{\mathbb{X}}$ defined in (20). According to Corollaries 1 & 2 by Pfister and Bresler (2018), if $\tilde{Q} \geq 2f_{\max} + 1$ — i.e., the minimum/Nyquist sampling rate is met — then we have for all $x \in \tilde{\mathbb{X}}$ that

$$\frac{1}{2} \left(\hat{B}_{\tilde{N}} + \check{B}_{\tilde{N}} - C_{\tilde{N}} \left(\hat{B}_{\tilde{N}} - \check{B}_{\tilde{N}} \right) \right) \leq B(x) \leq \frac{1}{2} \left(\hat{B}_{\tilde{N}} + \check{B}_{\tilde{N}} + C_{\tilde{N}} \left(\hat{B}_{\tilde{N}} - \check{B}_{\tilde{N}} \right) \right).$$

As B is periodic on $\tilde{\mathbb{X}}$, the same bounds hold for any arbitrary domain $\mathbb{X}$. $\square$

Via Lemma 1, we can obtain a finite-constraint relaxation of upper/lower bounds on the barrier; however, this holds only for constraints enforced on the entire periodic domain $\tilde{\mathbb{X}}$. In fact, not only must (23b) and (23c) be selectively enforced only on $\mathbb{X}_0$ and $\mathbb{X}_u$, respectively, but the state space $\mathbb{X}$ may differ from $\tilde{\mathbb{X}}$ too, affecting the relaxation of (23d) and (23e).

Local Relaxation. Without loss of generality, let $\mathbb{S} \in \{\mathbb{X}_0, \mathbb{X}_u, \mathbb{X}\}$. According to Pfister and Bresler (2018, Lemma 2) we have for $\tilde{Q} \geq 2f_{\max} + 1$ that for all $x \in \tilde{\mathbb{X}}$

$$B(x) = \frac{1}{\tilde{N}} \sum_{\bar{x} \in \Theta_{\tilde{N}}} B(\bar{x}) D_{f_{\max}, \tilde{Q}-f_{\max}}^n(x - \bar{x}),$$

with $D_{a,b}^n : \mathbb{R}^n \rightarrow \mathbb{R}$ the Vallée-Poussin kernel, defined as

$$D_{a,b}^n(z) := \frac{1}{(b-a)^n} \prod_{i=1}^n \frac{\sin(\frac{b+a}{2}z_i) \sin(\frac{b-a}{2}z_i)}{\sin^2(\frac{z_i}{2})}.$$

For the set $\mathbb{S}$, we have for all $x \in \tilde{\mathbb{X}}$

$$B(x) \leq \underbrace{\hat{B}_{\tilde{N}}^{\mathbb{S}} \frac{1}{\tilde{N}} \sum_{\bar{x} \in \Theta_{\tilde{N}}} D_{f_{\max}, \tilde{Q}-f_{\max}}^n(x - \bar{x})}_{(a)} + \underbrace{\left(\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \hat{B}_{\tilde{N}}^{\mathbb{S}} \right) \frac{1}{\tilde{N}} \sum_{\bar{x} \in \Theta_{\tilde{N}} \setminus \mathbb{S}} D_{f_{\max}, \tilde{Q}-f_{\max}}^n(x - \bar{x})}_{(b)}, \quad (25)$$

where $\hat{B}_{\tilde{N}}^{\mathbb{S}} := \max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} \{\phi_M(x)^\top b\}$ and $\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} := \max_{x \in \Theta_{\tilde{N}} \setminus \mathbb{S}} \{\phi_M(x)^\top b\}$ are computed on the lattice. We can bound (a) as in [Pfister and Bresler \(2018, Lemma 4\)](#):

$$\frac{1}{\tilde{N}} \sum_{\bar{x} \in \Theta_{\tilde{N}}} D_{f_{\max}, \tilde{Q}-f_{\max}}^n(x - \bar{x}) \leq C_{\tilde{N}},$$

where $C_{\tilde{N}}$ is the constraint-tightening coefficient from [Lemma 1](#). Term (b) can be computed numerically using, e.g., particle swarm optimization; we write

$$A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} := \frac{1}{\tilde{N}} \sum_{\bar{x} \in \Theta_{\tilde{N}} \setminus \mathbb{S}} D_{f_{\max}, \tilde{Q}-f_{\max}}^n(x - \bar{x}). \quad (26)$$

Based on this decomposition, we establish upper and lower bounds akin [Lemma 1](#) for subsets $\mathbb{S} \subset \tilde{\mathbb{X}}$. The following result is based on [Pfister and Bresler \(2018, Cors. 1 & 2\)](#), and the full proof can be found in [App. A](#).

LEMMA 2 (LOCAL TRIGONOMETRIC BOUNDS). *Consider the Fourier CBC $B: \mathbb{X} \rightarrow \mathbb{R}$ in (21) with maximum degree $f_{\max} \in \mathbb{N}_{\geq 0}$ and the sampling lattice $\Theta_{\tilde{N}}$ in [Lemma 1](#) on $\tilde{\mathbb{X}}$. For a subset $\mathbb{S} \subset \tilde{\mathbb{X}}$, let $\hat{B}_{\tilde{N}}^{\mathbb{S}} := \max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} B(x)$, $\check{B}_{\tilde{N}}^{\mathbb{S}} := \min_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} B(x)$, $\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} := \max_{x \in \Theta_{\tilde{N}} \setminus \mathbb{S}} B(x)$, and $\check{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} := \min_{x \in \Theta_{\tilde{N}} \setminus \mathbb{S}} B(x)$. Then, we have for all $x \in \mathbb{S}$,*

$$\frac{1}{2} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}} - C_{\tilde{N}} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right) \right) + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\check{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right) \leq B(x) \leq \frac{1}{2} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}} + C_{\tilde{N}} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right) \right) + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \hat{B}_{\tilde{N}}^{\mathbb{S}} \right),$$

with the constraint-tightening coefficients $C_{\tilde{N}}$ and $A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}}$ defined in (24) and (26), respectively.

Via [Lemma 2](#), we obtain a finite-constraint relaxation of the SILP (23) based on the following constraint-tightening reasoning ($\mathbb{S} \in \{\mathbb{X}_0, \mathbb{X}_u, \mathbb{X}\}$):

$$\frac{1}{2} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}} - C_{\tilde{N}} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right) \right) + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\check{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right) \geq \text{lb} \implies \forall x \in \mathbb{S}: B(x) \geq \text{lb},$$

for some desired lower bound $\text{lb} \in \mathbb{R}$. Thus, we obtain strengthened constraints of the form

$$\check{B}_{\tilde{N}}^{\tilde{\mathbb{X}}} \geq \frac{2\text{lb} + (C_{\tilde{N}} - 1)\hat{B}_{\tilde{N}}^{\mathbb{S}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}}\check{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} + 1}.$$

For an upper bound $\forall x \in \mathbb{S}: B(x) \leq \text{ub}$, we obtain analogous constraints of the form

$$\hat{B}_{\tilde{N}}^{\mathbb{S}} \leq \frac{2\text{ub} + (C_{\tilde{N}} - 1)\check{B}_{\tilde{N}}^{\mathbb{S}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}}\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} + 1}.$$

REMARK 4. Note that the separation of $\mathbb{S}$ and $\tilde{\mathbb{X}} \setminus \mathbb{S}$ in (25) is arbitrary. Instead, we can also choose, e.g., a superset of $\mathbb{S}$ for (a), shrinking (b) and yielding a potentially less conservative bound, as most of the functional mass of the Vallée-Poussin kernel is assigned locally around $\bar{x}$. In practice, we compute the terms for a superset of $\mathbb{S}$ inflated by, e.g., 2%, yielding an only marginal contribution of (25)(b) to the local bound on the barrier B .

7.3 Linear Program

We are ready to present the LP relaxation of the SILP (23). Additionally to $\Theta_{\tilde{N}}$ and $\Theta_{\tilde{N}}$ we form discrete sets $\{x_0^{(1)}, \dots, x_0^{(\hat{N}_0)}\} \subset \mathbb{X}_0$ and $\{x_u^{(1)}, \dots, x_u^{(\hat{N}_u)}\} \subset \mathbb{X}_u$ of cardinality $\hat{N}_0, \hat{N}_u \in \mathbb{N}$. For given $\bar{B}, C_{\tilde{N}}, A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}, A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_u}$, and $A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}$, we obtain the following LP:

$$\begin{aligned}
 & \min_{b, c, \eta} \quad \eta + cT, \\
 & \text{subject to} \quad \begin{aligned}
 & \hat{B}_{\tilde{N}}^{\mathbb{X}_0}, \hat{B}_{\tilde{N}}^{\mathbb{X}_u}, \hat{B}_{\tilde{N}}^{\mathbb{X}}, \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}, \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_u}, \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}, \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} \\
 & \hat{B}_{\tilde{N}}^{\mathbb{X}_0} \leq \phi_M(x_0^{(i)})^\top b \leq \hat{\eta}, \quad i = 1, \dots, \hat{N}_0, \\
 & \hat{\gamma} \leq \phi_M(x_u^{(i)})^\top b \leq \hat{B}_{\tilde{N}}^{\mathbb{X}_u}, \quad i = 1, \dots, \hat{N}_u, \\
 & \hat{B}_{\tilde{N}}^{\mathbb{X}} \leq \phi_M(x^{(i)})^\top (Hb - b) \leq \hat{\Delta}, \quad i = 1, \dots, \hat{N}, \\
 & \hat{\xi} \leq \phi_M(x^{(i)})^\top b \leq \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}}}, \quad i = 1, \dots, \hat{N}, \\
 & \phi_M(x^{(i)})^\top b \leq \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}, \quad i = 1, \dots, \tilde{N} - \hat{N}_0, \\
 & \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_u} \leq \phi_M(x^{(i)})^\top b, \quad i = 1, \dots, \tilde{N} - \hat{N}_u, \\
 & \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} \leq \phi_M(x^{(i)})^\top b, \quad i = 1, \dots, \tilde{N} - \hat{N}, \\
 & \phi_M(x^{(i)})^\top (Hb - b) \leq \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}, \quad i = 1, \dots, \tilde{N} - \hat{N}, \\
 & c \geq 0, \eta \in [0, 1), b \in \mathbb{R}^{2M+1},
 \end{aligned} \tag{27}
 \end{aligned}$$

with $\kappa \geq \sigma_f$, $\bar{B} \geq \|b\|_2$, and constraint-tightening coefficients

$$\begin{aligned}
 \hat{\eta} &:= \frac{2\eta + (C_{\tilde{N}} - 1)\hat{B}_{\tilde{N}}^{\mathbb{X}_0} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} + 1}, & \hat{\gamma} &:= \frac{2 + (C_{\tilde{N}} - 1)\hat{B}_{\tilde{N}}^{\mathbb{X}_u} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_u} \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_u}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_u} + 1}, \\
 \hat{\Delta} &:= \frac{2(c - \varepsilon \bar{B} \kappa) + (C_{\tilde{N}} - 1)\hat{B}_{\tilde{N}}^{\mathbb{X}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} + 1}, & \hat{\xi} &:= \frac{(C_{\tilde{N}} - 1)\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} + 1}.
 \end{aligned}$$

Any solution to the LP (27) that satisfies $\|b\|_2 \leq \bar{B}$ is a feasible solution to the SIP (15). We provide more details on the derivation of the LP (27) from the SILP (23) via the constraint-tightening procedure outlined in Subsection 7.2 in Appendix B.

REMARK 5 (COMPUTATIONAL COMPLEXITY). The LP in (27) scales most critically with the periodic lattice size $\tilde{N}$ and the chosen Fourier barrier basis size M (see Fig. 6b for an ablation study). The computational complexity of LPs differs between solution algorithms: While the simplex algorithm admits an exponential worst-case complexity, its practical performance is typically much better (Spielman and Teng 2004). Interior-point methods, on the other hand, have polynomial worst-case complexity in the problem size.

The matrix H in (22) is computed via an FFT-based decomposition. The FFT has complexity $O(Q \log Q)$, where Q denotes the lattice resolution, which depends on the maximum frequency $f_{\max}$ considered. In our setting, $Q \geq (2f_{\max} + 1)^2$ by the Nyquist–Shannon sampling theorem. Since H is computed only once during LP construction, this step contributes only a minor fraction of the overall runtime compared to solving the LP, which dominates the overall runtime in our experiments.

8 Experimental Results

We demonstrate the proposed Fourier barrier approach on two benchmarks. To this end, we developed the toolbox LUCID (Casablanca et al. 2026), which implements the framework and theory introduced in this paper and is used to carry out all reported experiments. For all examples, we draw $N = 1,000$ samples from the unknown system, which we found to provide a good trade-off between computational cost and empirical accuracy. The inverse of the Gram matrix is computed with a regularization $\lambda = 10^{-5}$, which we found to provide numerically stable results across the benchmarks. To construct the finite-constraint LP (27), we generate a discrete lattice of $\tilde{N}$ points on $\mathbb{X}$ (as reported in the individual benchmarks). For simplicity, we use a rectangular sampling lattice, leaving the exploration of more complex (optimal) sampling lattices for future work. The numerical results reported are obtained on a machine running an AMD Ryzen 9 5950X 16-core CPU and 64 GB of memory. LUCID uses a serial implementation to generate the LP. For solving the LP, we use GUROBI (Gurobi Optimization, LLC 2024), which is allowed to execute solution methods concurrently by default. We set its feasibility tolerance to the minimum of 10^{-8} , keeping the remaining settings at default.

8.1 Complex Safety Specification

The first benchmark is inspired by the example Barr_3 from the deterministic benchmarks by Abate, Ahmed, et al. (2021), which we extend by adding stochastic noise $w_t \sim \mathcal{N}(\cdot | 0, 0.01I_2)$ to arrive at the two-dimensional nonlinear stochastic dynamics

$$\begin{bmatrix} x_{1,t+1} \\ x_{2,t+1} \end{bmatrix} = \begin{bmatrix} x_{1,t} \\ x_{2,t} \end{bmatrix} + \tau \begin{bmatrix} x_{2,t} \\ \frac{1}{3}x_{1,t}^3 - x_{1,t} - x_{2,t} \end{bmatrix} + w_t,$$

with $\tau := 0.1$. We illustrate the setup in Fig. 5. The goal is to compute the probability that the system initialized in $[x_{1,0}, x_{2,0}]^\top \in \mathbb{X}_0$ (blue regions) does not enter the unsafe regions $\mathbb{X}_u$ (in red) within $T = 5$ time steps. Note that the specified regions $\mathbb{X}_0$ and $\mathbb{X}_u$ are in close proximity to each other, rendering the barrier synthesis particularly challenging.

Setup. We generate the LP based on a fixed lattice $\Theta_{\tilde{N}}$ with a resolution of $\tilde{Q} = 40 \cdot (2f_{\max} + 1)$ (i.e., 40-times Nyquist sampling rate), leading to a periodic lattice of size $\tilde{N} = \tilde{Q}^2$. To demonstrate the expressiveness of different spectral bases, we synthesize barriers for (a) $M = 6^2 - 1$ and (b) $M = 10^2 - 1$ wavenumbers (with $f_{\max} = 6$ and 15, respectively), resulting in barriers with $2M + 1 = 71$ and 199 coefficients $b_0, b_1, \dots \in \mathbb{R}$, respectively. We fix the kernel hyperparameters to $\sigma_f = 1$ and $\sigma_l = [2.993, 4.629]$ for the input kernel k_x , and obtain $\sigma_l = [0.143, 0.358]$ (resp. $\sigma_l = [0.326, 0.327]$) for the output kernel k_+ via hyperparameter optimization. For the computation of the constraint-tightening coefficients, we inflate the sets by 2% (cf. Remark 4). We further use an upper bound $\bar{B} = 7$ and $\varepsilon = 0.001$ for both experiments.

Results. We solve the corresponding LPs and obtain the barriers shown in Fig. 5. For (a), we find $\eta = 0.277$ and $c = 0.072$, meaning that the identified barrier certifies a lower bound on the safety probability of the system of at least 36.3% (via the formula in Theorem 3). The barrier basis (b) is richer, enabling us to find $\eta = 0.263$ and $c = 0.047$, leading to a robust lower bound on the safety probability of the system of at least 50.0%. For comparison, we estimate the true safety probability P_{safe} using Monte Carlo simulation. We run 10,000 simulations starting from a fixed initial state $[x_{1,0}, x_{2,0}] = [1, 0.5]$ with a maximum length of T time steps. Using Chebychev’s inequality




Fig. 5. Robustified CBCs and level sets at $B(x) = 1$ (in red) and η (in blue) computed for the benchmark in Section 8.1, robustified with $\varepsilon \bar{B} \sigma_f = 0.007$ and with (a) $M = 35$ wavenumbers or (b) $M = 99$ wavenumbers.

with a confidence level of 90% we obtain an estimated satisfaction probability of 85 – 88%. Note that even for an infinitely complex barrier (i.e., $M \rightarrow \infty$) and data $N \rightarrow \infty$, this latent probability might not be attainable due to the inherent conservatism of the general barrier approach. This is in contrast to abstraction-based approaches, for which the lower bound converges to the true solution as the partitioning of the state space is refined and more data is accumulated (Laurenti and Lahijanian 2023). The effect is particularly strong here due to the closeness of the safe and unsafe regions (see Fig. 5). See Fig. 6 for an ablation study on the influence of M and $\tilde{N}$ on the certified safety probability p_N and runtime ($\varepsilon = 0$).

8.2 Overtaking Scenario Featuring a Neural Network Controller

Now, we move to a system controlled by a NN controller. For this, consider the overtaking scenario in Example 1, visualized in Figures 1 and 7. The NN controller $\pi: \mathbb{X} \rightarrow \mathbb{U}$ was trained using a *twin-delayed deep deterministic* (TD3) policy gradient RL algorithm and features two layers of 200 neurons each and issues a steering wheel angle $u_t \in [-\pi, \pi]$. We set a fixed velocity $v = 1$. The goal is to certify that the controlled system remains safe by computing the probability that the system initialized in $[x_0, y_0, \phi_0]^\top \in \mathbb{X}_0$ does not enter the unsafe region $\mathbb{X}_u$ within $T = 5$ time steps.

Setup. The LP is generated based on a lattice $\Theta_{\tilde{N}}$ of size $\tilde{N} = 70^3$. We select a spectral basis of $M = 5^3 - 1$ wavenumbers, resulting in a barrier characterized by $2M + 1 = 249$ coefficients $b_0, b_1, \dots \in \mathbb{R}$. The kernel hyperparameters are selected to be $\sigma_f = 7$ and $\sigma_l = [0.054, 0.094, 5.078]$ for the input kernel k_x and $\sigma_l = [0.525, 0.05, 0.525]$ for the output kernel k_+ . For the computation of the constraint-tightening coefficients, we inflate the sets by 3% (cf. Remark 4). Throughout all experiments, we set $\bar{B} = 6.1$.

Results. For $\varepsilon = 0$, $c \approx 0$ (within GUROBI's tolerance), and $\eta = 0.463$, we obtain the barrier shown in Fig. 8a, which certifies safety with a probability of at least 53.7%. Note that this is only certifying safety w.r.t. the empirical distribution, i.e., the CME constructed from the observed data. In order to robustify the result to out-of-distribution




Fig. 6. Ablation study on the influence of M and $\tilde{N}$ for the benchmark in Section 8.1.




Fig. 7. Control performance of the NN controller for the overtaking scenario (Section 8.2), demonstrated on 10 random trajectories initialized in $\mathbb{X}_0$ (in blue), trained to avoid the unsafe regions $\mathbb{X}_u$ (in red) and reach a target set (in black).

behavior, we increase the robustness coefficient ε and report its influence on the lower bound on the safety probability, p_N^π , in Fig. 9. Intuitively, increasing ε tightens the Kushner constraint (15d) by enforcing the barrier to decay by at least $\varepsilon \bar{B} \sigma_f$ in every time step. Note that this is demanding the barrier to be more complex, as




Fig. 8. CBCs and level sets at $B(x) = 1$ (in red) and η (in blue) computed for the overtaking scenario (Section 8.2) shown for a fixed $\phi = 0$, (a) based only on the empirical data, and (b) robustified with $\varepsilon = 0.021$ and $\bar{B} = 6.1$. We indicate the intersections of the barrier and the level sets as dashed lines.

observed in an exponentially increasing complexity of the barrier, measured through $\|b\|$. Simultaneously, the distance between the level sets $B(x) = 1$ and η is shrinking, prompting the robust safety probability p_N^π to decline. For the highest reported $\varepsilon = 0.021$, the required decay reaches $\varepsilon \bar{B} \sigma_f = 0.897$, satisfied by the barrier shown in Fig. 8b, for which the robust safety probability is 5.07%.

9 Concluding Remarks

This work contributes to the verification and synthesis of stochastic systems without explicit model knowledge by leveraging the theory of conditional mean embedding and introducing a data-driven approach to control barrier certificates (CBCs). Through the reformulation of probabilistic CBC constraints into a distributionally robust optimization problem, we showed how to obtain the CBC characterization using data with probabilistic correctness guarantees. We showed how the approach can be generalized to temporal logic specifications beyond safety using an automata representation of the specification and the concept of Streett supermartingales. For the data-driven computation of safety barriers, we introduced a finite Fourier expansion to cast the optimization as a linear program.

Whilst this paper provides an algorithmic solution for safety *verification*, the resulting linear formulation does not generally extend to black-box control synthesis. In the synthesis setting, the control input enters the constraint nonlinearly through the input kernel, preventing the problem from being recast as a linear program as done here for verification. Only in special cases — such as systems linear in the control input (Santoyo et al. 2021) — a more tractable structure arises. In the general case, the synthesis problem may need to be addressed using gradient-based solvers or custom nonlinear global optimizers with deterministic convergence guarantees. Exploring this extension remains an important avenue for future work.

Further investigation is needed to reduce the conservatism of the approach and determine sharp ambiguity set radii ε with high confidence values $1 - \rho$, including the use of alternative concentration theorems tailored




Fig. 9. Robust safety probability p_N^π (in blue) and required barrier complexity $\|b\|$ (in black) for the overtaking scenario (Section 8.2) generated for different robustness radii ϵ .

to the class of barrier functions (e.g., via Rademacher complexity) and alternative algorithmic solutions to the Fourier barrier program (23). Albeit the presented spectral barrier design lifts the need for an exponential number of spatial support vectors, the number of samples required by the trigonometric sampling bounds remains exponential in the system dimensionality; the FFT only mitigates this effect. It will therefore be interesting to combine the Fourier barrier formulation with nonlinear global optimizers with deterministic convergence guarantees, potentially requiring the development of a custom semi-infinite solver to exploit the problem structure fully. Moreover, techniques such as kernel herding (see, e.g., [Chen, Welling, et al. 2010](#)) could be explored to alleviate the cubic complexity of the CME, although this lies outside the scope of the present work. Additionally, further exploration of hyperparameter tuning and an optimal selection of the spectral basis are needed.

Acknowledgments

The work of Sadeh Soudjani is supported by the EIC SymAware project 101070802 and the ERC Auto-CyPheR project 101089047. Zhengang Zhong is grateful for the support of the Leverhulme Trust through the Project Award “Robust Learning: Uncertainty Quantification, Sensitivity and Stability” (grant agreement RPG-2024-051).

The authors wish to thank Dominik Bongartz for his feedback on maximizing the potential of the MAiNGO solver. We also thank Arthur Gretton for helpful discussions and directing us to valuable references on concentration bounds for CMEs. We thank Ernesto Casablanca for his efforts in the development of the LUCID toolbox.

References

- A. Abate, D. Ahmed, A. Edwards, M. Giacobbe, and A. Peruffo. 2021. “FOSSIL: A Software Tool for the Formal Synthesis of Lyapunov Functions and Barrier Certificates using Neural Networks.” In: *Proceedings of the 24th International Conference on Hybrid Systems: Computation and Control* Article 24. ACM, 1–11. doi:[10.1145/3447928.3456646](https://doi.org/10.1145/3447928.3456646).
- A. Abate, S. Bogomolov, A. Edwards, K. Potomkin, S. Soudjani, and P. Zuliani. 2024. “Safe Reach Set Computation via Neural Barrier Certificates.” *IFAC-PapersOnLine*, 58, 11, 107–114. 8th IFAC Conference on Analysis and Design of Hybrid Systems. doi:[10.1016/j.ifacol.2024.07.433](https://doi.org/10.1016/j.ifacol.2024.07.433).

- A. Abate, M. Giacobbe, and D. Roy. 2024. “Stochastic Omega-Regular Verification and Control with Supermartingales.” In: *Computer Aided Verification*. Ed. by A. Gurfinkel and V. Ganesh. Springer, 395–419. doi:[10.1007/978-3-031-65633-0_18](https://doi.org/10.1007/978-3-031-65633-0_18).
- M. Anand, A. Lavaei, and M. Zamani. 2024. “Compositional Synthesis of Control Barrier Certificates for Networks of Stochastic Systems Against ω -Regular Specifications.” *Nonlinear Analysis: Hybrid Systems*, 51, 101427. doi:[10.1016/j.nahs.2023.101427](https://doi.org/10.1016/j.nahs.2023.101427).
- A. Arai. 2018. *Analysis on Fock Spaces and Mathematical Theory of Quantum Fields: An Introduction to Mathematical Analysis of Quantum Fields*. World Scientific.
- P.-C. Aubin-Frankowski and Z. Szabo. 2020. “Hard Shape-Constrained Kernel Machines.” In: *Advances in Neural Information Processing Systems 2020*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33, 384–395.
- J. A. Bagnell and A.-m. Farahmand. 2015. “Learning Positive Functions in a Hilbert Space.” In: *NIPS Workshop on Optimization (OPT2015)*. Vol. 20, 3240–3255.
- C. Baier and J.-P. Katoen. 2008. *Principles of Model Checking*. MIT Press. ISBN: 9780262026499.
- A. Banse, L. Romao, A. Abate, and R. M. Jungers. 2023. “Data-Driven Abstractions via Adaptive Refinements and a Kantorovich Metric.” In: *2023 62nd IEEE Conference on Decision and Control*. IEEE, 6038–6043. doi:[10.1109/CDC49753.2023.10383513](https://doi.org/10.1109/CDC49753.2023.10383513).
- C. Belta, B. Yordanov, and E. A. Gol. 2017. *Formal Methods for Discrete-Time Dynamical Systems*. Vol. 15. Springer.
- A. Berlinet and C. Thomas-Agnan. 2004. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. eng. Ed. by C. Thomas-Agnan. Kluwer Academic, Boston. ISBN: 9781461347927. doi:[10.1007/978-1-4419-9096-9](https://doi.org/10.1007/978-1-4419-9096-9).
- K. Bestuzheva et al.. 2021. *The SCIP Optimization Suite 8.0*. (2021). arXiv: [2112.08872](https://arxiv.org/abs/2112.08872).
- V. I. Bogachev. 2007. *Measure Theory*. Springer Science & Business Media.
- D. Bongartz, J. Najman, S. Sass, and A. Mitsos. 2018. “MAiNGO: McCormick Based Algorithm for Mixed Integer Nonlinear Global Optimization.” In: *Technical Report*. Process Systems Engineering (AVT. SVT), RWTH Aachen University.
- S. Boyd and L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.
- E. Casablanca, O. Schön, P. Zuliani, and S. Soudjani. 2026. “LUCID: Learning-Enabled Uncertainty-Aware Certification of Stochastic Dynamical Systems.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 24. Vol. 40, 19916–19924. doi:[10.1609/aaai.v40i24.39075](https://doi.org/10.1609/aaai.v40i24.39075).
- J. A. Chekan and C. Langbort. 2023. “Safety-Aware Learning-Based Control of Systems with Uncertainty Dependent Constraints.” In: *2023 American Control Conference*, 1264–1270. doi:[10.23919/ACC55779.2023.10156490](https://doi.org/10.23919/ACC55779.2023.10156490).
- Y. Chen, Y. Li, S. Li, and X. Yin. 2025. *Distributionally Robust Control Synthesis for Stochastic Systems with Safety and Reach-Avoid Specifications*. (2025). arXiv: [2501.03137](https://arxiv.org/abs/2501.03137).
- Y. Chen, M. Welling, and A. Smola. 2010. “Super-Samples from Kernel Herding.” In: *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 109–116.
- D. Cifuentes and P. A. Parrilo. 2017. “Sampling Algebraic Varieties for Sum of Squares Programs.” *SIAM Journal on Optimization*, 27, 4, 2381–2404. doi:[10.1137/15M1052548](https://doi.org/10.1137/15M1052548).
- C. Coey, L. Kapelevich, and J. P. Vielma. 2022. “Solving Natural Conic Formulations with Hypatia.jl.” *INFORMS Journal on Computing*, 34, 5, 2686–2699. doi:[10.1287/ijoc.2022.1202](https://doi.org/10.1287/ijoc.2022.1202).
- M. H. Cohen, C. Belta, and R. Tron. Aug. 2022. “Robust Control Barrier Functions for Nonlinear Control Systems with Uncertainty: A Duality-based Approach.” In: *Proceedings of the IEEE Conference on Decision and Control*. (Aug. 2022), 174–179. doi:[10.1109/CDC51059.2022.9992667](https://doi.org/10.1109/CDC51059.2022.9992667).
- R. K. Cosner, I. Sadalski, J. K. Woo, P. Culbertson, and A. D. Ames. 2024. “Generative Modeling of Residuals for Real-Time Risk-Sensitive Safety with Discrete-Time Control Barrier Functions.” In: *2024 IEEE International Conference on Robotics and Automation*. IEEE, i–viii. doi:[10.1109/ICRA57147.2024.10611355](https://doi.org/10.1109/ICRA57147.2024.10611355).
- T. Cunis and B. Legat. 2022. *Sequential Sum-of-Squares Programming for Analysis of Nonlinear Systems*. (2022). arXiv: [2210.02142](https://arxiv.org/abs/2210.02142).
- L. De Moura and N. Bjørner. 2008. “Z3: An Efficient SMT Solver.” In: *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340. doi:[10.1007/978-3-540-78800-3_24](https://doi.org/10.1007/978-3-540-78800-3_24).
- H. Djelassi, A. Mitsos, and O. Stein. 2021. “Recent Advances in Nonconvex Semi-Infinite Programming: Applications and Algorithms.” *EURO Journal on Computational Optimization*, 9, 100006. doi:[10.1016/j.ejco.2021.100006](https://doi.org/10.1016/j.ejco.2021.100006).
- A. Edwards, A. Peruffo, and A. Abate. 2024. “Fossil 2.0: Formal Certificate Synthesis for the Verification and Control of Dynamical Models.” In: *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control* Article 26. ACM, 10 pages. doi:[10.1145/3641513.3651398](https://doi.org/10.1145/3641513.3651398).
- S. Gao, J. Avigad, and E. M. Clarke. 2012. “ δ -Complete Decision Procedures for Satisfiability Over the Reals.” In: *International Joint Conference on Automated Reasoning*. Springer, 286–300. doi:[10.1007/978-3-642-31365-3_23](https://doi.org/10.1007/978-3-642-31365-3_23).
- S. Gao, S. Kong, and E. M. Clarke. 2013. “dReal: An SMT Solver for Nonlinear Theories Over the Reals.” In: *International Conference on Automated Deduction*. Ed. by M. P. Bonacina. Springer, 208–214. doi:[10.1007/978-3-642-38574-2_14](https://doi.org/10.1007/978-3-642-38574-2_14).
- S. Garatti and M. C. Campi. 2024. “Non-Convex Scenario Optimization.” *Mathematical Programming*, 209, 557–608. doi:[10.1007/s10107-024-02074-3](https://doi.org/10.1007/s10107-024-02074-3).
- I. Gracia, D. Boskos, L. Laurenti, and M. Mazo Jr. 2023. “Distributionally Robust Strategy Synthesis for Switched Stochastic Systems.” In: *Proceedings of the 26th ACM International Conference on Hybrid Systems: Computation and Control*, 1–10. doi:[10.1145/3575870.3587127](https://doi.org/10.1145/3575870.3587127).

- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. 2012. “A Kernel Two-Sample Test.” *Journal of Machine Learning Research*, 13, 25, 723–773.
- Gurobi Optimization, LLC. 2024. *Gurobi Optimizer Reference Manual*. Retrieved from www.gurobi.com. (2024). <https://www.gurobi.com>.
- A. Harapanahalli, S. Jafarpour, and S. Coogan. 2024. “immrax: A Parallelizable and Differentiable Toolbox for Interval Analysis and Mixed Monotone Reachability in JAX.” *IFAC-PapersOnLine*, 58, 11, 75–80. doi:[10.1016/j.ifacol.2024.07.428](https://doi.org/10.1016/j.ifacol.2024.07.428).
- P. Henriksen and A. Lomuscio. 2021. “DEEPSPLIT: An Efficient Splitting Method for Neural Network Verification via Indirect Effect Analysis.” In: *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2549–2555.
- R. Horst and H. Tuy. 2013. *Global Optimization: Deterministic Approaches*. Springer Science & Business Media.
- X. Huang, M. Kwiatkowska, S. Wang, and M. Wu. 2017. “Safety Verification of Deep Neural Networks.” In: *Computer Aided Verification: 29th International Conference*. Springer, 3–29. doi:[10.1007/978-3-319-63387-9_1](https://doi.org/10.1007/978-3-319-63387-9_1).
- P. Jagtap, G. J. Pappas, and M. Zamani. Oct. 2020. “Control Barrier Functions for Unknown Nonlinear Systems using Gaussian Processes.” In: *Proceedings of the IEEE Conference on Decision and Control*. (Oct. 2020), 3699–3704. doi:[10.1109/CDC42340.2020.9303847](https://doi.org/10.1109/CDC42340.2020.9303847).
- P. Jagtap, S. Soudjani, and M. Zamani. 2021. “Formal Synthesis of Stochastic Systems via Control Barrier Certificates.” *IEEE Transactions on Automatic Control*, 66, 7, 3097–3110. doi:[10.1109/TAC.2020.3013916](https://doi.org/10.1109/TAC.2020.3013916).
- N. Kalra and S. M. Paddock. 2016. “Driving to Safety: How Many Miles of Driving Would it Take to Demonstrate Autonomous Vehicle Reliability?” *Transportation Research Part A: Policy and Practice*, 94, 182–193. doi:[10.1016/j.tra.2016.09.010](https://doi.org/10.1016/j.tra.2016.09.010).
- M. Kanagawa, P. Hennig, D. Sejdinovic, and B. K. Sriperumbudur. 2018. *Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences*. (2018). arXiv: [1807.02582](https://arxiv.org/abs/1807.02582).
- G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. 2017. “Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks.” In: *Computer Aided Verification: 29th International Conference*. Springer, 97–117. doi:[10.1007/978-3-319-63387-9_5](https://doi.org/10.1007/978-3-319-63387-9_5).
- M. Kazemi, R. Majumdar, M. Salamati, S. Soudjani, and B. Wooding. 2024. “Data-Driven Abstraction-Based Control Synthesis.” *Nonlinear Analysis: Hybrid Systems*, 52, 101467. doi:[10.1016/j.nahs.2024.101467](https://doi.org/10.1016/j.nahs.2024.101467).
- M. Kazemi, M. Perez, F. Somenzi, S. Soudjani, A. Trivedi, and A. Velasquez. 2024. “Assume-Guarantee Reinforcement Learning.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 19. Vol. 38, 21223–21231. doi:[10.1609/aaai.v38i19.30116](https://doi.org/10.1609/aaai.v38i19.30116).
- M. Kazemi and S. Soudjani. 2020. “Formal Policy Synthesis for Continuous-State Systems via Reinforcement Learning.” In: *Integrated Formal Methods: 16th International Conference*. Springer, 3–21. doi:[10.1007/978-3-030-63461-2_1](https://doi.org/10.1007/978-3-030-63461-2_1).
- I. Klebanov, I. Schuster, and T. J. Sullivan. 2020. “A Rigorous Theory of Conditional Mean Embeddings.” *SIAM Journal on Mathematics of Data Science*, 2, 3, 583–606. doi:[10.1137/19M1305069](https://doi.org/10.1137/19M1305069).
- P. Kouvaros and A. Lomuscio. 2021. “Towards Scalable Complete Verification of ReLU Neural Networks via Dependency-Based Branching.” In: *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2643–2650.
- H. J. Kushner. 1967. *Stochastic Stability and Control*. Vol. 33. Academic Press New York.
- J. B. Lasserre. 2001. “Global Optimization with Polynomials and the Problem of Moments.” *SIAM Journal on Optimization*, 11, 3, 796–817. doi:[10.1137/S1052623400366802](https://doi.org/10.1137/S1052623400366802).
- L. Laurenti and M. Lahijanian. 2023. *Unifying Safety Approaches for Stochastic Systems: From Barrier Functions to Uncertain Abstractions via Dynamic Programming*. (2023). arXiv: [2310.01802](https://arxiv.org/abs/2310.01802).
- A. Lavaei, S. Soudjani, A. Abate, and M. Zamani. 2022. “Automated Verification and Synthesis of Stochastic Hybrid Systems: A Survey.” *Automatica*, 146, 110617. doi:[10.1016/j.automatica.2022.110617](https://doi.org/10.1016/j.automatica.2022.110617).
- Z. Li, D. Meunier, M. Mollenhauer, and A. Gretton. 2022. “Optimal Rates for Regularized Conditional Mean Embedding Learning.” *Advances in Neural Information Processing Systems*, 35, 4433–4445.
- J. Lofberg and P. A. Parrilo. 2004. “From Coefficients to Samples: A New Approach to SOS optimization.” In: *2004 43rd IEEE Conference on Decision and Control*. Vol. 3. IEEE, 3154–3159. doi:[10.1109/CDC.2004.1428957](https://doi.org/10.1109/CDC.2004.1428957).
- B. T. Lopez and J.-J. E. Slotine. 2023. “Unmatched Control Barrier Functions: Certainty Equivalence Adaptive Safety.” In: *2023 American Control Conference*, 3662–3668. doi:[10.23919/ACC55779.2023.10156496](https://doi.org/10.23919/ACC55779.2023.10156496).
- R. Majumdar, M. Salamati, and S. Soudjani. 2023. “Neural Abstraction-Based Controller Synthesis and Deployment.” *ACM Transactions on Embedded Computing Systems*, 22, 5s, 1–25. doi:[10.1145/3608104](https://doi.org/10.1145/3608104).
- A. Makdesi, A. Girard, and L. Fribourg. 2024. “Data-Driven Models of Monotone Systems.” *IEEE Transactions on Automatic Control*, 69, 8, 5294–5309. doi:[10.1109/TAC.2023.3346793](https://doi.org/10.1109/TAC.2023.3346793).
- U. Marteau-Ferey, F. Bach, and A. Rudi. 2020. “Non-Parametric Models for Non-Negative Functions.” *Advances in Neural Information Processing Systems*, 33, 12816–12826.
- P.-F. Massiani, S. Trimpe, and F. Solowjow. 2024. *On the Consistency of Kernel Methods with Dependent Observations*. (2024). arXiv: [2406.06101](https://arxiv.org/abs/2406.06101).
- F. B. Mathiesen, L. Romao, S. C. Calvert, L. Laurenti, and A. Abate. 2024. *A Data-Driven Approach for Safety Quantification of Non-Linear Stochastic Systems with Unknown Additive Noise Distribution*. (2024). arXiv: [2410.06662](https://arxiv.org/abs/2410.06662).
- R. Mazouz, J. Skovbekk, F. B. Mathiesen, E. Frew, L. Laurenti, and M. Lahijanian. 2024. *Data-Driven Permissible Safe Control with Barrier Certificates*. (2024). arXiv: [2405.00136](https://arxiv.org/abs/2405.00136).

- M. Mollenhauer, N. Mücke, and T. Sullivan. 2022. *Learning Linear Operators: Infinite-Dimensional Regression as a Well-Behaved Non-Compact Inverse Problem*. (2022). arXiv: [2211.08875](#).
- M. Nazeri, T. Badings, S. Soudjani, and A. Abate. 2025. *Data-Driven Yet Formal Policy Synthesis for Stochastic Nonlinear Dynamical Systems*. (2025). arXiv: [2501.01191](#).
- Y. Nemmour, H. Kremer, B. Schölkopf, and J.-J. Zhu. 2022. “Maximum Mean Discrepancy Distributionally Robust Nonlinear Chance-Constrained Optimization with Finite-Sample Guarantee.” In: *2022 IEEE 61st Conference on Decision and Control*, 5660–5667. doi:[10.1109/CDC51059.2022.9993212](#).
- J. Park and K. Muandet. 2020. “A Measure-Theoretic Approach to Kernel Conditional Mean Embeddings.” *Advances in Neural Information Processing Systems*, 33, 21247–21259.
- P. A. Parrilo. 2003. “Semidefinite Programming Relaxations for Semialgebraic Problems.” *Mathematical Programming*, 96, 293–320.
- D. P. Petersen and D. Middleton. 1962. “Sampling and Reconstruction of Wave-Number-Limited Functions in N-Dimensional Euclidean Spaces.” *Information and Control*, 5, 4, 279–323. doi:[10.1016/S0019-9958\(62\)90633-2](#).
- L. Pfister and Y. Bresler. 2018. “Bounding Multivariate Trigonometric Polynomials.” *IEEE Transactions on Signal Processing*, 67, 3, 700–707. doi:[10.1109/TSP.2018.2883925](#).
- S. Prajna, A. Jadbabaie, and G. J. Pappas. 2007. “A Framework for Worst-Case and Stochastic Safety Verification using Barrier Certificates.” *IEEE Transactions on Automatic Control*, 52, 8, 1415–1428. doi:[10.1109/TAC.2007.902736](#).
- “SOSTOOLS and its Control Applications.” *Positive Polynomials in Control*. Springer Berlin Heidelberg, 273–292. ISBN: 978-3-540-31594-0. doi:[10.1007/10997703_14](#).
- M. L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- A. Rahimi and B. Recht. 2007. “Random Features for Large-Scale Kernel Machines.” *Advances in Neural Information Processing Systems*, 20.
- C. E. Rasmussen and C. K. I. Williams. Nov. 2005. *Gaussian Processes for Machine Learning*. The MIT Press, (Nov. 2005). ISBN: 9780262256834. doi:[10.7551/mitpress/3206.001.0001](#).
- L. Romao, A. R. Hota, and A. Abate. 2023. “Distributionally Robust Optimal and Safe Control of Stochastic Systems via Kernel Conditional Mean Embedding.” In: *2023 62nd IEEE Conference on Decision and Control*, 2016–2021. doi:[10.1109/CDC49753.2023.10383997](#).
- N. V. Sahinidis. 1996. “BARON: A General Purpose Global Optimization Software Package.” *Journal of Global Optimization*, 8, 201–205. doi:[10.1007/BF00138693](#).
- A. Salamati, A. Lavaei, S. Soudjani, and M. Zamani. 2024. “Data-Driven Verification and Synthesis of Stochastic Systems via Barrier Certificates.” *Automatica*, 159, 111323. doi:[10.1016/j.automatica.2023.111323](#).
- C. Santoyo, M. Dutreix, and S. Coogan. 2021. “A Barrier Function Approach to Finite-Time Stochastic System Verification and Control.” *Automatica*, 125, 109439. doi:[10.1016/j.automatica.2020.109439](#).
- B. Schölkopf and A. J. Smola. 2002. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT press.
- O. Schön, B. van Huijgevoort, S. Haesaert, and S. Soudjani. 2024. “Bayesian Formal Synthesis of Unknown Systems via Robust Simulation Relations.” *IEEE Transactions on Automatic Control*, 70, 3, 1531–1546. doi:[10.1109/TAC.2024.3459308](#).
- O. Schön, S. Naseer, B. Wooding, and S. Soudjani. 2024. “Data-Driven Abstractions via Binary-Tree Gaussian Processes for Formal Verification.” *IFAC-PapersOnLine*, 58, 11, 115–122. doi:[10.1016/j.ifacol.2024.07.434](#).
- O. Schön, Z. Zhong, and S. Soudjani. 2024. “Data-Driven Distributionally Robust Safety Verification Using Barrier Certificates and Conditional Mean Embeddings.” In: *2024 American Control Conference*, 3417–3423. doi:[10.23919/ACC60939.2024.10644349](#).
- A. M. Schweidtmann, D. Bongartz, D. Grothe, T. Kerkenhoff, X. Lin, J. Najman, and A. Mitsos. 2021. “Deterministic Global Optimization with Gaussian Processes Embedded.” *Mathematical Programming Computation*, 13, 3, 553–581. doi:[10.1007/s12532-021-00204-y](#).
- A. Smola, A. Gretton, L. Song, and B. Schölkopf. 2007. “A Hilbert Space Embedding for Distributions.” In: *International Conference on Algorithmic Learning Theory*. Springer, 13–31. doi:[10.1007/978-3-540-75225-7_5](#).
- O. So, Z. Serlin, M. Mann, J. Gonzales, K. Rutledge, N. Roy, and C. Fan. 2023. *How to Train Your Neural Control Barrier Function: Learning Safety Filters for Complex Input-Constrained Systems*. (2023). arXiv: [2310.15478](#).
- L. Song, J. Huang, A. Smola, and K. Fukumizu. 2009. “Hilbert Space Embeddings of Conditional Distributions with Applications to Dynamical Systems.” In: *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 961–968. doi:[10.1145/1553374.1553497](#).
- S. Soudjani and A. Abate. 2013. “Adaptive and Sequential Gridding Procedures for the Abstraction and Verification of Stochastic Processes.” *SIAM Journal on Applied Dynamical Systems*, 12, 2, 921–956. doi:[10.1137/120871456](#).
- D. A. Spielman and S.-H. Teng. 2004. “Smoothed Analysis of Algorithms: Why the Simplex Algorithm Usually Takes Polynomial Time.” *Journal of the ACM (JACM)*, 51, 3, 385–463. doi:[10.1145/990308.990310](#).
- I. Steinwart and A. Christmann. 2009. “Fast Learning from Non-iid Observations.” *Advances in Neural Information Processing Systems*, 22.
- I. Steinwart and A. Christmann. 2008. *Support Vector Machines*. (1st ed.). Springer. ISBN: 0387772413.
- R. S. Streett. 1981. “Propositional Dynamic Logic of Looping and Converse.” In: *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing*. ACM, 375–383. doi:[10.1145/800076.802492](#).
- C. Wang, Y. Meng, J. Liu, and S. Smith. 2023. *Stochastic Control Barrier Functions with Bayesian Inference for Unknown Stochastic Differential Equations*. (2023). arXiv: [2312.12759](#).

- L. Wang, E. A. Theodorou, and M. Egerstedt. 2018. “Safe Learning of Quadrotor Dynamics using Barrier Certificates.” In: *2018 IEEE International Conference on Robotics and Automation*. IEEE, 2460–2465. doi:[10.1109/ICRA.2018.8460471](https://doi.org/10.1109/ICRA.2018.8460471).
- X. Wang, L. Knoedler, F. B. Mathiesen, and J. Alonso-Mora. 2024. “Simultaneous Synthesis and Verification of Neural Control Barrier Functions Through Branch-and-Bound Verification-in-the-Loop Training.” In: *2024 European Control Conference*. IEEE, 571–578. doi:[10.23919/ECC64448.2024.10591251](https://doi.org/10.23919/ECC64448.2024.10591251).
- K. Xu, Z. Shi, H. Zhang, Y. Wang, K.-W. Chang, M. Huang, B. Kailkhura, X. Lin, and C.-J. Hsieh. 2020. “Automatic Perturbation Analysis for Scalable Certified Robustness and Beyond.” *Advances in Neural Information Processing Systems*, 33, 1129–1141.
- X. Yin, B. Gao, and X. Yu. 2024. “Formal Synthesis of Controllers for Safety-Critical autonomous Systems: Developments and Challenges.” *Annual Reviews in Control*, 57, 100940. doi:[10.1016/j.arcontrol.2024.100940](https://doi.org/10.1016/j.arcontrol.2024.100940).
- Z. Zhang, C. Ma, S. Soudjani, and S. Soudjani. 2024. “Formal Verification of Unknown Stochastic Systems via Non-Parametric Estimation.” In: *International Conference on Artificial Intelligence and Statistics*. PMLR, 3277–3285.
- I. Ziemann and S. Tu. 2022. “Learning with Little Mixing.” *Advances in Neural Information Processing Systems*, 35, 4626–4637.

A Proof of Lemma 2

PROOF. The proof follows the same steps as Pfister and Bresler (2018, Corollaries 1 & 2). We first prove the upper bound and then the lower bound.

Upper bound: We start by defining the trigonometric polynomial $F: \tilde{\mathbb{X}} \rightarrow \mathbb{R}$,

$$F(x) := B(x) - \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2},$$

which is the barrier B with the output on $\mathbb{S}$ re-centered around zero. F satisfies

$$\max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} F(x) = \hat{B}_{\tilde{N}}^{\mathbb{S}} - \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} = \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}. \quad (28)$$

Via (25) and (28), we have for all $x \in \mathbb{S}$ that

$$\begin{aligned} B(x) &= F(x) + \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}, \\ &\leq \max_{x \in \mathbb{S}} F(x) + \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}, \\ &\stackrel{(25)}{\leq} C_{\tilde{N}} \max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} F(x) + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\max_{x \in \Theta_{\tilde{N}} \setminus \mathbb{S}} F(x) - \max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} F(x) \right) + \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}, \\ &\stackrel{(28)}{=} C_{\tilde{N}} \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} - \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} \right) + \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}, \\ &= C_{\tilde{N}} \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \hat{B}_{\tilde{N}}^{\mathbb{S}} \right) + \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}, \\ &= \frac{1}{2} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}} + C_{\tilde{N}} \left(\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right) \right) + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \hat{B}_{\tilde{N}}^{\mathbb{S}} \right), \end{aligned}$$

matching the upper bound in Lemma 2.

Lower bound: As for the upper bound, we start by defining a trigonometric polynomial,

$$F': \tilde{\mathbb{X}} \rightarrow \mathbb{R}, \quad F'(x) := \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} - B(x), \quad (29)$$

which, compared to F , is flipped in sign. F' satisfies

$$\max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} F'(x) \leq \left| \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} - \check{B}_{\tilde{N}}^{\mathbb{S}} \right| = \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2}. \quad (30)$$

Starting from (29), we have for all $x \in \mathbb{S}$ that

$$\begin{aligned} \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} &= B(x) + F'(x), \\ &\stackrel{(25)}{\leq} B(x) + C_{\tilde{N}} \max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} F'(x) + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\max_{x \in \Theta_{\tilde{N}} \setminus \mathbb{S}} F'(x) - \max_{x \in \Theta_{\tilde{N}} \cap \mathbb{S}} F'(x) \right), \\ &\stackrel{(30)}{=} B(x) + C_{\tilde{N}} \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} + A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} \left(\frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} + \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} - \check{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{S}} - \frac{\hat{B}_{\tilde{N}}^{\mathbb{S}} - \check{B}_{\tilde{N}}^{\mathbb{S}}}{2} \right). \end{aligned}$$

Reordering yields the upper bound in Lemma 2, concluding the proof. $\square$

B SILP to LP Relaxation

We provide further details on the finite-constraint relaxation of the SILP (23) to the LP (27). Let $\{x_0^{(1)}, \dots, x_0^{(\hat{N}_0)}\} \subset \mathbb{X}_0$, $\{x_u^{(1)}, \dots, x_u^{(\hat{N}_u)}\} \subset \mathbb{X}_u$, $\Theta_{\tilde{N}} := \{x^{(1)}, \dots, x^{(\hat{N})}\} \subset \mathbb{X}$, and $\Theta_{\tilde{N}} := \{x^{(1)}, \dots, x^{(\tilde{N})}\} \subset \tilde{\mathbb{X}}$ be sampling lattices with a common $C_{\tilde{N}}$ as in Lemma 1. Based on the constraint-tightening reasoning outlined in Subsection 7.2, we have that the first constraint (23b) given by $\forall x_0 \in \mathbb{X}_0: \phi_M(x_0)^\top b \leq \eta$ holds if

$$\hat{B}_{\tilde{N}}^{\mathbb{X}_0} \leq \frac{2\eta + (C_{\tilde{N}} - 1)\check{B}_{\tilde{N}}^{\mathbb{X}_0} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} + 1}$$

with the additional variables $\check{B}_{\tilde{N}}^{\mathbb{X}_0} \leq \phi_M(x_0^{(i)})^\top b$, $i = 1, \dots, \hat{N}_0$, $\hat{B}_{\tilde{N}}^{\mathbb{X}_0} \geq \phi_M(x_0^{(i)})^\top b$, $i = 1, \dots, \hat{N}_0$, and $\hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} \geq \phi_M(x^{(j)})^\top b$, $j = 1, \dots, \tilde{N}$. We call $\hat{B}_{\tilde{N}}^{\mathbb{X}_0} =: \hat{\eta}$ and thus get

$$\left. \begin{aligned} \check{B}_{\tilde{N}}^{\mathbb{X}_0} &\leq \phi_M(x_0^{(i)})^\top b \leq \hat{\eta}, & i = 1, \dots, \hat{N}_0, \\ \phi_M(x^{(i)})^\top b &\leq \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}, & i = 1, \dots, \tilde{N} - \hat{N}_0, \\ \hat{\eta} &:= \frac{2\eta + (C_{\tilde{N}} - 1)\check{B}_{\tilde{N}}^{\mathbb{X}_0} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} \hat{B}_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}_0} + 1} \end{aligned} \right\} \implies \forall x_0 \in \mathbb{X}_0: \phi_M(x_0)^\top b \leq \eta.$$

Similarly, we obtain for the constraint (23d) that

$$\left. \begin{aligned} \check{B}_{\tilde{\Delta}}^{\mathbb{X}} &\leq \phi_M(x^{(i)})^\top (Hb - b) \leq \hat{\Delta}, & i = 1, \dots, \hat{N}, \\ \phi_M(x^{(i)})^\top (Hb - b) &\leq \hat{B}_{\tilde{\Delta}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}, & i = 1, \dots, \tilde{N} - \hat{N}, \\ \hat{\Delta} &:= \frac{2(c - \varepsilon \bar{B}\kappa) + (C_{\tilde{N}} - 1)\check{B}_{\tilde{\Delta}}^{\mathbb{X}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}} \hat{B}_{\tilde{\Delta}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}}{C_{\tilde{N}} - 2A_{\tilde{N}}^{\tilde{\mathbb{X}} \setminus \mathbb{X}}} + 1 \end{aligned} \right\} \implies \forall x \in \mathbb{X}: \phi_M(x)^\top (Hb - b) \leq c - \varepsilon \bar{B}\kappa.$$

For the remaining lower bounds in the constraints (23c) and (23b), the constraint-tightening reasoning outlined in Subsection 7.2 gives

$$\left. \begin{aligned} \hat{\gamma} &\leq \phi_M(x_u^{(i)})^\top b \leq \hat{B}_{\tilde{N}}^{\mathbb{X}_u}, & i &= 1, \dots, \hat{N}_u, \\ \hat{B}_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}_u} &\leq \phi_M(x^{(i)})^\top b, & i &= 1, \dots, \tilde{N} - \hat{N}_u, \\ \hat{\gamma} &:= \frac{2+(C_{\tilde{N}}-1)\hat{B}_{\tilde{N}}^{\mathbb{X}_u}-2A_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}_u}\hat{B}_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}_u}}{C_{\tilde{N}}-2A_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}_u}} + 1 \end{aligned} \right\} \implies \forall x_u \in \mathbb{X}_u: \phi_M(x_u)^\top b \geq 1,$$

$$\left. \begin{aligned} \hat{\xi} &\leq \phi_M(x^{(i)})^\top b \leq \hat{B}_{\tilde{N}}^{\mathbb{X}}, & i &= 1, \dots, \hat{N}, \\ \hat{B}_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}} &\leq \phi_M(x^{(i)})^\top b, & i &= 1, \dots, \tilde{N} - \hat{N}, \\ \hat{\xi} &:= \frac{(C_{\tilde{N}}-1)\hat{B}_{\tilde{N}}^{\mathbb{X}}-2A_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}}\hat{B}_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}}}{\hat{B}_{\tilde{N}}^{\mathbb{X}}-2A_{\tilde{N}}^{\mathbb{X} \setminus \mathbb{X}}} + 1 \end{aligned} \right\} \implies \forall x \in \mathbb{X}: \phi_M(x)^\top b \geq 0,$$

respectively. Substituting the relaxed constraints into the SILP (23), we obtain the LP (27).

Received 19 January 2026; accepted 14 March 2026