

A data-driven approach for safety quantification of non-linear stochastic systems with unknown additive noise distribution

Frederik Baymler Mathiesen ^a, Licio Romao ^b, Simeon C. Calvert ^c, Luca Laurenti ^a,
and Alessandro Abate ^b

^a*Delft Center for Systems and Control, TU Delft.*

^b*Department of Computer Science, University of Oxford.*

^c*Department of Transport & Planning, TU Delft.*

Abstract

In this paper, we present a novel data-driven approach to quantify safety for non-linear, discrete-time stochastic systems with unknown noise distribution. We define safety as the probability that the system remains in a given region of the state space for a given time horizon and, to quantify it, we present an approach based on Stochastic Barrier Functions (SBFs). In particular, we introduce an inner approximation of the stochastic program to design a SBF in terms of a chance-constrained optimisation problem, which allows us to leverage the scenario approach theory to design a SBF from samples of the system with Probably Approximately Correct (PAC) guarantees. Our approach leads to tractable, robust linear programs, which enable us to assert safety for non-linear models that were otherwise deemed infeasible with existing methods. To further mitigate the computational complexity of our approach, we exploit the structure of the system dynamics and rely on spatial data structures to accelerate the construction and solution of the underlying optimisation problem. We show the efficacy and validity of our framework in several benchmarks, showing that our approach can obtain substantially tighter certificates compared to state-of-the-art with a confidence that is several orders of magnitude higher.

1 Introduction

Safety-critical applications, such as autonomous driving [39] and robotics [26] require provable guarantees of safety, as undesirable behaviours may lead to catastrophic outcomes with long-term economic costs. As a consequence, asserting the safety of complex non-linear noisy systems has been the focus of many recent approaches [32], [1], [36], [11], [23], [42]. However, these approaches generally suffer from exponential complexity with respect to the dimension of the state space. Besides, generally, these approaches require that the distribution of the noise affecting the system is either known and Gaussian or of bounded support [36], [32]. Unfortunately, in practice, the noise characteristics of the system are often not known [1], [4], [34]. This leads to the main question of this paper: how can we compute formal certificates of safety for non-linear stochastic systems with

unknown noise distribution?

Following the existing literature [36], [40], [2], we define safety as the probability that the system will remain within a given safe set for a given time horizon. Common approaches to quantify safety for non-linear stochastic systems either rely on formal abstraction methods [3], [8] or on Stochastic Barrier Functions (SBFs) [36], [40], [27]. Abstraction-based methods build a discrete representation (i.e., a variant of a Markov chain) of the underlying stochastic system via the discretisation of its state space [3], [8]. Then, value iteration is performed on such representation to verify properties and synthesise controllers, which can be mapped back on the original system by relying on simulation relations between the discrete representation and the underlying system [17], [14]. Various approaches have been proposed that combine abstractions and data-driven methods, including distributionally robust methods [15], the scenario approach [3], [4], Chernoff bounds [24], and Gaussian process regression [19]. However, a common drawback of these approaches is the need to partition (or discretise) the state space of the original stochastic system and to solve the value iteration on the resulting discrete abstraction, which leads

Email address: f.b.mathiesen@tudelft.nl (Frederik Baymler Mathiesen).

¹ Frederik Baymler Mathiesen and Licio Romao contributed equally to this work. Corresponding author is Frederik Baymler Mathiesen.

to the well-known state-space explosion problem [8].

In contrast, SBFs [40], [21], [31] are Lyapunov-like functions, whose level sets allow one to bound the probability that a dynamical system will remain safe for a given time horizon, without the need to explicitly evolve the system over time [36], [40], [20]. By not requiring an analytical solution to the system’s governing equations over time, SBFs represent a promising technique to efficiently quantify the safety of stochastic systems. However, one of the main challenges in SBF design, as we are interested in quantitative (rather than binary) safety evaluation, is in finding a SBF that does not lead to overly conservative results [22], [10]. In fact, designing a SBF requires the solution to a stochastic optimisation problem whose nature depends on the class of dynamics under study. In literature, synthesis of SBFs is usually performed with convex optimisation, in particular sum-of-squares (SoS) optimisation [32], [20], [36] and Linear Programming (LP) for piece-wise constant SBFs [29], or with deep learning [27]. However, these papers make restrictive assumptions on the noise distribution and on the system dynamics, e.g. Gaussian noise and linear or polynomial dynamics, are often enforced [15], [33]. More recently, some papers relax the assumption on the noise distribution by means of data-driven approaches [34], [35], which in addition to the level of safety also induces a (formally quantified) confidence. In particular, [34] uses Sample Average Approximation (SAA) to synthesise a SBF. This alternative approach, however, suffers from a sample complexity that is linear in the inverse of the confidence. Recent work on synthesising SBFs purely from trajectory data [38], i.e., without assuming partial knowledge of the dynamics, has used Conditional Mean Embeddings, whose bottleneck is a computational complexity of $O(N^3)$ where N is the number of samples.

Our approach departs from previous techniques to ensure safety for non-linear stochastic dynamical systems using SBFs. First, we present an inner approximation to the stochastic program commonly used to design SBF in terms of a chance-constrained optimisation problem. The feasible set of the latter is contained in the feasible set of the former. Then, by restricting to the class of piece-wise affine stochastic barrier functions² and relying on uncertain linear relaxations of non-linear systems [43], we show that the resulting chance-constrained problem can be reformulated into a robust LP problem [6]. This reformulation allows us to employ the scenario approach theory to devise a data-driven framework to synthesise a SBF, and consequently obtain safety guarantees to the trajectories of the system. The resulting approach is data-efficient, as it only requires a limited amount of samples from the noise distribution that is logarithmic in the negative inverse of the confidence, which

² It is known that with sufficiently many pieces a piece-wise affine function can formally approximate any continuous function arbitrarily well.

is in contrast with [34] where, as previously mentioned, the sample complexity is instead proportional in the inverse of the confidence. Furthermore, our approach is scalable due to its LP representation. We also introduce an a priori sample discarding procedure and spatial indexing for efficient model construction to improve scalability of the proposed framework. Our experiments show competitive performance on various systems, including a model of a vehicle in windy conditions and various Neural Network Dynamical Models (NNDMs) [30]. Our numerical analysis illustrates how our approach substantially outperforms state-of-the-art comparable methods in terms of both the tightness of bounds and the amount of data required. Overall, our main contributions are:

- We develop a data-driven technique for the design of barrier functions that relies on a chance-constrained inner approximation of stochastic programs.
- We use the scenario approach theory and Linear Bound Propagation techniques to synthesise a SBF for non-linear dynamical systems.
- We present an efficient computational architecture to construct the resulting optimisation problem using spatial indexing methods, such as R-trees, for faster set intersection computations.
- We benchmark the proposed framework, showing its advantages with respect to approaches in the literature, including instances of Neural Network Dynamical Models (NNDMs).

A conference version of this paper appeared in [28]. This paper significantly expands the preliminary results in [28], which focused on developing data-driven SBF synthesis techniques only for piece-wise affine (PWA) dynamics. In this paper, we deal with general non-linear dynamics, which necessitates non-trivial extensions of the techniques in [28] to handle uncertain PWA models. Furthermore, here we show how to leverage spatial indexing and sample discarding to enhance scalability of our framework. We also newly include detailed proofs, as well as a substantially extended empirical analysis, including experiments on NNDMs.

The structure of this paper is as follows. The problem statement is described in Section 2. Section 3 covers preliminary results from convex analysis, scenario optimisation, and piece-wise affine relaxations of non-linear functions. Section 4 defines piece-wise linear SBFs and how they can be used to certify safety for non-linear stochastic systems. Section 5 details the main theoretical results, including a reformulation of the stochastic program to synthesise a SBF in to an inner chance-constrained approximation. In Section 6, we apply scenario approach theory towards a data-driven synthesis of a PWA SBF. In Section 7 we show how to compute the necessary polyhedral over-approximations and two methods to improve the scalability, namely the convex hull over sample set and spatial indexing. Empirical studies are reported in Section 8.

2 Problem formulation

2.1 Notation

Let \mathbb{R} , $\mathbb{R}_{\geq 0}$, and \mathbb{N} represent, respectively, the set of real, non-negative real, and natural numbers. We denote by $\{x_1, \dots, x_m\}$ the elements in a finite set. Let Ω be an abstract space, \mathcal{F} be a σ -algebra defined on this set, and \mathbb{P} be a probability measure; we denote by $(\Omega, \mathcal{F}, \mathbb{P})$ the associated complete probability (or uncertainty) space. A random η variable with values in \mathbb{R}^n is a measurable function $\eta : \Omega \mapsto \mathbb{R}^n$, with \mathbb{R}^n equipped with its standard Borel σ -algebra. A realisation of η is denoted by $\eta(\omega)$, for some $\omega \in \Omega$. For a given set Ω , we denote by $\mathcal{P}(\Omega)$ the set of probability measures defined over Ω . The supremum norm $\|\cdot\|_\infty$ on a function $f : \mathbb{R}^n \mapsto \mathbb{R}$ is defined as $\|f\|_\infty = \sup_{x \in \mathbb{R}^n} |f(x)|$.

2.2 Problem formulation

We consider the following discrete-time, stochastic non-linear system

$$x(k+1) = f(x(k)) + \eta, \quad x(0) = \bar{x}, \quad (1)$$

where $k \in \mathbb{N}$ denotes time, $\bar{x} \in \mathbb{R}^n$ is the initial condition, and $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ is a Lipschitz continuous function. η is a random variable defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and, at each time instance, an independent realization is drawn and added to the nominal dynamics represented by the function f (independent and identically distributed (i.i.d.)). We further assume that the probability distribution of η is absolutely continuous with respect to the Lebesgue measure of \mathbb{R}^n for a well-defined probability density function. Throughout this paper, we assume that the probability distribution of η is unknown.

Because of the i.i.d. assumption on η , we can alternatively write the dynamics of System (1) using its kernel representation. A *stochastic kernel* is a measurable map from \mathbb{R}^n onto the space of probability measures $\mathcal{P}(\mathbb{R}^n)$, $T : \mathbb{R}^n \mapsto \mathcal{P}(\mathbb{R}^n)$ [5]. In particular, the stochastic kernel associated with System (1) is given by

$$T(X | x) = \int_{\Omega} \mathbf{1}_X(f(x) + \eta(\omega)) \mathbb{P}(d\omega), \quad (2)$$

where $X \subset \mathbb{R}^n$ is a Borel set of \mathbb{R}^n , and $\mathbf{1}_X$ is the indicator function, i.e., $\mathbf{1}_X(x) = 1$ if $x \in X$, and 0 otherwise. In other words, for a fixed $x \in \mathbb{R}^n$ representing the current state, the stochastic kernel associated with System (1) returns the probability distribution of the state in the next time step. For $K \in \mathbb{N}$ and $\bar{x} \in \mathbb{R}^n$, we denote by $(\mathbb{R}^n)^K = \mathbb{R}^n \times \dots \times \mathbb{R}^n$ the K -fold Cartesian product of \mathbb{R}^n . Then, the stochastic kernel (2) induces a unique measure on $(\mathbb{R}^n)^K$ given by the unique extension (due

to Kolmogorov's extension theorem [41, Theorem 2.4.3]) of the measure

$$\mathbb{P}^{\bar{x}}(X_1, \dots, X_K) = \int_{X_1} \left(\prod_{k=2}^K \int_{X_k} T(d\xi_k | \xi_{k-1}) \right) T(d\xi_1 | \bar{x}),$$

which represents the probability of a trajectory starting at \bar{x} , under the dynamics of System (1), being in the set X_k , $k = 1, \dots, K$ at the various time steps. We now have all the ingredients to define the notion of probabilistic safety that is crucial to our approach.

Definition 1 (Probabilistic safety [2]). Let K be a non-negative integer and $X_s \subseteq X$ be a compact set representing the safe set. Then, for a given initial state \bar{x} , we define probabilistic safety as

$$\zeta(X_s, K | \bar{x}) = \mathbb{P}^{\bar{x}}(X_s, \dots, X_s).$$

Our main objective in this paper is to obtain a uniform lower bound on the probabilistic safety of System (1) for all initial conditions in a given set X_0 .

Problem 1. Let $D = (\omega_1, \dots, \omega_N)$ be i.i.d. samples from η . Then, for a given compact set $X_0 \subseteq X$, a safe set $X_s \subseteq X$, and a time horizon $K \in \mathbb{N}$, find $\rho \in (0, 1]$ such that, with high confidence,

$$\zeta(X_s, K) = \inf_{\bar{x} \in X_0} \zeta(X_s, K | \bar{x}) \geq \rho.$$

Note that, as the noise is additive, the assumption of having i.i.d. noise samples in Problem 1 is equivalent to having i.i.d. full measurements of the state. Exact computation of probabilistic safety for System (1) is generally infeasible, even when η is restricted to being a Gaussian random variable [8]. Consequently, it is obvious that the more general setting considered in this paper, where the noise distribution is arbitrary and possibly unknown, requires approximations. We should also stress that, while in Problem 1 we focus on safety, the techniques developed in this paper to obtain its solution can also be applied to perform verification of more complex temporal properties. In fact, safety is the dual to, and can be reformulated as, reachability[2]³, and complex temporal specifications such as LTLf [12] can be reduced to reachability by representing the formula as an automaton and checking reachability to an accepting state of the product system between the automaton and the system [23], [20], [17].

Approach We leverage the availability of data D to synthesise a Stochastic Barrier Function (SBF) that,

³ The safety probability is one minus the probability of reaching the unsafe set.

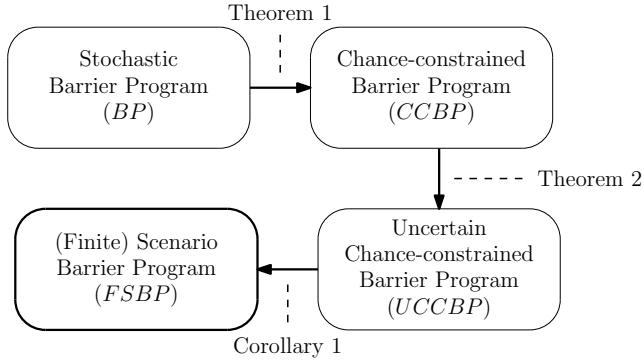


Fig. 1. Summary of the proposed framework for safety verification of nonlinear systems. Theorem 1 (Section 5) guarantees that the associated stochastic program to find a SBF can be solved as a chance-constrained program via constraint tightening. Non-linearities can be dealt with by means of a PWA abstraction as in Theorem 2 (Section 5.1). The scenario approach is applied to use available data for high-confidence certificates; strong duality is essential to relax the problem to Linear Programming in Corollary 1 (in Section 6).

with high confidence, bounds probabilistic safety. Our approach is summarised in Figure 1 and is based on a novel reformulation of the optimisation problem to find a SBF into a chance-constrained problem (detailed in Section 5). To deal with the non-linearity of a system, we develop in Section 5.1 a formal (i.e. with a quantified error) approximation of System (1) as an uncertain piece-wise affine (PWA) system. Such an approximation is then used in Section 6 to reformulate the chance-constrained problem into a robust linear program, which can be efficiently solved using duality.

3 Preliminaries

Our approach requires some results from convex optimisation and scenario optimisation, which we summarise below for convenience.

3.1 Convex analysis and linear programming

A class of convex sets that will be widely used in this paper is the class of polyhedral sets [25]. Given a matrix $H \in \mathbb{R}^{p \times n}$ and a vector $h \in \mathbb{R}^p$, a polyhedral set is denoted by

$$P = \{x \in \mathbb{R}^n : Hx \leq h\}. \quad (3)$$

Representation (3) is called the *half-space representation* of polyhedral sets. The vertex representation of P is $P = \text{conv}(x_1, \dots, x_m)$ where $\text{vert}(P) = \{x_1, \dots, x_m\}$ are the vertices⁴ of P .

⁴ Formally, an element x of a convex set C is called a vertex if whenever $x = \lambda x_1 + (1 - \lambda)x_2$ for some $\lambda \in (0, 1)$ and $x_1, x_2 \in C$, we have that $x_1 = x_2$.

The following class of robust, or semi-infinite, LPs is crucial to our developments:

$$\begin{aligned} \min_z \quad & c^\top z \\ \text{s. t.} \quad & (Az + a)^\top x \leq Bz + b, \quad \text{for all } x \in P, \end{aligned} \quad (4)$$

where $z \in \mathbb{R}^d$ is the vector of decision variables of size $d \in \mathbb{N}$, $c \in \mathbb{R}^d$ is the objective cost, $A \in \mathbb{R}^{n \times d}$, $a \in \mathbb{R}^n$, $B \in \mathbb{R}^{1 \times d}$, $b \in \mathbb{R}$ are constraint coefficients, and $P \subset \mathbb{R}^n$ is a polyhedral set. By relying on standard duality arguments, we can recast the semi-infinite optimisation problem (4) as a regular LP as shown in the following proposition.

Proposition 2 ([6, Exercise 5.17]). *Consider the semi-infinite LP problem (4), and denote by*

$$\mathcal{Z} = \{z \in \mathbb{R}^d : (Az + a)^\top x \leq Bz + b, \text{ for all } x \in P\},$$

its feasible set. Define the optimisation problem

$$\begin{aligned} \min_{z, \lambda} \quad & c^\top z \\ \text{s. t.} \quad & h^\top \lambda \leq Bz + b \\ & H^\top \lambda = Az + a, \quad \lambda \geq 0, \end{aligned} \quad (5)$$

whose feasible set is given by

$$\mathcal{Z}' = \{z : \exists \lambda \in \mathbb{R}_{\geq 0}^p, h^\top \lambda \leq Bz + b, H^\top \lambda = (Az + a)\}.$$

Then, it holds that $\mathcal{Z} = \mathcal{Z}'$.

3.2 Scenario optimisation

The scenario approach theory allows one to certify the probability of constraint violation associated with chance-constrained optimisation problems [7]. Let $\eta : \Omega \mapsto \mathbb{R}^q$ be a random variable on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $D = (\omega_1, \dots, \omega_N)$ be i.i.d. samples from \mathbb{P} , which live naturally in the space $(\Omega^N, \otimes_N \mathcal{F}, \mathbb{P}^N)$, where Ω^N is the N -fold Cartesian product of Ω , and $\otimes_N \mathcal{F}$ is the product sigma algebra generated by the sigma algebra \mathcal{F} , and \mathbb{P}^N represents the induced measure on Ω^N . Then, consider the scenario program

$$\begin{aligned} \min_z \quad & c^\top z \\ \text{s. t.} \quad & g(z, \eta(\omega)) \leq 0, \quad \text{for all } \omega \in D, \end{aligned} \quad (6)$$

where d is the dimension of the optimisation variables, $c \in \mathbb{R}^d$ is the objective cost, and $g(z, \eta) : \mathbb{R}^d \times \mathbb{R}^q$ is a function that is convex in z for each value of η and measurable in η for each value of z . Notice that the scenario program (6), by enforcing one convex constraint per sample in D , is convex program and consequently, it

can be solved using convex optimization tools [13], [18].

Assumption 1. We assume almost surely with respect to the measure \mathbb{P}^N that:

- The feasible set $\mathcal{Z} = \{z : g(z, \eta(\omega)) \leq 0, \forall \omega \in D\}$ has non-empty interior.
- The optimal solution of program (6) exists and is unique.

Both conditions in Assumption 1 are standard. In fact, uniqueness can always be enforced with a tie-break rule. Throughout this paper, we denote the unique solution of (6) by $z^*(D)$, which is a well-defined random variable on the space Ω^N . A key result within the scenario approach theory establishes an upper bound on the tail distribution of the constraint violation probability associated with $z^*(D)$.

Proposition 3 ([7]). *Consider the scenario program (6) and suppose Assumption 1 holds. Then, for any $\epsilon \in (0, 1)$, we have that*

$$\mathbb{P}^N \{D \in \Omega^N : V(z^*(D)) > \epsilon\} \leq \sum_{i=0}^{d-1} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i},$$

where $V(z) = \mathbb{P}\{\omega \in \Omega : g(z, \eta(\omega)) > 0\}$ is the violation probability of $z \in \mathbb{R}^d$.

The inequality in Proposition 3 holds with equality for the class of fully-supported scenario programs – the reader is referred to [7] for more details.

3.3 Uncertain piece-wise affine relaxations

Uncertain PWA relaxations, as depicted in Fig. 2, are key to our method. This type of relaxation allows one to treat complex non-linear functions as uncertain PWA functions, simplifying analysis and optimisation. An uncertain PWA relaxation for a function f is a collection of local linear relaxations, that is, $\underline{A}_i x + \underline{b}_i \leq f(x) \leq \bar{A}_i x + \bar{b}_i$ for all x in a convex region Q_i , given a partition $\mathcal{Q} = \{Q_1, \dots, Q_\ell\}$. We call a partition convex if each region in the partition is convex. The union of all regions $\bigcup_{i=1}^{\ell} Q_i$ is the domain of the relaxation. We note that any locally Lipschitz function f can be relaxed to an uncertain PWA function [9]. The idea is to partition the input domain into a number of regions and compute linear relaxations independently for each region. We formalise the existence of an uncertain PWA relaxation in what follows.

Proposition 4. *Let $\mathcal{Q} = \{Q_1, \dots, Q_\ell\}$ be a given convex partition of a compact set $X \subset \mathbb{R}^n$. Then, for any function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that is locally Lipschitz on X ,*




Fig. 2. Computation graph for the function $f(x) = x^3 + x$ with Linear Bound Propagation (LBP) annotation for the input regions $[-1, 0]$ and $[0, 1]$. LBP operates by propagating backward linear bounds on the computation graph.

there exists an uncertain PWA function

$$\hat{f}(x, \alpha) = \hat{f}_i(x, \alpha) = A_i(\alpha)x + b_i(\alpha), \quad \text{for } x \in Q_i \subseteq X$$

where $\alpha \in [0, 1]$ and

$$A_i(\alpha) = \alpha \underline{A}_i + (1 - \alpha) \bar{A}_i, \quad b_i(\alpha) = \alpha \underline{b}_i + (1 - \alpha) \bar{b}_i$$

with matrices $\underline{A}_i, \bar{A}_i \in \mathbb{R}^{n \times n}$ and vectors $\underline{b}_i, \bar{b}_i \in \mathbb{R}^n$ given for all $Q_i \in \mathcal{Q}$, such that it holds that $f(x) \in \{\hat{f}(x, \alpha) : \alpha \in [0, 1]\}$ for all $x \in X$.

A proof for the proposition can be found in Appendix A.2. Note that there are various possible approaches to select matrices $\underline{A}_i, \bar{A}_i$ and vectors $\underline{b}_i, \bar{b}_i$ for each region Q_i . In this paper, we use a state-of-the-art technique, called Linear Bound Propagation (LBP) [43]. The core idea of LBP is to recursively propagate linear bounds backward through the computation graph representing a function. An example of this bound propagation procedure can be seen in Fig. 2 where the cubic term is linearly bounded based on the input interval bounds and composed with the linear term.

4 Stochastic barrier functions

In this paper, we use Stochastic Barrier Functions [32] to provide safety guarantees for System (1).

Definition 2 (Stochastic Barrier Function). Given a safe set X_s and a set of initial conditions \bar{X} , a measurable function $B : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}$ is called a Stochastic Barrier Function (SBF) for System(1) if there exist non-negative constants γ, c satisfying

$$B(x) \leq \gamma, \quad \text{for all } x \in X_0, \quad (7)$$

$$B(x) \geq 1, \quad \text{for all } x \in \mathbb{R}^n \setminus X_s, \quad (8)$$

$$\mathbb{E}\{B(f(x) + \eta(\omega))\} \leq B(x) + c, \quad \text{for all } x \in X_s. \quad (9)$$

A pictorial representation of a SBF is presented in Fig-




Fig. 3. The figure is borrowed from [27]. A SBF $B(x)$ is a non-negative function that is greater than 1 in an unsafe region X_u , which is the complement of the safe set X_s . The variable γ is an upper bound for $B(x)$ over an initial region X_0 . The upper bound for the expected increase in $B(x)$ after one step of (1) over the safe set X_s is denoted c . Proposition 5 shows that $\zeta(X_s, T) \geq 1 - (\gamma + cT)$.

ure 3. Inequality (9) requires that for all $x \in X_s$, the expected value of the barrier function at the next step cannot increase more than c . As shown in Proposition 5 below, we can leverage results from martingale theory to obtain a lower bound on $\zeta(X_s, K)$.

Proposition 5 ([21, Chapter 3, Theorem 3], [40, Section 2.2]). *For a safe $X_s \subset \mathbb{R}^n$ and a set of initial conditions $X_0 \subset X_s$, let the function $B : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0}$ be a SBF for System (1), and let the positive integer K be a time horizon. Then, it holds that $\zeta(X_s, K) \geq 1 - (\gamma + cK)$.*

Thanks to Proposition 5, one can formulate the search for a barrier certificate as the following infinite-dimensional stochastic optimization problem

$$\begin{aligned} \min_{B \in \mathcal{M}, c, \gamma} \quad & \gamma + cK \\ \text{s.t.} \quad & (7), (8), (9), \end{aligned} \quad (\text{BP})$$

where \mathcal{M} represents the set of non-negative measurable functions in \mathbb{R}^n . Two challenges emerge when solving problem (BP) to solve Problem 1: (i) \mathcal{M} is an infinite-dimensional space, and (ii) constraint (9) involves computing an expectation over an unknown probability distribution. To address the first challenge, we restrict the search for barrier functions to the class of piece-wise affine functions given by

$$\begin{aligned} \mathcal{M}' = \{ & B : \mathbb{R}^n \mapsto \mathbb{R}_{\geq 0} : \\ & B(x, \theta) = \max\{B_1(x, \theta), \dots, B_\ell(x, \theta)\}, \end{aligned} \quad (10)$$

with $\theta \in \mathbb{R}^{\ell(n+1)}$ representing the set of parameters $(u_i, v_i) \in \mathbb{R}^{(n+1)}$ that define the barrier function, i.e.,

$$B_i(x, \theta) = \begin{cases} u_i^\top x + v_i, & \text{if } x \in Q_i, \\ 0, & \text{otherwise,} \end{cases}$$

where the polyhedral sets $\bar{Q}_i = \{x \in \mathbb{R}^n : H_i x \leq h_i\}$,

for all $i \in \{1, \dots, \bar{\ell}\}$, constitute a partition of \mathbb{R}^n . While various classes of barrier functions have been considered in the literature [32], [36], [40], in this work, we focus on PWA barrier functions, as they are expressive enough to approximate any non-linear function arbitrarily well and, as we will show in Section 6, this choice leads to a LP program to synthesise a SBF, thus guaranteeing efficiency. For the second challenge, in the next section we develop a new approach to create a chance-constrained approximation of (BP) using a novel constraint tightening technique.

5 An inner chance-constrained approximation of Problem (BP)

Solving Problem (BP) is challenging, even in the case that the probability distribution underpinning (9) is known. In this section, we show how to relax Problem (BP) using a reformulation in terms of a chance-constrained optimisation problem whose feasible set is contained in the set of feasible solutions of (BP). Such ideas have never been used in this context. Hence, we depart completely from the approaches taken by [20], [32], [34], [36], [40], and [44], which either rely on approximating Constraint (9) with the empirical distribution, make strong assumptions about the noise distribution to analytically compute the expectation and recast (9) as a convex constraint, or rely on convex over-approximations of the expectation to (conservatively) verify (9). Furthermore, due to the inner approximation of (BP) in terms of a chance-constrained problem, our approach opens the road to use the tools of scenario optimisation discussed in Section 3.2 to obtain strong sample complexity guarantees on the safety probability of System (1).

To this end, fix any $B \in \mathcal{M}$ and $x \in \mathbb{R}^n$, and let $E \in \mathcal{F}$ be a measurable set. Then, observe that

$$\begin{aligned} \mathbb{E}\{B(f(x) + \eta(\omega))\} &= \int_E B(f(x) + \eta(\omega)) \mathbb{P}(d\omega) \\ &+ \int_{E^c} B(f(x) + \eta(\omega)) \mathbb{P}(d\omega). \end{aligned} \quad (11)$$

Under the assumption that $B(x) \leq M$ for any $x \in \mathbb{R}^n$, which can always be enforced for SBFs, the second term on the right-hand side of (11) can be bounded by $M\mathbb{P}\{E^c\}$. Our main idea then is to choose a particular E^c that allows us to control the right-hand side of (11). Such an intuitive reasoning is made formal in the next lemma.

Lemma 1. For $B \in \mathcal{M}$ with $\|B\|_\infty = M$, let c and η be as in (BP). Define the set

$$E = \{\omega \in \Omega : B(f(x) + \eta(\omega)) + \xi \leq B(x) + c, \quad \text{for all } x \in X_s\}, \quad (12)$$

for some constant $\xi \geq 0$. If there exists an $\epsilon \in (0, 1)$ such that $\xi \geq M \frac{\epsilon}{1-\epsilon}$ and $\mathbb{P}\{E^c\} \leq \epsilon$, then we have $\mathbb{E}\{B(f(x) + \eta(\omega))\} \leq B(x) + c$ for all $x \in X_s$.

Proof. The proof of our result is based on (11). We first remark that E is a measurable set (see Appendix A.3). Next, fix any $x \in X_s$ and define the set

$$E_x = \{\omega \in \Omega : B(f(x) + \eta(\omega)) + \xi \leq B(x) + c\}. \quad (13)$$

Notice that, by definition, we have $E \subset E_x$. Furthermore, the set E_x is also measurable due to the assumptions on B , f and η , and the fact that measurability is closed under composition. Using (11), we then obtain

$$\mathbb{E}\{B(f(x) + \eta(\omega))\} \leq (B(x) + c - \xi)\mathbb{P}\{E_x\} + M\mathbb{P}\{E_x^c\}, \quad (14)$$

where the first term on the right-hand side of (14) follows from the definition of E_x in (13), and the second by the uniform boundedness condition on B . It remains to show that

$$(B(x) + c - \xi)\mathbb{P}\{E_x\} + M\mathbb{P}\{E_x^c\} \leq B(x) + c. \quad (15)$$

To this end, it suffices to show that

$$-\xi\mathbb{P}\{E_x\} + M\mathbb{P}\{E_x^c\} \leq 0 \quad (16)$$

holds, since $B(x) + c$ is non-negative and $\mathbb{P}\{E_x\} \leq 1$ by definition. Substituting the two inequalities $\mathbb{P}\{E_x\} \geq 1 - \epsilon$ and $\mathbb{P}\{E_x^c\} \leq \epsilon$ into the left-hand side of (16) we obtain

$$-\xi\mathbb{P}\{E_x\} + M\mathbb{P}\{E_x^c\} \leq -\xi(1 - \epsilon) + M\epsilon,$$

whose right-hand side is less than or equal to zero due to the fact that $\xi \geq M \frac{\epsilon}{1-\epsilon}$. This concludes the proof of the lemma. \square

Lemma 1 is enabled by the chance-constraint tightening variable ξ through the condition that $\xi \geq M \frac{\epsilon}{1-\epsilon}$. A reader may question how to choose ξ and ϵ . Choosing ϵ is a trade-off between assigning less probability mass to the uniform upper bound, which is desirable as the uniform upper bound represents a worst-case for barrier value at the next step, is and the amount of data required when applying the scenario approach theory. In our experiments (see Section 8), we employ $\epsilon = 0.005$. Once ϵ is chosen, the optimal choice of ξ to minimize c is $\xi = M \frac{\epsilon}{1-\epsilon}$, which follows from the fact that ξ is on the left-hand side of the inner inequality of E (and E_x) with c on the right-hand side and $M \frac{\epsilon}{1-\epsilon}$ is the smallest allowed value of ξ .

An immediate consequence of Lemma 1 is the fact that

we can obtain an inner approximation of the optimisation problem (BP) in terms of a chance-constrained optimisation problem.

Theorem 1. Consider the dynamical system given in System (1). Then, we have that for all $\epsilon \in (0, 1)$ and positive integers K and $M \geq 1$ such that $\xi \geq M \frac{\epsilon}{1-\epsilon}$, the feasible set of

$$\begin{aligned} \min_{B \in \mathcal{M}, c, \gamma} \quad & \gamma + cK \\ \text{s.t.} \quad & (7), (8), \quad \gamma \geq 0, \quad c \geq 0, \\ & B(x) \leq M, \quad \forall x \in \mathbb{R}^n \\ & \mathbb{P}\{\omega \in \Omega : B(f(x) + \eta(\omega)) + \xi \\ & \leq B(x) + c, \quad \forall x \in X_s\} \geq 1 - \epsilon, \end{aligned} \quad (\text{CCBP})$$

is contained in the feasible set of (BP).

Proof. Constraints (7) and (8) are shared between Problem (BP) and (CCBP). Thus, it is sufficient to show that the chance-constraint of Problem (CCBP) implies that constraint (9) is satisfied. Rewriting the chance-constraint as $\mathbb{P}\{E\} \geq 1 - \epsilon$ where E is defined as in Lemma 1, it holds that $\mathbb{P}\{E^c\} \leq \epsilon$. From this, the assumption $\xi \geq M \frac{\epsilon}{1-\epsilon}$, and the constraint that B is uniformly bounded by M , the conditions of Lemma 1 are satisfied. Thus, it holds that if the chance constraint of Problem (CCBP) is satisfied, then constraint (9) is satisfied. \square

Unfortunately, the computational burden of solving Problem (CCBP) is not negligible, mainly due to the quantification over all $x \in X_s$ and to the non-linear function f . In Section 5.1, we will show how the framework developed in Section 3.3 can alleviate this burden.

5.1 Over-approximating general non-linear systems with uncertain PWA dynamics

In Section 3.3, we showed that any locally Lipschitz function can be over-approximated by uncertain PWA functions, and by extension non-linear systems of the form of System (1) can be over-approximated by uncertain PWA systems. More formally, an uncertain PWA over-approximation of System (1) is described as follows.

$$x(k+1) \in F(x(k)) + \eta(k), \quad (17)$$

where $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is a set-valued mapping defined as

$$F(x(k)) = \left\{ \hat{f}(x(k), \alpha) : \alpha \in [0, 1] \right\}, \quad (18)$$

with \hat{f} being a set-valued PWA function in the uncertain parameter α and such that for any $x \in \mathbb{R}^n$ it holds that

$f(x) \in F(x)$. Such a definition of an uncertain PWA over-approximation of System (1) guarantees that between System (1) and PWA System (17) there is a *behavioural inclusion* relation, that is, for any $x \in \mathbb{R}^m$ there exists an $\alpha \in [0, 1]$ such that $\hat{f}(x, \alpha) = f(x)$.

Intuitively, we would like to synthesise a SBF for System (17) and rely on the behavioural relation described above to ensure that the synthesised SBF also is an SBF for System (1). This is what we do in the following Theorem, where we extend Theorem 1 to an uncertain PWA over-approximation of System (1).

Theorem 2. Consider the dynamical system given in (1) and assume access to an uncertain PWA over-approximation (17) of the system. Then, we have that for all $\epsilon \in (0, 1)$ and positive integer K , if there exist $M \geq 1$ and $\xi \geq M \frac{\epsilon}{1-\epsilon}$, then the feasible set of

$$\begin{aligned} \min_{B \in \mathcal{M}, c, \gamma} \quad & \gamma + cK \\ \text{s.t.} \quad & (7), (8), \quad \gamma \geq 0, \quad c \geq 0, \\ & B(x) \leq M, \quad \forall x \in \mathbb{R}^n \\ & \mathbb{P}\{\omega \in \Omega : B(y + \eta(\omega)) + \xi \\ & \leq B(x) + c, \quad \forall x \in X_s, \quad \forall y \in F(x)\} \geq 1 - \epsilon, \end{aligned} \quad (\text{UCCBP})$$

is contained in the feasible set of (BP).

Proof. Constraints (7) and (8) are imposed directly in (UCCBP). Since $f(x) \in F(x)$ implies that there exists an $\alpha \in [0, 1]$ such that $\hat{f}(x, \alpha) = f(x)$, it holds that $B(f(x) + \eta(\omega)) \leq \sup_{y \in F(x)} B(y + \eta(\omega))$. Therefore, the chance constraint of (UCCBP) implies the chance constraint of (CCBP). By Theorem 1 and transitivity of the subset relation, the feasible set of (UCCBP) is contained in the feasible set of (BP). \square

6 Data-driven approximation of (UCCBP) for the class of uncertain dynamical systems

To reformulate (UCCBP) as a robust LP problem, we need to introduce some mathematical notation. Let $\mathcal{Q} = \{Q_1, \dots, Q_\ell\}$ be a partition of the state space associated with a barrier function, as described in Section 4. Let us also denote four collections of indices by

$$\begin{aligned} I &= \{1, \dots, \ell\}, \\ I_{X_0} &= \{i \in I : Q_i \cap X_0 \neq \emptyset\}, \\ I_{X_s} &= \{i \in I : Q_i \cap X_s \neq \emptyset\}, \\ I_{X_u} &= \{i \in I : Q_i \cap X_u \neq \emptyset\}, \end{aligned} \quad (19)$$

which represent the collection of all indices, and the elements of \mathcal{Q} with non-empty intersections with the initial state, the safe set, and unsafe set, respectively. Finally,

for each pair $(i, j) \in I_{X_s} \times I$, we denote the set

$$Q_{ij}(\omega) = \{x \in Q_i : \exists y \in F(x), y + \eta(\omega) \in Q_j\}, \quad (20)$$

representing the subset of Q_i with i belonging to I_{X_s} and that is mapped to Q_j under a given realization of the noise. A pictorial example of $Q_{ij}(\omega)$ can be found in Figure 4. Leveraging the results of Theorem 2 and Proposition 3 and using the notation we have introduced so far, we obtain the following intermediate result. The goal of Lemma 2 below is two fold: (i) to impose piecewise constraints on the barrier synthesis and (ii) apply scenario approach theory to obtain a tractable solution in the face of an unknown noise distribution.

Lemma 2. Let $D = \{\omega_1, \dots, \omega_N\}$ be a collection of N independent samples from the noise distribution \mathbb{P} . Fix $\epsilon \in (0, 1)$, $M \geq 1$, and $\xi \geq M \frac{\epsilon}{1-\epsilon}$, and consider the scenario optimisation program

$$\begin{aligned} \min_z \quad & \gamma + cK \\ \text{s.t.} \quad & \gamma \geq 0, \quad c \geq 0, \\ & B_i(x, \theta) \in [0, M], \quad \forall x \in Q_i, i \in I, \\ & B_i(x, \theta) \leq \gamma, \quad \forall x \in Q_i, i \in I_{X_0}, \\ & B_i(x, \theta) \geq 1, \quad \forall x \in Q_i, i \in I_{X_u}, \\ & B_j(y + \eta(\omega), \theta) + \xi \leq B_i(x, \theta) + c, \\ & \forall (\omega, i, j) \in D \times I_{X_s} \times I, y \in F(x), x \in Q_{ij}(\omega) \end{aligned} \quad (\text{SBP})$$

where $d = 2 + \ell(n+1)$ is the number of decision variables in Problem (UCCBP), and $z = (\gamma, c, \theta)$ is the collection of optimisation variables. Then, with probability at least $1 - \beta$, where β is the right-hand side of the inequality in Proposition (3), the optimal solution of (SBP) satisfies the constraints of Definition 2.

Remark 1. The state-of-the-art literature on Sample Average Approximation (SAA) for SBF design relies on Chebyshev's inequality [34] to bound the probability of satisfaction, which yields a sample complexity proportional to $O(1/\beta)$. Instead, under the assumption that the barrier is uniformly upper bounded, which is always the case for our method, in Lemma 2 we can rely on Hoeffding's inequality, which yields a sample complexity proportional to $O(\log(1/\beta))$. \square

The relevance of Lemma 2 towards the general framework proposed in the paper can be summarised by two main points: (1) Lemma 2 establishes a sufficient condition to enforce the constraints of (UCCBP) by restricting the attention to each element of the partition \mathcal{Q} individually; (2) it allows us to use the duality results of Section 3.1 to obtain a computationally tractable reformulation of the optimisation problem in Lemma 2 into a robust LP. To illustrate the latter point, let $i \in I_0$ and consider its corresponding partition $Q_i = \{x : H_i x \leq h_i\}$.




Fig. 4. Given two regions Q_i, Q_j and a realisation of the noise ω , the set $Q_{ij}(\omega)$ represents the subset of $x \in Q_i$ such that $\hat{f}(x, \alpha) + \eta(\omega) \in Q_j$ for some $\alpha \in [0, 1]$. In other words, $Q_{ij}(\omega)$ is the subset of Q_i that can reach Q_j given the realisation of the noise ω . Unfortunately, $Q_{ij}(\omega)$ is not easily computed (the example above is an exception, see Fig. 5) and thus in Section 7.1, we will compute an over-approximation $\bar{Q}_{ij}(\omega)$.

Observe that $B(x, \theta) \leq \gamma$ for all $x \in Q_i$ can be written as $h_i^\top \lambda \leq v_i$ and $H_i^\top \lambda = -u_i$ hold where λ_i is a non-negative decision variable. This follows from the fact that within the region Q_i the function $B(x, \theta) = B_i(x, \theta)$ is affine and via strong duality (Prop. 2), we can rewrite this robust constraint to two regular linear constraints. The rewrite of the non-negativity, uniform upper bound, and unsafe set robust constraints (see Def. 2) to regular linear constraints follow a similar argument. Hence, for brevity, we omit the reformulation of these.

The remaining constraint $B(y + \eta(\omega), \theta) + \xi \leq B(x, \theta) + c$ for all $\omega \in D, x \in X_s, y \in F(x)$, requires more care, yet Lemma 2 also enables a computationally tractable reformulation of this. To this end, consider $(i, j) \in I_{X_s} \times I$ and $\omega \in D$. A challenge in reformulating the constraint as a linear constraint is that the set $Q_{ij}(\omega)$ is not a polyhedron or even convex (see Fig. 5). For now, we assume access to a polyhedral over-approximation $Q_{ij}(\omega) \subset \bar{Q}_{ij}(\omega) = \{x : H_{ij\omega} x \leq h_{ij\omega}\}$. Then if we impose the constraint for all $x \in \bar{Q}_{ij}(\omega)$, then it trivially follows that it also holds for all $x \in Q_{ij}(\omega)$. We will defer the discussion of how to compute $\bar{Q}_{ij}(\omega)$ to Section 7.1.

Proposition 6. *Let $(i, j) \in I_{X_s} \times I$ and $\omega \in D$ be given. Assume that there exists two affine functions such that $\underline{A}_i x + \underline{b}_i \leq F(x) \leq \bar{A}_i x + \bar{b}_i$ for all $x \in Q_i$. Then it holds that $B_j(y + \eta(\omega), \theta) + \xi \leq B_i(x, \theta) + c$ for all $x \in \bar{Q}_{ij}(\omega), y \in F(x)$ if and only if the following constraints*

hold

$$\begin{aligned} h_{ij\omega}^\top \underline{\lambda}_{ij\omega} &\leq v_i - v_j - u_j^\top (\underline{b}_i + \eta(\omega)) + c - \xi, \\ H_{ij\omega}^\top \underline{\lambda}_{ij\omega} &= \underline{A}_i^\top u_j - u_i, \\ h_{ij\omega}^\top \bar{\lambda}_{ij\omega} &\leq v_i - v_j - u_j^\top (\bar{b}_i + \eta(\omega)) + c - \xi, \\ H_{ij\omega}^\top \bar{\lambda}_{ij\omega} &= \bar{A}_i^\top u_j - u_i, \end{aligned}$$

where $\underline{\lambda}_{ij\omega}, \bar{\lambda}_{ij\omega}$ are a non-negative dual variables.

Collecting together all finite sets of constraints, a finite representation of the semi-infinite program (SBP) is given as

$$\begin{aligned} \min_z \quad & \gamma + cK \\ \text{s. t.} \quad & \gamma \geq 0, c \geq 0, \\ & \text{(Non-negativity)} \\ & h_i^\top \underline{v}_i \leq v_i, H_i^\top \underline{v}_i = -u_i, \text{ for all } i \in I, \\ & \text{(Uniform upper bound)} \\ & h_i^\top \bar{v}_i \leq M - v_i, H_i^\top \bar{v}_i = u_i, \text{ for all } i \in I, \\ & \text{(Initial set)} \\ & h_{i0}^\top \mu_i^0 \leq \gamma - v_i, H_{i0}^\top \mu_i^0 = u_i, \text{ for all } i \in I_{X_0}, \\ & \text{(Unsafe set)} \\ & h_i^\top \mu_i^u \leq v_i - 1, H_i^\top \mu_i^u = -u_i, \text{ for all } i \in I_{X_u}, \\ & \text{(One-step constraints)} \\ & h_{ij\omega}^\top \underline{\lambda}_{ij\omega} \leq v_i - v_j - u_j^\top (\underline{b}_i + \eta(\omega)) + c - \xi, \\ & H_{ij\omega}^\top \underline{\lambda}_{ij\omega} = \underline{A}_i^\top u_j - u_i, \\ & h_{ij\omega}^\top \bar{\lambda}_{ij\omega} \leq v_i - v_j - u_j^\top (\bar{b}_i + \eta(\omega)) + c - \xi, \\ & H_{ij\omega}^\top \bar{\lambda}_{ij\omega} = \bar{A}_i^\top u_j - u_i, \text{ for all } \omega \in D, \\ & \text{for all } (i, j) \in I_{X_s} \times I, \end{aligned} \tag{FSBP}$$

where $\underline{v}_i, \bar{v}_i, \mu_i^0, \mu_i^u, \underline{\lambda}_{ij\omega}, \bar{\lambda}_{ij\omega}$ are non-negative dual variables. (H_{i0}, h_{i0}) denotes the half-space representation of $Q_i \cap X_0$.

In Corollary 1, we put together the results we introduced and show how probabilistic safety can be computed via the scenario approach using samples of the random variable $\eta(\omega)$.

Corollary 1. *Assume that $D = \{\omega_1, \dots, \omega_N\}$ is a collection of N independent samples from the noise distribution \mathbb{P} . Fix $\epsilon \in (0, 1)$, $M \geq 1$, and $\xi \geq M \frac{\epsilon}{1-\epsilon}$, and let β be defined as Prop. 3 where $d = 2 + \ell(n + 1)$ is the number of (non-dual) decision variables in Problem (FSBP). Consider the optimal (primal) solution $z^*(D) = (c^*, \gamma^*, \theta^*)$ of Problem (FSBP). Then, with*

confidence $1 - \beta$, it holds that

$$\zeta(X_s, K) \geq 1 - (\gamma^* + c^*K).$$

Thus, by solving Problem (FSBP), we can certify probabilistic safety with high confidence. Note that naively trying to solve (FSBP) can soon become intractable on contemporary hardware, due to both memory requirements and computational time. In particular, the cardinality of the Cartesian product of I_{X_s} , I , and D can already be prohibitively large for relatively small systems. In the next section, we will discuss algorithmic strategies that make Problem (FSBP) computationally tractable.

7 Algorithms for program construction

In this section, we discuss three aspects that allows one to solve Problem (SBP) efficiently. Namely, in Section 7.1, we discuss how to compute a polyhedral over-approximation of $Q_{ij}(\omega)$. In Subsection 7.2, we introduce an a-priori sample discarding procedure for Problem (FSBP), which guarantees the same optimal solution, but allowing one to consider less samples in the optimisation problem. Finally, in Subsection 7.3, we employ spatial indexing methods to efficiently find the set of triplets with a non-empty $Q_{ij}(\omega)$. More specifically, we will rely on the fact that, often, for many triplets (i, j, ω) , the set $Q_{ij}(\omega)$ is empty, thus the martingale constraint is trivially satisfied. That is, $Q_{ij}(\omega)$ is empty if $\text{im}_i(Q_i, \omega) \cap Q_j = \emptyset$ where the image for region Q_i is defined as

$$\text{im}(Q_i, \omega) = \{y + \eta(\omega) : x \in Q_i, y \in F(x)\}. \quad (21)$$

7.1 Over-approximation of $Q_{ij}(\omega)$

As illustrated in Fig. 4, computing a polyhedral over-approximation of $Q_{ij}(\omega)$ is challenging, as $Q_{ij}(\omega)$ is (possibly) non-convex due to the uncertain⁵ affine transformation F . Furthermore, it is not sufficient to compute the convex hull for the vertices of the uncertainty variable α , that is,

$$\text{conv} \left(\left\{ x \in Q_i : \hat{f}_i(x, 0) + \eta(\omega) \in Q_j \right\} \cup \left\{ x \in Q_i : \hat{f}_i(x, 1) + \eta(\omega) \in Q_j \right\} \right), \quad (22)$$

as shown in Fig. 4. However, we note that by definition $Q_{ij}(\omega) \subseteq Q_i$, that is, Q_i is a, generally conservative, polyhedral over-approximation of $Q_{ij}(\omega)$. Hence,

⁵ If F is a deterministic affine transformation and Q_j is a polyhedron, then $Q_{ij}(\omega)$ is also a polyhedron and analytical methods for computation exist.

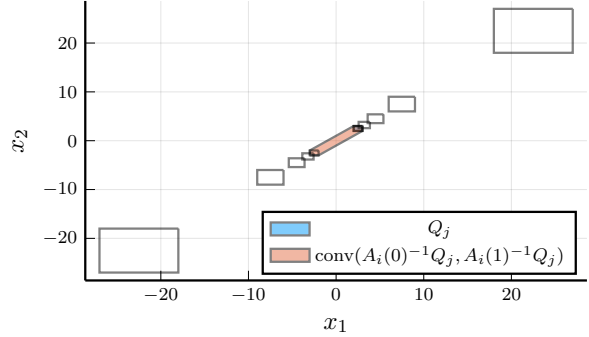


Fig. 5. A pictorial example that $Q_{ij}(\omega)$ is not necessarily convex. In this example, we have $\underline{A}_i = -I$, $\bar{A}_i = I$, and $\underline{b}_i = \bar{b}_i = 0$, which is a valid uncertain affine relaxation of the trivial function $f(x) = 0$ in the non-negative orthant. We consider the region $Q_j = [2, 3]^2$ and plot $A_i(\alpha)^{-1}Q_j$ for 10 different values of $\alpha \in [0, 1]$. Note that in this case $\bar{A}_i = I$, hence $Q_j = A_i(1)^{-1}Q_j$ and Q_j (blue) is contained in the convex hull (pink). $Q_{ij}(\omega)$ can take on complex shapes and no method exists for exactly computing $Q_{ij}(\omega)$. Therefore, we compute a sound over-approximation $\bar{Q}_{ij}(\omega) \subset Q_{ij}(\omega)$.

our approach to find a polyhedral over-approximation of $Q_{ij}(\omega)$, denoted $\bar{Q}_{ij}(\omega)$, is to start from Q_i and then iteratively removing subsets of $Q_i \setminus Q_{ij}(\omega)$. To accomplish this, we rely on repeated bisection. To simplify the presentation, in what follows, we assume that Q_i is a hyper-rectangle. Note, however, that the procedure generalises to compact polyhedra in half-space representation.

Fig. 6 shows an example of the bisection algorithm for two regions Q_i, Q_j and a given sample ω . The bisection is repeated twice along each axis, namely once to increase the lower bound, once to decrease the upper bound. Note that $Q_{ij}(\omega)$, although depicted in Fig. 6, is unknown and possibly non-convex, and further we cannot readily check if $Q_{ij}(\omega) \subset \bar{Q}_{ij}(\omega)$. However, recall that $Q_{ij}(\omega)$ is the subset of region Q_i that under a realisation of the noise ω reaches region Q_j in one time step. As a result, we can instead check if $\text{conv}(\bar{Q}_{ij}(\omega) + \eta(\omega))$ intersects with Q_j . By applying this algorithm, we compute a small over-approximation $\bar{Q}_{ij}(\omega)$ of $Q_{ij}(\omega)$. Indeed as $\bar{Q}_{ij}(\omega)$ is an over-approximation, using $\bar{Q}_{ij}(\omega)$ for the constraints in Proposition 6 yields a sound, although slightly conservative, solution.

7.2 Convex hull over the sample set

To reduce the number of constraints of Problem (FSBP), observe that the constraints of the problem are affine in $\eta(\omega)$. This implies that the active constraints, also known as support constraints, will always belong to the vertices of the convex hull over $\eta(D) := \{\eta(\omega_1), \dots, \eta(\omega_N)\}$, enabling a great reduction in the number of constraints in the program. That is, we can impose the constraint only on the following set

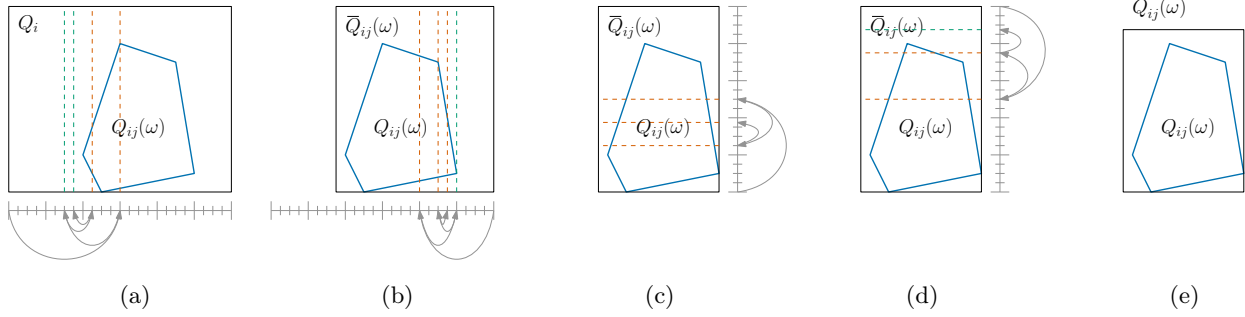


Fig. 6. An example of the bisection algorithm to compute the over-approximation $\bar{Q}_{ij}(\omega)$ of $Q_{ij}(\omega)$. The set $Q_{ij}(\omega)$ is unknown and possibly non-polyhedral, but Q_i is a sound over-approximation. By bisection from either side (first the lower then the upper bound) along each axis we obtain a smaller over-approximation. We start by bisecting for x_1 ((a) and (b)) followed by x_2 ((c) and (d)). This results in the over-approximation $\bar{Q}_{ij}(\omega)$ in (e).

of samples:

$$\bar{D} = \{\omega \in D : \eta(\omega) \in \text{vert}(\text{conv}(\eta(D)))\}. \quad (23)$$

In the experiments conducted (see Section 8), we find that generally, in practice, the cardinality of \bar{D} is orders of magnitude lower than the cardinality of D . Thus, this can greatly improve the efficiency of our approach.

Remark 2. The method presented in this subsection was discovered independently, but is similar to the method presented in [37] with the exception of that our method requires an exact convex hull rather than an approximate convex hull. The proposed method for sample reduction works for any scenario program that is affine in the random variable $\eta(\omega)$. \square

7.3 Spatial indexing for intersection search

Constructing Problem (FSBP) efficiently is also a non-trivial problem due to memory limits. In fact, a naïve approach to construct the problem is to iterate over all triplets (i, j, ω) in $I_{X_s} \times I \times D$, check if $Q_{ij}(\omega) \neq \emptyset$, and add a set of constraints if the test is positive. This approach is only tractable for smaller problems as it has time complexity $\mathcal{O}(|\mathcal{Q}|^3)$. To reduce the complexity, we can exploit methods from database theory; namely spatial indexing, which is the structuring and querying of spatially distributed data, such as maps, with improved computational complexity [16].

To apply spatial indexing to our setting, we must first establish the data and query. It holds that $Q_{ij}(\omega) \neq \emptyset$ only if $\text{im}_i(Q_i, \omega) \cap Q_j \neq \emptyset$. Hence, if we search for regions Q_j that intersect with the image $\text{im}_i(Q_i, \omega)$, we find all pairs (i, j) such that $Q_{ij}(\omega) \neq \emptyset$. For spatial indexing, we focus on R-trees as it is a well-studied and widely available method [16]. The idea is to structure the data in a tree structure and at each node store a Minimum Bounding Rectangle (MBR) for the nodes below. Then,

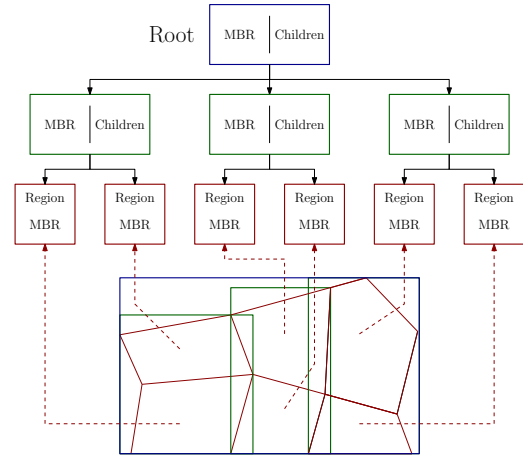


Fig. 7. An example of an R-tree applied to a partitioned state space to allow efficient search for regions intersecting with $\text{im}_i(Q_i, \omega)$. The rectangles are the Minimum Bounding Rectangle for each node in the tree.

querying the tree for the intersection with another region proceeds recursively down the tree, where it is only necessary to search down a branch if the MBR of the branch and the query intersect, which is an inexpensive operation by the separating hyperplane theorem [6]. Figure 7 shows an example of an R-tree for a partitioned state space. This method improves complexity by efficiently searching for relevant triplets (i, j, ω) .

In summary, to use the framework to compute data-driven safety certificates for non-linear systems: let the nominal dynamics f , initial and safe set X_0, X_s , a horizon T , and a dataset of samples D be given. Then start by abstracting the non-linear dynamics f to uncertain PWA dynamics \hat{f} using LBP techniques. Compute the vertices \bar{D} of the convex hull of D and discard all interior points. For each region Q_i , find, using an R-tree, regions that intersect with the image of the dynamics $\text{im}_i(Q_i, \omega)$ and add constraints accordingly. Solve the LP problem (FSBP), then the solution $z^*(\bar{D}) = (c^*, \gamma^*, \theta^*)$ is a safety certificate $\zeta(X_s, T) \geq 1 - (\gamma^* + c^*T)$ with

confidence $1 - \beta$ where β is defined as in Prop. 3.

8 Experiments

To support the theoretical results and investigate the efficacy of our approach, we implemented our framework in Julia⁶ and performed an empirical analysis on various benchmarks. The experiments have been conducted on a computer with an Intel i7-6700k CPU, Nvidia GTX1060 6GB GPU, and 16GB RAM, running Linux 5.10.211-1-MANJARO. We start by describing the benchmarks followed by the results. For a comparison with state-of-the-art, we consider the Sample Average Approximation (SAA) method proposed in [34], [35]. As this SBF synthesis method has been developed specifically for linear or polynomial systems, namely through SoS optimisation, in order to provide a general baseline, in the case of non-polynomial and/or uncertain systems, we combine it with the method proposed in [30], to find a valid SBF in the general case.

8.1 Benchmarks

The simplest system considered is an uncertain 1D linear system $x(k+1) = x(k) + b(\alpha) + \eta$ where $b(\alpha) = -0.05 + 0.1\alpha$, i.e. the uncertainty is in $b(\alpha)$ with uncertainty variable $\alpha \in [0, 1]$. Starting in a set around the origin $X_0 := \{|x| \leq 0.5\}$, the goal is to stay within a larger set $X_s := \{|x| \leq 2.5\}$ for a horizon $T = 10$. The distribution of the noise is a zero-mean normal distribution with standard deviation 0.01.

We also consider a 2D system from [3] representing the longitudinal dynamics of a drone. The coordinates x_1, x_2 are the position and velocity, respectively, and the system has the following dynamics

$$x(k+1) = \begin{pmatrix} 1 & \tau \\ 0 & 1 - \frac{0.1\tau}{m} \end{pmatrix} x(k) + \eta,$$

where $m \in [0.75, 1.25]$ and η has a zero-mean normal distribution with diagonal covariance of $[0.01 \ 0.01]$. The variable τ represents the discretisation step, which is set to $\tau = 1.0$. As with the 1D linear system, we certify safety for a horizon $T = 10$.

The third benchmark represents a model of a vehicle travelling down a straight road when it experiences an (uncertain) gust of wind. The coordinates x_1, x_2 represent, respectively, the longitudinal and lateral position of the vehicle. Similar to the drone system, τ represents

the discretisation step. The goal is certify the probability of staying on the road $X_s := \{|x_1| \leq 2.5\}$ for a horizon $T = 10$, when the system evolves according to the following dynamics

$$x(k+1) = \begin{pmatrix} 1 & 0 \\ 0 & 0.95 \end{pmatrix} x(k) + \begin{pmatrix} \frac{50}{3.6} \cdot \tau \\ \frac{1}{2} a_{lat} \cdot \tau^2 \end{pmatrix} + \eta,$$

where $\tau = 1$, and $a_{lat} = 0$ for regions where $x_1 \leq 80$ or $x_1 \geq 120$, and $a_{lat} \in [0.0913, 0.364]$ for regions where $x_1 \leq 80$ or $x_1 \geq 120$. η has a zero-mean normal distribution with diagonal covariance of $[0.01 \ 0.01]$.

While the previous models were linear, we also considered non-linear models. In particular, we consider Neural Network Dynamical Models (NNDMs) with 1 and 2 hidden layers of 64 neurons each modelling a pendulum taken from [30]. Finally, the last benchmark is the 3D model Dubin’s car from [27] for a time horizon $T = 10$. Dubin’s car is a unicycle model where the state is (x, y, ϕ) with ϕ being the heading of the vehicle. We consider a grid-based partitioning of 10 segments along each axis, i.e. 1000 regions. The noise is only applied to the last dimension and has a normal distribution with mean of $60 \cdot \frac{\pi}{180} \approx 1.053$ and standard deviation 0.1.

8.2 Results

Table 1 shows a list of results across all benchmarks. Both the safety probability and the computation time are reported as the mean over the 100 trials and for all cases the number of samples is selected to ensure a confidence $1 - \beta = 1 - 10^{-9}$. From Table 1 we observe that, depending on the system, the method can certify safety to $> 99\%$ with high confidence, e.g. 99.5% certified safety for the NNDM model of a pendulum with 2 layers and 64 neurons. This certification requires relatively few regions of 10-30 segments per axis. Comparing the NNDM pendulum model with 2 and 3 layers (1 and 2 hidden layers, respectively), the complexity of the nominal dynamics impacts both computation time and certifiable safety, e.g. 99.5% safety probability in 45.0s vs 97.6% safety probability in 78.5s for 480 regions, 2 and 3 layers respectively. This behavior can be explained by LBP computing wider uncertain affine transformations lead to more non-empty $Q_{ij}(\omega)$.

Remarkably, the linearity of the underlying system has little impact on the certifiable safety. This is observed in that both the 1D linear and drone systems exhibit uncertain linear behavior, yet the 1D linear system is certifiable to 50.8% safety while the drone is certifiable to 99.5%. Furthermore, the largest system considered, Dubin’s car, which includes trigonometric functions, safety is certified to 99.9%.

To compare against state-of-the-art, we report in Table 2

⁶ Code is available at <https://github.com/DAI-Lab-HERALD/scenario-barrier> under GNU GPLv3 license.

Table 1

Certified safety and computation time using the method explained in Sections 5-7. Results are reported as the mean over 100 iterations for each case study. n is the dimensionality of the system and ℓ is the number of regions. $\zeta(X_s, K)$ is the certified level of safety for $\beta = 10^{-9}$ where $1 - \beta$ is the level of confidence.

System	n	ℓ	$\zeta(X_s, K)$	Time (s)
Linear	1	27	0.508	0.239
Drone	2	37	0.995	41.8
Vehicle	2	18	0.607	0.716
		42	0.709	1.86
		54	0.827	2.20
		150	0.995	6.84
Pendulum (NNDM)	2	120	0.344	7.26
- 2 layers		240	0.756	17.6
- 64 neurons		480	0.995	45.0
Pendulum (NNDM)	2	120	0.254	14.3
- 3 layers		240	0.374	36.3
- 64 neurons		480	0.976	78.5
Dubin’s car		3	1000	0.999

Table 2

Comparison of our method against Sample Average Approximation (SAA) combined with the method presented in [30], to synthesise SBFs in a data-driven fashion for non-polynomial and uncertain systems. Results are reported as the mean over 100 iterations for each case study. $1 - \beta$ is the confidence in the certificate and $\zeta(X_s, K)$ is the certified level of safety. OOM means the certification procedure crashed with an out-of-memory error.

System	Method	β	$\zeta(X_s, K)$	Time (s)
Linear	Ours	10^{-2}	0.514	0.219
		10^{-3}	0.513	0.226
		10^{-4}	0.512	0.234
	SAA	10^{-2}	0.506	1.91
		10^{-3}	0.506	24.9
		10^{-4}	-	OOM
Pendulum (NNDM)	Ours	10^{-2}	0.995	42.0
		10^{-3}	0.995	42.0
		10^{-4}	0.995	42.3
	SAA	10^{-2}	0.903	16.5
		10^{-3}	0.903	47.3
		10^{-4}	-	OOM

the certified safety by using our method and using SAA, also as the mean over 100 trials. Because of the greater sample complexity of SAA having $\beta = 10^{-9}$ is not feasible. Hence, we compare for multiple, higher values of β to both make it tractable and find the limits of SAA. From the table, we observe that for our method certifying for orders of magnitude larger confidence (10^{-2} to 10^{-4}) negligibly increases the computation time (7% for the 1D linear system and 0.7% for the pendulum model)

and achieves similar levels of certified safety. In contrast, for SAA, going from 10^{-2} to 10^{-3} increases the computation time between 2.8x and 13x. Furthermore, for a confidence with $\beta = 10^{-4}$, the amount of memory required exceeds the 16GB available. The achieved level of certified safety is also marginally better with our proposed method (0.512 vs 0.506 for the 1D linear system and 0.995 vs 0.903 for the pendulum model).

9 Conclusion

We have presented a novel data-driven method to synthesise piece-wise affine Stochastic Barrier Function (SBF) based on a novel inner-approximation, which relies on the scenario theory. Our approach employs Linear Bound Propagation and stochastic approximations to guarantee that the search for a barrier reduces to solving a Linear Programming problem, which can be solved efficiently using the convex hull over the noise samples and spatial indexing for faster searching. As with any method, ours comes with limitations. In fact, in our approach we assume that the noise is additive and that we have i.i.d. full measurements of the state of the system. In particular, we rely on the additivity of the noise to achieve efficient algorithms in practice. How to extend our approach to non-additive noise represents an important open question. Another direction for future research to improve scalability is to consider more complex piece-wise templates for parameterising a SBF, e.g. piece-wise quadratic barrier candidates. Using more complex templates may require fewer pieces and thus improve synthesis performance. For more complex templates, one challenge is that strong duality, on which we have relied heavily, does not necessarily hold. Another direction of future work is to treat the safe control synthesis problem within the proposed framework, where a challenge is that the optimisation problem easily becomes bilinear.

References

- [1] Alessandro Abate. Formal verification of complex systems: Model-based and data-driven methods. In *MEMOCODE*, 2017.
- [2] Alessandro Abate, Maria Prandini, John Lygeros, and Shankar Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 2008.
- [3] Thom Badings, Licio Romao, Alessandro Abate, and Nils Jansen. Probabilities are not enough: Formal controller synthesis for stochastic dynamical models with epistemic uncertainty. In *AAAI*, 2023.
- [4] Thom Badings, Licio Romao, Alessandro Abate, David Parker, Hasan A Poonawala, Marielle Stoelinga, and Nils Jansen. Robust control for dynamical systems with non-gaussian noise via formal abstractions. *JAIR*, 2023.
- [5] Dimitri P Bertsekas and Steven E Shreve. *Stochastic optimal control: the discrete-time case*. Athena Scientific, 2004.

- [6] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- [7] Marco Campi and Simone Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 2008.
- [8] Nathalie Cauchi, Luca Laurenti, Morteza Lahijanian, Alessandro Abate, Marta Kwiatkowska, and Luca Cardelli. Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems. In *HSCC*, 2019.
- [9] Frank H Clarke. Generalized gradients of lipschitz functionals. *Advances in Mathematics*, 1981.
- [10] Ryan K. Cosner, Preston Culbertson, and Aaron D. Ames. Bounding stochastic safety: Leveraging freedman’s inequality with discrete-time control barrier functions, 2024.
- [11] Ryan K Cosner, Preston Culbertson, Andrew J Taylor, and Aaron D Ames. Robust safety under stochastic uncertainty with discrete-time control barrier functions, 2023.
- [12] Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI*, 2013.
- [13] Jared Lee Gearhart, Kristin Lynn Adair, Justin David Durfee, Katherine A Jones, Nathaniel Martin, and Richard Joseph Detry. Comparison of open-source linear programming solvers. Technical report, Sandia National Lab, 2013.
- [14] Antoine Girard and George Pappas. Approximation metrics for discrete and continuous systems. *IEEE TAC*, 2007.
- [15] Ibon Gracia, Dimitris Boskos, Luca Laurenti, and Manuel Mazo. Distributionally robust strategy synthesis for switched stochastic systems. *HSCC*, 2023.
- [16] Antonin Guttman. R-trees: A dynamic index structure for spatial searching. *SIGMOD Record*, 1984.
- [17] Sofie Haesaert and Sadegh Soudjani. Robust dynamic programming for temporal logic control of stochastic systems. *IEEE TAC*, 2021.
- [18] Q. Huangfu and J. A. J. Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 2017.
- [19] John Jackson, Luca Laurenti, Eric Frew, and Morteza Lahijanian. Safety Verification of Unknown Dynamical Systems via Gaussian Process Regression. *CDC*, 2020.
- [20] Pushpak Jagtap, Sadegh Soudjani, and Majid Zamani. Formal Synthesis of Stochastic Systems via Control Barrier Certificates. *IEEE TAC*, 2020.
- [21] Harold J Kushner. Stochastic stability and control. Technical report, Brown Univ Providence RI, 1967.
- [22] Luca Laurenti and Morteza Lahijanian. Unifying safety approaches for stochastic systems: From barrier functions to uncertain abstractions via dynamic programming, 2023.
- [23] Abolfazl Lavaei, Sadegh Soudjani, Alessandro Abate, and Majid Zamani. Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica*, 2022.
- [24] Abolfazl Lavaei, Sadegh Soudjani, Emilio Frazzoli, and Majid Zamani. Constructing MDP abstractions using data with formal guarantees. *IEEE Control Systems Letters*, 2023.
- [25] Benoît Legat, Robin Deits, Gustavo Goretkin, Twan Koolen, Joey Huchette, Daisuke Oyama, and Marcelo Forets. Juliapolyhedra/polyhedra.jl: v0.6.16, June 2021.
- [26] Scott C. Livingston, Richard M. Murray, and Joel W. Burdick. Backtracking temporal logic synthesis for uncertain environments. In *ICRA*, 2012.
- [27] Frederik Baymler Mathiesen, Simeon C. Calvert, and Luca Laurenti. Safety certification for stochastic systems via neural barrier functions. *IEEE Control Systems Letters*, 2023.
- [28] Frederik Baymler Mathiesen, Licio Romao, Simeon C. Calvert, Alessandro Abate, and Luca Laurenti. Inner approximations of stochastic programs for data-driven stochastic barrier function design. In *CDC*, 2023.
- [29] Rayan Mazouz, Frederik Baymler Mathiesen, Luca Laurenti, and Morteza Lahijanian. Piecewise stochastic barrier functions, 2024.
- [30] Rayan Mazouz, Karan Muvvala, Akash Ratheesh Babu, Luca Laurenti, and Morteza Lahijanian. Safety guarantees for neural network dynamic systems via stochastic barrier functions. In *NeurIPS*, 2022.
- [31] Robert D McAllister and James B Rawlings. Stochastic lyapunov functions and asymptotic stability in probability. Technical report, Texas - Wisconsin - California Control Consortium, 2020.
- [32] Stephen Prajna, Ali Jadbabaie, and George Pappas. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE TAC*, 2007.
- [33] Hamed Rahimian and Sanjay Mehrotra. Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization*, 2022.
- [34] Ali Salamati, Abolfazl Lavaei, Sadegh Soudjani, and Majid Zamani. Data-driven safety verification of stochastic systems via barrier certificates. *IFAC-PapersOnLine*, 2021.
- [35] Ali Salamati and Majid Zamani. Safety verification of stochastic systems: A repetitive scenario approach. *IEEE Control Systems Letters*, 2023.
- [36] Cesar Santoyo, Maxence Dutreix, and Samuel Coogan. A barrier function approach to finite-time stochastic system verification and control. *Automatica*, 2021.
- [37] Hossein Sartipizadeh and Behcet Acikmese. Approximate convex hull based sample truncation for scenario approach to chance constrained trajectory optimization. In *ACC*, 2018.
- [38] Oliver Schön, Zhengang Zhong, and Sadegh Soudjani. Data-driven distributionally robust safety verification using barrier certificates and conditional mean embeddings, 2024.
- [39] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. On a formal model of safe and scalable self-driving cars, 2018.
- [40] Jacob Steinhardt and Russ Tedrake. Finite-time regional verification of stochastic non-linear systems. *The International Journal of Robotics Research*, 2012.
- [41] T. Tao. *An Introduction to Measure Theory*. American Mathematical Society, 2013.
- [42] Birgit Van Huijgevoort, Oliver Schön, Sadegh Soudjani, and Sofie Haesaert. Syscore: Synthesis via stochastic coupling relations. In *HSCC*, 2023.
- [43] Kaidi Xu, Zhouxing Shi, Huan Zhang, Yihan Wang, Kai-Wei Chang, Minlie Huang, Bhavya Kaikhura, Xue Lin, and Cho-Jui Hsieh. Automatic perturbation analysis for scalable certified robustness and beyond. In *NeurIPS*, 2020.
- [44] Đorđe Žikelić, Mathias Lechner, Thomas A Henzinger, and Krishnendu Chatterjee. Learning control policies for stochastic systems with reach-avoid guarantees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 11926–11935, 2023.

A Proof and auxiliary results

A.1 Proof of Proposition 2

Proof. Let the sets \mathcal{Z} and

$$\mathcal{Z} = \{z : (Az + a)^\top x \leq Bz + b, \text{ for all } x \in P\}, \quad (\text{A.1})$$

$$\mathcal{Z}'' = \{(z, \lambda) : h^\top \lambda \leq Bz + b, H^\top \lambda = Az + a, \lambda \geq 0\}, \quad (\text{A.2})$$

be the feasible sets of (4) and (5), respectively. Our goal is to show that $\mathcal{Z} = \text{proj}_z(\mathcal{Z}'') = \mathcal{Z}'$, where

$$\mathcal{Z}' = \{z : \exists \lambda \in \mathbb{R}_{\geq 0}^p, h^\top \lambda \leq Bz + b, H^\top \lambda = (Az + a)\}.$$

is the projection of the set \mathcal{Z}'' onto its first d coordinates. At the core of this result is the strong duality [6, Section 5.2.1] between the linear programs

$$\begin{aligned} \max_x \quad & (Az + a)^\top x & \min_\lambda \quad & \lambda^\top h \\ \text{s. t.} \quad & Hx \leq h, & \text{s. t.} \quad & H^\top \lambda = (Az + a), \quad \lambda \geq 0, \end{aligned}$$

which states that for any element x_* in the optimal set of the maximisation problem there exist a λ_* in the optimal set of the minimisation problem such that $(Az + a)^\top x_* = \lambda_*^\top h$ and $\lambda_*^\top (H^\top x_* - h) = 0$. Vice-versa, for all λ_* in the optimal set of the minimisation problem there exists a x_* in the optimal set of the maximisation problem such that similar conclusions hold.

Pick any element $\bar{z} \in \mathcal{Z}$. Let \bar{x} be such that $(A\bar{z} + a)^\top \bar{x} = \sup_{x \in P} a(\bar{z})^\top x$. Hence, by strong duality, there exists a $\bar{\lambda} \geq 0$ with $H^\top \bar{\lambda} = (A\bar{z} + a)$ such that $\bar{\lambda}^\top h = (A\bar{z} + a)^\top \bar{x} \leq B\bar{z} + b$. In other words, there is a $\bar{\lambda}$ such that $(\bar{z}, \bar{\lambda}) \in \mathcal{Z}''$, which implies $\mathcal{Z} \subseteq \mathcal{Z}'$.

For the other direction, consider a tuple $(\bar{z}, \bar{\lambda}) \in \mathcal{Z}''$. For this given \bar{z} , pick

$$\lambda_*(\bar{z}) \in \arg \min_{H^\top \lambda = A\bar{z} + a, \lambda \geq 0} \lambda^\top h,$$

we notice that we have that $\bar{z} \in \mathcal{Z}'$ and

$$\sup_{x \in P} (A\bar{z} + a)^\top x = \lambda_*(\bar{z})^\top h \leq \bar{\lambda}^\top h \leq B\bar{z} + b.$$

The right-most inequality follows from feasibility of $(\bar{z}, \bar{\lambda})$, the middle inequality by our choice of $\lambda_*(\bar{z})$, the left-most equality by strong duality of linear programming. Then we conclude that $\bar{z} \in \mathcal{Z}$, thus concluding the proof of the proposition. \square

A.2 Proof of Proposition 4

Proof. Fix a dimension $j \in \{0, \dots, n\}$ and a convex region P_i . Then due to the mean value theorem for generalised gradients [9], there exists two hyperplanes $(\underline{A}_i x + \underline{b}_i)_j, (\bar{A}_i x + \bar{b}_i)_j$ such that it holds

$$(\underline{A}_i x + \underline{b}_i)_j \leq f(x)_j \leq (\bar{A}_i x + \bar{b}_i)_j$$

for all $x \in Q_i$. Combining all dimensions, it holds that

$$\underline{A}_i x + \underline{b}_i \leq f(x) \leq (\bar{A}_i x + \bar{b}_i)$$

for all $x \in Q_i$ for each region Q_i . As a result, by the definition of \hat{f} it holds that $f(x) \in \{\hat{f}(x, \alpha) : \alpha \in [0, 1]\}$ for all $x \in X$ concluding the proof. \square

A.3 Measurability issue of Theorem 1

We need to show that set

$$E = \{\omega \in \Omega : B(f(x) + \eta(\omega)) + \xi \leq B(x) + c, \text{ for all } x \in X_s\}, \quad (\text{A.3})$$

is Borel measurable. First of all, notice that due to the fact that $\{Q_1, \dots, Q_\ell\}$ is a finite partition of X_s , we have that

$$E = \bigcap_{Q_i \in X_s} \{\omega \in \Omega : \sup_{x \in Q_i} B(f(x) + \eta(\omega)) - B_i(x) \leq c - \xi\}. \quad (\text{A.4})$$

As finite intersections of measurable sets are still measurable and a measurable function maps measurable sets into measurable sets, it is enough to show that each set

$$E_i = \{\eta \in \mathbb{R}^n : \sup_{x \in Q_i} B(f(x) + \eta) - B_i(x) \leq c - \xi\}$$

is measurable. In order to do that note that when restricted to Q_i , B_i is a linear function, while by construction $B(x)$ is upper semi-continuous. Furthermore, as composition of an upper semi-continuous function with a continuous one is still upper semi-continuous, we have that both $B(f(x) + \eta)$ and $B(f(x) + \eta) - B_i(x)$ are upper semi-continuous functions. Consequently, by Proposition 7.32 in [5] we have that $g_i(\eta) = \sup_{x \in Q_i} B(f(x) + \eta) - B_i(x)$ is upper semi-continuous. As $g_i(\eta)$ is upper semi-continuous, hence Borel measurable, we have that set E_i is Borel measurable, thus concluding the proof.

A.4 Proof for Proposition 6

Proof. Start by fixing $\alpha \in [0, 1]$, the pair $(i, j) \in I_{X_s} \times I$, and the noise sample $\omega \in D$. Then the constraint $B(\hat{f}(x, \alpha) + \eta(\omega), \theta) + \xi \leq B(x, \theta) + c$ for all $x \in X_s$ can

be rewritten as

$$B_j(A_i(\alpha)x + b_i(\alpha) + \eta(\omega), \theta) + \xi \leq B_i(x, \theta) + c,$$

for all $x \in \bar{Q}_{ij}(\omega)$. Using Prop. 2, we rewrite this constraint into a dual formulation so that we can solve it in a lifted space. Thus, they become the following

$$\begin{aligned} h_{ij\omega}^\top \lambda_{ij\omega} &\leq v_i - v_j - u_j^\top (b_i(\alpha) + \eta(\omega)) + c - \xi, \\ H_{ij\omega}^\top \lambda_{ij\omega} &= A_i(\alpha)^\top u_j - u_i, \end{aligned}$$

where $\lambda_{ij\omega}$ is a non-negative dual variable. Since both constraints are affine in α , they will hold for all $\alpha \in [0, 1]$ if and only if they hold for the vertices, that is, $\alpha \in \{0, 1\}$. We conclude the proof by observing that for each noise sample $\omega \in D$, the union of $\bar{Q}_{ij}(\omega)$ for all $(i, j) \in I_{X_s} \times I$ is a superset of X_s . \square

B Algorithm for computing over-approximation of $Q_{ij}(\omega)$

Our approach to computing an over-approximation of $Q_{ij}(\omega)$ is detailed in Alg. 1, where we start with a hyperrectangle $\bar{Q}_{ij}(\omega) = Q_i = \{x \in \mathbb{R}^m : l \leq x \leq u\}$, with \leq interpreted element-wise. The idea to reduce the size of $\bar{Q}_{ij}(\omega)$ is to increase l (Line 5-14) and decrease u (Line 16-25) while maintaining the over-approximation of $Q_{ij}(\omega)$. Treating the axes sequentially (Line 3), denoting the current axis by k , we use bisection first to find the largest l_k and then the smallest u_k such that $\bar{Q}_{ij}(\omega) \cap \{x : l_k \leq x_k \leq u_k\}$ is an over-approximation of $Q_{ij}(\omega)$. Then we can replace $\bar{Q}_{ij}(\omega)$ with $\bar{Q}_{ij}(\omega) \cap \{x : l_k \leq x_k \leq u_k\}$ for the next axis (Line 27).

To perform the bisection for increasing l_k , we compute the midpoint c_k^l between l_k and u_k (Line 7), denoting them l_k^l and u_k^l respectively, and test if $Q_{ij}(\omega) \subset \bar{Q}_{ij}(\omega) \cap \{x : c_k^l \leq x_k\}$. If true, then we may let $l_k^l = c_k^l$ (Line 10), and if not, let $u_k^l = c_k^l$ (Line 12). This procedure repeats for a fixed number of iterations and is performed mutatis mutandis to decrease u_k . The question remains how to check if $Q_{ij}(\omega) \subset \bar{Q}_{ij}(\omega) \cap \{x : c_k^l \leq x_k\}$, since $Q_{ij}(\omega)$ is unknown. To this end, recall that $Q_{ij}(\omega)$ is the subset of region Q_i that under a realisation of the noise ω reaches region Q_j in one time step. Thus, if the image of the other subregion $\text{im}(\bar{Q}_{ij}(\omega) \cap \{x : x_k \leq c_k^l\}, \omega)$ under the realisation of the noise ω does not intersect Q_j (Line 9), then it necessarily holds that $Q_{ij}(\omega) \subset \bar{Q}_{ij}(\omega) \cap \{x : c_k^l \leq x_k\}$.

Algorithm 1 Bisection-based algorithm for computing a subset $\bar{Q}_{ij}(\omega)$ of region Q_i as an over-approximation of $Q_{ij}(\omega)$.

```

1: function POLYPREIMAGE( $Q_i, Q_j, \hat{f}, \omega, t$ )
2:    $\bar{Q}_{ij}(\omega) \leftarrow Q_i$ 
3:   for  $k \leftarrow 1$  to  $m$  do ▷ For each axis
4:
5:      $l_k^l, u_k^l \leftarrow l_k, u_k$  ▷ Increase lower bound
6:     for  $s \leftarrow 1$  to  $t$  do
7:        $c_k^l \leftarrow \frac{l_k^l + u_k^l}{2}$ 
8:        $\bar{Q}_{ij}(\omega)' \leftarrow \bar{Q}_{ij}(\omega) \cap \{x : x_k \leq c_k^l\}$ 
9:       if  $\text{im}_i(\bar{Q}_{ij}(\omega)', \omega) \cap Q_j = \emptyset$  then
10:         $l_k^l \leftarrow c_k^l$ 
11:       else
12:         $u_k^l \leftarrow c_k^l$ 
13:       end if
14:     end for
15:
16:      $l_k^u, u_k^u \leftarrow l_k, u_k$  ▷ Decrease upper bound
17:     for  $s \leftarrow 1$  to  $t$  do
18:        $c_k^u \leftarrow \frac{l_k^u + u_k^u}{2}$ 
19:        $\bar{Q}_{ij}(\omega)' \leftarrow \bar{Q}_{ij}(\omega) \cap \{x : x_k \geq c_k^u\}$ 
20:       if  $\text{im}_i(\bar{Q}_{ij}(\omega)', \omega) \cap Q_j = \emptyset$  then
21:         $u_k^u \leftarrow c_k^u$ 
22:       else
23:         $l_k^u \leftarrow c_k^u$ 
24:       end if
25:     end for
26:
27:      $\bar{Q}_{ij}(\omega) \leftarrow \bar{Q}_{ij}(\omega) \cap \{x : l_k^l \leq x_k \leq u_k^u\}$ 
28:   end for
29:   return  $\bar{Q}_{ij}(\omega)$  ▷ It holds that  $Q_{ij}(\omega) \subset \bar{Q}_{ij}(\omega)$ 
30: end function

```
