

QUANTITATIVE VERIFICATION WITH NEURAL NETWORKS

ALESSANDRO ABATE ^a, ALEC EDWARDS , MIRCO GIACOBBE ^b, HASHAN PUNCHIHEWA,
AND DIPTARKO ROY ^b

^a Department of Computer Science, University of Oxford, UK
e-mail address: alessandro.abate@cs.ox.ac.uk

^b School of Computer Science, University of Birmingham, UK
e-mail address: m.giacobbe@bham.ac.uk, d.s.roy@bham.ac.uk

ABSTRACT. We present a data-driven approach to the quantitative verification of probabilistic programs and stochastic dynamical models. Our approach leverages neural networks to compute tight and sound bounds for the probability that a stochastic process hits a target condition within finite time. This problem subsumes a variety of quantitative verification questions, from the reachability and safety analysis of discrete-time stochastic dynamical models, to the study of assertion-violation and termination analysis of probabilistic programs. We rely on neural networks to represent supermartingale certificates that yield such probability bounds, which we compute using a counterexample-guided inductive synthesis loop: we train the neural certificate while tightening the probability bound over samples of the state space using stochastic optimisation, and then we formally check the certificate’s validity over every possible state using satisfiability modulo theories; if we receive a counterexample, we add it to our set of samples and repeat the loop until validity is confirmed. We demonstrate on a diverse set of benchmarks that, thanks to the expressive power of neural networks, our method yields smaller or comparable probability bounds than existing symbolic methods in all cases, and that our approach succeeds on models that are entirely beyond the reach of such alternative techniques.

1. INTRODUCTION

Probabilistic programs extend imperative programs with the ability to sample from probability distributions [DS20, GHNR14, Koz81, MM05], which provides an expressive language to describe randomized algorithms, cryptographic protocols, and Bayesian inference schemes. Discrete-time stochastic dynamical models, characterised by stochastic difference equations, are a natural framework to describe auto-regressive time series, as well as sequential decision and planning problems in unknown environments. A fundamental quantitative verification problem for probabilistic programs and stochastic dynamical models is the quantitative

Key words and phrases: Data-driven Verification, Quantitative Verification, Probabilistic Programs, Stochastic Dynamical Models, Counterexample-guided Inductive Synthesis, Neural Networks.

* The conference version of this manuscript appeared at CONCUR 2023 [AEG⁺23].

This work was supported in part by Innovate UK under the HICLASS project, EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems (EP/S024050/1) and EPSRC Doctoral Training Partnership, Google DeepMind Scholarship, Advanced Research and Invention Agency under the Safeguarded AI programme.

reachability question, which amounts to finding the probability with which the system reaches a given target condition within a finite number of steps. Reachability is at the core of a variety of other important quantitative verification questions as, by selecting appropriate target conditions on the state space, we can express the probability that a probabilistic program terminates or that it violates an assertion, as well as the probability that a stochastic dynamical model satisfies an invariant or remains within a set of safe configurations.

Quantitative reachability verification has been studied extensively using theories and algorithms built upon symbolic reasoning techniques, such as quantitative calculi [MM02, MMS96], probabilistic model checking algorithms [KNP05, FKN⁺11], discrete abstractions of stochastic dynamical models [AKM11, TMKA17, SA12], and the synthesis of supermartingale-like certificates [CS13, CFG16, CNv17, CFNH18, LSAZ22]. Among the latter class, a method to provide a sound upper bound for the reachability probability of a system is to synthesise a supermartingale function that maps every reachable state to a non-negative real, whose value is never smaller than 1 inside the target condition, and such that it never increases in expectation as the system evolves outside of the target. This is referred to as a *non-negative repulsing supermartingale* or *stochastic invariant indicator* in the literature [TOUH21, CGMv22, CNv17], and its output over a given state provides an upper bound for the probability that the system reaches the target condition from that state. Symbolic methods for the synthesis of such certificates assume that the supermartingale function, as well as its post-expectation, which depends on the model constraints and distributions, are both in linear or polynomial form. This poses syntactic restrictions to their applicability.

Data-driven and counterexample-guided inductive synthesis (CEGIS) procedures, combined with machine learning approaches that leverage neural networks to represent certificates, have shown great promise in mitigating the aforementioned limitation [CRG19, AAGP21, GKP22, PAA21, MCL23, AGR21, LvCH22, vLHC23, vLV⁺23, GKPT24, NMTZ24, NGL25]. In particular, neural-based CEGIS decouples the task of guessing a certificate from that of checking its validity, delegating the guessing task to efficient machine learning algorithms that leverage the expressive power of neural networks, while confining symbolic reasoning to the checking part of the task, which is computationally easier to solve in isolation than the entire synthesis problem. CEGIS has been applied to the synthesis of *neural supermartingales* for the almost-sure termination analysis of probabilistic programs [AGR21], as well as their counterpart for stochastic dynamical models with applications to qualitative queries such as almost-sure safety and stability [MCL23, LvCH22].

In this paper, we present theory, methods, and an extensive experimental evaluation to demonstrate the efficacy and flexibility of neural supermartingales to solve *quantitative verification* questions for probabilistic programs and stochastic dynamical models, using machine learning combined with satisfiability modulo theories (SMT) technologies.

Theory: We use the theory of non-negative repulsing supermartingales [TOUH21, Section 4] to leverage neural networks as representations of supermartingale functions for quantitative verification, while exploiting SMT solving to guarantee soundness.

Methods: We present a CEGIS-based approach to train neural supermartingale functions that minimise an upper bound for the reachability probability over sample points from the state space, and check their validity over every possible state using SMT solving. We present a program-agnostic approach that relies on state samples in the training

$v \in \text{Vars}$	(variables)
$N \in \mathbb{R}$	(numerals)
$E ::= v \mid N \mid -E \mid E + E \mid E - E \mid E * E \mid \dots$	(arithmetic expressions)
$P ::= \text{Bernoulli}(E) \mid \text{Gaussian}(E, E) \mid \dots$	(probability distributions)
$B ::= \text{true} \mid !B \mid B \&\& B \mid B \mid\mid B \mid E == E \mid E < E \mid \dots$	(Boolean expressions)
$C ::= \text{skip}$	(update commands)
$\mid v = E$	(deterministic assignment)
$\mid v \sim P$	(probabilistic assignment)
$\mid C ; C$	(sequential composition)
$\mid \text{if } B \text{ then } C \text{ else } C \text{ fi}$	(conditional composition)

Figure 1: Grammar for update commands, Boolean and arithmetic expressions.

phase, and also a novel program-aware approach that embeds model information in the loss function to enhance the effectiveness of stochastic optimisation.

Experiments: We build a prototype implementation and compare the efficacy of our method with the state of the art in synthesis of linear supermartingales using symbolic reasoning [TOUH21]. We show that our program-aware approach computes tighter or comparable probability bounds than symbolic reasoning on existing benchmarks, while our program-agnostic approach matches it in over half of the instances. We additionally demonstrate that both our approaches can handle models beyond reach of purely symbolic methods.

2. PROBABILISTIC PROGRAMS AND STOCHASTIC DYNAMICAL MODELS

Probabilistic programs are computer programs whose execution is determined by random variables, and stochastic dynamical models describe discrete-time dynamical systems with probabilistic behaviour. The former enjoy the flexibility of imperative programming constructs and are used to describe randomised algorithms, and the latter are expressed as stochastic difference equations and are used to describe probabilistic systems that evolve over infinite time. The semantics of both can be described in terms of stochastic processes and, for this reason, verification questions for both can be solved with similar techniques.

The syntax of our modeling framework uses imperative constructs from probabilistic programs and defines executions over infinite time as dynamical systems. Specifically, we consider programs that operate over an ordered set of n real-valued variables, denoted by Vars , and update their values through the repeated execution of a command C whose grammar is described in Figure 1. Under this definition, a state of the system is an n -dimensional vector $s \in \mathbb{R}^n$ that assigns a value to each variable symbol. The update command C defines an update function $f: \mathbb{R}^n \times [0, 1]^m \rightarrow \mathbb{R}^n$, where m is the number of syntactic probabilistic assignment statements occurring in C , which maps the current state and m random variables uniformly distributed in $[0, 1]$ into the next state. Conceptually, within f , each random variable is mapped into its respective distribution by applying the appropriate inverse transformation. Altogether, our probabilistic program defines a stochastic process,

whose behaviour is determined by the following stochastic difference equation:

$$s_{t+1} = f(s_t, r_t), \quad r_t \sim \mathbb{U}^m, \quad (2.1)$$

where r_t is an m -dimensional random input sampled at time t from the uniform distribution \mathbb{U}^m over the m -dimensional hypercube $[0, 1]^m$. The initial state s_0 is either given as a deterministic assignment to constant values, or is non-deterministically chosen from a set of initial conditions S_0 characterized by a Boolean expression. This setting can be seen as an assertion about the initial conditions followed by the probabilistic program `while true do C od`. As we show in Section 3, this allows us to characterise verification questions such as termination, non-termination, invariance and assertion-violation for while loops with general guards, as well as reachability and safety verification questions for dynamical models with general stochastic disturbances.

The semantics of our model is defined as a stochastic process induced by the Markov chain over the probability space of infinite words of random samples. This is defined by the probability space triple $(\Omega, \mathcal{F}, \mathbb{P})$ [BS96, HLL96], where

- Ω is the set of infinite sequences $([0, 1]^m)^\omega$ of m -dimensional tuples of values in $[0, 1]$,
- \mathcal{F} is the extension of the Borel σ -algebra over the unit interval $\mathcal{B}([0, 1])$ to Ω ,
- \mathbb{P} is the extension of the Lebesgue measure on $[0, 1]$ to Ω .

Every initialisation of the system on state $s \in \mathbb{R}^n$ induces a stochastic process $\{X_t^s\}_{t \in \mathbb{N}}$ over the state space \mathbb{R}^n . Let $\omega = r_0 r_1 r_2 \dots$ be an infinite sequence of random samples in $[0, 1]^m$, then the stochastic process is defined by the sequence of random variables

$$X_{t+1}^s(\omega) = f(X_t^s(\omega), r_t), \quad X_0^s(\omega) = s. \quad (2.2)$$

This defines the natural filtration $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$, which is the smallest filtration to which the stochastic process X_t^s is adapted. In other words [Kal06], this can be seen as another Markov chain with state space \mathbb{R}^n and transition kernel

$$T(s, S') = \text{Leb}(\{r \in [0, 1]^m \mid f(s, r) \in S'\}), \quad (2.3)$$

where S' is a Borel measurable subset of \mathbb{R}^n and Leb refers to the Lebesgue measure of a measurable subset of $[0, 1]^m$. In other words, kernel T denotes the probability to transition from state s into a set of states S' . The transition kernel also defines the *post-expectation operator* \mathbb{X} , also known as the next-time operator [TOUH21, Definition 2.16]. \mathbb{X} can be applied to an arbitrary Borel-measurable function $h: \mathbb{R}^n \rightarrow \mathbb{R}$ defining the *post-expectation of h* , denoted by $\mathbb{X}[h]$ and defined as the following function over states:

$$\mathbb{X}[h](s) = \int h(s') T(s, ds'). \quad (2.4)$$

This represents the expected value of h evaluated at the next state, given the current state being s . Computing the symbolic representation of a post-expectation for probabilistic programs and stochastic dynamical models is a core problem in probabilistic verification [GMV16, GSV20]. Indeed, our theoretical framework builds upon the post-expectation (cf. Eq. (5.1b)), and our verification procedure uses a symbolic representation of the post-expectation (cf. Section 6.2) as well as our program-aware learning procedure, while our program-agnostic learning procedure approximates it statistically (cf. Section 6.1).

We remark that our model encompasses general probabilistic program loops as well as stochastic dynamical systems with general disturbances. For example, a probabilistic loop

with guard condition B and body C in the form

$$\text{while } B \text{ do } C \text{ od} \quad (2.5)$$

can be expressed as a loop with guard `true` and body `if B then C else skip fi`. This expresses the fact that, after termination, the program will stay on the terminal state indefinitely. Also, discrete-time stochastic difference equations with a nonlinear vector field g and time-invariant input disturbance with an arbitrary distribution \mathcal{W} in the form

$$s_{t+1} = g(s_t, w_t), \quad w_t \sim \mathcal{W}, \quad (2.6)$$

comply with our model. It is sufficient to derive w_t with an appropriate inverse transformation from the uniform distribution and embed it in the update function f (Eq. (2.1)), to simulate a function g that accepts random inputs w_t distributed according to distribution \mathcal{W} . Note that this is without loss of generality; we do not require f to be continuous and we do not require it to be differentiable, and therefore we can represent any distribution using the inverse CDF method, as is standard in stochastic simulation [AG07, Chapter 2, p.36]. Our model is even more general, as it encompasses state-dependent distributions, whose parameters depend on the state and may depend on other distributions and thus define joint, multi-variate and hierarchically-structured distributions. Notably, our model comprises both continuous and discrete probability distributions and is able to model discrete-time stochastic hybrid systems [LSAZ22]. Our model also subsumes probabilistic models and programs over finite state spaces, where each state can be represented as a point of the Euclidean space \mathbb{R}^n . Note that this encoding of a finite state space ensures Borel-measurability of the update function, as the σ -algebra of this finite state space (i.e., its powerset) is a subset of the Borel σ -algebra.

3. QUANTITATIVE REACHABILITY VERIFICATION OF PROBABILISTIC MODELS

Quantitative verification treats the question of providing a quantity (probability) with which a system satisfies a property, as opposed to providing a definite (positive or negative) answer. In fact, it is sometimes too conservative or inappropriate to demand a definite outcome to a formal verification question. For instance, a system whose behaviour is probabilistic may violate a specification on rare corner cases and yet satisfy it with a probability that is deemed acceptable for the application domain. We address the quantitative verification of probabilistic systems, which is the problem of computing the probability for which a system satisfies a specification [Kwi07, vBW01, BHHK10, BdAFK18].

We consider the *quantitative reachability verification* question, which as we show below, is at the core of a variety of quantitative verification problems for probabilistic programs and stochastic dynamical models. Henceforth, we use $\mathbf{1}_S$ to denote the indicator function of set S , i.e., $\mathbf{1}_S(s) = 1$ if $s \in S$ and $\mathbf{1}_S(s) = 0$ if $s \notin S$; we also use $\lambda x.M$ to denote the anonymous function that takes an argument x and evaluates the expression M to produce its result.

We now provide a recursive characterisation for the probability that a stochastic process reaches a Borel measurable target set in exactly time t (Lemma 3.1), and in at most time t (Lemma 3.2). Finally, in Lemma 3.3 we employ these results to show that the probability of reaching a target set in *any* finite time is the limit of the time-bounded reachability probability as the time bound goes to infinity.

Lemma 3.1. *Let the event that a stochastic process $\{X_t^s\}_{t \in \mathbb{N}}$ over state space \mathbb{R}^n initialised in state $s \in \mathbb{R}^n$ reaches a target set $A \in \mathcal{B}(\mathbb{R}^n)$ in exactly time $t \in \mathbb{N}$ be*

$$\text{Reach}_t^s(A) = \{\omega \in \Omega \mid X_0^s(\omega) \notin A, \dots, X_{t-1}^s(\omega) \notin A, X_t^s(\omega) \in A\}, \quad (3.1)$$

with $\text{Reach}_0^s(A) = \Omega$ if $s \in A$, and $\text{Reach}_0^s(A) = \emptyset$ if $s \notin A$. Then, $\text{Reach}_t^s(A)$ is measurable and its probability measure can be expressed as follows:

$$\mathbb{P}[\text{Reach}_{t+1}^s(A)] = \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\lambda s'. \mathbb{P}[\text{Reach}_t^{s'}(A)]](s), \quad (3.2a)$$

$$\mathbb{P}[\text{Reach}_0^s(A)] = \mathbf{1}_A(s). \quad (3.2b)$$

Proof. For every i we have that the pre-image of any measurable set $R \in \mathcal{B}(\mathbb{R}^n)$ under $X_i^s : (\Omega, \mathcal{F}) \rightarrow (\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$, denoted $(X_i^s)^{-1}[R]$, is measurable since X_i^s is a measurable function [MTG09, Section 3.4]. It follows that $\text{Reach}_t^s(A)$ is measurable because it can be expressed as the following finite intersection:

$$\text{Reach}_t^s(A) = \bigcap \{(X_i^s)^{-1}[\mathbb{R}^n \setminus A] : i = 0, \dots, t-1\} \cap (X_t^s)^{-1}[A].$$

It is a standard result (see [MTG09, Theorem 3.4.1] and [Rev75, Theorem 2.8]) that the probability measure of the event $\text{Reach}_t^{s_0}(A)$ with $t > 0$ can be expressed as an integral with respect to the transition kernel of the Markov chain:

$$\begin{aligned} \mathbb{P}[\text{Reach}_t^{s_0}(A)] &= \mathbf{1}_{\mathbb{R}^n \setminus A}(s_0) \int_{s_1 \in \mathbb{R}^n \setminus A} \cdots \int_{s_{t-1} \in \mathbb{R}^n \setminus A} \int_{s_t \in A} \prod_{i=1}^t T(s_{i-1}, ds_i) \\ &= \int_{s_1 \in \mathbb{R}^n} \cdots \int_{s_t \in \mathbb{R}^n} \left(\prod_{i=0}^{t-1} \mathbf{1}_{\mathbb{R}^n \setminus A}(s_i) \right) \mathbf{1}_A(s_t) \prod_{i=1}^t T(s_{i-1}, ds_i) \end{aligned}$$

As a consequence, we can express the probability measure of $\text{Reach}_{t+1}^{s_0}(A)$ as follows:

$$\begin{aligned} \mathbb{P}[\text{Reach}_{t+1}^{s_0}(A)] &= \mathbf{1}_{\mathbb{R}^n \setminus A}(s_0) \int_{s_1 \in \mathbb{R}^n} \\ &\quad \underbrace{\int_{s_2 \in \mathbb{R}^n} \cdots \int_{s_{t+1} \in \mathbb{R}^n} \left(\prod_{i=1}^t \mathbf{1}_{\mathbb{R}^n \setminus A}(s_i) \right) \mathbf{1}_A(s_{t+1}) \left(\prod_{i=2}^{t+1} T(s_{i-1}, ds_i) \right)}_{\mathbb{P}[\text{Reach}_t^{s_1}(A)]} T(s_0, ds_1) \end{aligned}$$

Then, after substituting s_0 with s and s_1 with s' , we derive Eq. (3.2a) as follows:

$$\begin{aligned} \mathbb{P}[\text{Reach}_{t+1}^s(A)] &= \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \int_{\mathbb{R}^n} \mathbb{P}[\text{Reach}_t^{s'}(A)] T(s, ds') \\ &= \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\lambda s'. \mathbb{P}[\text{Reach}_t^{s'}(A)]](s) \quad \text{by (2.4)} \end{aligned}$$

Finally, Eq. (3.2b), which states that $\mathbb{P}[\text{Reach}_0^s(A)] = \mathbf{1}_A(s)$ directly follows from the definition of $\text{Reach}_0^s(A)$, as $\text{Reach}_0^s(A) = \Omega$ when $s \in A$ and $\text{Reach}_0^s(A) = \emptyset$ when $s \notin A$. \square

Lemma 3.2. *Let the event that a stochastic process $\{X_t^s\}_{t \in \mathbb{N}}$ over state space \mathbb{R}^n initialised in state $s \in \mathbb{R}^n$ reaches a target set $A \in \mathcal{B}(\mathbb{R}^n)$ in at most time $t \in \mathbb{N}$ be*

$$\text{Reach}_{\leq t}^s(A) = \cup \{ \text{Reach}_i^s(A) : 0 \leq i \leq t \}. \quad (3.3)$$

Then, $\text{Reach}_{\leq t}^s(A)$ is measurable and its probability measure can be expressed as follows:

$$\mathbb{P}[\text{Reach}_{\leq t+1}^s(A)] = \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\lambda s' \cdot \mathbb{P}[\text{Reach}_{\leq t}^{s'}(A)]](s), \quad (3.4a)$$

$$\mathbb{P}[\text{Reach}_{\leq 0}^s(A)] = \mathbf{1}_A(s). \quad (3.4b)$$

Proof. Since the sets in the family $\{\text{Reach}_t^s(A) : t \in \mathbb{N}\}$ are disjoint, we note that the union in (3.3) that defines $\text{Reach}_{\leq t}^s(A)$ is a union of disjoint sets. It follows that, by countable additivity of measure and disjointness of the union in (3.3), its probability measure is given by

$$\mathbb{P}[\text{Reach}_{\leq t}^s(A)] = \sum_{i=0}^t \mathbb{P}[\text{Reach}_i^s(A)]. \quad (3.5)$$

Consequently and thanks to Lemma 3.1, we have

$$\begin{aligned} \mathbb{P}[\text{Reach}_{\leq t+1}^s(A)] &= \mathbb{P}[\text{Reach}_0^s(A)] + \sum_{i=1}^{t+1} \mathbb{P}[\text{Reach}_i^s(A)] && \text{by (3.5)} \\ &= \mathbf{1}_A(s) + \sum_{i=1}^{t+1} \mathbb{P}[\text{Reach}_i^s(A)] && \text{by (3.2b)} \\ &= \mathbf{1}_A(s) + \sum_{i=0}^t \mathbb{P}[\text{Reach}_{i+1}^s(A)] \\ &= \mathbf{1}_A(s) + \sum_{i=0}^t \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\lambda s' \cdot \mathbb{P}[\text{Reach}_i^{s'}(A)]](s) && \text{by (3.2a)} \\ &= \mathbf{1}_A(s) + \sum_{i=0}^t \left(\mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \int_{\mathbb{R}^n} \mathbb{P}[\text{Reach}_i^{s'}(A)] T(s, ds') \right) && \text{by (2.4)} \\ &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \int_{\mathbb{R}^n} \left(\sum_{i=0}^t \mathbb{P}[\text{Reach}_i^{s'}(A)] \right) T(s, ds') \\ &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \int_{\mathbb{R}^n} \mathbb{P}[\text{Reach}_{\leq t}^{s'}(A)] T(s, ds') && \text{by (3.5)} \\ &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\lambda s' \cdot \mathbb{P}[\text{Reach}_{\leq t}^{s'}(A)]](s) && \text{by (2.4)}. \end{aligned}$$

Furthermore, we also have that

$$\mathbb{P}[\text{Reach}_{\leq 0}^{s_0}(A)] = \mathbb{P}[\text{Reach}_0^{s_0}(A)] = \mathbf{1}_A(s_0) \quad \text{by (3.5) and (3.2b)}. \quad \square$$

Lemma 3.3. *Let the event that a stochastic process $\{X_t^s\}_{t \in \mathbb{N}}$ over state space \mathbb{R}^n initialised in state $s \in \mathbb{R}^n$ reaches a target set $A \in \mathcal{B}(\mathbb{R}^n)$ in finite time be*

$$\text{Reach}_{\text{fin}}^s(A) = \cup \{\text{Reach}_t^s(A) : t \in \mathbb{N}\} = \cup \{\text{Reach}_{\leq t}^s(A) : t \in \mathbb{N}\}. \quad (3.6)$$

Then, $\text{Reach}_{\text{fin}}^s(A)$ is measurable and its probability measure (i.e. the reachability probability of the target set) may be expressed as follows:

$$\mathbb{P}[\text{Reach}_{\text{fin}}^s(A)] = \lim_{t \rightarrow \infty} \mathbb{P}[\text{Reach}_{\leq t}^s(A)]. \quad (3.7)$$

Proof. We recall from Eq. (3.3) that $\text{Reach}_{\leq t}^s(A) = \cup \{\text{Reach}_i^s(A) : 0 \leq i \leq t\}$, and we note that this implies that $\text{Reach}_{\leq t}^s(A) \subseteq \text{Reach}_{\leq t+1}^s(A)$. Then, we recall from Eq. (3.6) that

$\text{Reach}_{\text{fin}}^s(A) = \cup \{\text{Reach}_{\leq t}^s(A) : t \in \mathbb{N}\}$. Finally, we apply a standard result in measure theory regarding the measure of an increasing union [Ax19, Theorem 2.59] to conclude $\mathbb{P}[\text{Reach}_{\text{fin}}^s(A)] = \lim_{t \rightarrow \infty} \mathbb{P}[\text{Reach}_{\leq t}^s(A)]$. \square

The recursive characterisation of the time-bounded reachability probability (Lemma 3.2) and the limit characterisation of the probability of reaching the target set in any finite time (Lemma 3.3) provides the foundation for a fixpoint characterisation for the reachability probability (Section 4), in turn justifying the conditions that a supermartingale certificate must satisfy (Section 5) for it to provide a sound upper bound on the reachability probability.

Our method leverages neural networks to compute an upper bound for the reachability probability with respect to a given target set (cf. Section 5). Notice that an appropriate choice of target set allows us to express a variety of other quantitative verification questions, for which our method can provide an upper or lower bound. Next, by providing a suitable choice for the target set, we show how important quantitative verification questions for probabilistic programs and stochastic dynamical models can be characterised as instances of quantitative reachability.

Termination Analysis: Let $G \in \mathcal{B}(\mathbb{R}^n)$ be the guard set of a probabilistic while loop as in Eq. (2.5), i.e., the set of states for which that guard condition evaluates to true. The event that the loop terminates from initial state s_0 is $\text{Reach}_{\text{fin}}^{s_0}(\mathbb{R}^n \setminus G)$. Our method computes an upper bound for the probability that the loop terminates or, dually, it computes a lower bound for the probability of non-termination. Notably, when this lower bound is greater than 0, then almost-sure termination is refuted.

Assertion-violation Analysis: Let $G \in \mathcal{B}(\mathbb{R}^n)$ be the guard set of a probabilistic while loop and $A \in \mathcal{B}(\mathbb{R}^n)$ be the satisfying set of an assertion placed at the beginning of the loop body. Given initial state s_0 , the event that the assertion is eventually violated is $\text{Reach}_{\text{fin}}^{s_0}(G \setminus A)$. Our method computes an upper bound for the probability of assertion violation or, dually, a lower bound for its satisfaction. Note that assertions in other positions of the body can be modelled similarly by using additional Boolean variables.

Safety Verification: Let $B \in \mathcal{B}(\mathbb{R}^n)$ be a set of undesirable states in a stochastic dynamical model. The event that the system is safe when initialised in s_0 , i.e., it never reaches an undesirable state, is given by $\Omega \setminus \text{Reach}_{\text{fin}}^{s_0}(B)$. Our method computes a lower bound for the probability that the system is safe, which is $1 - \mathbb{P}[\text{Reach}_{\text{fin}}^{s_0}(B)]$.

Invariant Verification: Let $I \in \mathcal{B}(\mathbb{R}^n)$ be a candidate invariant set. The event that I is invariant when the system is initialised in s_0 is $\Omega \setminus \text{Reach}_{\text{fin}}^{s_0}(\mathbb{R}^n \setminus I)$. Our method computes a lower bound for the probability that set I is invariant, which is $1 - \mathbb{P}[\text{Reach}_{\text{fin}}^{s_0}(\mathbb{R}^n \setminus I)]$. Note that if $s_0 \notin I$, this definition yields a trivial lower bound of zero for the probability of invariance.

4. FIXPOINT CHARACTERISATION OF QUANTITATIVE REACHABILITY

Our supermartingale certificate proof-rule (Theorem 5.1) relies on the order-theoretic characterisation of reachability probabilities in probabilistic programs [TOUH18]. This previous seminal work considered a more general setting allowing for both non-determinism and arbitrary program invariants. For clarity and self-containment, and to build intuition about order-theoretic reasoning in our context, we present a streamlined analysis that mirrors the soundness proof for non-negative supermartingales [TOUH18, Proposition 4.2] in a setting without non-determinism or invariants.

We define the space of candidate certificates as the set of $(\mathcal{B}(\mathbb{R}^n), \mathcal{B}(\mathbb{R}_{\geq 0} \cup \{\infty\}))$ -measurable functions $\mathcal{V} \subset \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$. Furthermore, we define the ordering relation \leq between two functions $U, V \in \mathcal{V}$ pointwise as follows:

$$U \leq V \Leftrightarrow \forall s \in \mathbb{R}^n. U(s) \leq V(s). \quad (4.1)$$

The set of functions \mathcal{V} under the ordering \leq defines a pointed complete partially ordered set (pointed CPO) [DP02, Definition 8.1] with the constant zero function as its bottom element; we overload the symbol $\mathbf{0} \in \mathcal{V}$ to denote the constant zero function.

Lemma 4.1. *(\mathcal{V}, \leq) is a pointed CPO with bottom element $\mathbf{0}$, and with each countable increasing chain $\{V_i\}_{i \in \mathbb{N}}$ having least upper bound:*

$$(\sup \{V_i : i \in \mathbb{N}\})(s) = \sup \{V_i(s) : i \in \mathbb{N}\} = \lim_{i \rightarrow \infty} V_i(s). \quad (4.2)$$

Proof. The fact that the pointwise relation \leq on functions is reflexive, transitive and anti-symmetric follows from the properties of the natural ordering relation \leq on the set $\mathbb{R}_{\geq 0} \cup \{\infty\}$. This shows that (\mathcal{V}, \leq) is a partial order. The constant function $\mathbf{0}$ is the bottom element of the partial order because \mathcal{V} contains only non-negative functions.

To show that the partial order is also complete, we need to show that the least upper bound of every countable increasing chain $\{V_i\}_{i \in \mathbb{N}}$ in (\mathcal{V}, \leq) is also in \mathcal{V} ; in other words, that the least upper bound is $(\mathcal{B}(\mathbb{R}^n), \mathcal{B}(\mathbb{R}_{\geq 0} \cup \{\infty\}))$ -measurable. To this end, we note that for all $s \in \mathbb{R}^n$ we have $(\sup \{V_i : i \in \mathbb{N}\})(s) = \sup \{V_i(s) : i \in \mathbb{N}\}$ owing to a standard result [DP02, Lemma 8.5(ii), p.177], and that the pointwise supremum of a countable family of measurable functions is also measurable [Pol01, Example 8, p.23].

Finally, since for any $s \in \mathbb{R}^n$ we have $\{V_i(s)\}_{i \in \mathbb{N}}$ is a non-decreasing sequence in the extended non-negative reals (namely, the set $\mathbb{R}_{\geq 0} \cup \{\infty\}$), we have by the monotone convergence property of non-decreasing real sequences [RF10, Theorem 15, p.21] that

$$\lim_{i \rightarrow \infty} V_i(s) = \sup \{V_i(s) : i \in \mathbb{N}\}. \quad \square$$

The exact reachability probability is a measurable function from states to non-negative reals, and is thus an element of the pointed CPO \mathcal{V} . By Lemma 3.2 we have a recursive characterisation of increasingly tight approximations to this probability in terms of the higher-order endofunction $\Phi : \mathcal{V} \rightarrow \mathcal{V}$ defined as

$$\Phi(V)(s) = \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[V](s). \quad (4.3)$$

Furthermore, Lemma 3.3 states that the exact reachability probability is the limit of its approximations. We formalise this argument by an application of Kleene's Fixpoint Theorem (see [DP02, Section 8.15, p.183] and [Win93, Theorem 5.11]) over the pointed CPO (\mathcal{V}, \leq) over the space of certificates. A technical requirement of Kleene's Fixpoint Theorem is that Φ is ω -continuous [Win93, Definition p.71]. A function is ω -continuous if it is monotonic and preserves least upper bounds of countable increasing chains. To show that the higher-order function Φ is ω -continuous, we first show that the post-expectation operator \mathbb{X} (which is in turn a higher-order function in \mathcal{V}) is ω -continuous.

Lemma 4.2. *\mathbb{X} is ω -continuous.*

Proof. Recalling the definition of ω -continuity for an endofunction on a pointed CPO [Win93, Definition p.71], we show that \mathbb{X} satisfies the following two properties:

- \mathbb{X} is monotonic. Suppose $U, V \in \mathcal{V}$ with $U \leq V$. Then for all $s \in \mathbb{R}^n$:

$$\begin{aligned} \mathbb{X}[U](s) &= \int_{\mathbb{R}^n} U(s') T(s, ds') \\ &\leq \int_{\mathbb{R}^n} V(s') T(s, ds') && \text{by } U \leq V \\ &= \mathbb{X}[V](s). \end{aligned}$$

- \mathbb{X} preserves least upper bounds of countable increasing chains. Let $V_0 \leq V_1 \leq \dots$ be a countable increasing chain in \mathcal{V} . As \mathbb{X} is monotone, $\mathbb{X}[V_0] \leq \mathbb{X}[V_1] \leq \dots$ is also a countable increasing chain. We note that since \mathbb{X} is an integral, it satisfies the Monotone Convergence Theorem (MCT) [Pol01, Theorem 12], which can be invoked in the following since $\{V_i\}_{i \in \mathbb{N}}$ is a sequence of non-negative measurable functions that converge monotonically to the function $\sup \{V_i : i \in \mathbb{N}\}$. Thus, for all $s \in \mathbb{R}^n$, we have that

$$\begin{aligned} (\sup \{\mathbb{X}[V_i] : i \in \mathbb{N}\})(s) &= \lim_{i \rightarrow \infty} \mathbb{X}[V_i](s) && \text{by (4.2), Lemma 4.1} \\ &= \lim_{i \rightarrow \infty} \int_{\mathbb{R}^n} V_i(s') T(s, ds') && \text{by (2.4)} \\ &= \int_{\mathbb{R}^n} (\sup \{V_i : i \in \mathbb{N}\})(s') T(s, ds') && \text{(4.2) \& MCT} \\ &= \mathbb{X}[\sup \{V_i : i \in \mathbb{N}\}](s) \end{aligned} \quad \square$$

Lemma 4.3. Φ is ω -continuous.

Proof. We show that Φ satisfies the following two properties.

- Φ is monotonic. Let $s \in \mathbb{R}^n$ be arbitrary and assume $U \leq V$. Then

$$\begin{aligned} \Phi(U)(s) &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[U](s) \\ &\leq \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[V](s) && \text{by monotonicity of } \mathbb{X}, \text{ Lemma 4.2} \\ &= \Phi(V)(s), \end{aligned}$$

therefore $\Phi(U) \leq \Phi(V)$.

- Φ preserves least upper bounds of countable increasing chains. Suppose $\{V_i : i \in \mathbb{N}\}$ is a countable increasing chain in (\mathcal{V}, \leq) . Let $s \in \mathbb{R}^n$ be arbitrary. Then

$$\begin{aligned} \Phi(\sup \{V_i : i \in \mathbb{N}\})(s) &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\sup \{V_i : i \in \mathbb{N}\}](s) && \text{by (4.3)} \\ &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot (\sup \{\mathbb{X}[V_i] : i \in \mathbb{N}\})(s) && \text{by Lemma 4.2} \\ &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot (\sup \{\mathbb{X}[V_i](s) : i \in \mathbb{N}\}) && \text{by (4.2)} \\ &= \sup \{(\mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[V_i](s)) : i \in \mathbb{N}\} \\ &= \sup \{\Phi[V_i](s) : i \in \mathbb{N}\} && \text{by (4.3)} \\ &= (\sup \{\Phi(V_i) : i \in \mathbb{N}\})(s) && \text{by (4.2);} \end{aligned}$$

we note that the steps above relying on Eq. (4.2) use the property that suprema in the pointed CPO $(\mathcal{V}, \sqsubseteq)$ are defined pointwise. Finally, the derivation above establishes that $\Phi(\sup \{V_i : i \in \mathbb{N}\}) = \sup \{\Phi(V_i) : i \in \mathbb{N}\}$. \square

To demonstrate that the exact reachability probability is the least fixpoint of the higher-order function Φ , we leverage Kleene's Fixpoint Theorem which states that since Φ is

ω -continuous, its least fixpoint is given by

$$\sup \{\Phi^t(\mathbf{0}) : t \in \mathbb{N}\} \quad (4.4)$$

where we let Φ^t denote the t -fold application of Φ to its argument. To relate this to quantitative reachability, we show in Lemma 4.4 that for any $t \in \mathbb{N}$, the function $\lambda s. \mathbb{P}[\text{Reach}_{\leq t}^s(A)]$ is equal to $\Phi^{t+1}(\mathbf{0})$ (which is simply a re-statement of Lemma 3.2). In combination with Eq. (3.7), this yields the following result.

Lemma 4.4. $\lambda s. \mathbb{P}[\text{Reach}_{\text{fin}}^s(A)]$ is the least fixpoint of Φ .

Proof. We need to prove

$$\lambda s. \mathbb{P}[\text{Reach}_{\text{fin}}^s(A)] = (\sup \{\Phi^t(\mathbf{0}) : t \in \mathbb{N}\}). \quad (4.5)$$

To prove Eq. (4.5) we first show that for all $t \in \mathbb{N}$ and $s \in \mathbb{R}^n$, we have

$$\mathbb{P}[\text{Reach}_{\leq t}^s(A)] = \Phi^{t+1}(\mathbf{0})(s). \quad (4.6)$$

We proceed by induction on $t \in \mathbb{N}$, that is, we show

- $\lambda s. \mathbb{P}[\text{Reach}_{\leq 0}^s(A)] = \Phi^1(\mathbf{0})$, and
- $\lambda s. \mathbb{P}[\text{Reach}_{\leq t}^s(A)] = \Phi^{t+1}(\mathbf{0})$ implies $\lambda s. \mathbb{P}[\text{Reach}_{\leq t+1}^s(A)] = \Phi^{t+2}(\mathbf{0})$.

For the base case, we let $s \in \mathbb{R}^n$ and have

$$\begin{aligned} \mathbb{P}[\text{Reach}_{\leq 0}^s(A)] &= \mathbf{1}_A(s) && \text{by (3.4b)} \\ &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\mathbf{0}](s) && \text{because } \mathbb{X}[\mathbf{0}] = \mathbf{0} \\ &= \Phi^1(\mathbf{0})(s) && \text{by (4.3)}. \end{aligned}$$

For the inductive case, we let $s \in \mathbb{R}^n$ and leverage Lemma 3.2 as follows:

$$\begin{aligned} \mathbb{P}[\text{Reach}_{\leq t+1}^s(A)] &= \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[\lambda s'. \mathbb{P}[\text{Reach}_{\leq t}^{s'}(A)]](s) && \text{by (3.4a)} \\ &= \Phi(\lambda s'. \mathbb{P}[\text{Reach}_{\leq t}^{s'}(A)])(s) && \text{by (4.3)} \\ &= \Phi(\Phi^{t+1}(\mathbf{0}))(s) && \text{by IH} \\ &= \Phi^{t+2}(\mathbf{0})(s) \end{aligned}$$

Next, we show that (4.5) is a consequence of Lemma 4.1, noting that $\{\Phi^t(\mathbf{0})\}_{t \in \mathbb{N}}$ is a countable increasing chain in (\mathcal{V}, \leq) :

$$\begin{aligned} \mathbb{P}[\text{Reach}_{\text{fin}}^s(A)] &= \lim_{t \rightarrow \infty} \mathbb{P}[\text{Reach}_{\leq t}^s(A)] && \text{by (3.7)} \\ &= \lim_{t \rightarrow \infty} \Phi^{t+1}(\mathbf{0})(s) && \text{by (4.6)} \\ &= \lim_{t \rightarrow \infty} \Phi^t(\mathbf{0})(s) \\ &= (\sup \{\Phi^t(\mathbf{0}) : t \in \mathbb{N}\})(s) && \text{by (4.2), Lemma 4.1} \quad \square \end{aligned}$$

In summary, Lemma 4.4 characterises the exact reachability probability as the least fixpoint of a ω -continuous endofunction Φ on the set \mathcal{V} . We use this fixpoint characterisation in Section 5 to show how a supermartingale certificate establishes an upper bound on the probability of reaching a target state in any number of steps.

5. NEURAL SUPERMARTINGALES FOR QUANTITATIVE VERIFICATION

Supermartingale certificates provide a flexible and powerful theoretical framework for the formal verification of probabilistic models with infinite state spaces (see [Dur10, Section 5.2], [CS13, Definition 9 and Theorem 4]). Not only have supermartingales been applied to qualitative questions, such as almost-sure termination analysis of probabilistic programs and almost-sure stability analysis of stochastic dynamical models, but also to quantitative reachability verification, as in this work. Specifically, this is enabled by the theory of non-negative repulsing supermartingales and of stochastic invariants [TOUH21, CGMv22, CNv17].

Our method builds upon the theory of non-negative repulsing supermartingales and stochastic invariant indicators and applies it to take advantage of the expressive power of neural networks and the flexibility of machine learning and counterexample-guided inductive synthesis algorithms. We refer to a function V that is non-increasing in expectation outside of the target set as a *supermartingale*. This is because $V(X_t^s)$ defines a supermartingale process [Dur10, Section 5.2], where X_t^s is the stochastic process corresponding to the semantics of the probabilistic model, c.f. Eq. (2.2). Our method hinges on the following theorem.

Theorem 5.1. *Let $A \in \mathcal{B}(\mathbb{R}^n)$ be a target set, and let $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a non-negative function that satisfies the following two conditions:*

$$(indicating\ condition) \quad \forall s \in A: V(s) \geq 1, \quad (5.1a)$$

$$(non-increasing\ condition) \quad \forall s \notin A: \mathbb{X}[V](s) \leq V(s), \quad (5.1b)$$

Then for all $s \in \mathbb{R}^n$ we have $V(s) \geq \mathbb{P}[\text{Reach}_{\text{fin}}^s(A)]$.

Proof. Suppose that $V \in \mathcal{V}$ is a function that satisfies (5.1a) and (5.1b). We show that V is a pre-fixpoint of Φ (defined in Eq. 4.3), namely, that $\Phi(V) \leq V$. We proceed by cases:

$$\text{Case } s \in A: \quad \Phi(V)(s) = \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[V](s) = 1 \leq V(s) \quad \text{by (5.1a),}$$

$$\text{Case } s \notin A: \quad \Phi(V)(s) = \mathbf{1}_A(s) + \mathbf{1}_{\mathbb{R}^n \setminus A}(s) \cdot \mathbb{X}[V](s) = \mathbb{X}[V](s) \leq V(s) \quad \text{by (5.1b).}$$

This establishes that V is a pre-fixpoint of Φ .

Next, we recall the standard result [DP02, Section 8.20, p.186] that (if it exists) the least fixpoint of a monotone endofunction is equal to its least pre-fixpoint (and therefore, the least fixpoint is less than or equal to *any* pre-fixpoint). In our case, since the endofunction $\Phi : \mathcal{V} \rightarrow \mathcal{V}$ is ω -continuous, its least fixpoint exists and is equal (by Lemma 4.4) to the exact reachability probability, $\lambda s. \mathbb{P}[\text{Reach}_{\text{fin}}^s(A)]$. Thus, as V is a pre-fixpoint of Φ , we have that $\lambda s. \mathbb{P}[\text{Reach}_{\text{fin}}^s(A)] \leq V$ holds true, which corresponds to the statement of Theorem 5.1. \square

As we show in detail in Section 6, we compute a supermartingale that satisfies the two criteria (5.1a) and (5.1b), while also minimising its output over the initial state s_0 . When the initial state is chosen nondeterministically from a set S_0 , we instead minimise the output over all states in the set. As a consequence of Theorem 5.1, the maximum of V over S_0 is a sound upper bound for the reachability probability.

In this work, we template V as a neural network with n input neurons, l hidden layers with respectively h_1, \dots, h_l neurons in each hidden layer, and 1 output neuron. We guarantee a-priori that the function's output will be non-negative over the entire domain using the following architecture:

$$V(x) = \text{sum}(\sigma_l \circ \dots \circ \sigma_1(x)), \quad \sigma_i(z) = (\text{ReLU}(\mathbf{w}_{i,1}^T z + \mathbf{b}_{i,1}), \dots, \text{ReLU}(\mathbf{w}_{i,h_i}^T z + \mathbf{b}_{i,h_i})), \quad (5.2)$$

```

while x > 0 do
  assert(x <= 10);
  p ~ Bernoulli(0.5);
  if p == 1 then
    x -= 2
  else
    x += 1
  fi
od

```

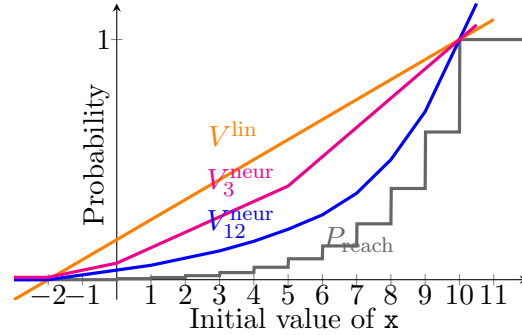


Figure 2: Comparison between linear and neural supermartingale functions in tightness of bounds for the assertion violation probability of the program `repulse`, shown on the left. On the right, the function V^{lin} indicates the tightest linear supermartingale (under the restriction that x is greater than -2). The functions V_3^{neur} and V_{12}^{neur} indicate single-layer neural supermartingales with 3 and 12 neurons respectively. The piecewise constant function P_{reach} is the true probability of assertion violation, and indicates the ideal lower bound for the value of any supermartingale function.

where $\text{sum}(z_{l,1}, \dots, z_{l,h_l}) = \sum_{k=1}^{h_l} z_{l,k}$, $\text{ReLU}(z) = \max\{0, z\}$, and $\mathbf{w}_{i,j}^T \in \mathbb{Q}^{h_{i-1}}$ (defining $h_0 = n$, the number of input neurons) and $\mathbf{b}_{i,j} \in \mathbb{Q}$ are respectively the weight vector and the bias parameter for the inputs to neuron j at layer i . Then, our method trains the neural network to satisfy the two criteria (5.1a) and (5.1b), while also minimising its output on the initial condition.

Using neural networks as templates of supermartingale functions introduces non-trivial advantages with respect to symbolic methods for the synthesis of supermartingales. Firstly, neural supermartingales are able to better approximate the true probability of reachability and thus attain tighter upper bounds on it. Secondly, symbolic methods require deterministic invariants that overapproximate the set of reachable states, and to suitably restrict the domain of the template. Figure 2 illustrates using an example the advantages of neural certificates with respect to linear supermartingales synthesised using Farkas’ lemma. This example shows that increasing the number of neurons provides greater flexibility and allows the certificate to more tightly approximate the true probability. Moreover, this example shows that linear supermartingales require their domain to be restricted with an appropriate deterministic invariant (to enforce non-negativity). Notably, symbolic methods for the synthesis of polynomial supermartingales based on Putinar’s Positivstellensatz also require compact deterministic invariants to be provided [CGMv22]. By contrast, neural supermartingales (whose output is always non-negative) achieve the same result while relaxing the requirement of providing an invariant beforehand.

6. DATA-DRIVEN SYNTHESIS OF NEURAL SUPERMARTINGALES

Our approach to synthesising neural supermartingales for quantitative verification utilises a counterexample-guided inductive synthesis (CEGIS) procedure [Sol08, STB⁺06] (cf. Fig. 3). This procedure consists of two components, a learner and a verifier, that work in opposition to each other. On the one hand, the learner seeks to synthesise a candidate supermartingale




Figure 3: Overview of the counterexample-guided inductive synthesis procedure used to synthesise neural supermartingales for quantitative reachability verification. Inputs to the procedure are a probabilistic program f , a set of initial states S_0 , and a target set A . The procedure outputs a valid neural supermartingale V and a probability bound p .

that meets the desired specification (cf. Eq. (5.1)) over a finite set of samples from the state space, while simultaneously optimising the tightness of the probability bound. On the other hand, the verifier seeks to disprove the validity of this candidate by searching for counterexamples, i.e., instances where the desired specification is invalidated, over the entire state space. If the verifier shows that no such counterexample exists, then the desired specification is met by the supermartingale and the procedure provides a sound probability bound for the reachability probability of interest, together with a neural supermartingale to certify it.

6.1. Training of Neural Supermartingales from Samples. Our neural supermartingale for quantitative reachability verification consists of a neural network with ReLU activation functions, with an arbitrary number of hidden layers (cf. Section 5). We train this neural network using gradient descent over a finite set $D = \{d^{(1)}, \dots, d^{(m)}\} \subseteq \mathbb{R}^n$ of states $d^{(i)}$ sampled over the state space \mathbb{R}^n . Initially, we sample uniformly within a bounded hyper-rectangle of \mathbb{R}^n , a technique that scales effectively to high dimensional state spaces and efficiently populates the initial dataset of state samples. Then, we construct a loss function that guides gradient descent to optimise the parameters (weight and biases) of the neural network V to satisfy the specification set out in Eq. (5.1) while also minimising the probability bound. We define a loss function $\mathcal{L}(D)$ that consists of three terms:

$$\mathcal{L}(D) = \beta_1 \mathcal{L}_{\text{ind}}(D) + \beta_2 \mathcal{L}_{\text{non-inc}}(D) + \beta_3 \mathcal{L}_{\text{min}}(D). \quad (6.1)$$

Components \mathcal{L}_{ind} and $\mathcal{L}_{\text{non-inc}}$ are responsible for encouraging satisfaction of the conditions in Eq. (5.1), while the component \mathcal{L}_{min} is responsible for tightening the probability bound. The parameters of this optimisation problem are the parameters of the neural network, which are initialised randomly. The dataset consists of the state samples, initially sampled randomly, and generated from counterexamples in subsequent CEGIS iterations (cf. Section 6.2). The coefficients β_1, β_2 and β_3 denote scale factors for each term, which we choose according to the priority that we want to assign to each condition (cf. Section 7).

First, consider the condition in Eq. (5.1a), which we refer to as the *indicating condition*. For this, we use the following loss function:

$$\text{(indicating loss)} \quad \mathcal{L}_{\text{ind}}(D) = \frac{1}{|D \cap A|} \sum_{d \in D \cap A} \text{ReLU}(1 - V(d)) \quad (6.2)$$

This adds a penalty for states $d \in D$ lying inside the target condition A , at which V fails to satisfy the indicating condition, whilst ignoring any states where V satisfies it. We average

this per-state penalty across all states in $D \cap A$ (where $|D \cap A|$ denotes the number of elements in this set).

We next consider the *non-increasing condition* in Eq. (5.1b), for which we use the following loss term:

$$\text{(non-increasing loss)} \quad \mathcal{L}_{\text{non-inc}}(D) = \frac{1}{|D \setminus A|} \sum_{d \in D \setminus A} \text{ReLU}(\mathbb{X}[V](d) - V(d)). \quad (6.3)$$

This penalises states $d \in D$ lying outside of the target set A , at which V fails to satisfy the non-increasing condition. Notably, this component is defined in terms of the post-expectation $\mathbb{X}[V]$ of our supermartingale. To embed this expression in our loss function we consider two alternative approaches, which we call *program-aware* and *program-agnostic*.

Program-aware Approach: The program-aware approach uses the source code of the program to generate a symbolic expression for the post-expectation of V . For this purpose, we exploit the symbolic inference algorithm introduced by the tool PSI [GMV16, GSV20], along with a symbolic representation of V . We construct a probabilistic program which represents the evaluation of V on the state resulting after the execution of the update function f . The expected value of this program’s result is precisely $\mathbb{X}[V]$. This results in a symbolic expression that is a function of the program state and parameters of V . In the non-increasing loss, states are instantiated to elements of $D \setminus A$, while the parameters are left as free-variables that the gradient descent engine differentiates with respect to.

Program-agnostic Approach: The program-agnostic approach provides an alternative formulation of the non-increasing loss term that does not require symbolic reasoning. Instead of directly leveraging the program’s source code, it assumes that for any state we can draw independent samples from its distribution over successor states, i.e., we assume access to a *generative model* [MMLG10, KMN99, JS21]. We therefore utilise a Monte Carlo scheme to estimate the post-expectation. For each state d in our dataset D , to obtain an estimate of $\mathbb{X}[V](d)$ we sample a number m' of successor states $D' = \{d^{(k)} : 1 \leq k \leq m'\}$. Each successor state $d^{(k)}$ is sampled by executing the program’s update function f (cf. Eq. (2.1)) at state d . Then $\mathbb{X}[V](d)$ is estimated as $\mathbb{X}[V](d) \approx \frac{1}{|D'|} \sum_{d' \in D'} V(d')$, namely the average of V over D' . Even though this is an approximation, we emphasise that this does not affect the soundness of our scheme, which is ensured by the verifier.

Finally, we introduce a tightness criterion that minimises the probability upper bound:

$$\text{(minimisation loss)} \quad \mathcal{L}_{\text{min}}(D) = \frac{1}{|D \cap S_0|} \sum_{d \in D \cap S_0} V(d). \quad (6.4)$$

This term encourages V to take smaller values over the set of initial states S_0 . Recall that the smaller the probability upper bound, the closer it is to the true value.

The loss function is provided to the gradient descent optimiser, whose performance benefits from a smooth objective. For this reason, to improve the performance of our learner we replace every ReLU activation within the neural network V with a smooth approximation, Softplus, which takes the form $\text{Softplus}(s) = \log(1 + \exp(s))$, whereas the occurrences of ReLU applied to the output of V in Eqs. (6.2), (6.3) and (6.3) are retained. Additionally, we improve the approximation of Softplus to ReLU at small values over the interval $[0, 1]$ by re-scaling V . This means that we modify the indicating condition to require that $V(x) \geq \alpha$

over the target set A , for some large $\alpha > 1$. In other words, we modify the indicating loss to $\mathcal{L}_{\text{ind}}(V) = \frac{1}{|D \cap A|} \sum_{d \in D \cap A} \text{ReLU}(\alpha - V(d))$. We remark that while Softplus is used as the activation function in the learning stage, for the verification stage we instead employ ReLU activation functions, ensuring soundness of the generated neural supermartingale. We remark that this has no effect on the soundness of our approach, which is ultimately guaranteed by the verifier.

6.2. Verification of Neural Supermartingales Using SMT Solvers. The purpose of the verification stage is to check that the neural supermartingale meets the requirements of Eq. (5.1) over the entire state space \mathbb{R}^n , and if that is determined to be the case to furthermore obtain a sound upper bound on the reachability probability. We achieve this by constructing a suitable formula in first-order logic, Eq. (6.5), and use SMT solving to decide its validity, or equivalently, to decide the satisfiability of its negation, Eq. (6.6).

The conditions pertaining to the validity of the neural supermartingale, given in (5.1a) and (5.1b), are encoded by the formulas φ_{ind} and $\varphi_{\text{non-inc}}$. We also note that for a constant $p \in [0, 1]$ to be a sound upper bound on the reachability probability, it is sufficient to require that p is an upper bound on the neural supermartingale’s value over the set of initial states S_0 (cf. Theorem 5.1), which is expressed by the formula φ_{bound} . A suitable choice for the bound p is determined by a binary search over the interval $[0, 1]$.

$$\forall s \in \mathbb{R}^n : \underbrace{(s \in A \rightarrow V(s) \geq 1)}_{\varphi_{\text{ind}}} \wedge \underbrace{(s \notin A \rightarrow \mathbb{X}[V](s) \leq V(s))}_{\varphi_{\text{non-inc}}} \wedge \underbrace{(s \in S_0 \rightarrow V(s) < p)}_{\varphi_{\text{bound}}}. \quad (6.5)$$

Here, $V(s)$ is a symbolic encoding of the candidate neural supermartingale proposed by the learner (Section 6.1), and S_0 and A are defined by Boolean predicates over program variables, all of which (in our setting of networks composed from ReLU activations) are expressible using expressions and constraints in non-linear real arithmetic. For this reason, we use Z3 as our SMT solver [dMB08].

The SMT solver is provided with the negation of Eq. (6.5), namely

$$\exists s \in \mathbb{R}^n : \underbrace{(s \in A \wedge V(s) < 1)}_{\neg\varphi_{\text{ind}}} \vee \underbrace{(s \notin A \wedge \mathbb{X}[V](s) > V(s))}_{\neg\varphi_{\text{non-inc}}} \vee \underbrace{(s \in S_0 \wedge V(s) \geq p)}_{\neg\varphi_{\text{bound}}}, \quad (6.6)$$

and decides its satisfiability, seeking an assignment d_{cex} of s that is a counterexample to the neural supermartingale’s validity, for which any of $\neg\varphi_{\text{ind}}$, $\neg\varphi_{\text{non-inc}}$ and $\neg\varphi_{\text{bound}}$ are satisfied. If no counterexample is found, this certifies the validity of the neural supermartingale. Alternatively, if a counterexample d_{cex} is found, it is added to the data set D , for the synthesis to incrementally resume.

7. EXPERIMENTAL EVALUATION

The previous section has presented a new method for synthesising neural supermartingales. This section presents an empirical evaluation of the method, by testing it against a series of benchmarks. Each benchmark is run ten times. A test is successful if a valid supermartingale is synthesised, and the proportion of successful tests is recorded. Further, the average bound from the valid supermartingales is also recorded, along with the average time taken by the learning and verification steps, respectively. This procedure is applied separately for program-aware and program-agnostic synthesis. We use values $\beta_1 = 10$, $\beta_2 = 10^5$, $\beta_3 = 1$ in Eq. (6.1) to define the priority ordering of the three terms in the loss function, which

Table 1: Results comparing neural supermartingales with Farkas Lemma for different benchmarks. Here, p is the average probability bound generated by the certificate; success ratio is the number of successful experiments, out of 10 repeats, generated by CEGIS with neural supermartingale; ‘-’ means no result was obtained. We also denote the architecture of the network: (h_1, h_2) denotes a network with 2 hidden layers consisting of h_1 and h_2 neurons respectively.

Benchmark	Farkas’ Lemma	Quantitative Neural Certificates				Network Arch.
		Program-Agnostic		Program-Aware		
		p	Success Ratio	p	Success Ratio	
<code>persist_2d</code>	-	≤ 0.1026	0.9	≤ 0.1175	0.9	(3, 1)
<code>faulty_marbles</code>	-	≤ 0.0739	0.9	≤ 0.0649	0.8	3
<code>faulty_unreliable</code>	-	≤ 0.0553	0.9	≤ 0.0536	1.0	3
<code>faulty_regions</code>	-	≤ 0.0473	0.9	≤ 0.0411	0.9	(3, 1)
<code>cliff_crossing</code>	≤ 0.4546	≤ 0.0553	0.9	≤ 0.0591	0.8	4
<code>repulse100</code>	≤ 0.0991	≤ 0.0288	1.0	≤ 0.0268	1.0	3
<code>repulse100_uniform</code>	≤ 0.0991	≤ 0.0344	1.0	-	-	2
<code>repulse100_2d</code>	≤ 0.0991	≤ 0.0568	1.0	≤ 0.0541	1.0	3
<code>faulty_varying</code>	≤ 0.1819	≤ 0.0864	1.0	≤ 0.0865	1.0	2
<code>faulty_concave</code>	≤ 0.1819	≤ 0.1399	1.0	≤ 0.1356	0.9	(3, 1)
<code>fixed_loop</code>	≤ 0.0091	≤ 0.0095	1.0	≤ 0.0094	1.0	1
<code>faulty_loop</code>	≤ 0.0181	≤ 0.0195	1.0	≤ 0.0184	1.0	1
<code>faulty_uniform</code>	≤ 0.0181	≤ 0.0233	1.0	≤ 0.0221	1.0	1
<code>faulty_rare</code>	≤ 0.0019	≤ 0.0022	1.0	≤ 0.0022	1.0	1
<code>faulty_easy1</code>	≤ 0.0801	≤ 0.1007	1.0	≤ 0.0865	1.0	1
<code>faulty_ndecr</code>	≤ 0.0561	≤ 0.0723	1.0	≤ 0.0630	1.0	1
<code>faulty_walk</code>	≤ 0.0121	≤ 0.0173	1.0	≤ 0.0166	1.0	1

we find beneficial across our set of benchmarks. To compare our method against existing work, we perform template-based synthesis of linear supermartingales using Farkas’ Lemma. This requires deterministic invariants to overapproximate the reachable set of states, which may either be generated by abstract interpretation, or provided manually [TOUH21]. In our experiments involving the Farkas’ Lemma, we provide a suitable invariant manually based on the guard of the loop, in some cases strengthening them with additional constraints by an educated guess.

It should be noted that our method is inherently stochastic. One reason is the random initialisation of the neural template’s parameters in the learning phase. In program-agnostic synthesis, an additional source of randomness is the sampling of successor states. So that the results accurately reflect the performance of our method, the random seed for these sources of randomness is selected differently for each test. An additional source of non-determinism arises from the SMT solver Z3 as it generates counterexamples: this cannot be controlled externally. Benchmarks are run on a machine with an Nvidia A40 GPU, and involve the assertion-violation analysis of programs created using the following two patterns:

Unreliable Hardware: These are programs that execute on unreliable hardware. The goal is to upper bound the probability that the program fails to terminate due to a hardware fault. A simple example is `faulty_loop`, whose source code is presented




Figure 4: Probability bounds generated using program-agnostic neural supermartingales and using Farkas’ Lemma. The reference line $y = x$ allows one to see which approach outperforms the other and by how much: above the line means that neural supermartingales outperform Farkas’ Lemma, below the line the opposite. Neural supermartingales can significantly outperform linear templates when a better bound exists, but otherwise achieve similar results. Our approach with program-aware neural supermartingales provides even better outcomes, compared with Farkas’ Lemma.

in Listing 1, and which consists of a loop which may violate an assertion with small probability, modelling a hardware fault.

Robot Motion: These programs model an agent (e.g., a robot) that moves within a physical environment. In these benchmarks, the uncertainty in control and sensing is modelled probabilistically. The environment contains a target region and a hazardous region. The goal is to upper bound the probability that the robot enters the hazardous region. We provide the program `repulse100` as an example (listing 2), which is a variant of `repulse` (Fig. 2) containing a modified assertion and initial state. The program models the motion of a robot in a one-dimensional environment, starting at $x = 10$. The target region is where $x < 0$, and hazardous region is where $x > 100$. As with `repulse`, in each iteration there is an equal probability of x being decremented by 2, and x being incremented by 1.

Several of the programs are based on benchmarks used in prior work focusing on other types of supermartingales [ADFG10, CS13]. Additionally, there are several benchmarks that are entirely new.

The results are reported in Table 1. The table is divided into three sections. The first section shows the benchmarks where Farkas’ Lemma cannot be applied, and where only our method is capable of producing a bound. The second section shows examples where both methods are able to produce a bound, but our method produces a notably better bound. The third section shows benchmarks where both methods produce comparable bounds.

Table 2: Results showing the time taken in seconds to synthesise supermartingales by our method and Farkas’ Lemma. For our method, we show the time taken during learning and verification.

Benchmark	Farkas’ Lemma	Quantitative Neural Certificates			
		Program-Agnostic		Program-Aware	
		Learn Time	Verify Time	Learn Time	Verify Time
persist_2d	-	169.14	85.31	44.96	74.90
faulty_marbles	-	114.24	29.23	15.86	28.68
faulty_unreliable	-	123.85	45.48	18.34	33.97
faulty_regions	-	17.92	35.85	17.55	32.38
cliff_crossing	0.11	134.61	19.02	21.27	29.07
repulse100	0.19	16.65	5.00	6.49	3.74
repulse100_uniform	0.19	21.28	14.18	-	-
repulse100_2d	0.12	122.92	64.54	15.75	47.70
faulty_varying	0.36	21.74	5.06	4.71	3.28
faulty_concave	0.39	49.12	13.37	13.49	7.82
fixed_loop	0.15	14.16	3.14	3.34	2.43
faulty_loop	0.16	25.52	3.81	3.73	2.66
faulty_uniform	0.34	20.20	1.91	6.75	1.33
faulty_rare	0.27	25.52	4.27	3.71	2.96
faulty_easy1	0.31	104.20	12.78	4.95	7.51
faulty_ndecr	0.33	104.89	9.06	5.37	4.66
faulty_walk	0.32	15.08	4.00	6.97	3.33

Dashes in the table indicate experiments where a valid supermartingale could not be obtained. In the case of Farkas’ Lemma, there are several cases where there is no linear supermartingale for the benchmark. By contrast, program-agnostic and program-aware synthesis could be applied to almost all benchmarks. The one exception is `repulse100_uniform` where only program-aware synthesis was unsuccessful: this is due to indicator functions in the post-expectation, which are not smooth and posed a problem for the optimiser. This benchmark underscores the value of program-agnostic synthesis, since it does not require embedding the explicit post-expectation in the loss function.

The first section of the benchmarks in Table 1 demonstrates that our method produces useful results on programs that are out-of-scope for existing techniques. Furthermore, the success ratio of our method is high on all the benchmarks, which indicates its robustness. For the second and third sections (which consist of benchmarks to which Farkas’ Lemma is applicable), the success ratio of our method is broadly maximal, which is to be expected, since these programs can also be solved by Farkas’ Lemma.

In the second section of Table 1, we find more complex benchmarks where our method was able to significantly improve the bound from Farkas’ Lemma. The smallest improvement was about 0.04, and the largest improvement was over 0.39. The intuition here is that neural templates allow more sophisticated supermartingales to be learnt, that can approximate how the reachability probability varies across the state space better than linear templates, and thereby yield tighter probability bounds.

The third section of Table 1 consists of relatively simple benchmarks, where our method produces results that are marginally less tight in comparison to Farkas’ Lemma. This is

not surprising since our method uses neural networks consisting of a single neuron for these examples, owing to their simplicity. The expressive power of these networks is therefore similar to linear templates.

In summary, the results show that our method does significantly better on more complex examples, and marginally worse on very simple examples. This is highlighted in Figure 4. Each point represents a benchmark. The position on the x -axis shows the probability bound obtained by our program-agnostic method, and the y -axis shows the probability bound obtained by Farkas’ Lemma. Points above the line indicate benchmarks where neural supermartingales outperform Farkas’ Lemma, and vice versa. The scale is logarithmic to emphasise order-of-magnitude differences.

Notice that in Table 1 the program-aware algorithm usually yields better bounds than the program-agnostic algorithm, but the improvement is mostly marginal. This is in fact a strength of our method: our data-driven approach performs almost as well as one dependent on symbolic representations, which is promising in light of questions of scalability to more complex programs. We also include a breakdown of computation time (Table 2) which allows distinguishing between learning and verification overheads. Notably, Farkas’ Lemma is significantly faster than our method, given that it relies on solving a convex optimisation problem via linear programming, whereas the synthesis of neural supermartingales is a non-convex optimisation problem that is addressed using gradient descent.

Having presented the experimental results, we shall further comment on some specific benchmarks. The `repulse100` program (listing 2) is a variation of `repulse` presented in Fig. 2, but with the assertion changed to `assert(x <= 100)`. While this is a small program, our method is still able to produce a significantly better result than Farkas’ lemma, using a neural supermartingale with a single hidden layer consisting of three ReLU components that are summed together, which allows a convex piecewise linear function to be learnt.

The `cliff_crossing` program (listing 3) is a further benchmark for which our method is capable of producing a significantly better bound. `cliff_crossing` (listing 3) models a robot moving along a road (defined by the region $(x, y) \in [0, 50] \times [0, 10]$) that is adjacent to a cliff (defined by the region with $y > 10$). This is a 2 dimensional benchmark, for which we use a neural supermartingale that consists of two input neurons and four ReLU components, leading to a clear improvement compared to the linear supermartingale in the tightness of the probability bounds generated, as illustrated by Fig. 5.

Both `repulse100` and `cliff_crossing` are benchmarks that use neural supermartingales with a single hidden layer. An example that uses two hidden layers is `faulty_concave` (listing 4), which simulates a program executing on unreliable hardware, where the probability of a fault varies across the state space. This results in two distinct regions of the state space, one of which has a significantly higher probability of assertion violation than the other. We find that neither a linear template nor a single-layer neural supermartingale is able to exploit this conditional behaviour, each of which yield an overly conservative certificate, but that a neural supermartingale with two hidden layers is able to more tightly approximate the reachability probability in each of the two regions.

8. RELATED WORK

The formal verification of probabilistic programs using supermartingales is a well-studied topic. Early approaches to introduce this technique applied them to almost-sure termination analysis of probabilistic programs [CS13], which allowed several extensions to polynomial

Listing 1: faulty_loop

```

error = 0;
i = 0;
while i < 10 do
  assert(error == 0);
  p ~ Bernoulli(0.999);
  if p == 1 then
    q ~ Bernoulli(0.5);
    if q == 1 then
      i += 1
    fi
  else
    error = 1
  fi
od

```

Listing 2: repulse100

```

x = 10;
while x >= 0 do
  assert(x <= 100);
  p ~ Bernoulli(0.5);
  if p == 1 then
    x -= 2
  else
    x += 1
  fi
od

```

Listing 3: cliff_crossing

```

x = 0;
y = 0;
while x <= 50 do
  assert(y <= 10);
  p ~ Bernoulli(0.9);
  x += 1;
  if p == 0 then
    y += 1
  else
    if y >= 1 then
      y -= 1
    fi
  fi
od

```

Listing 4: faulty_concave

```

i = 1;
error = 0;
while i < 10 do
  assert(error == 0);
  if i < 4 then
    p ~ Bernoulli(0.999)
  else
    p ~ Bernoulli(0.99)
  fi;
  if p == 1 then
    q ~ Bernoulli(0.5);
    if q == 1 then
      i += 1
    fi
  else
    error = 1
  fi
od

```

programs, programs with non-determinism, lexicographic and modular termination arguments, and persistence properties [ACN18, CFG16, CFNH18, CNv17, FC19, HFCG19]. All these methods relied on symbolic reasoning algorithms for synthesising supermartingales, that leveraged theories based on Farkas' lemma for the synthesis of linear certificates, and Putinar's Positivstellensatz and sum-of-square methods for the synthesis of polynomial certificates [CGG⁺25]. While these methods are the state-of-the-art for many existing problem instances in literature, to achieve the strong guarantees that they provide (such as

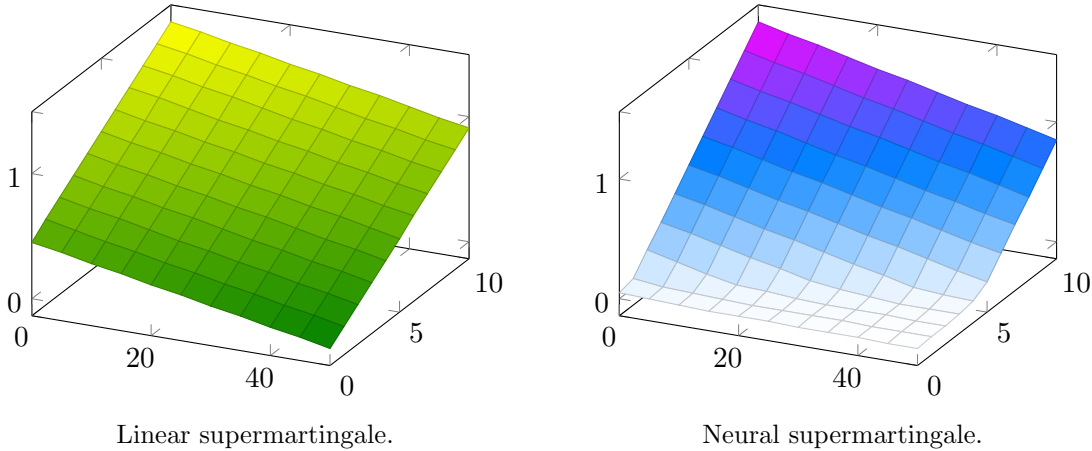


Figure 5: Supermartingale functions for the `cliff_crossing` benchmark as generated using Farkas’ Lemma (on the left) and using neural supermartingales (on the right). The right hand figure illustrates the tighter bounds obtainable through the use of neural templates.

completeness for the specific class of programs they target), they must necessarily introduce restrictions on the class of programs to which they are applicable, and the form of certificates that they derive. Moreover, symbolic methods need externally provided invariants that are stronger than \mathbb{R}^n to enforce non-negativity in the case of linear certificates, as we illustrate in Fig. 2. Also, symbolic methods for the synthesis of polynomial certificates require compact deterministic invariants to operate.

The use of neural networks to represent certificates has allowed many of these restrictions to be lifted. In the context of the analysis of probabilistic programs, neural networks were first applied to certify positive almost-sure termination [AGR21]. This approach lent itself to a wider range of formal verification questions for stochastic dynamical models, from stability and safety analysis to controller synthesis [LSAZ22, LvCH22, CHLv23, MCL23, NMTZ24, NGL25]. These data-driven inductive synthesis techniques for supermartingales have also been extended to machine learning techniques other than deep learning, such as piecewise linear regression and decision tree learning [BCJ⁺23, BTP⁺22]. Data-driven techniques for formal verification have also been applied to the construction of system abstractions, reducing the verification of the original system to a simpler yet equivalent verification task [AGS24, AEG22, AGMS25, MVF⁺25].

The present work focuses on verifying the system under analysis through the construction of proof certificates. In particular, our work extends the data-driven synthesis of neural supermartingale certificates to quantitative verification questions. The correctness of our approach builds upon the theory of non-negative repulsing supermartingales [TOUH21], as formulated in Theorem 5.1. Our experiments have demonstrated that neural certificates attain comparable results on programs that are amenable to symbolic analysis (such as those in the third section of Table 1), while surpassing symbolic methods on more complex programs that are either out-of-scope or yield overly conservative bounds when existing techniques are applied (such as those in the first and second section of Table 1).

9. CONCLUSION

We have presented a data-driven framework for the quantitative verification of probabilistic models that leverage neural networks to represent supermartingale certificates. Our experiments have shown that neural certificates are applicable to a wider range of probabilistic models than was previously possible using purely symbolic techniques. We also illustrate that on existing models our method yields certificates of better or comparable quality than those produced by symbolic techniques for the synthesis of linear supermartingales. This builds upon the ability of neural networks to approximate non-linear functions, while satisfying the constraints imposed by Theorem 5.1 without the need for supporting deterministic invariants to be provided externally. Our method applies to quantitative termination and assertion-violation analysis for probabilistic programs, as well as safety and invariant verification for stochastic dynamic models. We imagine extensions to further quantitative verification questions, such as temporal properties beyond reachability [BK08, AGR24, GKPT24, AGR25, GKPT25], and bounding expected accrued costs [NCH18, WFG⁺19, WHR21]. While neural networks of no greater than two layers suffice for our examples, it also remains open to investigate whether deeper network architectures are beneficial for quantitative verification, as well as their effect on the tightness of bounds obtained and scalability of the approach.

REFERENCES

- [AAGP21] Alessandro Abate, Daniele Ahmed, Mirco Giacobbe, and Andrea Peruffo. Formal synthesis of Lyapunov Neural Networks. *IEEE Control. Syst. Lett.*, 5(3):773–778, 2021.
- [ACN18] Sheshansh Agrawal, Krishnendu Chatterjee, and Petr Novotný. Lexicographic ranking supermartingales: an efficient approach to termination of probabilistic programs. *Proc. ACM Program. Lang.*, 2(POPL):34:1–34:32, 2018.
- [ADFG10] Christophe Alias, Alain Darte, Paul Feautrier, and Laure Gonnord. Multi-dimensional rankings, program termination, and complexity bounds of flowchart programs. In *SAS*, volume 6337 of *Lecture Notes in Computer Science*, pages 117–133. Springer, 2010.
- [AEG22] Alessandro Abate, Alec Edwards, and Mirco Giacobbe. Neural abstractions. In *NeurIPS*, 2022.
- [AEG⁺23] Alessandro Abate, Alec Edwards, Mirco Giacobbe, Hashan Punchihewa, and Diptarko Roy. Quantitative verification with neural networks. In *CONCUR*, volume 279 of *LIPICs*, pages 22:1–22:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [AG07] S. Asmussen and P.W. Glynn. *Stochastic Simulation: Algorithms and Analysis*. Stochastic Modelling and Applied Probability. Springer New York, 2007. URL: <https://books.google.com.sg/books?id=vMI2MdxchU0C>.
- [AGMS25] Alessandro Abate, Mirco Giacobbe, Christian Micheletti, and Yannik Schnitzer. Branching bisimulation learning. In *CAV (4)*, volume 15934 of *Lecture Notes in Computer Science*, pages 161–184. Springer, 2025.
- [AGR21] Alessandro Abate, Mirco Giacobbe, and Diptarko Roy. Learning probabilistic termination proofs. In *CAV (2)*, volume 12760 of *Lecture Notes in Computer Science*, pages 3–26. Springer, 2021.
- [AGR24] Alessandro Abate, Mirco Giacobbe, and Diptarko Roy. Stochastic omega-regular verification and control with supermartingales. In *CAV (3)*, volume 14683 of *Lecture Notes in Computer Science*, pages 395–419. Springer, 2024.
- [AGR25] Alessandro Abate, Mirco Giacobbe, and Diptarko Roy. Quantitative supermartingale certificates. In *CAV (2)*, volume 15932 of *Lecture Notes in Computer Science*, pages 3–28. Springer, 2025.
- [AGS24] Alessandro Abate, Mirco Giacobbe, and Yannik Schnitzer. Bisimulation learning. In *CAV (3)*, volume 14683 of *Lecture Notes in Computer Science*, pages 161–183. Springer, 2024.
- [AKM11] Alessandro Abate, Joost-Pieter Katoen, and Alexandru Mereacre. Quantitative automata model checking of autonomous stochastic hybrid systems. In *HSCC*, pages 83–92. ACM, 2011.

- [Axl19] S. Axler. *Measure, Integration & Real Analysis*. Graduate Texts in Mathematics. Springer International Publishing, 2019.
- [BCJ⁺23] Kevin Batz, Mingshuai Chen, Sebastian Junges, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. Probabilistic program verification via inductive synthesis of inductive invariants. In *TACAS (2)*, volume 13994 of *Lecture Notes in Computer Science*, pages 410–429. Springer, 2023.
- [BdAFK18] Christel Baier, Luca de Alfaro, Vojtech Forejt, and Marta Kwiatkowska. Model checking probabilistic systems. In *Handbook of Model Checking*, pages 963–999. Springer, 2018.
- [BHHK10] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Performance evaluation and model checking join forces. *Commun. ACM*, 53(9):76–85, 2010.
- [BK08] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- [BS96] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control: The discrete-time case*. Athena Scientific, 1996.
- [BTP⁺22] Jialu Bao, Nitesh Trivedi, Drashti Pathak, Justin Hsu, and Subhajit Roy. Data-driven invariant learning for probabilistic programs. In *CAV (1)*, volume 13371 of *Lecture Notes in Computer Science*, pages 33–54. Springer, 2022.
- [CFG16] Krishnendu Chatterjee, Hongfei Fu, and Amir Kafshdar Goharshady. Termination analysis of probabilistic programs through positivstellensatz’s. In *CAV (1)*, volume 9779 of *Lecture Notes in Computer Science*, pages 3–22. Springer, 2016.
- [CFNH18] Krishnendu Chatterjee, Hongfei Fu, Petr Novotný, and Rouzbeh Hasheminezhad. Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. *ACM Trans. Program. Lang. Syst.*, 40(2):7:1–7:45, 2018.
- [CGG⁺25] Krishnendu Chatterjee, Amir Kafshdar Goharshady, Ehsan Kafshdar Goharshady, Mehrdad Karrabi, Milad Saadat, Maximilian Seeliger, and Dorde Žikelić. Polyqent: A polynomial quantified entailment solver. In *ATVA*, volume 16145 of *Lecture Notes in Computer Science*, pages 411–424. Springer, 2025.
- [CGMv22] Krishnendu Chatterjee, Amir Kafshdar Goharshady, Tobias Meggendorfer, and Đorđe Žikelić. Sound and complete certificates for quantitative termination analysis of probabilistic programs. In *CAV (1)*, volume 13371 of *Lecture Notes in Computer Science*, pages 55–78. Springer, 2022.
- [CHLv23] Krishnendu Chatterjee, Thomas A. Henzinger, Mathias Lechner, and Đorđe Žikelić. A learner-verifier framework for neural network controllers and certificates of stochastic systems. In *TACAS (1)*, volume 13993 of *Lecture Notes in Computer Science*, pages 3–25. Springer, 2023.
- [CNv17] Krishnendu Chatterjee, Petr Novotný, and Đorđe Žikelić. Stochastic invariants for probabilistic termination. In *POPL*, pages 145–160. ACM, 2017.
- [CRG19] Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov Control. In *NeurIPS*, pages 3240–3249, 2019.
- [CS13] Aleksandar Chakarov and Sriram Sankaranarayanan. Probabilistic program analysis with martingales. In *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 511–526. Springer, 2013.
- [dMB08] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Z3: an efficient SMT solver. In *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
- [DP02] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2 edition, 2002. doi:10.1017/CB09780511809088.
- [DS20] Fredrik Dahlqvist and Alexandra Silva. Semantics of probabilistic programming: A gentle introduction. In Gilles Barthe, Joost-Pieter Katoen, and Alexandra Silva, editors, *Foundations of Probabilistic Programming*, pages 1–42. Cambridge University Press, 2020.
- [Dur10] Rick Durrett. *Probability: Theory and Examples, 4th Edition*. Cambridge University Press, 2010.
- [FC19] Hongfei Fu and Krishnendu Chatterjee. Termination of nondeterministic probabilistic programs. In *VMCAI*, volume 11388 of *LNCS*, pages 468–490. Springer, 2019.
- [FKN⁺11] Vojtech Forejt, Marta Z. Kwiatkowska, Gethin Norman, David Parker, and Hongyang Qu. Quantitative multi-objective verification for probabilistic systems. In *TACAS*, volume 6605 of *Lecture Notes in Computer Science*, pages 112–127. Springer, 2011.
- [GHNr14] Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. Probabilistic programming. In *FOSE*, pages 167–181. ACM, 2014.

- [GKP22] Mirco Giacobbe, Daniel Kroening, and Julian Parsert. Neural termination analysis. In *ESEC/SIGSOFT FSE*, pages 633–645. ACM, 2022.
- [GKPT24] Mirco Giacobbe, Daniel Kroening, Abhinandan Pal, and Michael Tautschnig. Neural model checking. In *NeurIPS*, 2024.
- [GKPT25] Mirco Giacobbe, Daniel Kroening, Abhinandan Pal, and Michael Tautschnig. Let a neural network be your invariant. In *NeurIPS*, 2025.
- [GMV16] Timon Gehr, Sasa Misailovic, and Martin T. Vechev. PSI: exact symbolic inference for probabilistic programs. In *CAV (1)*, volume 9779 of *Lecture Notes in Computer Science*, pages 62–83. Springer, 2016.
- [GSV20] Timon Gehr, Samuel Steffen, and Martin T. Vechev. λ PSI: exact inference for higher-order probabilistic programs. In *PLDI*, pages 883–897. ACM, 2020.
- [HFCG19] Mingzhang Huang, Hongfei Fu, Krishnendu Chatterjee, and Amir Kafshdar Goharshady. Modular verification for almost-sure termination of probabilistic programs. *Proc. ACM Program. Lang.*, 3(OOPSLA):129:1–129:29, 2019.
- [HLL96] O. Hernández-Lerma and J. B. Lasserre. *Discrete-time Markov control processes*, volume 30 of *Applications of Mathematics*. Springer-Verlag, 1996.
- [JS21] Yujia Jin and Aaron Sidford. Towards tight bounds on the sample complexity of average-reward mdps. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 5055–5064. PMLR, 2021.
- [Kal06] O. Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- [KMN99] Michael J. Kearns, Yishay Mansour, and Andrew Y. Ng. A sparse sampling algorithm for near-optimal planning in large markov decision processes. In *IJCAI*, pages 1324–1231. Morgan Kaufmann, 1999.
- [KNP05] Marta Z. Kwiatkowska, Gethin Norman, and David Parker. Quantitative analysis with the probabilistic model checker PRISM. In *QAPL*, volume 153 of *Electronic Notes in Theoretical Computer Science*, pages 5–31. Elsevier, 2005.
- [Koz81] Dexter Kozen. Semantics of probabilistic programs. *J. Comput. Syst. Sci.*, 22(3):328–350, 1981.
- [Kwi07] Marta Z. Kwiatkowska. Quantitative verification: models techniques and tools. In *ESEC/SIGSOFT FSE*, pages 449–458. ACM, 2007.
- [LSAZ22] A. Lavaei, S. Soudjani, A. Abate, and M. Zamani. Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica*, 146(12), 2022.
- [LvCH22] Mathias Lechner, Đorđe Žikelić, Krishnendu Chatterjee, and Thomas A. Henzinger. Stability verification in stochastic control systems via neural network supermartingales. In *AAAI*, pages 7326–7336. AAAI Press, 2022.
- [MCL23] Frederik Baymler Mathiesen, Simeon Craig Calvert, and Luca Laurenti. Safety certification for stochastic systems via neural barrier functions. *IEEE Control. Syst. Lett.*, 7:973–978, 2023.
- [MM02] Annabelle McIver and Carroll Morgan. Games, probability and the quantitative μ -calculus $qm\mu$. In *LPAR*, volume 2514 of *Lecture Notes in Computer Science*, pages 292–310. Springer, 2002.
- [MM05] Annabelle McIver and Carroll Morgan. *Abstraction, Refinement and Proof for Probabilistic Systems*. Monographs in Computer Science. Springer, 2005.
- [MMLG10] Odalric-Ambrym Maillard, Rémi Munos, Alessandro Lazaric, and Mohammad Ghavamzadeh. Finite-sample analysis of bellman residual minimization. In *ACML*, volume 13 of *JMLR Proceedings*, pages 299–314. JMLR.org, 2010.
- [MMS96] Carroll Morgan, Annabelle McIver, and Karen Seidel. Probabilistic predicate transformers. *ACM Trans. Program. Lang. Syst.*, 18(3):325–353, 1996.
- [MTG09] Sean Meyn, Richard L. Tweedie, and Peter W. Glynn. *Markov Chains and Stochastic Stability*. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2009.
- [MV⁺25] Frederik Baymler Mathiesen, Nikolaus Vertovec, Francesco Fabiano, Luca Laurenti, and Alessandro Abate. Certified neural approximations of nonlinear dynamics. *CoRR*, abs/2505.15497, 2025.
- [NCH18] Van Chan Ngo, Quentin Carbonneaux, and Jan Hoffmann. Bounded expectations: resource analysis for probabilistic programs. In *PLDI*, pages 496–512. ACM, 2018.
- [NGL25] Grigory Neustroev, Mirco Giacobbe, and Anna Lukina. Neural continuous-time supermartingale certificates. In *AAAI*, pages 27538–27546. AAAI Press, 2025.

- [NMTZ24] Alireza Nadali, Vishnu Murali, Ashutosh Trivedi, and Majid Zamani. Neural closure certificates. In *AAAI*, pages 21446–21453. AAAI Press, 2024.
- [PAA21] Andrea Peruffo, Daniele Ahmed, and Alessandro Abate. Automated and formal synthesis of neural barrier certificates for dynamical models. In *TACAS (1)*, volume 12651 of *Lecture Notes in Computer Science*, pages 370–388. Springer, 2021.
- [Pol01] David Pollard. *A User’s Guide to Measure Theoretic Probability*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2001.
- [Rev75] D. Revuz. *Markov Chains*. Mathematical Studies. North-Holland Publishing Company, 1975.
- [RF10] H.L. Royden and P. Fitzpatrick. *Real Analysis*. Prentice Hall, 4 edition, 2010.
- [SA12] S. Esmacil Zadeh Soudjani and A. Abate. Adaptive and sequential gridding for abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2012.
- [Sol08] Armando Solar-Lezama. *Program Synthesis by Sketching*. PhD thesis, University of California at Berkeley, USA, 2008.
- [STB⁺06] Armando Solar-Lezama, Liviu Tancau, Rastislav Bodík, Sanjit A. Seshia, and Vijay A. Saraswat. Combinatorial sketching for finite programs. In *ASPLOS*, pages 404–415. ACM, 2006.
- [TMKA17] Ilya Tkachev, Alexandru Mereacre, Joost-Pieter Katoen, and Alessandro Abate. Quantitative model-checking of controlled discrete-time markov processes. *Inf. Comput.*, 253:1–35, 2017.
- [TOUH18] Toru Takisaka, Yuichiro Oyabu, Natsuki Urabe, and Ichiro Hasuo. Ranking and repulsing supermartingales for reachability in probabilistic programs. In *ATVA*, volume 11138 of *Lecture Notes in Computer Science*, pages 476–493. Springer, 2018.
- [TOUH21] Toru Takisaka, Yuichiro Oyabu, Natsuki Urabe, and Ichiro Hasuo. Ranking and repulsing supermartingales for reachability in randomized programs. *ACM Trans. Program. Lang. Syst.*, 43(2):5:1–5:46, 2021.
- [vBW01] Franck van Breugel and James Worrell. Towards quantitative verification of probabilistic transition systems. In *ICALP*, volume 2076 of *Lecture Notes in Computer Science*, pages 421–432. Springer, 2001.
- [vLHC23] Đorđe Žikelić, Mathias Lechner, Thomas A. Henzinger, and Krishnendu Chatterjee. Learning control policies for stochastic systems with reach-avoid guarantees. In *AAAI*, pages 11926–11935. AAAI Press, 2023.
- [vLV⁺23] Đorđe Žikelić, Mathias Lechner, Abhinav Verma, Krishnendu Chatterjee, and Thomas A. Henzinger. Compositional policy learning in stochastic control systems with formal guarantees. In *NeurIPS*, 2023.
- [WFG⁺19] Peixin Wang, Hongfei Fu, Amir Kafshdar Goharshady, Krishnendu Chatterjee, Xudong Qin, and Wenjun Shi. Cost analysis of nondeterministic probabilistic programs. In *PLDI*, pages 204–220. ACM, 2019.
- [WHR21] Di Wang, Jan Hoffmann, and Thomas W. Reps. Central moment analysis for cost accumulators in probabilistic programs. In *PLDI*, pages 559–573. ACM, 2021.
- [Win93] Glynn Winskel. *The formal semantics of programming languages - an introduction*. Foundation of computing series. MIT Press, 1993.