

Neural Continuous-Time Supermartingale Certificates

Grigory Neustroev¹, Mirco Giacobbe², Anna Lukina¹

¹Delft University of Technology, the Netherlands

²University of Birmingham, UK

g.neustroev@tudelft.nl, m.giacobbe@bham.ac.uk, a.lukina@tudelft.nl

Abstract

We introduce for the first time a neural-certificate framework for continuous-time stochastic dynamical systems. Autonomous learning systems in the physical world demand continuous-time reasoning, yet existing learnable certificates for probabilistic verification assume discretization of the time continuum. Inspired by the success of training neural Lyapunov certificates for deterministic continuous-time systems and neural supermartingale certificates for stochastic discrete-time systems, we propose a framework that bridges the gap between continuous-time and probabilistic neural certification for dynamical systems under complex requirements. Our method combines machine learning and symbolic reasoning to produce formally certified bounds on the probabilities that a nonlinear system satisfies specifications of reachability, avoidance, and persistence. We present both the theoretical justification and the algorithmic implementation of our framework and showcase its efficacy on popular benchmarks.

Code — <https://doi.org/10.5281/zenodo.14537003>

Extended version —

<https://doi.org/10.48550/arXiv.2412.17432>

Introduction

Ensuring safety is paramount to the alignment and governance of autonomous learning systems. Providing verifiable guarantees of compliance with safety requirements is essential for building trust between a system, its users, and the regulators. This is especially critical when the system is deployed in the physical world, where it can cause harm. A central challenge in designing autonomous learning systems is developing control policies that ensure the desired spatial and temporal behavior. These objectives include reaching a specific target region (reachability), consistently avoiding unsafe regions (avoidance), reaching and remaining within a safe region (persistence), and combinations thereof.

Autonomous systems often face multiple probabilistic uncertainties, leading to noisy model parameters or stochastic system evolution. Safe control design for the most general models—nonlinear continuous-time stochastic dynamical systems—provides *formal* guarantees for systems operating in uncertain physical environments. Powerful techniques for achieving this are based on Lyapunov functions,

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: A neural supermartingale certificate for the continuous-time stochastic inverted pendulum. Darker colors indicate higher probability for the system trajectories (sampled in white) to reach and remain in the green rectangle, while avoiding the red rectangles.

which serve as a potential-energy characterization for differential equations and, under certain conditions, become formal proof certificates for desired behavioral properties. Initially introduced for the asymptotic stability analysis of deterministic systems with respect to an equilibrium (Kalman and Bertram 1959), Lyapunov functions have been generalized to stochastic dynamical systems (Blumenthal and Gettoor 1968) and extended into formal certificates for reachability, avoidance (known as barrier certificates), and persistence properties (Prajna, Jadbabaie, and Pappas 2004).

Computing proof certificates analytically is a hard problem for nonlinear systems, which has traditionally required manual effort and considerable expertise. This approach is feasible when a system design is familiar to an engineer. However, for autonomous learning systems with embedded neural networks, whose internals are often not human-comprehensible, the construction of proof certificates must be automated. The best-known algorithms for their construction are based on sum-of-square programming (Papachristodoulou and Prajna 2002; Topcu, Packard, and

Seiler 2008), and are limited to polynomial systems. Systems with dynamics involving trigonometric, exponential, or piecewise functions, which include neural building blocks, are beyond their reach. Therefore, systems that embed neural networks necessarily require more flexible techniques.

Neural networks as control policies for nonlinear dynamical systems have long proven extreme efficacy when in the presence of complex tasks. This has inspired the introduction of Lyapunov certification into learning (Richards, Berkenkamp, and Krause 2018; Dawson, Gao, and Fan 2023), and the growing demand for formally verified artificial intelligence has motivated their coupling with symbolic reasoning to guarantee soundness of the certificate (Chang, Roohi, and Gao 2019; Abate et al. 2021). These methods exploit the expressive capabilities of neural networks to represent proof certificates—referred to as *neural certificates*—and build upon the fact that algorithmically verifying neural networks is increasingly feasible (Katz et al. 2017; Gehr et al. 2018; Wu et al. 2024). This is particularly true in control applications, where the neural networks involved are typically much smaller compared to those used in computer vision or natural language processing (Tran et al. 2020; Bacci, Giacobbe, and Parker 2021; Akintunde et al. 2022).

Neural certificates for deterministic continuous-time dynamical systems implement Lyapunov-like conditions that are both amenable to gradient descent for learning and compatible with symbolic reasoning for formal verification. While Lyapunov conditions are well-defined for deterministic systems and worst-case reasoning, they fall short in quantifying the probability of compliance with a requirement for stochastic systems. Moreover, worst-case reasoning may consider measure-zero trajectories as plausible, imposing sure satisfaction of a requirement, which is unreasonably conservative in many practical scenarios. Due to the fundamental differences between (deterministic) Newton–Leibnitz calculus and (stochastic) Itô calculus, the safety assurance of stochastic dynamical systems demands quantitative probabilistic reasoning techniques.

Probabilistic reasoning for the safety assurance of stochastic systems is grounded in martingale theory, which establishes conditions for their compliance with behavioral requirements (Prajna, Jadbabaie, and Pappas 2007). More specifically, proof rules for probabilistic verification rely on the construction of supermartingale processes. Neural certificates based on supermartingale-like conditions have been successful in proving termination for probabilistic programs (Abate, Giacobbe, and Roy 2021; Abate et al. 2023) and, analogously, probabilistic reachability, avoidance, and reach-avoidance for discrete-time stochastic dynamical systems (Ansari pour et al. 2023; Mathiesen, Calvert, and Laurenti 2023; Žikelić et al. 2023a; Badings et al. 2024). However, all existing results on *neural supermartingale certificates* assume discretization of the time continuum, which is suitable for computer simulations of dynamical systems and other experiments *in silico*, but provide no formal guarantees for systems in continuous time.

We present for the first time a framework for the construction of neural supermartingale certificates in continuous time, with formal guarantees for reachability, avoidance

and persistence requirements. We provide (1) a quantitative proof rule for continuous-time reach-avoid-stay properties, (2) an algorithm for training neural supermartingale certificates from samples of drift and diffusion of the stochastic dynamics, and (3) a symbolic technique for the formal verification of our neural certificates based on interval-bound propagation. Our method produces a quantitative neural certificate that yields a sound lower bound on the probability of compliance of the system with the given requirement.

Figure 1 illustrates a valid quantitative neural certificate for a stochastic inverted pendulum subject to a neural controller, depicting its likelihood of avoiding the red rectangles while reaching and remaining within the green rectangle. It can be observed that any trajectory initialized in a given state satisfies the property with at least the probability indicated for that state. Specifically, our certificate shows that every trajectory initialized within the yellow rectangle satisfies the reach-stay-avoidance property with high probability.

We demonstrate the effectiveness of our method on well-known continuous-time stochastic system models from the literature. Our examples include systems with dynamics that incorporate trigonometric functions and neural controllers, which are beyond the capabilities of traditional algorithms for provably safe control of continuous-time stochastic dynamical systems. Our results establish a foundation for unsupervised and practically effective safeguarded stochastic control based on neural certificates, offering quantitative guarantees for continuous-time dynamical models.

In summary, our contribution is twofold. First, we introduce a comprehensive methodology for formally guaranteed neural reach-avoid-stay certification of continuous-time stochastic dynamical systems controlled by a given policy. Second, we develop a prototype based on our technique and demonstrate its practical effectiveness on examples that are beyond the capabilities of traditional methods for automated provably safe control.

Problem Statement

Consider a filtered probability space $(\Omega, \mathcal{F}, \mathfrak{F}, \mathbb{P})$ with σ -algebra \mathcal{F} , filtration $\mathfrak{F} = (\mathcal{F}_t)_{t \geq 0}$, and probability measure \mathbb{P} . Let $(\mathbb{R}^l, \mathcal{B}(\mathbb{R}^l))$ and $(\mathbb{U}, \mathcal{B}(\mathbb{U}))$ be Polish spaces of states and actions (here $\mathcal{B}(\cdot)$ denotes the Borel σ -algebra, i.e., the smallest σ -algebra that contains the open intervals). In other words, we assume that the states $x \in \mathbb{R}^l$ are real-valued vectors, but the controls $u \in \mathbb{U}$ are anything Borel-measurable.

Consider a continuous-time dynamical system described by the following stochastic differential equation (SDE):

$$dX_t = f(t, X_t, u) dt + g(t, X_t, u) dW_t \quad (1)$$

where $f : \mathbb{R}_+ \times \mathbb{R}^l \times \mathbb{U} \rightarrow \mathbb{R}^l$ and $g : \mathbb{R}_+ \times \mathbb{R}^l \times \mathbb{U} \rightarrow \mathbb{R}^l \times \mathbb{R}^k$ are vector- and matrix-valued *drift* and *diffusion* functions, $\mathbb{R}_+ \triangleq [0, \infty)$; they represent the deterministic and stochastic parts of the dynamics. $(W_t)_{t \geq 0}$ is a k -dimensional Wiener process with independent coordinates.

We assume that the state space \mathbb{R}^l contains a subset \mathbb{X} of interest that we restrict our attention to. For example, the system can be known to remain in \mathbb{X} *almost surely* (*a.s.*), that is, with probability one.

The control signal $u = \pi(t, x)$ is given by a function of time and state known as a *policy* $\pi : \mathbb{R}_+ \times \mathbb{R}^l \rightarrow \mathbb{U}$. When a policy is fixed, we denote

$$f_\pi(t, x) \triangleq f(t, x, \pi(t, x)), \quad g_\pi(t, x) \triangleq g(t, x, \pi(t, x)).$$

We use $(X_{\pi, t}^{s, x})_{t \geq s}$ for a stochastic state process issuing from a state $x \in \mathbb{R}^l$ at time $s \in \mathbb{R}_+$ when the system is controlled by a policy π . By its definition,

$$X_{\pi, t}^{t, x} \equiv x. \quad (2)$$

For a given policy π , our goal is to design an algorithm for verifying that the system's stochastic dynamics satisfy the conditions of the following definition.

Definition 1. A probabilistic reach-stay-avoid (RAS) specification is a tuple $(\mathcal{X}_*, \mathcal{X}_\circ, \mathbb{X}_0, \varepsilon, \delta)$, where

- $\mathcal{X}_* \triangleq (\mathbb{X}_*(t))_{t \geq 0}$ is a family of target sets $\mathbb{X}_*(t) \subset \mathbb{X}$,
- $\mathcal{X}_\circ \triangleq (\mathbb{X}_\circ(t))_{t \geq 0}$ is a family of unsafe sets $\mathbb{X}_\circ(t) \subset \mathbb{X}$,
- \mathbb{X}_0 is a set of initial states,
- $\varepsilon \in [0, 1)$ is an avoidance probability, and
- $\delta \in [0, 1)$ is a stay probability.

A policy π satisfies this specification if for every initial state $x_0 \in \mathbb{X}_0$ there exists a.s. a finite time $\tau < \infty$ such that the probability to reach the target set without ever entering the unsafe set by the time τ is at least ε , and the probability to remain in the target set afterwards is at least δ ; in other words, if $\mathbb{P}[E_{RA}] \geq \varepsilon$ and $\mathbb{P}[E_S] \geq \delta$ for the events

$$E_{RA} \triangleq \{X_{\pi, \tau}^{0, x_0} \in \mathbb{X}_*(\tau) \wedge \forall t < \tau : X_{\pi, t}^{0, x_0} \notin \mathbb{X}_\circ(t)\}, \quad (3)$$

$$E_S \triangleq \{\forall t \geq \tau : X_{\pi, t}^{0, x_0} \in \mathbb{X}_*(t)\}. \quad (4)$$

The dependence of the initial, target, and unsafe sets on time allows us to tackle problems that are not time-homogenous. For example, the unsafe set $\mathbb{X}_\circ(t)$ can move or grow over time or the policy can evolve during training.

Model Assumptions & Further Notation

A policy is assumed to be *admissible*, that is, the coefficients $f_\pi(t, x)$ and $g_{\pi, r}(t, x)$ are Lipschitz-continuous, have continuous derivatives with respect to x which are bounded uniformly in $t > 0$. We use $g_{\pi, r}, r = 1, \dots, k$ to denote the column vectors of g_π .

The set \mathbb{X} is compact. For $\mathbb{X}_\circ \in \{\mathbb{X}_*, \mathbb{X}_\circ\}$ the sets

$$\mathcal{X}_\circ \triangleq \{(t, x) \mid t \geq 0 \wedge x \in \mathbb{X}_\circ(t)\} \quad (5)$$

are Borel. With a slight abuse of notation, we denote these sets \mathcal{X}_* and \mathcal{X}_\circ . The initial set \mathbb{X}_0 is compact and belongs to the interior of safe states at time 0, $\mathbb{X}_0 \subset \text{int}(\mathbb{X} \setminus \mathbb{X}_\circ(0))$. The sets $\mathbb{X}_*(t)$ belong to the interiors of safe sets, $\mathbb{X}_*(t) \subset \text{int}(\mathbb{X} \setminus \mathbb{X}_\circ(t))$.

Without loss of generality, we assume that the equilibrium of the system is at zero, $x_* \equiv 0$ and always belongs to the set $\mathbb{X}_*(t)$; similarly, $u_* \equiv 0$. Moreover, $\pi(t, 0) \equiv 0$, $f_\pi(t, 0) \equiv 0$, and $g_{\pi, r}(t, 0) \equiv 0$.

The design of our certificate defines its properties within different sublevel sets. A *sublevel set* $L_\gamma^-(V)$ of level γ (a *sub- γ set*) of a function $V(t, x)$ is defined as

$$L_\gamma^-(V) \triangleq \{(t, x) \in \mathbb{R}_+ \times \mathbb{X} \mid V(t, x) \leq \gamma\}.$$

Remark 1. Note that we restrict sublevel sets to subsets of a bounded domain \mathbb{X} ensuring that they are bounded as well.

A Primer on Stochastic Calculus

Under our assumptions, the problem (1) has a solution that is a so-called Feller–Dynkin process¹. The dynamics of a Feller–Dynkin process can be described by a linear operator G_π called its *infinitesimal generator*, which in our case admits the following form (Khasminskii 2011):

$$G_\pi \triangleq \frac{\partial}{\partial t} + \sum_{i=1}^l f_\pi(t, x_i) \frac{\partial}{\partial x_i} + \frac{1}{2} \sum_{i=1}^l \sum_{r=1}^k g_{\pi, r}^2(t, x_i) \frac{\partial^2}{\partial x_i^2}. \quad (6)$$

Infinitesimal generator is the stochastic counterpart of gradient, describing the direction of fastest expected increase.

Next, let us recall the definition of a natural filtration and stopping time. The *natural filtration* of \mathcal{F} with respect to a stochastic process $(\eta_t)_{t \geq 0}$ is the smallest σ -algebra $(\mathcal{N}_t)_{t \geq 0}$ on Ω that contains all pre-images of $\mathcal{B}(\mathbb{X})$ -measurable subsets of \mathbb{X} for times s up to t ,

$$\mathcal{N}_t \triangleq \sigma\{\eta_s^{-1}(A) \mid s \in \mathbb{R}_+ \wedge s \leq t \wedge A \in \mathcal{B}(\mathbb{X})\}.$$

The natural filtration represents the information flow generated by a stochastic process: at any moment of time t , the σ -algebra \mathcal{N}_t includes all of the information generated by the process up to that point of time.

A random variable $\tau : \Omega \rightarrow \mathbb{R}_+$ is called an \mathfrak{F} -*stopping time* if $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \in \mathbb{R}_+$. Intuitively, τ is a stopping time if it is possible to determine whether it happened without looking into the future. We can “freeze” a stochastic process at a stopping time (hence the name). Formally, this results in a new version of a stochastic process called a stopped process that is defined for all $\omega \in \Omega$ as

$$\eta_{t \wedge \tau}(\omega) \triangleq \eta_t(\omega) \mathbf{1}_{\{t \leq \tau(\omega)\}} + \eta_\tau(\omega) \mathbf{1}_{\{t > \tau(\omega)\}}.$$

The stopped process follows the original process up to a random time τ , and after that point it stays the same. Stopped processes are of special interest to us because they are used to construct supermartingales.

An \mathfrak{F} -adapted (i.e., such that η_t is always \mathcal{F}_t -measurable) stochastic process $(\eta_t)_{t \geq 0}$ is called an \mathfrak{F} -*supermartingale* if $\mathbb{E}[\eta_t \mid \mathcal{F}_s] \leq \eta_s$ for every pair of times $t > s \geq 0$.

Comparison to Alternative Models

Continuous-time deterministic certificates do not consider the quadratic term in (6) and therefore fail to fully capture the stochastics of a system defined by (1). For example, the solutions to $\dot{x} = x$ and $dX_t = X_t dt + X_t dW_t$ are $x_t = x_0 e^t$ and $X_t = x_0 e^{t/2 + W_t}$; the growth rates (1 and 1/2) are different. Similarly, discrete-time stochastic systems represent the dynamics of a continuous-time system only approximately. The solutions to $x_{t+1} - x_t = -\frac{1}{2}x_t + \eta_t$ where $\eta_t \sim \mathcal{N}(0, 1)$ (i.i.d.) and $dX_t = -\frac{1}{2}X_t dt + dW_t$ are $x_t = 2^{-t}x_0 + \sum_{i=1}^t 2^{i-t}\eta_{i-1} = e^{-t \ln 2}x_0 + \frac{2\sqrt{3}}{3}\sqrt{1-4^{-t}}\tilde{\eta}_t$, $\tilde{\eta}_t \sim \mathcal{N}(0, 1)$ and $X_t = e^{-t/2}x_0 + W_{1-e^{-2t}}$; again, the decay rates ($\ln 2$ and 0.5) differ, leading to different dynamics.

These examples show that existing certificates cannot be directly applied to the continuous-time stochastic setting. To satisfy a RAS specification, we need to find a new type of certificate and prove its validity for our problem.

¹This roughly means that it is right-continuous, exhibits the strong Markov property, and has time-independent transitions.

Reach-Avoid-Stay Certificates

Definition 2. A function $V(t, x) : \mathbb{R}_+ \times \mathbb{R}^l \rightarrow \mathbb{R}$ is called a reach-avoid-stay certificate (RAS-C) if in the domain $\mathbb{R}_{++} \times \mathbb{X}$ it is bounded and twice continuously differentiable with respect to x and continuously differentiable with respect to t , and, moreover, the following conditions hold, given the sets \mathbb{X} and \mathbb{X}_0 , the families of sets $\mathcal{X}_\circlearrowleft$, and \mathcal{X}_* , as well as some positive constants $\alpha_S < \beta_S < \alpha_{RA} < \beta_{RA}$:

1. nonnegativity: $V(t, x) \geq 0$ for all $t \geq 0$ and $x \in \mathbb{X}$;
2. initial condition: $V(0, x_0) \leq \alpha_{RA}$ for all $x_0 \in \mathbb{X}_0$;
3. safety condition: $V(t, x) \geq \beta_{RA}$ for all $(t, x) \in \mathcal{X}_\circlearrowleft$;
4. decrease condition: there exists $\zeta(t)$ satisfying

$$\lim_{t \rightarrow \infty} \int_0^t \zeta(s) ds = \infty \quad \text{and} \quad \zeta(t) > 0, \quad (7)$$

and $G_\pi V(t, x) \leq -\zeta(t)$ for all $(t, x) \in L_{\beta_{RA}}^-(V) \setminus \text{int } \mathcal{X}_*$.

5. goal condition: in the target set \mathcal{X}_* there exists a sub- β_S set, $L_{\beta_S}^-(V)$;
6. stay condition: there exists $\xi(t)$ satisfying

$$\lim_{t \rightarrow \infty} \int_0^t \xi(s) ds = \infty \quad \text{and} \quad \xi(t) > 0, \quad (8)$$

and $G_\pi V(t, x) \leq -\xi(t)$ for all $(t, x) \in (\text{int } \mathcal{X}_* \cap L_{\beta_{RA}}^-(V)) \setminus \text{int } L_{\alpha_S}^-(V)$.

Remark 2. Conditions (7) and (8) are satisfied by any positive constants ζ and ξ . We use a more general definition, because sometimes it is possible to derive better bounds based on the domain knowledge of the problem at hand.

At the core of the definition is the following idea: every process issuing within the sub- α_{RA} set $L_{\alpha_{RA}}^-(V)$ is unlikely to leave the sub- β_{RA} set $L_{\beta_{RA}}^-(V)$, and similarly for β_S and α_S . Based on this property, we can show that the function V decreases in expectation along the sample paths of the state process until it reaches the target. If the value of V is initially at most α_{RA} , we have not left the sub- β_{RA} set with some high probability (to be defined later) and thus have avoided the unsafe states. Similarly, if within the target there exists a subset where the values of V are sufficiently low, and they are expected to decrease outside of this set, the process will be “trapped” within this smaller subset with high probability (again, to be defined later). These properties motivate the conditions of Definition 2, and are outlined in Figure 2. This intuitive explanation of reach-stay-avoid certification can be stated rigorously as the following theorem.

Theorem 1 (RAS-C certifies an RAS specification). A policy π satisfies an RAS specification $(\mathcal{X}_*, \mathcal{X}_\circlearrowleft, \mathbb{X}_0, \varepsilon, \delta)$ if there exists a reach-avoid-stay certificate such that

$$\varepsilon = 1 - \frac{\alpha_{RA}}{\beta_{RA}} \quad \text{and} \quad \delta = 1 - \frac{\alpha_S}{\beta_S}.$$

Remark 3. Note that there are two probabilities ε and δ in the definition of the RAS specification, but four numbers in the definition of RAS-C. Since the probabilities of an RAS specification are defined by the ratios, we can fix either α_{RA} or β_{RA} beforehand, and similarly for α_S or β_S .




Figure 2: A (time-homogenous) example of the sets of Definition 2. A sample trajectory $\bullet \longrightarrow$ first reaches the target set \mathbb{X}_* without staying, and then reaches it again and stays. Our goal is to find the latter moment (with high probability).

We prove Theorem 1 in the technical appendix available in the extended version of the paper. We present a sketch of the proof here.

Proof sketch for Theorem 1. First, we construct a random variable representing the first time the value of RAS-C for a given sample path leaves the interval $[\alpha_S, \beta_{RA}]$:

$$\psi \triangleq \inf_{t \geq 0} \{t \mid V(t, X_{\pi, t}^{0, x_0}) \notin [\alpha_S, \beta_{RA}]\}. \quad (9)$$

Next, we show that it is almost surely a finite stopping time adapted to the state process. Then, we construct the following stopped stochastic process that represents the evolution of the certificate and show that it is a supermartingale.

Definition 3. Given the natural filtration $\mathfrak{N} = (\mathcal{N}_t)_{t \geq 0}$ of a stochastic state process $(X_{\pi, t}^{0, x_0})_{t \geq 0}$, and a \mathfrak{N} -stopping time τ , a τ -stopped reach-stay-avoid supermartingale (τ -RAS-SM) is a stochastic process $(Y_t)_{t \geq 0}$ defined as

$$Y_t \triangleq V(t \wedge \tau, X_{\pi, t \wedge \tau}^{0, x_0}).$$

Lemma 1 (the proof can be found in the technical appendix). ψ -RAS-SM $(Y_t)_{t \geq 0}$ is an \mathfrak{N} -supermartingale.

Finally, we use maximal inequalities for supermartingales to bound the probabilities that when the chosen moment of time is reached, both of the events used in Definition 1 are guaranteed to happen with the required probabilities. \square

Moreover, it follows from the proof of the theorem that our approach can be used to certify either of the reach-avoid or stay properties by themselves due to the following corollary (proof outlined in the technical appendix).

Corollary 1. By omitting conditions 5 and 6, we can use an RAS-C to certify reach-avoidance requirement (3) without the staying property. Similarly, the stay property (4) by itself can be certified by omitting conditions 2 and 3.

Neural Certificate Training & Verification

Finding certificates analytically is a challenging task. Therefore, drawing on advancements in the new field of neural certificates, we employ a deep neural network to find the certificate V_θ as a function of an input state x and trainable

Algorithm 1: RAS Certificate Training

Data: a drift f and diffusion g of (1), a policy π , an RAS specification $(\mathcal{X}_*, \mathcal{X}_\circ, \mathbb{X}_0, \varepsilon, \delta)$, batch size n , regularization multiplier λ , number of cells per dimension m , maximum cell splitting depth k , frequency of verification q .

Result: satisfaction of the specification (Yes/No).

- 1 Initialize the levels $\beta_{\text{RA}} > \alpha_{\text{RA}} > \beta_{\text{S}} > \alpha_{\text{S}} > 0$;
 - 2 create AABB cells \mathcal{C} for the verifier;
 - 3 **for** N epochs **do**
 - 4 sample a batch \mathbb{B} of n points from \mathbb{X} ;
 - 5 find the loss L using (10)–(15);
 - 6 train the parameters θ using the gradient $\nabla_\theta L$;
 - 7 **if every** q steps **then**
 - 8 using IBP, compute lower $V_{\text{low}}(C)$ and upper $V_{\text{up}}(C)$ bounds on $V(x)$ for all cells C ;
 - 9 compute $\hat{\varepsilon}$ via (16) and continue if $\hat{\varepsilon} < \varepsilon$;
 - 10 compute $\hat{\delta}$ via (17) and continue if $\hat{\delta} < \delta$;
 - 11 using IBP, compute the upper bounds $[G_\pi V]_{\text{up}}(C)$ for cells in \mathcal{C}_2 defined by (18);
 - 12 set the verification depth $d \leftarrow 0$;
 - 13 find unverified cells $\mathcal{U} \leftarrow \{C \mid [G_\pi V]_{\text{up}}(C) \geq 0\}$;
 - 14 **while** $|\mathcal{U}| > 0$ and $d < k$ **do**
 - 15 split each cell in \mathcal{U} into 2^d smaller cells;
 - 16 increase the depth $d \leftarrow d + 1$;
 - 17 compute the upper bounds $[G_\pi V]_{\text{up}}(C)$;
 - 18 remove cells where the bound is negative;
 - 19 **if** $\hat{\varepsilon} \geq \varepsilon$ and $\hat{\delta} \geq \delta$ and $|\mathcal{U}| = 0$ **then**
 - 20 **return** Yes;
 - 21 **return** No.
-

network parameters θ . Our method is summarized in Algorithm 1 and Figure 3, and explained in this section.

While the theoretical contributions separate the time $t \in \mathbb{R}_+$ and the state $x \in \mathbb{X}$ variables due to different requirements on differentiability of the certificate with respect to them. When this distinction is of no concern, we can embed the time variable t into the space vector (for example, as $x = (e^{-t}, x_1, \dots, x_l)$) to map it into a compact interval). Moreover, if the system is time-homogeneous, we can drop the dependence of the certificate on time entirely. Technically, this makes the initial condition of Definition 2 more restrictive, as it will require that $V(0, x) \leq \alpha_{\text{RA}}$ for all $(t, x) \in \mathbb{R}_+ \times \mathbb{X}_0$ instead of just for $t = 0$, but the resulting dimensionality reduction simplifies the verification process. For these reasons, we opt to make the presentation more succinct by dropping the time variable from now on.

Training

This approach raises a few issues to be addressed.

First, we need to set the threshold levels. We elect:

$$\alpha_{\text{RA}} = 1, \quad \beta_{\text{RA}} = \frac{\kappa \alpha_{\text{RA}}}{1 - \varepsilon}, \quad \beta_{\text{S}} = 0.9, \quad \alpha_{\text{S}} = \frac{\beta_{\text{S}}(1 - \delta)}{\kappa}.$$

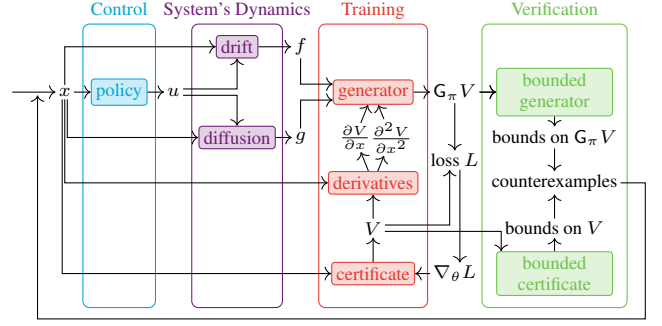


Figure 3: Architecture of the neural certificate training and verification system.

The constant $\kappa > 1$ is chosen to force the training to find larger probabilities ε and δ in order to give some slack to the verification procedure.

Next, to ensure that the certificate is twice continuously differentiable, we employ smooth activation functions such as the hyperbolic tangent and soft-plus. To ensure that the nonnegativity condition of Definition 2 holds, we add a non-negativity transformation to the last layer.

Then, armed with Theorem 1, we train the network as follows. To ensure that the conditions of Definition 2 hold, we sample batches of points from the various sets used therein. We then check how much the points violate the conditions, and sum up the violations into the loss function.

For the initial, safety, and goal conditions, this requires evaluating the certificate which can be done with a simple forward pass through the certificate network.

We combine the decrease and stay conditions by setting the thresholds $\zeta(t)$ and $\xi(t)$ to the same constant ζ . Then we check both conditions over the union of their respective sets. This requires evaluating the generator $G_\pi V(x)$ at each sampled point x as defined by (6). First, we obtain the control signal $u = \pi(x)$ from the policy network (or a function), and compute the drift $f_\pi(x, u)$ and diffusion $g_\pi(x, u)$ of the SDE (1). Next, we find the first and second derivatives of the certificate network analytically using the method of Singla and Feizi (2020). The details are provided in the technical appendix. Since evaluation of the derivatives is costly and only required for some of the points, we separate the certificate and derivative networks.

Finally, we regularize the loss by the product of operator norms of the neural network's weight matrices $\|\cdot\|_{\infty \rightarrow \infty}$ defined as the maximum ℓ_1 -norm of the rows of a matrix. This ensures that the weights of the certificate network remain small, making the verification task easier.

The resulting loss L is equal to

$$L \triangleq L_\circ + L_0 + L_\star + L_\downarrow + R, \quad \text{where} \quad (10)$$

$$L_\circ \triangleq \sum_{x \in \mathbb{B} \cap \mathbb{X}_\circ} (\beta_{\text{RA}} - V_\theta(x))^+, \quad (11)$$

$$L_0 \triangleq \sum_{x_0 \in \mathbb{B} \cap \mathbb{X}_0} (V_\theta(x_0) - \alpha_{\text{RA}})^+, \quad (12)$$

$$L_\star \triangleq \sum_{x \in \mathbb{B} \cap \mathbb{X}_\star} (\beta_{\text{S}} - V_\theta(x))^+, \quad (13)$$

$$L_\downarrow \triangleq \sum_{x \in \mathbb{B} \cap (L_{\beta_{\text{RA}}}^- \setminus L_{\alpha_{\text{S}}}^-)} (G_\pi V_\theta(x) + \zeta)^+, \quad (14)$$

$$R \triangleq \lambda \prod_{i=0}^{N-1} \|\mathbf{W}^{(i)}\|_{\infty \rightarrow \infty}. \quad (15)$$

Verification

To verify that all of the properties of Definition 2 hold, we employ interval bound propagation (IBP) (Xu et al. 2020). IBP, as well as other abstract interpretation methods (Gehr et al. 2018; Kouvaros and Lomuscio 2021; Zhang et al. 2018), have proven effective in verifying correctness of neural networks’ output. To this end, we construct bounded certificate and generator networks that take intervals rather than point values as their inputs. By propagating the input intervals through these networks, we obtain upper and/or lower bounds on the certificate and its generator.

To construct the interval inputs for IBP, we divide the set \mathbb{X} into cells C which are axis-aligned bounding boxes (AABBs) covering the whole space \mathbb{X} , that is, $\mathbb{X} \subseteq \bigcup_{C \in \mathcal{C}} C$. We do this by discretizing each dimension $[x_i^{\min}, x_i^{\max}]$, $i = 1, \dots, l$ of the state space \mathbb{X} into collections \mathcal{X}_i of m equal intervals, $\mathcal{X}_i \triangleq \{[x_i^{\min}, x_i^{\min} + \Delta_i], \dots, [x_i^{\max} - \Delta_i, x_i^{\max}]\}$, $\Delta_i \triangleq (x_i^{\max} - x_i^{\min})/l$, and using all of the possible Cartesian products thereof as m^l cells, $\mathcal{C} \triangleq \mathcal{X}_1 \times \dots \times \mathcal{X}_l$.

For each cell, we compute the lower $V_{\text{low}}(C)$ and upper $V_{\text{up}}(C)$ bounds on $V(x)$ within that cell, $x \in C$. Given the bounds, we find the empirical reach-avoid probability $\hat{\varepsilon}$

$$\hat{\varepsilon} \triangleq (1 - \hat{\alpha}_{\text{RA}}/\hat{\beta}_{\text{RA}})^+, \quad \text{where} \quad (16)$$

$$\hat{\alpha}_{\text{RA}} \triangleq \max_{C: C \cap \mathcal{X}_0 \neq \emptyset} V_{\text{up}}(C), \quad \hat{\beta}_{\text{RA}} \triangleq \min_{C: C \cap \partial \mathcal{X}_* \neq \emptyset} V_{\text{low}}(C).$$

This verifies that the initial set values are at most $\hat{\alpha}_{\text{RA}}$, and the unsafe set values are at least $\hat{\beta}_{\text{RA}}$. Therefore, if the decrease condition is satisfied, the reach-avoid event (3) occurs with probability at least $\hat{\varepsilon}$.

Similarly, for the goal condition, we compute a provable stay probability $\hat{\delta}$ as

$$\hat{\delta} \triangleq (1 - \hat{\alpha}_{\text{S}}/\hat{\beta}_{\text{S}})^+, \quad \text{where} \quad (17)$$

$$\hat{\alpha}_{\text{S}} \triangleq \min_{C: C \cap \mathcal{X}_* \neq \emptyset} V_{\text{up}}(C), \quad \hat{\beta}_{\text{S}} \triangleq \min_{C: C \cap \partial \mathcal{X}_* \neq \emptyset} V_{\text{low}}(C).$$

This guarantees that at least in one of the cells within the target the certificate values go below $\hat{\alpha}_{\text{S}}$, and on the boundary the values are at least $\hat{\beta}_{\text{S}}$.

Finally, we verify the decrease condition by computing upper bounds on the generator values $[G_{\pi} V]_{\text{up}}(C)$ in the set

$$\mathcal{C}_2 \triangleq \{C \mid V_{\text{low}}(C) > \hat{\alpha}_{\text{S}} \wedge V_{\text{up}}(C) \leq \hat{\beta}_{\text{RA}}\}. \quad (18)$$

We then check if any of the cells have non-negative upper bounds. If they do (and we observed this happening often, as the generator bounds involve quadratic terms and as a result they are much looser than the value bounds) the decrease condition in those cells cannot be verified. We split such cells in two along each of the dimensions. We then recompute the cell bounds for each of the 2^l sub-cells. We repeat this until either no counterexample is found, meaning that the decrease condition is verified (as it holds over a collection of cells fully covering the desired subspace), or the maximum iteration depth is achieved.

Finally, if the estimated probabilities $\hat{\varepsilon}$ and $\hat{\delta}$ exceed the specified ones ε and δ , and no cells with $G_{\pi} V_{\text{up}}(C) \geq 0$ are found, we obtained a provably correct RAS-C.

SDE	Time per verification		Number of verifications	
	mean	st. dev.	mean	st. dev.
Inverted pendulum	14.9 s	6.4 s	3.2	0.8
Bivariate GBM	21.1 s	14.4 s	11.2	4.0

Table 1: Runtimes for the experiments. Each data point is based on five randomly generated seeds. The seeds and data for each run are included in the code appendix.

Experiments

To illustrate the practical efficacy of our framework, we synthesize neural certificates for two continuous-time control benchmarks. To the extent of our knowledge, this is the first work on formal verification of neural certificates for continuous-time stochastic control using neural networks. Thus, we do not include comparisons to any other methods.

The computing infrastructure and values of the hyperparameters for all of the experiments and the neural network architectures are given in the technical appendix.

We employ `auto_LiRPA` (Xu et al. 2020) for IBP, which is a library for automatic bound computation with linear relaxation based perturbation analysis. We integrate (1) numerically with stochastic Runge–Kutta method using `torchsde` (Li et al. 2020). Note that this approximation scheme is not employed in Algorithm 1, but only to simulate the sample paths shown in the figures.

We successfully train RAS certificates for two controlled SDEs. The times and numbers of verification phases required for training are summarized in Table 1. More details on each of the problems are given in the technical appendix.

Stochastic Inverted Pendulum

Inverted pendulum (Chang, Roohi, and Gao 2019; Wu et al. 2023) is a classical nonlinear control problem. We disturb the angle by a Wiener process with scale σ , resulting in the following SDEs:

$$d\varphi_t = \left(\frac{g}{L} \sin \theta_t + \frac{M u_t - b \varphi_t}{m L^2} \right) dt + \sigma dW_t, \quad d\theta_t = \varphi_t dt,$$

where φ is the angular velocity. The drift and diffusion are

$$f_{\pi}(\varphi, \theta) = \begin{bmatrix} \frac{g}{L} \sin \theta + \frac{M \pi(\varphi, \theta) - b \varphi}{m L^2} \\ \varphi \end{bmatrix}, \quad g_{\pi}(\varphi, \theta) = \begin{bmatrix} \sigma \\ 0 \end{bmatrix}.$$

The policy $\pi(\varphi, \theta)$ is represented by a neural network. An example of a trained certificate is presented in Figure 1.

Bivariate Geometric Brownian Motion

Geometric Brownian motion (GBM) is one of the most studied stochastic processes and is frequently utilized in financial mathematics. We consider the following controlled version of bivariate GBM:

$$dX_t = (\mu X_t - u_t(X_t)) dt + \sigma(X_t) dW_t, \quad \text{where}$$

$$\mu = \begin{bmatrix} -0.5 & 1 \\ -1 & -0.5 \end{bmatrix} \quad \text{and} \quad \sigma = 0.2 \text{diag}(X_t).$$




Figure 4: An example of a neural supermartingale certificate for the bivariate GBM problem.

We use a stabilizing policy $\pi(x) = -x$ which pushes the system towards the equilibrium point. The resulting certificate is visualized in Figure 4.

Related Work

The concept of representing Lyapunov functions as neural networks was theoretically discussed in seminal work (Long and Bayoumi 1993; Prokhorov 1994). This idea led to the development of numerical machine learning algorithms for neural Lyapunov functions (Serpén 2005; Petridis and Petridis 2006; Noroozi et al. 2008). Building on this, learning algorithms have been extended to encompass barrier functions for avoidance control (Richards, Berkenkamp, and Krause 2018; Dawson et al. 2021), contraction metrics (Sun, Jha, and Fan 2020), and the compositional certification of multi-agent systems (Qin et al. 2021; Zhang et al. 2023b), and the simultaneous training of control policies and neural certificates (Kolter and Manek 2019; Wu et al. 2023), albeit initially without formal soundness guarantees. Our result provides formal soundness guarantees, drawing upon and extending the following related work.

Neural Continuous-Time Lyapunov Certificates. Coupling machine learning with symbolic reasoning techniques like satisfiability modulo theory (SMT) solving and bound propagation made a significant step forward for neural certification with formal guarantees (Chang, Roohi, and Gao 2019; Abate et al. 2021). Initially focused on Lagrangian and asymptotic stability, neural certificates with provable guarantees for continuous-time deterministic systems have since been extended to cover avoidance, reachability, invariance, persistence, and combinations of these properties (Zhao et al. 2020; Takeishi and Kawahara 2021; Zhang et al. 2023a; Edwards, Peruffo, and Abate 2023).

Neural Discrete-Time Lyapunov Certificates. Advances in neural certificates for differential equations have inspired the development of their discrete-time counterparts. Ini-

tially applied to stability and reachability analysis of difference equations and hybrid systems (Chen et al. 2021a), neural discrete-time Lyapunov certificates have been extended to control under avoidance and persistence specifications (Anand and Zamani 2023; Chen et al. 2021b), and have enabled symbolic reasoning with efficient neural network verifiers (Yang et al. 2024; Mandal et al. 2024). As reasoning about discrete-time systems is akin to reasoning about programs, neural certificates have also been applied to termination analysis (Giacobbe, Kroening, and Parsert 2022).

Neural Discrete-Time Supermartingale Certificates. In the presence of stochastic uncertainty, traditional Lyapunov-like certificates are overly conservative and unrealistic; relying on worst-case reasoning, they do not quantify the probability of events to occur. Probabilistic neural certificates, which build on supermartingale-like conditions, were initially developed for almost-sure termination and, analogously, discrete-time stability (Abate, Giacobbe, and Roy 2021; Lechner et al. 2022). These neural supermartingale certificates were later generalized to quantitative reachability, avoidance and reach-avoidance (Mathiesen, Calvert, and Laurenti 2023; Žikelić et al. 2023a; Chatterjee et al. 2023; Ansari pour et al. 2023; Badings et al. 2024). Further advancements have studied the optimization of probability bounds and compositional certification (Abate et al. 2023; Žikelić et al. 2023b).

Neural Continuous-Time Supermartingale Certificates. Neural supermartingale-based techniques have been studied primarily in the context of discrete-time systems. Their extension to continuous-time systems has been solely explored in the context of stability control, and even then, without provable guarantees (Zhang, Zhu, and Lin 2022).

Conclusion & Future Work

We have introduced the first neural supermartingale for continuous-time reasoning with provable guarantees, applicable to combinations of reachability, avoidance, and persistence properties. We have formulated a proof rule (Theorem 1) for a supermartingale certificate for continuous-time systems and, by integrating machine learning with symbolic reasoning, we have fully automated their construction. We have built a prototype and demonstrated that our method is practically effective on continuous-time systems with neural policies, extending the state of the art in algorithmic verification for provably safe stochastic nonlinear control.

Our framework is open to extension towards the automated synthesis of neural controllers alongside their certificates, without the need to pre-initialize a control policy or guide it with other methods. We foresee integration with advanced symbolic reasoning techniques based on adaptive discretization (Žikelić et al. 2023a; Badings et al. 2024), adversarial attack algorithms (Yang et al. 2024; Wu et al. 2023), and generalization to more expressive temporal logic requirements (Abate, Giacobbe, and Roy 2024; Nadali et al. 2024; Giacobbe et al. 2024).

Acknowledgements

This research was funded by the NWO Veni grant Explainable Monitoring (222.119), and the Advanced Research + Invention Agency (ARIA) under the Safeguarded AI programme. This work was done in part while Anna Lukina was visiting the Simons Institute for the Theory of Computing.

References

- Abate, A.; Ahmed, D.; Giacobbe, M.; and Peruffo, A. 2021. Formal Synthesis of Lyapunov Neural Networks. *IEEE Control. Syst. Lett.*, 5(3): 773–778.
- Abate, A.; Edwards, A.; Giacobbe, M.; Punchihewa, H.; and Roy, D. 2023. Quantitative Verification with Neural Networks. In *CONCUR*, volume 279 of *LIPICs*, 22:1–22:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Abate, A.; Giacobbe, M.; and Roy, D. 2021. Learning Probabilistic Termination Proofs. In *CAV (2)*, volume 12760 of *Lecture Notes in Computer Science*, 3–26. Springer.
- Abate, A.; Giacobbe, M.; and Roy, D. 2024. Stochastic Omega-Regular Verification and Control with Supermartingales. In *CAV (3)*, volume 14683 of *Lecture Notes in Computer Science*, 395–419. Springer.
- Akintunde, M. E.; Botoeva, E.; Kouvaros, P.; and Lomuscio, A. 2022. Formal verification of neural agents in non-deterministic environments. *Auton. Agents Multi Agent Syst.*, 36(1): 6.
- Anand, M.; and Zamani, M. 2023. Formally Verified Neural Network Control Barrier Certificates for Unknown Systems. *IFAC-PapersOnLine*, 56(2): 2431–2436.
- Ansari-pour, M.; Chatterjee, K.; Henzinger, T. A.; Lechner, M.; and Žikelić, D. 2023. Learning Provably Stabilizing Neural Controllers for Discrete-Time Stochastic Systems. In *ATVA (1)*, volume 14215 of *Lecture Notes in Computer Science*, 357–379. Springer.
- Bacci, E.; Giacobbe, M.; and Parker, D. 2021. Verifying Reinforcement Learning up to Infinity. In *IJCAI*, 2154–2160. ijcai.org.
- Badings, T.; Koops, W.; Junges, S.; and Jansen, N. 2024. Learning-Based Verification of Stochastic Dynamical Systems with Neural Network Policies. arXiv:2406.00826.
- Blumenthal, R. M. C.; and Getoor, R. K. 1968. *Markov Processes and Potential Theory*. Pure and applied mathematics: A series of monographs and textbooks. Academic Press. ISBN 9780121078508.
- Chang, Y.-C.; Roohi, N.; and Gao, S. 2019. Neural Lyapunov Control. In *NeurIPS*, 3240–3249.
- Chatterjee, K.; Henzinger, T. A.; Lechner, M.; and Žikelić, D. 2023. A Learner-Verifier Framework for Neural Network Controllers and Certificates of Stochastic Systems. In *TACAS (1)*, volume 13993 of *Lecture Notes in Computer Science*, 3–25. Springer.
- Chen, S.; Fazlyab, M.; Morari, M.; Pappas, G. J.; and Preciado, V. M. 2021a. Learning Lyapunov functions for hybrid systems. In *HSCC*, 13:1–13:11. ACM.
- Chen, S.; Fazlyab, M.; Morari, M.; Pappas, G. J.; and Preciado, V. M. 2021b. Learning Region of Attraction for Non-linear Systems. In *CDC*, 6477–6484. IEEE.
- Dawson, C.; Gao, S.; and Fan, C. 2023. Safe Control With Learned Certificates: A Survey of Neural Lyapunov, Barrier, and Contraction Methods for Robotics and Control. *IEEE Trans. Robotics*, 39(3): 1749–1767.
- Dawson, C.; Qin, Z.; Gao, S.; and Fan, C. 2021. Safe Non-linear Control Using Robust Neural Lyapunov-Barrier Functions. In *CoRL*, volume 164 of *Proceedings of Machine Learning Research*, 1724–1735. PMLR.
- Edwards, A.; Peruffo, A.; and Abate, A. 2023. A General Verification Framework for Dynamical and Control Models via Certificate Synthesis. *CoRR*, abs/2309.06090.
- Gehr, T.; Mirman, M.; Drachler-Cohen, D.; Tsankov, P.; Chaudhuri, S.; and Vechev, M. T. 2018. AI2: Safety and Robustness Certification of Neural Networks with Abstract Interpretation. In *IEEE Symposium on Security and Privacy*, 3–18. IEEE Computer Society.
- Giacobbe, M.; Kroening, D.; Pal, A.; and Tautschnig, M. 2024. Neural Model Checking. In *NeurIPS*.
- Giacobbe, M.; Kroening, D.; and Parsert, J. 2022. Neural Termination Analysis. In *ESEC/SIGSOFT FSE*, 633–645. ACM.
- Kalman, R. E.; and Bertram, J. E. 1959. Control system analysis and design via the second method of Lyapunov: (I) continuous-time systems (II) discrete time systems. *IRE Transactions on Automatic Control*, 4(3): 112.
- Katz, G.; Barrett, C. W.; Dill, D. L.; Julian, K.; and Kochenderfer, M. J. 2017. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks. In *CAV (1)*, volume 10426 of *Lecture Notes in Computer Science*, 97–117. Springer.
- Khasminskii, R. 2011. *Stochastic Stability of Differential Equations*. Stochastic Modelling and Applied Probability. Springer.
- Kolter, J. Z.; and Manek, G. 2019. Learning Stable Deep Dynamics Models. In *NeurIPS*, 11126–11134.
- Kouvaros, P.; and Lomuscio, A. 2021. Towards Scalable Complete Verification of ReLU Neural Networks via Dependency-based Branching. In *IJCAI*, 2643–2650.
- Lechner, M.; Žikelić, D.; Chatterjee, K.; and Henzinger, T. A. 2022. Stability Verification in Stochastic Control Systems via Neural Network Supermartingales. In *AAAI*, 7326–7336. AAAI Press.
- Li, X.; Wong, T.-K. L.; Chen, R. T. Q.; and Duvenaud, D. 2020. Scalable Gradients for Stochastic Differential Equations. In *AISTATS*, volume 108 of *Proceedings of Machine Learning Research*, 3870–3882. PMLR.
- Long, Y.; and Bayoumi, M. 1993. Feedback stabilization: control Lyapunov functions modelled by neural networks. In *CDC*, 2812–2814. IEEE.
- Mandal, U.; Amir, G.; Wu, H.; Daukantas, I.; Newell, F. L.; Ravaioli, U. J.; Meng, B.; Durling, M.; Ganai, M.; Shim, T.; Katz, G.; and Barrett, C. 2024. Formally Verifying Deep Reinforcement Learning Controllers with Lyapunov Barrier Certificates. arXiv:2405.14058.

- Mathiesen, F. B.; Calvert, S. C.; and Laurenti, L. 2023. Safety Certification for Stochastic Systems via Neural Barrier Functions. *IEEE Control. Syst. Lett.*, 7: 973–978.
- Nadali, A.; Murali, V.; Trivedi, A.; and Zamani, M. 2024. Neural Closure Certificates. In *AAAI*, 21446–21453. AAAI Press.
- Noroozi, N.; Karimaghaee, P.; Safaei, F.; and Javadi, H. 2008. Generation of Lyapunov functions by neural networks. In *Proceedings of the World Congress on Engineering*.
- Papachristodoulou, A.; and Prajna, S. 2002. On the construction of Lyapunov functions using the sum of squares decomposition. In *CDC*, 3482–3487. IEEE.
- Petridis, V.; and Petridis, S. 2006. Construction of Neural Network Based Lyapunov Functions. In *IJCNN*, 5059–5065. IEEE.
- Prajna, S.; Jadbabaie, A.; and Pappas, G. J. 2004. Stochastic safety verification using barrier certificates. In *CDC*, 929–934. IEEE.
- Prajna, S.; Jadbabaie, A.; and Pappas, G. J. 2007. A Framework for Worst-Case and Stochastic Safety Verification Using Barrier Certificates. *IEEE Trans. Autom. Control.*, 52(8): 1415–1428.
- Prokhorov, D. V. 1994. A Lyapunov machine for stability analysis of nonlinear systems. In *ICNN*, 1028–1031. IEEE.
- Qin, Z.; Zhang, K.; Chen, Y.; Chen, J.; and Fan, C. 2021. Learning Safe Multi-agent Control with Decentralized Neural Barrier Certificates. In *ICLR*. OpenReview.net.
- Richards, S. M.; Berkenkamp, F.; and Krause, A. 2018. The Lyapunov Neural Network: Adaptive Stability Certification for Safe Learning of Dynamical Systems. In *CoRL*, volume 87 of *Proceedings of Machine Learning Research*, 466–476. PMLR.
- Serpen, G. 2005. Empirical approximation for Lyapunov functions with artificial neural nets. In *IJCNN*, 735–740. IEEE.
- Singla, S.; and Feizi, S. 2020. Second-Order Provable Defenses against Adversarial Attacks. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 8981–8991. PMLR.
- Sun, D.; Jha, S.; and Fan, C. 2020. Learning Certified Control Using Contraction Metric. In *CoRL*, volume 155 of *Proceedings of Machine Learning Research*, 1519–1539. PMLR.
- Takeishi, N.; and Kawahara, Y. 2021. Learning Dynamics Models with Stable Invariant Sets. In *AAAI*, 9782–9790. AAAI Press.
- Topcu, U.; Packard, A. K.; and Seiler, P. J. 2008. Local stability analysis using simulations and sum-of-squares programming. *Autom.*, 44(10): 2669–2675.
- Tran, H.-D.; Yang, X.; Lopez, D. M.; Musau, P.; Nguyen, L. V.; Xiang, W.; Bak, S.; and Johnson, T. T. 2020. NNV: The Neural Network Verification Tool for Deep Neural Networks and Learning-Enabled Cyber-Physical Systems. In *CAV (1)*, volume 12224 of *Lecture Notes in Computer Science*, 3–17. Springer.
- Wu, H.; Isac, O.; Zeljic, A.; Tagomori, T.; Daggitt, M. L.; Kokke, W.; Refaeli, I.; Amir, G.; Julian, K.; Bassan, S.; Huang, P.; Lahav, O.; Wu, M.; Zhang, M.; Komendantskaya, E.; Katz, G.; and Barrett, C. W. 2024. Marabou 2.0: A Versatile Formal Analyzer of Neural Networks. In *CAV (2)*, volume 14682 of *Lecture Notes in Computer Science*, 249–264. Springer.
- Wu, J.; Clark, A.; Kantaros, Y.; and Vorobeychik, Y. 2023. Neural Lyapunov Control for Discrete-Time Systems. In *NeurIPS*.
- Xu, K.; Shi, Z.; Zhang, H.; Wang, Y.; Chang, K.-W.; Huang, M.; Kailkhura, B.; Lin, X.; and Hsieh, C.-J. 2020. Automatic Perturbation Analysis for Scalable Certified Robustness and Beyond. In *NeurIPS*.
- Yang, L.; Dai, H.; Shi, Z.; Hsieh, C.-J.; Tedrake, R.; and Zhang, H. 2024. Lyapunov-Stable Neural Control for State and Output Feedback: A Novel Formulation. In *ICML*. OpenReview.net.
- Zhang, H.; Weng, T.-W.; Chen, P.-Y.; Hsieh, C.-J.; and Daniel, L. 2018. Efficient Neural Network Robustness Certification with General Activation Functions. In *NeurIPS*, 4944–4953.
- Zhang, H.; Wu, J.; Vorobeychik, Y.; and Clark, A. 2023a. Exact Verification of ReLU Neural Control Barrier Functions. In *NeurIPS*.
- Zhang, J.; Zhu, Q.; and Lin, W. 2022. Neural Stochastic Control. In *NeurIPS*.
- Zhang, S.; Xiu, Y.; Qu, G.; and Fan, C. 2023b. Compositional Neural Certificates for Networked Dynamical Systems. In *LADC*, volume 211 of *Proceedings of Machine Learning Research*, 272–285. PMLR.
- Zhao, H.; Zeng, X.; Chen, T.; and Liu, Z. 2020. Synthesizing barrier certificates using neural networks. In *HSCC*, 25:1–25:11. ACM.
- Žikelić, D.; Lechner, M.; Henzinger, T. A.; and Chatterjee, K. 2023a. Learning Control Policies for Stochastic Systems with Reach-Avoid Guarantees. In *AAAI*, 11926–11935. AAAI Press.
- Žikelić, D.; Lechner, M.; Verma, A.; Chatterjee, K.; and Henzinger, T. A. 2023b. Compositional Policy Learning in Stochastic Control Systems with Formal Guarantees. In *NeurIPS*.