
Formal Control Synthesis for Stochastic Neural Network Dynamic Models

Steven Adams¹ Morteza Lahijanian² Luca Laurenti¹

Abstract

Neural networks (NNs) are emerging as powerful tools to represent the dynamics of control systems with complicated physics or black-box components. Due to complexity of NNs, however, existing methods are unable to synthesize complex behaviors with guarantees for *NN dynamic models* (NNDMs). This work introduces a control synthesis framework for stochastic NNDMs with performance guarantees. The focus is on specifications expressed in *linear temporal logic interpreted over finite traces* (LTLf), and the approach is based on finite abstraction. Specifically, we leverage recent techniques for convex relaxation of NNs to formally abstract a NNDM into an *interval Markov decision process* (IMDP). Then, a strategy that maximizes the probability of satisfying a given specification is synthesized over the IMDP and mapped back to the underlying NNDM. We show that the process of abstracting NNDMs to IMDPs reduces to a set of convex optimization problems, hence guaranteeing efficiency. We also present an adaptive refinement procedure that makes the framework scalable. On several case studies, we illustrate that our framework is able to provide non-trivial guarantees of correctness for NNDMs with architectures of up to 5 hidden layers and hundreds of neurons per layer.

1. Introduction

Autonomous systems are becoming increasingly complex, often including black-box components and performing complex tasks in the presence of uncertainty. In this context, because of their data efficiency and representation power, deep neural networks (NNs) can be a transformative technology:

^{*}Equal contribution ¹Delft Center for Systems and Control, TU Delft, The Netherlands ²Dept. of Aero. Eng. Sciences and Computer Science, Univ. of Colorado Boulder, USA. Correspondence to: steven.adams <s.j.l.adams@tudelft.nl>.

Presented at the 2nd Workshop on Formal Verification of Machine Learning, co-located with the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA., 2023. Copyright 2023 by the author(s).

NNs have already achieved state-of-the-art performance to model and control dynamical systems in various fields, including reinforcement learning (RL) (Nagabandi & et al., 2018). However, employing NNs in *safety-critical* applications, such as UAVs, where failures may have catastrophic effects, remains a major challenge due to limitations of existing methods to provide performance guarantees. This work focuses on this challenge and develops a correct-by-construction synthesis framework for systems with *NN dynamic models* (NNDMs)

To achieve complex behaviors with strong guarantees, formal synthesis for control systems have been well-studied in recent years (Tabuada, 2009; Lahijanian et al., 2015; Doyen et al., 2018). These methods use expressive formal languages such as *linear temporal logic* with infinite (LTL) (Baier & Katoen, 2008) or *finite* (LTLf) (De Giacomo & Vardi, 2013) interpretation over traces, to specify complex behaviors. For synthesis, control barrier functions (CBF) (Xu et al., 2015) and abstraction methods are proposed (Lahijanian et al., 2015). CBFs allow the use of continuous dynamics but are typically limited to invariant-set problems. The abstraction methods apply rigorous techniques to represent the dynamics as a finite (Kripke) models. Then, by utilizing model-checking-like algorithms on the abstraction, they synthesize controllers that achieve the specification. A key aspect is that both of these methods generally relies on (simple) analytical models. For modern systems, however, such models are often unavailable due to, e.g., complexity of the physics or black-box components.

For their accurate predictive ability, NNDMs are already employed to model dynamical systems (Raissi et al., 2019) and enhance controller training in RL frameworks (Nagabandi & et al., 2018). In those works, a NN model of the system is trained in closed-loop with a NN controllers, which can be concatenated in a single NN representing the dynamics of the closed-loop system. These benefits have motivated the recent development of methods for formal analysis of NNDM properties (Wei & Liu, 2021; Wicker et al., 2021; Fazlyab et al., 2021), extending verification algorithms for NNs (Xu et al., 2020) to support temporal properties. Nevertheless, these methods are limited to simple safety properties and often neglect noise in the dynamics. As a consequence, the state-of-the-art techniques for NNDM are still unable to achieve complex behaviors with guarantees.

In this work, we close the gap by introducing a control synthesis framework for stochastic NNDMs to achieve a complex specification with formal guarantees. Our approach is based on finite abstraction, and we use LTLf as the specification language which has the same expressivity as LTL, but specifies finite behaviors, making it an appropriate language for stochastic models. In particular, we leverage recent convex relaxation techniques for NNs (Xu et al., 2020) to build piece-wise linear functions that under- and over-approximate the NNDM and construct the abstraction as an interval Markov decision process (IMDP) (Givan et al., 2000). Critically, we show that this discretization-based method only requires solving a set of convex optimization problems, which can be reduced to evaluation of an analytical function on a finite set of points, resulting in efficient abstraction procedure. Then, we use existing tools to synthesize a control strategy that optimizes the probability of satisfying a given specification while guaranteeing robustness against uncertainties due to dynamics approximation and discretization. To ensure scalability, we present a refinement algorithm that iteratively reduces uncertainty in a targeted manner. Finally, we illustrate the efficacy of our framework in several case studies.

In summary, the contributions of this paper are: (i) a novel framework for formal synthesis for stochastic NNDMs with complex specifications, (ii) an efficient finite abstraction technique for NNDMs, (iii) an adaptive refinement algorithm for uncertainty reduction, and (iv) illustration of the efficacy and scalability of the framework on a set of rich case studies with complex NNDMs, whose architecture include up to five hidden layers and hundreds of neurons per layer.

2. Problem Formulation

We consider the following stochastic *neural network dynamic model* (NNDM):

$$\mathbf{x}_{k+1} = f_a^w(\mathbf{x}_k) + \mathbf{v}_k, \quad \mathbf{v}_k \sim \mathcal{N}(0, \text{Cov}_v), \quad (1)$$

where $k \in \mathbb{N}$, $\mathbf{x}_k, \mathbf{v}_k \in \mathbb{R}^n$, and $a \in A = \{a_1, \dots, a_m\}$ is a finite set of actions. For every a , $f_a^w : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a (trained) feed-forward NN with continuous activation functions, where w denotes the maximum likelihood weights. The noise term \mathbf{v}_k is a random variable with stationary non-degenerate Gaussian distribution with zero mean and covariance $\text{Cov}_v \in \mathbb{R}^{n \times n}$. Intuitively, \mathbf{x}_k is a discrete-time stochastic process whose time evolution is given by iterative predictions of various NNs. We remark that models such as Process (1) are increasingly employed in both robotics and biological systems for both model representation and NN controller training with, e.g., state-of-the-art model-based RL techniques (Nagabandi & et al., 2018; Zhao et al., 2020; Raissi et al., 2019). For instance, Process (1) can represent

a NN model in closed loop with different (possibly NN) feedback controllers, and the role of actions a is to switch between different controllers.

Let $\omega_{\mathbf{x}}^N = x_0 \xrightarrow{a_0} x_1 \xrightarrow{a_1} \dots \xrightarrow{a_{N-1}} x_N$ be a finite path of Process (1) of length $N \in \mathbb{N}$ and $\Omega_{\mathbf{x}}^{\text{fin}}$ be the set of all finite paths. Paths of infinite length and the set of all paths of infinite lengths are denoted by $\omega_{\mathbf{x}}$ and $\Omega_{\mathbf{x}}$, respectively, with $\omega_{\mathbf{x}}(k)$ denoting the state of $\omega_{\mathbf{x}}$ at time k . Given a finite path, a *switching strategy* $\pi_{\mathbf{x}} : \Omega_{\mathbf{x}}^{\text{fin}} \rightarrow A$ chooses the next action of Process (1). The set of all switching strategies is denoted by $\Pi_{\mathbf{x}}$. For $a \in A$, $X \subseteq \mathbb{R}^n$, and $x \in \mathbb{R}^n$, we call

$$T^a(X|x) = \int_X \mathcal{N}(\bar{x} | f_a^w(x), \text{Cov}_v) d\bar{x} \quad (2)$$

the *transition kernel* of Process (1) under action a , where $\mathcal{N}(\cdot | f_a^w(x), \text{Cov}_v)$ is a normal distribution with mean $f_a^w(x)$ and covariance Cov_v . For a strategy $\pi_{\mathbf{x}}$, Process (1) defines a probability measure P which is uniquely defined by T^a and by the initial conditions (Bertsekas & Shreve, 1996) s.t. for every $k > 0$,

$$P[\omega_{\mathbf{x}}(k+1) \in X | \omega_{\mathbf{x}}(k) = x, \pi_{\mathbf{x}}(\omega_{\mathbf{x}}(k)) = a] = T^a(X|x).$$

We are interested in the behavior of Process (1) in compact set $X_{\text{safe}} \subseteq \mathbb{R}^n$ with respect to the regions of interest in $R = \{\mathbf{r}_1, \dots, \mathbf{r}_R\}$, where $\mathbf{r}_i \subseteq X_{\text{safe}}$. To define properties over R , we associate to each region \mathbf{r}_i the atomic proposition \mathbf{p}_i such that \mathbf{p}_i is true iff $x \in \mathbf{r}_i$. The set of atomic propositions is given by AP , and the labeling function $L : X \rightarrow 2^{AP}$ returns the set of atomic propositions that are true at each state. Then, we define the observation of path $\omega_{\mathbf{x}}^N$ to be $\rho = \rho_0 \rho_1 \dots \rho_N$, where $\rho_i = L(\omega_{\mathbf{x}}^N(i))$ for all $i \leq N$.

To express the temporal properties of Process (1), we consider LTLf, which is an expressive language to specify finite behaviors, and hence, appropriate for stochastic systems.

Definition 2.1. An LTLf formula is built from a set of propositional symbols AP and is closed under the boolean connectives as well as the “until” operator \mathcal{U} : $\phi := \top | \mathbf{p} | \neg\phi | \phi_1 \wedge \phi_2 | \phi_1 \mathcal{U} \phi_2$, where $\mathbf{p} \in AP$.

The common temporal operators “eventually” (\mathcal{F}) and “globally” (\mathcal{G}) are defined as: $\mathcal{F}\phi = \top \mathcal{U} \phi$ and $\mathcal{G}\phi = \neg \mathcal{F} \neg \phi$. The semantics of LTLf can be found in (De Giacomo & Vardi, 2013). We say a path $\omega_{\mathbf{x}}$ satisfies ϕ , denoted by $\omega_{\mathbf{x}} \models \phi$, if a prefix of its observation satisfies ϕ (Wells et al., 2020).

Problem 2.2 (Control synthesis). *Given a NNDM as defined in Process (1), a compact set X_{safe} , and an LTLf formula ϕ defined over the regions of interest in X_{safe} , find a switching strategy $\pi_{\mathbf{x}}^*$ that maximizes the probability that a path $\omega_{\mathbf{x}} \in \Omega_{\mathbf{x}}$ of Process (1) satisfies ϕ while remaining in X_{safe} .*

To solve Problem 2.2, we abstract Process (1) into a finite Markov model, where the stochastic nature of Process (1)

and the error corresponding to the discretization of the space are formally modelled as uncertainties. The abstraction process involves the computation of bounds on the transition probabilities between different regions of the state space. In Section 4.1, we show that by using linear functions that locally under and overapproximate the NN-dynamics, these bounds can be efficiently computed by solving convex optimization problems. For the resulting Markov model, we then synthesize a strategy that maximizes the probability that the paths of the Markov model satisfy ϕ and can be mapped onto Process (1). Finally, in Section 5.1, we develop a refinement scheme that iteratively builds a finer abstraction based on the synthesis results by reducing the conservatism induced by the approximation bounds of f_a^w .

3. Preliminaries

Notation: We denote by $x^{(l)}$ the l -th element of vector $x \in \mathbb{R}^n$. Further, for convex region $X \subset \mathbb{R}^n$, we denote by $X^{[l]} \subset \mathbb{R}$ the interval of values of X in the l -th dimension, i.e., $X^{[l]} = \{x^{(l)} \mid x \in X\}$. Given a linear transformation function (matrix) $\mathcal{T} \in \mathbb{R}^{n \times n}$, the image of region $X \subset \mathbb{R}^n$ under \mathcal{T} is defined as $Im(X, \mathcal{T}) = \{\mathcal{T}x \mid x \in X\}$. The post image of region X under action a of Process (1) and \mathcal{T} is defined as $Post(X, \mathcal{T}, a) = \{\mathcal{T}f_a^w(x) \mid x \in X\}$. For vectors $x, x' \in \mathbb{R}^n$, we denote by $rect(x, x')$ the axis-aligned hyper-rectangle that is defined by the intervals $r_1 \times r_2 \times \dots \times r_n$, where $r_l = [\min(x^{(l)}, x'^{(l)}), \max(x^{(l)}, x'^{(l)})]$. In addition, for region $X \subset \mathbb{R}^n$, we denote by $rect(X)$ the hyper-rectangular overapproximation of X , i.e., $rect(X) = rect(\tilde{x}, \hat{x})$, where $\tilde{x}^{(l)} = \inf(X^{[l]})$ and $\hat{x}^{(l)} = \sup(X^{[l]})$ for every $l \in \{0, \dots, n\}$. Lastly, we define a proper linear transformation function as follows.

Definition 3.1 (Proper Transformation). For $X \subset \mathbb{R}^n$, the transformation matrix $\mathcal{T} \in \mathbb{R}^{n \times n}$ is proper w.r.t. X if $Im(X, \mathcal{T})$ is an axis-aligned hyper-rectangle.

Note that, as \mathcal{T} is a linear transformation, X must necessarily be a convex polytope in order for \mathcal{T} to be proper (Ralph Tyrell Rockafellar, 2015).

Interval Markov Decision Processes: We utilize Interval Markov Decision Processes (IMDPs), also called Bounded MDPs (Givan et al., 2000), to abstract Process (1). IMDPs are a generalized class of MDPs that allows for a range of transition probabilities between states.

Definition 3.2 (IMDP). An interval Markov decision process (IMDP) is a tuple $\mathcal{I} = (Q, A, \check{P}, \hat{P}, AP_{\mathcal{I}}, L_{\mathcal{I}})$, where

- Q is a finite set of states,
- A is a finite set of actions available in each state $q \in Q$.
- $\check{P} : Q \times A \times Q \rightarrow [0, 1]$ is a function, where $\check{P}(q, a, q')$ defines the lower bound of the transition probability from state $q \in Q$ to state $q' \in Q$ under action $a \in A$,

- $\hat{P} : Q \times A \times Q \rightarrow [0, 1]$ is a function, where $\hat{P}(q, a, q')$ defines the upper bound of the transition probability from state $q \in Q$ to state $q' \in Q$ under action $a \in A$.

- $AP_{\mathcal{I}}$ is a finite set of atomic propositions,

- $L_{\mathcal{I}} : Q \rightarrow 2^{AP}$ is a labeling function assigning to each state $q \in Q$ a subset of $AP_{\mathcal{I}}$.

For all $q, q' \in Q$ and $a \in A$, it holds that $\check{P}(q, a, q') \leq \hat{P}(q, a, q')$ and $\sum_{q' \in Q} \check{P}(q, a, q') \leq 1 \leq \sum_{q' \in Q} \hat{P}(q, a, q')$. A path $\omega_{\mathcal{I}}$ of an IMDP is a sequence of states $\omega_{\mathcal{I}} = q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \xrightarrow{a_2} \dots$ such that $\check{P}(q_k, a, q_{k+1}) > 0$ for all $k \in \mathbb{N}$. We denote the last state of a finite path $\omega_{\mathcal{I}}^{\text{fin}}$ by $last(\omega_{\mathcal{I}}^{\text{fin}})$ and the set of all finite and infinite paths by $\Omega_{\mathcal{I}}^{\text{fin}}$ and $\Omega_{\mathcal{I}}$, respectively. A strategy of an IMDP $\pi_{\mathcal{I}} : \Omega_{\mathcal{I}}^{\text{fin}} \rightarrow A$ maps a finite path $\omega_{\mathcal{I}}^{\text{fin}} \in \Omega_{\mathcal{I}}^{\text{fin}}$ of \mathcal{I} onto an action in A . The set of all strategies is denoted by $\Pi_{\mathcal{I}}$. Let $\mathcal{D}(Q)$ denote the set of discrete probability distributions over Q . Given a strategy $\pi_{\mathcal{I}}$, the IMDP reduces to a set of infinitely many Markov chains defined by the transition probability bounds of the IMDP. An adversary chooses a feasible distribution from this set at each state and reduces the IMDP to a Markov chain. In particular, an adversary is defined as a function $\xi : \Omega_{\mathcal{I}}^{\text{fin}} \times A \rightarrow \mathcal{D}(Q)$ that, for each finite path $\omega_{\mathcal{I}}^{\text{fin}} \in \Omega_{\mathcal{I}}^{\text{fin}}$, state $q = last(\omega_{\mathcal{I}}^{\text{fin}})$, and action $a \in A$, assigns a feasible distribution γ_q^a which satisfies $\check{P}(q, a, q') \leq \gamma_q^a(q') \leq \hat{P}(q, a, q')$. The set of all adversaries is denoted by Ξ .

4. IMDP Abstraction

In order to solve Problem 2.2, we first abstract Process (1) into an IMDP $\mathcal{I} = (Q, A, \check{P}, \hat{P}, AP_{\mathcal{I}}, L_{\mathcal{I}})$. To do that, similarly as in (Cauchi et al., 2019), we discretize X_{safe} in such way that the transition kernel in (2) can be computed analytically. Let \mathcal{T} be the Mahalanobis transformation $\mathcal{T} = \Lambda^{-\frac{1}{2}} \mathcal{V}^T$, where $\Lambda = \mathcal{V}^T \text{Cov}_v \mathcal{V}$ is a diagonal matrix whose entries are the eigenvalues of Cov_v , and \mathcal{V} is the corresponding orthogonal (eigenvector) matrix. Then, the distribution of $\mathcal{T}x_{k+1}$ given $x_k = x$ under action a becomes $\mathcal{N}(\cdot \mid \mathcal{T}f_a^w(x), I)$, where I is the identity matrix. Consequently, given a region $X \subset \mathbb{R}^n$ for which \mathcal{T} is a proper transformation (i.e., $Im(X, \mathcal{T}) = rect(\tilde{x}, \hat{x})$), we obtain that $T^a(X|x) = g(\mathcal{T}f_a^w(x))$, where

$$g(z) = \frac{1}{2^n} \prod_{l=1}^n \left(\text{erf} \left(\frac{z^{(l)} - \tilde{x}^{(l)}}{\sqrt{2}} \right) - \text{erf} \left(\frac{z^{(l)} - \hat{x}^{(l)}}{\sqrt{2}} \right) \right), \quad (3)$$

and $\text{erf}(\cdot)$ is the error function. Hence, we discretize X_{safe} by using a grid in $Im(X_{\text{safe}}, \mathcal{T})$, and denote by $Q_s = \{q_1, \dots, q_{|Q_s|}\}$ the resulting set of regions. To each cell q_i , we associate a state of the IMDP \mathcal{I} . We overload the notation by using q_i for both a region in X_{safe} and a state of \mathcal{I} . Then,

the set of states of \mathcal{I} is defined as $Q = Q_s \cup \{q_u\}$, where q_u denotes the remainder of the state space, i.e., $\mathbb{R}^n \setminus X_{\text{safe}}$.

We define the set of actions of \mathcal{I} to be the set of actions A of Process (1). To ensure a correct abstraction of Process (1), we assume a discretization of X_{safe} that respects the regions of interest in R , i.e., $\forall r \in R, \exists Q_r \subseteq Q$ such that $\cup_{q \in Q_r} q = r$. Under this assumption, the set of atomic propositions $AP_{\mathcal{I}}$ is equal to AP . We define the labeling function $L_{\mathcal{I}}$ with $L_{\mathcal{I}}(q) = L(x)$ for any choice of $x \in q$.

To compute the transition probability bounds \check{P} and \hat{P} for all $q, q' \in Q_s$ and $a \in A$, we need to derive the following bounds, which are the subject of Section 4.1 and 4.2:

$$\check{P}(q, a, q') \leq \min_{x \in q} T^a(q'|x), \quad \hat{P}(q, a, q') \geq \max_{x \in q} T^a(q'|x) \quad (4)$$

The probability interval for transitioning to the state $q_u \in Q$, i.e., the region outside of X_{safe} , is given by

$$\check{P}(q, a, q_u) \leq 1 - \max_{x \in q} T^a(X_{\text{safe}}|x)$$

$$\hat{P}(q, a, q_u) \geq 1 - \min_{x \in q} T^a(X_{\text{safe}}|x),$$

for all $a \in A$ and $q \in Q_s$. Since we are not interested in the behavior of Process (1) outside X_{safe} , we make q_u absorbing, i.e., $\check{P}(q_u, a, q_u) = \hat{P}(q_u, a, q_u) = 1$ for all $a \in A$.

4.1. Transition Probability Bounds Computation

We derive an efficient and scalable procedure for the computation of the bounds in (4). Recall that the discretization procedure described above enables to write transition kernel T^a as the product of erf in (3). Hence, the optimization problems in (4) can be performed on (3), i.e., for $q', q \in Q_s$ to bound $T^a(q'|x)$, we can optimize $g(z)$ over $z \in \text{Post}(q, \mathcal{T}, a)$. However, the exact computation of $\text{Post}(q, \mathcal{T}, a)$ is intractable since NN-dynamics are inherently nonconvex. Hence, we instead seek to overapproximate $\text{Post}(q, \mathcal{T}, a)$ by recursively finding linear functions on the NN-structure that under- and overapproximate $f_a^w(x)$ for all $x \in q \in Q$, as shown in (Xu et al., 2020). We can then use these linear functions to bound $\text{Post}(q, \mathcal{T}, a)$ for all $q \in Q_s$ as shown in the following proposition.

Proposition 4.1. *For Process (1) under action a , region $q \subset \mathbb{R}^n$, and proper transformation matrix \mathcal{T} w.r.t. q , let $\{v_1, \dots, v_{(2^n)}\} \in \mathbb{R}^n$ be the vertices of hyper-rectangle $\text{Im}(q, \mathcal{T})$, and $\check{f}, \hat{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be linear functions that bound $\mathcal{T}f_a^w(\mathcal{T}^{-1}z)$ for all $z \in \text{Im}(q, \mathcal{T})$. Define $H = \{\text{rect}(\check{f}(v), \hat{f}(v)) \mid v \in \{v_1, \dots, v_{(2^n)}\}\}$. Then, it holds that $\text{Post}(q, \mathcal{T}, a) \subseteq \text{conv}(H)$, where $\text{conv}(H)$ is the convex hull of hyper-rectangles in H .*

Proof. For $z \in \text{Im}(q, \mathcal{T})$, we have that $\mathcal{T}f_a^w(\mathcal{T}^{-1}z) \in \text{rect}(\check{f}(z), \hat{f}(z))$. Consequently, $\text{Post}(q, \mathcal{T}, a) \subseteq$




Figure 1. (a) For each point z in the black rectangle ($\text{Im}(q, \mathcal{T})$), $\mathcal{T}f_a^w(\mathcal{T}^{-1}z)$ is contained in the orange rectangle whose vertices are defined by $\check{f}(z)$ and $\hat{f}(z)$. (b) The blue rectangles contain $\check{f}(z)$ and $\hat{f}(z)$ for all $z \in \text{Im}(z, \mathcal{T})$. (c) The orange rectangles capture $\mathcal{T}f_a^w(\mathcal{T}^{-1}v)$ for all vertices v of $\text{Im}(q, \mathcal{T})$ and (d) fully define the red convex region that captures $\mathcal{T}f_a^w(\mathcal{T}^{-1}z)$ for all $z \in \text{Im}(q, \mathcal{T})$.

$\widetilde{\text{Post}}(q, \mathcal{T}, a) = \cup_{z \in \text{Im}(q, \mathcal{T})} \text{rect}(\check{f}(z), \hat{f}(z))$. Note that $\text{Im}(q, \mathcal{T})$ is a convex polytope and the construction of $\text{rect}(\check{f}(z), \hat{f}(z))$ only involves linear operations. As a consequence, $\widetilde{\text{Post}}(q, \mathcal{T}, a)$ is fully described by the vertices of $\text{Im}(q, \mathcal{T})$. Hence, $\widetilde{\text{Post}}(q, \mathcal{T}, a) \subseteq \text{conv}(H)$. \square

As a result of the above proposition, to construct the post image overapproximation induced by the local linear under- and overapproximations of the NN, we only have to check the vertices of the image as illustrated in Figure 1. Utilizing the analytical reformulation of the transition kernel as in (3) and the post-image overapproximation, we obtain that

$$\min_{x \in q} T^a(q'|x) \geq \min_{z \in \text{conv}(H)} g(z) \quad (5)$$

$$\max_{x \in q} T^a(q'|x) \leq \max_{z \in \text{conv}(H)} g(z) \quad (6)$$

where H is as defined in Proposition 4.1. Here, (6) is a log-concave maximization problem, which can be solved with convex optimization algorithms, such as gradient descent (Boyd et al., 2004). Although (5) is in general non-convex, the following result – a consequence of Corollary 32.3.4 in (Ralph Tyrell Rockafellar, 2015) – guarantees that to compute lower bound (5), i.e., the minimum of a log-concave problem, it suffices to check the vertices of $\text{conv}(H)$.

Lemma 4.2. *For $g(z)$ as defined in (3) it holds that $\min_{z \in \text{conv}(H)} g(z) = \min_{z \in V} g(z)$, where V is the set of vertices of $\text{conv}(H)$.*

4.2. Efficient Computation of Transition Probabilities

Note that, although solving for (5) and (6) reduces to the solution of convex maximization and minimization problems, to build the abstraction, we still need to solve $\mathcal{O}(|Q_s|^2)$ of these problems (one for each pair of states in Q_s). This becomes expensive for large $|Q_s|$, which is often the case for high-dimensional systems. In this section, we propose

an alternative approach to reduce this computational burden. In particular, the following theorem shows that if we overapproximate $\text{conv}(V)$ by an axis-aligned hyper-rectangle, to find solutions to (4), we only have to check a finite number of points at the boundary of the axis-aligned hyper-rectangle and perform $\mathcal{O}(|Q_s|^2)$ function evaluations, rather than optimizations.

Theorem 4.3. *For Process (1) under action a , regions $q, q' \subset \mathbb{R}^n$, and proper transformation matrix \mathcal{T} w.r.t. q, q' , construct H w.r.t. q per Proposition 4.1. Further, let vectors \hat{z}, \hat{z} define the vertices of $\text{rect}(\text{conv}(H))$, i.e., $\text{rect}(\text{conv}(H)) = \text{rect}(\hat{z}, \hat{z})$ such that for every $l \in \{0, \dots, n\}$ $\hat{z}^{(l)} \leq \hat{z}^{(l)}$, and denote by \bar{v} the center of $\text{Im}(q', \mathcal{T})$. Then, for $z_{\min}, z_{\max} \in \mathbb{R}^n$ defined such that for $l \in \{0, \dots, n\}$*

$$z_{\min}^{(l)} = \arg \max_{z^{(l)} \in \{\hat{z}^{(l)}, \hat{z}^{(l)}\}} |z^{(l)} - \bar{v}^{(l)}|,$$

$$z_{\max}^{(l)} = \begin{cases} \bar{v}^{(l)} & \text{if } \bar{v}^{(l)} \in [\hat{z}^{(l)}, \hat{z}^{(l)}] \\ \arg \min_{z^{(l)} \in \{\hat{z}^{(l)}, \hat{z}^{(l)}\}} |z^{(l)} - \bar{v}^{(l)}| & \text{otherwise} \end{cases}$$

it holds that

$$\max_{x \in q} T^a(q'|x) \geq g(z_{\min}), \quad \max_{x \in q} T^a(q'|x) \leq g(z_{\max}). \quad (7)$$

Proof. We consider the max case; the min case follows similarly. By construction $\text{conv}(H) \subseteq \text{rect}(\text{conv}(H))$, hence $\max_{x \in q} T^a(q'|x) \leq \max_{z \in \text{rect}(\text{conv}(H))} g(z)$. As $\text{rect}(\text{conv}(H))$ is an axis-aligned hyperrectangle, it holds that $\max_{z \in \text{rect}(\text{conv}(H))} g(z) = \prod_{l=1}^n \max_{z^{(l)} \in [\hat{z}^{(l)}, \hat{z}^{(l)}]} \int_{\hat{v}^{(l)}}^{\hat{v}^{(l)}} \mathcal{N}(\bar{x}^{(l)} \mid z^{(l)}, 1) d\bar{x}^{(l)}$, where $[\hat{v}^{(l)}, \hat{v}^{(l)}]$ is the interval of the l -th dimension of $\text{Im}(q', \mathcal{T})$. This is a product of n maximization problems that seek for the mean of a Gaussian distribution that maximizes its integral on a set. Each of these is maximized by minimizing the distance of $z^{(l)}$ to the center point $\bar{v}^{(l)}$ of the integration set. Hence, $z_{\max}^{(l)}$ is equal to $\bar{v}^{(l)}$ if $\bar{v}^{(l)} \in [\hat{z}^{(l)}, \hat{z}^{(l)}]$, else to one of the endpoints of $[\hat{z}^{(l)}, \hat{z}^{(l)}]$. \square

According to the above theorem, given $\text{conv}(H)$ there exist a finite number of potential z_{\max} and z_{\min} . Moreover, given a potential z_{\min} or z_{\max} we can immediately find all the regions that take optimal value for (7) at this point, based on the positions of the regions in the grid, as illustrated in Figure 2. As a consequence, we can simply check the finite sets of all possible z_{\min} and z_{\max} once to obtain z_{\max} and z_{\min} for all regions in the same group, and compute (7) by evaluating function g on z_{\min} and z_{\max} for each region. Although this dramatic reduction of computation comes at the cost of more conservative bounds compared




Figure 2. Regions q_1, q_2 share the same relative position w.r.t. $\text{rect}(\text{conv}(H))$, i.e., $\forall z \in q_1, q_2$ it holds that $z^{(l)} \leq \inf(\text{rect}(\text{conv}(H))^{[l]})$ for all $l \in \{0, \dots, n\}$. Consequently, they share the same z_{\min} and z_{\max} for problems (7). The white-blue coloring of the regions represents the grouping of regions accordingly. For overlapping region q_3 , the minimizing location for (5) is found at a vertex of $\text{conv}(H)$.

to solving (5) and (6), the introduced overapproximation can be reduced by a refinement algorithm as proposed in Section 5.1.

5. Control Synthesis & Abstraction Refinement

Given Process (1) and an LTLf property ϕ , our objective is to synthesize a strategy that maximizes the probability of satisfying ϕ . The IMDP abstraction \mathcal{I} as constructed above, captures the behavior of Process (1) w.r.t. the regions of interest. Therefore, we can focus on finding a strategy for \mathcal{I} that maximizes ϕ subject to being robust against all the uncertainties (errors) induced by the discretization of space and the NN-dynamics approximation process. This translates to assuming that the adversary's (uncertainty) objective is to minimize the probability of satisfaction. Hence,

$$\pi_{\mathcal{I}}^* = \arg \max_{\pi_{\mathcal{I}} \in \Pi_{\mathcal{I}}} \min_{\xi \in \Xi} \mathbb{P}[\omega_{\mathcal{I}} \models \phi \mid \pi_{\mathcal{I}}, \xi, \omega_{\mathcal{I}}(0) = q]. \quad (8)$$

We note that $\pi_{\mathcal{I}}^*$ can be computed using known algorithms with a computational complexity polynomial in the number of states in the IMDP (Lahijanian et al., 2015). To show that $\pi_{\mathcal{I}}^*$ maps to a robust strategy for Process (1), we need to introduce a mapping between the process and the IMDP. Let $M_{\mathbf{x}} : \mathbb{R}^n \rightarrow Q$ be a function that maps continuous states $x \in \mathbb{R}^n$ to their corresponding discrete regions in Q , i.e., $x \in q \implies M_{\mathbf{x}}(x) = q$. In addition, let $M_{\omega} : \Omega_{\mathbf{x}}^{\text{fin}} \rightarrow \Omega_{\mathcal{I}}^{\text{fin}}$ be a function that maps finite paths of Process (1) to the finite paths of IMDP \mathcal{I} , i.e., for a finite path $\omega_{\mathbf{x}}^N = x_0 \xrightarrow{a_0} x_1 \xrightarrow{a_1} \dots \xrightarrow{a_{N-1}} x_N$, $M_{\omega}(\omega_{\mathbf{x}}^N) = M_{\mathbf{x}}(x_0) \xrightarrow{a_0} M_{\mathbf{x}}(x_1) \xrightarrow{a_1} \dots \xrightarrow{a_{N-1}} M_{\mathbf{x}}(x_N)$. Then, we can map $\pi_{\mathcal{I}}^*$ to a switching strategy $\pi_{\mathbf{x}}$ through

$$\pi_{\mathbf{x}}^*(\omega_{\mathbf{x}}^N) = \pi_{\mathcal{I}}^*(M_{\omega}(\omega_{\mathbf{x}}^N)). \quad (9)$$

Further, we define the lower and upper bounds of the probability of satisfaction of ϕ under $\pi_{\mathcal{I}}^*$ as

$$\check{p}(q) = \min_{\xi \in \Xi} \mathbb{P}[\omega_{\mathcal{I}} \models \phi \mid \pi_{\mathcal{I}}^*, \xi, \omega_{\mathcal{I}}(0) = q], \quad (10)$$

$$\hat{p}(q) = \max_{\xi \in \Xi} \mathbb{P}[\omega_{\mathcal{I}} \models \phi \mid \pi_{\mathcal{I}}^*, \xi, \omega_{\mathcal{I}}(0) = q], \quad (11)$$

respectively. The following theorem shows that the satisfaction probability bounds also hold for Process (1) under $\pi_{\mathbf{x}}^*$.

Theorem 5.1. *Given Process (1), a compact set $X_{safe} \subset \mathbb{R}^n$, and an LTLf formula ϕ defined over the regions of interest in X_{safe} , let \mathcal{I} be the IMDP abstraction of Process (1) as described in Section 4. Further, let $\pi_{\mathcal{I}}^*$ be computed by (8) with probability bounds \check{p} and \hat{p} as in (10) and (11), respectively. Map $\pi_{\mathcal{I}}^*$ into a switching strategy $\pi_{\mathbf{x}}^*$ as in (9). Then for any initial state $x_0 \in X_{safe}$ it holds that*

$$P[\omega_{\mathbf{x}} \models \phi \mid \pi_{\mathbf{x}}^*, \omega_{\mathbf{x}}(0) = x_0] \in [\check{p}(M_{\mathbf{x}}(x_0)), \hat{p}(M_{\mathbf{x}}(x_0))].$$

The proof of this theorem follows similarly as the proof of Theorem 2 in (Jackson et al., 2021a). Theorem 5.1 guarantees that the probability that Process (1) satisfies ϕ is contained in the satisfaction probability bounds \check{p} and \hat{p} . The difference between \check{p} and \hat{p} can be viewed as the error induced by space discretization and local approximation of the NN dynamics with linear functions. This error monotonically decreases if the size of the discretization decreases. As a consequence, the synthesized strategy is optimal for an infinitely fine grid.

5.1. Synthesis driven refinement

Here, we present a discretization refinement scheme that aims to efficiently reduce the error induced by the space discretization. In each refinement step, we refine a predefined fixed number of states in Q_s , which we refer to as n_{ref} . To enable the use of Theorem 4.3 and Proposition 4.2, our refinement guarantees that all refined regions are axis-aligned hyper-rectangles in the transformed space. Hence, a region is refined by splitting the corresponding hyper-rectangle region $Im(q, \mathcal{T})$ over one dimension. To decide on which states to refine, we define a score function $\theta : Q \rightarrow \mathbb{R}^+$ as

$$\theta(q) = (\check{p}(q) - \hat{p}(q)) \sum_{a \in A} \sum_{q' \in Q} \left(\hat{P}(q', a, q) - \check{P}(q', a, q) \right)$$

where \check{p} and \hat{p} are the satisfaction probabilities as defined by (10) and (11). We refine the n_{ref} regions with the highest $\theta(q)$. The score function serves as a measure of uncertainty caused by state $q \in Q_s$ and closely resembles the uncertainty measure proposed in (Dutreix & Coogan, 2018) in a verification context.

The rationale behind our choice of which dimension to refine is based on the objective to reduce the conservatism




Figure 3. Region labeling and classification of each initial state $x \in X$ as Q^{yes} if $\check{p}(M_{\mathbf{x}}(x)) \geq 0.95$, Q^{no} if $\hat{p}(M_{\mathbf{x}}(x)) < 0.95$ and $Q^?$ otherwise, for the first, an intermediate and the final abstraction of Experiment 1. In blue, a simulated path labeled with the action chosen at each step.

introduced by the NN overapproximation process. In particular, for $q \in Q_s$, we want to find the dimension that minimizes the volume of $conv(H)$, as described in Proposition 4.1, for both regions created by splitting $Im(q, \mathcal{T})$ over this dimension. To do so, we transform all the edges of $Im(q, \mathcal{T})$ using the bounding functions and measure the expansion of the edges, i.e., the relative difference in distance between the vertices describing an edge before and after the transformation. As $Im(q, \mathcal{T})$ is an axis-aligned hyper-rectangle, we then take the dimension to refine equal to the dimension the largest expanded edge aligns to. The exact procedure to find the dimension to refine can be found in Appendix B

6. Case Studies

We consider 3 different NNDMs learned on non-linear datasets taken from the literature (for details see Appendix C). All transition probability bounds of the IMDP abstractions were computed using Theorem 4.3, except for the lower bounds on transitions to regions that overlap with the post image overapproximation of the region from which the transition starts. For those, we used Proposition 4.1, as illustrated in Figure 2 and explained in detail in Appendix A.2. Empirically, we found this approach to offer good results in balancing between precision and scalability. All experiments were run on an Intel Core i7-10610U CPU at 1.80GHz \times 2.30Hz with 16GB of RAM.

6.1. Efficient Control Synthesis by Iterative Refinement

We consider a 3-D car model from (Rajesh Rajamani, 2011), with state space representing position and orientation of the car and seven discrete actions switching between different feedback controllers that interact with the car and steer it to a given orientation. We are interested in synthesizing a strategy for a static overtaking scenario as shown in Figure 3. Here, the car should globally avoid an obstacle (“O”) and

eventually reach a desired (“D”) region, i.e., $\phi_1 = \mathcal{G}(\neg O) \wedge \mathcal{F}(D)$. To do so, we start with a very coarse abstraction and iteratively refine the discretization, which overall takes approximately 36 minutes. From Figure 3 we observe that the refinement procedure preserves the initial coarseness of the discretization for regions with small uncertainty on the satisfaction probability, whereas the critical regions, such as the corners around the obstacle, are further refined. Hereby, not only the lower bounds improve (the orange regions turn green), but also the upper bounds improve (red regions turn orange or green), and the controller strategies based on non-informative lower bounds are possibly updated (red regions turn green). Then, for the final abstraction, which is roughly one-fifth of the number of states (approximately 10,000) a standard uniform discretization of the domain using the finest grid of the final abstraction would contain, we obtain tight satisfaction probabilities.

6.2. Control Synthesis for Complex Specifications

To show that our framework can handle complex specifications, we use four nonlinear 2-D datasets generated by the nonlinear system considered in (Jackson et al., 2021b), and perform control synthesis given the same labeling of the domain and complex LTLf specification as in (Jackson et al., 2021b), i.e., $\phi_2 = \mathcal{G}(\neg O) \wedge \mathcal{F}(D1) \wedge \mathcal{F}(D2)$. The iterative abstraction and control synthesis procedure takes approximately 65 seconds, and the final abstraction consists of 1,500 states. Figure 4a shows that, although we assume a noisier dataset, we are able to compute informative satisfying probability bounds that resemble the result in (Jackson et al., 2021b).

6.3. Scalability: high-dimensional and complex NN-structures

Last, we test the scalability of our framework on a 5-D system with NNs of 5 hidden layers with 100 neurons per layer. Here, we consider the reach-avoid specification ϕ_1 with the labeling of the space as shown in Figure 4b. We again start with a coarse abstraction and iteratively improve the abstraction, which overall takes approximately 2 hours, from which 1.5 hours for generating the IMDP abstractions. The final abstraction consists of approximately 15,000 states. Figure 4b shows that we are able to guarantee for a large part of the domain that the initial states almost always (green regions) or almost never (red regions) satisfy the specification using complex controller strategies as indicated by the simulated paths and corresponding actions.

7. Conclusion

We introduced a formal control synthesis framework for a stochastic NNDMs with LTLf specifications. We showed that in practice the abstraction can be constructed very ef-

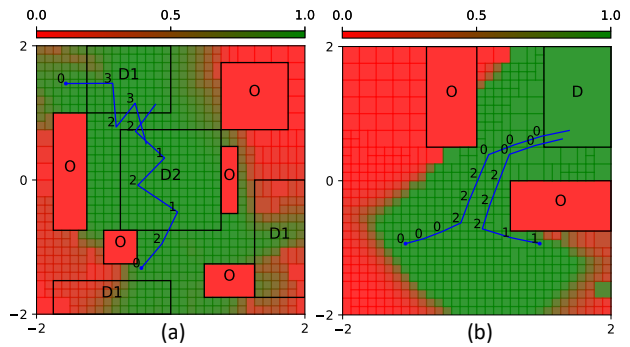


Figure 4. Region labeling and lower satisfaction probability bounds of the initial states for Experiment 2 (a) and 3 (b).

ficiently and developed an iterative refinement scheme to minimize the number of states of this discretization-based method. By experiments on various datasets, we showed that our framework enables efficient control synthesis of provably correct strategies for complex NNDM of several input dimensions on nontrivial control tasks.

References

- Baier, C. and Katoen, J. *Principles of model checking*. MIT press, 2008.
- Bertsekas, D. P. and Shreve, S. E. *Stochastic optimal control: the discrete-time case*. 1996.
- Boyd, S., Boyd, S. P., and Vandenberghe, L. *Convex optimization*. 2004.
- Cauchi, N., Laurenti, L., Lahijanian, M., Abate, A., Kwiatkowska, M., and Cardelli, L. Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems. *HSCC*, 2019. doi: 10.1145/3302504.3311805.
- De Giacomo, G. and Vardi, M. Y. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. *IJCAI*, 2013.
- Doyen, L., Frehse, G., Pappas, G. J., and Platzer, A. *Verification of Hybrid Systems*. 2018. ISBN 9783319105758. doi: 10.1007/978-3-319-10575-8.
- Dutreix, M. and Coogan, S. Efficient Verification for Stochastic Mixed Monotone Systems. In *ICCPs*, 2018.
- Fazlyab, M., Morari, M., and Pappas, G. J. An introduction to neural network analysis via semidefinite programming. In *CDC*, pp. 6341–6350. IEEE, 2021.
- Givan, R., Leach, S., and Dean, T. Bounded-parameter Markov decision processes. *Artificial Intelligence*, 2000.
- Jackson, J., Laurenti, L., Frew, E., and Lahijanian, M. Formal Verification of Unknown Dynamical Systems via Gaussian Process Regression. 2021a.

- Jackson, J., Laurenti, L., Frew, E., and Lahijanian, M. Strategy Synthesis for Partially-known Switched Stochastic Systems. In *HSCC*, 2021b.
- Lahijanian, M., Andersson, S. B., and Belta, C. Formal Verification and Synthesis for Discrete-Time Stochastic Systems. *TACON*, 2015. ISSN 00189286. doi: 10.1109/TAC.2015.2398883.
- Nagabandi, A. and et al. Neural Network Dynamics for Model-Based Deep Reinforcement Learning with Model-Free Fine-Tuning. In *ICRA*, 2018.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, 2019.
- Rajesh Rajamani. *Vehicle Dynamics and Control*. 2011.
- Ralph Tyrell Rockafellar. *Convex analysis*. 2015.
- Tabuada, P. *Verification and control of hybrid systems: A symbolic approach*. Springer US, 2009. ISBN 9781441902238. doi: 10.1007/978-1-4419-0224-5.
- Wei, T. and Liu, C. Safe Control with Neural Network Dynamic Models. 2021.
- Wells, A. M., Lahijanian, M., Kavraki, L. E., and Vardi, M. Y. LTLf synthesis on probabilistic systems. In *EPTCS*, 2020.
- Wicker, M., Laurenti, L., Patane, A., Paoletti, N., Abate, A., and Kwiatkowska, M. Certification of iterative predictions in bayesian neural networks. In *UAI*, pp. 1713–1723. PMLR, 2021.
- Xu, K., Shi, Z., Zhang, H., Wang, Y., Chang, K.-W., Huang, M., Kailkhura, B., Lin, X., and Hsieh, C.-J. Automatic Perturbation Analysis for Scalable Certified Robustness and Beyond. 2020.
- Xu, X., Tabuada, P., Grizzle, J. W., and Ames, A. D. Robustness of control barrier functions for safety critical control. *IFAC-PapersOnLine*, 48(27):54–61, 2015.
- Zhao, H., Zeng, X., Chen, T., Liu, Z., and Woodcock, J. Learning Safe Neural Network Controllers with Barrier Certificates. *FAOC*, 2020.

A. Efficient Computation of Transition Probabilities

A.1. Grouping Procedure

Recall that we discretized X_{safe} using a grid in the transformed space induced by transformation matrix \mathcal{T} . For $q \in Q_s$, let vectors \check{h} and \hat{h} define the vertices of the hyper-rectangular overapproximation of $\text{rect}(\text{conv}(H))$, i.e., $\text{rect}(\check{h}, \hat{h}) = \text{rect}(\text{conv}(H))$, such that $\check{h}^{(l)} \leq \hat{h}^{(l)}$ for $l \in \{0, \dots, n\}$, as in Proposition 4.1. Then, for the l -th dimension, we denote the set of discretization intervals by

$$\Delta^{[l]} = \{[\check{z}^{(l)}, \hat{z}^{(l)}] \mid \text{rect}(\check{z}, \hat{z}) = \text{Im}(q, \mathcal{T}), \check{z}^{(l)} \leq \hat{z}^{(l)}, q \in Q_s\} \quad (12)$$

Based on $\text{rect}(\check{h}, \hat{h})$, we split $\Delta^{[l]}$ in three subsets,

$$\begin{aligned} \Delta_L^{[l]} &= \{r \in \Delta^{[l]} \mid z < h, \forall z \in r, \forall h \in [\check{h}^{(l)}, \hat{h}^{(l)}]\} \\ \Delta_O^{[l]} &= \{r \in \Delta^{[l]} \mid r \cup [\check{h}^{(l)}, \hat{h}^{(l)}] \neq \emptyset\} \\ \Delta_U^{[l]} &= \{r \in \Delta^{[l]} \mid z > h, \forall z \in r, \forall h \in [\check{h}^{(l)}, \hat{h}^{(l)}]\}, \end{aligned} \quad (13)$$

and, we create a new set of intervals $\Delta_{\text{new}}^{[l]}$ consisting of the union of all intervals $\Delta_L^{[l]}$, the union of all intervals in $\Delta_U^{[l]}$, and the intervals of $\Delta_O^{[l]}$, i.e.,

$$\Delta_{\text{new}}^{[l]} = \left(\bigcup_{r \in \Delta_L^{[l]}} r \right) \cup \Delta_O^{[l]} \cup \left(\bigcup_{r \in \Delta_U^{[l]}} r \right)$$

By construction, it is guaranteed that each interval in $\Delta_{\text{new}}^{[l]}$ has a unique value for $z_{\text{min}}^{(l)}$ and $z_{\text{max}}^{(l)}$ as defined in Theorem 4.3. Next, we want to define a function $G : Q_s \rightarrow Q_s$ that given a region $q \in Q_s$ returns a grouping of the regions in Q_s such that all $q' \in g \in G$ have equal z_{min} and z_{max} as defined in Theorem 4.3. First, we introduce a labeling function $L_\Delta : \mathbb{R} \times \mathbb{R} \times \mathbb{N} \rightarrow 2$, such that

$$L_\Delta(r, \delta, l) = \begin{cases} \top & \text{if } r \cup \delta = r \wedge \delta \cup [\check{h}^{(l)}, \hat{h}^{(l)}] \neq \emptyset \\ \top & \text{if } r \subseteq \delta \\ \perp & \text{else.} \end{cases}$$

Then, grouping $G : Q_s \rightarrow Q_s$ is defined as

$$G(q) = \{Q_r \subseteq Q_s \mid \forall q' \in Q_r, \delta \in \Delta_{\text{new}}^{[l]}, \forall l \in \{0, \dots, n\}, L_\Delta(q', \delta, l)\}$$

where vectors \check{h} and \hat{h} define the vertices of the hyper-rectangular overapproximation of $\text{rect}(\text{conv}(H))$, i.e., $\text{rect}(\check{h}, \hat{h}) = \text{rect}(\text{conv}(H))$, for region q as in Proposition 4.1. The labeling function returns true (\top) for two intervals (r and δ) that have the same relative position w.r.t. $\text{conv}(H)$ for dimension l , where the possible positions for each dimension l are defined by $\Delta_{\text{new}}^{[l]}$. Note that even for a nonuniform discretization, the labeling function defines a complete and unique mapping function, i.e., for all $q \in Q_s$ it holds that $\bigcup_{g \in G(q)} g = Q_s$ and $\forall q' \in Q_s, \exists! g \in G(q)$ s.t. $q' \in g$.

A.2. Algorithm

Our approach to compute \check{P} and \hat{P} as described in Section 4.2 is summarized in Algorithm 1.

B. Refinement Procedure

Here, we denote by $\text{vert}(\text{rect}(\check{v}, \hat{v}))$ the set of vertices of hyper-rectangle $\text{rect}(\check{v}, \hat{v})$, i.e.,

$$\begin{aligned} \text{vert}(\text{rect}(\check{v}, \hat{v})) &= \{([v^{(1)}, v^{(2)}, \dots, v^{(n)}]^T \mid \\ &v^{(l)} \in \{\check{v}, \hat{v}\}, l \in \{0, \dots, n\}\} \end{aligned} \quad (14)$$

Algorithm 1: Computation of \tilde{P} and \hat{P} .

```

1 for  $q \in Q_s, a \in A$  do
2   Define  $H$  as in Proposition 4.1 and construct  $rect(conv(H))$ .
3   Find grouping  $G(q)$  as described in Appendix A.1.
4   for  $Q_r \in G(q)$  do
5     Pick an element  $q_{ref}$  from  $Q_r$ .
6     Find  $z_{max}$  that upper-bounds  $T^a(q_{ref}|x)$  for  $x \in q$  as in Theorem 4.3.
7     if  $q_{ref} \cup rect(conv(H)) = \emptyset$  then
8       Find  $\tilde{v}, \hat{v}$  s.t.  $rect(\tilde{v}, \hat{v}) = rect(q_{ref})$ 
9       Find  $z_{min}$  that lower-bounds  $T^a(q_{ref}|x)$  for  $x \in q$  as in Theorem 4.3.
10      for  $q' \in Q_r$  do
11         $\tilde{P}(q, a, q') = g(z_{min})$ 
12         $\hat{P}(q, a, q') = g(z_{max})$ 
13      else
14        for  $q' \in Q_r$  do
15          Find  $\tilde{v}, \hat{v}$  s.t.  $rect(\tilde{v}, \hat{v}) = rect(q_{ref})$ 
16           $z_{min} = \arg \min_{z \in V} g(z)$ 
17           $\tilde{P}(q, a, q') = g(z_{min})$ 
18           $\hat{P}(q, a, q') = g(z_{max})$ 

```

	ϕ	n	States	X_{safe}	Cov_v	$ A $	#L	#N/L	Proj.
1	ϕ_1	2	$[x_1, x_2]$	$[-2, 2] \times [-2, 2]$	$0.2I$	4	3	20	$[x_1, x_2]$
2	ϕ_2	3	$[x, y, \phi]$	$[0 \times 10] \times [0 \times 2] \times [-0.5 \times 0.5]$	$diag([0.1, 0.1, 0.01])$	7	4	50	$[x, y]$
3	ϕ_2	5	$[x_1, x_2, x_3, x_4, x_5]$	$[-2, 2] \times [-2, 2] \times [-0.4, 0.4] \times [-0.4, 0.4] \times [-0.4, 0.4]$	$0.05I$	3	5	100	$[x_1, x_2]$

Table 1. Overview of the LTLf specification, number of dimensions of the system, state space representation, compact set, noise characteristics, number of modes, number of layers and neurons per layer of the NNs considered per experiment. Last, the states spanning the projected space as shown in the figures 3 and 4 (Proj.), are specified.

Let the vertices of hyper-rectangle $Im(q, \mathcal{T})$ be defined by vectors \tilde{v} and \hat{v} (i.e., $rect(\tilde{v}, \hat{v}) = rect(q)$), and \check{f}_a, \hat{f}_a be the linear functions that bound $\mathcal{T}f_a^w(\mathcal{T}^{-1})$ for $z \in Im(q, \mathcal{T})$ as defined by Proposition 4.1, and $\check{A}_a, \hat{A}_a \in \mathbb{R}^{n \times n}$ be the matrices that capture the linear transformation of \check{f}_a and \hat{f}_a , respectively.

To measure the relative deformation of the edges of $Im(q, \mathcal{T})$, we first denote the set of vertices of $Im(q, \mathcal{T})$ by $vert(q)$ as in (14). Further, we define a mapping $M_\xi : \mathbb{R}^n \rightarrow \{0, \dots, n\}$ that returns the matching dimension(s) for two vectors $v, v' \in \mathbb{R}^n$, i.e.,

$$M_\xi(v, v') = l \iff v^{(l)} = v'^{(l)}$$

Note that, because $Im(q, \mathcal{T})$ is an axis-aligned hyper-rectangle, for all $v, v' \in vert(q)$ and $v \neq v'$, $M_v(v, v')$ is an unique and complete mapping. The maximum expansion of an edge that is described by vertices $v, v' \in vert(q)$ under transformation of \check{A}_a and \hat{A}_a for all $a \in A$, which we refer to as $\xi_{max}(v, v') \in \mathbb{R}$, is then defined as

$$\xi_{max}(v, v') = \max_{a \in A} \left(\frac{\|\check{A}_a \tilde{v} - \check{A}_a \hat{v}\|}{\|\tilde{v} - \hat{v}\|}, \frac{\|\hat{A}_a \tilde{v} - \hat{A}_a \hat{v}\|}{\|\tilde{v} - \hat{v}\|} \right)$$

Then, for $q \in Q_s$ the dimension to refine, which we call $l_c \in \{0, \dots, n\}$, is defined as the dimension corresponding to vectors $v, v' \in vert(q)$ that the maximize $\xi_{max}(v, v')$, i.e.,

$$l_c = M_\xi \left(\arg \max_{v, v' \in vert(q)} \xi_{max}(v, v') \right)$$

C. Case Studies

The details on the system and NN-characteristics considered for each system are reported in Table 1.