

NNSynth: Neural Network Guided Abstraction-Based Controller Synthesis for Stochastic Systems

Xiaowu Sun and Yasser Shoukry

Abstract—In this paper, we introduce NNSynth, a new framework that uses machine learning techniques to guide the design of abstraction-based controllers with correctness guarantees. NNSynth utilizes neural networks (NNs) to guide the search over the space of controllers. The trained neural networks are “projected” and used for constructing a “local” abstraction of the system. An abstraction-based controller is then synthesized from such “local” abstractions. If a controller that satisfies the specifications is not found, then the best found controller is “lifted” to a neural network for further training. Our experiments show that this neural network-guided synthesis leads to more than 50× or even 100× speedup in high dimensional systems compared to the state-of-the-art.

I. INTRODUCTION

Abstraction-based control synthesis techniques have gained considerable attention in the past decade. These techniques provide tools for automated, correct-by-construction controller synthesis from complex specifications, typically given in the form of a Linear Temporal Logic (LTL) formulae [1]. It is then unsurprising the vast amount of developed software tools that can handle a wide variety of nonlinear control systems including Pessoa [2], CoSyMa [3], SCOTS [4], QUEST [5], FAUST [6], StocHy [7], and AMYTISS [8]. At the heart of all these tools is the need to obtain discrete abstraction of continuous-time dynamical systems using various quantization methods for state and input spaces. The resulting discrete abstraction is then traversed to search for a feedback controller that conforms to the required LTL specification. A significant drawback of discrete abstraction is the vast number of combinations of quantized states and inputs that need to be considered. The problem is exacerbated in high-dimensional state and input spaces, leading to the so-called *curse of dimensionality*.

Motivated by the recent success of machine learning techniques in efficiently searching over the space of feedback controllers (e.g., imitation learning and reinforcement learning), we ask the following question: *Can machine learning techniques be used to accelerate the process of synthesizing abstraction-based controllers from LTL specifications?* On the one hand, machine learning techniques enjoy favorable scalability properties and eliminate the dependency on state-space quantization. On the other hand, these learning-based feedback controllers (or policies) do not come with the guarantee that they conform to the LTL specifications.

This work was partially sponsored by the NSF awards #CNS-2002405, #CNS-2013824, and #EECS-2139781 and C3.AI Digital Transformation Institute.

Xiaowu Sun and Yasser Shoukry are with the Department of Electrical Engineering and Computer Science, University of California, Irvine {xiaowus,yshoukry}@uci.edu

This motivates the need to closely integrate the scalability of learning-based techniques with the provable guarantees provided by the abstraction-based techniques.

Toward this end, we propose NNSynth, a new framework for synthesizing abstraction-based controllers from LTL specifications. Unique to NNSynth is the use of machine learning techniques to train a neural network (NN) based controller, which will guide the synthesis of the final abstraction-based controller. The advantages of the proposed NN-guided abstraction-based controller synthesis are multi-fold. First, it utilizes the empirically proven advantages of machine learning algorithms to search the space of feedback controllers without relying on expensive quantizations of state and input spaces. Second, it limits the search over the quantized spaces only to the neighborhood of the controllers proposed by the NN training. That is, our approach uses NN training to guide the search over the quantized abstract system and eliminates the need to consider all combinations of quantized states and inputs. Third, the use of neural networks to guide the design of the abstraction-based controller opens the door to encoding human preference on how a dynamical system should act. Such human preference is crucial in several real-world settings where a human user or an operator interacts with an autonomous dynamical system. Current research found that human preference can be efficiently captured using expert demonstrations and preference-based learning, but is hard to be accurately captured in the form of logical formulae or reward functions [9]. We demonstrate these advantages using several key applications and show that NNSynth scales more favorably compared to the state-of-the-art techniques by achieving more than 50× or even 100× speedup in high dimensional systems.

Related Work. The closest results to our work are those reported in [10], [11], which propose a neurosymbolic framework to train control policies represented by short programs in a symbolic language. Similar to our approach, the work in [10], [11] trains a NN controller, projects it to the space of symbolic controllers, analyzes the symbolic controller, and lifts it back to the space of NN policies for further training. Differently, our approach focuses on designing a finite-state, abstraction-based controller instead of short programs in a symbolic language. This difference (short programs versus finite-state controllers) manifests itself in all the framework steps, particularly the NN training, projection, and lifting. We confine our focus to synthesizing finite-state controllers due to the extensive literature on analyzing such controllers in tandem with the controlled physical systems [1]. Another line of related work is reported in [12], [13] which study the

problem of extracting a finite-state controller from a recurrent neural network controller. We note that in our framework, NN policies are not the final produced controllers, but are used to guide the search for abstraction-based controllers.

II. PROBLEM FORMULATION

A. Notation

We denote the set of real numbers, positive real numbers, and natural numbers by \mathbb{R} , \mathbb{R}^+ , \mathbb{N} , respectively. Let $|X|$ be the cardinality of a set X and $\text{Int}(X)$ be the interior of a set X . Let $\|x\|$ be the Euclidean norm of a vector $x \in \mathbb{R}^n$ and x^\top be the transpose of $x \in \mathbb{R}^n$. Let the inner product of two functions $h_1 : X \rightarrow \mathbb{R}^m$ and $h_2 : X \rightarrow \mathbb{R}^m$ be defined as $\langle h_1, h_2 \rangle = \int_X h_1(x)^\top h_2(x) dx$, which induces a norm $\|h_1\| = \sqrt{\langle h_1, h_1 \rangle}$. We use ∇J to denote the Fréchet gradient of a functional J , and use the big O notation for upper bounds.

B. Dynamical Model

We consider discrete-time nonlinear dynamical systems of the form:

$$x^{(t+1)} = f(x^{(t)}, u^{(t)}) + g(x^{(t)}, u^{(t)}), \quad (1)$$

where $x^{(t)} \in X \subset \mathbb{R}^n$ is the state and $u^{(t)} \in U \subset \mathbb{R}^m$ is the control input at time step $t \in \mathbb{N}$. The dynamical model consists of the priori known nominal model f and the unknown model-error g capturing unmodeled dynamics. Both functions f and g are assumed to be locally Lipschitz continuous. As a well-studied technique to learn unknown functions from data, we assume the model-error g can be learned using Gaussian Process (GP) regression [14]. We use $\mathcal{GP}(\mu_g, \sigma_g^2)$ to denote a GP regression model with the posterior mean and variance functions be μ_g and σ_g^2 , respectively. Given the dynamical system (1) with the model-error g learned by $\mathcal{GP}(\mu_g, \sigma_g^2)$, let $\tau : X \times X \times U \rightarrow [0, 1]$ be the corresponding conditional stochastic kernel. Specifically, given the current state $x \in X$ and input $u \in U$, the distribution $\tau(\cdot|x, u)$ is given by the Gaussian distribution $\mathcal{N}(f(x, u) + \mu_g(x, u), \sigma_g^2(x, u))$.

We treat the nonlinear system (1) with the model-error g learned by $\mathcal{GP}(\mu_g, \sigma_g^2)$ as a continuous Markov Decision Process (MDP) denoted by a tuple $\Sigma \triangleq (X, U, \tau)$. We denote by $T(A|x, u)$ the transition probability of reaching a subset $A \subset X$ in one step from state $x \in X$ with input $u \in U$:

$$T(A|x, u) = \int_A \tau(x'|x, u) dx'. \quad (2)$$

This integral can be easily computed since $\tau(\cdot|x, u)$ is a Gaussian distribution.

C. Abstraction-based Controller and Specification

We consider to control the continuous MDP Σ (i.e., the nonlinear system (1) with the model-error learned by GP) using abstraction-based controllers. An abstraction-based controller considers to partition the continuous state space $X \subset \mathbb{R}^n$ into a finite set of abstract states $\hat{X} = \{q_1, \dots, q_N\}$, where each abstract state $q_i \in \hat{X}$ is an infinity-norm ball in

\mathbb{R}^n . The partitioning satisfies $X = \bigcup_{q \in \hat{X}} q$ and $\text{Int}(q_i) \cap \text{Int}(q_j) = \emptyset$ if $i \neq j$. We denote by $\lambda \in \mathbb{R}^+$ the pre-specified grid size used for partitioning the state space. Let $\text{abs} : X \rightarrow \hat{X}$ map a state $x \in X$ to the abstract state $q = \text{abs}(x)$ that contains x , i.e., $x \in q$, and $\text{ct} : \hat{X} \rightarrow X$ map an abstract state $q \in \hat{X}$ to its center $\text{ct}(q) \in X$, which is well-defined since abstract states are infinity-norm balls. With some abuse of notation, we denote by q both an abstract state, i.e., $q \in \hat{X}$, and a subset of states, i.e., $q \subset X$.

Given a partitioning of the state space, an abstraction-based controller $\Psi : X \rightarrow U$ assigns the same control input to all the states in the same abstract state, i.e., $\Psi(x_1) = \Psi(x_2)$ if $\text{abs}(x_1) = \text{abs}(x_2)$. We denote by \mathcal{S} the set of all abstraction-based controllers, where the underlying partitioning of the state space can be different for different abstraction-based controllers in \mathcal{S} .

For the high-level specifications, though our framework can be easily extended to general Linear Temporal Logic (LTL) specifications in a bounded time horizon, we describe our algorithms with safety and liveness specifications for simplicity. Let $\xi_{x_0, \Psi} : \{1, \dots, H\} \rightarrow X$ denote a closed-loop trajectory of the system (1) that starts from the state $x_0 \in X$ and evolves under the control law Ψ in a bounded time horizon H . Then, we use $\xi_{x_0, \Psi} \models \phi_{\text{safety}}$ and $\xi_{x_0, \Psi} \models \phi_{\text{liveness}}$ to denote a trajectory $\xi_{x_0, \Psi}$ satisfying the safety and liveness specifications, respectively, i.e.,

$$\begin{aligned} \xi_{x_0, \Psi} \models \phi_{\text{safety}} &\iff \forall t \in \{1, \dots, H\}, \xi_{x_0, \Psi}(t) \notin X_{\text{obst}}, \\ \xi_{x_0, \Psi} \models \phi_{\text{liveness}} &\iff \exists t \in \{1, \dots, H\}, \xi_{x_0, \Psi}(t) \in X_{\text{goal}}, \end{aligned}$$

where $X_{\text{goal}} \subset X$ and $X_{\text{obst}} \subset X$ represent the goal and the obstacles, respectively. Given a specification $\phi = \phi_{\text{safety}} \wedge \phi_{\text{liveness}}$, we denote by $\text{Pr}(\Sigma_\Psi \models \phi)$ the average probability that the continuous MDP Σ controlled by Ψ satisfies the specification ϕ (averaged over initial states).

D. Main Problem

The goal of this paper is to synthesize an abstraction-based controller $\Psi \in \mathcal{S}$ for the continuous system (1) to satisfy the given specifications ϕ while minimizing some given cost. The cost functional of a controller Ψ is defined as $J(\Psi) = \int_X c(x, \Psi(x)) d\mu_\Psi(x)$, where $c(x, u)$ is the state-action cost and μ_Ψ is the distribution of states induced by the controller Ψ . Now, we can define the problem of interest as follows:

Problem 2.1: Given a cost functional J , a high-level specification ϕ and a user defined threshold p , we are interested in synthesizing an abstraction-based controller $\Psi : X \rightarrow U$ for the continuous MDP Σ to minimize the cost $J(\Psi)$ while satisfying the specification ϕ with probability at least p :

$$\Psi^* = \operatorname{argmin}_{\Psi \in \mathcal{S}} J(\Psi) \quad \text{s.t. } \text{Pr}(\Sigma_\Psi \models \phi) \geq p. \quad (3)$$

III. NNSYNTH FRAMEWORK

Our framework is featured by the use of neural networks to guide the search for abstraction-based controllers satisfying the specification ϕ , and the ability to utilize policy gradient approaches for abstraction-based controllers to minimize the cost functional J . In this section, we first give an overview




Fig. 1. A cartoon figure summarizing the NNSynth framework. NNSynth starts by training a neural network controller \mathcal{NN} using the data provided by the expert \mathcal{D} . The obtained neural network is then projected to an abstraction-based controller by evaluating the neural network on the representative points of abstract states, i.e. using the control actions $\mathcal{NN}(\text{ct}(q))$. The obtained finite-state abstraction is then augmented with control actions in the neighborhood of the actions proposed by the neural network $\mathcal{NN}(\text{ct}(x)) \pm i\delta$. A controller is then synthesized from the augmented model. In case a controller was not found, the “best” controller so far is then lifted to a neural network controller which is further trained using the expert data \mathcal{D} to obtain a new \mathcal{NN} . The loop continues until an abstraction-based controller is found.

Algorithm 1 NNSYNTH ($\mathcal{D}_{\text{exp}}, \phi, H, p, \varsigma, \eta$)

```

1: Initialize  $\mathcal{NN}_{\text{init}}$  with random weights
2:  $\mathcal{NN}_{\text{init}} = \text{UPDATE}(\mathcal{NN}_{\text{init}}, \mathcal{D}_{\text{exp}}, \eta)$ 
3:  $\Psi_0, V_{\text{avg}} = \text{PROJECT-BY-SYNTH}(\mathcal{NN}_{\text{init}}, \phi, H)$ 
4: for  $k = 0, \dots, K-1$  do
5:   if  $V_{\text{avg}} \geq p + \varsigma$  then
6:     Return  $\Psi_k, V_{\text{avg}}$ 
7:    $\mathcal{NN}_k = \text{LIFT}(\Psi_k)$ 
8:    $\mathcal{NN}_{k+1} = \text{UPDATE}(\mathcal{NN}_k, \mathcal{D}_{\text{exp}}, \eta)$ 
9:    $\Psi_{k+1}, V_{\text{avg}} = \text{PROJECT-BY-SYNTH}(\mathcal{NN}_{k+1}, \phi, H)$ 
10:  Return  $\Psi_K, V_{\text{avg}}$ 

```

of our framework and then present each step separately in the following subsections.

The overview of the proposed NNSynth is depicted in Figure 1. Algorithm 1 outlines the framework. After initializing an abstraction-based controller Ψ_0 (line 1-3 in Algorithm 1), NNSynth lifts the abstraction-based controller Ψ_k to a neural network \mathcal{NN}_k through imitation learning of the data generated by Ψ_k (line 7 in Algorithm 1), updates the neural network through either imitation learning of the expert dataset \mathcal{D}_{exp} or reinforcement learning (by providing the state-action cost function c instead of expert data) with learning rate η (line 8 in Algorithm 1), and finally synthesizes a new abstraction-based controller Ψ_{k+1} under the guidance of \mathcal{NN}_{k+1} (line 9 in Algorithm 1). This loop iterates until the satisfaction probability V_{avg} is no less than the pre-specified threshold $p + \varsigma$ (line 5 in Algorithm 1).

A. Step 1: NN Training

Starting from the expert-provided trajectories $\mathcal{D}_{\text{exp}} = \{\xi_1, \xi_2, \dots\}$, we use imitation learning to train a neural network controller \mathcal{NN} for the continuous MDP Σ . Alternatively, the NN controller can be trained by reinforcement learning, which requires the expert to provide the state-action cost $c : X \times U \rightarrow \mathbb{R}$ instead of the dataset \mathcal{D}_{exp} . Neural networks are highly parameterized and can be updated using gradient-based approaches $\mathcal{NN}_{k+1} = \mathcal{NN}_k - \eta \nabla J(\mathcal{NN}_k)$, where $\eta \in \mathbb{R}^+$ is the learning rate. The gradient $\nabla J(\mathcal{NN}_\theta)$ of a neural network parameterized by weights θ can be approximated using sampled trajectories:

$$\nabla J(\mathcal{NN}_\theta) \approx \frac{1}{M} \sum_{i=1}^M \sum_{t=1}^H \nabla_\theta \mathcal{NN}_\theta(u_{i,t} | x_{i,t}) \hat{Q}_i^t \quad (4)$$

where M is the number of trajectories, H is the bounded time horizon, and \hat{Q}_i^t is the estimated cost-to-go. We use the neural network to improve the controller’s performance (i.e., minimizing the cost functional J) although the gradient of an abstraction-based controller, denoted by $\nabla J(\Psi)$, does not exist. Detailed optimality analysis is given in Section IV.

B. Step 2: NN Projection

Regardless of the use of imitation learning or reinforcement learning, the resulting neural network \mathcal{NN} is not guaranteed to satisfy the specification ϕ and hence can not be used directly as a controller. Nevertheless, the neural network contains relevant control actions that can be used to obtain the final controller. To that end, NNSynth constructs a finite-state abstraction guided by \mathcal{NN} . Given a partitioning of the state space, we denote by $\widehat{X} = \{q_1, \dots, q_N\}$ the corresponding set of abstract states, where the partitioning grid size $\lambda \in \mathbb{R}^+$ is determined based on the theoretical guarantees to be achieved (see Section IV). Then, the finite-state abstraction induced by \mathcal{NN} is given as a tuple $\widehat{\Sigma}^{\mathcal{NN}} \triangleq (\widehat{X}, \widehat{U}^{\mathcal{NN}}, \widehat{T}^{\mathcal{NN}})$ with $\widehat{X} = \{q_1, \dots, q_N\}$, $\widehat{U}^{\mathcal{NN}} = \{\mathcal{NN}(\text{ct}(q)) \mid q \in \widehat{X}\}$, and

$$\widehat{T}^{\mathcal{NN}}(q' | q, u) = \begin{cases} T(q' | \text{ct}(q), u) & \text{if } u = \mathcal{NN}(\text{ct}(q)) \\ 0 & \text{otherwise,} \end{cases}$$

where the transition probabilities $T(q' | \text{ct}(q), u)$ can be computed as (2). In other words, the finite-state abstraction $\widehat{\Sigma}^{\mathcal{NN}}$ considers only one control action $\mathcal{NN}(\text{ct}(q))$ at each abstract state q and discards all other possible control actions. Computing such abstraction $\widehat{\Sigma}^{\mathcal{NN}}$ is straightforward and entails evaluating the NN controller at the center of each abstract state and computing the transition probabilities associated with these actions.

C. Step 3: System Augmentation

As shown in Figure 1, the finite-state abstraction $\widehat{\Sigma}^{\mathcal{NN}}$ may contain transitions that violate the given specification ϕ . This stems from the fact that $\widehat{\Sigma}^{\mathcal{NN}}$ considers only the actions taken by the trained network \mathcal{NN} . Therefore, the next step is to “augment” $\widehat{\Sigma}^{\mathcal{NN}}$ with additional transitions corresponding to control actions that are close to those given by \mathcal{NN} . This augmentation will provide the controller synthesis algorithm

with more freedom to choose other control actions. Given a precision $\delta \in \mathbb{R}^+$ and a range parameter $I \in \mathbb{N}$ (δ and I are determined based on theoretical guarantees in Section IV), we construct the augmented finite-state abstraction $\widehat{\Sigma}^{\mathcal{NN}+\delta} \triangleq (\widehat{X}, \widehat{U}^{\mathcal{NN}+\delta}, \widehat{T}_{\mathcal{NN}+\delta})$ where:

$$\widehat{X} = \{q_1, \dots, q_n\}, \quad (5)$$

$$\widehat{U}^{\mathcal{NN}+\delta} = \{\mathcal{NN}(\text{ct}(q)) \pm i\delta \mid q \in \widehat{X}, i = 0, 1, \dots, I\}, \quad (6)$$

$$\widehat{T}^{\mathcal{NN}+\delta}(q'|q, u) = \begin{cases} T(q'|ct(q), u) & \text{if } u \in \{\mathcal{NN}(\text{ct}(q)) \pm i\delta \mid i = 0, 1, \dots, I\} \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where with some abuse of notation, we use $\mathcal{NN}(\text{ct}(q)) \pm i\delta$ to denote $\mathcal{NN}(\text{ct}(q)) + [\pm i_1\delta, \pm i_2\delta, \dots, \pm i_m\delta]^\top$ with $i_1, i_2, \dots, i_m \in \{0, 1, \dots, I\}$. In other words, the augmented abstraction $\widehat{\Sigma}^{\mathcal{NN}+\delta}$ takes into account all the control actions that are $\delta, 2\delta, \dots, I\delta$ away from those given by the neural network \mathcal{NN} , where the distance is considered for each dimension of the control input $u \in \mathbb{R}^m$.

D. Step 4: Controller Synthesis

The next step is to synthesize a controller $\widehat{\Psi}$ for the augmented abstraction $\widehat{\Sigma}^{\mathcal{NN}+\delta}$ to satisfy the specification ϕ . We emphasize that though the controller $\widehat{\Psi} : \widehat{X} \rightarrow U$ is synthesized for the finite-state abstraction $\widehat{\Sigma}^{\mathcal{NN}+\delta}$, it yields an abstraction-based controller $\Psi : X \rightarrow U$ for the continuous system Σ by letting $\Psi(x) = \widehat{\Psi}(\text{abs}(x))$. In other words, the controller Ψ controls the continuous system Σ by applying the same control action $\widehat{\Psi}(q)$ at all states $x \in q$, where $q \in \widehat{X}$. The difference in the probabilities of satisfying the specification ϕ for the finite-state abstraction $\widehat{\Sigma}^{\mathcal{NN}+\delta}$ controlled by $\widehat{\Psi}$ and the continuous MDP Σ controlled by Ψ can be bounded [15] (see Section IV).

With the notations introduced above, let $\widehat{\Sigma}_{\widehat{\Psi}}^{\mathcal{NN}+\delta}$ be the finite-state abstraction $\widehat{\Sigma}^{\mathcal{NN}+\delta}$ controlled by $\widehat{\Psi}$. Given the bounded time horizon H , we define the value function $V : \widehat{X} \times \{0, \dots, H\} \rightarrow [0, 1]$ by letting $V(q, t)$ be the probability of satisfying the given specification ϕ in $H - t$ time steps when the system $\widehat{\Sigma}_{\widehat{\Psi}}^{\mathcal{NN}+\delta}$ starts from $q \in \widehat{X}$. Then, the average probability of satisfying the specification ϕ is given by:

$$V_{\text{avg}} \triangleq \Pr\left(\widehat{\Sigma}_{\widehat{\Psi}}^{\mathcal{NN}+\delta} \models \phi\right) = \frac{1}{|\widehat{X}|} \sum_{q \in \widehat{X}} V(q, 0). \quad (8)$$

Algorithm 2 presents details on the abstraction-based controller synthesis, which summarizes Subsections III-B, III-C, and III-D. To maximize the probability of satisfying ϕ -liveness (similarly for ϕ_{safety}), we solve the following dynamic programming (DP) recursion:

$$Q_t(q, u) = \sum_{q' \in \widehat{X}} V_{t+1}^*(q') \widehat{T}^{\mathcal{NN}+\delta}(q'|q, u) \quad (9)$$

$$V_t^*(q) = \max_{u \in \{\mathcal{NN}(\text{ct}(q)) \pm i\delta \mid i = 0, \dots, I\}} Q_t(q, u) \quad (10)$$

with the initial condition $V_H^*(q) = 1$ if $q \subseteq X_{\text{goal}}$ and 0 otherwise, where $t = H - 1, \dots, 0$, and the transition probability matrix $\widehat{T}^{\mathcal{NN}+\delta}$ is given by (7). Critical to the

Algorithm 2 PROJECT-BY-SYNTH (\mathcal{NN} , ϕ , H)

```

1: if  $\phi == \text{Safety}$  then
2:    $V(q, H) = 1$  for all  $q \in \widehat{X}$ 
3: else
4:    $V(q, H) = 0$  for all  $q \in \widehat{X}$ 
5:    $U_{\text{buffer}}(q) = \text{set}()$  for all  $q \in \widehat{X}$ 
6:   for  $t = H - 1, \dots, 0$  do
7:     for  $q \in \widehat{X} \setminus (\widehat{X}_{\text{goal}} \cup \widehat{X}_{\text{obst}})$  do
8:       for  $u \in \{\mathcal{NN}(\text{ct}(q), t) \pm i\delta \mid i = 0, \dots, I\} \setminus U_{\text{buffer}}(q)$  do
9:          $U_{\text{buffer}}(q).\text{add}(u)$ 
10:        Compute  $\widehat{T}(q'|q, u)$  for all  $q' \in \widehat{B}_\rho(f(\text{ct}(q), u))$ 
11:        if  $\phi \neq \text{Safety}$  then
12:          Compute transition prob. to the goal  $\widehat{T}(X_{\text{goal}}|q, u)$ 
13:           $\widehat{T}(q'|q, u) = 0$  for all  $q' \in \widehat{X}_{\text{goal}}$ 
14:        for  $q \in \widehat{X} \setminus (\widehat{X}_{\text{goal}} \cup \widehat{X}_{\text{obst}})$  do
15:           $V_{\text{max}} = 0$ 
16:          for  $u \in \{\mathcal{NN}(\text{ct}(q), t) \pm i\delta \mid i = 0, \dots, I\}$  do
17:             $Q(q, u) = \sum_{q' \in \widehat{B}_\rho(f(\text{ct}(q), u))} \widehat{T}(q'|q, u) V(q', t+1)$ 
18:            if  $\phi \neq \text{Safety}$  then
19:               $Q(q, u) = Q(q, u) + \widehat{T}(X_{\text{goal}}|q, u)$ 
20:            if  $Q(q, u) > V_{\text{max}}$  then
21:               $V_{\text{max}} = Q(q, u)$ 
22:               $\widehat{\Psi}(q, t) = u$ 
23:             $V(q, t) = V_{\text{max}}$ 
24:           $V_{\text{avg}} = \frac{1}{|\widehat{X}|} \sum_{q \in \widehat{X}} V(q, 0)$ 
25:         $\Psi(x, t) = \widehat{\Psi}(\text{abs}(x), t)$ 
26:      Return  $\Psi, V_{\text{avg}}$ 

```

speedup of NNSynth is that entries $\widehat{T}^{\mathcal{NN}+\delta}(q'|q, u)$ are nonzero only when $u \in \{\mathcal{NN}(\text{ct}(q)) \pm i\delta \mid i = 0, \dots, I\}$, i.e., the control actions are close to that suggested by the neural network. This avoids computing all the transition probabilities $\widehat{T}(q'|q, u)$, and searching for the optimal actions (that maximize $Q_t(q, u)$ in (10)) over the whole discretized input space, which are the computational bottlenecks for abstraction-based controller synthesis.

In Algorithm 2, NNSynth first computes entries of \widehat{T} that are suggested by the neural network \mathcal{NN} (line 7-13 of Algorithm 2). In particular, line 8 of Algorithm 2 checks if control action u is close to the action given by \mathcal{NN} , and computes the corresponding entries of \widehat{T} only if u has not been considered before at q , i.e., $u \notin U_{\text{buffer}}(q)$. The optimal control action at each state is determined by maximizing the Q-function (line 14-23 in Algorithm 2). Unique to NNSynth, it only searches the local action space that contains $\mathcal{NN}(q, t)$ at q (line 16 in Algorithm 2). Since the optimal policy is in general time-dependent, we explicitly include the time steps t in the input feature to the neural network \mathcal{NN} . In line 10 and 17 of Algorithm 2, $\widehat{B}_\rho(f(\text{ct}(q), u))$ denotes the subset of abstract states that are in a ball centered at $f(\text{ct}(q), u)$ with radius ρ , where ρ is a user-provided probability cut-off (when probability is smaller than the cut-off, the probability is treated as zero), which allows further speedup by limiting the transitions due to the model-error [8].

E. Step 5: Lift to NN

To further minimize the cost $J(\Psi_k)$, NNSynth “lifts” the abstraction-based controller Ψ_k found in the previous step to a neural network \mathcal{NN}_k , which allows us to employ the well-developed deep policy gradient approaches to update the

controller. Such lifting can be done by imitation learning with sampled trajectories of the continuous MDP Σ controlled by Ψ_k . The obtained neural network is then used as an initialization for further training by either reinforcement learning or imitation learning of the expert dataset \mathcal{D} . In Section IV, we analyze the performance of the synthesized controllers by taking into account the error due to the lift. This loop of training a NN, obtaining a local abstract model, synthesizing a controller, and lifting back to a NN is then continued until a controller is found.

IV. THEORETICAL ANALYSIS

A. Specification Satisfaction Guarantee

We provide theoretical guarantees of NNSynth in both satisfying the given specification ϕ and minimizing the cost functional J in this section. The satisfaction of ϕ with pre-specified probability is correct-by-construction. In particular, the procedure PROJECT-BY-SYNTH (Algorithm 2) maximizes the probability for the finite-state abstraction $\widehat{\Sigma}_{\widehat{\Psi}}^{\mathcal{N}+\delta}$ to satisfy ϕ , and the difference in the satisfaction probability is bounded between the finite-state abstraction and the original continuous system [15, Theorem 2.1]:

$$\left| \Pr\left(\widehat{\Sigma}_{\widehat{\Psi}}^{\mathcal{N}+\delta} \models \phi\right) - \Pr(\Sigma_{\Psi} \models \phi) \right| \leq \lambda H \mathcal{A} L_{\tau}, \quad (11)$$

where $\Psi(x) = \widehat{\Psi}(\text{abs}(x))$, λ is the grid size in partitioning the state space, H is the bounded time horizon, \mathcal{A} is the Lebesgue measure of the state space X , and L_{τ} is the Lipschitz constant of the stochastic kernel τ . Therefore, we only need to set the margin $\varsigma = \lambda H \mathcal{A} L_{\tau}$ in Algorithm 1 to ensure that the continuous MDP Σ_{Ψ} satisfies ϕ with the pre-specified probability p .

Theorem 4.1: Consider Algorithm 1 returns an abstraction-based controller Ψ_k with average probability $V_{\text{avg}} > p + \varsigma$, where $\varsigma = \lambda H \mathcal{A} L_{\tau}$. Then, the continuous MDP Σ controlled by Ψ_k is guaranteed to satisfy the given specification ϕ with probability at least p , i.e., $\Pr(\Sigma_{\Psi_k} \models \phi) \geq p$.

B. Optimality Guarantee

Now, we focus on the performance analysis of NNSynth, i.e., the optimality of controllers returned by Algorithm 1 in terms of minimizing the cost functional J . In Algorithm 1, the procedure UPDATE (line 8 in Algorithm 1) improves the neural network \mathcal{N}_k using its gradient, i.e., $\mathcal{N}_{k+1} = \mathcal{N}_k - \eta \nabla J(\mathcal{N}_k)$, where η is the learning rate and $\nabla J(\mathcal{N}_k)$ can be evaluated as (4). This can be treated as an approximation of updating the abstraction-based controller Ψ_k directly through $\Upsilon_{k+1} = \Psi_k - \eta \nabla J(\Psi_k)$, where $\Upsilon_{k+1} : X \rightarrow U$ is not necessarily an abstraction-based controller and need to be projected back to the abstraction-based controller space \mathcal{S} . We take into account this gradient approximation error, along with the lift and projection errors corresponding to line 7 and line 9 in Algorithm 1, to provide the overall performance guarantee of NNSynth in terms of regret as follows:

Theorem 4.2: Consider the loop (line 4-9) in Algorithm 1 executes K iterations and the abstraction-based controller obtained at the end of each iteration is Ψ_k , $k = 1, \dots, K$.

TABLE I
COMPARISON BETWEEN NNSYNTH AND AMYTISS.

Benchmark	2-d Robot	5-d Room Temp.	5-d Traffic
Specification ϕ	Reach-avoid	Safety	Safety
Specification horizon H	16	8	7
Problem complexity $ X \times \hat{U} $	705600	3429216	1.25×10^8
Satisfaction Probability V_{avg}	96%	95%	80%
NNSynth (time) [s]	49.0	319.1	367.7
AMYTISS (time) [s]	108.4	34640.0	23100.0
Speedup	2 x	108 x	62 x

Let Ψ^* be the optimal abstraction-based controller, i.e., $\Psi^* = \operatorname{argmin}_{\Psi \in \mathcal{S}} J(\Psi)$ s.t. $\Pr(\Sigma_{\Psi} \models \phi) \geq p$. Then, the regret over K iterations is upper bounded as follows:

$$\frac{1}{K} \sum_{k=1}^K J(\Psi_k) - J(\Psi^*) = O\left(\frac{1}{\eta K} + \frac{\delta I + \lambda L_{nn}}{\eta} + \lambda L_{nn} + \eta\right). \quad (12)$$

In the above theorem, η is the learning rate, λ is the grid size in partitioning the state space, δ and I are the precision and range parameters in system augmentation, respectively (see (6)), and L_{nn} is the Lipschitz constant of the trained neural network \mathcal{N} , i.e., $\|\mathcal{N}(x_1) - \mathcal{N}(x_2)\| \leq L_{nn} \|x_1 - x_2\|$, $\forall x_1, x_2 \in X$. Due to the space limit, we present the proof of Theorem 4.2 in the extended version [16].

V. RESULTS

We implemented NNSynth in Python and evaluated its performance on a Macbook Pro 15 with 32 GB RAM and Intel Core i9 2.4-GHz CPU. To compare with existing tools, we run all experiments on a single CPU core without using GPUs to accelerate neural network training.

We compared the performance of NNSynth with the state-of-the-art tool in synthesizing controllers for stochastic systems, AMYTISS [8], on three benchmarks with increasing complexity. Table I summarizes the comparison results. For each of the benchmarks, we list the specification ϕ used in the experiment along with its horizon H , the complexity of the problem measured by the number of abstract states times the number of discretized control actions $|X| \times |\hat{U}|$, the average probability of satisfying the specification (averaged over the state space) V_{avg} , the execution time for each of the two tools, and the corresponding speedup. Indeed, the last row in Table I empirically proves that using neural networks to guide the controller synthesis provides significant improvement to the overall execution time. Below, we provide more details about each of the benchmarks.

Experiment #1: 2-d Robot. Consider a 2-dimensional robot model given by:

$$\begin{aligned} x_1^{(t+1)} &= x_1^{(t)} + u_1^{(t)} \cos(u_2^{(t)}) + \varsigma_1^{(t)} \\ x_2^{(t+1)} &= x_2^{(t)} + u_2^{(t)} \sin(u_2^{(t)}) + \varsigma_2^{(t)}, \end{aligned}$$

where the state space $X = [-10, 10] \times [-10, 10]$, control input space $U = [-1, 1] \times [-1, 1]$, and the noise $(\varsigma_1, \varsigma_2)$ follows a Gaussian distribution with covariance matrix $\Sigma = \text{diag}(0.75, 0.75)$. We are interested in the task of steering the robot into a goal set $[5, 7] \times [5, 7]$ in 16 time steps, while avoiding the obstacle set $[-2, 2] \times [-2, 2]$ (see Figure 3).

To construct the abstraction-based controller, we partition the state space with discretization parameters $(0.5, 0.5)$, and the input space with $(0.1, 0.1)$. NNSynth starts by training a




Fig. 2. State trajectories sampled from different initial conditions using the synthesized controller in Experiment #2.

neural network using imitation learning with a total of 121 expert trajectories. The neural network consists of two hidden layers and ten neurons per hidden layer. We used Keras to train the neural network with the default adaptive learning rate optimization algorithm ADAM. The controller synthesis is then executed to find a controller $\hat{\Psi}$ that maximizes the probability of satisfying the specification, and one was found in 49.0 seconds with an average satisfaction probability of 96%. In Figure 3, we present 8 example trajectories under the control of $\hat{\Psi}$, by sampling some initial states. Using the same discretization parameters, AMYTISS was able to find a controller that satisfies the specs with 93% probability in 108.4 seconds. This shows a $2.2\times$ speedup of our tool with an increase in the satisfaction probability.

Experiment #2: 5-d Room Temperature Control. This example considers temperature regulation of 5 rooms each equipped with a heater and connected on a circle [8]. The state variables are temperatures of individual rooms, and the evolution of the 5 room temperatures is described as:

$$\begin{aligned} T_i^{(t+1)} &= a_{ii}T_i^{(t)} + \gamma T_h u_i^{(t)} + \eta w_i^{(t)} + \beta T_{ei} + 0.01\zeta_i^{(t)}, \quad i \in \{1, 3\} \\ T_i^{(t+1)} &= b_{ii}T_i^{(t)} + \eta w_i^{(t)} + \beta T_{ei} + 0.01\zeta_i^{(t)}, \quad i \in \{2, 4, 5\} \end{aligned}$$

where $a_{ii} = (1 - 2\eta - \beta - \gamma u_i^{(t)})$, $b_{ii} = (1 - 2\eta - \beta)$, and $w_i^{(t)} = T_{i-1}^{(t)} + T_{i+1}^{(t)}$ (with $T_0 = T_5$ and $T_6 = T_1$), and the parameters $\eta = 0.3$, $\beta = 0.022$, $\gamma = 0.05$, $T_{ei} = -1$, $T_h = 50$.

We consider a safety specification that requires the temperature of each room to maintain in the safe set [18.8, 21.2] for at least 8 time steps. As shown in Table I, NNSynth achieves a satisfaction probability of 95% and $108\times$ speedup compared to AMYTISS. In Figure 2, we sample 100 initial states and present the evolution of the 5 state variables, which are all maintained within the safe set for at least 8 steps under the abstraction-based controller provided by NNSynth.

Experiment #3: 5-d Road Traffic Network. This example considers a road traffic network divided into 5 cells, and state variables x_i denote the number of vehicles per cell [8]. The 5-d road traffic network is modeled as:

$$\begin{aligned} x_1^{(t+1)} &= (1 - \frac{\tau v_1}{L_1})x_1^{(t)} + \frac{\tau v_5}{L_5}w_1^{(t)} + 6u_1^{(t)} + 0.7\zeta_1^{(t)} \\ x_i^{(t+1)} &= (1 - \frac{\tau v_i}{L_i} - q)x_i^{(t)} + \frac{\tau v_{i-1}}{L_{i-1}}w_i^{(t)} + 0.7\zeta_i^{(t)}, \quad i \in \{2, 4\} \\ x_3^{(t+1)} &= (1 - \frac{\tau v_3}{L_3})x_3^{(t)} + \frac{\tau v_2}{L_2}w_3^{(t)} + 8u_2^{(t)} + 0.7\zeta_3^{(t)} \\ x_5^{(t+1)} &= (1 - \frac{\tau v_5}{L_5})x_5^{(t)} + \frac{\tau v_4}{L_4}w_5^{(t)} + 0.7\zeta_5^{(t)} \end{aligned}$$

where $w_i^{(t)} = x_{i-1}^{(t)}$ (with $x_0 = x_5$). Given the state space $X = [0, 10]^5$, the input space $U = [0, 1]^2$, and a noise covariance matrix $\Sigma = \text{diag}(0.7, 0.7, 0.7, 0.7, 0.7)$, we are interested in designing a control strategy that keeps the number of vehicles per cell in a safety set $[0, 10]$ for at least 7 steps. As shown in Table I, NNSynth was able to solve this problem in 367.7 seconds achieving more than $60\times$ speedup compared with AMYTISS.




Fig. 3. Closed-loop trajectories sampled from different initial states using the synthesized controller in Experiment #1.

REFERENCES

- [1] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.
- [2] M. Mazo, A. Davitian, and P. Tabuada, “Pessoa: a tool for embedded controller synthesis,” in *International conference on computer aided verification*. Springer, 2010, pp. 566–569.
- [3] S. Mouelhi, A. Girard, and G. Gössler, “CoSyMA: a tool for controller synthesis using multi-scale abstractions,” in *Proceedings of the 16th international conference on Hybrid systems: computation and control*, 2013, pp. 83–88.
- [4] M. Rungger and M. Zamani, “SCOTS: a tool for the synthesis of symbolic controllers,” in *Proceedings of the 19th international conference on hybrid systems: Computation and control*, 2016, pp. 99–104.
- [5] P. Jagtap and M. Zamani, “QUEST: a tool for state-space quantization-free synthesis of symbolic controllers,” in *International conference on quantitative evaluation of systems*. Springer, 2017, pp. 309–313.
- [6] S. E. Z. Soudjani, C. Gevaerts, and A. Abate, “FAUST²: Formal abstractions of uncountable-state stochastic processes,” in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2015, pp. 272–286.
- [7] N. Cauchi and A. Abate, “StocHy-automated verification and synthesis of stochastic processes,” in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, 2019, pp. 258–259.
- [8] A. Lavaei, M. Khaled, S. Soudjani, and M. Zamani, “AMYTISS: parallelized automated controller synthesis for large-scale stochastic systems,” in *International Conference on Computer Aided Verification*. Springer, 2020, pp. 461–474.
- [9] M. Palan, G. Shevchuk, N. Charles Landolfi, and D. Sadigh, “Learning reward functions by integrating human demonstrations and preferences,” in *Robotics: Science and Systems*, 2019.
- [10] G. Anderson, A. Verma, I. Dillig, and S. Chaudhuri, “Neurosymbolic reinforcement learning with formally verified exploration,” *Advances in neural information processing systems*, 2020.
- [11] A. Verma, H. Le, Y. Yue, and S. Chaudhuri, “Imitation-projected programmatic reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 32, pp. 15 752–15 763, 2019.
- [12] G. Weiss, Y. Goldberg, and E. Yahav, “Extracting automata from recurrent neural networks using queries and counterexamples,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 5247–5256.
- [13] S. Carr, N. Jansen, and U. Topcu, “Verifiable rnn-based policies for pomdps under temporal logic constraints,” in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2020, p. 4121–4127.
- [14] C. E. Rasmussen and C. Williams, “Gaussian processes for machine learning,” *the MIT Press*, 2006.
- [15] A. Lavaei, S. Soudjani, A. Abate, and M. Zamani, “Automated verification and synthesis of stochastic hybrid systems: A survey,” in *Automatica*, 2021.
- [16] X. Sun and Y. Shoukry, “NNSynth: Neural network guided abstraction-based controller synthesis for stochastic systems,” *arXiv:2111.08853*, 2022.