

Document Version

Final published version

Citation (APA)

Solanki, P., Vertovec, N., Schnitzer, Y., van Beers, J. J., de Visser, C. C., & Abate, A. (2025). Certified Approximate Reachability (CARE): Formal Error Bounds on Deep Learning of Reachable Sets. In *Proceedings of the IEEE 64th Conference on Decision and Control (CDC 2025)* (pp. 3907-3912). (Proceedings of the IEEE Conference on Decision and Control). IEEE. <https://doi.org/10.1109/CDC57313.2025.11312709>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

Certified Approximate Reachability (CARE): Formal Error Bounds on Deep Learning of Reachable Sets

Prashant Solanki^{*1}, Nikolaus Vertovec^{*2}, Yannik Schnitzer²,
Jasper Van Beers¹, Coen de Visser¹, Alessandro Abate²

Abstract—Recent approaches to leveraging deep learning for computing reachable sets of continuous-time dynamical systems have gained popularity over traditional level-set methods, as they overcome the curse of dimensionality. However, as with level-set methods, considerable care needs to be taken in limiting approximation errors, particularly since no guarantees are provided during training on the accuracy of the learned reachable set. To address this limitation, we introduce an ϵ -approximate Hamilton-Jacobi partial differential equation (HJ-PDE), which establishes a relationship between training loss and accuracy of the *true* reachable set. To formally certify this approximation, we leverage Satisfiability Modulo Theories (SMT) solvers to bound the residual error of the HJ-based loss function across the domain of interest. Leveraging Counter Example Guided Inductive Synthesis (CEGIS), we close the loop around learning and verification, by fine-tuning the neural network on counterexamples found by the SMT solver, thus improving the accuracy of the learned reachable set. To the best of our knowledge, Certified Approximate Reachability (CARE) is the first approach to provide soundness guarantees on learned reachable sets of continuous dynamical systems.

I. INTRODUCTION

Ensuring the safety of autonomous systems under uncertainty is a fundamental challenge in control and verification. Reachability analysis plays a key role in addressing safety concerns across various domains, including spacecraft trajectory design [33], [35], ground transportation systems [20], [22], air traffic management [19], [31], and flight control [21], [23], [36]. The theoretical foundations of reachability stem from viability theory [3], [4], while computational techniques have been developed for both exact and approximate reachable set computations in hybrid systems [2], [7], [8], [27], [28]. These advances have formulated reachability analysis as an optimal control problem, where reachable, viable, or invariant sets are represented as level sets of a value function satisfying a Hamilton-Jacobi (HJ) partial differential equation (PDE) [5], [9], [10], [21], [24].

Traditionally, level set methods are used to compute the unique viscosity solution of the HJ equation [28]. However, even with high-order numerical schemes such as WENO methods, the non-differentiability of the value function introduces numerical inaccuracies [34]. While refining the grid used in finite element methods such as level set methods improves numerical accuracy, it is hard to predict what

constitutes a sufficiently fine grid. Moreover, these grid-based approaches suffer from the curse of dimensionality, as the number of required grid points grows exponentially with system dimensionality.

Deepreach [6], a recent approach leveraging neural networks, has enabled solving the HJ equation without relying on finite difference methods. If the loss function enforcing the PDE conditions converges to zero in the region of interest, the learned value function is the unique viscosity solution of the HJ-PDE, ensuring that its zero-level set correctly represents the reachable set. This approach improves scalability to higher-dimensional systems by mitigating the curse of dimensionality.

In this work, we demonstrate an additional advantage of using neural networks to approximate the value function: the ability to provide a formally bounded ϵ -accurate reachable set. While both level-set methods and DeepReach yield empirically accurate solutions, their adherence to the HJ-PDE conditions—and thus the validity of the reachable set—has not been quantified. Follow-up work to DeepReach uses probabilistic methods to recover the reachable set [17], [18], [30]. In contrast, we provide a formal approach to verification that does not rely on recovery of the reachable set, but rather provides a sound over- and under-approximation of the reachable set. Our key contributions are:

- 1) We establish a formal bound ϵ on the residual error, ensuring that the absolute value of the HJ-based loss function remains within a specified threshold across the domain of interest.
- 2) We demonstrate that the neural network-based value function, constrained by ϵ , provides a certified approximation of the true value function, allowing for over- and under-approximation of the reachable set.

The remainder of the paper is structured as follows: Section II provides an overview of the reachability problems and introduces the Hamilton-Jacobi (HJ) partial differential equation (PDE) whose solution we aim to learn. Section III presents the deep learning approach for computing the reachable set, while Section IV discusses the formal verification method used to derive and certify an upper bound on the loss across the entire domain of interest. In Section V, we introduce our main contribution, establishing a connection between the residual training loss and an under- and over-approximation of the reachable set. Finally, Section VI provides implementation details and a case study.

^{*}Authors contributed equally to this article.

¹Department of Aerospace Engineering, TU Delft, Netherlands {p.solanki, j.j.vanbeers, c.c.devisser}@tudelft.nl

²Department of Computer Science, University of Oxford, OX1 3PJ, UK {nikolaus.vertovec, yannik.schnitzer, alessandro.abate}@cs.ox.ac.uk

II. PROBLEM SETUP

In reachability theory, a key objective is to compute the backward reachable set (BRS) of a dynamical system, which consists of all states from which trajectories can reach a given target set at the end of the time horizon. In contrast, the backward reachable tube (BRT) represents the set of states that can reach the target set within a specified time horizon [5]. If the target set represents unsafe states, the BRS/BRT identifies states that may lead to unsafe conditions and should be avoided. To account for adversarial disturbances, safety-critical scenarios are often modeled as a two-player game, where Player 1 represents the control input and Player 2 represents the disturbance input [25].

Mathematically, consider a dynamical system with state $x \in \mathbb{R}^m$ governed by the ordinary differential equation (ODE)

$$\dot{x}(s) \in f(s, x(s), u(s), d(s)), \quad t_0 \leq s \leq T, \quad (1)$$

where $T \geq t_0$ is the fixed time horizon. The initial state is given by $x(t_0) := x_0$, and the control and disturbance inputs are measurable functions

$$u : [t_0, T] \rightarrow \mathcal{U}, \quad d : [t_0, T] \rightarrow \mathcal{D},$$

with $\mathcal{U} \subset \mathbb{R}^k$ and $\mathcal{D} \subset \mathbb{R}^l$ being compact sets. The dynamics $f : [t_0, T] \times \mathbb{R}^m \times \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}^m$ satisfy the following conditions for some constants C_1, C_2 :

$$|f(t, x, u, d)| \leq C_1, \quad (2)$$

$$|f(t, x, u, d) - f(t, \hat{x}, u, d)| \leq C_2|x - \hat{x}|, \quad (3)$$

for all $t \in [t_0, T]$, $x, \hat{x} \in \mathbb{R}^m$, $u \in \mathcal{U}$, and $d \in \mathcal{D}$. These assumptions ensure that the ODE (1) admits a unique solution

$$x(t) = \phi(t, t_0, x_0, u(\cdot), d(\cdot)). \quad (4)$$

Next, we define the sets of control and disturbance policies:

$$\mathcal{M}_{[t, T]} \equiv \{u : [t, T] \rightarrow \mathcal{U} \mid u \text{ is measurable}\},$$

$$\mathcal{N}_{[t, T]} \equiv \{d : [t, T] \rightarrow \mathcal{D} \mid d \text{ is measurable}\}.$$

A nonanticipative strategy is a mapping $\beta : \mathcal{M}_{[t, T]} \rightarrow \mathcal{N}_{[t, T]}$ such that for all $s \in [t, T]$ and for all $u, \hat{u} \in \mathcal{M}_{[t, T]}$, if $u(\tau) = \hat{u}(\tau)$ for almost every $\tau \in [t, s]$, then $\beta[u](\tau) = \beta[\hat{u}](\tau)$ for almost every $\tau \in [t, s]$. The class of such strategies is denoted as $\Delta_{[t, T]}$.

Given a target set $\mathcal{G}$ with a signed distance function $g(x)$ such that $g(x) \leq 0 \iff x \in \mathcal{G}$, the BRT is defined as

$$\text{BRT}_{\mathcal{G}}([t_0, T]) = \{x \mid \exists \beta \in \mathcal{N}_{[t_0, T]}, \forall u \in \mathcal{M}_{[t_0, T]}, \exists s \in [t_0, T] : \phi(s, t_0, x_0, u(\cdot), \beta(\cdot)) \in \mathcal{G}\}. \quad (5)$$

Equivalently, the BRT is given by the subzero level set of the value function $V^*(t, x)$, where

$$V^*(t, x) = \inf_{\beta \in \Delta_{[t, T]}} \sup_{u \in \mathcal{M}_{[t, T]}} \inf_{\tau \in [t, T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))). \quad (6)$$

Let us introduce the shorthand notation $D_x V(t, x) = \frac{\partial V(t, x)}{\partial x}$ and $D_t V(t, x) = \frac{\partial V(t, x)}{\partial t}$. Then, a standard result of HJ

reachability [5], [11], [25] allows us to reduce the optimization problem in Equation (6) from an infinite-dimensional optimization over policies to an optimization over control and disturbance inputs by formulating the value function as the viscosity solution of the HJ-PDE:

$$D_t V^*(t, x) + \min\{0, \mathcal{H}(t, x, D_x V^*)\} = 0, \quad V^*(T, x) = g(x), \quad (7)$$

with the Hamiltonian defined as

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p].$$

The HJ-PDE can be solved using level-set methods [21], [27], [32] or through curriculum training of a neural network [6].

III. DEEP LEARNING APPROACH TO HJ REACHABILITY

Inspired by the self-supervised framework of Physics-Informed Neural Networks (PINNs) [29], the solution of the PDE (7) can be effectively learned, as demonstrated in [6].

Let V_θ denote the approximate value function learned by a neural network. To ensure that the zero level set of V_θ correctly represents the BRT, the network must achieve zero loss for the loss function defined as

$$\mathcal{L}(t_i, x_i; \theta) = \mathbb{1}(t_i = T) \mathcal{L}_1(x_i; \theta) + \lambda \mathcal{L}_2(t_i, x_i; \theta), \quad (8)$$

$$\mathcal{L}_1(x_i; \theta) = \|V_\theta(T, x_i) - g(x_i)\|, \quad (9)$$

$$\mathcal{L}_2(t_i, x_i; \theta) = \|D_t V_\theta(t_i, x_i) + \min[0, \mathcal{H}(t_i, x_i, D_x V_\theta)]\|, \quad (10)$$

where $\mathbb{1}(t_i = T)$ is an indicator function ensuring that $\mathcal{L}_1$ is only applied at $t_i = T$, and λ controls the balance between the two loss terms. The nonlinear nature of the neural network might result in areas of the state space with significantly higher loss than empirically observed during training. This acute shortcoming necessitates a formal approach to verifying the approximation errors. To this end, in the subsequent sections, we will

- 1) establish a verification method to ensure $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$ for all x and t within the domain of interest; and
- 2) derive upper and lower bounds on V_θ in terms of viscosity solutions to the HJ-PDE, thus enabling under- and over-approximations of the BRT/BRS.

IV. FORMAL SYNTHESIS OF NEURAL VALUE FUNCTIONS

Our approach to synthesizing the value function solution follows a two-phase process—a learning phase and a certification phase—that alternates in a Counter Example Guided Inductive Synthesis (CEGIS) loop [1]. The learning phase employs a curriculum training approach, discussed in detail in Section VI, terminating when we empirically satisfy $|\mathcal{L}_1(x_i; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t_i, x_i; \theta)| < \epsilon_2$, for N randomly sampled points $(t_1, x_1), \dots, (t_n, x_n)$.

To further ensure that $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$ for all x in the domain of interest $\mathcal{X}$ and for all $t \in [t_0, T]$, we design a sound certification phase: this certification is performed via SMT (Satisfiability Modulo Theories) solving,

which symbolically reasons over the continuous domain $[t_0, T] \times \mathcal{X}$, hence generalizing across the sample-based loss error. While computationally intensive, formal verification provides soundness guarantees, allowing us to rigorously validate the learned value function as an approximate solution to the HJ-PDE.

Several approaches exist for certifying the accuracy of the learned value function. To remain general in our choice of dynamical systems, we employ an SMT solver capable of handling quantifier-free nonlinear real arithmetic formulae [12]. This allows us to incorporate arbitrary nonlinear activation functions into the neural network. Furthermore, it enables us to address systems with nonlinear dynamics, which induce nonlinearities in the Hamiltonian. Specifically, we employ `dReal` [14], which supports both polynomial and non-polynomial terms, such as transcendental functions (e.g., trigonometric and exponential functions).

After generating a symbolic representation of the neural network, the SMT solver searches for an assignment of variables (t, x) that satisfies the quantifier-free formula

$$(x \in \mathcal{X} \wedge \|V_\theta(t, x) - g(x)\| > \epsilon_1) \vee \quad (11)$$

$$(x \in \mathcal{X} \wedge t \in [t_0, T] \wedge \|D_t V_\theta(t, x) + \min[0, \mathcal{H}(t, x, D_x V_\theta)]\| > \epsilon_2) \quad (12)$$

The logical disjunction between (11) and (12) allows us to split the SMT call into separate parallel queries. If one query identifies a valid assignment, i.e., a counterexample to either $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ or $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$, the remaining SMT calls terminate. A counterexample indicates the necessity to finetune the neural network, as it does not yet approximate a valid value function with sufficient accuracy. We leverage the counterexample for targeted training specifically at inputs where the network is insufficiently accurate. We sample additional training data points for finetuning in close proximity to the counterexample, thus reducing the need for extensive sampling in regions of the input space that are already well approximated. Finetuning and certification are alternately repeated within an inductive synthesis loop (CEGIS), progressively improving the neural value function in targeted regions of the input space, until the certifier determines the Formulae (11) and (12) to be unsatisfiable. This implies that there exists no counterexample to $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$, formally proving the validity of the learned value function over the entire domain of interest. In the following section, we derive results establishing how this formal certification can be leveraged to obtain sound over- and under-approximations of the BRT and BRS from the neural value function.

V. ϵ -ACCURATE VALUE FUNCTIONS

The certification of the neural value function in the previous section established that for all $(t, x) \in [t_0, T] \times \mathcal{X}$, there exists some $\hat{\epsilon} \in [0, \epsilon_2]$ such that

$$\|D_t V_\theta(t, x) + \min[0, \mathcal{H}(t, x, D_x V_\theta)]\| = \hat{\epsilon},$$

which can equivalently be rewritten as

$$\begin{aligned} \|D_t V_\theta(t, x) + \min[\pm \hat{\epsilon}, \tilde{\mathcal{H}}(t, x, D_x V_\theta)]\| &= \hat{\epsilon}, \\ \tilde{\mathcal{H}}(t, x, p) &= \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p \pm \hat{\epsilon}]. \end{aligned}$$

We consider the worst-case realisation of $\hat{\epsilon}$, i.e., $\hat{\epsilon} = \epsilon_2$. We will now show that, for $\epsilon = (\epsilon_1 + \epsilon_2(T - t))$,

$$V_\theta(t, x) - \epsilon \leq V^*(t, x) \leq V_\theta(t, x) + \epsilon, \quad (13)$$

which allows us to use the neural value function to reason about the true BRT and BRS.

A. Bounding the True Value Function

To derive these bounds, we introduce the modified value functions $\bar{V}$ and $\underline{V}$, which are the unique viscosity solutions under the worst-case realizations of $\hat{\epsilon}$:

$$\begin{aligned} \bar{V}(t, x) &= \inf_{\beta(\cdot) \in \Delta_{[t, T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t, T]}} \\ &\quad \inf_{\tau \in [t, T]} \left[\int_t^\tau \epsilon_2 ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right], \end{aligned} \quad (14)$$

$$\begin{aligned} \underline{V}(t, x) &= \inf_{\beta(\cdot) \in \Delta_{[t, T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t, T]}} \\ &\quad \inf_{\tau \in [t, T]} \left[\int_t^\tau -\epsilon_2 ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right]. \end{aligned} \quad (15)$$

Since ϵ_2 is a non-negative constant, it follows that

$$\underline{V}(t, x) \leq V^*(t, x) \leq \bar{V}(t, x). \quad (16)$$

Furthermore, using the inequality $-\epsilon_2(T - t) \leq \int_t^T -\epsilon_2 ds$, we obtain:

$$\begin{aligned} &\inf_{\tau \in [t, T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) - \epsilon_2(T - t) \\ &\leq \inf_{\tau \in [t, T]} \left[g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) - \int_t^\tau \epsilon_2 ds \right]. \end{aligned}$$

Similarly, using $\epsilon_2(T - t) \geq \int_t^T \epsilon_2 ds$, we obtain:

$$\begin{aligned} &\inf_{\tau \in [t, T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) + \epsilon_2(T - t) \\ &\geq \inf_{\tau \in [t, T]} \left[g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) + \int_t^\tau \epsilon_2 ds \right]. \end{aligned}$$

Thus, we derive the key bound:

$$\bar{V}(t, x) - \epsilon_2(T - t) \leq V^*(t, x) \leq \underline{V}(t, x) + \epsilon_2(T - t). \quad (17)$$

B. Relating Bounds to the Neural Value Function

To relate $\bar{V}$ and $\underline{V}$ to the neural value function, we need to establish that these auxiliary value functions are the unique viscosity solutions of the ϵ -modified HJ-PDE. This result is formalized in the following theorem.

Theorem 1: The function $\bar{V}(t, x)$ is the unique viscosity solution of the ϵ -modified HJ-PDE:

$$D_t \bar{V}(t, x) + \min \{ \epsilon_2, \mathcal{H}(t, x, D_x \bar{V}) \} = 0,$$

with the Hamiltonian defined as

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p + \epsilon_2],$$

and the terminal condition

$$\bar{V}(T, x) = g(x).$$

Similarly, $\underline{V}(t, x)$ is the unique viscosity solution of the ϵ -modified HJ-PDE:

$$D_t \underline{V}(t, x) + \min \{-\epsilon_2, \mathcal{H}(t, x, D_x \underline{V})\} = 0,$$

with the Hamiltonian

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p - \epsilon_2],$$

and the terminal condition

$$\underline{V}(T, x) = g(x).$$

The proof of Theorem 1 is given in the extended version of the paper¹.

As a consequence of Theorem 1, we obtain

$$\underline{V}(t, x) \leq V_\theta(t, x) \leq \bar{V}(t, x). \quad (18)$$

Finally, incorporating ϵ_1 , which perturbs the boundary condition, we have:

$$V_\theta(T, x) = g(x) \pm \epsilon_1. \quad (19)$$

Thus, by combining Equation (17) with Equations (18) and (19), we recover the bound in Equation (13).

VI. CASE STUDIES

For clarity, we focus on outlining the essential steps involved in the training and verification process, thereby excluding a comprehensive comparison of different reach/avoid scenarios. The provided toolbox² includes additional examples, such as an evader–pursuer scenario, forward and backward reachability problems, and implementations of both the reachable tube and the reachable set. Although our theoretical discussion throughout the paper is framed in terms of backward reachable sets and tubes (BRS/BRT), the same guarantees extend directly to forward reachability. We therefore illustrate the approach on a forward reachability problem for the double integrator, as this canonical system provides a clear and concise benchmark to demonstrate the interaction between training and SMT-based certification.

Let us consider the simple double integrator example, a canonical second-order control system, $\dot{x}_1 = x_2, \dot{x}_2 = u$, where x_1 represents position, x_2 represents velocity, and u is the control input. We consider the problem of learning the forward reachable set, with the initial set defined by

$$g(x) = x_1^2 + x_2^2 - R^2, \quad (20)$$

with $R^2 = 0.5$.

¹<https://arxiv.org/abs/2503.23912>

²<https://github.com/nikover/CARe>

A. Training of $V_\theta(t, x)$

The training procedure, outlined in Algorithm 1 and illustrated in Fig. 1 follows a curriculum learning strategy with three phases:

- **Pretraining phase:** Focuses on boundary conditions at the final time $t = T$.
- **Curriculum phase:** Gradually extends the time horizon from $[T, T]$ to $[t_0, T]$.
- **Finetuning phase:** Further refines the model to minimize loss before certification.

Each training epoch involves generating a random batch of samples and computing the loss $\mathcal{L}(t_i, x_i, \theta)$ for each sample $i \in \{1, \dots, N\}$. The formal error bounds on the reachable set can be conservative as the results in Theorem 1 assume the worst-case realisation of ϵ_2 along any trajectory. Consequently, even when the mean training loss is low, the existence of regions of the state space with high network loss can lead to conservative over- or under-approximations of the reachable set. To mitigate this, we not only minimize the mean error over the batch but also the maximum error, leading to the training objective of minimizing:

$$\frac{1}{N} \sum_i \mathcal{L}(t_i, x_i, \theta) + \lambda_{\max} \max_i \mathcal{L}(t_i, x_i, \theta).$$

Here, $\lambda_{\max}$ is set to 0.1 during the curriculum phase and 0.3 during the finetuning phase. The finetuning phase uses a patience counter ($p_{\text{patience}} = 1000$), terminating only when the loss consistently stays below a threshold $\lambda_\epsilon \epsilon_2$, with $\lambda_\epsilon = 0.95$. This modified approach to the curriculum training scheme presented in [6] reduces the total training time while ensuring sufficient confidence in the model prior to progressing to the certification phase.

To improve performance and enable smaller networks to approximate the value function effectively, we introduce a polynomial layer as the first layer of the network. This allows subsequent layers to operate not only on the original inputs (t, x) but also on polynomial transformations, such as (t^2, x^2) .

B. Formal Verification

To formally verify that the trained network represents an ϵ -accurate value function, we generate a symbolic representation $V_\theta(t, x)$ to instantiate the quantifier-free SMT query described in Section IV. To obtain the corresponding derivative terms $D_t V_\theta(t, x)$ and $\mathcal{H}(t, x, D_x V_\theta)$ we apply symbolic differentiation. To further parallelize the SMT call, we decompose the absolute value expressions of Equations (11)–(12) into separate calls, e.g., we seek x s.t.

$$\begin{aligned} x &\in \mathcal{X} \wedge V_\theta(T, x) - g(x) > \epsilon_1, \\ \vee x &\in \mathcal{X} \wedge V_\theta(T, x) - g(x) < -\epsilon_1. \end{aligned}$$

As SMT with nonlinear real arithmetic is undecidable in general, no exact decision procedure can exist for arbitrary HJ-PDEs [16]. However, dreal provides a δ -complete decision procedure that allows for user-specified δ deviations in the satisfying assignments [13], [15]. Consequently, while




Fig. 1. Flowchart of the training and verification Procedure

Algorithm 1 Training of $V_\theta(t, x)$

```

1: Input: Thresholds  $\epsilon_1$  and  $\epsilon_2$ 
2: Initialize: Model  $V_\theta$ ,  $t_{\text{current}} \leftarrow T$ 
3: while Training do
4:   Generate samples  $x_i$  and  $t_i \in [t_{\text{current}}, T]$ 
5:   Compute model output  $\hat{y} = V_\theta(t_i, x_i)$ 
6:   Compute loss  $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{mean}} + \lambda_{\text{max}} \cdot \mathcal{L}_{\text{max}}$ 
7:   Backpropagate  $\mathcal{L}_{\text{total}}$  and update parameters  $\theta$ 
8:   if  $\mathcal{P}_{\text{pre-training}}$  and  $\max(\mathcal{L}_1) < \epsilon_1$  then
9:     Progress to  $\mathcal{P}_{\text{curriculum}}$ 
10:  else if  $\mathcal{P}_{\text{curriculum}}$  and  $\max(\mathcal{L}_2) < \epsilon_2$  then
11:    Expand time horizon:  $t_{\text{current}} \leftarrow t_{\text{current}} - \Delta_t$ 
12:    if  $t_{\text{current}} = 0$  then
13:      Progress to  $\mathcal{P}_{\text{finetune}}$ 
14:       $p \leftarrow 0$ 
15:    end if
16:  else if  $\mathcal{P}_{\text{finetune}}$  and  $\max(\mathcal{L}_2) < \lambda_\epsilon \epsilon_2$  then
17:     $p \leftarrow p + 1$ 
18:    if  $p > p_{\text{patience}}$  then
19:      break ▷ Stop training
20:    end if
21:  end if
22: end while
23: Output: Trained model  $V_\theta$ 

```




Fig. 2. The over-approximated reachable set is shown in red, while the learned reachable set is depicted in blue. The initial set at time $t_0 = 0$ is a circle around the centre of radius 0.25, the time horizon is $T = 1$.

satisfying assignments may be spurious, unsatisfiability carries over to the exact problem. Hence, if the SMT query is unsatisfiable, it formally establishes a global upper bound on the loss function across the entire domain of interest.

TABLE I
TRAINING AND VERIFICATION RESULTS OF THE CEGIS LOOP

Iter.	ϵ	Training (s)	Verif. (s)	Result
1	0.30	836.92	1367	Certified
2	0.27	-	22	Counterexample found
3	0.27	181.27	2735	Certified
4	0.243	-	5537	Certified
5	0.2187	-	237	Counterexample found

Checking SMT queries with nonlinear real arithmetic can be computationally expensive, especially as unsatisfiability constitutes a full proof for validity. In practice, however, if a counterexample exists, the SMT query typically terminates within a reasonable timeframe.

For illustration, we consider training and verifying the double integrator system using a single hidden layer with 16 neurons, employing sine activation functions. After training with $\epsilon := (\epsilon_1 + \epsilon_2(T - t_0))$, and $T = 1, t_0 = 0, \epsilon_2 = 0.95\epsilon = 0.285, \epsilon_1 = 0.05\epsilon = 0.015$ we progress to certification, which takes 1367 seconds resulting in no counterexample being found. Subsequently, we attempt to reduce ϵ by 10%. Within 22 seconds, one of the parallelized SMT queries finds a counterexample. Thus, at this stage, our trained neural network satisfies $\mathcal{L}(t, x) < 0.3$ but not $\mathcal{L}(t, x) < 0.27$. Following a Counterexample-Guided Inductive Synthesis (CEGIS) approach, we leverage the identified counterexample to refine our model (Fig. 1). Specifically, we re-enter the finetuning phase of training with the reduced ϵ this time ensuring that 10% of our samples are drawn near the counterexample. This refinement improves the model such that a subsequent verification attempt fails to find a new counterexample, thereby establishing a tighter bound on the loss. We continue the CEGIS loop for a total of 5 iterations with the results shown in Table I. The fine-tuning phase of

the final iteration is unable to obtain an empirical loss below 0.2187, thus, the best certifiable result with $\epsilon = 0.243$ is used to produce the final over-approximated reachable set, shown in Fig. 2.

Our approach is applicable to arbitrary continuous-time reachability problems. However, the use of *dreal* practically limits the network size, as increasing the depth and number of neurons adds complexity to the formulas (11)-(12) [26]. If an empirical result is satisfactory, Theorem 1 can provide an empirical estimate of the over- or under-approximation of the reachable set.

VII. CONCLUSIONS

The primary objective of this paper is to establish formal bounds on the approximation error induced from computing Reachable Sets/Tubes with neural networks. To achieve this, we introduce a modified HJ-PDE that ties bounds on the training loss to the true solutions of the reachability problem. This development represents a crucial step toward providing soundness guarantees for learned reachable sets of continuous dynamical systems. To obtain bounds on the final training loss—and thereby deliver the formal assurances required in safety-critical contexts—we employ SMT-based verification. However, the inherent scalability limitations of SMT-based methods remain a significant challenge, and addressing this bottleneck constitutes an important direction for future research.

REFERENCES

- [1] A. Abate, C. David, P. Kesseli, D. Kroening, and E. Polgreen. *Counterexample Guided Inductive Synthesis Modulo Theories*, pages 270–288. Springer International Publishing, 2018.
- [2] E. Asarin, O. Bournez, T. Dang, O. Maler, and A. Pnueli. Effective synthesis of switching controllers for linear systems. *Proceedings of the IEEE*, 88(7):1011–1025, 2000.
- [3] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre. *Viability theory: new directions*. Springer Science & Business Media, 2011.
- [4] J.-P. Aubin and H. Frankowska. Viability kernel of control systems. In *Nonlinear synthesis*, pages 12–33. Springer, 1991.
- [5] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 2242–2253. IEEE, 2017.
- [6] S. Bansal and C. J. Tomlin. Deepreach: A deep learning approach to high-dimensional reachability. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1817–1824. IEEE, 2021.
- [7] P. Cardaliaguet. A differential game with two players and one target. *SIAM Journal on Control and Optimization*, 34(4):1441–1460, 1996.
- [8] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. Differential games through viability theory: Old and recent results. In *Advances in dynamic game theory*, pages 3–35. Springer, 2007.
- [9] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking. *IEEE Transactions on Automatic Control*, 66(12):5861–5876, 2021.
- [10] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. Decomposition of reachable sets and tubes for a class of nonlinear systems. *IEEE Transactions on Automatic Control*, 63(11):3675–3688, 2018.
- [11] L. C. Evans and P. E. Souganidis. Differential games and representation formulas for solutions of hamilton-jacobi-isaacs equations. *Indiana University mathematics journal*, 33(5):773–797, 1984.
- [12] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure1. *Journal on Satisfiability, Boolean Modeling and Computation*, 1(3–4):209–236, May 2007.
- [13] S. Gao, J. Avigad, and E. M. Clarke. δ -complete decision procedures for satisfiability over the reals. In *IJCAR*, volume 7364 of *Lecture Notes in Computer Science*, pages 286–300. Springer, 2012.
- [14] S. Gao, S. Kong, and E. M. Clarke. *dReal: An SMT Solver for Nonlinear Theories over the Reals*, page 208–214. Springer Berlin Heidelberg, 2013.
- [15] S. Gao, S. Kong, and E. M. Clarke. dreal: An smt solver for nonlinear theories over the reals. In *International conference on automated deduction*, pages 208–214. Springer, 2013.
- [16] D. Kroening and O. Strichman. *Decision Procedures - An Algorithmic Point of View*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2008.
- [17] A. Lin and S. Bansal. Generating formal safety assurances for high-dimensional reachability. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, page 10525–10531. IEEE, May 2023.
- [18] A. Lin and S. Bansal. Verification of neural reachable tubes via scenario optimization and conformal prediction. In *Proceedings of the 6th Annual Learning for Dynamics and Control Conference*, volume 242 of *Proceedings of Machine Learning Research*, pages 719–731. PMLR, 15–17 Jul 2024.
- [19] C. Livadas, J. Lygeros, and N. A. Lynch. High-level modeling and analysis of the traffic alert and collision avoidance system (tcas). *Proceedings of the IEEE*, 88(7):926–948, 2000.
- [20] C. Livadas and N. A. Lynch. Formal verification of safety-critical hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 253–272. Springer, 1998.
- [21] J. Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.
- [22] J. Lygeros, D. N. Godbole, and S. Sastry. Verified hybrid controllers for automated vehicles. *IEEE transactions on automatic control*, 43(4):522–539, 1998.
- [23] J. Lygeros, C. Tomlin, and S. Sastry. Controllers for reachability specifications for hybrid systems. *Automatica*, 35(3):349–370, 1999.
- [24] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. Lagrangian methods for approximating the viability kernel in high-dimensional systems. *Automatica*, 49(7):2017–2029, 2013.
- [25] K. Margellos and J. Lygeros. Hamilton-jacobi formulation for reach-avoid differential games. *IEEE Transactions on automatic control*, 56(8):1849–1861, 2011.
- [26] F. B. Mathiesen, N. Vertovec, F. Fabiano, L. Laurenti, and A. Abate. Certified neural approximations of nonlinear dynamics. *arXiv preprint: 2505.15497*, 2025.
- [27] I. Mitchell, A. M. Bayen, and C. J. Tomlin. Validating a hamilton-jacobi approximation to hybrid system reachable sets. In *International workshop on hybrid systems: Computation and control*, pages 418–432. Springer, 2001.
- [28] I. M. Mitchell et al. A toolbox of level set methods. *UBC Department of Computer Science Technical Report TR-2007-11*, 1:6, 2007.
- [29] M. Raissi, P. Perdikaris, and G. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, Feb. 2019.
- [30] M. Tayal, A. Singh, S. Kolathaya, and S. Bansal. A physics-informed machine learning framework for safe and optimal control of autonomous systems. *arXiv preprint: 2502.11057*, 2025.
- [31] C. Tomlin, I. Mitchell, and R. Ghosh. Safety verification of conflict resolution manoeuvres. *IEEE Transactions on Intelligent Transportation Systems*, 2(2):110–120, 2001.
- [32] C. J. Tomlin, J. Lygeros, and S. S. Sastry. A game theoretic approach to controller design for hybrid systems. *Proceedings of the IEEE*, 88(7):949–970, 2000.
- [33] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Multi-objective minimum time optimal control for low-thrust trajectory design. In *2021 European Control Conference (ECC)*, pages 1975–1980, 2021.
- [34] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Verification of safety critical control policies using kernel methods. In *2022 European Control Conference (ECC)*, pages 1870–1875, 2022.
- [35] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Multi-objective low-thrust spacecraft trajectory design using reachability analysis. *European Journal of Control*, 69:100758, 2023.
- [36] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Safety-aware hybrid control of airborne wind energy systems. *Journal of Guidance, Control, and Dynamics*, 47(2):326–338, 2024.