Automated and Formal Synthesis of Neural Barrier Certificates for Dynamical Models

Andrea Peruffo¹, Daniele Ahmed², Alessandro Abate¹

¹ Department of Computer Science, University of Oxford,
Oxford, UK

`{name.surname}@cs.ox.ac.uk`

² Amazon Inc, London, UK

Abstract. We introduce an automated, formal, counterexample-based approach to synthesise Barrier Certificates (BC) for the safety verification of continuous and hybrid dynamical models. The approach is underpinned by an inductive framework: this is structured as a sequential loop between a learner, which manipulates a candidate BC structured as a neural network, and a sound verifier, which either certifies the candidate's validity or generates counter-examples to further guide the learner. We compare the approach against state-of-the-art techniques, over polynomial and non-polynomial dynamical models: the outcomes show that we can synthesise sound BCs up to two orders of magnitude faster, with in particular a stark speedup on the verification engine (up to three orders less), whilst needing a far smaller data set (up to three orders less) for the learning part. Beyond improvements over the state of the art, we further challenge the new approach on a hybrid dynamical model and on larger-dimensional models, and showcase the numerical robustness of our algorithms and codebase.

1 Introduction

Barrier Certificates (BC) are an effective and powerful technique to prove safety properties on models of continuous dynamical systems, as well as hybrid models (featuring both continuous and discrete states) [21,22]. Whenever found, a BC partitions the state space of the model into two parts, ensuring that all trajectories starting from a given initial set, located within one side of the BC, cannot reach a given set of states (deemed to be unsafe), located on the other side. Thus a successful synthesis of a BC (which is in general not a unique object) represents a formal proof of safety for the dynamical model. BC find various applications spanning robotics, multi-agent systems, and biology [7,32].

This work addresses the safety of dynamical systems modelled in general by non-linear ordinary differential equations (ODE), and presents a novel method for the automated and formal synthesis of BC. The approach leverages Satisfiability Modulo Theory (SMT) and inductive reasoning (CEGIS, Figure 1, introduced later), to guarantee the correctness of the automated synthesis procedure: this rules out both algorithmic and numerical errors related to BC synthesis [10].

Background and Related Work A few techniques have been developed to synthesise BC. For polynomial models, sum-of-squares (SOS) and semi-definite programming relaxations [14,16,29] convert the BC synthesis problem into constraints expressed as linear or bilinear matrix inequalities: these are numerically solved as a convex optimisation problem, however unsoundly. To increase scalability and to enhance expressiveness, numerous barrier formats have been considered: BC based on exponential conditions are presented in [14]; BC based on Darboux polynomials are outlined in [33]; [30] newly introduces a multi-dimensional generalisation of BC, thus broadening their scope and applicability. BC can also be used to verify safety of uncertain (e.g. parametric) models [20]. Let us remark that SOS approaches are typically *unsound*, namely they rely on iterative and numerical methods to synthesise the BC. [10] a-posteriori verifies SOS candidates via computer-aided design (CAD) techniques [15].

Model *invariants* (namely, regions that provably contain model trajectories, such as *basins of attractions* [28]) can be employed as BC, though their synthesis is less general, as it does not comprise an unsafe set and tacitly presupposes the initial set to be “well placed” within the state space (that is, within the aforementioned basin): [19] introduces a fixpoint algorithm to find algebraic-differential invariants for hybrid models; invariants can be characterised analytically [4] or synthesised computationally [8]. Invariants can be alternatively studied by *Lyapunov theory* [5], which provides *stability* guarantees for dynamical models, and thus can characterise invariants (and barriers) as side products: however this again requires that initial conditions are positioned around stable equilibria, and does not explicitly encompass unsafe sets in the synthesis. Whilst Lyapunov theory is classically approached either analytically (explicit synthesis) or numerically (with unsound techniques), an approach that is relevant for the results of this work looks at automated and sound Lyapunov function synthesis: in [27] Lyapunov functions are soundly found within parametric templates, by constructing a system of linear inequality constraints over unknown coefficients. [23,24,25] employ a counterexample-based approach to synthesise control Lyapunov functions, which inspires this work, using a combination of SMT solvers and convex optimisation engines: however unlike this work, SMT solvers are never used for verification, which is instead handled by solving optimisation problems that are numerically unsound. As argued above, let us emphasise again that the BC synthesis problem, as studied in this work, cannot in general be reduced to a problem of Lyapunov stability analysis, and is indeed more general.

Fig. 1. Schematic representation of the CEGIS loop.

Core approach We introduce a method that efficiently exploits machine learning, whilst guaranteeing formal proofs of correctness via SMT. We leverage a CounterExample-Guided Inductive Synthesis (CEGIS) procedure [31], which is structured as an inductive loop between a *Learner* and a *Verifier* (cf. Fig. 1). A learner numerically (and unsoundly) trains a neural network (NN) to fit over a finite set of samples the requirements for a BC, which are expressed through a loss function; then a verifier either formally proves the validity of the BC or provides (a) counter-example(s) through an SMT solver: the counter-examples indicate where the barrier conditions are violated, and are passed back to the learner for further training. This synthesis method for neural BC is formally sound and fully automated, and thanks to its specific new features, is shown to be much faster and to clearly require less data than state-of-the-art results.

Contributions beyond the State of the Art Cognate work [34] presents a method to compute BC using neural networks and to verify their correctness a-posteriori: as such, it does not generate counter-examples within an inductive loop, as in this work. [34] considers large sample sets that are randomly divided into batches and fed to a feed-forward NN; the verification at the end of the (rather long) training either validates the candidate, or invalidates it and the training starts anew on the same dataset. In Section 4 the method in [34] is shown to be slower (both in the training and in the verification), and to require more data than the CEGIS-based approach of this work, which furthermore introduces numerous bespoke optimisations, as outlined in Section 3: our CEGIS-based technique exploits fast learning, verification simplified by the candidates passed by the Learner, and an enhanced communication between Learner and Verifier. Our approach further showcases numerical robustness and scalability features.

Related to the work on BC is the synthesis of Lyapunov functions, mentioned above. The construction of *Lyapunov Neural Networks* (LNNs) has been studied with approaches based on simulations and numerical optimisation, which are in general unsound [26]. Formal methods for Lyapunov synthesis are introduced in [5], together with a counterexample-based approach using polynomial candidates. The work is later extended in [2], which employs NN as candidates over polynomial dynamical models. The generation of control Lyapunov functions using counterexample-based NN is similarly considered in [9], however this is done by means of differing architectural details and does not extend to BC synthesis. Beyond the work in [5], this contribution is not limited to a specific polynomial template, since it supports more general mixtures of polynomial functions obtained through the NN structure, as well as the canonical tanh, sigmoid, ReLU activations (we provide one example of BC using tanh activations). Compared to [5], where we use LP programming to synthesise Lyapunov functions, in this work: a) we use a template-free procedure, thanks to the integration of NNs - these are needed since template-based SOS-programming approaches are not sufficient to provide BCs for several of the presented benchmarks (see Section 4 and [34]); b) we provide an enhanced loss function (naturally absent from [5]), enriched counter-example generation, prioritised check of the verification constraints, and c) we newly synthesise verified barrier certificates for hybrid models, which are

generated using counterexample-based, neural architectures. Finally, beyond [5] the new approach is endowed with numerical robustness features.

SOS programming solutions [14,16,29] are not quite comparable to this work. Foremost, they are not sound, i.e. do not offer a formal guarantee of numerical and algorithmic correctness. The exception is [10], which verifies SOS candidates a-posteriori by means of CAD [15] techniques that are known not to scale well. Furthermore, they can be hardly embedded within a CEGIS loop - we experimentally show that SOS candidates are handled with difficulty by SMT solvers. Finally, they hardly cope with the experiments we have considered, as already observed in [34]. We instead use SMT solvers (Z3 [11] and dReal [13]) within CEGIS to provide sound outcomes based on NN candidates, proffering a new approach that synthesises and formally verifies candidate BCs altogether, with minimum effort from the user.

Organisation The remainder of the paper is organised as follows: Section 2 presents preliminary notions on BCs and outlines the problem. Section 3 describes the approach: training of the NN in Sec. 3.1 and verification in Sec. 3.2. Section 4 presents case studies, Section 5 delineates future work.

2 Safety Analysis with Barrier Certificates

We address the safety verification of continuous-time dynamical models by designing barrier certificates (BC) over the continuous state space X of the model. We consider n -dimensional dynamical models described by

$$\dot{x}(t) = \frac{dx}{dt} = f(x), \quad x(0) = x_0 \in X_0 \subset X, \quad (1)$$

where $f : X \rightarrow \mathbb{R}^n$ is a continuous vector field, $X \subseteq \mathbb{R}^n$ is an open set defining the state space of the system, and X_0 represents the set of initial states. Given model (1) and an unsafe set $X_u \subset X$, the safety verification problem concerns checking whether or not all trajectories of the model originating from X_0 reach the unsafe region X_u . BC offer a sufficient condition asserting the safety of the model, namely when no trajectory enters the unsafe region.

Definition 1. *The Lie derivative of a continuously differentiable scalar function $B : X \rightarrow \mathbb{R}$, with respect to a vector field f , is defined as follows*

$$\dot{B}(x) = \nabla B(x) \cdot f(x) = \sum_{i=1}^n \frac{\partial B}{\partial x_i} \frac{dx_i}{dt} = \sum_{i=1}^n \frac{\partial B}{\partial x_i} f_i(x). \quad (2)$$

Intuitively, this derivative denotes the rate of change of function B along the model trajectories.

Proposition 1 (Barrier Certificate for Safety Verification, [21]). *Let the model in (1) and the sets X , X_0 and X_u be given. Suppose there exists a function*

$B : X \rightarrow \mathbb{R}$ that is differentiable with respect to its argument and satisfies the following conditions:

$$B(x) \leq 0 \ \forall x \in X_0, \quad B(x) > 0 \ \forall x \in X_u, \quad \dot{B}(x) < 0 \ \forall x \in X \text{ s.t. } B(x) = 0, \quad (3)$$

then the safety of the model is guaranteed. That is, there is no trajectory of the model contained in X , starting from the initial set X_0 , that ever enters set X_u .

Consider a trajectory $x(t)$ starting in $x_0 \in X_0$ and the evolution of $B(x(t))$ along this trajectory. Whilst the first of the three conditions guarantees that $B(x_0) \leq 0$, the last condition asserts that the value of $B(x(t))$ along a trajectory $x(t)$ must decrease. Hence such a trajectory $x(t)$ cannot enter the set X_u , where $B(x) > 0$ (second condition), thus ensuring the safety of the model.

3 Synthesis of Neural Barrier Certificates via Learning and Verification

We introduce an automated and formal approach for the construction of barrier certificates (BC) that are expressed as feed-forward neural networks (NN). The procedure leverages CEGIS (see Fig. 1) [31], an automated and sound procedure for solving second-order logic synthesis problems, which comprises two interacting parts. The first component is a *Learner*, which provides candidate BC functions by training a NN over a finite set of sample inputs. The network is then translated into a logical formula in an appropriate theory, by evaluating it with symbolic inputs, instead of canonical floating point numbers. The details of this conversion are outlined in [2]. This encoded candidate is passed to the second component, a *Verifier*, which acts as an oracle: either it proves that the solution is valid, or it finds one (or more) instance (called a counter-example) where the candidate BC does not comply with required conditions. The verifier consists of an SMT solver [15], namely an algorithmic decision procedure that extends Boolean SAT problems to richer, more expressive theories, such as non-linear arithmetics.

More precisely, the learner trains a NN composed of n input neurons (this matches the dimension of the model f), k hidden layers, and one output neuron (recall that $B(x)$ is a scalar function): this NN candidate B is required to closely match the conditions in Eq. (3) over a discrete set of samples S , which is initialised randomly. The verifier checks whether the candidate B violates any of the conditions in Eq. (3) over the entire set X and, if so, produces one (or more, as in this work) counter-examples c . We add c to the samples set S as the loop restarts, hence forcing the NN to be trained *also* over the generated counter-examples c . Note that the NN retains its old weights, and restarts the training from the weights obtained at the end of the previous session. This loop repeats until the SMT verifier proves that no counter-examples exist or until a timeout is reached. CEGIS offers a scalable and flexible alternative for BC synthesis: on the one hand, the learner does not require soundness, and ensures a rapid synthesis exploiting the training of NN architectures; on the other, the algorithm is *sound*, i.e. a valid output from the SMT-based verifier is provably correct; of course we

cannot claim any *completeness*, since CEGIS might in general not terminate with a solution because it operates over a continuous model.

The performance of the CEGIS algorithm in practice hinges on the effective exchange of information between the learner and the verifier [3]. A core contribution of this work is to tailor the CEGIS architecture to the problem of BC synthesis: we devise several improvements to NN training, such as a bespoke loss function and a multi-layer NN architecture that ensures robustness and outputs a function that is tailored to the verification engine. Over consecutive loops, the verifier may return similar counter-examples: we thus propose a more informative counter-examples generation by the SMT verifier that is adapted to the candidate BC and the underlying dynamical model. These tailored architectural details generate in practice a rapid, efficient, and robust CEGIS loop, which is shown in this work to clearly outperform state-of-the-art methods.

3.1 Training of the Barrier Neural Network

The learner instantiates the candidate BC using the hyper-parameters k and h (depth and width of the NN), trains it over the N samples in the set S , and later refines its training whenever the verifier adds counter-examples to the set S . The class of candidate BC comprises multi-layered, feed-forward NN with *polynomial* and non-polynomial activation functions. Unlike most learning applications, the choice of polynomial activations comes from the need for interpretable outputs from the NN, whose analytical expression must be readily processed by the verifier. The order γ of the polynomial activations is a hyper-parameter fed at the start of the procedure: we split the i -th hidden layer into γ portions and apply polynomial activations of order j to the neurons of the j -th portion.

Example 1 (Polynomial Activations). Assume a NN composed of an input x , 3 hidden neurons and 1 activation-free output, with γ -th order polynomial activation, $\gamma = 3$. We split the hidden layer in γ sub-vectors, each containing one neuron. The hidden layer after the activation results in

$$z = \begin{bmatrix} W_1^{(1)}x + b_1 & (W_2^{(1)}x + b_2)^2 & (W_3^{(1)}x + b_3)^3 \end{bmatrix}^T,$$

where the $W_i^{(1)}$ are the i -th row of the first-layer weight matrix, and the b_i form the bias vector. $\square$

The learning process updates the NN parameters to improve the satisfaction of the BC conditions in (3): $B(x) \leq 0$ for $x \in X_0$, $B(x) > 0$ for $x \in X_u$, and a negative Lie derivative $\dot{B}$ (Eq. (2)) over the set implicitly defined by $B(x) = 0$. The training minimises a loss comprising three terms, namely

$$L = L_0 + L_u + L_d = \frac{1}{N} \sum_{i=1}^N \left(\max_{s_i \in X_0} \{\tau_0, B(s_i)\} + \max_{s_i \in X_u} \{\tau_u, -B(s_i)\} \right. \\ \left. + \max_{s_i: B(s_i)=0} \{\tau_d, \dot{B}(s_i)\} \right), \quad (4)$$

where s_i , $i = 1, \dots, N$ are the samples taken from the set S . The constants τ_0 , τ_u , τ_d are offsets, added to improve the numerical stability of the training. Notably, $B(x) = 0$ can be a set with small volume, thus it is highly unlikely that a single sample s will satisfy $B(s) = 0$. We thus relax this last condition and consider a belt $\mathcal{B}$ around $B(s) = 0$, namely $\mathcal{B} = |B(x)| \leq \beta$, which depends on the hyper-parameter β . Note that we must use continuously differentiable activations throughout, as we require the existence of Lie derivatives (cf. Eq. (2)), and thus cannot leverage simple ReLUs.

Enhanced Loss Functions The loss function in Eq. (4) experimentally yields possible drawbacks, which suggests a few ameliorations. Terms L_0 and L_u solely penalise samples with incorrect value of $B(x)$ without further providing a reward for samples with a correct value. The NN thus stops learning when the samples return correct values of $B(x)$ without further increasing the positivity of B over X_u or the negativity over X_0 . As such, the training often returns a candidate $B(x)$ with values just below τ_0 in X_0 or above τ_u in X_u . These candidates are easily falsified, thus potentially leading to a large number of CEGIS iterations.

We improve the learning by adopting a (saturated) *Leaky* ReLU, hence rewarding samples that evaluate to a correct value of $B(x)$. Noting that

$$\text{LeakyReLU}(\alpha, x) = \text{ReLU}(x) - \alpha \text{ReLU}(-x), \quad (5)$$

where α is a small positive constant, we rewrite term L_0 as

$$L_0 = \frac{1}{N} \sum_{s_i \in X_0} \text{ReLU}(B(s_i) - \tau_0) - \alpha \cdot \text{satReLU}(-B(s_i) + \tau_0), \quad (6)$$

where satReLU is the saturated ReLU function³. The term L_u is similarly modified. The composite loss function works as follows. Incorrect samples account for the main contribution to the loss function, leading the NN to correct those first via the ReLU term in Eq. (6). At a second stage, the network finds a direction of improvement by following the *leaky* portion of the loss function. This is saturated to prevent the training from following only one of these directions, without improving the other loss terms.

Another possible drawback of the loss function in (4) derives from the term L_d : it solely accounts for a penalisation of the sample points within $\mathcal{B}$. To quickly and myopically improve the loss function, the training can generate a candidate BC for which no samples are within $\mathcal{B}$ - we experimentally find that this behaviour persists, regardless of the value of β . Similarly to L_0 and L_u , we reward the points within a belt fulfilling the BC condition: namely, we solely apply the satReLU function to reward samples s with a negative $\dot{B}(s)$, whilst not penalising values $\dot{B}(s) \geq 0$. The training is driven to include more samples in $\mathcal{B}$, guiding towards a negative $\dot{B}(s)$, and finally enhancing learning. The expression of L_d results in

$$L_d = -\frac{1}{N} \sum_{s \in \mathcal{B}} \text{satReLU}(-\dot{B}(s) + \tau_d). \quad (7)$$

³ Let us define M to be an arbitrary upper bound, then $\text{satReLU}(x) = \min(\max(0, x), M)$.

Finally, we choose an asymmetric belt $\mathcal{B} = -\beta_1 \leq B(s) \leq \beta_2$, with $\beta_2 > \beta_1 > 0$ to both ensure a wider sample set and a stronger safety certificate.

Multi-layer Networks Polynomial activation functions generate interpretable barrier certificates with analytical expressions that are readily verifiable by an SMT solver. However, when considering polynomial networks, the use of multi-layer architectures quickly increases the order of the barrier function: a k -layer network with γ -th order activations returns a polynomial of $k\gamma$ degree. We have experienced that deep NN provide numerical robustness to our method, although the verification complexity increases with the order of the polynomial activation functions used and with the depth of the NN. As a consequence, our procedure leverages a deep architecture whilst maintaining a low-order polynomial by interchanging linear and polynomial activations over adjacent layers. We have observed that the use of linear activations, particularly in the output layer, positively affects the training: they provide robustness that is needed to the synthesis of BC (see Experimental results), without increasing the order of the network with new polynomial terms.

Learning in Separate Batches The structure of the conditions in (3) and the learning loss in (4) naturally suggests a separate approach to training. We then split the dataset S into three batches S_0 , S_u and S_x , each including samples belonging to X_0 , X_u and X , respectively. For training, we compute the loss function in a parallel fashion. Similarly, for the verifier, generated counter-examples are added to the relevant batch.

3.2 Certification of the Barrier Neural Network, or Falsification via Counter-examples

Every candidate BC function $B(x)$ which the learner generates requires to be certified by the verifier. Equivalently, in practice the SMT-based verifier aims at finding states that violate the barrier conditions in (3) over the continuous domain X . To this end, we express the *negation* of such requirements, and formulate a nonlinear constrained problem over real numbers, as

$$(x \in X_0 \wedge B(x) > 0) \vee (x \in X_u \wedge B(x) \leq 0) \vee (B(x) = 0 \wedge \dot{B}(x) \geq 0). \quad (8)$$

The verifier searches for solutions of the constraints in Eq. (8), which in general requires manipulating non-convex functions. This can be cumbersome and time-consuming, hence simple expressions of B can enhance the verification procedure. On the one hand, the soundness of our CEGIS procedure heavily relies on the correctness of SMT solving: an SMT solver never fails to assert the absence of solutions for (8). As a result, when it states that formula (8) is unsatisfiable, i.e. returns `unsat`, $B(x)$ is formally guaranteed to fulfil the BC conditions in Eq. (3). On the other hand, the CEGIS algorithm offers flexibility in the choice of the verifier, hence we implement and discuss two SMT solvers: dReal [13] and Z3 [11]. dReal is a δ -complete solver, namely the `unsat` decision is correct [12],

whereas when a solution for (8) is found, this comes with a δ -error bound. The value of δ characterises the procedure precision. In our setting, it is acceptable to return spurious counter-examples: indeed, these are then used as additional samples and do not invalidate the sound outcomes of the procedure, but rather help synthesising a more robust barrier candidate. dReal is capable of handling non-polynomial terms, such as exponentials or trigonometric vector fields f for some of the models considered in Section 4. Z3 is a powerful, sound and complete SMT solver, namely its conclusions are provably correct both when it determines the validity of a BC candidate and when it provides counter-examples. The shortcoming of Z3 is that it is unable to handle non-polynomial formulae.

Prioritisation and Relaxation of Constraints The effectiveness of the CEGIS framework is underpinned by rapid exchanges between the learner and the verifier, as well as by quick NN training and SMT verification procedures. We have experienced that the bottleneck resides in the handling of the constraint $\eta_d = (B(x) = 0 \wedge \dot{B}(x) \geq 0)$ by the SMT solver, since the formula contains the high-order expression $\dot{B}(x)$ and because it is defined over the thin region of the state space implicitly characterised by $B(x) = 0$. As a consequence, we have prioritised constraints $\eta_0 = (x \in X_0 \wedge B(x) > 0)$ and $\eta_u = (x \in X_u \wedge B(x) \leq 0)$: that is, if either clause is satisfied, i.e. a counter-example is found for at least one of them, the verifier omits testing η_d whilst the obtained counter-examples are passed to the learner. The constraint η_d is thus checked solely if η_0 and η_u are both deemed to be **unsat**. Whenever this occurs, and the verification of η_d times out, the solver searches for a solution of a relaxed constraint $(|B(x)| < \tau_v \wedge \dot{B}(x) \geq 0)$, similarly to the improved learning conditions discussed in Eq. (7). Whilst this constraint is arguably easier to solve in general, it may generate spurious counter-examples, namely a sample $\bar{x}$ that satisfy the relaxed constraint, but such that $B(\bar{x}) \neq 0$. The generation of these samples does not contradict the soundness of the procedure, and indeed improve the robustness of the next candidate BC – this of course comes with the cost of increasing the number of CEGIS iterations.

Increased Information from Counter-examples The verification task encompasses an SMT solver attempting to generate a counter-example, namely a (single) instance satisfying Eq. (8). However, a lone sample might not always provide insightful information for the learner to process. Naïvely asking the SMT solver to generate more than one counter-example can be in general expensive. Specifically, the verifier solves Eq. (8) to find a first counter-example $\bar{x}$; then, to find any additional sample, we include the statement $(x \neq \bar{x})$ and solve again for the resulting formula. We are interested in finding numerous points invalidating the BC conditions and feed them to the learner as a batch, or in increasing the information generated by the verifier by finding a sample that maximises the violation of the BC conditions. To this end, firstly we randomly generate a *cloud* of points around the generated counter-example: in view of the continuity of the candidate function B , samples around a counter-example are also likely to invalidate the BC conditions. Secondly, for the original counter-example, we compute the gradient of B (or of $\dot{B}$) and follow the direction that maximises the

violation of the BC constraints. As such, we follow the B (resp. $\dot{B}$) maximisation when considering $x \in X_0$ (x s.t. $|B(x)| < \tau_v$), and vice versa when $x \in X_u$. This gradient computation is extremely fast as it exploits the neural architecture, and it provides more informative samples for further use by the learner.

Algorithm 1 Synthesis of Neural Barrier Certificate

<pre> function LEARNER(S, f) repeat $B(S) \leftarrow \text{NN}(S)$ $\dot{B}(S) \leftarrow \nabla B(S) \cdot f(S)$ compute loss L, update NN until convergence return NN end function function VERIFIER($B, \dot{B}$) encode conditions in (8) Cex or unsat $\leftarrow \text{SMTcheck}(B, \dot{B})$ return Cex or unsat end function </pre>	<pre> function CEGIS(f) initialise NN, S repeat NN $\leftarrow \text{LEARNER}(S, f)$ $B(x), \dot{B}(x) \leftarrow \text{Translate}(\text{NN}, f)$ Cex or unsat $\leftarrow \text{VERIFIER}(B, \dot{B})$ $S \leftarrow S \cup \text{Cex}$ until unsat return $B(x), \dot{B}(x)$ end function </pre>
--	---

4 Case Studies and Experimental Results

All experiments are performed on a laptop workstation with 8 GB RAM, running on Ubuntu 18.04. We demonstrate that the proposed method finds provably correct BCs on benchmarks from literature comprising both polynomial and non-polynomial dynamics: we compare our approach against the work [34], as this is the only work on sound synthesis of BCs with NNs to the best of our knowledge, and against the SOS optimisation software SOSTOOLS [18]. Beyond the benchmarks proposed in [34], we newly tackle a hybrid model as well as larger, (up to) 8-dimensional models, which push the boundaries of the verification engine and display a significant extension to the state of the art. To confirm the flexibility of our architecture, we employ SMT-solver dReal in the first four benchmarks, whereas we study the last four using Z3. In all the examples, we use a learning rate of 0.1 for the NN and the loss function in Section 3.1 with $\alpha = 10^{-4}$, $\tau_0 = \tau_u = \tau_d = 0.1$. The region in Eq. (7) is limited by $\beta_1 = 0.1$, whilst $\beta_2 = \infty$. Accordingly, the training over a large set $\mathcal{B}$ results in a candidate B with a negative derivative over this large region, which validity is more likely to be certified by the verifier. We set a verification parameter $\tau_v = 0.05$ (cf. Sec. 3.2), a timeout (later denoted as OOT) of 60 seconds and the precision for dReal to $\delta = 10^{-6}$. Table 1 summarises the outcomes. We emphasise that our approach supports any network depth and width. The presented results seek a tradeoff between speed (low order, small networks) and expressiveness (high order, larger networks): a different architecture may result in a slower or faster synthesis.

For the first four benchmarks, we compare our procedure, denoted as CEGIS, with the repeated results from [34], which however does not handle the hybrid model in the fifth benchmark. We have run the algorithm in [34] and reported the cumulative synthesis time under the ‘Learn’ column. However the verification is not included in the repeatability package, hence we report the results from [34], which are generated with much more powerful hardware. Due to this issue of lack of repeatability, we have not run [34] on the larger models. Compared to [34], the outcomes suggest that we obtain *much faster* synthesis and verification times, whilst requiring up to only 0.1% (see Obstacle Avoidance Problem) of the training data: [34] performs a uniform sampling of the space X , hence suffers especially in the 3-D case, where the learning runs *two orders of magnitude* faster. Evidently this gap in performance derives from the different synthesis procedure: it appears to be more advantageous to employ a smaller, randomly sampled initial dataset that is progressively augmented with counter-examples, rather than to uniformly sample the state space to then train the neural network.

Next, we have implemented the SOS optimisation problems in [10] within the software SOSTOOLS [18] to generate barrier candidates, which are polynomials up to order 4 (this is the maximum order of the polynomial candidates generated by our Learner). In a few instances we ought to conservatively approximate the expression of X_0 or X_u in order to encode them as SOS program - this makes their applicability less general. SOSTOOLS has successfully found BC candidates for five of the eight benchmarks, and they were generated consistently fast, in view of the convex structure of the underlying optimisation problem. However, recall that these techniques lack soundness (also due to numerical errors), which is instead a core asset of our approach. Consequently, we have passed them to the Z3 SMT solver, which should easily handle polynomial formulae: only one of them (‘Hybrid Model’) has been successfully verified; instead, the candidate for the ‘Polynomial Model’ has been invalidated (namely Z3 has found a counter-example for it), whereas the verification of the remaining BC candidates has run out of time. For the latter instances, we have experienced that SOSTOOLS generally returns numerically ill-conditioned expressions, namely candidates with coefficients of rather different magnitude, with many decimal digits: even after rounding, expressions with this structure are known to be hardly handled by SMT solvers [2,5], which results in long time needed to return an answer - this explains the experienced timeouts. These experiments suggest that the use of SOS programs within a CEGIS loop appears hardly attainable.

Notice that all the case studies are solved with a *small number* of iterations (up to 9) of the CEGIS loop: this feature, along with the limited runtimes, is promising towards tackling synthesis problems over larger models.

For the eight case studies, we report below the full expressions of the dynamics of the models, the spatial domain X (as a set of constraints), the set of initial conditions $X_0 \subset X$, and the unsafe set $X_u \subset X$. We add a detailed analysis of the CEGIS iterations involved in the synthesis of the corresponding BCs.

Benchmark	CEGIS (this work)				BC from [34]			SOS from [18]	
	Learn	Verify	Samples	Iters	Learn	Verify	Samples	Synth	Verify
Darboux	31.6	0.01	0.5 k	2	54.9	20.8	65 k	×	—
Exponential	15.9	0.07	1.5 k	2	234.0	11.3	65 k	×	—
Obstacle	55.5	1.83	2.0 k	9	3165.3	1003.3	2097 k	×	—
Polynomial	64.5	4.20	2.3 k	2	1731.0	635.3	65 k	8.10	×
Hybrid	0.58	2.01	0.5 k	1	—	—	—	12.30	0.11
4-d ODE	29.31	0.07	1 k	1	—	—	—	12.90	OOT
6-d ODE	89.52	1.61	1 k	3	—	—	—	16.60	OOT
8-d ODE	104.5	82.51	1 k	3	—	—	—	26.10	OOT

Table 1. Outcomes of the case studies: Cumulative time for Learning and Verification steps are given in seconds; ‘Samples’ indicates the size of input data for the Learner (in thousands); ‘Iters’ is the number of iterations of the CEGIS loop (which is specific to our work); × indicates a synthesis or verification failure; OOT denotes a verification timeout. The Hybrid and the three ODE Models are newly introduced in this work.

Darboux Model This 2-dimensional model is approached using polynomial BCs. Its analytical expression is

$$\begin{cases} \dot{x} = y + 2xy, \\ \dot{y} = -x + 2x^2 - y^2, \end{cases} \quad \text{with domains} \quad \begin{aligned} X &= \{-2 \leq x, y \leq 2\}, \\ X_0 &= \{0 \leq x \leq 1, 1 \leq y \leq 2\}, \\ X_u &= \{x + y^2 \leq 0\}. \end{aligned}$$

The work [33] reports that methods based on linear matrix inequalities fail to verify this model using polynomial templates of degree 6. Our approach generates the BC shown in Fig. 2 (left) in approximately 30 seconds, roughly half as much as in [34], and using only 500 initial samples vs more than 65000. The initial and unsafe sets are depicted in green and red, respectively, whereas the level set $B(x) = 0$ is outlined in black. The BC is derived from a single-layer architecture of 10 nodes, with linear activations.

Exponential Model This model from [17] shows that our approach extends to non-polynomial systems encompassing exponential and trigonometric functions:

$$\begin{cases} \dot{x} = e^{-x} + y - 1, \\ \dot{y} = -\sin^2 x, \end{cases} \quad \text{with domains} \quad \begin{aligned} X &= \{-2 \leq x, y \leq 2\}, \\ X_0 &= \{(x + 0.5)^2 + (y - 0.5)^2 \leq 0.16\}, \\ X_u &= \{(x - 0.7)^2 + (y + 0.7)^2 \leq 0.09\}. \end{aligned}$$

Our algorithm provides a valid BC in 16 seconds, around 7% of the results in [34], again using solely 1500 initial samples. The BC, depicted in Fig. 2 (centre), results from a single-layer neural architecture of 10 nodes, with polynomial ($\gamma = 3$) activation function.

Obstacle Avoidance Problem This 3-dimensional model, originally presented in [6], describes a robotic application: the control of the angular velocity of a

Fig. 2. The BC for the Darboux (top left), Exponential (middle left), and Obstacle Avoidance (the 3-D study, bottom left) models with corresponding vector fields (right column). Initial and unsafe sets are represented in green and red, respectively; the black line outlines the level curve $B(x) = 0$.

two-dimensional airplane, aimed at avoiding a still obstacle. The details are

$$\begin{cases} \dot{x} = v \sin \varphi, \\ \dot{y} = v \cos \varphi, \\ \dot{\varphi} = u, \end{cases} \quad \text{where} \quad u = -\sin \varphi + 3 \cdot \frac{x \sin \varphi + y \cos \varphi}{0.5 + x^2 + y^2}, \quad \text{with domains}$$

$$\begin{aligned}
X &= \{-2 \leq x, y \leq 2, -\pi/2 < \varphi < \pi/2\}, \\
X_0 &= \{-0.1 \leq x \leq 0.1, -2 \leq y \leq -1.8, -\pi/6 < \varphi < \pi/6\}, \\
X_u &= \{x^2 + y^2 \leq 0.04\}.
\end{aligned}$$

The BC is obtained from a single-layer NN comprising 10 neurons, using ($\gamma = 3$) polynomial activations. Fig. 2 (right) plots the vector field on the plane $z = 0$. Our procedure takes 1% of the computational time in [34], providing a valid BC with 9 iteration starting from an initial dataset of 2000 samples.

Polynomial Model This model describes a polynomial system [22] and presents initial and unsafe sets with complex, non convex shapes [34], as follows:

$$\begin{aligned}
&\begin{cases} \dot{x} = y, \\ \dot{y} = -x + 1/3 x^3 - y, \end{cases} \text{ with domains} \\
X &= \{-3.5 \leq x \leq 2, -2 \leq y \leq 1\}, \\
X_0 &= \{(x - 1.5)^2 + y^2 \leq 0.25 \vee (x \geq -1.8 \wedge x \leq -1.2 \wedge y \geq -0.1 \wedge y \leq 0.1) \\
&\quad \vee (x \geq -1.4 \wedge x \leq -1.2 \wedge y \geq -0.5 \wedge y \leq 0.1)\}, \\
X_u &= \{(x + 1)^2 + (y + 1)^2 \leq 0.16 \vee (x \geq 0.4 \wedge x \leq 0.6 \wedge y \geq 0.1 \wedge y \leq 0.5) \\
&\quad \vee (x \geq 0.4 \wedge x \leq 0.8 \wedge y \geq 0.1 \wedge y \leq 0.3)\}.
\end{aligned}$$

SOS-based procedures [16,29], have required high-order polynomial templates, which has suggested the use of alternative activation functions. The BC, shown in Fig. 3, is generated using a 10-neuron, two-layer NN with polynomial ($\gamma = 3$) and tanh activations. Needing just around 1 min and only 2300 initial samples, the overall procedure is 30 times faster than that in [34].

Hybrid Model We challenge our procedure with a 2-dimensional hybrid model, which extends beyond the capability of the results in [34]. This hybrid framework partitions the set X into two non-overlapping subsets, X_1 and X_2 . Each subset is associated to different model dynamics, respectively f_1 and f_2 . In other words, the model trajectories evolve according to the f_1 dynamics when in X_1 , and according to f_2 when in X_2 .

$$f_1 = \begin{cases} \dot{x} = y, \\ \dot{y} = -x - 0.5x^3, \end{cases} \quad f_2 = \begin{cases} \dot{x} = y, \\ \dot{y} = x - 0.25y^2, \end{cases}$$

with domain for $f_1 = \{(x, y) : x < 0\}$, domain for $f_2 = \{(x, y) : x \geq 0\}$, and sets

$$\begin{aligned}
X &= \{x^2 + y^2 \leq 4\}, & X_0 &= \{(x + 1)^2 + (y + 1)^2 \leq 0.25\}, \\
X_u &= \{(x - 1)^2 + (y - 1)^2 \leq 0.25\}.
\end{aligned}$$

The structure of this model represents a non-trivial task for the verification engine, for which we employ the Z3 SMT solver. The learning phase has instead been quite fast. The BC (Fig.3) is obtained from a single-layer NN comprising 3 neurons, using polynomial activations with $\gamma = 2$, overall in less than 3 seconds, starting with an initial dataset of 500 samples.

Fig. 3. The BC for the polynomial model (top left) and the hybrid model (top right) with the respective vector field (below).

Larger-dimensional Models We finally challenge our procedure with three high-order ODEs, respectively of order four, six and eight, to display the general applicability of our counter-example guided BC synthesis. We consider dynamical models described by the following differential equations:

$$x^{(4)} + 3980x^{(3)} + 4180x^{(2)} + 2400x^{(1)} + 576 = 0, \quad (9)$$

$$x^{(6)} + 800x^{(5)} + 2273x^{(4)} + 3980x^{(3)} + 4180x^{(2)} + 2400x^{(1)} + 576 = 0, \quad (10)$$

$$x^{(8)} + 20x^{(7)} + 170x^{(6)} + 800x^{(5)} + 2273x^{(4)} + 3980x^{(3)} + 4180x^{(2)} + 2400x^{(1)} + 576 = 0, \quad (11)$$

where we denote the i -th derivative of variable x by $x^{(i)}$. We translate the ODE into a state-space model with variables $x_1, \dots, x_j$, where $j = \{4, 6, 8\}$, respectively. In all three instances, we select as spatial domain X an hyper-sphere centred at the origin of radius 4; an initial set X_0 as hyper-sphere⁴ centred at $+1^{[j]}$ of radius 0.25; an unsafe set X_u as an hyper-sphere centred at $-2^{[j]}$ of radius 0.16. For the synthesis, we employ for all case studies a single-layer, 5-node architecture with polynomial ($\gamma = 1$) activation function. Whilst in particular the verification engine is challenged from the high dimensionality of the models, the CEGIS procedure returns a valid barrier certificate in up to 3 iterations and with very reasonable run times.

⁴ We denote $1^{[j]}$ the point of a j -dimensional state-space that has all its components equal to 1. For instance, $1^{[3]}$ is the 3-dimensional point $[1, 1, 1]$. Similarly for $2^{[j]}$.

Codebase Robustness The results in Table 1 are obtained setting the NN initialisation seed manually for repeatability. We now test the robustness of the overall algorithm by randomising the initialisation seed. We report in Table 2 the percentage of successful runs, the average time and iterations count, along with minimum and maximum values, over 50 runs. We set timeouts as a max running time of 10 minutes, or as 12 CEGIS loops. Notice that small architectures are highly susceptible to initialisations, which renders this test rather challenging. Compared to Table 1, we notice similar performances for the *Darboux*, *Exponential* and *Hybrid* models, vouching for the robustness of our approach. However, the performance decreases when tackling the most challenging models. Still, we highlight that the procedure can synthesise a valid BC very rapidly for every benchmark (notice the lower bounds of the computational times). This outcome suggests that a parallel approach - i.e. the procedure running on several networks simultaneously - may be suited to quickly synthesise candidates. Overall, the table shows a high degree of variance, possibly indicating the need for larger architectures to enhance robustness.

Benchmark	Success [%]	Iters	Avg Time
Darboux Model	84.0	4.76 [1, 12]	75.33 [15.00, 189.25]
Exponential Model	76.0	5.20 [1, 12]	9.50 [3.17, 21.59]
Obstacle Avoidance	28.0	9.88 [1, 11]	129.24 [16.17, 549.41]
Polynomial Model	8.0	5.56 [5, 9]	335.32 [230.86, 377.91]
Hybrid Model	84.0	4.20 [1, 12]	36.75 [0.43, 102.03]
4-d ODE Model	32.0	9.00 [1, 12]	362.14 [29.42, 681.41]
6-d ODE Model	40.0	8.60 [1, 12]	310.45 [30.65, 562.67]
8-d ODE Model	12.0	11.00 [2, 12]	495.23 [111.50, 698.93]

Table 2. Percentage of successful runs, average number of iterations and average computational times (in seconds) of the CEGIS procedure, over 50 runs. The square brackets contain the minimum and maximum values obtained.

5 Conclusions and Future Work

We have presented a new inductive, formal, automated technique to synthesise neural-based barrier certificates for polynomial and non-polynomial, continuous and hybrid dynamical models. Thanks to a number of architectural choices for the new procedure, our method requires less training data and thus displays faster learning, as well as quicker verification time, than state-of-the-art techniques.

Ongoing work is porting presented and related [5,2] theoretical results into a software tool [1]. Towards increased automation, future work includes the development of an automated selection of activation functions that are tailored to the dynamical models of interest.

References

1. Alessandro Abate, Daniele Ahmed, Alec Edwards, Mirco Giacobbe, and Andrea Peruffo. FOSSIL: A Software Tool for the Formal Synthesis of Lyapunov Functions and Barrier Certificates using Neural Networks. In *HSCC*. ACM, 2021.
2. Alessandro Abate, Daniele Ahmed, Mirco Giacobbe, and Andrea Peruffo. Formal Synthesis of Lyapunov Neural Networks. *IEEE Control Systems Letters*, 5(3):773–778, 2021.
3. Alessandro Abate, Cristina David, Pascal Kesseli, Daniel Kroening, and Elizabeth Polgreen. Counterexample Guided Inductive Synthesis Modulo Theories. In *Proceedings of CAV, LNCS 10981*, pages 270–288, 2018.
4. Alessandro Abate, Ashish Tiwari, and Shankar Sastry. Box Invariance in Biologically-inspired Dynamical Systems. *Automatica*, 45(7):1601–1610, 2009.
5. Daniele Ahmed, Andrea Peruffo, and Alessandro Abate. Automated and Sound Synthesis of Lyapunov Functions with SMT Solvers. In *TACAS (1)*, volume 12078 of *LNCS*, pages 97–114. Springer, 2020.
6. Andrew J Barry, Anirudha Majumdar, and Russ Tedrake. Safety Verification of Reactive Controllers for UAV Flight in Cluttered Environments using Barrier Certificates. In *2012 IEEE International Conference on Robotics and Automation*, pages 484–490. IEEE, 2012.
7. Urs Borrmann, Li Wang, Aaron D Ames, and Magnus Egerstedt. Control Barrier Certificates for Safe Swarm Behavior. *IFAC-PapersOnLine*, 48(27):68–73, 2015.
8. Dario Cattaruzza, Alessandro Abate, Peter Schrammel, and Daniel Kroening. Unbounded-Time Safety Verification of Guarded LTI Models with Inputs by Abstract Acceleration. *Journal of Automated Reasoning*, 2020.
9. Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov Control. In *NeurIPS*, pages 3240–3249, 2019.
10. Liyun Dai, Ting Gan, Bican Xia, and Naijun Zhan. Barrier Certificates Revisited. *Journal of Symbolic Computation*, 80:62–86, 2017.
11. Leonardo de Moura and Nikolaj Bjørner. Z3: An Efficient SMT Solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.
12. Sicun Gao, Jeremy Avigad, and Edmund M Clarke. δ -complete Decision Procedures for Satisfiability over the Reals. In *International Joint Conference on Automated Reasoning*, pages 286–300. Springer, 2012.
13. Sicun Gao, Soonho Kong, and Edmund M Clarke. dReal: An SMT Solver for Nonlinear Theories over the Reals. In *International conference on automated deduction*, pages 208–214. Springer, 2013.
14. Hui Kong, Fei He, Xiaoyu Song, William NN Hung, and Ming Gu. Exponential-condition-based Barrier Certificate Generation for Safety Verification of Hybrid Systems. In *International Conference on Computer Aided Verification*, pages 242–257. Springer, 2013.
15. Daniel Kroening and Ofer Strichman. *Decision Procedures - An Algorithmic Point of View*. Springer Verlag, 2016.
16. Benoît Legat, Paulo Tabuada, and Raphaël M Jungers. Sum-of-Squares Methods for Controlled Invariant Sets with Applications to Model-predictive Control. *Nonlinear Analysis: Hybrid Systems*, 36:100858, 2020.
17. Jiang Liu, Naijun Zhan, Hengjun Zhao, and Liang Zou. Abstraction of Elementary Hybrid Systems by Variable Transformation. In *International Symposium on Formal Methods*, pages 360–377. Springer, 2015.

18. A. Papachristodoulou, J. Anderson, G. Valmorbida, S. Prajna, P. Seiler, and P. A. Parrilo. *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*. <http://arxiv.org/abs/1310.4716>, 2013.
19. André Platzer and Edmund M Clarke. Computing Differential Invariants of Hybrid Systems as Fixedpoints. In *International Conference on Computer Aided Verification*, pages 176–189. Springer, 2008.
20. Stephen Prajna. Barrier Certificates for Nonlinear Model Validation. *Automatica*, 42(1):117–126, 2006.
21. Stephen Prajna and Ali Jadbabaie. Safety Verification of Hybrid Systems Using Barrier Certificates. In *International Workshop on Hybrid Systems: Computation and Control*, pages 477–492. Springer, 2004.
22. Stephen Prajna, Ali Jadbabaie, and George J Pappas. A Framework for Worst-case and Stochastic Safety Verification Using Barrier Certificates. *IEEE Transactions on Automatic Control*, 52(8):1415–1428, 2007.
23. Hadi Ravanbakhsh and Sriram Sankaranarayanan. Counter-example guided synthesis of control lyapunov functions for switched systems. In *IEEE Control and Decision Conference (CDC)*, pages 4232–4239, 2015.
24. Hadi Ravanbakhsh and Sriram Sankaranarayanan. Robust Controller Synthesis of Switched Systems Using Counterexample Guided Framework. In *ACM/IEEE Conference on Embedded Software (EMSOFT)*, pages 8:1–8:10, 2016.
25. Hadi Ravanbakhsh and Sriram Sankaranarayanan. Learning Control Lyapunov Functions from Counterexamples and Demonstrations. *Autonomous Robots*, pages 1–33, 2018.
26. Spencer M. Richards, Felix Berkenkamp, and Andreas Krause. The Lyapunov Neural Network: Adaptive Stability Certification for Safe Learning of Dynamical Systems. In *CoRL*, volume 87 of *Proceedings of Machine Learning Research*, pages 466–476. PMLR, 2018.
27. Sriram Sankaranarayanan, Xin Chen, and Erika Abraham. Lyapunov Function Synthesis using Handelman Representations. *IFAC Proceedings Volumes*, 46(23):576–581, 2013.
28. Shankar Sastry. *Nonlinear Systems: Analysis, Stability and Control*. Springer Verlag, 1999.
29. Christoffer Sloth, George J Pappas, and Rafael Wisniewski. Compositional Safety Analysis using Barrier Certificates. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control*, pages 15–24, 2012.
30. Andrew Sogokon, Khalil Ghorbal, Yong Kiam Tan, and André Platzer. Vector Barrier Certificates and Comparison Systems. In *International Symposium on Formal Methods*, pages 418–437. Springer, 2018.
31. Armando Solar-Lezama, Liviu Tancau, Rastislav Bodik, Sanjit Seshia, and Vijay Saraswat. Combinatorial sketching for finite programs. In *Proceedings of the 12th international conference on Architectural support for programming languages and operating systems*, pages 404–415, 2006.
32. Li Wang, Aaron D Ames, and Magnus Egerstedt. Safety Barrier Certificates for Collisions-free Multirobot Systems. *IEEE Transactions on Robotics*, 33(3):661–674, 2017.
33. Xia Zeng, Wang Lin, Zhengfeng Yang, Xin Chen, and Lilei Wang. Darboux-type Barrier Certificates for Safety Verification of Nonlinear Hybrid Systems. In *Proceedings of the 13th International Conference on Embedded Software*, pages 1–10, 2016.

34. Hengjun Zhao, Xia Zeng, Taolue Chen, and Zhiming Liu. Synthesizing Barrier Certificates Using Neural Networks. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, HSCC '20, New York, NY, USA, 2020. Association for Computing Machinery.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.