

Automated and Sound Synthesis of Lyapunov Functions with SMT Solvers

Daniele Ahmed^{1,2}, Andrea Peruffo¹, and Alessandro Abate¹

¹ Department of Computer Science, University of Oxford, OX1 3QD Oxford, UK.

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

² Amazon Inc, London, UK.

Abstract. In this paper we employ SMT solvers to soundly synthesise Lyapunov functions that assert the stability of a given dynamical model. The search for a Lyapunov function is framed as the satisfiability of a second-order logic formula, asking whether there exists a function satisfying a desired specification (stability) for all possible initial conditions of the model. We synthesise Lyapunov functions for linear, non-linear (polynomial), and for parametric models. For non-linear models, the algorithm also determines a region of validity for the Lyapunov function. We exploit an inductive framework to synthesise Lyapunov functions, starting from parametric templates. The inductive framework comprises two elements: a “learner” proposes a Lyapunov function, and a “verifier” checks its validity - its lack is expressed via a counterexample (in practice, a point over the state space), of further use by the learner. Whilst the verifier uses the SMT solver Z3, thus ensuring the overall soundness of the procedure, we examine two alternatives for the learner: a numerical approach based on the optimisation tool Gurobi, and a sound one based again on Z3. The overall technique is evaluated over a broad set of benchmarks, which show that this methodology not only scales to 10-dimensional models within reasonable computational time, but also offers a novel soundness proof for the generated Lyapunov functions and their domains of validity.

Keywords: Lyapunov functions, automated synthesis, inductive synthesis, counter-example guided synthesis

1 Introduction

Dynamical systems represent a major modelling framework in both theoretical and applied sciences: they describe how objects move by means of the laws governing their dynamics in time. Often they encompass a system of ordinary differential equations (ODE) with nontrivial solutions.

This work aims at studying the stability property of general ODEs, without knowledge of their analytical solution. Stability analysis via Lyapunov functions is a known approach to assert such property. As such, the problem of constructing relevant Lyapunov functions for stability analysis has drawn much attention in

39 the literature [1,2]. By and large, existing approaches leverage Linear Algebra or
40 Convex Optimisation solutions, and are not automated nor numerically sound.

41 **Contributions** We apply an inductive synthesis framework, known as Counter-
42 Example Guided Inductive Synthesis (CEGIS) [3,4], to construct Lyapunov func-
43 tions for linear, polynomial and parametric ODEs, and to constructively charac-
44 terise their domain of validity. CEGIS, originally developed for program synthesis
45 based on the satisfiability of second-order logical formulae, is employed in this
46 work with templates Lyapunov functions and in conjunction with a Satisfiability
47 Modulo Theory (SMT) solver [5]. Our results offer a formal guarantee of
48 correctness in combination with a simple algorithmic implementation.

49 The synthesis of a Lyapunov function V can be written as a second-order logic
50 formula $F := \exists V \forall x : \psi$, where x represents the state variables and ψ represents
51 requirements that V needs to satisfy in order to be a Lyapunov function.

52 The CEGIS architecture is structured as a loop between two components, a
53 “learner” and a “verifier”. The learner provides a candidate function V and the
54 verifier checks the validity of ψ over the set of x ; if the function is not valid, the
55 verifier provides a counterexample, namely a point \bar{x} in the state space where the
56 candidate function does not satisfy ψ . The learner incorporates the generated
57 counterexample \bar{x} , subsequently computes a new candidate function, and loops
58 it back to the verifier.

59 We exploit SMT solvers to (repeatedly) assert the validity of ψ , given V , over
60 a domain in the space of x . Satisfiability Modulo Theory (SMT) is a powerful
61 tool to assert the existence of such a function. An SMT problem is a decision
62 problem – a problem that can be formulated as a yes/no question – for logical
63 formulae within one or more theories, e.g. the theory of arithmetics over real
64 numbers. The generation of simple counterexamples \bar{x} is a key new feature of
65 our technique.

66 Furthermore, in this work we provide two alternative CEGIS implementa-
67 tions: 1) a numerical learner and an SMT-based verifier, 2) an SMT-based learner
68 and verifier. The numerical generation of Lyapunov functions is based on the op-
69 timisation tool Gurobi [6], whereas the SMT-based leverages Z3 [7].

70 **Related Work** The construction of Lyapunov functions is recognisably an
71 important yet hard problem, particularly for non-linear models, and has been
72 the objective of classical studies [8,9,10]. A know constructive result has been
73 introduced in [11], which additionally provides an estimate of the domain of at-
74 traction. It has led to further work based on recursive procedures. Broadly, these
75 approaches are numerical and based on the solution of optimisation problems.
76 For instance, linear programming is exploited in [12] to iteratively search for
77 stable matrices inside a predefined convex set, resulting in an approximate Lya-
78 punov function for the given model. Alternative approximate methods include [1]
79 ε -bounded numerical methods, techniques leveraging series expansion of a func-
80 tion, the construction of functions from trajectory samples, and the framework

81 of linear matrix inequalities. The approach in [13] uses sum-of-squares (SOS)
82 polynomials to synthesise Lyapunov functions, however its scalability remains
83 an issue. The work in [14] uses SOS decomposition to synthesise Lyapunov func-
84 tions for (non-polynomial) non-linear systems: the algorithmic implementation
85 is known as SOSTOOLS [15,16]. [17] focuses on an analytical result involving a
86 summation over finite time interval, under a stability assumption. Recent devel-
87 opments are in [18] and subsequent work. Surveys on this topic are in [1,2].

88 In conclusion, existing constructive approaches either rely on complex candi-
89 date functions (whether rational or polynomial), on semi-analytical results, or
90 alternatively they involve state-space partitions (for which scalability with the
91 state-space dimension is problematic) accompanied by correspondingly complex
92 or large optimisation problems. These approximate methods evidently lack either
93 numerical robustness, being bound by machine precision, or algorithmic sound-
94 ness: they cannot provide formal certificates of reliability which, in safety-critical
95 applications, can be an evident limit.

96 In [19] Lyapunov functions are soundly found within a parametric frame-
97 work, by constructing a system of linear inequality constraints over unknown
98 coefficients. A twofold linear programming relaxation is made: it includes in-
99 terval evaluation of the polynomial form and “Handelman representations” for
100 positive polynomials. Simulations are used in [20] to generate constraints for a
101 template Lyapunov function, which are then resolved via LP, resulting in candi-
102 date solutions. Whilst the authors refer to traces as counterexamples, they do
103 not employ the CEGIS framework, as in this work. When no counterexamples are
104 found, [20] further uses dReal [21] and Mathematica [22] to verify the obtained
105 candidate Lyapunov functions. The sound technique, which is not complete, is
106 tested on low-dimensional models with non-linear dynamics.

107 The cognate work in [23,24,25] is the first to employ a CEGIS-based ap-
108 proach to synthesise Lyapunov functions. [23,24] focuses on such synthesis for
109 switching control models - a more general setup than ours. [23] employs an SMT
110 solver for the learner, and towards scalability solves an optimisation problem
111 over LMI constraints for the verifier over a given domain (unlike our approach).
112 As such, counterexamples are matrices, not points over the state space, and fur-
113 thermore the use of LMI solvers does not in principle lead to sound outcomes.
114 Along the above line, [24] expands this approach towards robust synthesis; [25]
115 instead employs MPC within the learner to suggest template functions, which
116 are later verified via semi-definite programming relaxations (again, possibly gen-
117 erating counterexamples by solving optimisation problems over a given domain).
118 Whilst inspired by this line of work, our contribution provides a simple (with
119 interpretable counterexamples that are points over the state space) yet effective
120 (scalable to at least 10-dimensional models) SAT-based CEGIS implementation,
121 which automates the construction of Lyapunov functions and associated validity
122 domains, which is sound, and also applicable to parameterised models.

123 The remainder of the paper is organised as follows. In Section 2 we present the
124 SMT Z3 solver and the inductive synthesis (IS) framework. The implementation
125 of CEGIS, for both linear and non-linear models, is explained in Section 3.

126 Experiments and case studies are in Section 4. Finally, conclusions are drawn in
127 Section 5.

128 2 Formal Verification – Concepts and Techniques

129 In this work we use Z3, an SMT solver, and the CEGIS architecture, to build
130 and to verify Lyapunov functions.

131 2.1 Satisfiability Modulo Theory

132 A Satisfiability Modulo Theory problem is a decision problem formulated within
133 a theory, e.g. first-order logic with equality [26]. The aim is to check whether a
134 first-order logical formula within such theory, referred to as an SMT instance, is
135 satisfied. For example, a formula can be the inequality $3x_0 + x_1 > 0$ evaluated
136 within the theory of linear inequalities. An SMT solver is a software that checks
137 the satisfiability of an SMT instance, i.e. whether there exists an instantiation
138 of the formula that evaluates to **True**. SMT solvers can be useful for function
139 synthesis, namely to mechanically construct a function, given requirements on
140 its output.

141 2.2 The Z3 SMT Solver

142 Z3 [7,27] is a powerful SMT solver that integrates SAT solvers, theory solvers for
143 equalities and interpreted functions, satellite solvers for arithmetic, real, array,
144 and other theories, and an abstract machine to handle quantifiers. Receiving an
145 input formula, Z3 represents it as an abstract syntax tree and processes it with
146 its SAT solver core, until it returns **SAT** if the formula is satisfiable, or **UNSAT**
147 otherwise.

148 *Example 1 (Operation of Z3).* Consider the formula $a = b \wedge f(a) = f(b)$ in the
149 theory of equality. To verify its satisfiability, Z3 constructs a syntax tree, with
150 nodes for each variable (a, b) and formulae ($a = b, f(a), f(b), f(a) = f(b)$). Once
151 the tree is built, Z3 merges a with b and $f(a)$ with $f(b)$ to represent the equality
152 operation and, in order to verify the correctness of the assertion, applies the
153 congruence rule $\bigwedge_{i=0}^{n-1} x_i = y_i \Rightarrow f(x_0, \dots, x_{n-1}) = f(y_0, \dots, y_{n-1})$ to conclude
154 that $a = b \Rightarrow f(a) = f(b)$. Finally, nodes $a = b$ and $f(a) = f(b)$ are merged and
155 Z3 returns **SAT**. \square

156 Of particular interest for the synthesis of Lyapunov functions, is the ability of
157 Z3 to solve polynomial constraints. Z3 stores and exactly manipulates algebraic
158 real numbers that are roots of rational univariate polynomials: this is done for
159 an algebraic real α , by storing a polynomial $p(x)$ for which $p(\alpha) = 0$ and two
160 rationals l, u such that $p(x) = 0$ for $x \in (l, u)$ if and only if $x = \alpha$. In this work,
161 Z3 has been used through its Python APIs, named Z3Py. An example of a simple
162 assertion verification follows.

163 *Example 2 (Assertion in Z3).* Consider the (valid) formula $x \geq 0 \Rightarrow 3x + 1 > 0$.
164 The code using Z3Py results in:

```
165 x = Real('x')
166 s = Solver()
167 s.add(Implies(x >= 0, 3 * x + 1 > 0))
168 print(s.check())
```

169 which evaluates (as expected) to SAT. □

170 2.3 Inductive Synthesis - CEGIS

171 An approach to solve second-order logic problems, such as those characterising
172 the synthesis of Lyapunov functions, is *inductive synthesis* (IS). IS infers general
173 rules (or functions) from specific examples (observations), entailing the process of
174 generalisation. Within the IS procedure, a synthesiser attempts the construction
175 from a (usually small) subset of the original specifications. It then generalises to
176 the complete specification by identifying patterns in the input data.

177 An exemplar of IS is the CEGIS framework. Fig. 1 depicts the relation be-
178 tween its two main components. It sets off with a given specification ψ over a
179 set \mathcal{I} for the synthesis. The synthesis engine (a component that will be also de-
180 noted as *learner*) provides a candidate solution for ι , a subset of \mathcal{I} , the space of
181 possible inputs. This candidate solution is passed to a second component, called
182 *verifier*, that acts as an oracle: either it approves the solution over the entire \mathcal{I} ,
183 so that the process terminates, or it finds an instance \bar{x} (a counterexample in
184 \mathcal{I}) where the candidate solution does not comply with the specifications. The
185 learner takes \bar{x} and adds it to ι , computing a new (more general) candidate solu-
186 tion for the problem. This cycle is repeated. Note that this algorithm might not
187 terminate, depending on the structure of \mathcal{I} , or might take many cycles to find
188 a proper solution: in those instances, tailored candidate solutions and insightful
189 counterexamples are necessary. In this work, the IS is implemented using SMT-
190 solvers. The verifier finds counterexamples \bar{x} by seeking a witness of the negated
191 formula $\neg\psi$, namely trying to prove that a violation of the formula exists. The
192 learner might employ SMT solvers to solve the system of constraints generated
193 by the counterexamples, i.e. to find a valid instance of such constraints, however
194 in general it does not need to be sound, as it is the verifier that guarantees
195 the soundness of the proposed solution. Section 3.1 illustrates the two CEGIS
196 components, the learner L and the verifier Z in relation to Lyapunov function
197 synthesis.

Example 3 (CEGIS Operation). Assume the task is the synthesis of a function $g(x)$ that satisfies the following formula $F(g(x))$:

$$\exists g(x) \forall x \in \mathbb{R} : \psi, \text{ where } \psi(g(x)) = g(x) + 1 > 0.$$

The learner L offers an initial (often naïve, random or default) candidate, e.g. $g(x) = x$, and passes it to the verifier Z . The verifier checks the validity of

$\psi(x) = x + 1 > 0, \forall x \in \mathbb{R}$, by searching an instance \bar{x} that might invalidate the formula. Z finds that $\bar{x} = -1$ invalidates the formula, thus sends \bar{x} to L , which incorporates this counterexample to synthesise a new $g(x)$. The learner now adds a constraint on the next candidate, as

$$C := g(\bar{x} = -1) + 1 > 0, \quad \forall x \in \mathbb{R},$$

198 such that the new candidate solution satisfies the formula at $\bar{x} = -1$. The
 199 learner now proposes $g(x) = x^2$, which satisfies C , and passes it to Z . The
 200 verifier searches for a counterexample to $\psi(x^2)$, but cannot find any. Thus, it
 201 exits the loop with an UNSAT answer, which proves that the synthesised function
 202 $g(x) = x^2$ is valid $\forall x \in \mathbb{R}$. \square




Fig. 1. CEGIS-based inductive synthesis. The iterative procedure loops between a learner L and a verifier Z . L provides a candidate solution S to the verifier Z , which asserts its validity or outputs a counterexample \bar{x} . The learner provides a new solution encompassing also \bar{x} . The procedure stops once no counterexamples are found.

203 3 Automated and Sound Synthesis of Lyapunov 204 Functions via CEGIS and SMT

205 Consider a dynamical system $\dot{x} = f(x)$, where $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, and assume that
 206 point $x_e \in \mathbb{R}^n$ is an equilibrium, namely such that $f(x_e) = 0$ – without loss of
 207 generality, we assume that $x_e = 0$. The goal is assessing the stability of such
 208 equilibrium point via the synthesis of a Lyapunov function $V(x) : \mathbb{R}^n \rightarrow \mathbb{R}$. The
 209 stability of an equilibrium is a significant property to study, as it guarantees that
 210 trajectories starting by the equilibrium remain close to it at all times (how close
 211 can often be quantified, as done later in this work). If $V(x)$ fulfils the following
 212 two conditions, $\forall x \in \mathcal{D}$,

$$V(x) > 0, \quad \dot{V}(x) = \nabla V(x) \cdot f(x) \leq 0, \quad (1)$$

213 where \mathcal{D} is a domain of interest containing x_e then the Lyapunov function ensures
 214 that for every initial point contained in \mathcal{D} , the trajectories of the models do not
 215 escape \mathcal{D} (with reference to notations introduced above, the condition in (1)

216 represents the requirement ψ , and \mathcal{D} denotes the set of inputs \mathcal{I}). We use the
 217 following polynomial expression for the Lyapunov function

$$V(x) = \sum_{l=1}^c (x^l)^T P_l x^l, \quad (2)$$

218 where x^l represents the element-wise exponentiation of vector x , i.e. element $x(j)$
 219 to the power l , $\forall j = 1, \dots, n$; $P_l \in \mathbb{R}^{n \times n}$ is a weighting matrix associated with
 220 x^l , and c is the order of the polynomial function. In order to obtain a proper
 221 Lyapunov function $V(x)$, the synthesiser is asked to verify the specification ex-
 222 pressed by the formula

$$F(V(x)) : \forall x \in \mathcal{D}, V(x) > 0 \wedge \dot{V}(x) \leq 0. \quad (3)$$

This specification requires the Lyapunov function to be positive definite, and
 not to increase along the trajectories of the model. For linear systems, unless
 otherwise stated, we consider $\mathcal{D} = \mathbb{R}^n \setminus \{0\}$ and $c = 1$, as it is known that
 quadratic functions are sufficient to prove the stability of linear models over the
 whole state space. Formula (3) keeps the elements of P uninterpreted, and thus
 they are parameters to be found. Notice that the second-order formula

$$\exists P \in \mathbb{R}^{n \times n} : \forall x \in \mathcal{D}, V(x) > 0 \wedge \dot{V}(x) \leq 0,$$

223 would return a boolean value, i.e. **True** or **False**: to obtain the synthesised $V(x)$
 224 function, we remove the existential quantifier.

225 3.1 The CEGIS Architecture for Lyapunov Function Synthesis

226 We introduce the CEGIS architecture to find Lyapunov functions. To better il-
 227 lustrate the methodology, we start by considering linear models (the non-linear
 228 case is further discussed in Section 3.2). As mentioned earlier, two components
 229 characterise the CEGIS approach: a learner and a verifier. The CEGIS architec-
 230 ture takes the system matrix A and it outputs a matrix P as the key component
 231 of the function $V(x)$, verifying the conditions in Eq. (1). We denote by \bar{P}_i ,
 232 $i = 0, 1, 2, \dots$ the *candidate* matrices yet to be verified, i.e. the outputs of the
 233 learner. As anticipated earlier, referring to Eq. (2), we set $c = 1$ and $\mathcal{D} = \mathbb{R}^n \setminus \{0\}$.

234 **Verifier** The scope of a verifier is twofold: generate a counterexample to the
 235 validity of the candidate Lyapunov function, or certify its validity over a domain
 236 of interest. We implement the verifier in Z3.

237 The methodology to assert the correctness of a Lyapunov function is as fol-
 238 lows. Assume the learner computes a candidate Lyapunov function $V(x)$ and
 239 passes it to the verifier (in case of a linear function, the learner offers a matrix
 240 \bar{P}_i). The goal of the verifier is to assert the validity of formula F from (3) ac-
 241 cording to the specification ψ in (1). The check is performed by negating F : if
 242 there exists a vector \bar{x} that satisfies $\neg F$, it is a counterexample for F ; if it does

243 not exist, formula F is valid and the candidate Lyapunov function is an actual
 244 Lyapunov function. The domain \mathcal{D} is encoded as an additional formula. Assume,
 245 as an example, the domain is an hyper-sphere of radius one: \mathcal{D} can be written
 246 formally as $d: \|x\|^2 \leq 1$. The final formula thus results in $\neg F \wedge d$.

247 A counterexample \bar{x} can satisfy either $V(\bar{x}) \leq 0$, $\dot{V}(\bar{x}) > 0$, or both con-
 248 ditions. Reasoning on either condition, it is easy to show that if there exists a
 249 counterexample \bar{x} invalidating a matrix \bar{P} , then there exists an infinite number of
 250 counterexamples for this \bar{P} . Thus, particularly for high-dimensional models the
 251 generation of meaningful counterexamples is crucial to find a Lyapunov function
 252 quickly.

253 Let us denote \bar{x}_i , $i = 1, 2, \dots$, the series of counterexamples provided by
 254 the verifier and \bar{P}_i the series of candidate Lyapunov function matrices provided
 255 by the learner. In this setting, the learner proposes the first default candidate
 256 matrix \bar{P}_0 ; the verifier will (possibly) provide a counterexample \bar{x}_0 ; the learner
 257 then includes \bar{x}_0 in the set of constraints (cf. Section 3.1) and offers a new
 258 candidate \bar{P}_1 .

259 In this work, we let Z3 generate counterexamples without any further goals.
 260 However, more generally counterexamples can be generated adding constraints,
 261 e.g. linear independence or orthogonality. Intuitively, more constraints might
 262 generate “better” candidates by the learner, albeit at an increase in computa-
 263 tional cost.

264 As intuition suggests, if we were to work with models having a diagonal ma-
 265 trix A , then the synthesis of diagonal candidates \bar{P}_i and of a diagonal solution P
 266 would reduce the number of variables needed, thus speeding up the computation.
 267 As such, if A is not diagonal but diagonalisable, the algorithm pre-computes the
 268 system diagonalisation and feeds it to the CEGIS architecture returning a ma-
 269 trix P for the diagonal system, which is then converted to a solution for the
 270 original model.

271 **Learner** A learner is the CEGIS component designated to suggest a candidate
 272 solution for the problem under consideration. Within our framework, a learner
 273 solves linear inequalities derived from $F(V(\bar{x}))$ as per Eq. (3), while memorising
 274 the set of counterexamples $\{\bar{x}_i \mid \neg F(\bar{x}_i)\}$ generated by the verifier. Whilst the
 275 verifier works over continuous domains, note that the learner only considers a
 276 *finite* number of points to synthesise the candidate Lyapunov function. At each
 277 iteration i , the learner is tasked to solve $2i$ linear inequalities: i inequalities for
 278 $V \geq 0$ and i for $\dot{V} \leq 0$ – this is two inequalities per counterexample, so a set of
 279 useful counterexamples is vital to achieve efficiency.

280 We implement two learners, for comparison: 1) a numerical and 2) a Z3-
 281 based learner. However, our CEGIS architecture can in principle accommodate
 282 any learner. The first learner uses Gurobi [6], a fast, commercial optimisation
 283 solver for, among others, linear and quadratic programming problems, support-
 284 ing continuous variables. Notice that the synthesis is a linear program: variables
 285 $p_{i,j}$, the entries of matrix P , appear linearly within the inequalities in $F(V(\bar{x}_i))$.
 286 Gurobi is thus expected to outperform an SMT solver in this specific task. How-

287 ever these variables do not represent real numbers, but floating point numbers
 288 that are approximated at machine precision. The second learner instead em-
 289 ploys Z3, which is numerically sound and not affected by machine precision. Z3
 290 solves an SMT instance to synthesise $V(x)$: it asserts the satisfiability of Eq. (3)
 291 $F(V(\bar{x}_i))$ for all collected counterexamples \bar{x}_i .

292 As mentioned earlier, the number of inequalities to be solved depends on the
 293 number of counterexamples, which can grow to be quite large. Whilst the verifier
 294 ought to generate useful counterexamples, the learner is optimised to output a
 295 matrix \bar{P}_i that is easy to handle. From the comparison between a numerical
 296 learner (running on Gurobi) and a sound one (based on Z3), the compromise
 297 between speed and soundness results is evident (cf. Section 4). Z3 is sound, yet
 298 slower when compared to the numerical learner.

299 3.2 Lyapunov Function Synthesis for Non-linear Models

300 The problem of synthesizing Lyapunov functions and their region of validity for
 301 a general non-linear system $\dot{x} = f(x(t))$ is approached via linearisation or via
 302 direct computation.

The linearisation approach consists of three steps for the learner: we first
 linearise the $f(x(t))$, obtaining

$$\dot{\tilde{x}}(t) = A_L \tilde{x}(t),$$

where A_L is the Jacobian of $f(x(t))$ evaluated at x_e ; we then compute matrix
 P – and quadratic Lyapunov function $V(x) = x^T P x$ – on the linearised system;
 finally, we find \mathcal{R} , defined as the set in which the linear Lyapunov function
 is valid. Next, we detail the synthesis of region \mathcal{R} . Consider, without loss of
 generality, an autonomous non-linear system with (at least one) equilibrium
 point $x_e = 0$. Assume the CEGIS procedure is successful, i.e. it finds a Lyapunov
 function $V_L(x) = x^T P x$ that guarantees the asymptotic stability of system $\dot{\tilde{x}} =$
 $A_L \tilde{x}$ around x_e . We now compute the region where $V_L(x)$ guarantees stability
 with the original system, i.e. $\dot{x} = f(x)$. In view of the existence of $V_L(x)$ and by
 definition of linearisation, there exists a neighbourhood of the origin \mathcal{B}_0 in which
 the derivative of the Lyapunov function $\dot{V}(x)$ is non-positive; formally such set
 is defined as

$$\mathcal{B}_0 = \{x \in \mathbb{R}^n \setminus \{0\} \mid \dot{V}(x) \leq 0\},$$

where $\dot{V}(x)$ is computed on the original system, namely

$$\dot{V}(x) = \nabla V(x) \cdot f(x).$$

Let us define the boundary of \mathcal{B}_0 as $\partial\mathcal{B}_0 = \{x \in \mathbb{R}^n \setminus \{0\} \mid \dot{V}(x) = 0\}$. This set
 may be composed by single points or regions of the state space: in this case, we
 find r , the closest point to the equilibrium that belongs to $\partial\mathcal{B}_0$, as

$$r = \min_{x \in \partial\mathcal{B}_0} \sum_l x(l)^2.$$

303 We finally compute region \mathcal{R} as a hyper-sphere of radius r ,

$$\mathcal{R} = \{x \in \mathbb{R}^n \setminus \{0\} \mid \|x\|^2 < r\}, \quad (4)$$

304 defining the region where the Lyapunov function is valid. Finally, region \mathcal{R} is
 305 tested with the verifier: formula $F(V(x))$ from Eq. (3) is passed to Z3 with
 306 $\mathcal{D} = \mathcal{R}$. Our implementation uses a numerical optimisation technique to com-
 307 pute a value for r that is passed to Z3, as Z3 does not natively handle non-linear
 308 optimisation problems. With this selection, the region \mathcal{R} represents a sound
 309 under-approximation of the maximal stability region. The linearisation method
 310 is used in view of its rapid and effective synthesis capability. However, it pro-
 311 duces a Lyapunov function that does not ensure global stability when one of
 312 the eigenvalues of A_L is equal to zero. This is a well-known limitation of the
 313 linearisation, which suggests a more formal approach, called *direct computation*
 314 *method*.

The direct computation method, as the name suggests, analytically computes
 $V(x)$ and $\dot{V}(x)$ from a template $V(x)$ as in Eq. (2). The learner is tasked with
 resolving conditions ψ obtained by a light relaxation of the two inequalities in
 (1), namely

$$V(x) \geq 0, \quad \dot{V}(x) = \nabla V(x) \cdot f(x) \leq 0.$$

315 Note that the first inequality is not strict: this relaxation allows for a faster
 316 computation of a candidate. The verifier, on the other hand, produces coun-
 317 terexamples for $V(x) > 0$, thus retaining soundness of the overall procedure .
 318 The CEGIS framework allows the separation between synthesis and verification.
 319 So whilst the learner might propose candidates being completely independent
 320 from domain \mathcal{D} , the verifier is responsible to assert or to find the domain of
 321 validity \mathcal{D} . Our implementation establishes that at first the verifier checks the
 322 validity of $V(x)$ on the whole state space $\mathcal{D} = \mathbb{R}^n$; if the computation is not suc-
 323 cessful – namely, the computational time is greater than a predefined timeout –
 324 the verifier checks its validity over a smaller region, e.g. $\mathcal{D} = [-1, 1]^n$, and so on.

325 3.3 Lyapunov Function Synthesis for Parametric Models

326 Parametric models represent a challenge for both sound and numerical solvers.
 327 Let us remark that both Gurobi and Z3 can not synthesise functions in the
 328 presence of uncertainty, whereas Z3 can provide counterexamples using one (or
 329 more) variables as fixed parameters, using the quantifier **ForAll**.

330 Let us consider variable x , a parameter μ and a formula $\psi(x, \mu)$: Z3 can find
 331 a counterexample for all values of μ by validating **ForAll**(μ, ψ). If μ belongs
 332 to a range $[l, u]$, Z3 can find a counterexample by checking $\psi \wedge \mu \geq l \wedge \mu \leq u$.
 333 This provides a counterexample $(\bar{x}, \bar{\mu})$ for x and μ , respectively.

334 The synthesis procedure is split into two steps, in view of the inability of
 335 Z3 and Gurobi to propose parametric solutions. The first step synthesises a
 336 candidate Lyapunov function solely using the constraint $V(x) > 0$, in which no
 337 parameter appears. The second step evaluates the constraint $\dot{V} \leq 0$ to propose
 338 a parametric Lyapunov function exploiting the results from the first step. The
 339 following example details the procedure.

Example 4. Consider a two-dimensional linear parametric system [19] and a candidate Lyapunov function

$$\begin{cases} \dot{x} = y \\ \dot{y} = -(2 + \mu)x - y \end{cases}, \quad V(x, y) = p_1x^2 + p_2y^2.$$

Assume the first guess of the learner is invalid, i.e. the verifier finds a counterexample for the validity of $V(x, y)$. The counterexample (\bar{x}, \bar{y}) is then sent to the learner. The synthesis procedure is split into two steps: the first step entails the synthesis solely accounting for $V(\bar{x}, \bar{y}) > 0$. The learner is tasked to solve

$$V(\bar{x}, \bar{y}) = p_1\bar{x}^2 + p_2\bar{y}^2 > 0,$$

where p_1, p_2 are the variables of the inequality. The learner will propose values \bar{p}_1 and \bar{p}_2 satisfying the inequality. Two second step removes one of the synthesised \bar{p}_i , e.g. \bar{p}_1 , in order to re-synthesise it including the parameters found in \dot{V} . In practical terms, the expression of \dot{V} is evaluated at \bar{x}, \bar{y} and \bar{p}_2 , as

$$\dot{V} = 2p_1\bar{x}\bar{y} - 2\bar{p}_2\bar{y}^2 - 2(\mu + 2)\bar{x}\bar{y} \leq 0 \implies p_1 \leq \bar{p}_2 \left(\frac{\bar{y}}{\bar{x}} + 2 + \mu \right).$$

340 We choose the value p_1 that satisfies the equality. The candidate Lyapunov
 341 function thus results in $V(x, y) = \bar{p}_2 \left(\frac{\bar{y}}{\bar{x}} + 2 + \mu \right) \cdot x^2 + \bar{p}_2 \cdot y^2$. This procedure
 342 holds as long as $\bar{x} \neq 0$: if this is not the case, we can either choose to synthesise
 343 a new value for p_2 or simply maintain the numerical values obtained after the
 344 first step. In the latter case, once the candidate Lyapunov function is passed to
 345 the verifier, a new counterexample will be generated and the procedure can be
 346 repeated until a parametric Lyapunov function is found and verified. Another
 347 possible approach is based on the mixed-terms removal: p_1 is synthesised so
 348 that the terms carrying $\bar{x}\bar{y}$ cancel out. Further, the choice of p_1 satisfying the
 349 equality is arbitrary: we can add a negative constant to its value to solve the
 350 strict inequality instead. Finally, more than one parameter \bar{p}_i can be removed
 351 in the second step: this can spread the parametric coefficients among more than
 352 one p_i . However, this is likely to increase the computational cost in view of the
 353 inequality being a function of more than one variable. \square

354 4 Case Studies and Experiments

355 In this Section we outline a few experiments to challenge the validity of our
 356 approach. Our technique is coded in Python 2.7 [28], using external libraries as
 357 the numerical solver Gurobi and the SMT solver Z3 (cf. Section 2). Specifically,
 358 we compare two CEGIS architectures:

- 359 1. Gurobi learner and Z3 verifier,
- 360 2. Z3 learner and Z3 verifier,

361 later denoted as *Gurobi-CEGIS* and *Z3-CEGIS*, respectively. Whilst Z3 is an effi-
 362 cient verifier, it carries the weight of exact representations. We therefore compare
 363 its use within the learner to that of a numerical solver such as Gurobi - recall
 364 that the learner does not need to be sound. A relevant feature of the synthesis
 365 procedure is its *linearity* in the entries of matrix P : we expect an efficient LP
 366 solver to outperform an SMT solver. As such, we study the expected tradeoff
 367 between speed and precision. As specified earlier, the initial candidate for the
 368 learner \bar{P}_0 is arbitrary: we challenge the procedure by setting $\bar{P}_0 = -I$, which
 369 does not satisfy the first positivity condition for Lyapunov functions, thus show-
 370 ing that even with an ill-suited initial guess the procedure can rapidly synthesise
 371 a valid Lyapunov function.

372 We consider linear, non-linear and parametric ODEs with the origin as (one
 373 of) the equilibrium(a), and aim to obtain a Lyapunov function guaranteeing the
 374 stability of such equilibrium point. The procedure entails the following steps:

- 375 a) a function $f(x)$, $x \in \mathbb{R}^n$, is fed as the input;
- 376 b) a Lyapunov function $V(x)$, as in Eq. (2), is computed;
- 377 c) in the linearisation case, the stability region \mathcal{R} in Eq. (4) for $V(x)$ is found.

378 Let us emphasise that Z3 is unable to handle non-polynomial terms, which rep-
 379 represents the only limitation of our approach. Unlike most of the literature, coun-
 380 terexamples are not limited to a finite set but searched over the whole \mathbb{R}^n .

381 Linear models are certainly an easier task than polynomial systems. The
 382 study with linear models focuses mainly on the scalability of the method, en-
 383 compassed by the average and maximum/minimum computational time, and the
 384 number of iterations performed. We generate $N = 100$ random linear models of
 385 dimension $n \in [3, 10]$. For each linear system, the entries of matrix A range
 386 within $[-1000, 1000] \in \mathbb{R}$. For each test we set $c = 1$ (cf. Eq. (2)), namely we
 387 impose a quadratic structure to the Lyapunov function, and collect the num-
 388 ber of iterations of the procedure, i.e. the number of counterexamples needed
 389 to compute a valid Lyapunov function, and the total elapsed time. Recall that
 390 the initial synthesiser’s candidate is $\bar{P}_0 = -I$, which challenges the reliability
 391 of our method with a bad initial condition. A 180 seconds time out is set for
 392 every run. Results comparing the numerical learner using Gurobi and the sound
 393 learner using Z3 are reported in Table 1. The average values, as well as the min-
 394 imum and maximum value among the N random systems, are computed on the
 395 synthesis tests that have not timed out. The number of timed out procedures
 396 are also listed in the Table.

397 With regards to non-linear and parametric models, we assess our approach
 398 over a suite of examples taken from related work on Lyapunov function synthesis
 399 [14], [15], [16], [19], which are reported in the following. The value c from Eq.
 400 (2) is set heuristically as $\text{ceil}(d/2)$, where d is the order of the system, in view
 401 of the interpretation of Lyapunov functions as storage functions. Due to ease of
 402 implementation, only Z3-CEGIS performs the synthesis with $c > 1$ and in the
 403 case of parametric models. Results in terms of computational time and iterations
 404 are reported in Table 2. Experiments are run on a 4-core Dell laptop with Fedora
 405 30 and 8GB RAM.

Example 5. Consider the model [14]

$$\begin{aligned}\dot{x}_1 &= -x_1^2 - 4x_2^3 - 6x_3x_4, & \dot{x}_4 &= x_1x_3 + x_3x_6 - x_4^3, \\ \dot{x}_2 &= -x_1 - x_2 + x_5^3, & \dot{x}_5 &= -2x_2^3 - x_5 + x_6, \\ \dot{x}_3 &= x_1x_4 - x_3 + x_4x_6, & \dot{x}_6 &= -3x_3x_4 - x_5^3 - x_6.\end{aligned}$$

406 Z3-CEGIS finds the Lyapunov function $V(x) = 2x_1^2 + 4x_2^4 + x_3^2 + 11x_4^2 + 2x_5^4 + 4x_6^2$,
407 ensuring stability over the whole state space. \square

Example 6. Consider the model [19]

$$\begin{cases} \dot{x} = -x^3 + y \\ \dot{y} = -x - y. \end{cases}$$

408 Gurobi-CEGIS finds the Lyapunov function $V(x) = 5 \cdot 10^{-5}x^2 + 5 \cdot 10^{-5}y^2$,
409 whereas Z3-CEGIS finds $V(x) = 0.5x^2 + 0.5y^2$, both ensuring global stability.
410 The linearised Gurobi-CEGIS finds $V(x) = 3.2 \cdot 10^{-3}x^2 + 3.2 \cdot 10^{-3}y^2$ also ensuring
411 stability on the whole state space. \square

Example 7. Consider the system [16]

$$\begin{cases} \dot{x}_1 = -x_1^3 - x_1x_3^2, \\ \dot{x}_2 = -x_2 - x_1^2x_2, \\ \dot{x}_3 = -x_3 - \frac{3x_3}{x_3^2 + 1} + 3x_1^2x_3. \end{cases}$$

412 Note that the term $x_3^2 + 1$ is always non-negative, therefore we can consider
413 $\dot{V}(x) \cdot (x_3^2 + 1) \leq 0$. Gurobi-CEGIS finds the Lyapunov function $V(x) = 32 \cdot$
414 $10^{-4}x_1^2 + 32 \cdot 10^{-4}x_2^2 + 8 \cdot 10^{-4}x_3^2$, whereas Z3-CEGIS finds $V(x) = 3x_1^2 + x_2^2 + x_3^2$,
415 both ensuring global stability. \square

Example 8. Consider the system [19]

$$\begin{cases} \dot{x} = -x - 1.5x^2y^3, \\ \dot{y} = -y^3 + 0.5x^3y^2. \end{cases}$$

416 Z3-CEGIS finds $V(x) = 1/3x^2 + y^2$, valid on the whole \mathbb{R}^2 . Gurobi-CEGIS
417 returns an error, as it finds $V(x) = 1.00066454641347x^2 + 2.99933545358653y^2$
418 that is *not* a valid Lyapunov function. The correct solution, $V(x) = x^2 + 3y^2$,
419 can not be attained in view of lack of convergence of the optimisation algorithm.
420 On the other hand, the linearised Gurobi-CEGIS delivers $V(x) = 32 \cdot 10^{-4}x^2 +$
421 $2 \cdot 10^{-4}y^2$ with a radius $r = 1.25$. \square

Example 9. Consider the system [19]:

$$\begin{aligned} \dot{x}_1 &= -x_1 + x_2^3 - 3x_3x_4, & \dot{x}_2 &= -x_1 - x_2^3, \\ \dot{x}_3 &= x_1x_4 - x_3, & \dot{x}_4 &= x_1x_3 - x_4^3. \end{aligned}$$

422 Z3-CEGIS finds the Lyapunov function $V(x) = 2x_1^2 + x_2^4 + 3201/1024x_3^2 +$
 423 $2943/1024x_4^2$, ensuring global stability. \square

Example 10. Consider the parametric linear system [19]

$$\begin{cases} \dot{x} = y, \\ \dot{y} = -(2 + \mu)x - y, \end{cases}$$

424 where $\mu \in (-2, 5]$. Z3-CEGIS discovers the Lyapunov function $V(x) = (\mu +$
 425 $2)x^2 + y^2$, ensuring stability on the whole state space. \square

Example 11. Consider the parametric system [19]

$$\begin{cases} \dot{x} = -(1 + \mu_1)x + (4 + \mu_2)y, \\ \dot{y} = -(1 + \mu_3)x - \mu_4y^3, \end{cases}$$

426 where $\mu_i \in [0, 100]$ for $i = 1, \dots, 4$. Z3-CEGIS discovers the Lyapunov function
 427 $V(x) = \frac{\mu_3 + 1}{\mu_2 + 4}x^2 + y^2$ that asserts stability on the whole state space. \square

428 As expected, Gurobi is faster than Z3 in terms of iterations and computa-
 429 tional time. The gap becomes larger with a high-dimensional system, as the SMT
 430 learner does not implement any optimisation techniques. The Z3-CEGIS synthe-
 431 sis is performed via an SMT call, which grows in complexity as the number of
 432 constraints, i.e. counterexamples, increases. Gurobi, on the other hand, using
 433 optimisation techniques converges faster to a candidate solution that is closer to
 434 the actual solution.

435 Notice that the coefficients of the Lyapunov function synthesised by Gurobi
 436 are small in magnitude, as the linear programming problem can encompass the
 437 minimisation of coefficients in its setup. On the other hand those obtained from
 438 Z3 (rational fractions) are arguably more interpretable. A very interesting result
 439 comes from Example 8. Gurobi-CEGIS converges towards the correct Lyapunov
 440 function, yet it can not reach the exact numerical values in view of the algorith-
 441 mic precision. Gurobi numerical guidelines [6] suggest that, as a rule of thumb,
 442 the ratio of the largest to the smallest coefficient of the LP problem should
 443 be less than 10^9 . In our setting, the coefficients are the counterexamples found
 444 by Z3, which might require high precision. In this case, the issue is (probab-
 445 ily) caused by a counterexample $\bar{x} \simeq [-755145, 1/8]$, where the first element
 446 is actually represented as a (very long) ratio between two integers. The ratio
 447 between the two \bar{x} coefficient is in the order of 10^7 . Roughly speaking, the coun-
 448 terexamples generated by Z3 depend on the complexity of the tested model: a
 449 high-order system might generate numerically ill-conditioned counterexamples,

450 as this example shows. It is also significant how the numerical algorithm tries
451 to converge to a correct solution. The first candidate Lyapunov function results
452 $V(x) = 1.07079661938449x^2 + 2.92920338061551y^2$ and it takes 99 counterexam-
453 ples to reach the final value (cf. Example 8), until the procedure stops, resulting
454 in an infeasible problem. Even enveloping the numerical values with the Python
455 types `Rational`, `Decimal`, `Fraction`, or the function `simplify` do not help in
456 this context, the limitation being Gurobi’s numerical precision.

n	Gurobi-CEGIS			Z3-CEGIS		
	Iterations	Time [sec]	Oot	Iterations	Time [sec]	Oot
3	3 [3, 3]	0.48 [0.33, 0.77]	–	3.03 [3, 4]	0.49 [0.4, 0.70]	–
4	3.10 [3, 4]	0.53 [0.36, 1.20]	–	5.93 [4, 7]	0.68 [0.54, 1.07]	–
5	4.15 [4, 5]	1.33 [1.08, 1.97]	–	7.38 [5, 12]	1.67 [1.10, 3.03]	–
6	6.99 [4, 10]	3.88 [2.41, 4.97]	–	9.10 [6, 10]	7.48 [2.40, 54.44]	–
7	8.56 [4, 12]	12.64 [2.9, 62.3]	–	12.88 [5, 17]	17.63 [5.41, 20.3]	1
8	9.14 [3, 13]	21.50 [3.9, 114.16]	1	16.2 [3, 25]	23.91 [4.05, 35.08]	1
9	15.72 [3, 32]	29.98 [3.87, 78.5]	2	22.47 [4, 35]	34.41 [5.67, 48.96]	5
10	18.45 [3, 41]	40.63 [6.17, 46.65]	5	27.25 [5, 47]	44.63 [6.32, 101.2]	7

Table 1. Comparison between Gurobi-CEGIS and Z3-CEGIS over n -dimensional linear models. The first values are the average performance on the $N = 100$ randomly generated models, and within brackets the minimum and maximum values. Oot is the number of runs (out of N) not finishing after 180 [sec].

Example #	Gurobi-CEGIS		Z3-CEGIS	
	Time [sec]	Iterations	Time [sec]	Iterations
5	–	–	18.38	4
6	0.32	2	1.27	5
7	0.37	4	0.60	3
8	0.16	2	0.27	2
9	–	–	9.26	3
10	–	–	0.14	3
11	–	–	0.23	3

Table 2. Comparison between Gurobi-CEGIS and Z3-CEGIS for non-linear models (see Examples description in main text). The result for Gurobi-CEGIS in Example 8 is obtained via linearisation.

457 5 Conclusions and Future Work

458 In this work, we have studied the problem of automated and sound synthesis
459 of Lyapunov functions. We have exploited a CEGIS framework, equipped with

460 a sound verifier (the Z3 SMT solver) and with either a numerical LP solver
461 (Gurobi) or a sound (Z3) learner.

462 We have provided a simple – yet effective – methodology to synthesise Lyapunov
463 functions for linear, polynomial and parametric systems and shown evidence
464 of scalability and reliability of our method using benchmarks from the
465 literature. We have in particular synthesised quadratic Lyapunov functions for
466 linear models and verified their validity on the whole state space. We have tackled
467 non-linear models following two approaches: either 1) the computation of
468 Lyapunov functions over the linearised system and the synthesis of its validity
469 region; or 2) the direct computation of a higher-order Lyapunov function.

470 Future work includes the implementation of synthesis techniques for Gurobi-
471 CEGIS for high-order and parametric models, together with the study of optimisation
472 techniques for the synthesis in Z3-CEGIS: the tuning of the SMT solvers
473 leaves much room, for example in order to provide insightful counterexamples
474 or to additionally optimise an objective function. Further, we aim at embedding
475 CEGIS with neural networks (simpler function approximators) to replace the
476 learner, whilst maintaining the verification in the hands of an SMT solver.

477 References

- 478 1. P. Giesl and S. Hafstein, “Review on Computational Methods for Lyapunov Functions,” *Discrete and Continuous Dynamical Systems-Series B*, vol. 20, no. 8, pp. 2291–2331, 2015.
- 479 2. C. M. Kellett, “Classical converse theorems in lyapunov’s second method,” *Discrete Continuous Dyn. Syst. Series B*, vol. 20, no. 8, pp. 2333–2360, 2015.
- 480 3. A. Solar-Lezama, L. Tancau, R. Bodik, S. Seshia, and V. Saraswat, “Combinatorial sketching for finite programs,” *ACM Sigplan Notices*, vol. 41, no. 11, pp. 404–415, 2006.
- 481 4. C. David and D. Kroening, “Program Synthesis: Challenges and Opportunities,” *Phil. Trans. R. Soc. A*, vol. 375, no. 2104, p. 20150403, 2017.
- 482 5. D. Kroening and O. Strichman, *Decision Procedures: An Algorithmic Point of View*, ser. Texts in Theoretical Computer Science. An EATCS Series. Springer Berlin Heidelberg, 2016.
- 483 6. Gurobi Optimization, LLC, “Gurobi Optimizer Reference Manual,” 2018. [Online]. Available: <http://www.gurobi.com>
- 484 7. L. De Moura and N. Bjørner, “Z3: An Efficient SMT solver,” in *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2008, pp. 337–340.
- 485 8. R. Kalman and J. Bertram, “Control system analysis and design via the second method of lyapunov: Part i continuous-time systems,” *Trans. AMSE Series D J. Basic Eng.*, vol. 82, no. 2, pp. 371–393, 1960.
- 486 9. N. N. Krasovskii, *Stability of Motion: Applications of Lyapunov’s Second Method to Differential Systems and Equations With Delay*. Stanford Univ. Press, 1963.
- 487 10. J. LaSalle and S. Lefschetz, *Stability by Liapunov’s Direct Method With Applications*. Academic Press, 1961.
- 488 11. V. I. Zubov, *Methods of A. M. Lyapunov and Their Application*. Noordhoff, 1964.

- 504 12. R. Brayton and C. Tong, “Stability of Dynamical Systems: A Constructive Ap-
505 proach,” *IEEE Transactions on Circuits and Systems*, vol. 26, no. 4, pp. 224–234,
506 1979.
- 507 13. P. A. Parrilo, “Structured Semidefinite Programs and Semialgebraic Geometry
508 Methods in Robustness and Optimization,” Ph.D. dissertation, California Institute
509 of Technology, 2000.
- 510 14. A. Papachristodoulou and S. Prajna, “On the Construction of Lyapunov Func-
511 tions using the Sum of Squares Decomposition,” in *Proceedings of the 41st IEEE
512 Conference on Decision and Control, 2002.*, vol. 3. IEEE, 2002, pp. 3482–3487.
- 513 15. S. Prajna, A. Papachristodoulou, and P. A. Parrilo, “SOSTOOLS: Sum of squares
514 Optimization Toolbox for MATLAB–Users Guide,” *Control and Dynamical Sys-
515 tems, California Institute of Technology, Pasadena, CA*, vol. 91125, 2004.
- 516 16. A. Papachristodoulou, J. Anderson, G. Valmorbida, S. Prajna, P. Seiler, and P. Par-
517 rilo, “SOSTOOLS Version 3.03. Sum of Squares Optimization Toolbox for MAT-
518 LAB,” 2018.
- 519 17. R. Geiselhart, R. H. Gielen, M. Lazar, and F. R. Wirth, “An alternative converse
520 lyapunov theorem for discrete-time systems,” *Syst. Control Lett.*, vol. 70, pp. 49–
521 59, 2014.
- 522 18. S. F. Hafstein, “An algorithm for constructing lyapunov functions,” *Electron. J.
523 Differ. Equ. Monograph*, vol. 8, 207.
- 524 19. S. Sankaranarayanan, X. Chen, and E. Abraham, “Lyapunov Function Synthesis
525 using Handelman Representations,” *IFAC Proceedings Volumes*, vol. 46, no. 23,
526 pp. 576–581, 2013.
- 527 20. J. Kapinski, J. V. Deshmukh, S. Sankaranarayanan, and N. Arechiga, “Simulation-
528 guided Lyapunov Analysis for Hybrid Dynamical Systems,” in *Proceedings of the
529 17th international conference on Hybrid systems: computation and control.* ACM,
530 2014, pp. 133–142.
- 531 21. S. Gao, S. Kong, and E. M. Clarke, “dReal: An SMT Solver for Nonlinear Theories
532 over the Reals,” in *International Conference on Automated Deduction.* Springer,
533 2013, pp. 208–214.
- 534 22. Wolfram Research, Inc., “Mathematica, Version 12.0,” champaign, IL, 2019.
- 535 23. H. Ravanbakhsh and S. Sankaranarayanan, “Counter-example guided synthesis of
536 control lyapunov functions for switched systems,” in *IEEE Control and Decision
537 Conference (CDC)*, 2015, pp. 4232–4239.
- 538 24. —, “Robust controller synthesis of switched systems using counterexample
539 guided framework,” in *ACM/IEEE Conference on Embedded Software (EMSOFT)*,
540 2016, pp. 8:1–8:10.
- 541 25. —, “Learning control lyapunov functions from counterexamples and demonstra-
542 tions,” *Autonomous Robots*, pp. 1–33, 2018.
- 543 26. E. M. Clarke, T. A. Henzinger, H. Veith, and R. Bloem, *Handbook of model check-
544 ing.* Springer, 2018, vol. 10.
- 545 27. Microsoft Research, “The Z3 Theorem Prover,” <https://github.com/Z3Prover/z3>,
546 accessed: 2018-07-25.
- 547 28. Python Software Foundation, “Python Language Reference, version 2.7,”
548 <http://www.python.org>.