

This item is the archived peer-reviewed author-version of:

A framework for the competitive analysis of model predictive controllers

Reference:

Bellis Stijn, Denil Joachim, Krishnamurthy Ramesh, Leys Tim, Pérez Guillermo Alberto, Raha Ritam.- A framework for the competitive analysis of model predictive controllers

Lecture notes in computer science - ISSN 1611-3349 - Cham, Springer, 2023, 14235, p. 141-154

Full text (Publisher's DOI): https://doi.org/10.1007/978-3-031-45286-4_11

To cite this reference: <https://hdl.handle.net/10067/1997690151162165141>

A Framework for the Competitive Analysis of Model Predictive Controllers

Stijn Bellis, Joachim Denil, Ramesh Krishnamurthy, Tim Leys, Guillermo A. Pérez, and Ritam Raha

University of Antwerp – Flanders Make, Belgium
ramesh.krishnamurthy@uantwerpen.be

Abstract. This paper presents a framework for the competitive analysis of Model Predictive Controllers (MPC). Competitive analysis means evaluating the relative performance of the MPC as compared to other controllers. Concretely, we associate the MPC with a regret value which quantifies the maximal difference between its cost and the cost of any alternative controller from a given class. Then, the problem we tackle is that of determining whether the regret value of is at most some given bound. Our contributions are both theoretical as well as practical: (1) We reduce the regret problem for controllers modeled as hybrid automata to the reachability problem for such automata. We propose a reachability-based framework to solve the regret problem. Concretely, (2) we propose a novel CEGAR-like algorithm to train a deep neural network (DNN) to clone the behavior of the MPC. Then, (3) we leverage existing reachability analysis tools capable of handling hybrid automata with DNNs to check bounds on the regret value of the controller.

Keywords: Competitive analysis · Hybrid automata

1 Introduction

An optimal control problem (OCP) deals with finding a function $u(t)$, called a *control law* that assigns values to control variables for every time step $t \in \mathbb{R}_{\geq 0}$. The control law should minimize a given cost function $J[x(\cdot), u(\cdot), t_0, t_f]$ evaluated for a time interval (t_0, t_f) and subject to the state-equation constraints $\dot{x}(t) = f[x(t), u(t), t]$. Model predictive controllers (MPC) solve such a control problem for a given f . This paper presents an approach for the competitive analysis of MPC. Competitive analysis, in this context, means evaluating the relative performance of the MPC as compared to other controllers. Referring to the OCP, our approach assumes that a control law $u(t)$ is given to us. Further, we associate to $u(t)$ a *regret value*, which quantifies the maximal difference between its cost and the cost of any alternative control law from a given class \mathcal{C} . Formally, the regret of $u(t)$ is: $Reg(u) := \sup_{c \in \mathcal{C}} \sup_{t_f \in \mathbb{R}_{\geq 0}} J[x(\cdot), u(\cdot), t_0, t_f] - J[x(\cdot), c(\cdot), t_0, t_f]$. If $Reg(u) < r$, then we say that the control law $u(t)$ is r -competitive.

In this work, we first show that the r -competitiveness problem for controllers modeled as hybrid automata is interreducible with the reachability problem for

39 hybrid automata. It follows that the r -competitiveness problem is undecidable.
 40 Fortunately, this also points to using approximate reachability analysis tools
 41 to realize approximate competitive analysis. Based on the latter, we propose
 42 a counterexample-guided abstraction refinement (CEGAR) framework that ab-
 43 stracts a given MPC using a deep neural network (DNN) trained to clone the
 44 behavior of the MPC. This abstraction allows us to use reachability analysis
 45 tools such as Verisig [14] to overapproximate the regret value of the abstracted
 46 controller. As usual with CEGAR approaches, the refinement step is the main
 47 challenge: If the regret is deemed too high (and Verisig finds a real example of
 48 this), then this might be due to our abstraction of the controller as a DNN,
 49 the overapproximation incurred by the reachability tool, or it might be a real
 50 problem with the MPC. In our proposal, when we cannot match the high-regret
 51 example to a behavior of the MPC, we use the output of the reachability analysis
 52 tool to augment the dataset used for training the DNN.

53 As a final contribution, we report on a prototype implementation of our CE-
 54 GAR framework using Verisig. We have used this prototype to analyze MPC for
 55 two well-known control problems. While the approach is promising, we conclude
 56 that further tooling support is required for the full automation of the framework.

57 *Related work.* Chen et al. 2022 [5] conducted a survey on recent advancements in
 58 verifying cyber-physical systems and identified as understudied the verification
 59 of control systems whose performance is measured using cost functions. Indeed,
 60 we did not find many works on the verification of controllers with respect to the
 61 cost functions used to obtain them from an OCP instance. Further, to the best
 62 of our knowledge, there have been no previous works on the formal analysis of
 63 regret in hybrid systems. A notable exception is the recent work of Muvvala et
 64 al. [18] who propose regret minimization as a less pessimistic objective for robots
 65 involved in collaborations (e.g., with humans), as opposed to a sole emphasis
 66 on worst-case optimization. However, their regret analysis focuses on a higher
 67 planning level, distinct from the hybrid-dynamics level of the system, making it
 68 closer to the work of Hunter et al. [13] rather than the present one.

69 Behavioral cloning, also known as imitation learning, is a topic of increasing
 70 interest within artificial intelligence (see, e.g. [3,19,20]). We do not claim to
 71 have a new behavioral cloning algorithm. Rather, we have integrated a data
 72 aggregation step into our CEGAR algorithm for the competitive analysis of
 73 hybrid automata. Interestingly, contrary to previous uses of DNNs as proxies for
 74 MPC [6,14], we have observed that a successful competitive analysis (i.e., the
 75 tool says the controller is r -competitive for a small enough r) suggests one can
 76 use the DNN instead of the MPC! Although this does not guarantee that the
 77 MPC itself is r -competitive, the DNN demonstrates competitiveness. Moreover,
 78 evaluating the DNN to compute the control law proves to be relatively efficient.

79 2 Hybrid Automata and Competitive Analysis

80 A *hybrid automaton* (HA, for short) is an extension of a finite-state automaton
 81 equipped with a finite set of real-valued variables. The values of the variables

82 change *discretely* along the transitions and they do so *continuously*, over time
83 while staying in a state.

84 Formally a HA is a tuple $(Q, I, T, \Sigma, X, \text{jump}, \text{flow}, \text{inv})$, where:

- 85 – Q is a finite set of states and $I \subseteq Q$ is the subset of initial states,
- 86 – Σ is a finite alphabet,
- 87 – $T \subseteq Q \times \Sigma \times Q$ is a set of transitions, and
- 88 – X is a finite set of real-valued variables. We write $V \subseteq \mathbb{R}^X$ to denote the set
89 of all possible valuations of X .
- 90 – $\text{jump} : T \rightarrow \text{Op}$ maps transitions to a set of *guards* and *effects* on the values
91 of the variables. That is, $\text{Op} \subseteq V^2$ and for a transition $\delta \in T$, $\text{jump}(\delta) =$
92 $(\text{guard}, \text{effect})$ implies that δ is “enabled” if the current valuation is *guard*
93 and *effect* is the valuation after the transition. Intuitively, *jump* denotes the
94 discrete changes in the variables along transitions. Usually, the guards and
95 effects are encoded as first-order predicates over the reals, e.g. $\text{jump}(\delta) =$
96 $(x > 2, x + 4)$ denotes the set $\{(v, v') \in V^2 \mid v(x) > 2 \text{ and } v'(x) = v(x) + 4\}$.
- 97 – $\text{flow} : Q \rightarrow F$, with $F \subseteq \{f : \mathbb{R}_{>0} \rightarrow V\}$, maps each state $q \in Q$ to a set F of
98 functions f_q that give the continuous change in the valuation of the variables
99 while in state q . Usually, the functions f_q are encoded as systems of first-
100 order differential equations, e.g. $\dot{x} = 5$ denotes functions¹ $f(t)(x) = 5t + c$,
101 where $c \in \mathbb{R}_{>0}$ is the value of x at time $t = 0$.
- 102 – $\text{inv} : Q \rightarrow 2^V$ maps each state $q \in Q$ to an invariant that constrains the
103 possible valuations of the variables in q . Similar to *jump*, *inv* is usually
104 encoded as first-order predicates over the reals.

105 *Configurations and runs.* A configuration is a pair (q, v) where $q \in Q$ and $v \in V$
106 is a valuation of the variables in X . A configuration (q, v) is *valid* if $v \in \text{inv}(q)$.
107 Let (q, v) and (q', v') be two valid configurations. We say (q', v') is a *discrete*
108 *successor* of (q, v) if $\delta = (q, a, q') \in T$ for some $a \in \Sigma$ and $(v, v') \in \text{jump}(\delta)$.
109 Similarly, (q', v') is a *continuous successor* of (q, v) if $q = q'$ and there exist
110 $t_0, t_1 \in \mathbb{R}_{>0}$ and $f_q \in \text{flow}(q)$ such that $f_q(t_0) = v, f_q(t_1) = v'$ and for all
111 $t_0 \leq t \leq t_1, f_q(t) \in \text{inv}(q)$.

112 A *run* ρ is a sequence of configurations $(q_0, v_0)(q_1, v_1) \dots (q_n, v_n)$ such that
113 $q_0 \in I, v_0$ assigns 0 to all variables and, for all $0 \leq i < n, (q_{i+1}, v_{i+1})$ is a discrete
114 or continuous successor of (q_i, v_i) . The REACH decision problem asks, for given
115 A and (q, v) , whether there is a run of A whose last configuration is (q, v) .

116 *Parallel composition.* Let $A_i = (Q_i, I_i, T_i, \Sigma_i, X_i, \text{jump}_i, \text{flow}_i, \text{inv}_i)$ for $i = 1, 2$
117 be two HA. Then, $A = (Q, I, T, \Sigma, X, \text{jump}, \text{flow}, \text{inv})$ is the *parallel composition*
118 of A_1 and A_2 , written $A = A_1 \parallel A_2$, if and only if:

- 119 – $Q = Q_1 \times Q_2$ and $I = I_1 \times I_2$,
- 120 – $\Sigma = \Sigma_1 \cup \Sigma_2$ and $X = X_1 \cup X_2$.
- 121 – The transition set T contains $(\langle q_1, q_2 \rangle, \sigma, \langle q'_1, q'_2 \rangle)$ if and only if there are
122 $i, j \in \{1, 2\}$ such that $i \neq j$ and:

¹ Note that if X contains more variables than just x , this function is not unique.

- 123 • either $\sigma \in \Sigma_i \setminus \Sigma_j$, $(q_i, \sigma, q'_i) \in T_i$, and $q_j = q'_j$;
- 124 • or $\sigma \in \Sigma_i \cap \Sigma_j$, $(q_i, \sigma, q'_i) \in T_i$, and $(q_j, \sigma, q'_j) \in T_j$.
- 125 – The *jump* function is such that, for $\delta = (\langle q_1, q_2 \rangle, \sigma, \langle q'_1, q'_2 \rangle)$, we have that:
- 126 • either $\sigma \in \Sigma_i \setminus \Sigma_j$ and $jump(\delta) = jump_i(\langle q_i, \sigma, q'_i \rangle)$ for some $i, j \in \{1, 2\}$
- 127 with $i \neq j$,
- 128 • or $\sigma \in \Sigma_i \cap \Sigma_j$ and $jump(t) = jump_1(\langle q_1, \sigma, q'_1 \rangle) \cap jump_2(\langle q_2, \sigma, q'_2 \rangle)$.
- 129 – Finally, $flow(\langle q_1, q_2 \rangle) = flow_1(q_1) \cap flow_2(q_2)$, and
- 130 – $inv(\langle q_1, q_2 \rangle) = inv_1(q_1) \cap inv_2(q_2)$.

131 2.1 The cost of control

132 In this work, we use HA to model *hybrid systems* and *controllers*. In particular,
 133 we henceforth assume any HA $A = (Q, I, T, \Sigma, X, jump, flow, inv)$ modelling a
 134 hybrid system has a designated *cost* variable $J \in X$. We make no such assump-
 135 tion for HA used to model controllers. Observe that from the definition of parallel
 136 composition, it follows that if A models a hybrid system, then $B = A \parallel C$ also
 137 models a hybrid system — i.e. it has the cost variable J — for any HA C .

138 The following notation will be convenient: For a run $\rho = \dots(q_n, v_n)$ we write
 139 J_ρ to denote the value $v_n(J)$. Further, we write $\rho \in A$, where ρ is a run of the
 140 hybrid automaton A . Now, the *maximal and minimal cost of a HA* A respectively
 141 are $\overline{J}(A) := \sup_{\rho \in A} J_\rho$ and $\underline{J}(A) := \inf_{\rho \in A} J_\rho$.

142 2.2 Regret

143 Fix a hybrid-system HA $A = (Q, I, T, \Sigma, X, jump, flow, inv)$. We define the
 144 (*worst-case*) *regret* $Reg(U)$ of a controller HA U as the maximal difference be-
 145 tween the (maximal) cost incurred by the parallel composition of A and U — i.e.
 146 the controlled system — and the (minimal) cost incurred by an alternative con-
 147 troller HA from a set \mathcal{C} : $Reg(U) := \sup_{U' \in \mathcal{C}} (J(A \parallel U) - J(A \parallel U'))$. The REGRET
 148 problem asks, for given A, U, \mathcal{C} , and $r \in \mathbb{Q}$, whether $\overline{Reg}(U) \geq r$.

149 3 Reachability and Competitive Analysis

150 In this section, we establish that the reachability and regret problems are interre-
 151 ducible. While this implies an exact algorithm for the competitive analysis of
 152 hybrid automata does not exist, it suggests the use of approximation algorithms
 153 for reachability as a means to realize an approximate analysis.

154 **Theorem 1.** *Let \mathcal{C} be the set of all possible controllers. Then, the REGRET*
 155 *problem reduces in polynomial time to the REACH problem.*

156 *Proof (of Theorem 1).* Given a hybrid-system HA A , a controller U , a set of all
 157 possible controllers \mathcal{C} and a regret bound $r \in \mathbb{Q}$, we will construct another HA
 158 $A' = (Q', I', T', \Sigma, X', jump', flow', inv')$ and a target configuration (q', v) of A'
 159 such that, (q', v) is reachable in A' if and only if $Reg(U) < r$ in $A \parallel U$. Let us
 160 write $A = (Q, I, T, \Sigma, X, jump, flow, inv)$ and note that $J \in X$ because A is a




Fig. 1. Gadget for simulating any possible controller

161 hybrid-system HA. We extend the automaton $A \parallel U$ with a gadget to obtain A' .
 162 The idea is as follows: for every variable $y \in Y$ of $A \parallel U$, we add a copy of it
 163 in the variable set X' of A' that simulates any possible choice of value for that
 164 variable by an alternative controller U' . The variable $J' \in X'$ calculates the cost
 165 of that alternative controller. Formally, $X' = Y \cup \{y' \mid y \in Y\}$.

166 To simulate any possible valuation of the variables, we introduce the gadget
 167 given in Figure 1. For every variable x'_i such that x_i is a variable in U , the gadget
 168 contains two states $q_{x'_i}^+$ and $q_{x'_i}^-$. Then, $flow'(q_{x'_i}^+)$ contains $\dot{x}'_i = 1$ and $\dot{x}'_j = 0$ for
 169 all $j \neq i$. Intuitively, this state allows us to positively update the value of x'_i to
 170 any arbitrary value. Similarly, $flow'(q_{x'_i}^-)$ contains $\dot{x}'_i = -1$ and $\dot{x}'_j = 0$, $\forall j \neq i$,
 171 which allows it to negatively update the value of x'_i .

172 Now, we add a “sink” state q_{reach} and make it reachable from all the other
 173 states using transitions $\delta'_i \in T'$ such that $jump'(\delta'_i)$ contains guard of the form
 174 $J - J' \geq r$. Finally, from every state $q' \in Q'$, we add the option to go into its
 175 own copy of the gadget, set the values of the variables to any desired value and
 176 come back to the same state.

Note that if $(q_{\text{reach}}, \mathbf{0})$ is reachable in A' , via a run $\rho \in A'$, then $J_\rho - J'_\rho \geq r$.
 As the gadget does not update the value of J and J' , it is easy to see that
 $Reg(U) \geq r$. Now, if $(q_{\text{reach}}, \mathbf{0})$ is not reachable that means, $J_\rho - J'_\rho < r$ for all
 $\rho \in A'$. Now, as all possible controllers (in fact, all possible configurations of
 variables from U) can be simulated in A' , it is easy to see that $Reg(U) < r$. \square

177 Interestingly, the construction presented above does not preserve the prop-
 178 erty of being *initialized*. Intuitively, an initialized hybrid automaton is one that
 179 “resets” a variable x on transitions between states which have different flows for
 180 x . Alas, we do not know whether an alternative proof exists which does preserve
 181 the property of being initialized (and also being rectangular, a property which we
 182 do not formally define here). Such a reduction would imply the regret problem
 183 is decidable for rectangular and initialized hybrid automata.

184 We now proceed to stating and proving the converse reduction.

185 **Theorem 2.** *The REACH problem reduces in polynomial time to the REGRET*
 186 *problem.*

187 *Proof (of Theorem 2).* Given a HA A and a target configuration (q, v) , we will
 188 construct a HA A' and a controller U such that $Reg(U) \geq 2$ with respect to
 189 $A' \parallel U$ if and only if (q, v) is reachable in A . The reduction works for any set \mathcal{C}
 190 of controllers that contains at least one controller that sets c to 0 all the time.




Fig. 2. Reduction from REACH to REGRET

191 First, we add two states to A' so that $Q' = Q \cup \{q_i, q_T\}$. In A' , q_i has a
 192 self-loop that can be taken if the value of c is 0 and the effect is that $J = 0$.
 193 From q_i , we can also transition to the initial states of A if $c \geq 1$, and in doing
 194 so, we set J to 1. Finally, from the target state q in A , we can go to the new
 195 state q_T if the target valuation \mathbf{v} is reached, and that changes the valuation of
 196 J to 2. The valuation of J does not change within A .

Note that the minimum cost incurred by a controller that constantly sets c to 0 in A' is 0, which is achieved by the run that loops on q_i . Now, if (q, \mathbf{v}) is reachable in A via run $\rho \in A$, then the maximum cost incurred by a controller that sets c to 1 occurs along a run $q_i \cdot \rho \cdot q_T$ and is 2, making $Reg(U) \geq 2$. On the other hand, if (q, \mathbf{v}) is not reachable in A , then the maximal value of J along any such run is 1, resulting in $Reg(U) < 2$. Our constructed controller U is such that it sets c to 1 all the time, and the above arguments give the desired result. \square

197 Since the reachability problem is known to be undecidable for hybrid au-
 198 tomata in general [11], it follows that our regret problem is also undecidable.
 199

200 **Corollary 1.** *The REGRET problem is undecidable.*

201 4 CEGAR-Based Competitive Analysis

202 We present our CEGAR approach to realize approximate competitive analysis.
 203 To keep the discussion simple, we focus on continuous systems, specifically single-
 204 state hybrid automata. Since our goal is to approximate the regret of MPCs, we
 205 model controllers as hybrid automata that sample variable values at discrete-
 206 time intervals and determine control variable values using a deep neural network
 207 (DNN) trained to behave as the MPC. Concretely, our approach specializes the
 208 reduction in the proof of Theorem 1: We will work with a hybrid automaton
 209 \mathcal{D} that abstracts the behavior of the controller using a DNN, and a hybrid
 210 automaton \mathcal{N} that abstracts the behaviors of all alternative controllers. The
 211 overview of our framework is depicted in Figure 3a. In section 5, we present a
 212 toolchain implementing this CEGAR-based approach.

213 4.1 Initial abstraction and analysis

214 Our proposed framework begins with the abstraction of the controller as a hybrid
 215 automaton \mathcal{D} and the alternative controllers as \mathcal{N} . Each of these automata are
 216 assumed to have a cost variable, say $J_{\mathcal{D}}$ for \mathcal{D} and $J_{\mathcal{N}}$ for \mathcal{N} . For a given

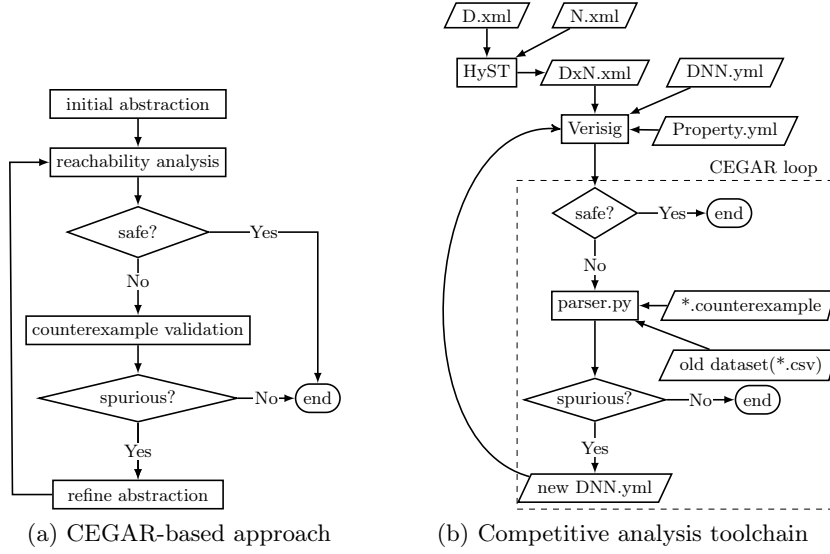


Fig. 3. Flowchart depictions of our approach and our toolchain implementing it; We use ANSI/ISO standard flowchart symbols: the parallelogram blocks represent inputs/outputs, and the rectangular blocks represent processes or tools

217 value $r \in \mathbb{R}$, if we want to determine whether \mathcal{D} is r -competitive then we add to
 218 $\mathcal{A} = \mathcal{D} \parallel \mathcal{N}$ a new cost variable $J = J_{\mathcal{N}} - J_{\mathcal{D}}$. As is argued in Theorem 1, \mathcal{D} should
 219 be r -competitive if and only if \mathcal{A} can reach a configuration where the value of J
 220 is larger than r . Hence, we can apply any reachability set (overapproximation)
 221 tool to determine the feasibility of such a configuration.

222 4.2 Reachability status

223 If the reachability tool finds that a configuration with $J \geq r$ is reachable in
 224 \mathcal{A} , we say it concludes \mathcal{A} is *unsafe*. In that case, we will have to process the
 225 reachability witness. Otherwise, \mathcal{A} is *safe*, and we can stop and conclude that \mathcal{D} is
 226 r -competitive. Interestingly, \mathcal{D} can now be used as an r -competitive replacement
 227 of the original controller! It is important to highlight that behavior cloning does
 228 not provide any guarantees regarding the relationship between the MPC and
 229 the DNN within \mathcal{D} . Consequently, even if we have evidence supporting the r -
 230 competitiveness of \mathcal{D} , we cannot infer the same for the MPC itself.

231 In the context of MPCs, this result is already quite useful. This is because
 232 MPCs have a non-trivial *latency* and memory usage before choosing a next valua-
 233 tion for the control variables (see, e.g. [12,15]). In our implementation described
 234 in the following section, \mathcal{D} takes the form of a DNN. As DNNs can be evaluated
 235 rather efficiently, using the DNN instead of the original MPC is desirable.

236 4.3 Counterexample analysis and refinement

237 If \mathcal{A} is deemed unsafe, we expect the reachability tool to output a counterex-
 238 ample in the form of a run. There is one main reason why such a run could
 239 be *spurious*, i.e. it does not correspond to a witness of the MPC not being
 240 r -competitive. Namely, the abstractions \mathcal{D} (representing the MPC) or \mathcal{N} (rep-
 241 resenting alternative controllers) might be too coarse. For the specific case of \mathcal{D} ,
 242 where a DNN is used to model the MPC, we describe sufficient conditions to
 243 determine if the counterexample is indeed spurious. If the counterexample is in-
 244 deed deemed spurious, we can refine our abstraction by incorporating new data
 245 obtained from the counterexample and retraining the DNN. In general, though,
 246 refining \mathcal{D} and \mathcal{N} falls into one of the tasks for which our framework does not
 247 rely on automation.

248 4.4 Human in the loop

249 There are three points in the framework, where human intervention is needed.

250 *Modelling and specification.* First, the task of obtaining initial abstractions \mathcal{D}
 251 and \mathcal{N} of the controller and all alternative controllers, respectively, does require a
 252 human in the loop. Indeed, crafting hybrid automata is not something we expect
 253 from every control engineer. In our prototype described in the next section, we
 254 mention partial support for obtaining \mathcal{D} and \mathcal{N} automatically when the MPC is
 255 given in the language of a particular OCP and optimization library.

256 *Reachability analysis.* Second, reachability being an undecidable problem, most
 257 reachability analysis tools can not only output safe and unsafe as results. Ad-
 258 ditionally, they might output an “unknown” status. In this case, revisiting the
 259 abstractions \mathcal{D} and \mathcal{N} , or even changing the options with which the tool is be-
 260 ing used may require human intervention. In fact, we see this as an additional
 261 abstraction-refinement step which is considerably harder to automate since there
 262 is an absence of a counterexample to work with.

263 *Abstraction refinement.* Finally, our framework does not say what to do if the
 264 counterexample being spurious is due to \mathcal{N} being too coarse an approximation.
 265 This scenario can occur when \mathcal{N} is purposefully modeled to discretize or ap-
 266 proximate certain behaviors of alternative controllers to facilitate reachability
 267 analysis. However, for \mathcal{D} , we offer automation support by proposing the retrain-
 268 ing of our DNN in the implementation. It might actually be needed to change
 269 the architecture of the DNN to obtain a better abstraction. This process can be
 270 automated, as increasing the number of layers is often sufficient according to the
 271 universal approximation theorem [4].

272 5 Implementation and Evaluation

273 We now present our implementation of the CEGAR-based competitive analysis
 274 method presented in the previous section, along with two case studies used for
 275 evaluation: the cart pendulum and an instance of motion planning.

276 5.1 Competitive analysis toolchain

277 Figure 3b gives a visual depiction of the toolchain in the form of a flowchart.
 278 Starting from the top, `D.xml`, `N.xml` are XML files encoding hybrid automata
 279 \mathcal{D} and \mathcal{N} , respectively, in the SpaceEx modeling language[8]. The automaton \mathcal{D}
 280 represents the controller, which could be a model predictive controller (MPC),
 281 and \mathcal{N} represents a class of controllers that the MPC is compared against —
 282 see also Section 4.1. We use the HyST [2] translation tool for hybrid automata
 283 to generate the parallel composition $\mathcal{D} \parallel \mathcal{N}$ (encoded in `DxN.xml`, again in the
 284 SpaceEx language). The composed automaton, along with the trained DNN and
 285 the property to be verified, are fed as inputs to Verisig. Verisig [14] is a tool
 286 that verifies the safety properties of closed-loop systems with neural network
 287 components. The tool takes a hybrid automaton, a trained neural network, and
 288 property specification files as inputs. It performs the reachability analysis and
 289 provides safety verification result. We then parse the output of Verisig to deter-
 290 mine whether \mathcal{D} is competitive enough (`parser.py`). If this is not the case, we
 291 realize a sound check to determine if the counterexample is spurious, in which
 292 case we use it to extend our dataset and further train the DNN.

293 5.2 Initial abstraction and training

294 Our toolchain is finetuned to work well for hybrid systems modeled in a tool
 295 called *Rockit* and MPCs obtained using the same tool. Rockit, which stands for
 296 Rapid Optimal Control Kit, is a tool designed to facilitate the rapid prototyp-
 297 ing of optimal control problems, including iterative learning, model predictive
 298 control, system identification, and motion planning [10].

299 Our toolchain includes a utility that interfaces with the API of Rockit to
 300 automatically generate the hybrid automata \mathcal{D} and \mathcal{N} from a model of a control
 301 problem. While the use of Rockit is convenient, it is not required by our toolchain.

302 Based on a dataset (in our examples, we obtain it from Rockit), we train
 303 a DNN using *behavioral cloning*: we try to learn the behavior of an *expert* (in
 304 our case, the MPC) and replicate it. For this, we make use of the *Dagger algo-*
 305 *rithm*[20], which, after an initial round of training on the dataset from Rockit,
 306 will simulate traces using the DNN. The points that the neural network visits
 307 along these traces are then given to the expert, and the output of the expert is
 308 recorded. These new points and outputs are appended to the first dataset, and
 309 this new dataset is used to train a second DNN. This iterative process is done
 310 multiple times to make the DNN more robust. In all of our experiments, the
 311 TensorFlow framework [1] was used for the creation and training of the DNN.

312 5.3 Reachability status

313 The regret property, encoded as a reachability property as is done in the proof
 314 of Theorem 1, is specified in the property file `Property.yml`, which also includes
 315 the initial states of $\mathcal{D} \parallel \mathcal{N}$. Verisig provides three possible results: “safe” if no
 316 property violation is found, “unsafe” if there is a violation, and “unknown” if the

317 property could not be verified, potentially due to a significant approximation
 318 error. In the latter two cases, a counterexample file (CE file) is generated.

319 5.4 Counterexample analysis and retraining

320 If the result is “unsafe”, the next step is to compare the counterexample tra-
 321 jectory against the dataset generated from the controller code. If a matching
 322 trajectory is found, it indicates a real counterexample, meaning that this tra-
 323 jectory could potentially occur in the actual controller, and no further action is
 324 required. If a matching trajectory is not found, then it is a spurious counterex-
 325 ample that requires either retraining the DNN or fix(es) in $\mathcal{D} \parallel \mathcal{N}$. Our toolchain
 326 automatically validates the counterexample by comparing the trajectories from
 327 Verisig and the controller as implemented in Rockit. To do so, since Rockit uses
 328 the floating-point representation of real numbers, we choose a decimal precision
 329 of $\epsilon = 10^{-3}$ for the comparison. In the case of a spurious counterexample that
 330 requires retraining the DNN, we update the existing dataset using Rockit to
 331 obtain additional labeled data based on the trajectory from the CE file.

332 The CE file from Verisig represents state variable values using interval arith-
 333 metic, while the controller dataset contains state variable values in \mathbb{R} without
 334 intervals. To accommodate this difference, we choose to append to the dataset
 335 new entries: (a) the lower bounds of input intervals, (b) the upper bounds, and
 336 (c) a range² of intermediate input values within the intervals. For each of these,
 337 we also include the corresponding controller outputs. The generation of the up-
 338 dated dataset and the retraining of the DNN are performed automatically by our
 339 toolchain. A DNN trained on the new dataset is then fed to Verisig again along
 340 with `DxN.xml` and the `Property.yml`. This way, the CEGAR loop is repeated
 341 until one of the following conditions is true: (a) the counterexample is real, or (b)
 342 a maximum number of retraining iterations (determined by the user) is reached.

343 5.5 Experiments

344 In the sequel, we use our tool to analyze two control problems that have been
 345 implemented using the Rockit framework. The research questions we want to
 346 answer with the forthcoming empirical study are the following.

347 **RQ1** Can we have a fully automated tool to perform the competitive analysis?

348 **RQ2** Is the toolchain scalable? Why or why not?

349 **RQ3** Does the approach help to improve confidence in (finite-horizon) competi-
 350 tivity of controllers?

351 **RQ4** Does the approach help find bugs in controller design?

352 We now briefly introduce the two case studies, their dynamics, and how each
 353 of them are modeled so that our toolchain can be used to analyze them.

² Our toolchain splits each interval into n equally large segments and adds all points
 in the resulting lattice. In our experiments, we use $n = 4$.

354 **Cart pendulum.** The cart pendulum problem is a classic challenge in control
 355 theory and dynamics [7]. In it, an inverted pendulum is mounted on a cart
 356 that can move horizontally via an electronic servo system. The objective is to
 357 minimize a cost $J = F^2 + 100 * pos^2$, where F represents the force applied to
 358 the cart and pos indicates the position of the cart. The values of F and pos are
 359 constrained within the range of $[-2, 2]$. The dynamics of the cart correspond to
 360 the physics of the system and make use of parameters including the mass of the
 361 cart and the pendulum and the length of the pendulum (see Appendix A).

362 While the proof of Theorem 1 provides a sound way to model all alternative
 363 controllers in the form of \mathcal{N} , the construction combines continuous dynamics
 364 and non-determinism. Current hybrid automata tools do not handle non-trivial
 365 combinations of these two elements very well. Hence, we have opted to discretize
 366 the choice of control values for alternative controllers. Intuitively, this means
 367 that every time the DNN is asked for new control variable values in \mathcal{D} , the
 368 automaton \mathcal{N} non-deterministically chooses new alternative values from a finite
 369 subset fixed by us a priori (see Figure 5b in appendix for an example).

370 **Motion planning** The case study involves computing a series of actions to
 371 move an object from one point to another while satisfying specific constraints
 372 [16]. In our case study, an MPC is used to plan the motion of an autonomous
 373 bicycle that is expected to move along a curved path on a 2D plane using a
 374 predefined set of waypoints. To prevent high-speed and skidding, the velocity
 375 (V) and the turning rate (δ , in radians) are constrained in the ranges $0 \leq V \leq 1$
 376 and $-\pi/6 \leq \delta \leq \pi/6$. The objective is to minimize the sum of squared estimate
 377 of errors between the actual path taken by the bicycle and the reference path.
 378 Intuitively, the more the controller deviates from the reference path, the higher
 379 its cost (see Appendix B).

380 Like in the cart pendulum case study, we discretize the alternative control
 381 variable valuations. A big difference is that the cost has both a *Mayer term* and
 382 a *Lagrangian* that depend on the location of the bicycle and the waypoints in an
 383 intricate way. In terms of modelling, this means that \mathcal{D} and \mathcal{N} have to “compute”
 384 closest waypoints relative to the current position of the bicycle (see Figure 6).

385 **Discussion.** Towards an answer for **RQ1**, we can say that while our toolchain³
 386 somewhat automates our CEGAR, it still requires manual work (e.g. the initial
 387 training and choice of DNN architecture). Moreover, in the described case stud-
 388 ies, we did not observe an MPC DNN that is labeled as competitive. This may be
 389 due to (over)approximations incurred by our framework and our use of Verisig.
 390 Despite this, we can answer **RQ4** positively as our toolchain allowed us to spot
 391 a bug hidden in the Rokit MPC solution for the cart pendulum. We observed in
 392 early experiments that the MPC was not competitive and short (run) examples
 393 of this were quickly found by Verisig. We then found that the objective function
 394 in Rokit was indeed not as intended by the developers.

³ All graphs and numbers can be reproduced using scripts from: <https://github.com/competitive-analysis-toolchain/competitive-analysis>.

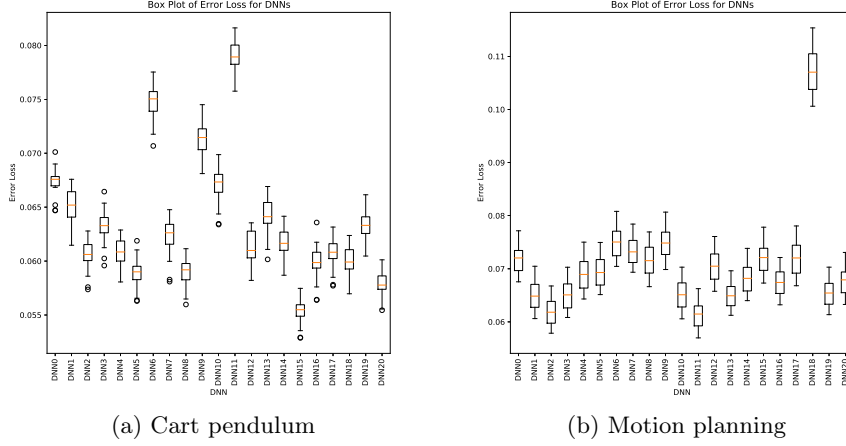


Fig. 4. Boxplots showing the training losses of all DNNs against all test sets

395 The DNNs do show a trend towards copying the behavior of the MPC (see
 396 Figure 4) even though we retrain a new DNN from scratch after each (spurious)
 397 counterexample obtained via Verisig and we (purposefully) randomize the choice
 398 of test and training set in each iteration. We do this to increase variability in
 399 the set of behaviors and the counterexamples used to extend the dataset. In
 400 the cart pendulum case study, we observe that in the iterations 2, 7, and 11,
 401 the number of discrete time steps during which the corresponding DNN can act
 402 while remaining competitive is larger than in the initial iteration. Hence, for
 403 **RQ3**, we conclude our toolchain can indeed help increase reliability in the DNN
 404 proxy being competitive, albeit only for a finite horizon. On the negative side,
 405 experiments for 20 iterations of retraining from spurious counterexamples take
 406 more than 90min in both our case studies. This leads us to conclude that our
 407 toolchain does not yet scale as required for industrial-size case studies (**RQ2**).

408 6 Conclusion

409 Based on our theoretical developments to link the regret problem with the clas-
 410 sical reachability problem, we proposed a CEGAR-based approach to realize the
 411 competitive analysis of MPCs via neural networks as proxies. We also presented
 412 an early proof-of-concept implementation of the approach. Now that we have a
 413 baseline, we strongly believe improvements in the form of algorithms and dedi-
 414 cated tools will allow us to improve our framework to the point where it scales
 415 for interesting classes of hybrid systems.

416 **References**

- 417 1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghe-
418 mawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S.,
419 Murray, D.G., Steiner, B., Tucker, P.A., Vasudevan, V., Warden, P., Wicke, M., Yu,
420 Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: Keeton,
421 K., Roscoe, T. (eds.) 12th USENIX Symposium on Operating Systems Design and
422 Implementation, OSDI 2016, Savannah, GA, USA, November 2-4, 2016. pp. 265–
423 283. USENIX Association (2016), [https://www.usenix.org/conference/osdi16/
424 technical-sessions/presentation/abadi](https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi)
- 425 2. Bak, S., Bogomolov, S., Johnson, T.T.: Hyst: a source transformation and transla-
426 tion tool for hybrid automaton models. In: Proceedings of the 18th International
427 Conference on Hybrid Systems: Computation and Control. pp. 128–133 (2015)
- 428 3. Bratko, I., Urbančič, T., Sammut, C.: Behavioural cloning: phenomena, results and
429 problems. IFAC Proceedings Volumes **28**(21), 143–149 (1995)
- 430 4. Chen, T., Chen, H.: Universal approximation to nonlinear operators by neural net-
431 works with arbitrary activation functions and its application to dynamical systems.
432 IEEE Transactions on Neural Networks **6**(4), 911–917 (1995)
- 433 5. Chen, X., Sankaranarayanan, S.: Reachability analysis for cyber-physical systems:
434 Are we there yet? In: NASA Formal Methods Symposium. pp. 109–130. Springer
435 (2022)
- 436 6. Clavière, A., Dutta, S., Sankaranarayanan, S.: Trajectory tracking control for
437 robotic vehicles using counterexample guided training of neural networks. In: Ben-
438 ton, J., Lipovetzky, N., Onaindia, E., Smith, D.E., Srivastava, S. (eds.) Proceed-
439 ings of the Twenty-Ninth International Conference on Automated Planning and
440 Scheduling, ICAPS 2018, Berkeley, CA, USA, July 11-15, 2019. pp. 680–688. AAAI
441 Press (2019), <https://ojs.aaai.org/index.php/ICAPS/article/view/3555>
- 442 7. Fantoni, I., Lozano, R., Lozano, R.: Non-linear control for underactuated mechan-
443 ical systems. Springer Science & Business Media (2002)
- 444 8. Frehse, G., Le Guernic, C., Donzé, A., Cotton, S., Ray, R., Lebeltel, O., Ripado, R.,
445 Girard, A., Dang, T., Maler, O.: Spaceex: Scalable verification of hybrid systems.
446 In: Computer Aided Verification: 23rd International Conference, CAV 2011, Snow-
447 bird, UT, USA, July 14-20, 2011. Proceedings 23. pp. 379–395. Springer (2011)
- 448 9. Gillis, J.: rockit. [https://gitlab.kuleuven.be/meco-software/rockit/-/blob/
449 master/examples/motion_planning_MPC.py](https://gitlab.kuleuven.be/meco-software/rockit/-/blob/master/examples/motion_planning_MPC.py) (2020)
- 450 10. Gillis, J., Vandewal, B., Pipeleers, G., Swevers, J.: Effortless modeling of optimal
451 control problems with rockit. In: 39th Benelux Meeting on Systems and Control,
452 Date: 2020/03/10-2020/03/12, Location: Elspeet, The Netherlands (2020)
- 453 11. Henzinger, T.A., Kopke, P.W., Puri, A., Varaiya, P.: What’s decidable
454 about hybrid automata? J. Comput. Syst. Sci. **57**(1), 94–124 (1998).
455 [https://doi.org/10.1006/jcss.1998.
456 1581](https://doi.org/10.1006/jcss.1998.1581)
- 457 12. Hertneck, M., Köhler, J., Trimpe, S., Allgöwer, F.: Learning an approximate
458 model predictive controller with guarantees. IEEE Control. Syst. Lett. **2**(3), 543–
459 548 (2018). <https://doi.org/10.1109/LCSYS.2018.2843682>, [https://doi.org/10.
460 1109/LCSYS.2018.2843682](https://doi.org/10.1109/LCSYS.2018.2843682)
- 461 13. Hunter, P., Pérez, G.A., Raskin, J.: Reactive synthesis without regret. Acta In-
462 formatica **54**(1), 3–39 (2017). <https://doi.org/10.1007/s00236-016-0268-z>, <https://doi.org/10.1007/s00236-016-0268-z>
463

- 464 14. Ivanov, R., Carpenter, T., Weimer, J., Alur, R., Pappas, G., Lee, I.: Verisig 2.0:
465 Verification of neural network controllers using taylor model preconditioning. In:
466 International Conference on Computer Aided Verification. pp. 249–262. Springer
467 (2021)
- 468 15. Julian, K.D., Lopez, J., Brush, J.S., Owen, M.P., Kochenderfer, M.J.: Policy com-
469 pression for aircraft collision avoidance systems. In: 2016 IEEE/AIAA 35th Digital
470 Avionics Systems Conference (DASC). pp. 1–10. IEEE (2016)
- 471 16. LaValle, S.M.: Planning algorithms. Cambridge university press (2006)
- 472 17. Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., Talwalkar, A.: Hyperband: A
473 novel bandit-based approach to hyperparameter optimization. *Journal of Machine
474 Learning Research* **18**(185), 1–52 (2018), [http://jmlr.org/papers/v18/16-558.
475 html](http://jmlr.org/papers/v18/16-558.html)
- 476 18. Muvvala, K., Amorese, P., Lahijanian, M.: Let’s collaborate: Regret-based reactive
477 synthesis for robotic manipulation. In: 2022 International Conference on Robotics
478 and Automation, ICRA 2022, Philadelphia, PA, USA, May 23-27, 2022. pp. 4340–
479 4346. IEEE (2022). <https://doi.org/10.1109/ICRA46639.2022.9812298>, [https://
480 doi.org/10.1109/ICRA46639.2022.9812298](https://doi.org/10.1109/ICRA46639.2022.9812298)
- 481 19. Ross, S., Bagnell, D.: Efficient reductions for imitation learning. In: Teh, Y.W.,
482 Titterton, D.M. (eds.) *Proceedings of the Thirteenth International Conference
483 on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sar-
484 dinia, Italy, May 13-15, 2010*. *JMLR Proceedings*, vol. 9, pp. 661–668. *JMLR.org*
485 (2010), <http://proceedings.mlr.press/v9/ross10a.html>
- 486 20. Ross, S., Gordon, G.J., Bagnell, D.: A reduction of imitation learning and struc-
487 tured prediction to no-regret online learning. In: Gordon, G.J., Dunson, D.B.,
488 Dudík, M. (eds.) *Proceedings of the Fourteenth International Conference on Ar-
489 tificial Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April
490 11-13, 2011*. *JMLR Proceedings*, vol. 15, pp. 627–635. *JMLR.org* (2011), [http://proceedings.mlr.press/v15/ross11a/ross11a.pdf
491](http://proceedings.mlr.press/v15/ross11a/ross11a.pdf)

492 **A Cart pendulum with physics-based cost**

493 The cart’s dynamics are summarized in Table 1, with parameters including the
 494 cart’s mass (m_{cart}), the pendulum mass (m), the pendulum length (L), and the
 495 gravitational constant (g).

Use case	Initial state	Dynamics	Objective	Control variable	Path constraints	Sample time(s)
cart pendulum	$pos = 0.5$ $\theta = 0$ $\dot{pos} = 0$ $\dot{\theta} = 0$	$\dot{pos} = dpos$ $\dot{\theta} = d\theta$ $dpos = (-mL\sin(\theta)(d\theta)^2 + mg\cos(\theta)\sin(\theta) + F) / (m_{cart} + m - m * (\cos(\theta))^2)$ $d\theta = (-mL\cos(\theta)\sin(\theta)(d\theta)^2 + F\cos(\theta) + (m_{cart} + m)g\sin(\theta)) / (L(m_{cart} + m - m(\cos(\theta))^2))$	$2F + 100pos^2$	Force (F)	$-2 \leq F \leq 2$ $-2 \leq pos \leq 2$	0.04
motion planning	$x = 0$ $y = 10$ $\theta = 0$	$x = V\cos(\theta)$ $y = V\sin(\theta)$ $\theta = V/L\tan(\delta)$	square error between position and reference path	Turning rate (δ)	$0 \leq V \leq 1$ $-\pi/6 \leq \delta \leq \pi/6$	$1.3 \leq t \leq 1.9$

Table 1. The initial states, the dynamics and other control parameters of the cart pendulum and motion planning.

496 To determine the sample time for the MPC, we calculate it as the ratio
 497 of the control horizon (Tf) to the number of control intervals ($Nhor$). In this
 498 case, Tf is set to 2 seconds and $Nhor$ is 50, resulting in a sample time of
 499 $dt = Tf/Nhor = 0.04s$. Additionally, the initial conditions for the system are
 500 specified as $[pos, \theta, dpos, d\theta] = [0.5, 0, 0, 0]$. The controller is said to have found
 501 an optimal solution, respecting the constraints, for which the cost is minimum.
 502 That is, the more the controller deviates from the set constraints for F and pos ,
 503 the more it is "punished", thereby setting a higher value to the cost variable.

504 For the DNN abstracting the MPC, we trained a fully connected model with
 505 4 inputs nodes, 1 output node⁴ and 4 hidden layers. The number of nodes in
 506 the hidden layers as well as the activation function and the learning rate were
 507 chosen using hypertuning. For hypertuning, the hyperband algorithm was used
 508 [17]. The options that the hyperband algorithm had for each hyperparameter
 509 were the following:

- 510 – The amount of nodes in each hidden layer: 16, 32, 48, 64, 80 or 96.
- 511 – The learning rate was sampled between 1e-5 and 1e-1 with a logarithmic sampling.
- 512 – The activation function: sigmoid or a tangent hyperbolic.

514 The sigmoid and tangent hyperbolic activation functions were chosen because
 515 they can be used in Verisig. The factor used in the hyperband algorithm was 3.

⁴ The last layer of the DNN has a shifted and scaled sigmoid function as an activation function. This function restricts the output of the DNN between -2 and 2. This makes sure that the DNN never outputs a parameter that is outside the bounds of the problem. This reduces the number of spurious counterexamples that are found by the CEGAR loop.

516 We begin the modeling process by representing the given controller as automaton \mathcal{D} , depicted in Figure Figure 5a. This automaton consists of three
 517 modes, each suffixed with a ‘D’ (e.g., initD, environmentD), distinguishing them
 518 from the modes of automaton \mathcal{N} . The initial mode, initD, and the DNN mode
 519 do not involve any time elapse. The incoming edges to the DNN mode labeled as
 520 $_f1$, $_f2$, etc., indicate the inputs to the DNN. Verisig uses the input DNN.yml
 521 in place of the DNN mode. The DNN “mimics” the MPC which takes pos , θ ,
 522 $dpos$, and $d\theta$ as inputs and outputs F . The output of the DNN ($_f1$) is set
 523 to the Force variable FD with a scaling factor that corresponds to the sigmoid
 524 activation function.
 525

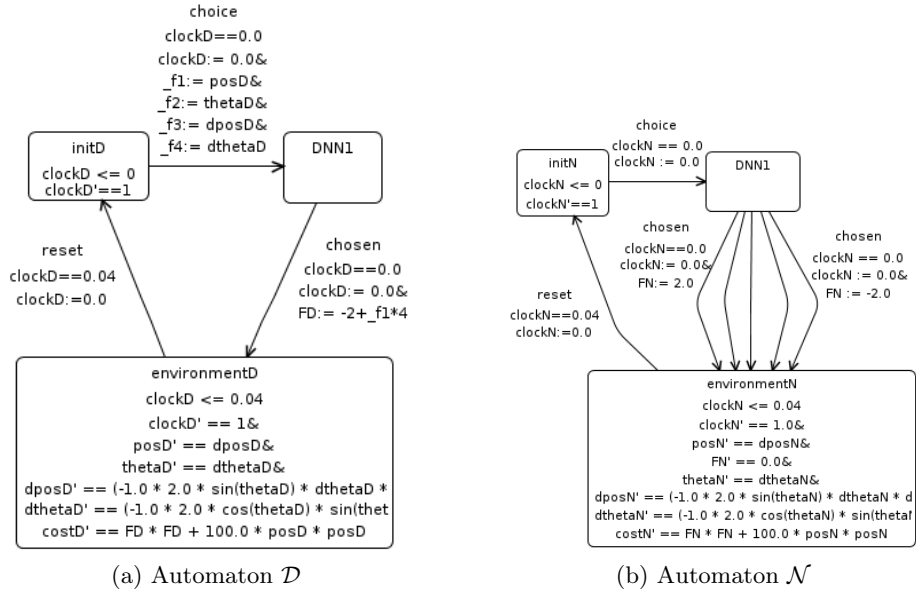


Fig. 5. Automaton \mathcal{D} and automaton \mathcal{N} for cart pendulum

526 The environment mode contains the control dynamics and the cost function,
 527 expressed as ordinary differential equations. To model the controller’s sample
 528 time, we employ a combination of setting the invariant of clockD to 0.04 and
 529 resetting clockD to 0.04 on the outgoing edge from the environment. This ensures
 530 that exactly 0.04 seconds are spent in the environment mode, reflecting the
 531 desired sample time of the controller.

532 Automaton \mathcal{N} models the class of controllers to be compared with MPC (or
 533 alike). Unlike \mathcal{D} , \mathcal{N} does not use a trained DNN – as \mathcal{N} corresponds to a class
 534 of controllers other than the MPC (or alike), there is no need to have a DNN.
 535 So, the DNN mode in \mathcal{N} is just a dummy mode used for sync and composition
 536 generation. Also, unlike \mathcal{D} , \mathcal{N} has discrete behaviour with deterministic
 537 transitions where each transition has F set between -2 and 2 (same constraint as in

538 \mathcal{D}). That is, by discretizing the environment in automaton \mathcal{N} we make the set of
 539 controllers we compare against as finite. Furthermore, the discretization helps in
 540 having a tighter reachability set, thereby reducing verification time. We assume
 541 \mathcal{D} and \mathcal{N} have the same sampling time of 0.04s. Note that the behaviour of \mathcal{D}
 542 and \mathcal{N} is identical, for e.g., a transition in \mathcal{D} from init mode to DNN mode is
 543 identical to the transition from init mode to DNN mode in \mathcal{N} . This ensures a
 544 fair comparison between the MPC and other controllers. Synchronization labels
 545 (choice, chosen, reset) are used on each edge to establish this identical behaviour,
 546 which also facilitates composition. The parallel composition $\mathcal{D}||\mathcal{N}$ generated by
 547 HyST.

548 The next step is to feed the generated composition $\mathcal{D}||\mathcal{N}$ to Verisig along with
 549 the property file and the trained DNN. The competitive analysis is performed for
 550 the regret property $costD - costN \geq 0.25$. That is, we ask Verisig if there is an
 551 ‘unsafe’ trajectory in the composed automaton with the difference in the costs is
 552 greater than 0.25. Verisig returned an ‘unsafe’ result along with a CE file, which
 553 is then automatically parsed to obtain the unsafe trajectory. The trajectory from
 554 Verisig is compared against the controller’s dataset and it is found that it is a
 555 spurious counterexample. As the next step, the dataset file is appended with
 556 new data from the CE file, and the DNN is retrained. The retrained DNN, along
 557 with the composed automaton and the property file is fed to Verisig again. The
 558 DNN retraining (a.k.a the refinement) is repeated for 20 iterations.

559 B Motion planning with waypoint-based cost

560 Motion planning involves computing a series of optimal steps to move an object
 561 from one point to another while satisfying specific constraints [16]. In our case
 562 study, an MPC is used to plan the motion of an autonomous bicycle that is
 563 expected to move on a curved path on a 2D plane using a predefined set of
 564 waypoints. To prevent high speed and skidding, the velocity(v) and the turning
 565 rate(δ) of the bicycle are constrained within the ranges of $0 \leq V \leq 1$ and $-\pi/6 \leq$
 566 $\delta \leq \pi/6$ respectively. The objective is to minimize the sum of squares error
 567 between the actual path taken by the bicycle and the reference path. That is, the
 568 more the controller deviates from the reference path, the more it is “punished”
 569 with a higher value of sum squared error. The bicycle dynamics are summarized
 570 in Table 1, where L represents the bicycle length. Unlike the cart pole case study
 571 where the sample time remains constant, the controller dynamically calculates
 572 the sample time (t_s second) within the range of $1.3 \leq t_s \leq 1.9$. The Python
 573 implementation of the use case is available at [9].

574 To model this case study, four automaton were employed. \mathcal{D} and \mathcal{N} represent
 575 the given controller and the class of controllers for comparison, respectively. Ad-
 576 ditionally, we use two more automaton $\mathcal{C1}$ and $\mathcal{C2}$, one for the given controller
 577 and one for the class of controllers. Since these two automaton are identical,
 578 automata $\mathcal{C1}$ is shown in Figure 6. These automata model three functions from
 579 the path planning algorithm that compare the bicycle’s current position with
 580 the reference path. The first function involves finding the closest waypoint on

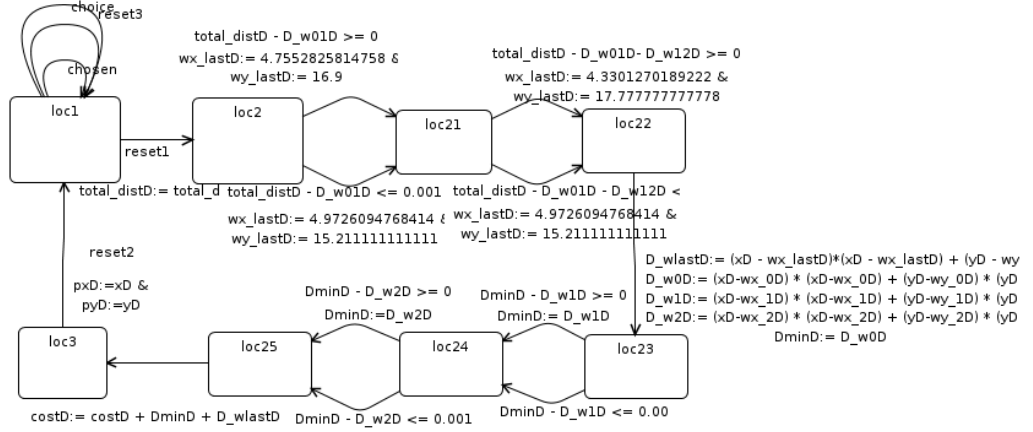


Fig. 6. Automaton C1 for motion planning

581 the reference path and comparing it with the current position. This is achieved
582 by computing the x- and y-coordinate distances between the current position
583 and each point on the reference path, starting from a specific index. The index
584 of the closest waypoint is saved using the distance formula⁵. The second function
585 determines the index of the last waypoint on the reference path that is located at
586 a certain distance from the current position. The last index is obtained by calculat-
587 ing the cumulative distance starting from the index of the first waypoint and
588 continuing until the index of the last waypoint. The third function creates a list
589 of N waypoints. It begins by using the index of the last waypoint obtained from
590 the second step to determine the number of indices (referred to as *delta_index*)
591 between the last index and the first index. If *delta_index* is greater than the
592 desired number of waypoints N , it means there are more than N path points
593 available. In this case, we consider the indices of the first N waypoints. However,
594 if *delta_index* is less than N , indicating that there are fewer than N waypoints
595 available, we consider the index of the last waypoint and repeat it multiple times.

596 In contrast to the physics-based cost utilized in the cart pole case study, the
597 cost in C1 and C2 is modeled as the sum of the accumulated cost, the distance
598 to the closest waypoint, and the distance to the last waypoint. This cumulative
599 sum represents the deviation between the path taken by the controller and the
600 reference path. Subsequently, we conduct a competitive analysis, as illustrated
601 in Figure 3b, utilizing the costs obtained from C1 and C2.

602 For the DNN abstracting this MPC, we trained a fully connected model with
603 3 input nodes, 3 output nodes and 4 hidden layers. The 3 input nodes correspond
604 to the parameters X , Y and θ that are used in the MPC. The output nodes

⁵ We employ the squared difference between the x- and y-coordinates without taking the square root. This choice is made due to the current limitation of flow*, which does not support the square root of variables (only the square root of a constant is supported).

605 correspond with the control variables Delta, V and a timestep. The timestep
606 is used by the simulation function of the system because the timestep is not
607 constant for this MPC. This MPC does not use an activation function on the
608 last layer to restrict the output. The rest of the construction and training of this
609 DNN is analogous to the previous DNN.

610 The generated composition $\mathcal{D} \parallel \mathcal{N} \parallel \mathcal{C}1 \parallel \mathcal{C}2$ is fed to Verisig along with
611 the property file and the trained DNN. The competitive analysis is performed
612 for the regret property $costD - costN \geq 0.25$. Verisig returned an ‘unsafe’
613 result along with a CE file. The trajectory from Verisig is compared against the
614 controller’s dataset and it is found that it is a spurious counterexample. As the
615 next step, the dataset file is appended with new data from the CE file, and the
616 DNN is retrained. The retrained DNN, along with the composed automaton and
617 the property file is fed to Verisig again. The DNN retraining is repeated for 20
618 iterations.