

# Learning-Based Synthesis of Safety Controllers

Daniel Neider

Max Planck Institute for Software Systems  
Kaiserslautern, Germany  
Email: neider@mpi-sws.org

Oliver Markgraf

Max Planck Institute for Software Systems  
Kaiserslautern, Germany

**Abstract**—We propose a machine learning framework to synthesize reactive controllers for systems whose interactions with their adversarial environment are modeled by infinite-duration, two-player games over (potentially) infinite graphs. Our framework targets safety games with infinitely many vertices, but it is also applicable to safety games over finite graphs whose size is too prohibitive for conventional synthesis techniques. The learning takes place in a feedback loop between a teacher component, which can reason symbolically about the safety game, and a learning algorithm, which successively learns an approximation of the winning region from various kinds of examples provided by the teacher. We develop a novel decision tree learning algorithm for this setting and show that our algorithm is guaranteed to converge to a reactive safety controller if a suitable approximation of the winning region can be expressed as a decision tree. Finally, we empirically compare the performance of a prototype implementation to existing approaches, which are based on constraint solving and automata learning, respectively.

## I. INTRODUCTION

Reactive synthesis offers an effective and promising way to solve a crucial practical problem: constructing correct and verified controllers for safety-critical systems. Rather than designing and implementing controllers by hand, reactive synthesis techniques construct controllers in an automatic fashion, thus, freeing engineers from this complex and error-prone task. In addition to being fully automatic, synthesis techniques produce correct-by-construction controllers that guarantee to satisfy the given specification, or they report that no such controller exist.

Typically, reactive synthesis is modeled as an infinite-duration game on a graph that is played by two antagonistic players: the system, which seeks to satisfy the specification, and the environment, which wants to violate it. More precisely, the specification and a model of the environment are in a first step converted into an infinite game. Then, one computes a winning strategy for the system, which prescribes how the system needs to play in order to win against every move of the environment. Finally, the winning strategy is translated into hard- or software, resulting in a reactive controller that satisfies the given specification.

In this paper, we focus on safety games, a class of infinite games that arises from safety specifications. Such specifications are in fact among the most important in practice (e.g., see Dwyer, Avrunin, and Corbett [1] for a survey of specification patterns) and capture many other interesting properties, including bounded-horizon reachability. In contrast to the classical setting, however, we consider safety games not only over

finite graphs but also over graphs with infinitely many (even uncountably many) vertices. Such games arise naturally, for instance, when the interaction between the controlled system and its environment is too complex to be modeled by finite graphs (e.g., in motion planning over unbounded environments) or when the environment has access to dynamic data structures, such as lists, stacks, or queues.

When the number of vertices of a game graph is infinite, traditional methods, which typically rely on an explicit exploration of the (whole) graph, are no longer applicable. To enable the computation of winning strategies for such games, our first contribution is a machine learning framework that constructs winning strategies by learning a proxy object called *winning set*. Intuitively, a winning set is an approximation of the vertices from which the system player can win and that permits to extract a winning strategy in a simple manner.

Our learning framework defines a feedback loop, akin to counterexample-guided inductive synthesis (CEGIS) [2], consisting of two entities: a teacher, who can reason symbolically about the game, and a learning algorithm, whose goal is to learn a winning set from information provided by the teacher. In every iteration of the loop, the learning algorithm constructs a set of vertices, which it proposes to the teacher. The teacher, on the other hand, checks whether the proposed set is a winning set and stops the learning process if so. If this is not the case, the teacher returns a counterexample. Upon receiving a counterexample, the learner refines its conjecture and proceeds with the next iteration.

Motivated by recent success in using decision trees as a concise and effective representation of strategies in infinite games [3], our second contribution is a new learning algorithm for decision trees, which is tailored specifically to the framework sketched above. Our algorithm builds upon a recent learning algorithm for decision trees that has been proposed in the context of software verification [4] and that puts a strong emphasis on learning “small” trees. As a consequence of the latter, our algorithm can in many situations guarantee to learn a winning set if one can be expressed as a decision tree.

Even though a game graph is infinite or prohibitively large, a reactive controller with a compact representation might already realize the specification. In a motion planning scenario, for instance, the system often only needs to consider a small subset of possible interactions with the environment to satisfy the specification. Based on this observation, our learning-based approach possesses various desirable properties: (i) it

leverages machine learning as an effective means to focus on the important parts of a game, (ii) it often learns “small” winning sets (and, by extension, small winning strategies) in a rule-like format, which tend to be relatively easy for humans to understand (see Brázdil et al. [3]), and (iii) besides operating over infinite graphs, it guarantees in many situations to learn a winning set if one exists. In addition, we demonstrate empirically that our approach is highly competitive to existing tools on two sets of benchmarks taken from the literature.

**Related Work:** Games over various types of infinite graphs have been studied, predominantly in the context of pushdown graphs [5]. For more general classes of graphs, a constraint-based approach [6], relying on constraint solvers such as Z3 [7], and various learning-based approaches have been proposed [8], [9] (we discuss these approaches in Section V). In the context of safety games over finite graphs, recent work [10] has demonstrated the ability of learning-based techniques to extract small controllers from precomputed controllers with a potentially large number of states.

Our learning framework is an extension of an earlier framework by Neider and Topcu [9]. Their work considers so-called rational safety games (defined in terms of finite automata) and proposes an automaton learning approach to infer winning strategies. By contrast, we do not fix a specific representation of the game graph and only require that certain operations can be performed symbolically. Many common formalisms such as finite automata, various types of decision diagrams, as well as formulas in the first-order theories of linear integer and real arithmetic satisfy these requirements. However, we consider only finitely-branching game graphs, whereas Neider and Topcu also consider graphs with infinite branching.

The algorithm we design for learning decision trees builds on top of a learning algorithm recently proposed by Ezudheen et al. [4], which learns from data in form of Horn clauses. For this setting, other learning algorithms have been developed as well [11], [12]. We have chosen Ezudheen et al.’s algorithm specifically for its property to guarantee convergence to a solution in many practical scenarios.

## II. CONTROLLER SYNTHESIS AND SAFETY GAMES

We follow the game-theoretic approach to controller synthesis as popularized by McNaughton [13] and view the problem as an infinite-duration, two-player game on a directed graph. Such games are played by two antagonistic players: Player 0, who embodies the system, and Player 1, who embodies the environment. In this setting, the type of specification dictates the type of game, and a winning strategy for Player 0 translates immediately into a controller that satisfies the given specification (we refer the reader to Grädel, Thomas, and Wilke [14] for a comprehensive discussion of this connection). Since we are interested in synthesizing controllers from safety specifications, the remainder of this paper is concerned with so-called safety games. However, before we introduce these types of games formally, let us first fix basic notations.

**Basic Notations:** Let  $\mathbb{B} = \{0, 1\}$  denote the set of Boolean values (0 representing *false* and 1 representing *true*),  $\mathbb{N}$  the

set of natural numbers,  $\mathbb{Z}$  the set of integers, and  $\mathbb{R}$  the set of real numbers. Given a set  $A$ , we denote the set of all finite sequences of elements of  $A$  by  $A^*$  and the set of all infinite sequences by  $A^\omega$ . Moreover, for a binary relation  $R \subseteq X \times X$  and two sets  $A, B \subseteq X$ , the *image of A under R* is the set  $R(A) = \{b \in B \mid \exists a \in A: (a, b) \in R\}$  and the *preimage of B under R* is the set  $R^{-1}(B) = \{a \in A \mid \exists b \in B: (a, b) \in R\}$ .

**Safety Games:** Our definitions and notations mainly follow those of Grädel, Thomas, and Wilke [14], and we refer the reader to this textbook for further details on infinite-duration games. Formally, a *safety game* is a five-tuple  $\mathcal{G} = (V_0, V_1, E, I, F)$  consisting of two disjoint sets  $V_0, V_1$  of vertices controlled by Player 0 and Player 1, respectively (we denote their union by  $V = V_0 \cup V_1$  and assume  $V \neq \emptyset$ ), a directed edge relation  $E \subseteq V \times V$ , a nonempty set  $I \subseteq V$  of *initial vertices*, and a set  $F \subseteq V$  of *safe vertices*. The directed graph  $(V, E)$  is typically called *game graph*. In contrast to the classical setting, we do not restrict  $V$  to be finite but allow even uncountable sets. However, we do make the following two restrictions to the edge relation: we assume that (1) every vertex has at least one outgoing edge (i.e.,  $E(\{v\}) \neq \emptyset$  for each  $v \in V$ ), and (2)  $E(\{v\})$  is finite for every  $v \in V$ , though not necessarily bounded. Note that the first restriction is standard and simply avoids situations in which the game gets stuck. The second restriction, on the other hand, is required by our learning framework and ensures that the data to learn from is always a finite object.

A safety game is played in rounds: initially, a token is placed on one of the initial vertices  $v_0 \in I$ ; in each round, the player controlling the current vertex then moves the token to the next vertex along one of the outgoing edges. This process of moving the token is repeated ad infinitum and results in an infinite sequence  $\pi = v_0 v_1 \dots \in V^\omega$  with  $v_0 \in I$  and  $(v_i, v_{i+1}) \in E$  for every  $i \in \mathbb{N}$ , which is called a *play*. The winner of a play is determined by the winning condition  $F$  in that a play  $\pi = v_0 v_1 \dots$  is *winning for Player 0* if  $v_i \in F$  for every  $i \in \mathbb{N}$ —otherwise it is *winning for Player 1*.

In the framework of infinite games, synthesizing a controller amounts to computing a so-called winning strategy for Player 0, which prescribes how Player 0 needs to move in order to win a play. Formally, a *strategy for Player 0* is a function  $\sigma: V^* \times V_0 \rightarrow V$  such that  $(v_n, \sigma(v_0 \dots v_n)) \in E$  for every  $v_0 \dots v_n \in V^* V_0$ . A strategy is called *winning* if every play that is *played according to  $\sigma$*  (i.e., that satisfies  $v_{n+1} = \sigma(v_0 \dots v_n)$  for all  $n \in \mathbb{N}$  with  $v_n \in V_0$ ) is winning for Player 0. It is well known that safety games permit memoryless winning strategies where the choice of the next vertex depends only on the vertex the play has currently reached. Such a strategy can then easily be implemented as a controller: the controller tracks the current vertex of a play and chooses the next move according to the strategy. Hence, the objective in the remainder of this paper is to compute a memoryless winning strategy for Player 0. We refer to this as *solving a game*.

If the game graph underlying a safety game is finite, memoryless winning strategies can be computed in linear time using a simple fixed-point computation [14]. For infinite game




Fig. 1: A robot on the one-dimensional grid world described in Example 1. The division of the world into intervals is indicated by dotted “boxes”. The number in the lower-right corner of each box displays the player having control over the robot in this interval. Safe vertices are surrounded by a bold line  $\square$ , initial vertices are indicated by a diagonal pattern  $\text{\textbackslash}$ .

graphs, on the other hand, this is no longer an option as a fixed-point computation might not converge in finite time. To overcome this problem, we propose a novel learning-based approach that learns a (memoryless) winning strategy via a proxy object named winning set.

**Winning Sets:** Intuitively, a winning set is a subset of the safe vertices that contains all initial vertices and is a trap for Player 1 (i.e., Player 0 can force any play to stay inside this set regardless of how Player 1 plays). Formally, we define winning sets as follows.

**Definition 1** (Winning set). *Let  $\mathcal{G} = (V_0, V_1, E, I, F)$  be a safety game. A winning set is a set  $W \subseteq V$  satisfying (1)  $I \subseteq W$ , (2)  $W \subseteq F$ , (3)  $E(\{v\}) \cap W \neq \emptyset$  for all  $v \in W \cap V_0$  (existential closedness), and (4)  $E(\{v\}) \subseteq W$  for all  $v \in W \cap V_1$  (universal closedness).*

A winning set  $W$  immediately provides a winning strategy for Player 0: starting in  $I \subseteq W$ , Player 0 simply moves to a (fixed) successor vertex inside  $W$  whenever it is his turn (note that this is possible since  $W$  is existentially closed). As  $W$  is also universally closed, a straightforward induction over the length of plays proves that every play that starts inside  $I$  and is played according to this strategy stays inside  $W$ , no matter how Player 1 plays. Thus, Player 0 wins since  $W \subseteq F$ .

In the remainder, we encourage the reader to think of a winning set as a (symbolic) representation of a winning strategy. The following example, inspired by robotic motion planning, illustrates the concept of winning sets.

**Example 1.** *Consider a robot in an unbounded, one-dimensional grid world as depicted in Figure 1. Moreover, let us assume that the robot’s position, indicated by an “x” in Figure 1, can be modeled as a single real-valued coordinate  $x \in \mathbb{R}$ , allowing for an uncountable number of positions.*

*The robot’s movement is controlled by two players, the system (Player 0) and the environment (Player 1). To decide which player is currently in control of the robot, we divide the world in infinitely many intervals  $[n, n + 1) \subset \mathbb{R}$  with  $n \in \mathbb{Z}$ : if the robot resides inside an interval  $[n, n + 1)$  with  $n$  even, then Player 0 is in control; otherwise, Player 1 is in control. In every turn, we allow the system or the environment to move the robot one unit to the left (i.e., decreasing its position  $x$  by one) or one unit to the right (i.e., increasing its position  $x$  by one). Note that this implies that the system and the environment are taking turns controlling the robot. We initially place the robot*

*at an arbitrary position  $x \in [0, 1)$  and define all positions  $x \geq 0$  to be safe. The objective of the system is to stay inside these safe positions at all times.*

*We can model the above situation as a safety game  $\mathcal{G} = (V_0, V_1, E, I, F)$  as follows. The set of Player 0 vertices is  $V_0 = \bigcup_{n \in \mathbb{Z}} [2n, 2n + 1)$ , while the set of Player 1 vertices is  $V_1 = \bigcup_{n \in \mathbb{Z}} [2n - 1, 2n)$ ; note that  $V = \mathbb{R}$ . The set of safe vertices is  $F = \{x \in \mathbb{R} \mid x \geq 0\}$  and the set of initial vertices is  $I = [0, 1)$ . Finally, the edge relation is defined as  $E = \{(x, x + 1) \in \mathbb{R} \times \mathbb{R} \mid x \in \mathbb{R}\} \cup \{(x, x - 1) \in \mathbb{R} \times \mathbb{R} \mid x \in \mathbb{R}\}$ .*

*An example of a winning set is  $W = [0, 3)$ . It clearly contains the set  $I = [0, 1)$  of initial vertices. Furthermore, if the position of the robot is in the interval  $[0, 1)$  or  $[2, 3)$ , Player 0 can stay inside  $W$  by moving the robot into the interval  $[1, 2)$ . Similarly, if the robot is in the interval  $[1, 2)$ , Player 1 has no choice but to move the robot either into the interval  $[0, 1)$  or into  $[2, 3)$ , thus staying inside  $W$ . Since  $W$  is also contained in  $F = \{x \in \mathbb{R} \mid x \geq 0\}$ , it is in fact a winning set. As sketched above, a winning strategy is easy to derive.  $\square$*

Finally, let us note that safety games as defined above subsume games on finite graphs, but one needs to choose a suitable symbolic representation in order to handle (un-)countably infinite game graphs. We comment on this in detail after having introduced our machine learning framework in the next section.

### III. A MACHINE LEARNING FRAMEWORK FOR SYNTHESIZING SAFETY CONTROLLERS

We now propose a machine learning framework for learning winning sets in safety games and, thus, reactive safety controllers. Our framework is a generalization of earlier work by Neider and Topcu [9], which encodes (countably) infinite game graphs using finite automata and uses automaton learning to learn winning sets. By contrast, the framework proposed here allows for game graphs with uncountably many vertices and is not restricted to a symbolic representation in terms of finite automata.

As illustrated in Figure 2, the learning takes place in a counterexample-guided feedback loop (CEGIS) [2] with two entities: a *teacher*, who has knowledge about the safety game, and a *learner* (or learning algorithm), whose objective is to learn a winning set, but who is agnostic to the game. In every iteration of the loop, the learner conjectures a set  $H \subseteq V$ , called *hypothesis*, based on the information about the game it has accumulated so far. Then, the teacher checks whether this set  $H$  is in fact a winning set—queries of this type are often called equivalence or correctness queries. Although the teacher does not know a winning set (the task is to learn one after all), it can verify whether the hypothesis is one by checking Conditions (1) to (4) of Definition 1. If the hypothesis satisfies these conditions, then  $H$  is a winning set and the learning stops. If this is not the case, the teacher replies with a counterexample that witnesses the violation of one of these conditions. Then, the feedback loop continues until a winning set has been found. The definition below formalizes the concept of counterexamples and fixes the communication between the teacher and the learner.




Fig. 2: Learning framework for synthesizing safety controllers

**Definition 2** (Teacher for Safety Games). Let  $\mathcal{G} = (V_0, V_1, E, I, F)$  be a safety game. Given a hypothesis  $H \subseteq V$ , the teacher replies as follows (whereby the order in which the checks are performed is arbitrary):

- 1) If  $I \not\subseteq H$ , then the teacher returns a positive counterexample  $v \in I \setminus H$ .
- 2) If  $H \not\subseteq F$ , then the teacher returns a negative counterexample  $v \in H \setminus F$ .
- 3) If there exists a  $v \in H \cap V_0$  with  $E(\{v\}) \cap H = \emptyset$ , then the teacher returns an existential counterexample  $v \rightarrow (v_1 \vee \dots \vee v_n)$  with  $\{v_1, \dots, v_n\} = E(\{v\})$ .
- 4) If there exists a  $v \in H \cap V_1$  with  $E(\{v\}) \not\subseteq H$ , then the teacher returns a universal counterexample  $v \rightarrow (v_1 \wedge \dots \wedge v_n)$  with  $\{v_1, \dots, v_n\} = E(\{v\})$ .

If  $H$  passes all four checks, then the teacher returns “yes”.

Each case of Definition 2 corresponds to one condition of Definition 1. Thus, it is not hard to verify that a given hypothesis is in fact a winning set if the teacher replies “yes” (as it satisfies Definition 1). Counterexamples, on the other hand, witness the violation of one of these conditions and guide the learner towards a winning set by communicating exactly why the hypothesis is incorrect. For instance, the meaning of a positive counterexample is that any future hypothesis needs to include this vertex (as it is initial), whereas a negative counterexample must be excluded (as it is not a safe vertex). An existential counterexample  $v \rightarrow (v_1 \vee \dots \vee v_n)$  signals that the hypothesis is not existentially closed and requires that if a future hypothesis contains  $v$ , it also needs to contain at least one vertex of the vertices  $v_1, \dots, v_n$ . Similarly, a universal counterexample  $v \rightarrow (v_1 \wedge \dots \wedge v_n)$  signals that the hypothesis is not universally closed and requires that if a future hypothesis contains  $v$ , it needs to contain all vertices  $v_1, \dots, v_n$ . Note that existential and universal counterexamples are always finite objects since we assume  $E(\{v\})$  to be finite for every  $v \in V$ .

Let us illustrate the overall learning process with an example.

**Example 2.** We continue Example 1. All hypotheses produced in the course of the learning process are depicted in Figure 3. Gray shaded areas indicate vertices in the hypothesis, safe vertices are surrounded by a bold line, and initial vertices are indicated by diagonal lines.

Let us assume that the learner proposes the hypothesis  $H_1 = \emptyset$  in the first iteration of the loop, which is shown in Figure 3a. Since this hypothesis does not include any initial vertex, it does not satisfy Condition (1) of Definition 1. Thus, the teacher returns a positive counterexample, say  $0 \in I \setminus H$ .

Next, suppose that the learner proposes the hypothesis  $H_2 = V$ , shown in Figure 3b. This hypothesis is consistent with the positive counterexample. However, it also includes unsafe




Fig. 3: Hypotheses produced in the course of Example 2.

vertices and, thus, does not satisfy Condition (2) of Definition 1. Hence, the teacher replies with a negative counterexample, say  $-1 \in H \setminus F$ .

Let us now assume that the learner conjectures  $H_3 = [0, 1]$ , depicted in Figure 3c. This conjecture is consistent with both the positive and the negative counterexample. However, it is not existentially closed because Player 0 has no choice but to move the robot outside of  $H_3$ . Therefore, the teacher replies with an existential counterexample, say  $0 \rightarrow (-1 \vee 1)$ .

Suppose now that the learner proposes the hypothesis  $H_4 = [0, 2]$  in the fourth iteration, shown in Figure 3d. Although this conjecture is consistent with all counterexamples received so far, it is not universally closed because Player 1 can move the robot into the interval  $[2, 3)$ , which is not included in  $H_4$ . Thus, the teacher returns a universal counterexample, say  $2 \rightarrow (1 \wedge 3)$ .

Finally, the learner proposes the hypothesis  $H_5 = [0, 3]$ , depicted in Figure 3e. This hypothesis satisfies all conditions of Definition 2. Thus,  $H_5$  is a winning set and the learning terminates.  $\perp$

Our learning framework is straightforward to implement if the underlying game graph is finite, in which case the teacher can be built on top of an explicit representation of the game. However, if the underlying game graph becomes too large or is infinite, one has to choose a suitable representation for sets of vertices and the edge relation that allows performing operations on the graph symbolically. More precisely, the chosen symbolic representation must feature Boolean operations (i.e., union, intersection, and complementation), and the image  $E(A)$  and preimage  $E^{-1}(A)$  of symbolically represented sets  $A \subseteq V$  need to be computable. Moreover, the emptiness problem (i.e., “given a set  $A$ , decide whether  $A = \emptyset$ ”) needs to be decidable, and it must be possible to extract an element from  $A$  if it is nonempty. Examples of such symbolic representations include many common formalisms such as finite automata, various types of decision diagrams, and first-order formulas in linear integer arithmetic and real arithmetic.

Furthermore, hypotheses must also be expressible in the chosen symbolic formalism. However, to build efficient learning algorithms, it is often necessary to restrict the class of

hypotheses—typically called the *hypothesis space* and denoted by  $\mathcal{H}$ —even further. Examples of such restricted hypothesis spaces are conjunctive formulas (e.g., as used by the popular Houdini algorithm [15]) or decision trees, which are common, for instance, in learning-based software verification [11], [16], [17], [12] and also used in this work. Note that it might happen that a winning set exists, though it cannot be expressed as a hypothesis in the hypothesis space. Thus, the choice of the hypothesis space and, hence, the learning algorithm needs to be made carefully in order for the learning to succeed.

A second important property of learning algorithms is what we call “consistency”. To make this notion precise, let us assume that the learner accumulates counterexamples in a so-called *game sample*  $\mathcal{S}_G = (Pos, Neg, Ex, Un)$  consisting of a finite set  $Pos$  of positive counterexamples, a finite set  $Neg$  of negative counterexample, a finite set  $Ex$  of existential counterexamples, and a finite set  $Un$  of universal counterexamples. Then, we say that a hypothesis  $H \subseteq V$  is *consistent* with a game sample  $\mathcal{S}_G = (Pos, Neg, Ex, Un)$  if (1)  $v \in H$  for each  $v \in Pos$ , (2)  $v \notin H$  for each  $v \in Neg$ , (3)  $v \in H$  implies  $\{v_1, \dots, v_n\} \cap H \neq \emptyset$  for each  $v \rightarrow (v_1 \vee \dots \vee v_n) \in Ex$ , and (4)  $v \in H$  implies  $\{v_1, \dots, v_n\} \subseteq H$  for each  $v \rightarrow (v_1 \wedge \dots \wedge v_n) \in Un$ . Moreover, we call a learner *consistent* if it always produces a consistent hypothesis. Consistency is an important property as it prevents the learner from making the same mistake twice and ensures progress towards a winning set.

In fact, the notion of consistency allows us to show that our framework is sound in the sense that any consistent learner learns a winning set in the limit if one exists in the chosen hypothesis space. This is formalized in the next theorem.

**Theorem 1.** *Let a teacher for a safety game (as described in Definition 2) and a consistent learner over an hypothesis space  $\mathcal{H}$  be given. If there exists a winning set expressible as a hypothesis in  $\mathcal{H}$ , then there exists an ordinal  $\alpha \in \mathbb{O}$ , where  $\mathbb{O}$  denotes the class of all ordinals, such that the learner proposes a winning set after at most  $\alpha$  iterations.*

Theorem 1 is a consequence of the fact that our learning framework is an instance of an *abstract learning framework for synthesis (ALF)*, as introduced by Löding et al. [18]. The proof roughly proceeds as follows. Since the teacher of Definition 2 allows “progress” (i.e., every counterexample refutes the current hypothesis) and we assume the learner to be consistent, the learner never conjectures the same hypothesis twice. In the worst case, the learner will have exhausted all incorrect hypotheses after  $\alpha$  iterations for an ordinal  $\alpha \in \mathbb{O}$  with cardinality less or equal to  $|\mathcal{H}|$ . Since the teacher is also “honest” (i.e., it does not return spurious counterexamples), the learner necessarily produces a winning set in the subsequent iteration if one exists.

Finally, let us point out that the safety games as defined in Section II are very general and even allow encoding computations of Turing machines. Consequently, determining the winner of such safety games is undecidable in general, and any algorithm for computing winning sets can be a

semi-algorithm at best (i.e., an algorithm that, on termination, gives the correct answer but does not guarantee to halt). The algorithm we design in the next section is of this kind.

#### IV. LEARNING DECISION TREES FROM GAME SAMPLES

We now fix the hypothesis space  $\mathcal{H}$  to be the class of all decision trees (as defined shortly) and describe an algorithm to learn consistent decision trees from game samples. Our algorithm builds on top of a learning algorithm recently proposed by Ezudheen et al. [4] in the context of software verification. To ease the presentation in this section, we abstract from the setting of infinite games and assume that the data to learn from is taken from an abstract domain  $\mathcal{D}$ , whose elements we call *data points*. We encourage the reader to think of data points and vertices as synonyms and define concepts such as game samples and consistency analogously for data points.

An example of a decision tree is shown in Figure 4. In general, decision trees are binary trees whose inner nodes are labeled with predicates from an a priori fixed set  $\mathcal{P}$  and whose leaves are labeled with Boolean values. In this context, each predicate is a function  $p: \mathcal{D} \rightarrow \mathbb{B}$  that maps data points to a Boolean




Fig. 4: A decision tree over  $\mathcal{D} = \mathbb{R}$  and  $\mathcal{P} = \{x \geq 0, y < 2\}$ .

values and corresponds to a property of interest. Typically, the set  $\mathcal{P}$  is finite, but Ezudheen et al.’s algorithm can build decision trees even from infinite sets of predicates if the underlying domain is numeric. The algorithm we design in this section retains this feature.

In the remainder of this section, we view a decision tree  $t$  as a representation of an (infinite) set  $D(t) \subseteq \mathcal{D}$  of data points. Whether a data point  $d \in \mathcal{D}$  belongs to this set depends on its *valuation*  $t(d) \in \mathbb{B}$ , which is defined as follows: starting at the root node, we recursively descend left (right) if  $d$  satisfies (does not satisfy) the predicate at the current node and define  $t(d)$  to be the label of the leaf node that is ultimately reached by this procedure. The set of data points represented by  $t$  is then simply the set  $D(t) = \{d \in \mathcal{D} \mid t(d) = 1\}$ . Since the learner proposes this set as a hypothesis to the teacher, we call a decision tree  $t$  consistent with a game sample  $\mathcal{S}_G$  if  $D(t)$  is consistent with  $\mathcal{S}_G$ .

Ezudheen et al.’s algorithm has been developed in the context of software verification and expects so-called Horn samples as input. Formally, a *Horn sample* is a finite set  $\mathcal{S}_H$  containing *Horn constraints* of the form  $(d_1 \wedge \dots \wedge d_n) \rightarrow d$  or  $(d_1 \wedge \dots \wedge d_n) \rightarrow false$  where  $i \in \mathbb{N}$  and  $d, d_1, \dots, d_n \in \mathcal{D}$  are data points. Note that the left-hand-side of a Horn constraint might be empty (i.e.,  $i = 0$ ), in which case it is interpreted as *true*.

Given a Horn sample  $\mathcal{S}_H$ , Ezudheen et al.’s algorithm learns a decision tree  $t$  that is consistent with  $\mathcal{S}_H$  in the sense that (1)  $\{d_1, \dots, d_n\} \subseteq D(t)$  implies  $d \in D(t)$  for each Horn constraint of the form  $(d_1 \wedge \dots \wedge d_n) \rightarrow d$  in  $\mathcal{S}_H$  and (2)  $\{d_1, \dots, d_n\} \not\subseteq D(t)$  for each Horn constraint of the form  $(d_1 \wedge \dots \wedge d_n) \rightarrow false$  in  $\mathcal{S}_H$ ; by extension, we then also say that  $D(t)$  is consistent with  $\mathcal{S}_H$ . Should no consistent decision

---

**Algorithm 1:** Learning decision trees from game samples

---

**Input:** A game sample  $\mathcal{S}_G = (Pos, Neg, Ex, Un)$  and a (finite) set  $\mathcal{P}$  of predicates (both over  $\mathcal{D}$ )

- 1 Construct a Horn sample  $\mathcal{S}_H$  as follows:
    - for each positive example  $d \in Pos$ , add the Horn constraint  $d \rightarrow false$ ;
    - for each negative example  $d \in Neg$ , add the Horn constraint  $true \rightarrow d$ ;
    - for each existential implication  $d \rightarrow (d_1 \vee \dots \vee d_n) \in Ex$ , add the Horn constraint  $(d_1 \wedge \dots \wedge d_n) \rightarrow d$ ; and
    - for each universal implication  $d \rightarrow (d_1 \wedge \dots \wedge d_n) \in Un$ , add the Horn constraints  $d_1 \rightarrow d, \dots, d_n \rightarrow d$  to  $\mathcal{S}_H$ .
  - 2 Apply Ezudheen et al.’s learning algorithm to learn a decision tree  $t_H$  over  $\mathcal{P}$  that is consistent with  $\mathcal{S}_H$ .
  - 3 Swap the labels of all leaves of  $t_H$  to obtain the decision tree  $t_G$  satisfying  $D(t_G) = \mathcal{D} \setminus D(t_H)$  (i.e., the decision trees  $t_H$  and  $t_G$  are structurally identical but every label  $b \in \mathbb{B}$  in  $t_H$  is replaced with  $1 - b$ ).
- 

tree exist (e.g., because the set  $\mathcal{P}$  does not contain sufficient predicates to separate positive and negative examples), the algorithm aborts with an error. The theorem below summarizes these key properties.

**Theorem 2** (Ezudheen et al. [4]). *Let  $\mathcal{S}_H$  be a Horn sample and  $\mathcal{P}$  a finite set of predicates, both over the domain  $\mathcal{D}$ . Moreover, let  $n$  be the number of data points and  $k$  the number of Horn constraints in  $\mathcal{S}_H$ . If a decision tree over  $\mathcal{P}$  exists that is consistent with  $\mathcal{S}_H$ , then Ezudheen et al.’s algorithm learns one in time  $\mathcal{O}(n|\mathcal{P}| + n^2k)$ , assuming that predicates can be evaluated in constant time.*

We are now ready to present our algorithm for learning decision trees from game samples, which is shown in pseudo code as Algorithm 1. It builds on top of Ezudheen et al.’s algorithm and proceeds in three steps: first, it translates a game sample  $\mathcal{S}_G$  into an “equivalent” Horn sample  $\mathcal{S}_H$ ; then, it applies Ezudheen et al.’s learning algorithm to obtain a decision tree  $t_H$  that is consistent with  $\mathcal{S}_H$ ; finally, it translates  $t_H$  into a decision tree  $t_G$ , which is consistent with  $\mathcal{S}_G$ . More precisely, the Horn sample  $\mathcal{S}_H$  constructed in Step 1 has the property that for every decision tree  $t_H$  the set  $D(t_H)$  is consistent with  $\mathcal{S}_H$  if and only if its complement  $\mathcal{D} \setminus D(t_H)$  is consistent with the game sample  $\mathcal{S}_G$ . Since the decision tree  $t_G$  obtained in Step 3 of our algorithm satisfies  $D(t_G) = \mathcal{D} \setminus D(t_H)$ , it is thus consistent with  $\mathcal{S}_G$ . This property is formalized next and proven in the extended version of this paper [19].

**Lemma 1.** *Let  $\mathcal{S}_G$  be a game sample and  $\mathcal{P}$  a finite set of predicates, both over the domain  $\mathcal{D}$ . Moreover, let  $\mathcal{S}_H, t_H,$*

*and  $t_G$  be as in Algorithm 1. Then,  $D(t_H)$  is consistent with  $\mathcal{S}_H$  if and only if  $D(t_G)$  is consistent with  $\mathcal{S}_G$ .*

The correctness of Algorithm 1, stated in Theorem 3 below, is now a direct consequence of Lemma 1 and Theorem 2. Note that the term  $n|Un|$  in the runtime estimation stems from the fact that each universal counterexample might have  $n$  data points on its right-hand-side, resulting in  $n$  Horn constraints.

**Theorem 3.** *Let  $\mathcal{S}_G = (Pos, Neg, Ex, Un)$  be a game sample and  $\mathcal{P}$  a finite set of predicates, both over the domain  $\mathcal{D}$ . Moreover, let  $n$  be the number of data points in  $\mathcal{S}_G$ . If a decision tree over  $\mathcal{P}$  exists that is consistent with  $\mathcal{S}_G$ , then Algorithm 1 learns one in time  $\mathcal{O}(n|\mathcal{P}| + n^2(|Pos| + |Neg| + |Ex| + n|Un|))$ , assuming that predicates can be evaluated in constant time.*

If the abstract domain  $\mathcal{D}$  is numeric, say  $\mathcal{D} \subseteq \mathbb{R}^n$  for some  $n \in \mathbb{N}$ , a simple extension of Ezudheen et al.’s algorithm, and hence Algorithm 1, can be used to learn decision trees even over infinite sets of predicates. This extension was originally proposed by Quinlan [20] and assumes the set of predicates to be  $\mathcal{P}^* = \{d[i] \leq c \mid c \in \mathbb{R}, i \in \{1, \dots, n\}\}$  where  $d[i]$  denotes the  $i$ -th component of the data point  $d$ . The key idea is in fact straightforward: to separate two data points  $d_1, d_2 \in \mathcal{D}$  (e.g., a positive and negative data point), it suffices to consider only such predicates  $d[i] \leq c$  for which the constant  $c$  occurs as an actual value in  $d_1$  or  $d_2$ . Thus, Algorithm 1 can always restrict itself to a finite subset of  $\mathcal{P}^*$ , which only depends on the values of the data points in the given sample.

In many situations, Algorithm 1 in fact guarantees that the overall feedback loop of Section III converges to a winning set in finite time if one can be expressed as a decision tree. For instance, if the set  $\mathcal{P}$  is finite, then there exist only finitely many semantically different decision trees (in terms of the set of data points they represent). Since Algorithm 1 always produces consistent decision trees, it will have exhausted all incorrect trees after a finite amount of time. If a winning set exists and can be expressed as a decision tree over  $\mathcal{P}$ , then the subsequent hypothesis will necessarily be one. On the other hand, if the underlying domain is numeric and Algorithm 1 operates over the infinite set  $\mathcal{P}^*$ , a technique proposed by Ezudheen et al. can be used to guarantee that a winning set will be learned in finite time (if one can be expressed as a decision tree over  $\mathcal{P}^*$ ). Due to the limited space, we have to refer the reader to Ezudheen et al. [4] for details and can here only state our main result.

**Theorem 4.** *Let  $\mathcal{G}$  be a safety game. If the learner of Section III uses Algorithm 1 over a finite set  $\mathcal{P}$  of predicates, then it is guaranteed to learn a winning set after a finite number of iterations if there exists one that is expressible as a decision tree over  $\mathcal{P}$ . The analogous statement holds for the set  $\mathcal{P}^*$ .*

In practice, the overall runtime depends not only on the learner but also on the teacher (which might be adversarial). Hence, an average case analysis requires further assumptions on the teacher. We leave such an analysis for future work and turn to an experimental evaluation instead.

## V. EXPERIMENTAL EVALUATION

To assess the performance of our learning framework and the decision tree learning algorithm, we have implemented a prototype named *DT-Synth*.<sup>1</sup> In this section, we describe *DT-Synth* in detail and compare it to three existing tools: *CONSYNTH* [6], which is based on constraint solving, as well as two tools based on learning finite automata [9], for brevity here called *SAT-Synth* and *RPNI-Synth*.

*The Tools:* *DT-Synth* takes games as input that are encoded as quantifier-free, first-order formulas in the theory of linear integer arithmetic (LIA). More precisely, variables  $x_1, \dots, x_n$  of type integer encode vertices of a game graph, while formulas  $\varphi(x_1, \dots, x_n)$  over those variables represent (infinite) sets of vertices. Additionally, the edge relation of a game graph is encoded by a formula  $\psi(x_1, \dots, x_n, x'_1, \dots, x'_n)$  where  $x_1, \dots, x_n$  represents a source-vertex and  $x'_1, \dots, x'_n$  a destination-vertex. We believe that this encoding of games is highly accessible and allows modeling many real-world synthesis tasks (e.g., robotic motion planning) in a natural way. Moreover, the decision trees computed by *DT-Synth* are generally easy to comprehend, as we demonstrate in the full version of this paper [19]. Note, however, that the choice of LIA restricts our game graphs to countably many vertices. We have made this choice deliberately to be able to reuse Ezudheen et al.’s implementation of the decision tree learner [4]. *DT-Synth* implements the teacher and learner as follows:

*Teacher:* The teacher builds on top of the Z3 SMT solver [7]. Upon receiving a hypothesis in the form of a formula  $\varphi_H(x_1, \dots, x_n)$ , it performs a series of satisfiability checks according to Definition 2 in order to search for counterexamples. If a satisfiability check succeeds, the teacher derives a counterexample from the model returned by Z3 (potentially triggering a finite number of additional satisfiability checks to compute the successors of a vertex in the case of existential and universal counterexamples).

*Learner:* The learner builds on top of code originally developed by Ezudheen et al. [4]. In addition to the set  $\mathcal{P}^*$  of predicates described in Section IV, it additionally uses octagonal constraints of the form  $x_i \pm x_j \leq c$  (which capture distances between vertices in the Euclidean space). After a consistent decision tree has been learned, the learner converts it into a formula  $\varphi_H := \bigvee_{\pi \in \Pi} \bigwedge_{\psi \in \pi} \psi$  where  $\Pi$  is the set of all paths from the root to a leaf labeled with *true*, and  $\psi \in \pi$  denotes that the predicate  $\psi$  occurs on  $\pi$  (negated if the path descends to the right). Then, it hands  $\varphi_H$  over to the teacher. Following the description in Section IV, it is not hard to verify that a vertex satisfies the formula  $\varphi_H$  obtained from the tree  $t$  if and only if the vertex belongs to  $D(t)$ .

We compared *DT-Synth* to three other tools: *CONSYNTH*, *SAT-Synth* and *RPNI-Synth*. Table I summarizes the main similarities and differences of these tools.

*CONSYNTH* [6] is a tool for synthesizing winning strategies in games over infinite graphs with  $\omega$ -regular winning conditions

TABLE I: Properties of the compared tools

| Tool                  | Easy to model games | Easy to interpret solution | Guarantees to find a strategy | No help of user required |
|-----------------------|---------------------|----------------------------|-------------------------------|--------------------------|
| <i>DT-Synth</i>       | ✓                   | ✓                          | ✓                             | ✓                        |
| <i>CONSYNTH</i> [6]   | ✓                   | ✗                          | ✓                             | ✗                        |
| <i>SAT-Synth</i> [9]  | ✗                   | ✗                          | ✓                             | ✓                        |
| <i>RPNI-Synth</i> [9] | ✗                   | ✗                          | ✗                             | ✓                        |

(which are more general than the safety games considered in this paper). The tool reduces the computation of a winning strategy into a satisfiability problem for Constrained Horn Clauses (CHCs) [21], which allows it to leverage the power of modern SMT solvers. Similar to *DT-Synth*, *CONSYNTH*’s input is a game (with potentially unaccountably many vertices) encoded as an SMT formula. However, to remove existential quantifiers that arise from translating the synthesis problem into CHCs, *CONSYNTH* relies on so-called *Skolem templates*, which have to be specified by the user (we discuss the impact of Skolem templates in the full version of this paper [19]). Intuitively, these templates constrain the search space of potential strategies and, hence, require insight into what a winning strategy might be (and the ability to express this insight as a Skolem template). By contrast, *DT-Synth* does not require additional help from the user (it synthesizes winning sets/strategies based on the game alone), but it can also not benefit from the user’s knowledge. The output of *CONSYNTH* is values for the parameters in the Skolem template (encoding a winning strategy) and a corresponding winning set in form of an SMT formula. Both objects can be challenging for humans to understand. Similar to *DT-Synth*, *CONSYNTH* guarantees to find a winning strategy if one can be expressed in terms of the Skolem template.

*SAT-Synth* and *RPNI-Synth* [9] are two techniques for synthesizing winning strategies of safety games over game graphs with countably many vertices. Both tools expect the game encoded by means of finite automata/transducers and leverage automata learning techniques to compute winning sets: *SAT-Synth* internally uses a SAT solver to learn automata, whereas *RPNI-Synth* is based on the popular RPNI learning algorithm [22]. Similar to *DT-Synth* and unlike *CONSYNTH*, both tools do not require additional information from the user. However, both the input and output of *SAT-Synth* and *RPNI-Synth* are finite automata, which is not a very natural encoding of games and can be difficult for humans to understand. Finally, *SAT-Synth* guarantees to find a winning set (if it can be expressed as an automaton), whereas *RPNI-Synth* does not.

*The Benchmarks:* We have evaluated the performance of all four tools on two benchmark suits, both featuring safety games over infinite game graphs.<sup>2</sup> The first benchmark suite accompanies *SAT-Synth*/*RPNI-Synth* [9] and consists of seven safety games, which are motivated by robotic motion

<sup>2</sup>Note that our approach is designed for infinite graphs but not to compete with highly-optimized synthesis engines on games over finite graphs. Hence, we did not consider benchmarks from the various synthesis competitions.

<sup>1</sup>Code/benchmarks available at <https://github.com/OliverMal/DT-Synth>.

TABLE II: Experimental results on the *SAT-Synth*/*RPNI-Synth* benchmark suite [9] (upper part) and *CONSYNTH* benchmark suite [6] (lower part). “Iter.” refers to the number of iterations in the counterexample-guided feedback loop. “Size” measures the size of the final decision tree learned by *DT-Synth* in terms of inner nodes and the size of the final automata produced by *SAT-Synth* and *RPNI-Synth* in terms of the number of states, respectively. “—” indicates a timeout after 900 s.

| Game                   | <i>DT-Synth</i> |       |      | <i>SAT-Synth</i> |       |      | <i>RPNI-Synth</i> |       |      | <i>CONSYNTH</i> |
|------------------------|-----------------|-------|------|------------------|-------|------|-------------------|-------|------|-----------------|
|                        | Time in s       | Iter. | Size | Time in s        | Iter. | Size | Time in s         | Iter. | Size | Time in s       |
| Box                    | 0.77            | 9     | 5    | 1.22             | 45    | 5    | 0.41              | 16    | 6    | 3.71            |
| Box Limited            | 0.28            | 4     | 2    | 0.74             | 36    | 4    | 0.32              | 15    | 5    | 0.44            |
| Diagonal               | 2.04            | 23    | 5    | 1.49             | 64    | 4    | 1.01              | 64    | 4    | 1.93            |
| Evasion                | 0.63            | 6     | 3    | 89.15            | 237   | 7    | 1.40              | 82    | 11   | 1.50            |
| Follow                 | 0.86            | 11    | 5    | 95.90            | 300   | 7    | 7.36              | 352   | 16   | —               |
| Solitary Box           | 0.24            | 4     | 2    | 5.81             | 76    | 6    | 0.42              | 16    | 6    | 0.42            |
| Square 5x5             | 7.77            | 61    | 12   | —                | —     | —    | 0.82              | 39    | 14   | —               |
| Cinderella ( $c = 2$ ) | —               | —     | —    | —                | —     | —    | —                 | —     | —    | —               |
| Cinderella ( $c = 3$ ) | —               | —     | —    | —                | —     | —    | —                 | —     | —    | 765.30          |
| Program-repair         | 0.99            | 14    | 11   | 1.42             | 69    | 3    | 0.15              | 7     | 3    | 2.50            |
| Repair-critical        | 23.74           | 237   | 14   | —                | —     | —    | 130.54            | 1772  | 11   | 19.52           |
| Synth-Synchronization  | 63.30           | 513   | 35   | —                | —     | —    | 42.74             | 858   | 26   | 10.01           |

planning. The second suite of benchmarks is shipped with *CONSYNTH* [6] and consists of five safety games: two versions of a combinatorial puzzle (Cinderella game), two program repair problems, and a synchronization problem for multi-threaded programs. Due to limited space, we have to refer to the full version of this paper [19] for a detailed description of these benchmarks. For the first benchmark suite, we have equipped *CONSYNTH* with moderately restrictive Skolem templates; templates for the second suite were provided by Beyene et al. The representations of winning sets/strategies for all four tools are expressive enough for all games in these benchmark suits.

*Results:* Table II lists the experimental results for all four tools on the *SAT-Synth*/*RPNI-Synth* benchmark suite [9] (upper part) and the *CONSYNTH* benchmark suite [6] (lower part).<sup>3</sup> In the case of *DT-Synth*, *RPNI-Synth*, and *SAT-Synth*, Table II also shows the number of iterations as well as the size of the result (measured in the number of inner nodes in a decision tree and the number of states of an automaton, respectively). We have conducted all experiments on an Intel Xeon E7-8857 v2 CPU with 4 GB of RAM running a 64-bit Debian operating system. The timeout was 900 s.

We rank the performance of the tools based on the number of games they solved. To break ties, we consider the aggregate runtime on games that the tools were able to solve (not accounting for time-outs). With this scheme, *DT-Synth* performed best and solved ten out of twelve games with an aggregated runtime of 100.62 s. *RPNI-Synth* ranked second, having solved the same ten games but with an aggregated runtime of 185.17 s (i.e., 1.8 times slower than *DT-Synth*). *CONSYNTH* ranked third and solved nine games with an aggregated runtime of 805.33 s. Finally, *SAT-Synth* performed worst, having solved seven games with an aggregate runtime of 195.73 s.

<sup>3</sup>Note that we were not able to reproduce the exact results for *CONSYNTH* reported by Beyene et al. [6]. *CONSYNTH* has numerous options, which influence the performance of the benchmarks drastically. We have been in contact with the authors to obtain the exact options that were used in their experiments. Unfortunately, some of these options resulted in crashes and incorrect results on our machine. Hence, we resorted to the default options.

Compared to *RPNI-Synth* and *SAT-Synth*, *DT-Synth* required far fewer iterations and, hence, fewer interactions with the typically computationally expensive teacher. Moreover, the size of the final output is smaller, which makes it easier for humans to interpret (see the full version of this paper [19]). It is also important to emphasize that despite *RPNI-Synth*’s good performance, the tool does not guarantee to find a winning set. By contrast, *DT-Synth* provides such a guarantee.

Note that *CONSYNTH* was the only tool able to solve at least one version of the Cinderella game. We believe that this is due to the somewhat restrictive Skolem template that Beyene et al. have provided: if equipped with less restrictive templates, *CONSYNTH* times out as well (see the full version [19]).

*Summary:* In conclusion, *DT-Synth* is competitive to the state-of-the-art tools for solving safety games over infinite graphs. It does not require any user guidance, guarantees to find a winning set (if one can be expressed as a decision tree), and features easy-to-understand input/output formats.

## VI. CONCLUSION

We have developed a machine learning framework for synthesizing reactive safety controllers whose interaction with their environment is modeled by games over infinite graphs. Moreover, we have designed a learning algorithm for decision trees that learns winning sets/strategies and is in many situations guaranteed to find a solution if one exists. Our experimental evaluation shows that our approach is highly competitive and promises applicability to a wide range of interesting practical problems, specifically due to its ease of use.

A promising direction for future work would be to apply our technique to distributed synthesis problems and other, more complex synthesis settings. Moreover, we plan to extend our learning-based framework to more general winning conditions, such as reachability and liveness.

## ACKNOWLEDGMENT

This work was partially funded by the ERC Starting Grant AV-SMP (grant agreement no. 759969)



## REFERENCES

- [1] M. B. Dwyer, G. S. Avrunin, and J. C. Corbett, "Patterns in property specifications for finite-state verification," in *Proceedings of the 1999 International Conference on Software Engineering, ICSE '99, Los Angeles, CA, USA, May 16-22, 1999*. ACM, 1999, pp. 411–420.
- [2] A. Solar-Lezama, "The sketching approach to program synthesis," in *Programming Languages and Systems, 7th Asian Symposium, APLAS 2009, Seoul, Korea, December 14-16, 2009. Proceedings*, ser. Lecture Notes in Computer Science, vol. 5904. Springer, 2009, pp. 4–13.
- [3] T. Brázdil, K. Chatterjee, J. Kretínský, and V. Toman, "Strategy representation by decision trees in reactive synthesis," in *Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018. Proceedings, Part I*, ser. Lecture Notes in Computer Science, vol. 10805. Springer, 2018, pp. 385–407.
- [4] P. Ezudheen, D. Neider, D. D'Souza, P. Garg, and P. Madhusudan, "Horn-ice learning for synthesizing invariants and contracts," *PACMPL*, vol. 2, no. OOPSLA, pp. 131:1–131:25, 2018.
- [5] O. Kupferman, N. Piterman, and M. Y. Vardi, "An automata-theoretic approach to infinite-state systems," in *Time for Verification, Essays in Memory of Amir Pnueli*, vol. 6200. Springer, 2010, pp. 202–259.
- [6] T. A. Beyene, S. Chaudhuri, C. Popeea, and A. Rybalchenko, "A constraint-based approach to solving games on infinite graphs," in *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*. ACM, 2014, pp. 221–234.
- [7] L. M. de Moura and N. Björner, "Z3: an efficient SMT solver," in *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, ser. Lecture Notes in Computer Science, vol. 4963. Springer, 2008, pp. 337–340.
- [8] D. Neider, "Reachability games on automatic graphs," in *Implementation and Application of Automata - 15th International Conference, CIAA 2010, Winnipeg, MB, Canada, August 12-15, 2010. Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 6482. Springer, 2010, pp. 222–230.
- [9] D. Neider and U. Topcu, "An automaton learning approach to solving safety games over infinite graphs," in *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, ser. LNCS, vol. 9636. Springer, 2016, pp. 204–221.
- [10] D. Neider, "Small strategies for safety games," in *Automated Technology for Verification and Analysis, 9th International Symposium, ATVA 2011, Taipei, Taiwan, October 11-14, 2011. Proceedings*, ser. Lecture Notes in Computer Science, vol. 6996. Springer, 2011, pp. 306–320.
- [11] A. Champion, T. Chiba, N. Kobayashi, and R. Sato, "Ice-based refinement type discovery for higher-order functional programs," in *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, ser. LNCS, vol. 10805. Springer, 2018, pp. 365–384.
- [12] H. Zhu, S. Magill, and S. Jagannathan, "A data-driven CHC solver," in *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, 2018, pp. 707–721.
- [13] R. McNaughton, "Infinite games played on finite graphs," *Ann. Pure Appl. Logic*, vol. 65, no. 2, pp. 149–184, 1993.
- [14] E. Grädel, W. Thomas, and T. Wilke, Eds., *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, ser. Lecture Notes in Computer Science, vol. 2500. Springer, 2002.
- [15] C. Flanagan and K. R. M. Leino, "Houdini, an annotation assistant for esc/java," in *FME 2001: Formal Methods for Increasing Software Productivity*, ser. LNCS, vol. 2021. Springer, 2001, pp. 500–517.
- [16] P. Garg, C. Löding, P. Madhusudan, and D. Neider, "ICE: A robust framework for learning invariants," in *Computer Aided Verification (CAV)*, ser. LNCS, vol. 8559. Springer, 2014, pp. 69–87.
- [17] P. Garg, D. Neider, P. Madhusudan, and D. Roth, "Learning invariants using decision trees and implication counterexamples," in *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 2016, pp. 499–512.
- [18] C. Löding, P. Madhusudan, and D. Neider, "Abstract learning frameworks for synthesis," in *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, ser. LNCS, vol. 9636. Springer, 2016, pp. 167–185.
- [19] D. Neider and O. Markgraf, "Learning-based synthesis of safety controllers," *CoRR*, vol. abs/1901.06801, 2019. [Online]. Available: <http://arxiv.org/abs/1901.06801>
- [20] J. R. Quinlan, *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.
- [21] T. A. Beyene, C. Popeea, and A. Rybalchenko, "Solving existentially quantified horn clauses," in *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, ser. Lecture Notes in Computer Science, vol. 8044. Springer, 2013, pp. 869–882.
- [22] J. Oncina and P. García, "Inferring regular languages in polynomial update time," *Pattern Recognition and Image Analysis*, pp. 49–61, 1992.