



Analysis of recurrent neural networks via property-directed verification of surrogate models

Igor Khmelnitsky^{1,2} · Daniel Neider⁸ · Rajarshi Roy³ · Xuan Xie⁹ · Benoît Barbot⁴ · Benedikt Bollig¹ · Alain Finkel^{1,7} · Serge Haddad^{1,2} · Martin Leucker⁵ · Lina Ye^{1,2,6}

Accepted: 18 October 2022 / Published online: 16 November 2022
© The Author(s) 2022

Abstract

This paper presents a property-directed approach to verifying recurrent neural networks (RNNs). To this end, we learn a deterministic finite automaton as a *surrogate model* from a given RNN using active automata learning. This model may then be analyzed using *model checking* as a verification technique. The term *property-directed* reflects the idea that our procedure is guided and controlled by the given property rather than performing the two steps separately. We show that this not only allows us to discover *small* counterexamples fast, but also to generalize them by pumping toward faulty flows hinting at the underlying error in the RNN. We also show that our method can be efficiently used for *adversarial robustness certification* of RNNs.

Keywords Explainable AI · Neural network verification · Active learning

The first four authors contributed equally, the remaining authors are ordered alphabetically. This work was partly supported by: the PHC PROCOPE 2020 Project LeaRNNify (Number 44707TK), funded by DAAD and Campus France; the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Grant Numbers 434592664 and 459419731.

✉ Martin Leucker
leucker@isp.uni-luebeck.de

- ¹ Université Paris-Saclay, CNRS, ENS Paris-Saclay, LMF, Gif-sur-Yvette, France
- ² Inria, Paris, France
- ³ Max Planck Institute for Software Systems, Kaiserslautern, Germany
- ⁴ Université Paris-Est Créteil, Créteil, France
- ⁵ Institute for Software Engineering and Programming Languages, Universität zu Lübeck, Lübeck, Germany
- ⁶ CentraleSupélec, Université Paris-Saclay, Gif-sur-Yvette, France
- ⁷ Institut Universitaire de France, Paris, France
- ⁸ Safety and Explainability of Learning Systems Group, Carl von Ossietzky University of Oldenburg, Oldenburg, Germany
- ⁹ University of Alberta, Edmonton, Canada

1 Introduction

Rather than programming manually, it seems charming to simply provide examples of the intended input–output-behavior of a given function and derive the implementation of the function using algorithmic means. That is the promise of machine learning, in which often some form of classification problem is addressed by adjusting the parameters of some (deep) neural network until it fits the sample set appropriately.

While machine learning has shown to provide reasonable solutions in many cases, it may be expected that this approach also comes with a lot of deficiencies. Starting with the question of whether the examples are characteristic, it is unclear to which extent the learning algorithm considers the right aspects of the examples, whether the resulting system really realizes or closely approximates the right function, and whether it meets privacy standards. As such, sophisticated verification techniques for the learned artifacts seem extremely important.

In verification, the goal is to show that an implementation meets its specification. A huge number of verification algorithms have been developed over the past 50 years, mostly for program verification, as so-called formal methods. However, it has been noted [31] that formal specifications are often not

available when machine learning is used. In fact, the given set of examples, the training set, can be considered as (an approximation of) the specification. That said, many verification procedures can be considered as analysis algorithms parameterized by a formal specification. For example, while originally model checking [6] answers the question whether system S satisfies its specification ϕ , one can consider the specifications ϕ as a query (of some query language) and the model checking procedure applied on S as a generic analysis routine.

As such, it seems promising to apply the enormous contributions in program verification also for the analysis of neural networks. To do so, two general approaches seem possible. First, one could try to adapt the procedures developed in formal methods to analyze the artifacts encountered in machine learning. Second, one may translate the artifacts found in machine learning, e.g., the neural network, into formal models well studied in program verification. In this paper, which is an extended version of [28], we are following the latter approach. More precisely, we consider recurrent neural networks as the object of study and model checking as verification technique.

Recurrent neural networks (RNNs) are a state-of-the-art tool to represent and learn sequence-based models. They have applications in time-series prediction, sentiment analysis, and many more. In particular, they are increasingly used in safety-critical applications and act, for example, as controllers in cyber-physical systems [3]. Thus, there is a growing need for formal verification. However, research in this domain is only at the beginning. While model checking has been successfully used in practice and reached a certain level of industrial acceptance [25], a transfer to machine-learning algorithms has yet to take place. We will apply it on machine-learning artifacts rather than on the algorithm.

An emerging research stream aims at extracting state-based *surrogate models* from RNNs, such as finite automata [5,34,36,39,40,47], and, in general, we follow this approach in this paper as well. Finite automata turned out to be useful for understanding and analyzing all kinds of systems using testing or model checking. In other words, such models are also beneficial as an *explanation* of the underlying RNN.

A popular approach for extracting an automaton model from a given RNN is using active automata learning, based on the pioneering work by Angluin's L^* algorithm [4]. The general idea is to ask so-called *membership queries* to the underlying system (here the RNN) and *equivalence queries* whether the learned system is the right or a good enough approximation of the system to learn. Angluin's L^* has been improved in several ways especially regarding when to ask queries and how to process and store the information obtained by the queries, starting from [42] and [26], and resulting in [23] in which especially the space consumption is optimized. For further developments in automata learning using

L^* , we refer the readers to the work by Vaandrager [45] and for hints on choosing which learning algorithm for maximal efficiency, we refer to [1]. While our approach does not exploit all discussed optimizations to L^* , it is rather easy to incorporate them to improve performance.

The challenging step in L^* is the check whether the learned automaton is a good enough approximation of the RNN. A common technique follows statistical testing techniques and answers this question by comparing the two artifacts based on a random set of words. The work by Mayr and Yovine [36] uses probably approximately correct (PAC) learning [46]. In this paper, we provide an approach based on Hoeffding's inequality bound [20] also used in statistical model checking [30]. For sampling, we use several approaches, one being a mixture of A^* and plain sampling as described in [7].

In the field of formal verification, it has proven to be beneficial to run the extraction and verification process simultaneously. Moreover, the state space of RNNs tends to be prohibitively large, or even infinite, and so do incremental abstractions thereof. Motivated by these facts, we propose an intertwined approach to verifying RNNs, where, in an incremental fashion, grammatical inference and model checking go hand-in-hand. Our approach is inspired by black-box checking [41], which *exploits* the property to be verified *during* the verification process. Our procedure can be used to find misclassified examples or to verify a system that the given RNN controls, and we call the approach *property directed verification*.

Property-directed verification. Let us give a glimpse of our method. We consider an RNN R as a binary classifier of finite sequences over a finite alphabet Σ . In other words, R represents the set of strings that are classified as positive. We denote this set by $L(R)$ and call it the *language* of R . Note that $L(R) \subseteq \Sigma^*$. We would like to know whether R is compatible with a given specification A , written $R \models A$. Here, we assume that A is given as a (deterministic) finite automaton. Finite automata are algorithmically feasible, albeit having a reasonable expressive power: many abstract specification languages such as temporal logics or regular expressions can be compiled into finite automata [18].

But what does $R \models A$ actually mean? In fact, there are various options. If A provides a complete characterization of the sequences that are to be classified as positive, then \models refers to language equivalence, i.e., $L(R) = L(A)$. Note that this would imply that $L(R)$ is supposed to be a regular language, which may rarely be the case in practice. Therefore, we will focus on checking inclusion $L(R) \subseteq L(A)$, which is more versatile as we explain next.

Suppose N is a finite automaton representing a negative specification, i.e., R must classify words in $L(N)$ as negative at any cost. In other words, R does not produce false positives. This amounts to checking that $L(R) \subseteq L(\bar{N})$

where \bar{N} is the “complement automaton” of N . For instance, assume that R is supposed to recognize valid XML documents over a finite predefined set of tags. Seen as a set of strings, this is not a regular language. However, we can still check whether $L(R)$ only contains words where every opening tag $\langle \text{tag-name} \rangle$ is eventually followed by a closing tag $\langle / \text{tag-name} \rangle$ (while the number of opening and the number of closing tags may differ). As negative specification, we can then take an automaton N accepting the corresponding *regular* set of strings. For example, $\langle \text{book} \rangle \langle \text{author} \rangle \langle / \text{author} \rangle \langle \text{author} \rangle \langle / \text{book} \rangle \in L(N)$, since the second occurrence of $\langle \text{author} \rangle$ is not followed by some $\langle / \text{author} \rangle$ anymore. On the other hand, we have $\langle \text{book} \rangle \langle \text{author} \rangle \langle \text{author} \rangle \langle / \text{author} \rangle \langle / \text{book} \rangle \in L(\bar{N})$ because $\langle \text{book} \rangle$ and $\langle \text{author} \rangle$ are always eventually followed by their closing counterpart.

Symmetrically, suppose P is a finite automaton representing a *positive* specification so that we can find false negative classifications: If P represents the words that R must classify as *positive*, we would like to know whether $L(P) \subseteq L(R)$. Our procedure can be run using the complement of P as specification and inverting the outputs of R , i.e., we check, equivalently, $L(\bar{R}) \subseteq L(\bar{P})$.

An important instance of this setting is *adversarial robustness certification*, which measures a neural network’s resilience against adversarial examples. Given a (regular) set of words L classified as positive by the given RNN, the RNN is *robust* wrt. L if slight modifications in a word from L do not alter the RNN’s judgment. This notion actually relies on a distance function. Then, P is the set of words whose distance to a word in L is bounded by a predefined threshold, which is regular for several popular distances such as the *Hamming* or *Levenshtein distance*. Similarly, we can also check whether the neighborhood of a regular set of words preserves a negative classification.

In all these cases, we are faced with the question of whether the language of an RNN R is contained in the (regular) language of a finite automaton A . Our approach to this problem relies on black-box checking [41], which has been designed as a combination of model checking and testing in order to verify finite-state systems and is based on Angluin’s L^* learning algorithm [4]. L^* produces a sequence of *hypothesis* automata based on queries to R . Every such hypothesis \mathcal{H} may already share some structural properties with R . So, instead of checking conformance of \mathcal{H} with R , it is worthwhile to first check $L(\mathcal{H}) \subseteq L(A)$ using classical model-checking algorithms. If the answer is affirmative, we apply statistical model checking to check $L(R) \subseteq L(\mathcal{H})$ to confirm the result. Otherwise, a counterexample is exploited to refine \mathcal{H} , starting a new cycle in L^* . Just like in black-box checking, our experimental results suggest that the process of interweaving automata learning and model checking is beneficial in the verification of RNNs and offers advantages over

more obvious approaches such as (pure) statistical model checking or running automata extraction and model checking in sequence. A further key advantage of our approach is that, unlike in statistical model checking, we often find a *family* of counterexamples, in terms of loops in the hypothesis automaton, which testify conceptual problems of the given RNN.

Note that, though we only cover the case of binary classifiers, our framework is in principle applicable to multiple labels using one-vs-all classification.

Related Work. Mayr and Yovine describe an adaptation of the PAC variant of Angluin’s L^* algorithm that can be applied to neural networks [36]. As L^* is not guaranteed to terminate when facing non-regular languages, the authors impose a bound on the number of states of the hypotheses and on the length of the words for membership queries. In [34,37], Mayr et al. propose *on-the-fly property checking* where one learns an automaton approximating the intersection of the RNN language and the complement of the property to be verified. Like the RNN, the property is considered as a black box, only decidability of the word problem is required. Therefore, the approach is suitable for non-regular specifications.

Weiss et al. introduce a different technique to extract finite automata from RNNs [47]. It also relies on Angluin’s L^* but, moreover, uses an orthogonal abstraction of the given RNN to perform equivalence checks between them.

The paper [3] studies formal verification of systems where an RNN-based agent interacts with a linearly definable environment. The verification procedure proceeds by a reduction to feed-forward neural networks (FFNNs). It is complete and fully automatic. This is at the expense of the expressive power of the specification language, which is restricted to properties that only depend on bounded prefixes of the system’s executions. In our approach, we do not restrict the kind of regular property to verify. The work [24] also reduces the verification of RNNs to FFNN verification. To do so, the authors calculate inductive invariants, thereby avoiding a blowup in the network size. The effectiveness of their approach is demonstrated on audio signal systems. Like in [3], a time interval is imposed in which a given property is verified.

For adversarial robustness certification, Ryou et al. [43] compute a convex relaxation of the nonlinear operations found in the recurrent cells for certifying the robustness of RNNs. The authors show the effectiveness of their approach in speech recognition. Besides, MARBLE [16] builds a probabilistic model to quantize the robustness of RNNs. However, these approaches are white-box based and demand the full structure and information of neural networks. Instead, our approach is based on learning with black-box checking.

Elboher et al. present a counter-example guided verification framework whose workflow shares similarities with our property-guided verification [17]. However, their approach addresses FFNNs rather than RNNs. For recent progress in

the area of safety and robustness verification of deep neural networks, see [29].

Outline. In Sect. 2, we recall basic notions such as RNNs and finite automata. Section 3 describes two basic algorithms for the verification of RNNs, before we present property-directed verification in Sect. 4. How to handle adversarial robustness certification is discussed in Sect. 5. The experimental evaluation and a thorough discussion can be found in Sect. 6. This paper extends [28] by a more comprehensive introduction and overview to verification of neural networks, by more elaborated explanations, full proofs of all theorems and lemmas and by using an A*-based heuristics for equivalence checks as well as an enriched evaluation.

2 Preliminaries

In this section, we provide definitions of basic concepts such as languages, recurrent neural networks, finite automata, and Angluin's L* algorithm.

Words and Languages. Let Σ be an alphabet, i.e., a non-empty finite set, whose elements are called *letters*. A (finite) word w over Σ is a sequence $a_1 \dots a_n$ of letters $a_i \in \Sigma$. The length of w is defined as $|w| = n$. The unique word of length 0 is called the *empty word* and denoted by λ . We let Σ^* refer to the set of all words over Σ . Any set $L \subseteq \Sigma^*$ is called a *language* (over Σ). Its complement is $\bar{L} = \{w \in \Sigma^* \mid w \notin L\}$. For two languages $L_1, L_2 \subseteq \Sigma^*$, we let $L_1 \setminus L_2 = L_1 \cap \bar{L}_2$. The symmetric difference of L_1 and L_2 is defined as $L_1 \oplus L_2 = (L_1 \setminus L_2) \cup (L_2 \setminus L_1)$.

Probability Distributions. In order to sample words over Σ , we assume a probability distribution $(p_a)_{a \in \Sigma}$ on Σ (by default, we pick the uniform distribution) and a "termination" probability $p \in (0, 1]$. Together, they determine a natural probability distribution on Σ^* given, for $w = a_1 \dots a_n \in \Sigma^*$, by $\Pr(w) = p_{a_1} \cdot \dots \cdot p_{a_n} \cdot (1 - p)^n \cdot p$. According to the geometric distribution, the expected length of a word is $(1/p) - 1$, with a variance of $(1 - p)/p^2$. Let $0 < \varepsilon < 1$ be an error parameter and $L_1, L_2 \subseteq \Sigma^*$ be languages. We call L_1 ε -approximately correct wrt. L_2 if $\Pr(L_1 \setminus L_2) = \sum_{w \in L_1 \setminus L_2} \Pr(w) < \varepsilon$.

Finite Automata and Recurrent Neural Networks. We employ two kinds of language acceptors: finite automata and recurrent neural networks.

Recurrent neural networks (RNNs) are a generic term for artificial neural networks that process sequential data. They are particularly suitable for classifying sequences of varying length, which is essential in domains such as natural language processing (NLP) or time-series prediction. For the purposes of this paper, we follow recent literature on extract-

ing surrogate models from RNNs [8,35–37,48] and make two assumptions on RNNs:

1. We assume that the inputs to an RNN are a finite set of symbols. While usually the symbols are vectors in one-hot encoding, we abstract away from such implementation details and simply rely on a finite alphabet Σ .
2. We assume that the RNNs are a binary (or a one-vs-all) classifier.

One typical application of RNNs with such assumptions is sentimental analysis [33] where the task is to predict whether a text (e.g., a movie review) expresses positive or negative opinion.

The above assumptions, mathematically speaking, render an RNN R to be an effective function $R : \Sigma^* \rightarrow \{0, 1\}$ with a language defined as $L(R) = \{w \in \Sigma^* \mid R(w) = 1\}$. Its complement \bar{R} is defined by $\bar{R}(w) = 1 - R(w)$ for all $w \in \Sigma^*$. There are several ways to effectively represent R . Among the most popular architectures are (simple) Elman RNNs, long short-term memory (LSTM) [19], and GRUs [13]. Their expressive power depends on the exact architecture, but generally goes beyond the power of finite automata, i.e., the class of regular languages.

A *deterministic finite automaton (DFA)* over Σ is a tuple $A = (Q, \delta, q_0, F)$ where Q is a finite set of states, $q_0 \in Q$ is the initial state, $F \subseteq Q$ is the set of final states, and $\delta : Q \times \Sigma \rightarrow Q$ is the transition function. We assume familiarity with basic automata theory and leave it at mentioning that the language $L(A)$ of A is defined as the set of words from Σ^* that δ guides into a final state when starting in q_0 . That is, for the complement DFA $\bar{A} = (Q, \delta, q_0, Q \setminus F)$, we get $L(\bar{A}) = \bar{L}(A) = \Sigma^* \setminus L(A)$. It is well known that high-level specifications such as LTL formulas over finite words [18] or regular expressions can be compiled into corresponding DFAs.

We sometimes use RNNs and DFAs synonymously for their respective languages. For example, we say that R is ε -approximately correct wrt. A if $L(R)$ is ε -approximately correct wrt. $L(A)$.

Angluin's Algorithm. Angluin introduced L*, a classical instance of a learning algorithm in the presence of a minimally adequate teacher (MAT) [4]. We do not detail the algorithm here but only define the interfaces that we need to embed L* into our framework. Given any regular language $L \subseteq \Sigma^*$, the algorithm L* eventually outputs the unique minimal DFA \mathcal{H} such that $L(\mathcal{H}) = L$. The crux is that, while Σ is given, L is a priori unknown and can only be accessed through *membership queries (MQ)* and *equivalence queries (EQ)*:

(MQ) $w \stackrel{?}{\in} L$ for a given word $w \in \Sigma^*$. Thus, the answer is either yes or no.

(EQ) $L(\mathcal{H}) \stackrel{?}{=} L$ for a given DFA \mathcal{H} . Again, the answer is either yes or no. If the answer is no, one also gets a counterexample word from the symmetric difference $L(\mathcal{H}) \oplus L$.

Essentially, L^* asks MQs until it considers that it has a consistent dataset to come up with a hypothesis DFA \mathcal{H} , which then undergoes an EQ. If the latter succeeds, then the algorithm stops. Otherwise, the counterexample and possibly more membership queries are used to refine the hypothesis. The algorithm provides the following guarantee: If MQs and EQs are answered according to a given regular language $L \subseteq \Sigma^*$, then the algorithm eventually outputs, after polynomially¹ many steps, the unique minimal DFA \mathcal{H} such that $L(\mathcal{H}) = L$.

3 Verification approaches

Before we present (in Sect. 4) our method of verifying RNNs, we here describe two simple approaches. The experiments will later compare all three algorithms wrt. their performance.

Statistical model checking (SMC). One obvious approach for checking whether the RNN under test R satisfies a given specification A , i.e., to check whether $L(R) \subseteq L(A)$, is by a form of random testing. The idea is to generate a finite test suite $T \subset \Sigma^*$ and to check, for each $w \in T$, whether for $w \in L(R)$ also $w \in L(A)$ holds. If not, each such w is a *counterexample*. On the other hand, if none of the words turns out to be a counterexample, the property holds on R with a certain error probability. The algorithm is sketched as Algorithm 1.

Note that the test suite is sampled according to a probability distribution on Σ^* . Recall that our choice depends on two parameters: a probability distribution on Σ and a “termination” probability, both are described in Sect. 2.

Algorithm 1: SMC

Input: RNN R , DFA A , $\varepsilon, \gamma \in (0, 1)$

```

1 for  $i = 1, \dots, \log(2/\gamma)/(2\varepsilon^2)$  do
2    $w \leftarrow \text{sampleWord}()$ 
3   if  $w \in L(R) \setminus L(A)$  then
4     return “Counterexample  $w$ ”
5
6 end
7 return “Property satisfied”

```

¹ In the index of the right congruence associated with L and in the size of the longest counterexample obtained as a reply to an EQ.

Algorithm 2: AAMC

Input: RNN R and DFA A

```

1  $A_R \leftarrow \text{Approximation}(R)$ 
2 if  $\exists w \in L(A_R) \setminus L(A)$  then
3   return “Counterexample  $w$ ”
4 else return “Property satisfied”

```

Algorithm 3: PDV

Input: RNN R , DFA A , $\varepsilon, \gamma \in (0, 1)$

```

1 Initialize  $L^*$ 
2 while true do
3    $\mathcal{H} \leftarrow$  hypothesis provided by  $L^*$ 
4   Check  $L(\mathcal{H}) \subseteq L(A)$ 
5   if  $L(\mathcal{H}) \subseteq L(A)$  then
6     Check  $L(R) \subseteq L(\mathcal{H})$  using Alg. 1
7     if  $L(R) \subseteq L(\mathcal{H})$  then
8       return “Property satisfied”
9     else Feed counterexample to  $L^*$ 
10  else
11    Let  $w \in L(\mathcal{H}) \setminus L(A)$ 
12    if  $w \in L(R)$  then
13      return “Counterexample  $w$ ”
14    else Feed counterexample  $w$  to  $L^*$ 
15  end
16 end

```

Theorem 1 (Correctness of SMC) *If Algorithm 1, with $\varepsilon, \gamma \in (0, 1)$, terminates with “Counterexample w ”, then w is mistakenly classified by R as positive. If it terminates with “Property satisfied”, then R is ε -approximately correct wrt. A with probability at least $1 - \gamma$.*

Proof If the algorithm terminates with “Counterexample w ”, we have $w \in L(R) \setminus L(A)$. Thus, w is mistakenly classified. Using the sampling described in Sect. 2, denote by \hat{p} the probability to pick $w \in \Sigma^*$ such that $w \in L(R)$ and $w \notin L(A)$. Taking $n = \log(2/\varepsilon)/(2\gamma^2)$ random samples where m of them are counter examples, by Hoeffding’s inequality bound [20] we get that $P(\hat{p} \notin [\frac{m}{n} - \varepsilon, \frac{m}{n} + \varepsilon]) < \gamma$. Therefore, if Algorithm 1 terminates without finding any counterexamples we get that R is ε -approximately correct wrt. A with probability at least $1 - \gamma$. \square

While the approach works in principle, it has several drawbacks for its practical application. The size of the test suite may be quite huge and it may take a while both finding a counterexample or proving correctness.

Moreover, the correctness result and the algorithm assume that the words to be tested are chosen according to a random distribution that somehow also has to take into account the RNN as well as the property automaton.

It has been reported that this method does not work well in practice [47] and our experiments support these findings. *Automaton Abstraction and Model Checking (AAMC)*. As model checking is mainly working for finite-state systems, a

straightforward idea would be to (a) *approximate* the RNN R by a finite automaton A_R such that $L(R) \approx L(A_R)$ and (b) to check whether $L(A_R) \subseteq L(A)$ using model checking. The algorithmic schema is depicted in Algorithm 2.

Here, we can instantiate Approximation() by the DFA-extraction algorithms from [36] or [47]. In fact, for approximating an RNN by a finite-state system, several approaches have been studied in the literature, which can be, roughly, divided into two approaches: (a) *abstraction* and (b) *automata learning*. In the first approach, the state space of the RNN is mapped to equivalence classes according to certain predicates. The second approach uses automata-learning techniques such as Angluin's L^* . The approach [47] is an intertwined version combining both ideas.

Therefore, there are different instances of AAMC, varying in the approximation approach. Note that, for verification as language inclusion, as considered here, it actually suffices to learn an over-approximation A_R such that $L(R) \subseteq L(A_R)$.

While the approach seems promising at first hand, its correctness has two glitches. First, the result "Property satisfied" depends on the quality of the approximation. Second, any returned counterexample w may be *spurious*: w is a counterexample with respect to A_R satisfying A but may not be a counterexample for R satisfying A . If $w \in L(R)$, then it is indeed a counterexample, but if not, it is spurious—an indication that the approximation needs to be refined. If the automaton is obtained using abstraction techniques (such as predicate abstraction) that guarantee over-approximations, well-known principles like CEGAR [14] may be used to refine it. In the automata-learning setting, w may be used as a counterexample for the learning algorithm to improve the approximation. Repeating the latter idea suggests an interplay between automata learning and verification—and this is the idea that we follow in the next section. However, rather than starting from some approximation with a certain quality that is later refined according to the RNN and the property, we perform a direct, *property-directed* approach.

4 Property-directed verification of RNNs

We are now ready to present our algorithm for property-directed verification (PDV). The underlying idea is to replace the EQ in Angluin's L^* algorithm with a combination of classical model checking and statistical model checking, which are used as an alternative to EQs. This approach, which we call *property-directed verification of RNNs*, is outlined as Algorithm 3 and works as follows.

After initialization of L^* and the corresponding data structure, L^* automatically generates and asks MQs to the given RNN R until it comes up with a first hypothesis DFA \mathcal{H} (Line 3). In particular, the language $L(\mathcal{H})$ is consistent with the MQs asked so far.

At an early stage of the algorithm, \mathcal{H} is generally small. However, it already shares some characteristics with R . So it is worth checking, using standard automata algorithms, whether there is no mismatch yet between \mathcal{H} and A , i.e., whether $L(\mathcal{H}) \subseteq L(A)$ holds (Line 4). Because otherwise (Line 10), a counterexample word $w \in L(\mathcal{H}) \setminus L(A)$ is already a candidate for being a misclassified input for R . If indeed $w \in L(R)$, w is mistakenly considered positive by R so that R violates the specification A . The algorithm then outputs "Counterexample w " (Line 13). If, on the other hand, R happens to agree with A on a negative classification of w , then there is a mismatch between R and the hypothesis \mathcal{H} (Line 14). In that case, w is fed back to L^* to refine \mathcal{H} .

Now, let us consider the case that $L(\mathcal{H}) \subseteq L(A)$ holds (Line 5). If, in addition, we can establish $L(R) \subseteq L(\mathcal{H})$, we conclude that $L(R) \subseteq L(A)$ and output "Property satisfied" (Line 8). This inclusion test (Line 6) relies on statistical model checking using given parameters $\varepsilon, \gamma > 0$ (cf. Algorithm 1). If the test passes, we have some statistical guarantee of correctness of R (cf. Theorem 1). Otherwise, we obtain a word $w \in L(R) \setminus L(\mathcal{H})$ witnessing a discrepancy between R and \mathcal{H} that will be exploited to refine \mathcal{H} (Line 9).

Overall, in the event that the algorithm terminates, we have the following theorem that assures the soundness of a returned counterexample and provides the statistical guarantees on the property satisfaction, depending on the result of the algorithm:

Theorem 2 (Correctness of PDV) *Suppose Algorithm 3 terminates, using SMC for inclusion checking with parameters ε and γ . If it outputs "Counterexample w ", then w is mistakenly classified by R as positive. If it outputs "Property satisfied", then R is ε -approximately correct wrt. A with probability at least $1 - \gamma$.*

Proof Suppose the algorithm outputs "Counterexample w " in Line 13. Due to Lines 11 and 12, we have $w \in L(R) \setminus L(A)$. Thus, w is a counterexample.

Suppose the algorithm outputs "Property satisfied" in Line 8. By Lines 6 and 7, R is ε -approximately correct wrt. \mathcal{H} with probability at least $1 - \gamma$. That is, $P(L(R) \setminus L(\mathcal{H})) < \varepsilon$ with high probability. Moreover, by Line 4, $L(\mathcal{H}) \subseteq L(A)$. This implies that $L(R) \setminus L(A) \subseteq L(R) \setminus L(\mathcal{H})$ and, therefore, $P(L(R) \setminus L(A)) \leq P(L(R) \setminus L(\mathcal{H}))$. We deduce that R is ε -approximately correct wrt. A with probability at least $1 - \gamma$. \square

Although we cannot hope that Algorithm 3 will always terminate, we demonstrate empirically that it is an effective way for the verification of RNNs.

5 Adversarial robustness certification

Our method can especially be used for *adversarial robustness certification*, which is parameterized by a distance function $dist : \Sigma^* \times \Sigma^* \rightarrow [0, \infty]$ satisfying, for all words $w_1, w_2, w_3 \in \Sigma^*$: (1) $dist(w_1, w_2) = 0$ iff $w_1 = w_2$, (2) $dist(w_1, w_2) = dist(w_2, w_1)$, and (3) $dist(w_1, w_3) \leq dist(w_1, w_2) + dist(w_2, w_3)$. Popular distance functions are *Hamming distance* and *Levenshtein distance*. The Hamming distance between $w_1, w_2 \in \Sigma^*$ is the number of positions in which w_1 differs from w_2 , provided $|w_1| = |w_2|$ (otherwise, the distance is ∞). The Levenshtein distance (edit distance) between w_1 and w_2 is the minimal number of operations among substitution, insertion, and deletion that are required to transform w_1 into w_2 . For $L \subseteq \Sigma^*$ and $r \in \mathbb{N}$, we let $\mathcal{N}_r(L) = \{w' \in \Sigma^* \mid dist(w, w') \leq r \text{ for some } w \in L\}$ be the r -neighborhood of L . If L is regular and $dist$ is the Hamming or Levenshtein distance, then $\mathcal{N}_r(L)$ is regular (for efficient constructions of *Levenshtein automata* when L is a singleton, see [44]).

Let R be an RNN, $L \subseteq \Sigma^*$ be a regular language such that $L \subseteq L(R)$, $r \in \mathbb{N}$, and $0 < \varepsilon < 1$. We call R ε -adversarially robust (wrt. L and r) if $\Pr(\mathcal{N}_r(L) \setminus L(R)) < \varepsilon$. Accordingly, every word from $\mathcal{N}_r(L) \setminus L(R)$ is an *adversarial example*. Thus, checking adversarial robustness amounts to checking the inclusion $L(\bar{R}) \subseteq \overline{\mathcal{N}_r(L)}$ through one of the above-mentioned algorithms.

Note that, even when L is a finite set, $\mathcal{N}_r(L)$ can be too large for exhaustive exploration so that PDV, in combination with SMC, is particularly promising, as we demonstrate in our experimental evaluation.

From the definitions and Theorem 2, we get:

Lemma 1 *Suppose Algorithm 3, for input \bar{R} and a DFA A recognizing $\mathcal{N}_r(L)$, terminates, using SMC for inclusion checking with parameters ε and γ . If it outputs “Counterexample w ,” then w is an adversarial example. Otherwise, R is ε -adversarially robust (wrt. L and r) with probability at least $1 - \gamma$.*

Similarly, we can handle the case where $L \cap L(R) = \emptyset$. Then, R is ε -adversarially robust if $\Pr(L(R) \cap \mathcal{N}_r(L)) < \varepsilon$, and every word in $L(R) \cap \mathcal{N}_r(L)$ is an *adversarial example*. Overall, this case amounts to checking $L(R) \subseteq \overline{\mathcal{N}_r(L)}$.

6 Experimental evaluation

We now present an experimental evaluation of the three algorithms SMC, AAMC, and PDV, and provide a comparison of their performance on LSTM networks [19] (a variant of RNNs using LSTM units). The algorithms have been imple-

mented² in Python 3.6 using PyTorch 19.09 and Numpy library. The experiments of adversarial robustness certification were run on Macbook Pro 13 with the macOS. The other experiments were run on NVIDIA DGX-2 with an Ubuntu OS.

Optimization For Equivalence Queries. In [36], the authors implement AAMC but with an optimization that was originally shown in [4]. This optimization concerns the number of samples required for checking the equivalence between the hypothesis and the taught language. This number depends on ε, γ and the number of previous equivalence queries n and is calculated by $\frac{1}{\varepsilon} \left(\log \frac{1}{\gamma} + \log(2)(n+1) \right)$. We adopt this optimization in AAMC and PDV as well (Algorithm 2 in Line 1 and Algorithm 3 in Line 6).

6.1 Evaluation on randomly generated DFAs

Synthetic Benchmarks. To compare the algorithms, we implemented the following procedure, which generates a random DFA A_{rand} , an RNN R that learned $L(A_{\text{rand}})$, and a finite set of specification DFAs: (1) choose a random DFA $A_{\text{rand}} = (Q, \delta, q_0, F)$, with $|Q| \leq 30$, over an alphabet Σ with $|\Sigma| = 5$; (2) randomly sample words from Σ^* as described in Sect. 2 in order to create a training set and a test set; (3) train an RNN R with hidden dimension $20|Q|$ and $1 + |Q|/10$ layers—if the accuracy of R on the training set is larger than 95%, continue, otherwise restart the procedure; (4) choose randomly up to five sets $F_i \subseteq Q \setminus F$ to define specification DFAs $A_i = (Q, \delta, q_0, F \cup F_i)$. Using this procedure, we created 30 DFAs/RNNs and 138 specifications.

Experimental Results. Given an RNN R and a specification DFA A , we checked whether R satisfies A using Algorithms 1–3, i.e., SMC, AAMC, and PDV, with $\varepsilon, \gamma = 5 \cdot 10^{-4}$.

Table 1 summarizes the executions of the three algorithms on our 138 random instances. The columns of the table are as follows: (1) *Avg time* was counted in seconds and all the algorithms were timed out after 10 min; (2) *Avg len* is the average length of the found counterexamples (if one was found); (iii) *#Mistakes* is the number of random instances for which a mistake was found; (iv) *Avg MQs* is the average number of membership queries asked to the RNN.

Note that not only is PDV faster and finds more errors than AAMC, the average number of states of the final DFA is also much smaller: **26** states with PDV and **319** with AAMC. Furthermore, it asked more than 10 times less MQs to the RNN. Comparing PDV to SMC, it is 4.5 times faster and the average length of counterexamples it found is 10 times smaller, even though with a little fewer mistakes discovered.

² Available at <https://github.com/LeaRNNify/Property-directed-verification>.

Table 1 Comparison of verification algorithms

Type	Avg time(s)	Avg len	#Mistakes	Avg MQs
SMC	92	111	122	286063
AAMC	444	7	30	3701916
PDV	21	11	109	28318

6.2 Comparing equivalence queries

The PDV algorithm heavily depends on the procedure for checking the language inclusion $L(R) \subseteq L(\mathcal{H})$ between the hypothesized DFA \mathcal{H} and the RNN model R . Checking whether $L(R)$ is included in $L(\mathcal{H})$, however, is generally computationally infeasible, and thus, we resort to statistical model-checking that ensures PAC guarantees.

In statistical model-checking, one of the crucial steps is the technique used for random sampling of words from Σ^* . Thus, to determine how random sampling affects statistical model-checking, we investigate three different natural sampling techniques. We discuss them below.

1. *Random*: The first technique is to randomly sample words based on the natural probability distribution on Σ^* introduced in Sect. 2.
2. *DFA-based*: The second technique exploits the hypothesis DFA \mathcal{H} for random generation of words. To this end, we rely on the work by Bernardi and Giménez [11] who provide a linear algorithm for sampling words from DFAs. We built on top of their algorithm to generate words both accepted and rejected by \mathcal{H} . As a heuristic, in our implementation, we incorporate modifications to reduce the chances of sampling the same word multiple times.
3. *RNN-based*: The third technique exploits the RNN R for the random sampling of words. To this end, we rely on a technique similar to the one used by Barbot et al. [9]. The technique, in essence, is an A^* exploration in the rooted directed tree of all words Σ^* , where each vertex is a word $w \in \Sigma^*$ and its children are wa for $a \in \Sigma$. The exploration is guided by a scoring function $f: \Sigma^* \rightarrow \mathbb{R}$ that indicates how likely a word is to be accepted by the RNN. For our experiments, we define the scoring function to be as follows:

$$f(w) = \frac{1}{|\text{val}_R(w) - 0.5|}$$

where $\text{val}_R(w)$ is a value assigned by an RNN R to a word w for determining its acceptance. Precisely, the RNN R accepts w if and only if $\text{val}_R(w) > 0.5$. The scoring function f , defined above, prefers words w for which $\text{val}_R(w)$ is close to 0.5, since they can lead to words that can be accepted.

Table 2 Comparison of different equivalence queries (EQs) for PDV

EQ type	Avg time (s)	Mistakes	Avg MQs
Random	89	94	14647.2
DFA-based	37.6	109	12857.4
RNN-based	176.6	30	34259.6

To compare the performances, we run PDV using all of the sampling techniques on the synthetic benchmarks introduced in Sect. 6.1. Table 2 summarizes the comparison results of the sampling techniques. We compare them based on the average runtime of inclusion checks, the number of mistakes found, and the number of membership queries (MQs) required. The comparison was run on a machine with an Intel Core i7 processor (using up to 1.80 Ghz), with 24GB of RAM. The timeout for each run was set to be 300 s.

From the above table, we observe that the sampling technique DFA-based performs the best in terms of the runtime, the number of mistakes identified and the number of MQs required. The sampling technique Random, on the other hand, spends more resources to find mistakes since it samples words simply based on a probability distribution. The RNN-based performs worst in our experiments because the function f , as we defined, does not direct the search toward appropriate words that could be potential mistakes. A better choice of function f , and consequently, a better understanding of the RNN R can improve this sampling technique.

In summary, we conclude that the random sampling technique for inclusion checks in PDV can greatly affect the search for mistakes in an RNN.

Faulty Flows. One of the advantages of extracting DFAs in order to detect mistakes in a given RNN is the possibility to find not only one mistake but a “faulty flow.” For example, Fig. 1 shows one hypothesis DFA extracted with PDV, based on which we found a mistake in the corresponding RNN. The counterexample we found was *abcee*. One can see that the word *abce* is a loop in the DFA. Hence, we can suspect that this could be a “faulty flow.” Checking the words $w_n = (abce)^n e$ for $n \in \{1, \dots, 100\}$, we observed that, for any $n \in \{1, \dots, 100\}$, the word w_n was in the RNN language but not in the specification.

To automate the reasoning above, we did the following: Given an RNN R , a specification A , the extracted DFA \mathcal{H} , and the counterexample w : (1) build the cross product DFA $\mathcal{H} \times \bar{A}$; (2) for every prefix w_1 of the counterexample $w = w_1 w_2$, denote by s_{w_1} the state to which the prefix w_1 leads in $\mathcal{H} \times \bar{A}$ —for any loop ℓ starting from s_{w_1} , check if $w_n = w_1 \ell^n w_2$ is a counterexample for $n \in \{1, \dots, 100\}$; (3) if w_n is a counterexample for more than 20 times, declare a “faulty flow.” Using this procedure, we managed to find faulty flows in 81/109 of the counterexamples that were found by PDV.

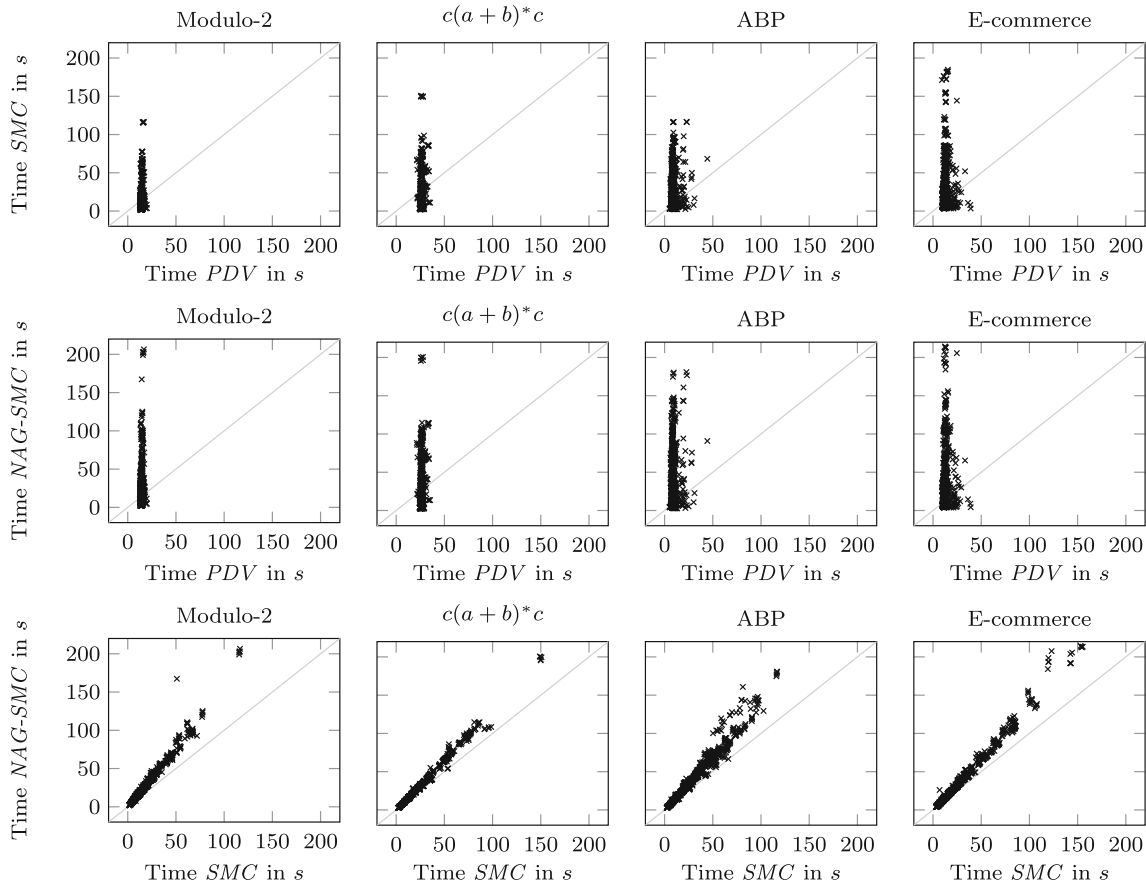


Fig. 2 Comparison of three algorithms on the regular languages




Fig. 1 Faulty flow in DFA extracted through PDV

6.3 Adversarial robustness certification

We also examined PDV for adversarial robustness certification, following the ideas explained in Sect. 5, both on synthetic and real-world examples.

Synthetic Benchmarks. For a given DFA (representing one of the languages described below), we randomly sampled words from Σ^* by using the DFA and created a training set and a test set. For RNN training, we proceeded like in step (3) for the benchmarks in Sect. 6.1. Moreover, for certification, we randomly sampled 100 positive words and 100 negative

words from the test set. For a given word w , we then let $L = \{w\}$ and considered $\mathcal{N}_r(L)$ where $r = 1, \dots, 5$.

Given an RNN R , we checked whether R satisfies adversarial robustness using the certification methods PDV, SMC, and neighborhood-automata generation SMC (NAG-SMC), with $\epsilon, \gamma = 0.01$. In SMC, we randomly modified the input word within a certain distance to generate words in the neighborhood. In NAG-SMC, on the other hand, we first generated a neighborhood automaton of the input word, and sampled words that are accepted by the automaton. Here, we followed the algorithm by Bernardi and Giménez [11], who introduce a method for generating a uniformly random word of length n in a given regular language with mean time bit-complexity $O(n)$.

Figure 2, which is a set of scatter plots, shows the results of the average time of executing the algorithms on the languages that we describe below. The x-axis and y-axis are both time in seconds, and each data point represents one adversarial robustness certification procedure. The length of words is from 50 to 500 and follows the normal distribution.

Simple Regular Languages. As a sanity check of our approach, we considered the following two regular languages and distance functions:




Fig. 3 Automaton for ABP

- $L_1 = ((a + b)(a + b))^*$ (also called *modulo-2 language*) with Hamming distance;
- $L_2 = c(a + b)^*c$ with distance function *dist* such that $\text{dist}(w_1, w_2)$ is the Hamming distance if $w_1, w_2 \in L_2$ and $|w_1| = |w_2|$, and $\text{dist}(w_1, w_2) = \infty$ otherwise.

The size of the Hamming neighborhood will exponentially grow with the distance.

The accuracies of the trained RNNs reached 100%. All three approaches successfully reported “adversarially robust” for the certified RNNs.

The first two diagrams on the first row of Fig. 2 compare the runtimes of PDV and SMC on the two regular-language datasets, resp., whereas the first two diagrams on the second row compare the runtimes of PDV and NAG-SMC. We make two main observations. First, on average, the running time of PDV (avg. 15.70 s) is faster than SMC (avg. 24.04 s) and NAG-SMC (avg. 32.5 s), which shows clearly that combining symbolically checking robustness on the extracted model and statistical approximation checking is more efficient than pure statistical approaches. Second, although SMC and NAG-SMC are able to certify short words (whose length is smaller than 30) faster, when the length of words is greater, they have to spend more time (which is more than 60 s) for certification. This is because, for short words, statistical approaches can easily explore the whole neighborhood, but when the neighborhood becomes larger and larger, this becomes infeasible.

The first two diagrams on the third row of Fig. 2 compare the running time of SMC and NAG-SMC, respectively. In general, SMC is faster than NAG-SMC. This is mainly because, for sampling random words from the neighborhood, using the algorithm proposed by Bernardi et al. [11] is slower than combining the *random.choice* function in the Python library and the corresponding modification.

Real-World Dataset. We used two real-world examples considered by Mayr and Yovine [36]. The first one is the alternating-bit protocol (ABP) shown in Fig. 3. However, we add a special letter *dummy* in the alphabet and a self-loop transition labeled with *dummy* on every state. In the figure, for readability, we replace the letter *dummy* using let-

ter *d*. The second example is a variant of an example from an e-commerce website [38], shown in Fig. 4. There are seven letters in the original automaton. Similarly, we also add letter *dummy* in the alphabet and also, in the self-loop transitions in every state (represented using *d* in the figure). In both the examples, we use the number of insertions of the letter *dummy* as the distance function.

The accuracies of the trained RNNs also reach 100%. For certification, the three approaches can certify the adversarial robustness for the RNNs as well.

The last two diagrams on the first (resp. second) row of Fig. 2 compare the runtime of PDV and SMC (resp. PDV and NAG-SMC) on the ABP and the E-commerce dataset. The data points in the first and second row have a vertical shape. The reason is that the running time of PDV is usually relatively stable (10–20 s), while the running time of SMC and NAG-SMC increases linearly with the word length.

The last two diagrams on the third row of Fig. 2 compare the runtimes of SMC and NAG-SMC on the two datasets. Here, the data points have a diagonal shape, but for NAG-SMC, when the word length is long (more than 300), it usually spends more time than SMC. This is mainly because it is inefficient to construct the neighborhood automaton and sample random words from the neighborhood.

6.4 RNNs identifying contact sequences

Contact tracing [27] has proven to be increasingly effective in curbing the spread of infectious diseases. In particular, analyzing contact sequences—sequences of individuals who have been in close contact in a certain order—can be crucial in identifying individuals who might be at risk during an epidemic. We, thus, look at RNNs which can potentially aid contact tracing by identifying possible contact sequences. However, in order to deploy such RNNs in practice, one would require them to be verified adequately. One does not want to alert individuals unnecessarily even if they are safe or overlook individuals who could be at risk.

In a real-world setting, one would obtain contact sequences from contact-tracing information available from, for instance, contact-tracing apps. However, such data is often difficult to procure due to privacy issues. Thus, in order to mimic a real-life scenario, we use data available from www.sociopatterns.org, which contains information about interaction of humans in public places (hospitals, schools, etc.) presented as temporal networks.

Formally, a *temporal network* $G = (V, E)$ [21] is a graph structure consisting of a set of vertices V and a set of labeled edges E , where the labels represent the timestamp during which the edge was active. Figure 5 is a simple temporal network, which can be perceived as contact graph of four workers in an office where edge labels represent the time of meeting between them. A *time-respecting path* $\pi \in V^*$ —a

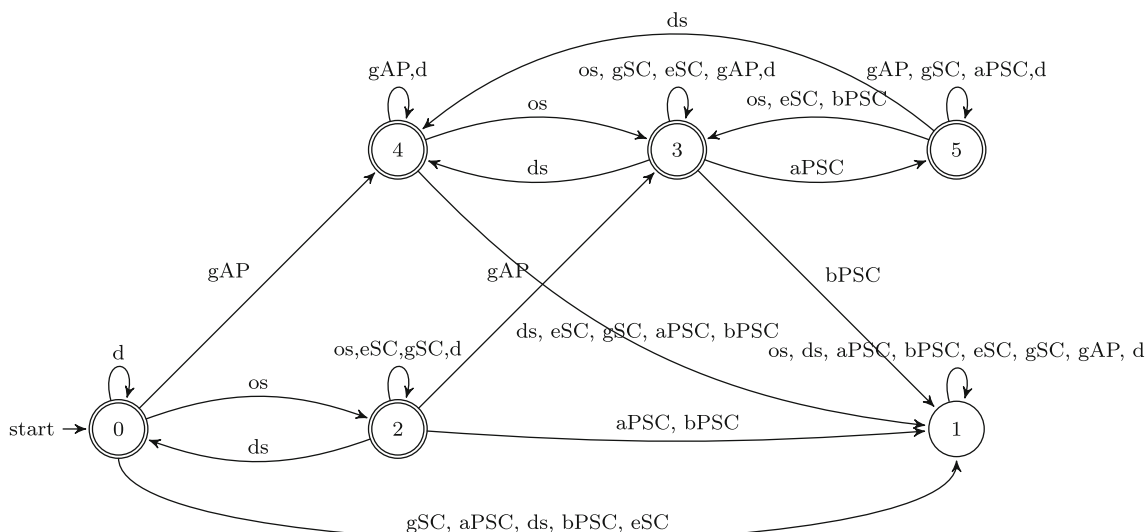


Fig. 4 Automaton for e-commerce example

sequence of vertices such that there exists a sequence of edges with increasing time labels—depicts a contact sequence in such a network. In the above example, *CDAB* is a time-respecting path while *ABCD* is not.

Benchmarks. For our experiment, given a temporal network *G*, we generated an RNN *R* recognizing contact sequences as follows:

1. We created training and test data for the RNN by generating (1) valid time-respecting paths (of lengths between 5 and 15) using labeled edges from *G*, and (2) invalid time-respecting paths, by considering a valid path and randomly introducing breaks in the path. The number of time-respecting paths in the training set is twice the size of the number of labeled edges in *G*, while the test set is one-fifth the size of the training set.
2. We trained RNN *R* with hidden dimension $|V|$ (minimum 100) as well as $\lfloor 2 + |V|/100 \rfloor$ layers on the training data. We considered only those RNNs that could be trained within 5 h with high accuracy (avg. 99%) on the test data.
3. We used a DFA that accepts all possible paths (disregarding the time labels) in the network as the specification, which would allow us to check whether the RNN learned unwanted edges between vertices.

Using this process, from the seven temporal networks, we generated seven RNNs and seven specification DFAs. We ran SMC, PDV, and AAMC on the generated RNNs, using the same parameters as used for the random instances.

Results. Table 3 notes the length of counterexample, the extracted DFA size (only for PDV and AAMC), and the running time of the algorithms. We make three main observations. First, the counterexamples obtained by PDV and




Fig. 5 Temporal network for contact between 4 people

AAMC (avg. length 2) are much more succinct than those by SMC (avg. length 13.1). Small counterexamples help in identifying the underlying error in the RNN, while long and random counterexamples provide much less insight. For example, from the counterexamples obtained from PDV and AAMC, we learned that the RNN overlooked certain edges or identified wrong edges. This result highlights the demerit of SMC, which has also been observed by [47]. Second, the running time of SMC and PDV (avg. 0.48 s and 0.41 s) is comparable, while that of AAMC is prohibitively large (avg. 655.68 s), indicating that model checking on small and rough abstractions of the RNN produces superior results. Third, the extracted DFA size, in case of AAMC (avg. size 124.14), is always larger compared to PDV (avg. size 2), indicating that RNNs are quite difficult to be approximated by small DFAs and this slows down the model-checking process as well. Again, our experiments confirm that PDV produces succinct counterexamples reasonably fast.

7 Conclusion

We proposed property-directed verification (PDV) as a new verification method for formally verifying RNNs with respect

Table 3 Results of model-checking algorithm on RNN identifying contact sequences

Case	Alg.	Counter-example len.	Extracted DFA size	Time (s)
Across	SMC	3		0.3
Kenyan	AAMC	2	328	624.76
Household	PDV	2	2	0.22
Workplace	SMC	2		0.23
	AAMC	2	111	604.99
	PDV	2	2	0.77
Highschool 2011	SMC	5		0.33
	AAMC	2	91	627.30
	PDV	2	2	0.19
Hospital	SMC	7		0.24
	AAMC	2	36	614.76
	PDV	2	2	0.006

Case	Alg.	Counter-example len.	Extracted DFA size	Time (s)
Within	SMC	2		0.28
Kenyan	AAMC	2	178	620.30
Household	PDV	2	2	0.27
Conference	SMC	71		1.51
	AAMC	2	38	876.19
	PDV	2	2	0.33
Workplace 2015	SMC	3		0.48
	AAMC	2	87	621.44
	PDV	2	2	1.11

to regular specifications, with adversarial robustness certification as one important application. It is straightforward to extend our ideas to the setting of Moore/Mealy machines supporting the setting of richer classes of RNN classifiers, but this is left as part of future work.

Recurrent neural networks have also often been employed for language processing. (Controlled) natural languages often have a context free nature and a context-free grammar might be the right object of study rather than finite automata. The work by Barbot et al. [8] presents an approach where instead of a finite automaton, a context-free grammar is learned as a surrogate model.

As future work, we plan to extend the PDV algorithm for the formal verification of RNN-based agent environment systems, and to compare it with the existing results [2,3]. Moreover, in the this paper, we define RNNs over a finite alphabet, while several applications of RNN, including speech [32] and hand-writing recognition [10], require defining them over an infinite (or very large) alphabet. To handle such RNNs, we plan to explore the possibility of using register automata that can classify data words over potentially infinite data domains as surrogate models [12,15,22].

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Aichernig, B.K., Tappler, M., Wallner, F.: Benchmarking combinations of learning and testing algorithms for active automata learning. In: Ahrendt, W., Wehrheim, H. (eds) Tests and Proofs—14th International Conference, TAP@STAF 2020, Bergen, Norway, June 22–23, 2020, Proceedings [postponed]. Lecture Notes in Computer Science, vol. 12165, pp. 3–22. Springer (2020). https://doi.org/10.1007/978-3-030-50995-8_1
2. Akintunde, M.E., Botoeva, E., Kouvaros, P., Lomuscio, A.: Formal verification of neural agents in non-deterministic environments. *Auton. Agents Multi Agent Syst.* **36**(1), 6 (2022)
3. Akintunde, M.E., Kevorchian, A., Lomuscio, A., Pirovano, E.: Verification of rnn-based neural agent-environment systems. In:

- Proceedings of AAAI 2019. pp. 6006–6013. AAAI Press (2019). <https://doi.org/10.1609/aaai.v33i01.33016006>
4. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987)
 5. Ayache, S., Eyraud, R., Goudian, N.: Explaining black boxes on sequential data using weighted automata. In: Proceedings of ICGI 2018. Proceedings of Machine Learning Research, vol. 93, pp. 81–103. PMLR (2018)
 6. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press, New York (2008)
 7. Barbot, B., Bollig, B., Finkel, A., Haddad, S., Khmelnitsky, I., Leucker, M., Neider, D., Roy, R., Ye, L.: Extracting context-free grammars from recurrent neural networks using tree-automata learning and a* search. In: Chandlee, J., Eyraud, R., Heinz, J., Jardine, A., van Zaanen, M. (eds) Proceedings of the Fifteenth International Conference on Grammatical Inference. Proceedings of Machine Learning Research, vol. 153, pp. 113–129. PMLR (23–27 Aug 2021). <https://proceedings.mlr.press/v153/barbot21a.html>
 8. Barbot, B., Bollig, B., Finkel, A., Haddad, S., Khmelnitsky, I., Leucker, M., Neider, D., Roy, R., Ye, L.: Extracting context-free grammars from recurrent neural networks using tree-automata learning and a* search. In: Chandlee, J., Eyraud, R., Heinz, J., Jardine, A., Zaanen, M. (eds) Proceedings of the 15th International Conference on Grammatical Inference, 23–27 August 2021, Virtual Event. Proceedings of Machine Learning Research, vol. 153, pp. 113–129. PMLR (2021). <https://proceedings.mlr.press/v153/barbot21a.html>
 9. Barbot, B., Bollig, B., Finkel, A., Haddad, S., Khmelnitsky, I., Leucker, M., Neider, D., Roy, R., Ye, L.: Extracting context-free grammars from recurrent neural networks using tree-automata learning and a* search. In: ICGI. Proceedings of Machine Learning Research, vol. 153, pp. 113–129. PMLR (2021)
 10. Bengio, Y., LeCun, Y., Nohl, C.R., Burges, C.J.C.: Lerec: a NN/HMM hybrid for on-line handwriting recognition. *Neural Comput.* **7**(6), 1289–1303 (1995)
 11. Bernardi, O., Giménez, O.: A linear algorithm for the random sampling from regular languages. *Algorithmica* **62**(1–2), 130–145 (2012)
 12. Bollig, B., Habermehl, P., Leucker, M., Monmege, B.: A robust class of data languages and an application to learning. *Log. Methods Comput. Sci.* (2014). [https://doi.org/10.2168/LMCS-10\(4:19\)2014](https://doi.org/10.2168/LMCS-10(4:19)2014)
 13. Cho, K., van Merriënboer, B., Gülçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using RNN encoder–decoder for statistical machine translation. In: Proceedings of the EMNLP. pp. 1724–1734. ACL (2014)
 14. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: Proceedings of CAV 2000. Lecture Notes in Computer Science, vol. 1855, pp. 154–169. Springer (2000)
 15. Decker, N., Habermehl, P., Leucker, M., Thoma, D.: Learning transparent data automata. In: Ciardo, G., Kindler, E. (eds) Application and Theory of Petri Nets and Concurrency—35th International Conference, PETRI NETS 2014, Tunis, Tunisia, June 23–27, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8489, pp. 130–149. Springer (2014). https://doi.org/10.1007/978-3-319-07734-5_8
 16. Du, X., Li, Y., Xie, X., Ma, L., Liu, Y., Zhao, J.: Marble: model-based robustness analysis of stateful deep learning systems. In: ASE 2020. pp. 423–435. IEEE (2020)
 17. Elboher, Y.Y., Gottschlich, J., Katz, G.: An abstraction-based framework for neural network verification. In: Proceedings of CAV 2020, Part I. Lecture Notes in Computer Science, vol. 12224, pp. 43–65. Springer (2020)
 18. Giacomo, G.D., Vardi, M.Y.: Synthesis for LTL and LDL on finite traces. In: Proceedings of IJCAI 2015. pp. 1558–1564. AAAI Press (2015)
 19. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
 20. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.* **58**(301), 13–30 (1963)
 21. Holme, P.: Temporal networks. In: Encyclopedia of Social Network Analysis and Mining, pp. 2119–2129. Springer (2014)
 22. Howar, F., Jonsson, B., Vaandrager, F.W.: Combining black-box and white-box techniques for learning register automata. In: Computing and Software Science, Lecture Notes in Computer Science, vol. 10000, pp. 563–588. Springer (2019)
 23. Isberner, M., Howar, F., Steffen, B.: The TTT algorithm: a redundancy-free approach to active automata learning. In: Bonakdarpour, B., Smolka, S.A. (eds) Runtime Verification—5th International Conference, RV 2014, Toronto, ON, Canada, September 22–25, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8734, pp. 307–322. Springer (2014). https://doi.org/10.1007/978-3-319-11164-3_26
 24. Jacoby, Y., Barrett, C.W., Katz, G.: Verifying recurrent neural networks using invariant inference. *CoRR* **abs/2004.02462** (2020)
 25. Karna, A.K., Chen, Y., Yu, H., Zhong, H., Zhao, J.: The role of model checking in software engineering. *Frontiers Comput. Sci.* **12**(4), 642–668 (2018). <https://doi.org/10.1007/s11704-016-6192-0>
 26. Kearns, M.J., Vazirani, U.V.: An Introduction to Computational Learning Theory. MIT Press (1994). <https://mitpress.mit.edu/books/introduction-computational-learning-theory>
 27. Keck, C.: Principles of Public Health Practice. Cengage Learning (2002)
 28. Khmelnitsky, I., Neider, D., Roy, R., Xie, X., Barbot, B., Bollig, B., Finkel, A., Haddad, S., Leucker, M., Ye, L.: Property-directed verification and robustness certification of recurrent neural networks. In: Hou, Z., Ganesh, V. (eds) Automated Technology for Verification and Analysis—19th International Symposium, ATVA 2021, Gold Coast, QLD, Australia, October 18–22, 2021. Proceedings. Lecture Notes in Computer Science, vol. 12971, pp. 364–380. Springer (2021). https://doi.org/10.1007/978-3-030-88885-5_24
 29. Kwiatkowska, M.Z.: Safety verification for deep neural networks with provable guarantees (invited paper). In: Proceedings of CONCUR 2019. Leibniz International Proceedings in Informatics (LIPIcs), vol. 140, pp. 1:1–1:5. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2019)
 30. Legay, A., Lukina, A., Traounez, L., Yang, J., Smolka, S.A., Grosu, R.: Statistical model checking. In: Steffen, B., Woeginger, G.J. (eds) Computing and Software Science—State of the Art and Perspectives, Lecture Notes in Computer Science, vol. 10000, pp. 478–504. Springer (2019). https://doi.org/10.1007/978-3-319-91908-9_23
 31. Leucker, M.: Formal verification of neural networks? In: Carvalho, G., Stolz, V. (eds) Formal Methods: Foundations and Applications—23rd Brazilian Symposium, SBMF 2020, Ouro Preto, Brazil, November 25–27, 2020. Proceedings. Lecture Notes in Computer Science, vol. 12475, pp. 3–7. Springer (2020). https://doi.org/10.1007/978-3-030-63882-5_1
 32. Lippmann, R.P.: Review of neural networks for speech recognition. *Neural Comput.* **1**(1), 1–38 (1989)
 33. Liu, B.: Sentiment Analysis—Mining Opinions, Sentiments, and Emotions. Cambridge University Press, Cambridge (2015)
 34. Mayr, F., Visca, R., Yovine, S.: On-the-fly black-box probably approximately correct checking of recurrent neural networks. In: Proceedings of CD-MAKE 2020. Lecture Notes in Computer Science, vol. 12279, pp. 343–363. Springer (2020)

35. Mayr, F., Visca, R., Yovine, S.: On-the-fly black-box probably approximately correct checking of recurrent neural networks. In: Proceedings of CD-MAKE 2020. Lecture Notes in Computer Science, vol. 12279, pp. 343–363. Springer (2020)
36. Mayr, F., Yovine, S.: Regular inference on artificial neural networks. In: Holzinger, A., Kieseberg, P., Tjoa, A.M., Weippl, E.R. (eds) Proceedings of CD-MAKE 2018. LNCS, vol. 11015, pp. 350–369. Springer (2018)
37. Mayr, F., Yovine, S., Visca, R.: Property checking with interpretable error characterization for recurrent neural networks. *Mach. Learn. Knowl. Extr.* **3**(1), 205–227 (2021)
38. Merten, M.: Active automata learning for real life applications. Ph.D. thesis, Dortmund University of Technology (2013)
39. Okudono, T., Waga, M., Sekiyama, T., Hasuo, I.: Weighted automata extraction from recurrent neural networks via regression on state spaces. In: Proceedings of AAAI 2020. pp. 5306–5314. AAAI Press (2020)
40. Omlin, C.W., Giles, C.L.: Extraction of rules from discrete-time recurrent neural networks. *Neural Netw.* **9**(1), 41–52 (1996)
41. Peled, D.A., Vardi, M.Y., Yannakakis, M.: Black box checking. *J. Autom. Lang. Comb.* **7**(2), 225–246 (2002)
42. Rivest, R.L., Schapire, R.E.: Inference of finite automata using homing sequences. *Inf. Comput.* **103**(2), 299–347 (1993). <https://doi.org/10.1006/inco.1993.1021>
43. Ryou, W., Chen, J., Balunovic, M., Singh, G., Dan, A.M., Vechev, M.T.: Fast and effective robustness certification for recurrent neural networks. CoRR [arXiv:2005.13300](https://arxiv.org/abs/2005.13300) (2020)
44. Schulz, K.U., Mihov, S.: Fast string correction with levenshtein automata. *Int. J. Doc. Anal. Recognit.* **5**(1), 67–85 (2002)
45. Vaandrager, F.W.: Model learning. *Commun. ACM* **60**(2), 86–95 (2017)
46. Valiant, L.G.: A theory of the learnable. *Commun. ACM* **27**(11), 1134–1142 (1984). <https://doi.org/10.1145/1968.1972>
47. Weiss, G., Goldberg, Y., Yahav, E.: Extracting automata from recurrent neural networks using queries and counterexamples. In: Proceedings of ICML 2018. Proceedings of Machine Learning Research, vol. 80, pp. 5244–5253. PMLR (2018)
48. Weiss, G., Goldberg, Y., Yahav, E.: Extracting automata from recurrent neural networks using queries and counterexamples. In: Dy, J.G., Krause, A. (eds) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10–15, 2018. Proceedings of Machine Learning Research, vol. 80, pp. 5244–5253. PMLR (2018). <http://proceedings.mlr.press/v80/weiss18a.html>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.