Deep Reinforcement Learning for Black-box Testing of Android Apps

ANDREA ROMDHANA, DIBRIS - Università degli Studi di Genova, FBK-ICT, Security & Trust Unit

ALESSIO MERLO, DIBRIS - Università degli Studi di Genova

MARIANO CECCATO, Università di Verona

PAOLO TONELLA, Università della Svizzera italiana

The state space of Android apps is huge, and its thorough exploration during testing remains a significant challenge. The best exploration strategy is highly dependent on the features of the app under test. Reinforcement Learning (RL) is a machine learning technique that learns the optimal strategy to solve a task by trial and error, guided by positive or negative reward, rather than explicit supervision. Deep RL is a recent extension of RL that takes advantage of the learning capabilities of neural networks. Such capabilities make Deep RL suitable for complex exploration spaces such as one of Android apps. However, state-of-the-art, publicly available tools only support basic, Tabular RL. We have developed ARES, a Deep RL approach for black-box testing of Android apps. Experimental results show that it achieves higher coverage and fault revelation than the baselines, including state-of-the-art tools, such as TimeMachine and Q-Testing. We also investigated the reasons behind such performance qualitatively, and we have identified the key features of Android apps that make Deep RL particularly effective on them to be the presence of chained and blocking activities. Moreover, we have developed FATE to fine-tune the hyperparameters of Deep RL algorithms on simulated apps, since it is computationally expensive to carry it out on real apps.

CCS Concepts: • **Software verification and validation** → **Empirical software validation**;

Additional Key Words and Phrases: Deep reinforcement learning, Android testing

ACM Reference format:

Andrea Romdhana, Alessio Merlo, Mariano Ceccato, and Paolo Tonella. 2022. Deep Reinforcement Learning for Black-box Testing of Android Apps. *ACM Trans. Softw. Eng. Methodol.* 31, 4, Article 65 (July 2022), 29 pages. <https://doi.org/10.1145/3502868>

This work was partially supported by: (1) “Istituto Poligrafico e Zecca dello Stato”; (2) Unige Curiosity Driven 2018 “User-defined Data Privacy in Android”; (3) H2020 project PRECRIME, funded under the ERC Advanced Grant 2017 Program (ERC Grant Agreement no. 787703); and (4) project MIUR 2018-2022 “Dipartimenti di Eccellenza.”

Authors’ addresses: A. Romdhana, DIBRIS - Università degli Studi di Genova, FBK-ICT, Security & Trust Unit; email: andrea.romdhana@dibris.unige.it; A. Merlo, DIBRIS - Università degli Studi di Genova; email: alessio.merlo@unige.it; M. Ceccato, Università di Verona; email: mariano.ceccato@univr.it; P. Tonella, Università della Svizzera italiana; email: paolo.tonella@usi.ch.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1049-331X/2022/07-ART65 \$15.00

<https://doi.org/10.1145/3502868>

1 INTRODUCTION

The complexity of mobile applications (hereafter, apps) keeps growing, as apps always provide more advanced services to the users. Nonetheless, it is of utmost importance that they work properly once they are published on app markets, as most of their success/failure depends on the user's evaluation. Therefore, an effective testing phase is fundamental to minimize the likelihood of app failures during execution.

However, automated testing of mobile apps is still an open problem, and the complexity of current apps makes their exploration trickier than in the past, as they can contain states that are difficult to reach and events that are hard to trigger.

There exist several approaches to automated testing of mobile apps that aim to maximize code coverage and bug detection during testing. Random testing strategies [21, 34] stimulate the **App Under Test (AUT)** by producing pseudo-random events. However, random exploration with no guidance may get stuck when dealing with complex transitions. Model-Based strategies [5, 22, 44] extract test cases from navigation models built employing static or dynamic analysis. If the model accurately reflects the AUT, then a Deep exploration can be achieved. Nonetheless, automatically constructed models tend to be incomplete and inaccurate. Structural strategies [7, 16, 35] generate coverage-oriented inputs using symbolic execution or evolutionary algorithms. These strategies are more powerful, since a specific coverage target guides them. However, they do not take advantage of past exploration successes to dynamically learn the most compelling exploration strategy.

Reinforcement Learning (RL) is a machine learning approach that does not require a labeled training set as input, since the learning process is guided by the positive or negative reward experienced during the tentative execution of the task. Hence, it represents a way to dynamically build an optimal exploration strategy by taking advantage of the past successful or unsuccessful moves.

RL has been extensively applied to the problem of GUI and Android testing [38, 41]. However, only the most basic RL (i.e., Tabular RL) has been applied to testing problems so far. In Tabular RL, the value of the state-action associations is stored in a fixed table. The advent of **Deep Neural Networks (DNN)** replaced Tabular approaches with Deep Learning ones, in which the action-value function is learned from the past positive and negative experiences made by one or more neural networks. When the state space to explore is extremely large (e.g., when an app has a significant amount of widgets), Deep RL has proved to be substantially superior to Tabular RL [9, 30, 42]. In this respect, we argue that the state space of Android apps is a good candidate for the successful adoption of Deep RL instead of Tabular RL for testing purposes.

This article presents the first Deep RL approach, ARES, for automated *black-box* testing of Android apps. ARES uses a DNN to learn the best exploration strategy from previous attempts. Thanks to such DNN, it achieves high scalability, general applicability, and the capability to handle complex app behaviors.

ARES implements multiple Deep RL algorithms that come with a set of configurable, often critical, hyperparameters. To speed up selecting the most appropriate algorithm for the AUT and the fine-tuning of its hyperparameters, we have developed another tool, FATE, which integrates with ARES.

FATE is a simulation environment that supports a rapid assessment of Android testing algorithms by running *synthetic* Android apps (i.e., abstract navigational models of real Android apps). The execution of a testing session on a FATE synthetic app is, on average, 10 to 100 times faster than the execution of the same session on the corresponding real Android app.

We applied ARES to two benchmarks made by 41 and 68 Android apps, respectively. The first benchmark compares the performance of the ARES algorithms, while the latter evaluates ARES w.r.t. the state-of-the-art testing tools for Android.

Experimental results confirmed the hypothesis that Deep RL outperforms Tabular RL in exploring the state space of Android apps, as ARES exposed the highest number of faults and obtained the highest code coverage. Furthermore, we carried out a qualitative analysis showing that the features of Android apps that make Deep RL particularly adequate include, among others, the presence of concatenated activities and blocking activities protected by authentication.

To sum up, this article provides the following advancements to state-of-the-art:

- ARES, the first publicly available testing approach based on Deep Reinforcement Learning, released as open-source;
- FATE, a simulation environment for fast experimentation of Android testing algorithms, also available as open-source;
- A thorough empirical evaluation of the proposed approach, whose replication package is publicly available to the research community.

2 BACKGROUND

After a general overview on RL, this section presents in more detail Tabular RL and Deep RL.

2.1 Overview on Reinforcement Learning

The objective of Reinforcement Learning [46] is to train an *agent* that interacts with some environment to achieve a given goal. The agent is assumed to be capable of sensing the *current state* of the *environment* and to receive a feedback signal, named *reward*, each time the agent takes an *action*.

At each timestep t , the agent receives an observation x_t , takes an action a_t that causes the transition of the environment from state s_t to state s_{t+1} . A state s_t is a complete description of the state of the environment. An observation x_t is a partial representation of the state, which may omit information. The agent also receives a scalar reward $R(x_t, a_t, x_{t+1})$, that quantifies the goodness of the last transition.

For simplicity, let us assume $x_t = s_t$ [4]. The behavior of an agent is represented by a *policy* π , i.e., a rule for making the decision on what action to take, based on the perceived state s_t . A policy can be:

- Deterministic: $a_t = \pi(s_t)$, i.e., a direct mapping between states and actions;
- Stochastic: $\pi(a_t|s_t)$, a probability distribution over actions, given their state.

DDPG [32] and TD3 [15] are examples of RL algorithms that learn a deterministic policy, while SAC [24] is an RL algorithm that learns a stochastic policy.

The standard mathematical formalism used to describe the agent environment is a **Markov Decision Process (MDP)**. An MDP is a 5-tuple, $\langle S, A, R, P, \rho_0 \rangle$, where :

- S is the set of all valid states,
- A is the set of all valid actions,
- $R : S \times A \rightarrow \mathbb{R}$ is the reward function, with $r_t = R(s_t, a_t, s_{t+1})$,
- $P : S \times A \rightarrow P(s)$ is the transition probability function, with $P(s_{t+1}|s_t, a_t)$ being the probability of transitioning into state s_{t+1} starting from state s_t and taking action a_t ,
- $\rho_0(s)$ is the starting state distribution.

Markov Decision Processes obey the *Markov property*: A transition only depends on the most recent state and action (and not on states/actions that precede the most recent ones).

The goal in RL is to learn a policy π that maximizes the so-called *expected return*, which can be computed as:

$$G(\tau) = \sum_{t=0}^{\infty} \gamma^t r_t.$$

A discount factor $\gamma \in (0, 1)$ is needed for convergence. When $t \rightarrow \infty$, if $\gamma = 1$, then the expected return and the RL algorithm would not converge. The discount factor determines how much the agent cares about rewards in the distant future relative to those in the immediate future: If γ is small, then the agent will tend to be myopic and only learn about actions that produce an immediate reward. If γ is close to 1, then the agent will evaluate each of its actions based on its estimated future rewards. τ is a sequence of states and actions in the environment $\tau = (a_0, s_0, a_1, s_1 \dots)$, named *trajectory* or *episode*. Testing an Android app can be seen as a task divided into finite-length episodes. Almost all reinforcement learning algorithms involve estimating *value functions*—functions of states (or *action-value functions*—functions of state-action pairs) [46] that estimate the goodness to perform a given action in a given state. The notion of goodness is defined in terms of expected future return. The rewards the agent can receive in the future depend on what actions it will take. The value function $V^\pi(s)$ is defined as the expected return starting in a state s and then acting according to a given policy π :

$$V^\pi(s) = E[G_t | s_0 = s].$$

The action-value function $Q^\pi(s, a)$ can be used to describe the expected return after taking an action a in state s and thereafter following the policy π :

$$Q^\pi(s, a) = E[G_t | s, a].$$

Correspondingly, we can define the *optimal value function*, $V^*(s)$, as the $V^\pi(s)$ that gives the highest expected return when starting in state s and acting according to the optimal policy in the environment. The *optimal action-value function*, $Q^*(s, a)$, gives the highest achievable expected return under the constraints that the process starts at state s , takes an action a and then acts according to the optimal policy in the environment.

A policy that chooses greedy actions only with respect to Q^* is optimal, i.e., knowledge of Q^* alone is sufficient for finding the optimal policy. As a result, if we have Q^* , then we can directly obtain the optimal action, $a^*(s)$, via $a^*(s) = \operatorname{argmax}_a Q^*(s, a)$. The optimal value function $V^*(s)$ and action-value function $Q^*(s, a)$ can be computed by means of a recursive relationship known as the *optimal Bellman equations*:

$$\begin{aligned} V^*(s_t) &= \max_a E[r(s_t, a_t) + \gamma V^*(s_{t+1})], \\ Q^*(s_t, a_t) &= E[r(s_t, a_t) + \gamma \max_{a_{t+1}} [Q^*(s_{t+1}, a_{t+1})]]. \end{aligned}$$

The optimal Bellman equations constrains the value of a state under an optimal policy to be equal to the expected return for the best action from that state. The optimal Bellman equations are actually a system of equations, one for each state. If the MDP of the environment is known, then the system of equations can be solved analytically, finding the optimal value function (or the optimal action-value function) and, at last, the optimal policy. Otherwise approximate solutions are found iteratively.

2.2 Tabular RL

Tabular techniques refer to RL algorithms where the state and action spaces are approximated using value functions defined by means of tables. In particular, *Q-Learning* [48] is one of the most

adopted algorithms of this family. Q-Learning aims to learn a so-called *Q-Table*, i.e., a table containing the value of each state-action pair. A Q-Table represents the current estimate of the action-value function $Q(s, a)$. Every time an action a_t is taken and a state s_t is reached, the associated state-action value in the Q-Table is updated as follows:

$$Q(s_t, a_t) := Q(s_t, a_t) + \alpha(r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)),$$

where α is the learning rate (between 0 and 1). If the learning rate is close to 0, then the agent exploits only prior knowledge, while when close to 1, the agent considers only the most up-to-date information. γ is the discount factor, applied to the future reward. Typically, γ ranges from 0.8 to 0.99 [15, 32, 40], while $\max_a Q(s_{t+1}, a)$ gives the maximum value for future rewards across all actions. It is used to update the reward for the current state-action pair.

RL algorithms based on the Tabular approach do not scale to high-dimensional problems, because in such cases it is difficult to manually define a good initial Q-Table. In case a good initial Q-Table is not available, convergence to the optimal table by means of the update rule described above is too slow [32].

2.3 Deep Reinforcement Learning

In large or unbounded discrete spaces, where representing all states and actions in a Q-Table is impractical, Tabular methods become highly unstable and incapable of learning a successful policy [40]. The rise of Deep Learning, relying on the powerful function approximation properties of Deep Neural Networks, has provided new tools to overcome these limitations. One of the first Deep Reinforcement Learning algorithms is **DQN (Deep Q-Networks)** [40].

DQN uses convolutional neural networks to approximate the computation of the action-value function Q^π . Training of such neural networks is achieved by using the so-called *experience replay* [33]. With experience replay the agent's experience e_t is stored within a buffer D named *replay buffer*. The experience e_t is defined as the tuple:

$$e_t = (s_t, a_t, r_t, s_{t+1}, d).$$

The tuple contains the state s_t of the environment, the action a_t taken in s_t , the reward r_t , and the next state of the environment s_{t+1} . The parameter d represents a binary value that indicates whether the transition has led the agent to a terminal state. To train the neural network, the algorithm occasionally retrieves random experience samples from the replay buffer. If the network learns only from consecutive samples of experience as they occurred sequentially in the environment, then the samples would be highly correlated and lead to inefficient learning [40]. Instead, taking random samples from the memory replay breaks this correlation.

While DQN can indeed solve problems with high-dimensional observation spaces, it can only handle discrete and low-dimensional action spaces. The recent advancements over DQN described in the following paragraphs (namely, DDPG, TD3, and SAC) overcome such limitations and allow dealing with high-dimensional action spaces.

Deep Deterministic Policy Gradient (DDPG). DDPG [32] is an *Actor-Critic* algorithm, i.e., it includes two roles: the *Critic*, which estimates the value function, and the *Actor*, which updates the policy π in the direction suggested by the Critic. It is based on a deterministic policy gradient [43] that can operate over continuous action spaces.

More specifically, the Critic has the objective of learning an approximation of the function $Q^*(s)$.

Suppose that the approximator is a neural network $Q_\phi(s, a)$ and that we have a set D of past experiences $e_t (s_t, a_t, r_t, s_{t+1}, d)$. The parameter list ϕ represents the coefficients of the neural network model, and the objective of the training phase is to optimize them, while set D is the previously

seen *replay buffer*. The replay buffer should be large enough to avoid overfitting. The loss function of the neural approximator is the so-called **Mean-Squared Bellman Error (MSBE)**, which measures how close $Q_\theta(s, a)$ is to satisfying the Bellman equation:

$$L(\phi, D) = E[(Q_\phi(s_t, a_t) - (r_t + \gamma(1 - d) \max_{a_{t+1}} Q_\phi(s_{t+1}, a_{t+1})))^2].$$

The term subtracted from $Q_\phi(s_t, a_t)$ is named the *target*, because minimization of MSBE makes the Q-function as close as possible to this value. Since the target depends recursively on the same parameter ϕ to train, MSBE minimization can become unstable. The solution is to use a second neural network, called *target network*, whose parameters are updated to be close to ϕ , but with some time delay that gives stability to the process.

The Actor's goal is to learn a deterministic policy $\pi_\theta(s)$ that maximizes $Q_\phi(s, a)$. Because the action space is continuous and we assume the Q-function is differentiable with respect to the action, we can use gradient ascent with respect to the policy parameter. In case the action space is non-differentiable, DDPG degenerates, performing like DQN [32].

Twin Delayed DDPG (TD3). Although DDPG can often achieve good performance, it tends to be susceptible to critical tuning of its hyperparameters. In fact, Fujimoto et al. [15] demonstrated that in Actor-Critic algorithms, such as DDPG, the policy update introduces errors, and overestimation of parameters ϕ . The researchers state that the introduced errors may be minimal, but raise two concerns: (1) the error may develop into a more significant bias over many updates if left unchecked, (2) an inaccurate value estimate may lead to poor policy update. TD3 [15] is an algorithm that addresses this issue by introducing three major changes, mostly on the Critic side: (1) *Clipped Double Q-Learning*: TD3 learns two Q-functions instead of one and uses the smaller of the two Q-values as the target in the MSBE function. Using the smaller Q-value for the target, and regressing towards that, helps mitigate overestimation in the Q-function. (2) *Delayed Policy Update*: TD3 updates the policy and the target networks less frequently than the Q-function. Delaying policy updates reduces per-update error and further improves performance. (3) *Target Policy Smoothing*: TD3 adds noise to the target action to make it harder for the policy to exploit Q-function errors by smoothing out Q across changes of the action.

Soft Actor Critic (SAC). The central feature of SAC [24] is entropy regularization. Using the entropy-regularized method, an agent gets a bonus reward at each timestep that is proportional to the entropy of the policy at that timestep. In fact, differently from TD3, the policy of SAC is non deterministic and inclusion of entropy in the reward aims at promoting policies with a wider spread of alternatives to choose stochastically from. The RL problem becomes the problem of finding the optimal policy π^* according to the following equation:

$$\pi^* = \operatorname{argmax}_{\pi} E \left[\sum_{t=0}^{\infty} \gamma^t (R(s_t, a_t, s_{t+1}) + \alpha H(\pi(\cdot|s_t))) \right].$$

The first term of the equation $E[\sum_{t=0}^{\infty} \gamma^t (R(s_t, a_t, s_{t+1}))]$ comes from standard RL algorithms, in which the objective is to maximize the expected sum of rewards. The second term $\alpha H(\pi(\cdot|s_t))$ contains the entropy H , which is directly controlled by the entropy regularization coefficient $\alpha > 0$. This parameter explicitly controls the explore-exploit tradeoff. With higher α the exploration is encouraged, while lower α corresponds to more exploitation.

3 ARES: APPROACH

This section describes **ARES (Application of REinforcement learning to android Software testing)**, our approach to black-box Android GUI testing based on Deep RL. Figure 1 shows an

Fig. 1. The ARES testing workflow. Information about widgets is extracted from the GUI and used to compute the state s_t and the reward r_t . b_i indicates the clickable buttons, $scroll_j$ indicates a portion of the GUI that we can scroll, and $edit_k$ an element that can be filled with text.

overview of the approach. The RL *environment* is represented by the Android **application under test (AUT)**, which is subject to several interaction steps. At each timestep, assuming the GUI state is s_t , and the reward is r_t , ARES first takes an action a_t . Then, it receives the new GUI state s_{t+1} of the AUT, and a reward r_{t+1} (not shown in Figure 1). Intuitively, if the new state s_{t+1} is similar to the prior state s_t , then the reward is negative. Otherwise, the reward is a large positive value. In this way, ARES promotes the exploration of new states in the AUT, under the assumption that this is useful to test the application more thoroughly.

The reward is used to update the neural network, which learns how to guide the Deep RL algorithm to explore the AUT in depth. The actual update strategy depends on the Deep RL algorithm used (either DDPG, TD3, or SAC).

3.1 Problem Formulation

To apply RL, we have to map the problem of Android black-box testing to the standard mathematical formalization of RL: an MDP, defined by the 5-tuple, $\langle S, A, R, P, \rho_0 \rangle$. Moreover, we have to map the testing problem onto an RL task divided into several finite-length episodes.

State Representation. Our approach is black-box, because it does not access the source code of the AUT. It only relies on the GUI of the AUT. The state $s_t \in S$ is defined as a combined state $(a_0, \dots, a_n, w_0, \dots, w_m)$. The first part of the state $a_0, \dots, a_n$ is a one-hot encoding of the current activity, i.e., a_i is equal to 1 only if the currently displayed activity is the i th activity, it is equal to 0 for all the other activities. In the second part of the state vector, w_j is equal to 1 if the j th widget is available in the current activity; it is equal to 0 otherwise.

Action Representation. User interaction events in the AUT are mapped to the action set A of the MDP. ARES infers executable events in the current state by analyzing the dumped widgets and their attributes (i.e., *clickable*, *long-clickable*, and *scrollable*). In addition to the widget-level actions, we use two system-level actions, namely, *toggle internet connection* and *rotate screen*. These system-level actions are the only system actions that can be easily tested. In fact, since Android version

7.0, testing other system-level actions (like those implemented in Stocat [44]) would depend on the Android version used [20, 45] and would require a rooted device [17].

Moreover, certain apps, such as apps in the Finance and Entertainment categories (e.g., Netflix, SkyGo, Amazon Prime, and Lloyds), do not start on a rooted device due to the root checks that are on the apps [18, 39, 47], compromising their testing.

Each action a is three-dimensional: The first dimension represents the widget ARES is going to interact with or the identifier of a system action. The second dimension specifies a string to be used as text input if needed. Actually, an index pointing to an entry in a dictionary of predefined strings is used for this dimension. The third dimension depends on the context: When the selected widget is both *clickable* and *long-clickable*, the third action determines which of the two actions to take. When ARES interacts with a *scrollable* object, the third dimension determines the scrolling direction.

Transition Probability Function. The transition function P determines which state the application can transit to after ARES has taken an action. In our case, this is decided solely by the execution of the AUT: ARES observes the process passively, collecting the new state after the transition has taken place.

Reward Function. The RL algorithm used by ARES receives a reward $r_t \in R$ every time it executes an action a_t . We define the following reward function:

$$r_t = \begin{cases} \Gamma_1 & \text{if } act(s_{t+1}) \notin act(E_t) \text{ or } crash, \\ -\Gamma_2 & \text{if } pack(act(s_{t+1})) \neq pack(AUT), \\ -\Gamma_3 & \text{otherwise,} \end{cases} \quad (1)$$

with $\Gamma_1 \gg \Gamma_2 \gg \Gamma_3$ (in our implementation $\Gamma_1 = 1,000$, $\Gamma_2 = 100$, $\Gamma_3 = 1$). To select the values of the reward, we conducted a preliminary experiment on a sub-sample of five apps (i.e., Antennapod, RedReader, Opentasks, Simple-Solitaire, and YalpStore). On these apps, we tested several types of rewards that are commonly used in literature: combination I ($\Gamma_1 = 1,000$, $\Gamma_2 = 100$, $\Gamma_3 = 1$), combination II ($\Gamma_1 = 1$, $\Gamma_2 = 1$, $\Gamma_3 = 0$), and combination III ($\Gamma_1 = 100$, $\Gamma_2 = 10$, $\Gamma_3 = 1$). We found that combination I gives the best results.

The exploration of ARES is divided into *episodes*. At time t , the reward r_t is high (Γ_1) if ARES was able to transition to an activity never observed during the current episode E_t (i.e., the next activity $act(s_{t+1})$ does not belong to the set of activities encountered so far in E_t): If a new episode is started at $t + 1$, then its set of activities is reset: $act(E_{t+1}) = \emptyset$.

Resetting the set of encountered activities at the beginning of each new episode is a technique that encourages ARES to visit and explore the highest number of activities in each episode to reinforce its explorative behaviors continuously. In contrast, if we provide the algorithm a significant reward only a few times (i.e., “sparsely”), then the information to learn the optimal state-action combinations might be insufficient. The algorithm might fail to reproduce the sequence of actions leading to a high reward in the future. In that case, the performance of the algorithm results poor. On the contrary, rewarding any activity change, regardless of their novelty, would encourage cycling behaviors [12].

The reward is high (Γ_1) also when a faulty behavior (*crash*) occurs. It is very low ($-\Gamma_2$) when the displayed activity does not belong to the AUT (i.e., the package of the current activity, $pack(act(s_{t+1}))$, is not the package of the AUT), as we aim to explore the current AUT only. In all other cases, the reward is moderately negative ($-\Gamma_3$), as the exploration remains inside the AUT, even if nothing new has been discovered.

4 ARES: IMPLEMENTATION

ARES features a custom environment based on the OpenAI Gym [10] interface, which is a de facto standard in the RL field. OpenAI Gym is a toolkit for designing and comparing RL algorithms and includes a number of built-in environments. It also includes guidelines for the definition of custom environments. Our custom environment interacts with the Android emulator [19] using the Appium Test Automator Framework [3].

4.1 Tool Overview

As soon as it is launched, ARES leverages Appium to dump the widget hierarchy of the GUI in the starting activity of the AUT. The widget hierarchy is analyzed by searching for clickable, long-clickable, and scrollable widgets. Afterward, these widgets are stored in a dictionary containing several associated attributes (e.g., resource-id, clickable, long-clickable, scrollable) and compose the *action vector*, i.e., the vector of executable actions in the current state. At each timestep, ARES takes an action according to the behavior of the exploration algorithm. Once the action has been fully processed, ARES extracts the new widget hierarchy from the current GUI and calculates its MD5 hash value. If it has the same MD5 value of the previous state, then ARES leaves the action vector unchanged. If the MD5 value does not match, then ARES updates the action vector. ARES performs the observation of the AUT state and returns the combined vector of activities and widgets. ARES organizes the testing of each app as a task divided into finite-length episodes. The goal of ARES is to maximize the total reward received during each episode. Every episode lasts at least 250 timesteps. Its duration is shorter only if the app crashes. To select the ideal episode boundaries, we conducted a preliminary experiment on a sample app. On this app, we trained the same algorithm by varying the episode length. Training characterized by short episodes results in poor performance due to the impossibility of exploring the application. Similarly, long episodes suffered from poor performance due to a low number of episodes completed. Once an episode comes to an end, the app is restarted, and ARES uses the acquired knowledge to explore the app more thoroughly in the next episode.

4.2 Application Environment

The application environment is responsible for handling the actions to interact with the AUT. Since the environment follows the guidelines of the Gym interface, it is structured as a class with two key functions. The first function `init(desired_capabilities)` is the initialization of the class. The additional parameter `desired_capabilities` consists of a dictionary containing the emulator setup and the application to be tested. The second function is the `step(a)` function, which takes an action `a` as a command and returns a list of objects, including observation (current AUT state) and reward.

4.3 Algorithm Implementation

ARES is a modular framework that adopts a plugin architecture to integrate the RL algorithm to use. Hence, extension with a new exploration algorithm can be easily achieved. In the current implementation, ARES provides five different exploration strategies: (1) random, (2) Q-Learning, (3) DDPG, (4) SAC, (5) TD3. The random algorithm interacts with the AUT by randomly selecting an action from those in the action vector. Compared to Monkey [21], our random approach performs better, since it selects only actions from the action vector. In fact, Monkey generates random, low-level events on the whole GUI, which could target no actual widget and then be discarded.

Our Q-Learning strategy implements the algorithm proposed by Watkins and Dayan [48]. The Deep RL algorithms available in ARES are DDPG, SAC, and TD3. Their implementation comes

Fig. 2. Prototype of a FATE model.

from the Python library *Stable Baselines* [25] and allows ARES to save the status of the neural network as a policy file at the end of the exploration. In this way, the policy can be loaded and reused on a new version of the AUT at a later stage, rather than restarting ARES from scratch each time a new AUT version is released. ARES is publicly available as open source software at <https://github.com/H2SO4T/ARES>.

4.4 Compatibility

ARES has been successfully tested on Windows 10, macOS 11.1 (and older), Ubuntu 20 (and older), and Scientific Linux 7.5. ARES is fully compliant with parallel execution and enables parallel experiments to be performed on emulators or real devices, handling each instance in a completely separate manner. ARES is also compatible with several Android versions (i.e., it has been successfully tested on Android 6.0, 7.0, 7.1, 8.0, 8.1, 9.0, and 10.0). Moreover, since ARES is based on the standard OpenAIGym, new algorithms and exploration strategies can be easily added to the tool.

5 FAST ANDROID TEST ENVIRONMENT (FATE)

Deep RL algorithms require fine-tuning, which is expensive on real apps. Therefore, we developed FATE, a simulation environment for fast Android testing. FATE models only the navigation constraints of Android apps, so it can efficiently compare alternative testing algorithms and quickly tune their corresponding hyperparameters. After this algorithm selection and tuning phase through FATE is completed, the selected algorithms and their configurations are ported to ARES to test real apps.

5.1 Model-based Prototyping

5.1.1 FATE Design. In FATE, developers model an Android app by means of a deterministic **Finite State Machine (FSM)** $\mathcal{F} = (\Sigma, S, s_0, \delta, F)$, where Σ is a set of events, S a set of states with s_0 the initial state and F the set of final states, and δ the state transition function $\delta : S \times \Sigma \longrightarrow 2^S$. The *states* S of the FSM correspond to the activities of the app, while the *events* Σ trigger the transitions between activities, which in turn are modeled as a transition table δ . Events represent the clickable widgets (e.g., buttons) available in each activity. Transitions have access to a set of *global variables* and possess, among others, the following attributes: *ID*, *type*, *active* (Boolean attribute), *guard* (Boolean expression that prevents the transition from being taken if it evaluates to false), *set* (new values to be assigned to global variables), *destination* (target activity, i.e., value of δ). A prototype of a FATE model is shown in Figure 2.

Fig. 3. FATE model of *Social Network*: Global variables are shown on the top-left; inputs on the bottom-left; red edges indicate non deterministic transitions.

5.1.2 FATE Implementation. Figure 3 shows the FATE model of the prototypical app *Social Network*. To build such a model, developers can use Ptolemy [29] to graphically draw an FSM that mimics the behavior of the application. While creating an FSM with Ptolemy is not mandatory in FATE, it simplifies the job of designing a logically correct model, thanks to the checks it performs. Then, FATE automatically translates the Ptolemy model (saved in XML format) into a JSON file that replicates the structure and behavior of the Ptolemy model. The JSON model translation has two main fields: `global_vars` and `nodes`. The first contains a list of global variables organized by name and value. The latter contains a list of all the activities. Each activity is characterized by a `node_id` and a list of corresponding node transitions, each including all the respective transition attributes.

The model of *Social Network* in Figure 3 contains a transition from `login` to `main_act`, which is subjected to the guard `user_pass == real_pass`, i.e., the entered password must be correct for the transition to be taken. In the JSON model, such transition is coded as:

```
1 "transition": {
2   "transition_id": 0,
3   "type": "button",
4   "active": true,
5   "guard": "user_pass == real_pass",
6   "set": null,
7   "destination": "main_act"
8 }
```

Another example of guarded transition is the one between `messages` and `chat_x`. The guard `count_messages >= 1` checks whether there exists at least one message from which a chat thread can be started. In the JSON model this is coded as:

```
1 "transition": {
2   "transition_id": 0,
3   "type": "button",
4   "active": true,
5   "guard": "count_messages >= 1",
6   "set": null,
7   "destination": "main_act"
8 }
```

In FATE, a Python environment compliant with the OpenAI Gym standard takes as input the JSON app model and tests it automatically using the selected algorithm. The available algorithms

are: (1) Random, (2) Q-Learning, (3) DDPG, (4) SAC, (5) TD3. FATE was built with modularity in mind, and new exploration algorithms can be easily added to the tool. Compared to testing an Android app through Espresso or Appium, FATE makes test case execution significantly faster, because there is no need to interact with the app via its GUI. Moreover, the application navigation logic is simulated by the transition function δ , making it usually much faster to execute. Consequently, developers can run a large number of experiments, evaluate multiple algorithms, check various algorithm or application configurations, and find the optimal set of hyperparameters, all of which would be prohibitively expensive to execute on a standard Android testing platform. A limitation of FATE is that its effectiveness in hyperparameter tuning depends on the fidelity of the app models created by the developers. Despite not being the developers of the apps under test, we have been able to define models for them that turned out to be sufficiently faithful. In fact (see Section 6.1), the results obtained on the real apps with ARES and on their models with FATE are very close to each other. FATE is publicly available as open source software at <https://github.com/H2SO4T/ARES>.

5.2 Representative Family of Models

For fast evaluation of the Deep RL algorithms implemented in ARES, we modeled four Android apps using FATE. Each model represents the generalization of the apps belonging to a specific family, such as Shopping category. To obtain a set of app models representing the most common apps used in everyday life, we inspected AppBrain (a website that aggregates Google Play Store statistics and rankings) [2]. We selected four different and representative categories from the top ten: Music & Audio, Lifestyle, Business, and Shopping. From each category, we then selected and modeled in FATE one prototypical app: *Player*, *Social Network*, *Bank*, and *Market Place*. Each model is configurable with a variable degree of complexity.

The simplest scenario is *Player*. It features a wide number of activities arranged in a tree-like structure. It reflects the generalization of various applications, including apps or app components, to manage the settings and stream/add/remove media contents. *Social Network* (see Figure 3) starts by prompting a login activity with fields for username and password. Following the login activity, we have several activities that replicate a standard social network behavior, including a basic chat activity. The presence of inner password-protected operations characterizes the *Bank* model. *Market Place* models a typical app for e-commerce: The user can search for goods, login, purchase products, and monitor the orders. The four representative app models used in this work are publicly available inside the FATE tool.

6 EVALUATION

We seek to address the following research questions, split between the following two studies:

Study 1 (FATE):

- **RQ1:** *Are the results of synthetic apps comparable to those of their translated counterparts?*
- **RQ2:** *Which Deep RL algorithm and which algorithm configuration performs better on the synthetic apps?*
- **RQ3:** *How does activity coverage vary as the model of the AUT becomes increasingly difficult to explore?*
- **RQ4:** *What are the features of the synthetic apps that allow Deep RL to perform better than Q-Learning?*

In Study 1, we want to understand if results obtained on synthetic app models run by FATE correlate with those obtained when executing the same apps (RQ1), once they are translated into

real Android apps in Java code that can be executed by ARES. In particular, we translated three synthetic apps to Java/Android: Social Network, Bank, and Market Place. We decided not to translate the “Player” model due to its simplicity. We compare the rankings of the algorithms produced, respectively, by FATE and by ARES. Text inputs are chosen from a dictionary of 40 strings, which include credentials necessary to pass through the login activities. Since in this study we also use synthetic apps generated from models (i.e., Player, Social Network, Bank, and Market Place), coverage is measured at the granularity of Android activities. In fact, there is no source code implementing the business logic. Each run has a length of 4,000 timesteps, close to an hour of testing in a real Android test setting. With FATE, 4,000 times steps are executed approximately in 200 seconds.

To answer RQ2, we take advantage of the fast execution granted by FATE to compare alternative RL algorithms on all synthetic apps and determine their optimal configuration (see Appendix A.1). Text inputs are chosen from a dictionary of 20 strings, which include credentials necessary to pass through the login activities. To account for non-determinism, we executed each algorithm 60 times for each hyperparameter configuration of the algorithms and applied the Wilcoxon non-parametric statistical test to draw conclusions on the difference between algorithms and configurations, adopting the conventional p -value threshold at $\alpha = 0.05$. Since multiple pairwise comparisons are performed with overlapping data, the chance to reject true null hypotheses may increase (type I error). To control this problem, we adopt the Holm-Bonferroni correction [26], which consists of using more strict significance levels when the number of pairwise tests increases.

To answer RQ3, we consider the best-performing configuration of the Deep RL algorithms, as selected from RQ2, and gradually increase the exploration complexity of the apps. Specifically, *20_strings*, *40_strings*, *80_strings* indicate an increasing size of the *string pool*. Such *string pool* is a dictionary of 20, 40, or 80 strings containing numbers and words, including the app’s username and password, to use with a login activity. The string pool does not contain duplicates. *augmented_5* and *augmented_10* indicate an increasing size of the self navigation links (with 5 or 10 “dummy” buttons that do nothing) within the login activities.

For the assessment, we adopt the widely used metric **AUC (Area Under the Curve)**, measuring the area below the activity coverage plot over time. To account for the non-determinism of the algorithms, we repeated each experiment 30 times and applied the Wilcoxon non-parametric statistical test. In RQ4, we investigate qualitatively the cases where Deep RL is superior to Tabular Q-Learning.

Study 2 (ARES):

- **RQ5:** *How do code coverage and time-dependent code coverage compare between Random, Q-Learning, DDPG, and SAC?*
- **RQ6:** *What are the fault exposure capabilities of the alternative approaches?*
- **RQ7:** *What features of the real apps make Deep RL perform better than Q-Learning?*
- **RQ8:** *How does ARES compare with state-of-the-art tools in terms of coverage and bug detection?*

In Study 2, we use real apps and compare the alternative Deep RL algorithms between each other and with Random and Tabular Q-Learning. At last, we compare ARES to state-of-the-art testing tools.

To address RQ5–RQ6 and RQ7, we randomly selected 100 apps among the 500 most starred F-Droid apps available on GitHub, and 41 successfully compiled. We consider coverage at the source code level and compare both the final coverage and the coverage increase over time (RQ5). To obtain coverage data at the instruction level, we instrumented each app using JaCoCo [37]. As in Study 1, we measured AUC with respect to the code coverage and compared AUC values using the Wilcoxon statistical test with a significance level set to 0.05 (with correction). We exclude TD3 from

Table 1. Ranking of Algorithms Produced by ARES vs. FATE; AUC Values Below Ranked Algorithms

App	Tool	Ranking/AUC				
Social	ARES	Q-Learn	DDPG	Rand	SAC	TD3
		4: 7788	5: 6,802	3: 9,547	2: 9,594	1: 15,101
	FATE	Q-Learn	DDPG	Rand	SAC	TD3
		4:7737	5:7,363	3:9,291	2:10,361	1:14,451
Bank	ARES	Q-Learn	DDPG	Rand	SAC	TD3
		4: 8614	5: 7,976	3: 9,344	1: 12,138	2: 10,932
	FATE	Q-Learn	DDPG	Rand	SAC	TD3
		4: 7750	5: 6,458	2: 9,746	1: 16,535	3: 9,305
Market	ARES	Q-Learn	DDPG	Rand	SAC	TD3
		1: 16866	2: 16,788	4: 15,936	5: 15,930	3: 15,944
	FATE	Q-Learn	DDPG	Rand	SAC	TD3
		1: 16496	2: 16,318	4: 15,943	5: 15,936	3: 15,949

the comparison, since it performed consistently worse than the other RL algorithms on synthetic apps.

In addition to code coverage, we also report the number of failures (unique app crashes) triggered by each approach (RQ6). To measure the number of unique crashes observed, we parsed the output of Logcat and (1) removed all crashes that do not contain the package name of the app; (2) extracted the stack trace; (3) computed the hash code of the sanitized stack trace, to uniquely identify it. With RQ7, we analyze the different performances of Deep RL vs. Q-Learning on real apps qualitatively.

To address RQ8, we evaluate and compare ARES and state-of-the-art tools in terms of code coverage and the number of crashes, using two different sets of apps under test, RQ8-a and RQ8-b, that accommodate the different requirements and constraints of the tools being compared. As state-of-the-art tools, we selected Monkey [21], Sapienz [36], TimeMachine [14], and Q-Testing [41]. In RQ8-a, we compare ARES, Monkey, Sapienz, and TimeMachine on a set of 68 apps coming from AndroTest [11]. These apps are instrumented using Emma [1], the same coverage tool that is used in Sapienz and TimeMachine. In RQ8-b, we compare ARES to Q-Testing on a set of 10 apps instrumented using JaCoCo, the coverage tool supported by Q-Testing.

All experimental data were generated and processed automatically. Each experiment was conducted with a one-hour timeout and was repeated 10 times for a total of 4,560 hours (≈ 190 days). The emulators involved in the study are equipped with 2 GB of RAM and Android 10.0 (API Level 29) or Android 4.4 only for the tool comparison.

6.1 Experimental Results: Study 1

Table 1 shows the ranking of the algorithms produced by ARES vs. FATE on the three apps translated from the synthetic FATE models to Java/Android. Below the ranking, Table 1 shows the AUC values obtained by the respective algorithms. The behaviors of the considered algorithms on synthetic (FATE) vs. translated (ARES) apps are very similar. The AUC values are quite close, and Spearman's correlation between AUC values across algorithms is 0.99 for Social, 0.89 for Bank, and 0.99 for Market; it is 0.95 overall. All correlations are statistically significant at level 0.05. ARES required 450 hours to complete the experiments. FATE required around 10 hours, reducing the computation time by a factor of 45 while producing similar results as ARES.

Fig. 4. Activity coverage of the Social synthetic app in FATE.

RQ1: The results obtained on synthetic apps are comparable to those obtained on their translated counterparts.

Figure 4 shows the coverage growth for the synthetic app Social. Each curve shows the mean of 60 runs. The shaded area around the mean represents the **Standard Error of the Mean** ($SEM = \sigma/\sqrt{n}$, where σ is the standard deviation and $n = 60$ the number of points). The highest activity coverage is obtained consistently by Deep RL algorithms, which have higher AUC values. Table 2 reports the AUC obtained on the synthetic apps in all tested configurations. Table 2 also shows the Vargha-Delaney effect size in the case of a statistically significant p -value $< \alpha/k$ where k is computed from the Holm-Bonferroni correction for multiple tests, between the winner algorithm (highest AUC) and the remainders.

Results show that Deep RL algorithms achieve higher coverage in most experiments. DDPG performs better in the simplest configuration, *20_strings*, while SAC performs better in almost all other configurations, including the most complex ones. Q-Learning prevails in only two scenarios belonging to *Market Place*, but the difference from the other algorithms is not statistically significant (p -value $> \alpha$).

RQ2: Results lead us to select DDPG and SAC as best performing Deep RL algorithms to be involved in Study 2 experiments.

DDPG is selected due to its high performance in relatively simple scenarios; SAC because of its ability to adapt and maintain good performance in the majority of scenarios.

RQ3: While in simple situations (e.g., the Player app), all algorithms achieve a high level of coverage, when things get more complex (e.g., when the string pool increases), Deep RL algorithms retain higher coverage than all other algorithms.

Table 2. Mean AUC for Synthetic Apps: Effect Size between the Winner (Shaded Cell) and Other Algorithms Is Reported Only When p -value is Statistically Significant (S = Small; M = Medium; L = Large)

App	Config	Rand	Q-Learn	TD3	SAC	DDPG	Effect Size
Player	20_str	88,719	89,022	89,903	89,943	90,337	–
Social	20_str	15,840	20,387	22,809	25,463	30,008	L(Rand), M(Q)
	40_str	9,291	7,737	14,451	10,361	7,363	S(DDPG)
	80_str	4,535	5,640	5,730	7,254	4,774	–
	aug_5	13,960	15,400	13,094	17,402	13,385	–
	aug_10	5,291	3,998	13,737	11,559	8,870	M(Q, Rand)
Bank	20_str	22,894	21,622	29,159	28,016	36,977	M(Q, Rand)
	40_str	9,746	7,750	9,305	16,535	6,458	S(Q, DDPG)
	80_str	3,998	4,843	4,776	5,621	4,798	–
	aug_5	12,815	8,634	8,702	14,914	11,472	–
	aug_10	4,121	6,289	13,289	14,195	15,361	M(Q, Rand)
Market	20_str	19,236	18,471	20,980	23,403	25,923	–
	40_str	15,943	16,496	15,949	15,936	16,318	–
	80_str	15,944	15,945	15,935	15,937	15,932	–
	aug_5	18,917	16,377	16,500	21,208	16,027	–
	aug_10	4,121	6,289	13,289	14,195	15,361	–

We have manually inspected the step-by-step exploration performed by Q-Learning and by the Deep RL algorithms. We found that login activities complicate substantially the exploration performed by Q-Learning. In fact, it is more difficult to reproduce the right username-password combination for a Tabular Q-Learning algorithm, which has limited adaptation capabilities. In contrast, Deep RL algorithms memorize the right combination in the DNN used to guide the exploration. In addition, large action spaces make it challenging for Q-Learning to learn an effective exploration strategy. The DNNs used by Deep RL algorithms can easily cope with large spaces of alternatives to choose from. The performance degradation of Q-Learning confirms this as the string pool increases in dimension or as new interactive elements (“dummy” buttons) are added, which confuse Q-Learning during its exploration.

RQ4: The performance of Q-Learning declines in the presence of blocking activities that require specific input combinations that must be learned from past interactions or when the input/action space becomes excessively large, while Deep RL can learn how to cope with such obstacles thanks to the DNN employed to learn the exploration strategy.

6.2 Experimental Results: Study 2

Table 3 shows coverage and crashes produced by each algorithm deployed in ARES. The highest average coverage and average number of crashes over 10 runs are shaded in gray for each app. We grouped the apps into three different size categories (Low-Medium, Medium, and High), depending on their **ELOC (Executable Lines Of Code)**. Results show that the Deep RL algorithms arise more often as winners when the ELOC increase. Usually, larger-size apps are more sophisticated and offer a richer set of user interactions, making their exploration more challenging for automated tools. We already know from Study 1 that when the action space or the observation space of the

Table 3. Average Coverage and Number of Crashes Observed on 41 Real Open-source Apps in 10 Runs of ARES

Applications	ELOC	Rand Q SAC DDPG				Rand Q SAC DDPG			
		%Coverage(mean)				#Crashes(mean)			
Silent-ping-sms	263	41	41	41	41	0	0	0	0
Drawablenotepad	452	20	21	26	25	0.7	0.6	0.7	0.1
SmsMatrix	466	23	20	24	22	1.2	0	1.2	0.9
Busybox	540	75	73	74	76	0	0	0	0
WiFiKeyShare	627	37	36	37	37	0	0	0	0
Talalarmo	1,094	69	71	71	71	0	0.5	0.5	0
AquaDroid	1,157	55	55	55	55	1.0	0.4	0.8	0.3
Lexica	1,215	72	72	74	75	0.3	0.1	1.5	1.2
Loyalty-card-locker	1,228	41	37	50	41	0.5	0.4	0.8	0.1
Dns66	1,264	58	58	58	58	0.1	0	0	0.2
Gptest	1,311	47	46	47	46	0	0	0	0
Memento	1,336	77	76	74	77	0	0	0	0
Editor	1,547	50	46	51	50	0	0	0	0
AndOTP	1,560	20	25	27	20	0.5	0.5	0.7	0.2
BookyMcBookface	1,595	26	25	25	24	0	0	0	0
Tuner	2,207	80	74	79	75	0	0	0	0
WifiAnalyzer	2,511	78	75	80	79	0	0	0	0
AdAway	3,064	38	37	45	40	0	0	0.1	0.1
Gpslogger	3,201	36	31	32	28	0	0	0	0.1
Connectbot	3,904	26	25	28	18	0	0	0	0
Neurolab	3,954	29	28	29	28	0	0.4	0.3	0.6
Anuto	4,325	46	46	47	47	0	0	0	0
PassAndroid	4,569	1	1	1	1	0	0	0	0
Markor	4,607	51	43	53	41	0.3	0	0.4	0
Vanilla	4,747	29	34	41	33	0	0	0	0
Average		45	43.84	46.76	44.32	0.15	0.12	0.28	0.15
Afwall	5,130	12	12	16	13	0	0	0	0
OpenTracks	5,260	45	42	44	45	0	0	0	0
Opentasks	5,772	43	50	53	44	0	0	0.2	0
UserLAnd	5,901	60	60	60	60	0.1	0.2	0.4	0.2
Simple-Solitaire	5,907	10	30	31	31	0	0.4	0.4	0.2
Authorizer	5,923	5	5	5	5	0	0	0	0
YalpStore	6,734	35	34	38	33	0	0	0	0
CameraRoll	6,836	32	31	31	32	0.8	0.1	1.6	0.1
AntennaPod	7,975	46	40	48	38	0.5	0.1	0.8	0.4
Phonograph	8,758	16	16	16	16	0	0	0	0
Average		30.4	30.5	34.2	31.7	0.14	0.08	0.34	0.09
MicroMathematics	10,506	35	35	47	41	0	0	0	0
LightningBrowser	11,961	35	36	43	37	0	0	0.4	0.1
Firefox-focus	12,482	33	34	41	35	0.5	0.3	0.8	0.1
RedReader	12,958	42	42	44	46	0	0	0.1	0
Wikipedia	23,543	42	43	44	41	0	0	0	0
Slide	30,483	19	17	18	19	0.8	0.3	1.2	0.3
Average		34.33	34.50	39.5	36.5	0.21	0.1	0.38	0.1
Total Average		39.62	39.58	42.61	40.09	0.17	0.1	0.3	0.12
Unique crashes		73	43	102	52

Fig. 5. Instruction coverage over time for the app *Loyalty-card-locker*.

apps increase, Deep RL can infer the complex actions needed to explore such apps more easily than other algorithms. Study 2 confirms the same trend.

Overall, SAC achieves the best performance, with 42.61% instruction coverage and 0.3 faults detected on average. DDPG comes next, with 40.09% instruction coverage and 0.12 faults detected on average. To further investigate these results, we computed the AUCs reached by each algorithm, and we applied the Wilcoxon test to each pair of algorithms. Table 4 shows the AUCs achieved by the four algorithms and the Vargha-Delaney effect size between the winner and the other algorithms when the p -value is less than α . SAC results as the winner on 56% of the considered real apps, followed by Random (34%). Moreover, Table 4 confirms the trend observed in Table 3: As ELOC increase, a higher proportion of Deep RL algorithms produces the highest AUC. Figure 5 shows an example of code coverage over time for the app *Loyalty-card-locker*, averaged on 10 runs. SAC increases its coverage almost until the end of the exploration, while the other algorithms reach a plateau after around 35 minutes of exploration.

RQ5: SAC reached the highest coverage in 24/41 apps, followed by DDPG (11 apps), Random (10 apps), and Q-Learning(1 app). SAC also has the highest AUC in 24/41 apps, followed by Random (13 apps), Q-Learning (11 apps), and DDPG (5 apps).

We suspect that the higher performance of SAC is related to its entropy regularization parameter. Thanks to the entropy regularization, in contrast to the other Deep RL algorithms that do not contain such parameter, SAC can maintain a high level of exploration even at the end of the testing phase, preventing the policy from converging to a bad local optimum.

Table 3 shows that SAC exposed the highest number of unique crashes (102), followed by Random (73), DDPG (52), and Q-Learning (43). The most common types of error exposed by SAC during testing are: *RuntimeException* (34 occurrences), *NullPointerException* (14), *IllegalArgument* (13). Figure 6 shows around a 30% overlap between the crashes found by SAC and

Table 4. AUCs Achieved on Real Apps; Effect Size between Winner and Others When p -value $< \alpha$

App	Rand	Q	SAC	DDPG	Effect Size
Silent-ping-sms	0.36	0.36	0.38	0.37	S(DDPG),M(Rand),L(Q)
Drawable-notepad	0.13	0.14	0.20	0.18	L(Rand,Q)
SmsMatrix	0.13	0.13	0.11	0.11	-
Busybox	0.54	0.56	0.68	0.68	L(Q,Rand)
WiFiKeyShare	0.29	0.29	0.33	0.30	L(DDPG,Q,Rand)
TalalarMO	0.62	0.60	0.64	0.64	L(Q)
AquaDroid	0.529	0.526	0.531	0.522	L(Rand, Q, DDPG)
Lexica	0.63	0.61	0.66	0.65	L(Q,Rand)
Loyalty-card-locker	0.23	0.23	0.34	0.21	L(DDPG,Q,Rand)
Dns66	0.51	0.51	0.47	0.45	L(DDPG,SAC)
Gptest	0.40	0.40	0.39	0.36	L(DDPG)
Memento	0.64	0.65	0.64	0.65	-
Editor	0.42	0.37	0.37	0.37	L(DDPG,Q,SAC)
AndOTP	0.16	0.18	0.23	0.15	M(Rand), L(DDPG)
BookyMcBookface	0.22	0.20	0.20	0.20	M(DDPG), L(SAC)
Tuner	0.68	0.66	0.60	0.60	L(DDPG,SAC)
WifiAnalyzer	0.56	0.56	0.67	0.58	L(DDPG,Q,Rand)
AdAway	0.25	0.25	0.27	0.25	-
Gpslogger	0.28	0.28	0.23	0.20	L(DDPG,SAC)
Connectbot	0.19	0.19	0.22	0.09	L(DDPG,Q,Rand)
Neurolab	0.23	0.23	0.22	0.22	-
Anuto	0.35	0.40	0.43	0.33	L(DDPG,Q,Rand)
PassAndroid	0.018	0.017	0.017	0.017	-
Markor	0.40	0.40	0.43	0.25	L(DDPG,Q,Rand)
Vanilla	0.17	0.23	0.26	0.23	L(Rand)
Afwall	0.09	0.09	0.13	0.10	L(DDPG,Q,Rand)
OpenTracks	0.37	0.35	0.35	0.35	-
Opentasks	0.18	0.46	0.46	0.29	L(DDPG,Rand)
UserLAnd	0.49	0.49	0.49	0.47	-
Simple-Solitaire	0.06	0.21	0.19	0.18	L(Rand)
Authorizer	0.05	0.046	0.047	0.049	S(SAC), M(Q)
YalpStore	0.28	0.28	0.31	0.26	L(DDPG,Q,Rand)
Camera-Roll	0.26	0.37	0.25	0.25	L(Rand,DDPG,SAC)
AntennaPod	0.33	0.33	0.35	0.22	L(DDPG)
Phonograph	0.085	0.077	0.075	0.076	S(Q,DDPG,SAC)
MicroMathematics	0.17	0.17	0.30	0.18	L(DDPG,Q,Rand)
Lightning-Browser	0.28	0.28	0.36	0.29	L(DDPG,Q,Rand)
Firefox-focus	0.27	0.28	0.43	0.35	L(Rand, Q)
RedReader	0.31	0.31	0.31	0.33	-
Wikipedia	0.30	0.30	0.33	0.28	-
Slide	0.13	0.13	0.12	0.14	-

Fig. 6. Comparison of total number of unique crashes on the 41 apps involved in RQ5-6: Dark gray areas indicate the proportion of crashes found by both techniques.

the other algorithms. The overlap with Random is the highest. SAC discovers about 40% of unique crashes found by Random; however, SAC found many new crashes that Random did not find.

RQ6: The SAC algorithm implemented in ARES generates the highest number of crashes, 102, in line with the results on coverage (RQ5), where SAC was also the best performing algorithm.

We have manually inspected the coverage progress of the different algorithms on some of the real apps considered in Study 2. We have identified two structural patterns: concatenated activities (i.e., a sequence of nested activities possibly requiring some precondition to move from one to the next) and blocking activities (activities that require a specific input combination to enable the transition to the next activity). We observed that Deep RL algorithms achieve higher coverage than the other algorithms when it is necessary to replicate complex behaviors to: (1) overcome blocking activities, e.g., to create an item to be able to access its properties later or to successfully authenticate within the app; (2) to pass through concatenated activities without being distracted by already seen activities or ineffective buttons (high dimensional action/observation space); (3) reach an activity located deeply in the app. Such behaviors are possible, thanks to the learning capabilities of the DNNs used by Deep RL algorithms, while they are hardly achieved by the other existing approaches, including Tabular Q-Learning.

RQ7: In the presence of blocking activities or complex concatenated activities (activities with a high number of widgets or located in depth) that require the capability to reuse knowledge acquired in previous explorations, the learning capabilities of Deep RL algorithms make them the most effective and efficient exploration strategies.

6.2.1 Comparison between ARES and State-of-the-art Tools. RQ8-a. Table 5 shows the coverage reached and the faults exposed by each testing tool on 68 Android apps from AndroTest. Coverage data are summarized employing boxplots in Figure 7. The highest average coverage and average number of crashes over 10 runs are highlighted with a gray background. ARES achieved 54.2% coverage and detected 0.48 crashes on average. TimeMachine achieved 50.4% code coverage and 0.42 faults on average. Sapienz reached a mean code coverage of 48.8% and discovered 0.22 faults

Table 5. Results on 68 Open-source Apps Coming from AndroTest

App	Coverage				Faults			
	Monkey	Sapienz	TM	ARES	Monkey	Sapienz	TM	ARES
a2dp	38	44	42	42	0	0.4	0.1	0.3
aagtl	16	18	17	19	0.3	1.0	1.7	1.5
aarddict	13	15	17	17	0	0	0.2	0.2
acal	18	27	27	28	0.5	0.8	1.0	1.0
addi	19	20	17	20	0.4	0.3	0.4	0.6
adsdroid	30	36	36	35	0.1	0.4	0.5	0.7
aGrep	45	-	59	56	0.1	-	0.3	0.2
aka	65	84	77	84	0.1	0.7	0.1	0.4
alarmclock	64	41	60	71	0.5	0.5	0.6	0.8
aLogCat	67	71	75	87	0	0	0	0
Amazed	36	66	63	89	0.1	0.3	0.2	0.2
AnyCut	63	65	63	73	0	0	0	0
anymemo	31	50	43	53	0.2	0.8	0.3	1.0
autoanswer	12	16	21	15	0	0	0.5	0.5
batterydog	63	67	62	69	0	0.1	0.4	0.5
battery	73	78	77	92	0	0.4	0.5	0.4
bites	34	41	45	43	0.1	0.2	0.9	0.5
blokish	55	52	68	45	0	0.2	0	0.3
bomber	76	73	77	84	0	0	0	0
Book-Catalogue	27	29	27	25	0.1	0.2	0.8	0.5
CountdownTimer	74	62	77	84	0	0	0	0
dalvik-explorer	66	72	70	72	0.3	0.2	0.7	0.8
dialer2	39	42	42	44	0	0	0.3	0.4
DivideAndConquer	84	83	82	80	0	0.2	1.0	0.9
fileexplorer	41	49	55	64	0	0	0	0
frozenbubble	80	76	75	70	0	0	0	0
gestures	37	52	51	55	0	0	0	0
hndroid	7	15	18	18	0.1	0.4	1.1	1.3
hotdeath	75	75	72	74	0.1	0.2	0.8	0.9
importcontacts	40	39	40	42	0.1	0	0.6	0.8
jamendo	53	41	54	63	0	0.4	1.4	1.6
k9mail	6	7	8	8	0.4	0	1.8	1.2
LNM	47	-	-	75	0	-	-	0.2
lockpatterngenerator	75	79	74	78	0	0	0	0
LolcatBuilder	26	25	29	26	0	0	0.1	0
manpages	40	73	70	74	0	0.4	0.3	0.4
mileage	38	45	48	45	0.3	1.0	2.3	1.8
Mirrored	57	59	62	59	0.4	0	0.8	0.7
mnv	41	60	43	56	0.5	0.3	1.0	1.1
multismssender	34	59	61	73	0	0.2	0.3	0.4
MunchLife	67	72	71	88	0	0	0	0
MyExpenses	41	60	50	63	0	0	0.2	0.2
myLock	25	31	50	30	0	0	0.5	0.2
Nectroid	34	66	58	57	0	1.0	0.3	0.9
netcounter	43	70	58	69	0	0	0.4	0.5
PasswordMaker	53	58	55	59	0.3	0.8	0.6	0.9
passwordmanager	7	8	17	18	0	0	0	0
Photostream	30	34	35	29	0.1	0	0.4	0.8
QuickSettings	50	45	46	52	0	0.2	0	0.4
RandomMediaPlayer	53	58	58	63	0	0	0.6	0.8
Ringdroid	22	29	48	30	0	0.1	0.3	0.2
sanity	26	19	31	22	0.2	0.3	0.5	0.4
soundboard	42	32	59	61	0	0.6	0	0.4
SpriteMethodTest	58	80	73	88	0	0	0	0
SpriteText	60	60	57	60	0	0.4	0	0.5
swift	12	14	13	17	0	0.4	-	0.6
SyncMyPix	21	21	23	25	0	0	0	0
tippy	75	83	74	85	0	0.4	0.3	0.4
tomdroid	47	46	51	69	0	0.3	0	0.3
Translate	49	48	48	50	0	0	0	0
Triangle	-	-	-	-	-	-	-	-
weight-chart	63	67	66	71	0	0	0	0
whohasmystuff	61	68	66	81	0.1	0	0.9	1.0
wikipedia	31	32	33	35	0	0	0	0
Wordpress	4	5	7	8	0	0.5	1.5	1.0
worldclock	83	88	86	90	0	0	0.6	0.4
yahtzee	51	57	56	69	2	0.2	0.5	0.5
zooborns	30	16	35	37	0	0	0.1	0.5
Average	43.9	48.8	50.4	54.2	0.11	0.22	0.42	0.48
Sum					51	103	179	171

Fig. 7. Code coverage achieved by ARES, TimeMachine, Sapienz, and Monkey.

Fig. 8. Comparison of total number of unique crashes involved in RQ8-a: Dark gray areas indicate the proportion of crashes found by both testing tools.

on average. Monkey achieved 43.9% code coverage and discovered 0.11 faults. ARES achieved the highest code coverage on average, followed by TimeMachine, Sapienz, and Monkey. TimeMachine detected most unique crashes (179), followed by ARES (171), Sapienz (103), and Monkey (51). Actually, ARES discovered less crashes than TimeMachine mostly because TimeMachine uses a system-level event generator, taken from Stoa [44], which ARES does not support. However, system events present two major drawbacks: (a) they vastly change depending on the Android version [20, 45] (despite TimeMachine is compatible with Android 7.1, it uses only system-level actions compatible with Android 4.4 [13]); and (b) to execute them, root privileges are required. ARES does not require root privileges to work properly on any app (i.e., we recall that certain apps do not execute on rooted devices [47]). Analyzing the execution traces of each tool, we searched and identified the faults immediately after the generation of system events not related to the AUT. More than a third (63) of the crashes generated by TimeMachine come from system-level actions. Figure 8 shows a pairwise comparison of detected crashes among evaluated techniques. TimeMachine also finds only 20% of unique crashes found by ARES. For example, in the app *mnv*, only ARES generated a crash of the type *NullPointerException*, in which a missing control on input generates the failure of the conversion function *CharSequence.toString()*. The text field from which the bug can be generated is not immediately available, but several interaction steps with the app are required.

This shows that ARES can be used together with other state-of-the-art Android testing techniques (in particular, TimeMachine) to cover more code and discover more crashes jointly.

The good results of ARES in code coverage and exposed faults are due to the reinforcement mechanisms of the RL algorithms and the reward function that drives the testing tool through states of the app leading to hard-to-reach states. Search-based techniques such as Sapienz typically observe the program behavior over an event sequence that may be very long. Hence, the associated coverage feedback, used to drive search-based exploration, does not have enough fine-grained details to support good local choices of the actions to be taken. The fitness function used by these algorithms evaluates an entire action sequence and does not consider the individual steps in the sequence. TimeMachine improved this weakness by relying on the coverage feedback obtained at an individual state to determine which portion of the app is not still explored. The drawback of this kind of approach is a higher computational cost that requires a higher testing time. In fact, while the time to dump the GUI is in common both to TimeMachine and ARES, measuring the coverage as feedback at each timestep implies three steps

- generation of the coverage files concerning the AUT,
- retrieval of coverage files from the Android device,
- coverage computation and processing.

These steps together take on average a second and a half to be completed. The strategy of ARES relies on monitoring the transition between activities taking on average 0.1 milliseconds rather than computing the code coverage at each timestep. Hence, the latter approach offers a better tradeoff between the granularity of the feedback and the computational cost required to obtain it.

RQ8-a: ARES achieved the highest code coverage in 41/68 apps, followed by TimeMachine (12/68), Sapienz (6/68), and Monkey (2/68). ARES triggered crashes more often than other tools in 23/68 apps, followed by TimeMachine (15/68), Sapienz (5/68), and Monkey (0/68). However, TimeMachine generated the highest number of unique crashes (179), but 63 of them come from system-level events. ARES generates 171 faults, Sapienz 103, and Monkey 51.

RQ8-b. Table 6 shows the coverage reached and the faults exposed by ARES and Q-Testing on 10 Android apps instrumented with JaCoCo. ARES achieved 64.3% coverage and exposed 0.41 faults on average; it detected 17 unique crashes. Q-Testing achieved 52.5% code coverage and 0.27 crashes on average, and it detected 16 unique crashes. ARES achieved the highest code coverage on almost all apps, and on average ARES covered 12% more code than Q-Testing. Q-Testing generated six faults in common with ARES, while the other four faults are generated using the system-level events of Stoot. In the app *antennapod*, only ARES generated a *NumberFormatException* with a text field located deeply in the Settings submenu. In this comparison, the main advantage of ARES seems to be a better reward function that encourages the tool to visit the greatest number of activities within the same episode. Instead, Q-Testing determines the reward of the Q-Learning algorithm by computing the similarity between Android app states, which does not guarantee an efficient way to overcome blocking activities.

RQ8-b: ARES reached the highest code coverage on 9/10 apps, the highest average of crashes in 7/10 apps, and the highest number of unique crashes. Q-Testing, the state-of-the-art RL Android testing tool, reached the same level of exploration of ARES only in 1/10 app, and the highest average number of crashes on only 1/10 apps.

Fig. 9. Code coverage achieved by ARES and Q-Testing.

Table 6. Comparison between Q-Testing and ARES

App	Coverage		Faults	
	ARES	Q-Testing	ARES	Q-Testing
Alogcat	84	76	0	0
Antennapod	48	42	0.8	0.4
AnyCut	72	67	0	0
batterydog	65	49	0.5	0.3
Jamendo	64	46	0.4	0.7
Multismssender	71	45	0.4	0
Myexpanses	63	36	0.4	0.2
talalarma	74	74	0.8	0.5
Tomdroid	61	50	0.3	0.2
vanilla	41	40	0.5	0.4
Average	64.3	52.5	0.41	0.27
Sum			17	16

6.3 Threats To Validity

We adopted several strategies to enhance the internal validity of our results. We chose apps coming from a standard testing benchmark used in previous studies to mitigate risks of selection bias. We used default settings, given the same starting condition, and ran each tool several times under the same workload to ensure that no testing tool was at a disadvantage. We followed the Stoa protocol to identify unique crashes and, also, we manually checked the crashes found. To measure coverage, we used JaCoCo, a standard coverage tool.

7 IMPLICATIONS AND LIMITATIONS

Using ARES as testing tool involves several benefits:

- **black-box automated testing:** ARES relies only on the GUI of the AUT. This allows developers to test their production apps with no modifications. Available state-of-the-art tools, such as Sapienz and TimeMachine, rely on code coverage to drive exploration, and as recognized by the researchers who developed them, this makes testing less efficient.

- **wide compatibility:** ARES works on Windows, Linux, and MacOS. ARES can test apps in parallel on emulators or real devices with Android from 6.0 to 10.0.
- **policy reuse:** At the end of the testing phase, ARES saves the status of the neural network as a policy file. Instead of restarting ARES from scratch each time a new version of the AUT is launched, the policy can be loaded and reused on a later version of the AUT.
- **modularity:** Within ARES, the app environment is decoupled from the RL algorithm used during the testing phase, allowing to deploy new algorithms easily.

Despite the advantages given by ARES there are also some limitations:

- **benefits on easy apps:** Performance on simpler apps sometimes align with the performance of dummy methods such as random algorithms and does not justify the use of Deep RL.
- **system-level events:** ARES implements a limited set of events that act at the system-level, including the most commonly supported ones, such as *toggle Internet connection* and *rotate screen*. The other system-level events require rooted devices to work, which brings some drawbacks, among which the inability to work with apps that, for security reasons, perform a “root-check” and stop working if the device is rooted.

8 RELATED WORK

As suggested by Reference [27], existing research on Android testing can be classified by the methodologies that the testing approaches adopt.

8.1 Random Testing

Testing tools in this category generate random events on the AUT GUI. Monkey [21] is one of the most popular black-box Android testing tools. It triggers events by interacting randomly with screen coordinates. This simple random approach worked relatively well on some benchmark applications [11]. Nonetheless, Monkey tests involve many ineffective or repeated events, as there is no guidance to make the exploration efficient. ARES implements a smarter version of Monkey, which we used as a baseline (Random). Such an improved version selects only actions that are possible in a given GUI state, thus making the exploration strategy a bit more efficient.

8.2 Model-based Testing

Model-based tools [5, 6, 23] first build navigation models of the Android app by means of static or dynamic analysis, used to explore efficiently the application states, and then they extract test cases from such models to eventually expose bugs. AndroidRipper [5], MobiGUITAR [6] try to maximize the exploration by using the ripping technique. Guo et al. [23] use static analysis to effectively improve GUI exploration performance. Stoa [44] uses a stochastic FSM to model the app behavior. The app model is built using dynamic analysis, enhanced with a weighted UI exploration strategy, and with the help of static analysis. Compared to Stoa, ARES can be viewed as computing a navigation model implicitly: The MDP model used by the Deep RL algorithms. One key advantage of using an implicit model is that we do not have to deal with the combinatorial explosion of its size, which Stoa controls using model compaction heuristics.

8.3 Structural Testing

Structural strategies [7, 14, 16, 35, 36] generate coverage-oriented inputs using symbolic execution or evolutionary algorithms. Sapienz [36] maximizes code coverage and bug revelation using a Pareto-optimal multi-objective search-based approach, which applies genetic operators such as mutation and crossover to produce new test cases. It can generate specific input for text fields by reverse-engineering the APK. This process occasionally results in invalid sequences discarded

by the fitness functions that reward test cases with high coverage. TimeMachine [14] improves Sapienz by identifying interesting states in the past and restarting the search process from them when the search stagnates. While coverage-oriented approaches require the possibility to instrument the app to measure coverage and possibly execute it symbolically, our approach is black-box and does not suffer from the scalability issues affecting, e.g., symbolic execution.

8.4 Machine Learning-based Testing

Some ML-based testing approaches [8, 28, 31] use an explicit, supervised training process to learn from previous test executions. They can reuse previous knowledge acquired on different apps or past versions of the app under test. Approaches as QBE [28] make the transfer of knowledge to new apps possible by abstracting the app state in a form that is supposed to hold across different domains and implementations. However, the effectiveness of such a transfer learning process depends on the similarity between new and old apps.

One of the first works proposing RL for GUI testing is AutoBlackTest [38]. This approach is based on the simplest form of RL, Tabular Q-Learning, whose effectiveness is strongly dependent on the initial values in the Q-Table. On the contrary, ARES learns the action-value function from scratch during the exploration of the AUT. One of the most recent approaches to Android testing based on Deep Learning is Q-Testing [41]. However, it also uses Tabular Q-Learning as a backbone. At the same time, learning is limited to the computation of the similarity between Android app states, which determines the reward of the Q-Learning algorithm. ARES instead learns both the state similarity and the action-value function during its interactions with the AUT.

9 CONCLUSION AND FUTURE WORK

We have proposed an approach based on Deep RL for the automated exploration of Android apps. The best exploration strategy is learned automatically as the test progresses. The approach is implemented in the open-source tool ARES, which is complemented by FATE, a model-based Android testing tool that we developed to support fast execution and configuration of the alternative Deep RL algorithms of ARES. The resulting configuration of ARES, particularly when running the SAC Deep RL algorithm, outperformed all the considered baselines in terms of coverage achieved over time and exposed bugs.

In our future work, we plan to investigate specific fault categories that are particularly relevant for Android apps, such as security vulnerabilities. In fact, we think that the adaptation and reward mechanisms used by Deep RL algorithms to learn the optimal exploration strategy could be particularly effective when the fault to be exposed is a security fault. We also plan to port ARES and FATE to iOS.

A APPENDIX

A.1 Study 1- Deep RL

Certain parameters of the algorithms are omitted for simplicity and can be found in the documentation of Stable Baselines [25].

Control Policy:

Deep RL algorithms rely on a policy based on an MLP composed of 2 layers and 64 neurons each.

Learning rates:

- DDPG: 0.0001
- SAC: 0.0003
- TD3: 0.0003

DDPG									
Configuration		1	2	3	4	5	6	7	8
random_exploration		0.5	0.5	0.6	0.6	0.7	0.7	0.8	0.8
nb_train_steps		5	25	5	25	5	25	5	25

TD3									
Configuration		1	2	3	4	5	6	7	8
random_exploration		0.5	0.5	0.6	0.6	0.7	0.7	0.8	0.8
train_frequency		25	100	25	100	25	100	25	100

SAC									
Configuration		1	2	3	4	5	6	7	8
target_update_interval		1	1	2	2	5	5	10	10
train_frequency		1	5	1	5	1	5	1	5

A.2 Study 1- Tabular RL

Q-Learning Configuration	1	2	3	4	5	6	7	8
epsilon	0.5	0.6	0.7	0.8	0.5	0.6	0.7	0.8
gamma	0.99	0.99	0.99	0.99	0.9	0.9	0.9	0.9
alpha	0.628	0.628	0.628	0.628	0.628	0.628	0.628	0.628

A.3 Study 2

Some parameters used in Study 2 have been selected from the results of Study 1 (see Section 6).

DDPG:

- Control Policy: MLP, 2 layers, 64 neurons
- Learning rate: 0.0001
- nb_train_steps: 10
- random_exploration: 0.7

SAC:

- Control Policy: MLP, 2 layers, 64 neurons
- Learning rate: 0.0003
- train_freq: 5
- target_update_interval: 10

TD3:

- Control Policy: MLP, 2 layers, 64 neurons
- Learning rate: 0.0003
- train_freq: 10
- random_exploration: 0.8

Q-Learning:

- epsilon: 0.8
- gamma: 0.9
- alpha: 0.628

REFERENCES

- [1] 2006. Emma. Retrieved from <http://emma.sourceforge.net>.
- [2] 2020. Appbrain. Retrieved from <https://www.appbrain.com>.
- [3] 2020. Appium. Retrieved from <http://appium.io>.
- [4] Josh Achiam. 2018. Key Concepts in RL. Retrieved from https://spinningup.openai.com/en/latest/spinningup/rl_intro.html.
- [5] Domenico Amalfitano, Anna Rita Fasolino, Porfirio Tramontana, Salvatore De Carmine, and Atif M. Memon. 2012. Using GUI ripping for automated testing of Android applications. In *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering*. IEEE, 258–261.
- [6] Domenico Amalfitano, Anna Rita Fasolino, Porfirio Tramontana, Bryan Dzung Ta, and Atif M. Memon. 2014. Mobi-GUITAR: Automated model-based testing of mobile apps. *IEEE Softw.* 32, 5 (2014), 53–59.
- [7] Saswat Anand, Mayur Naik, Mary Jean Harrold, and Hongseok Yang. 2012. Automated concolic testing of smartphone apps. In *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*. 1–11.
- [8] Nataniel P. Borges, Maria Gómez, and Andreas Zeller. 2018. Guiding app testing with mined interaction models. In *Proceedings of the IEEE/ACM 5th International Conference on Mobile Software Engineering and Systems (MOBILESoft)*. IEEE, 133–143.
- [9] Justin A. Boyan and Andrew W. Moore. 1995. Generalization in reinforcement learning: Safely approximating the value function. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*. 369–376.
- [10] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai Gym. *arXiv preprint arXiv:1606.01540* (2016).
- [11] Shaunik Roy Choudhary, Alessandra Gorla, and Alessandro Orso. 2015. Automated test input generation for Android: Are we there yet?(e). In *Proceedings of the 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 429–440.
- [12] Jack Clark and Dario Amodei. 2016. Faulty Reward Functions in the Wild. Retrieved from <https://openai.com/blog/faulty-reward-functions/>.
- [13] Zhen Dong, Marcel Bohme, Lucia Cojocaru, and Abhik Roychoudhury. 2020. Github Repository: TimeMachine. Retrieved from https://github.com/DroidTest/TimeMachine/blob/master/fuzzingandroid/sys_event_generator/sys_event.py.
- [14] Zhen Dong, Marcel Bohme, Lucia Cojocaru, and Abhik Roychoudhury. 2020. Time-travel Testing of Android apps. In *Proceedings of the 42nd International Conference on Software Engineering (ICSE)*. 481–492.
- [15] Scott Fujimoto, Herke van Hoof, and David Meger. 2018. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477* (2018).
- [16] Xiang Gao, Shin Hwei Tan, Zhen Dong, and Abhik Roychoudhury. 2018. Android testing via synthetic symbolic execution. In *Proceedings of the 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 419–429.
- [17] Google. 2019. Broadcasts. Retrieved from <https://developer.android.com/guide/components/broadcasts>.
- [18] Google. 2019. Safety Net. Retrieved from <https://developer.android.com/training/safetynet/attestation>.
- [19] Google. 2020. Android Emulator. Retrieved from <https://developer.android.com/studio/run/emulator/>.
- [20] Google. 2020. System-Level Events API 25. Retrieved from https://cs.android.com/android/platform/superproject/+/android-7.1.2_r36:frameworks/base/core/res/AndroidManifest.xml.
- [21] Google. 2020. UI/Application Exerciser Monkey. Retrieved from <https://developer.android.com/studio/test/monkey>.
- [22] Tianxiao Gu, Chengnian Sun, Xiaoxing Ma, Chun Cao, Chang Xu, Yuan Yao, Qirun Zhang, Jian Lu, and Zhendong Su. 2019. Practical GUI testing of Android applications via model abstraction and refinement. In *Proceedings of the IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 269–280.
- [23] Wunan Guo, Liwei Shen, Ting Su, Xin Peng, and Weiyang Xie. 2020. Improving automated GUI exploration of Android apps via static dependency analysis. In *Proceedings of the IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 557–568. DOI: <https://doi.org/10.1109/ICSME46990.2020.00059>
- [24] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290* (2018).
- [25] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. 2018. Stable Baselines. Retrieved from <https://github.com/hill-a/stable-baselines>.
- [26] Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scand. J. Statist.* 6, 2 (1979), 65–70.
- [27] Pingfan Kong, Li Li, Jun Gao, Kui Liu, Tegawendé F. Bissyandé, and Jacques Klein. 2018. Automated testing of Android apps: A systematic literature review. *IEEE Trans. Reliab.* 68, 1 (2018), 45–66.

- [28] Yavuz Koroglu, Alper Sen, Ozlem Muslu, Yunus Mete, Ceyda Ulker, Tolga Tanriverdi, and Yunus Donmez. 2018. QBE: QLearning-based exploration of Android applications. In *Proceedings of the IEEE 11th International Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 105–115.
- [29] Edward A. Lee. 2009. *Finite State Machines and Modal Models in Ptolemy II*. Technical Report UCB/EECS-2009-151. EECS Department, University of California, Berkeley. Retrieved from <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-151.html>.
- [30] Yuxi Li. 2017. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274* (2017).
- [31] Yuanchun Li, Ziyue Yang, Yao Guo, and Xiangqun Chen. 2019. Humanoid: A deep learning-based approach to automated black-box Android app testing. In *Proceedings of the 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 1070–1073.
- [32] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015).
- [33] Long-Ji Lin. 1993. *Reinforcement Learning for Robots Using Neural Networks*. Technical Report. Carnegie-Mellon University School of Computer Science, Pittsburgh PA.
- [34] Aravind Machiry, Rohan Tahlilani, and Mayur Naik. 2013. Dynodroid: An input generation system for Android apps. In *Proceedings of the 9th Joint Meeting on Foundations of Software Engineering*. 224–234.
- [35] Riyadh Mahmood, Nariman Mirzaei, and Sam Malek. 2014. EvoDroid: Segmented evolutionary testing of Android apps. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 599–609.
- [36] Ke Mao, Mark Harman, and Yue Jia. 2016. Sapienz: Multi-objective automated testing for Android applications. In *Proceedings of the 25th International Symposium on Software Testing and Analysis*. 94–105.
- [37] Evgeny Mandrikov, Marc R. Hoffmann, Brock Janiczak. 2020. Jacoco Code Coverage. Retrieved from <https://www.eclemma.org/jacoco/>.
- [38] Leonardo Mariani, Mauro Pezze, Oliviero Riganelli, and Mauro Santoro. 2012. AutoBlackTest: Automatic black-box testing of interactive applications. In *Proceedings of the IEEE 5th International Conference on Software Testing, Verification and Validation*. IEEE, 81–90.
- [39] Microsoft. 2020. Your Device is Rooted and You Can't Connect-Android. Retrieved from <https://docs.microsoft.com/it-it/mem/intune/user-help/your-device-is-rooted-and-you-cant-connect-android>.
- [40] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [41] Minxue Pan, An Huang, Guoxin Wang, Tian Zhang, and Xuandong Li. 2020. Reinforcement learning based curiosity-driven testing of Android applications. In *Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 153–164.
- [42] Martin Riedmiller. 2005. Neural fitted Q iteration—First experiences with a data efficient neural reinforcement learning method. In *Proceedings of the European Conference on Machine Learning*. Springer, 317–328.
- [43] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. 2014. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning*, Vol. 32, Eric P. Xing, and Tony Jebara (Eds.). PMLR, 387–395. <https://proceedings.mlr.press/v32/silver14.html>.
- [44] Ting Su, Guozhu Meng, Yuting Chen, Ke Wu, Weiming Yang, Yao Yao, Geguang Pu, Yang Liu, and Zhendong Su. 2017. Guided, stochastic model-based GUI testing of Android apps. In *Proceedings of the 11th Joint Meeting on Foundations of Software Engineering*. 245–256.
- [45] Ting Su, Guozhu Meng, Yuting Chen, Ke Wu, Weiming Yang, Yao Yao, Geguang Pu, Yang Liu, and Zhendong Su. 2017. System-Level Events API 19. Retrieved from <https://sites.google.com/site/stoat2017/evaluation/stoat-s-system-level-events>.
- [46] Richard S. Sutton and Andrew G. Barto. 2018. Reinforcement learning: An introduction. MIT press.
- [47] Vincent F. Taylor and Ivan Martinovic. 2017. Short paper: A longitudinal study of financial apps in the Google Play store. In *Proceedings of the International Conference on Financial Cryptography and Data Security*. Springer, 302–309.
- [48] Christopher J. C. H. Watkins and Peter Dayan. 1992. Q-learning. *Mach. Learn.* 8, 3–4 (1992), 279–292.

Received January 2021; revised August 2021; accepted November 2022