

A survey of benchmarking frameworks for reinforcement learning

Belinda Stapelberg , Katherine M. Malan 

Department of Decision Sciences, University of South Africa, Pretoria, South Africa

ABSTRACT

Reinforcement learning has recently experienced increased prominence in the machine learning community. There are many approaches to solving reinforcement learning problems with new techniques developed constantly. When solving problems using reinforcement learning, there are various difficult challenges to overcome.

To ensure progress in the field, benchmarks are important for testing new algorithms and comparing with other approaches. The reproducibility of results for fair comparison is therefore vital in ensuring that improvements are accurately judged. This paper provides an overview of different contributions to reinforcement learning benchmarking and discusses how they can assist researchers to address the challenges facing reinforcement learning. The contributions discussed are the most used and recent in the literature. The paper discusses the contributions in terms of implementation, tasks and provided algorithm implementations with benchmarks.

The survey aims to bring attention to the wide range of reinforcement learning benchmarking tasks available and to encourage research to take place in a standardised manner. Additionally, this survey acts as an overview for researchers not familiar with the different tasks that can be used to develop and test new reinforcement learning algorithms.

Keywords: reinforcement learning, benchmarking

Categories: • Computing methodologies ~ Reinforcement learning

Email:

Belinda Stapelberg belinda.stapelberg@up.ac.za (CORRESPONDING),
Katherine M. Malan malankm@unisa.ac.za

Article history:

Received: 16 September 2019
Accepted: 25 November 2020
Available online: 08 December 2020

1 INTRODUCTION

Reinforcement learning (RL) is a subfield of machine learning, based on rewarding desired behaviours and/or punishing undesired ones of an agent interacting with its environment (Sutton & Barto, 2018). The agent learns by taking sequential actions in its environment, observing the state of the environment and receiving a reward. The agent needs to learn a strategy, called a policy, to decide which action to take in any state. The goal of RL is to find the policy that maximises the long-term reward of the agent.

In recent years RL has experienced dramatic growth in research attention and interest due to promising results in areas including robotics control (Lillicrap et al., 2015), playing Atari

Belinda Stapelberg and Katherine M. Malan (2020). A survey of benchmarking frameworks for reinforcement learning. *South African Computer Journal* 32(2), 258–292. <https://doi.org/10.18489/sacj.v32i2.746>

Copyright © the author(s); published under a [Creative Commons NonCommercial 4.0 License \(CC BY-NC 4.0\)](https://creativecommons.org/licenses/by-nc/4.0/). SACJ is a publication of the South African Institute of Computer Scientists and Information Technologists. ISSN 1015-7999 (print) ISSN 2313-7835 (online).

2600 (Mnih et al., 2013; Mnih et al., 2015), competitive video games (Silva & Chaimowicz, 2017; Vinyals et al., 2017), traffic light control (Arel et al., 2010) and more. In 2016, RL came into the general spotlight when Google DeepMind's AlphaGo (Silver et al., 2016) program defeated the Go world champion, Lee Sedol. Even more recently, Google DeepMind's AlphaStar AI program defeated professional StarCraft II players (considered to be one of the most challenging real-time strategy games) and OpenAI Five defeated professional Dota 2 players.

Progress in machine learning is driven by new algorithm development and the availability of high-quality data. In supervised and unsupervised machine learning fields, resources such as the UCI Machine Learning repository¹, the Penn Treebank (Marcus et al., 1993), the MNIST database of handwritten digits², the ImageNet large scale visual recognition challenge (Rusakovsky et al., 2015), and Pascal Visual Object Classes (Everingham et al., 2010) are available. In contrast to the datasets used in supervised and unsupervised machine learning, progress in RL is instead driven by research on agent behaviour within challenging environments. Games have been used for decades to test and evaluate the performance of artificial intelligence systems. Many of the benchmarks that are available for RL are also based on games, such as the Arcade Learning Environment for Atari 2600 games (Bellemare et al., 2013) but others involve tasks that simulate real-world situations, such as locomotion tasks in Garage (originally rllab) (Duan et al., 2016). These benchmarking tasks have been used extensively in research and significant progress has been made in using RL in ever more challenging domains.

Benchmarks and standardised environments are crucial in facilitating progress in RL. One advantage of the use of these benchmarking tasks is the reproducibility and comparison of algorithms to state-of-the-art RL methods. Progress in the field can only be sustained if existing work can be reproduced and accurately compared to judge improvements of new methods (Henderson et al., 2018; Machado et al., 2018). The existence of standardised tasks can facilitate accurate benchmarking of RL performance.

This paper provides a survey of the most important and most recent contributions to benchmarking for RL. These are OpenAI Gym (Brockman et al., 2016), the Arcade Learning Environment (Bellemare et al., 2013), a continuous control benchmark rllab (Duan et al., 2016), RoboCup Keepaway soccer (Stone & Sutton, 2001) and Microsoft TextWorld (Côté et al., 2018). When solving RL problems, there are many challenges that need to be overcome, such as the fundamental trade-off problem between exploration and exploitation, partial observability of the environment, delayed rewards, enormous state spaces and so on. This paper discusses these challenges in terms of important RL benchmarking contributions and in what manner the benchmarks can be used to overcome or address these challenges.

The rest of the paper is organised as follows. Section 2 introduces the key concepts and terminology of RL, and then discusses the approaches to solving RL problems and the challenges for RL. Section 3 provides a survey on the contributions to RL benchmarking and Section 4 discusses the ways that the different contributions to RL benchmarking deal with or contribute to the challenges for RL. A conclusion follows in Section 5.

¹<http://archive.ics.uci.edu/ml/index.php>

²<http://yann.lecun.com/exdb/mnist/>

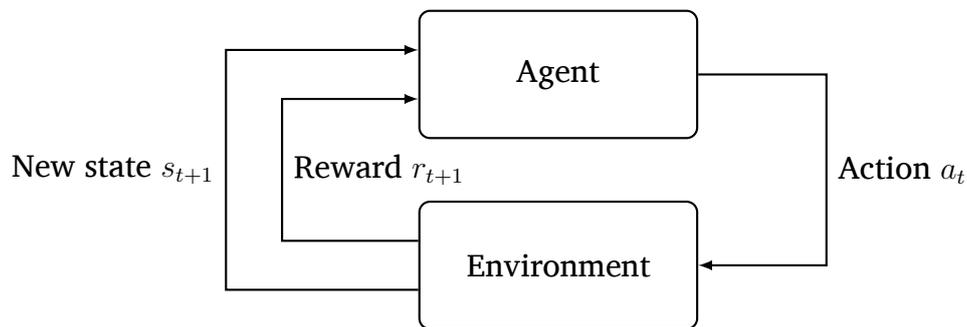


Figure 1: Illustration of an RL system.

2 REINFORCEMENT LEARNING

RL focuses on training an agent by using a trial-and-error approach. Figure 1 illustrates the workings of an RL system. The agent evaluates a current situation (state), takes an action, and receives feedback (reward) from the environment after each act. The agent is rewarded with either positive feedback (when taking a “good” action) or negative feedback as punishment for taking a “bad” action. An RL agent learns how to act best through many attempts and failures. Through this type of trial-and-error learning, the agent’s goal is to receive the best so-called long-term reward. The agent gets short-term rewards that together lead to the cumulative, long-term reward. The key goal of RL is to define the best sequence of actions that allow the agent to solve a problem while maximizing its cumulative long-term reward. That set of optimal actions is learned through the interaction of the agent with its environment and observation of rewards in every state.

This section provides the key concepts and terminology of RL used throughout this paper. The challenges of RL are also discussed.

2.1 Concepts and terminology

The core idea behind RL is to learn from the environment through interactions and feedback, and find an optimal strategy for solving the problem. The agent takes actions in its environment based on a (possibly partial) observation of the state of the environment and the environment provides a reward for the actions, which is usually a scalar value. The set of all valid actions is referred to as the action space, which can be either discrete (as in Atari and Go) or continuous (controlling a robot in the physical world). The goal of the agent is to maximise its long-term cumulative reward.

2.1.1 Policy

A policy of an agent is the control strategy used to make decisions, and is a mapping from states to actions. A policy can be deterministic or stochastic and is denoted by π . A deterministic

policy maps states to actions without uncertainty while a stochastic policy is a probability distribution over actions for a given state. Therefore, when an agent follows a deterministic policy it will always take the same action for a given state, whereas a stochastic policy may take different actions in the same state. The immediate advantage of a stochastic policy is that an agent is not doomed to repeat a looped sequence of non-advancing actions.

2.1.2 On-policy and off-policy learning

There are two types of policy learning methods. On-policy learning is when the agent “learns on the job”, i.e. it evaluates or improves the policy that is used to make the decisions directly. Off-policy learning is when the agent learns one policy, called the target policy, while following another policy, called the behaviour policy, which generates behaviour. The off-policy learning method is comparable to humans learning a task by observing others performing the task.

2.1.3 Value functions

Having a value for a state (or state-action pair) is often useful in guiding the agent towards the optimal policy. The value under policy π is the expected return if the agent starts in a specific state or state-action pair, and then follows the policy thereafter. So the state-value function v_π is a mapping from states to real numbers and represents the long-term reward obtained by starting from a particular state and executing policy π . The action-value function q_π is a mapping from state-action pairs to real numbers. The action-value $q_\pi(s, a)$ of state s and action a (where a is an arbitrary action and not necessarily in line with the policy) is the expected return from starting in state s , taking action a and then following policy π . The optimal value function v_* gives the expected return starting in a state and then following the optimal policy π_* . The optimal action-value function q_* is the expected return starting in some state, taking an arbitrary action and then following the optimal policy π_* .

These state-value and action-value functions all obey so-called Bellman equations, where the idea is that the value of the agent’s starting point is the reward that is expected to be obtained from being there, plus the value of wherever the agent lands next. These Bellman equations are used in most RL approaches where the Bellman-backup is used, i.e. for a state or state-action pair the Bellman-backup is the (immediate) reward plus the next value.

2.1.4 Function approximators

In many RL problems the state space can be extremely large. Traditional solution methods where value functions are represented as arrays or tables mapping all states to values are therefore very difficult (Sutton & Barto, 2018). One approach to this shortcoming is to use features to generalise an estimation of the value of states with similar features. Methods that compute these approximations are called function approximators. There are many techniques

used for implementing function approximators including linear combinations of features, neural networks, decision trees, nearest neighbours, etc.

2.1.5 Monte Carlo methods

Monte Carlo methods are a class of learning methods where value functions are learned (Sutton & Barto, 2018). The value of a state, s_i , is estimated by running many trials starting from s_i and then averaging the total rewards received on those trials.

2.1.6 Temporal difference algorithms

Temporal difference (TD) learning algorithms are a class of learning methods that are based on the idea of comparing temporally successive predictions. These methods are a fundamental idea in RL and use a combination of Monte Carlo learning and dynamic programming (Sutton & Barto, 2018). TD methods learn value functions directly from experience by using the so-called TD error and bootstrapping (not waiting for a final outcome).

2.1.7 Markov decisions processes

The standard formalism for RL settings is called a Markov decision process (MDP). MDPs are used to define the interaction between an agent and its environment in terms of states, actions, and rewards. For an RL problem to be an MDP, it has to satisfy the Markov property: “The future is independent of the past given the present”. This means that once the current state is known, then the history encountered so far can be discarded and that state completely characterises all the information needed as it captures all the relevant information from the history. Mathematically, an MDP is a tuple: $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma \rangle$, where \mathcal{S} is a (finite) set of states, \mathcal{A} is a (finite) set of actions, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, \mathcal{P} is a state transition probability matrix and $\gamma \in [0, 1]$ is a discount factor included to control the reward.

2.1.8 Model-free and model-based reinforcement learning approaches

There are different aspects of RL systems that can be learnt. These include learning policies (either deterministic or stochastic), learning action-value functions (so-called Q-functions or Q-learning), learning state-value functions, and/or learning a model of the environment. A model of the environment is a function that predicts state transitions and rewards, and is an optional element of an RL system. If a model is available, i.e. if all the elements of the MDP are known, particularly the transition probabilities and the reward function, then a solution can be computed using classic techniques before executing any action in the environment. This is known as planning: computing the solution to a decision-making problem before executing an actual decision.

When an agent does not know all the elements of the MDP, then the agent does not know how the environment will change in response to its actions or what its immediate reward will be. In this situation the agent will have to try out different actions, observe what happens and

in some way find a good policy from doing this. One approach to solve a problem without a complete model is for the agent to learn a model of how the environment works from its observations and then plan a solution using that model. Methods that use the framework of models and planning are referred to as model-based methods.

Another way of solving RL problems without a complete model of the environment is to learn through trial-and-error. Methods that do not have or learn a model of the environment and do not use planning are called model-free methods. The two main approaches to represent and train agents with model-free RL are policy optimisation and Q-learning. In policy optimisation methods (or policy-iteration methods) the agent learns the policy function directly. Examples include policy gradient methods, asynchronous advantage actor-critic (A3C) (Mnih et al., 2016), trust region policy optimization (TRPO) (Schulman et al., 2015) and proximal policy optimization (PPO) (Schulman et al., 2017). Q-Learning methods include deep Q-networks (DQN) (Mnih et al., 2013), C51 algorithm (Bellemare et al., 2017) and Hindsight Experience Replay (HER) (Andrychowicz et al., 2017). Hybrid methods combining the strengths of Q-learning and policy gradients exist as well, such as deep deterministic policy gradients (DDPG) (Lillicrap et al., 2015), soft actor-critic algorithm (SAC) (Haarnoja et al., 2018) and twin delayed deep deterministic policy gradients (TD3) (Fujimoto et al., 2018).

In the current literature, the most used approaches incorporate a mixture of model-based and model-free methods, such as Dyna and Monte Carlo tree search (MCTS) (Sutton & Barto, 2018), and temporal difference search (Silver et al., 2012).

2.2 Challenges for reinforcement learning

This section discusses some of the challenges faced by RL. These challenges will be discussed in terms of how they are addressed by different contributions in Section 4.

2.2.1 Partially observable environment

How the agent observes the environment can have a significant impact on the difficulty of the problem. In most real-world environments the agent does not have a complete or perfect perception of the state of its environment due to incomplete information provided by its sensors, the sensors being noisy or some of the state being hidden. However, for learning methods that are based on MDPs, the complete state of the environment should be known. To address the problem of partial observability of the environment, the MDP framework is extended to the partially observable Markov decision process (POMDP) model.

2.2.2 Delayed or sparse rewards

In an RL problem, an agent's actions determine its immediate reward as well as the next state of the environment. Therefore, an agent has to take both these factors into account when deciding which action to take in any state. Since the goal is to learn which actions to take that will give the most reward in the long-run, it can become challenging when there is little or

no immediate reward. The agent will consequently have to learn from delayed reinforcement, where it may take many actions with insignificant rewards to reach a future state with full reward feedback. The agent must therefore be able to learn which actions will result in an optimal reward, which it might only receive far into the future.

In line with the challenge of delayed or sparse rewards is the problem of long-term credit assignment (Minsky, 1961): how must credit for success be distributed among the sequence of decisions that have been made to produce the outcome?

2.2.3 Unspecified or multi-objective reward functions

Many tasks (especially real-world problems) have multiple objectives. The goal of RL is to optimise a reward function, which is commonly framed as a global reward function, but tasks with more than one objective could require optimisation of different reward functions. In addition, when an agent is training to optimise some objective, other objectives could be discovered which might have to be maintained or improved upon. Work on multi-objective RL (MORL) has received increased interest, but research is still primarily devoted to single-objective RL.

2.2.4 Size of the state and action spaces

Large state and action spaces can result in enormous policy spaces in RL problems. Both state and action spaces can be continuous and therefore infinite. However, even discrete states and actions can lead to infeasible enumeration of policy/state-value space. In RL problems for which state and/or action spaces are small enough, so-called tabular solutions methods can be used, where value functions can be represented as arrays or tables and exact solutions are often possible. For RL problems with state and/or action spaces that are too large, the goal is to instead find good approximate solutions with the limited computational resources available and to avoid the curse of dimensionality (Bellman, 1957).

2.2.5 The trade-off between exploration and exploitation

One of the most important and fundamental overarching challenges in RL is the trade-off between exploration and exploitation. Since the goal is to obtain as much reward as possible, an agent has to learn to take actions that were previously most effective in producing a reward. However, to discover these desirable actions, the agent has to try actions that were not tried before. It has to exploit the knowledge of actions that were already taken, but also explore new actions that could potentially be better selections in the future. The agent may have to sacrifice short-term gains to achieve the best long-term reward. Therefore, both exploration and exploitation are fundamental in the learning process, and exclusive use of either will result in failure of the task at hand. There are many exploration strategies (Sutton & Barto, 2018), but a key issue is the scalability to more complex or larger problems. The exploration vs.

exploitation challenge is affected by many of the other challenges that are discussed in this section, such as delayed or sparse rewards, and the size of the state or action spaces.

2.2.6 Representation learning

Representation (or feature) learning involves automatically extracting features or understanding the representation of raw input data to perform tasks such as classification or prediction. It is fundamental not just to RL, but to machine learning and AI in general, even with a conference dedicated to it: International Conference on Learning Representations (ICLR).

One of the clearest challenges that representation learning tries to solve in an RL context is to effectively reduce the impact of the curse of dimensionality, which results from very large state and/or action spaces. Ideally an effective representation learning scheme will be able to extract the most important information from the problem input in a compressed form.

2.2.7 Transfer learning

Transfer learning (Pan & Yang, 2010; Weiss et al., 2016) uses the notion that, as in human learning, knowledge gained from a previous task can improve the learning in a new (related) task through the transfer of knowledge that has already been learned. The field of transfer learning has recently been experiencing growth in RL (Taylor & Stone, 2009) to accelerate learning and mitigate issues regarding scalability.

2.2.8 Model learning

Model-based RL methods (Section 2.1.8) are important in problems where the agent's interactions with the environment are expensive. These methods are also significant in the trade-off between exploration and exploitation, since planning impacts the need for exploration. Model learning can reduce the interactions with the environment, something which can be limited in practice, but introduces additional complexities and the possibility of model errors. Another challenge related to model learning is the problem of planning using an imperfect model, which is also a difficult challenge that has not received much attention in the literature.

2.2.9 Off-policy learning

Off-policy learning methods (e.g. Q-learning) scale well in comparison to other methods and the algorithms can (in principle) learn from data without interacting with the environment. An agent is trained using data collected by other agents (off-policy data) and data it collects itself to learn generalisable skills.

Disadvantages of off-policy learning methods include greater variance and slow convergence, but are more powerful and general than on-policy learning methods (Sutton & Barto, 2018). Advantages of using off-policy learning is the use of a variety of exploration strategies, and learning from training data that are generated by unrelated controllers, which includes manual human control and previously collected data.

2.2.10 Reinforcement learning in real-world settings

The use of RL in real-world scenarios has been gaining attention due to the success of RL in artificial domains. In real-world settings, more challenges become apparent for RL. Dulac-Arnold et al. (2019) provide a list of nine challenges for RL in the real-world, many of which have been mentioned in this section already. Further challenges not discussed here include *safety constraints*, *policy explainability* and *real-time inference*. Many of these challenges have been studied extensively in isolation, but there is a need for research on algorithms (both in artificial domains and real-world settings) that addresses more than one or all of these challenges together, since many of the challenges are present in the same problem.

2.2.11 A standard methodology for benchmarking

A diverse range of methodologies is currently common in the literature, which brings into question the validity of direct comparisons between different approaches. A standard methodology for benchmarking is necessary for the research community to compare results in a valid way and accelerate advancement in a rigorous scientific manner.

3 CONTRIBUTIONS TO REINFORCEMENT LEARNING BENCHMARKING

This section discusses some important reinforcement learning benchmarks currently in use. The list of contributions is by no means exhaustive, but includes the ones that are most in use currently in the RL research community.

3.1 OpenAI Gym

Released publicly in April 2016, OpenAI Gym (Brockman et al., 2016) is a toolkit for developing and comparing reinforcement learning algorithms. It includes a collection of benchmark problems which is continuing to grow as well as a website where researchers can share their results and compare algorithm performance. It provides a tool to standardise reporting of environments in research publications to facilitate the reproducibility of published research. OpenAI Gym has become very popular since its release, with Brockman et al. (2016) having over 1300 citations on Google Scholar to date.

3.1.1 Implementation

The OpenAI Gym library is a collection of test problems (environments) with a common interface and makes no assumptions about the structure of an agent. OpenAI Gym currently supports Linux and OS X running Python 2.7 or 3.5 – 3.7. Windows support is currently experimental, with limited support for some problem environments. OpenAI Gym is compatible

with any numerical computation library, such as TensorFlow or Theano. To get started with OpenAI Gym, visit the documentation site³ or the actively maintained GitHub repository⁴.

3.1.2 Benchmark tasks

The environments available in the library are diverse, ranging from easy to difficult and include a variety of data. A brief overview of the different environments is provided here with the full list and descriptions of environments available on the main site³.

Classic control and toy text: These small-scale problems are a good starting point for researchers not familiar with the field. The classic control problems include balancing a pole on a moving cart (Figure 2a), driving a car up a steep hill, swinging a pendulum and more. The toy text problems include finding a safe path across a grid of ice and water tiles, playing Roulette, Blackjack and more.

Algorithmic: The objective here is for the agent to learn algorithms such as adding multi-digit numbers and reversing sequences, purely from examples. The difficulty of the tasks can be varied by changing the sequence length.

Atari 2600: The Arcade Learning Environment (ALE) (Bellemare et al., 2013) has been integrated into OpenAI Gym in easy-to-install form, where classic Atari 2600 games (see Figure 2b for an example) can be used for developing agents (see Section 3.2 for a detailed discussion). For each game there are two versions: a version which takes the RAM as input and a version which takes the observable screen as the input.

MuJoCo: These robot simulation tasks use the MuJoCo proprietary software physics engine (Todorov et al., 2012), but free trial and postgraduate student licences are available. The problems include 3D robot walking or standing up tasks, 2D robots running, hopping, swimming or walking (see Figure 2c for an example), balancing two poles vertically on top of each other on a moving cart, and repositioning the end of a two-link robotic arm to a given spot.

Box2D: These are continuous control tasks in the Box2D simulator, which is a free open source 2-dimensional physics simulator engine. Problems include training a bipedal robot (Figure 2d) to walk (even on rough terrain), racing a car around a track and navigating a lunar lander to its landing pad.

Roboschool: Most of these problems are the same as in MuJoCo, but use the open-source software physics engine, Bullet. Additional tasks include teaching a 3D humanoid robot to walk as fast as possible (see Figure 2e) as well as a continuous control version of Atari Pong.

Robotics: Released in 2018, these environments are used to train models which work on physical robots. It includes four environments using the Fetch⁵ research platform and four environments using the ShadowHand⁶ robot. These manipulation tasks are significantly more difficult than the MuJoCo continuous control environments. The tasks for the Fetch robot are

³<https://gym.openai.com>

⁴<https://github.com/openai/gym>

⁵<https://fetchrobotics.com/>

⁶<https://www.shadowrobot.com/products/dexterous-hand/>

to move the end-effector to a desired goal position, hitting a puck across a long table such that it slides and comes to rest on the desired goal, moving a box by pushing it until it reaches a desired goal position, and picking up a box from a table using its gripper and moving it to a desired goal above the table. The tasks for the ShadowHand are reaching with its thumb and a selected finger until they meet at a desired goal position above the palm, manipulating a block (see Figure 2f), an egg, and a pen, until the object achieves a desired goal position and rotation.

Alongside these new robotics environments, OpenAI also released code for Hindsight Experience Replay (HER), a reinforcement learning algorithm that can learn from failure. Their results show that HER can learn successful policies on most of the new robotics problems from only sparse rewards. A set of requests for research has also been released⁷ in order to encourage and facilitate research in this area, with a few ideas of ways to improve HER specifically.

3.2 The Arcade Learning Environment

The Atari 2600 gaming console was released in September 1977, with over 565 games developed for it over many different genres. The games are considerably simpler than modern era video games. However, the Atari 2600 games are still challenging and provide interesting tasks for human players.

The Arcade Learning Environment (ALE) (Bellemare et al., 2013) is an object-oriented software framework allowing researchers to develop AI agents for the original Atari 2600 games. It is a platform to empirically assess and evaluate AI agents designed for general competency. ALE allows interfacing through the Atari 2600 emulator Stella and enables the separation of designing an AI agent and the details of emulation. There are currently over 50 game environments supported in the ALE.

The ALE has received a lot of attention since its release in 2013 (over 1200 citations on Google Scholar to date), perhaps the most note-worthy being the success of Deep Q-networks (DQN), which was the first algorithm to achieve human-level control performance in many of the Atari 2600 games (Mnih et al., 2015).

3.2.1 Implementation

The Stella emulator interfaces with the Atari 2600 games by receiving joystick movements and sending screen and/or RAM information to the user. For the reinforcement learning context, ALE has a game-handling layer to provide the accumulated score and a signal for whether the game has ended. The default observation of a single game screen or frame is made up of a two-dimensional array of 7-bit pixels, 160 pixels wide by 210 pixels high. The joystick controller defines 18 discrete actions, which makes up the action space of the problem. Only some actions are needed to play a game and the game-handling layer also provides the minimum set of actions needed to play any particular game. The simulator generates 60 frames per

⁷<https://openai.com/blog/ingredients-for-robotics-research/>

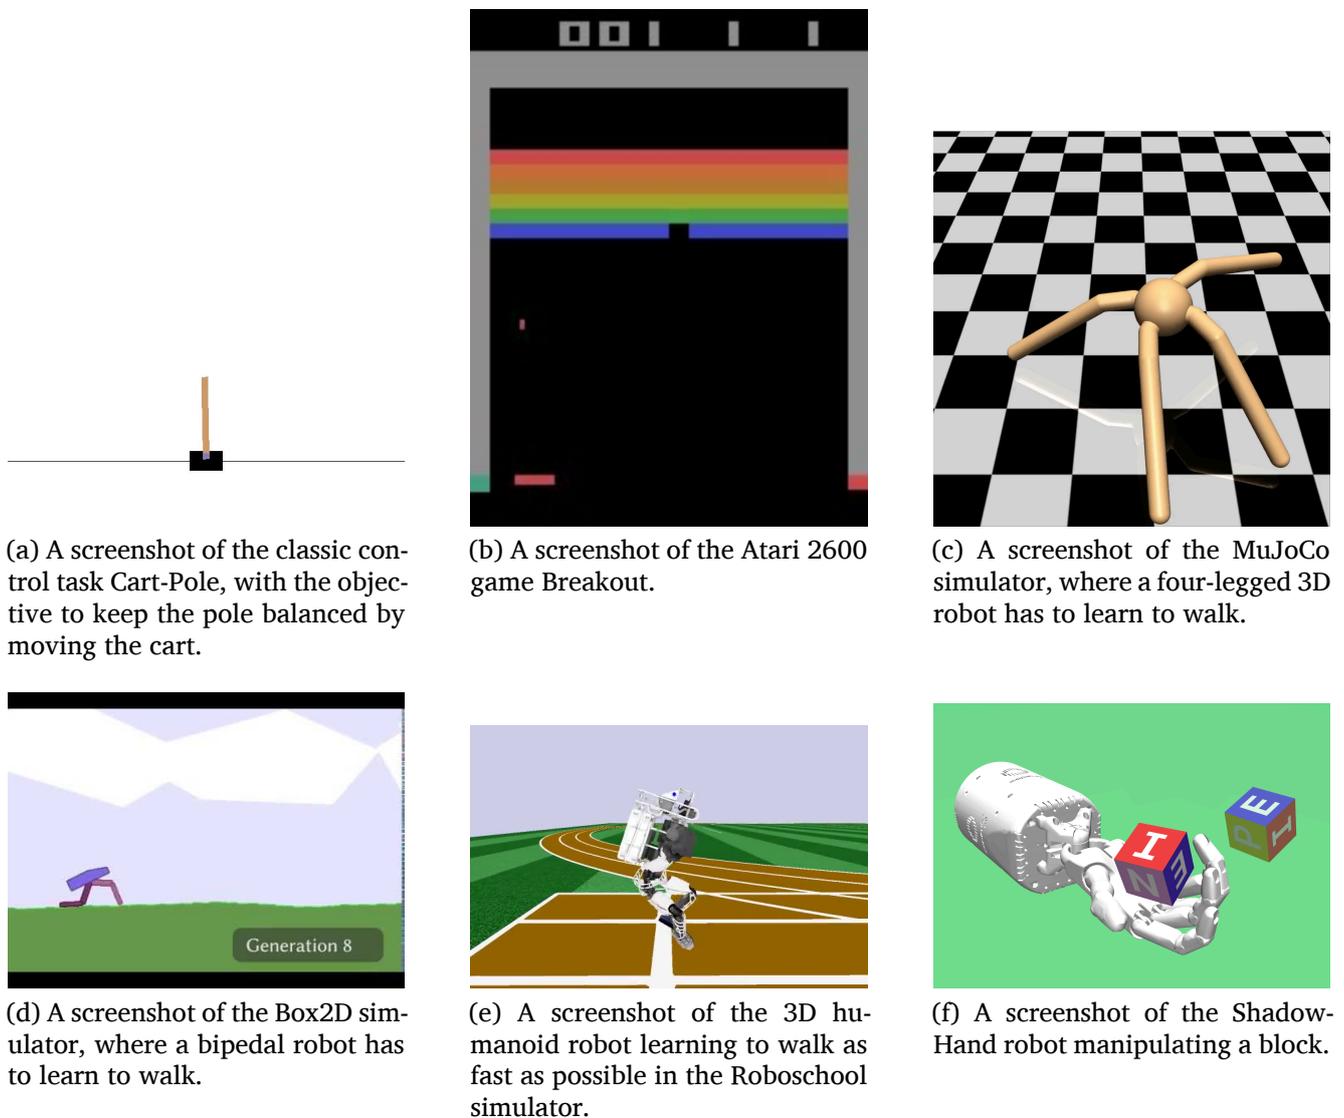


Figure 2: Some examples of the environments used in OpenAI Gym.

second in real-time and up to 6000 frames per second at full speed. The reward the agent receives depends on each game, but is generally the score difference between frames. A game episode starts when the first frame is shown and ends when the goal of the game has been achieved or after a predefined number of frames. The ALE therefore offers access to a variety of games through one common interface.

The ALE also has the functionality of saving and restoring the current state of the emulator. This functionality allows the investigation of topics including planning and model-based reinforcement learning.

ALE is free, open-source software⁸, including the source code for the agents used in associated research studies (Bellemare et al., 2013). ALE is written in C++, but there are many interfaces available that allow the interaction with ALE in other programming languages, with detail provided in (Bellemare et al., 2013).

Due to the increase in popularity and importance in the AI literature, another paper was published in 2018 by some of the original proposers of the ALE (Machado et al., 2018), providing a broad overview of how the ALE is used by researchers, highlighting overlooked issues and discussing propositions for maximising the future use of the testbed. Concerns are raised at how agents are evaluated in the ALE and new benchmark results are provided.

In addition, a new version of the ALE was introduced in 2018 (Machado et al., 2018), which supports multiple game modes and includes so called *sticky actions*, providing some form of stochasticity to the controller. When sticky actions are used, there is a possibility that the action requested by the agent is not executed, but instead the agent's previous action is used, emulating a sticky controller. The probability that an action will be sticky can be specified using a pre-set control parameter. The original ALE is fully deterministic and consequently it is possible for an agent to memorise a good action sequence, instead of learning how to make good decisions. Introducing sticky actions therefore increases the robustness of the policy that the agent has to learn.

Originally the ALE only allowed agents to play games in their default mode and difficulty. In the latest version of the ALE (Machado et al., 2018) it is possible to select among different game modes and difficulty levels for single player games, where each mode-difficulty pair is referred to as a *flavour*. Changes in the mode and difficulty of the games can impact game dynamics and introduce new actions.

3.2.2 Published benchmark results

Bellemare et al. (2013) provide performance results on the ALE tasks using an augmented version of the SARSA(λ) algorithm (Sutton & Barto, 2018), where linear function approximation is used. For comparison, the performance results of a non-expert human player and three baseline agents (Random, Const and Perturb) are also provided. A set of games is used for training and parameter tuning, and another set for testing. The ALE can also be used to study planning techniques. Benchmark results for two traditional search methods (Breadth-first search and UCT: Upper Confidence Bounds Applied to Trees) are provided, as well as the performance results of the best learning agent and the best baseline policy.

Machado et al. (2018) provide benchmark results for 60 Atari 2600 games with sticky actions for DQN and SARSA(λ) + Blob-PROST (Liang et al., 2016) (an algorithm that includes a feature representation which enables SARSA(λ) to achieve performance that is comparable to that of DQN).

⁸<http://arcadelalearningenvironment.org>

3.3 Continuous control: rllab

The Arcade Learning Environment (Section 3.2) is a popular benchmark to evaluate algorithms which are designed for tasks with discrete actions. Duan et al. (2016) present a benchmark of 31 continuous control tasks, ranging in difficulty, and also implement a range of RL algorithms on the tasks.

The benchmark as well as the implementations of the algorithms are available at the rllab GitHub repository⁹, however this repository is no longer under development but is currently actively maintained at the **garage** GitHub repository¹⁰, which includes many improvements. The documentation¹¹ for garage is a work in progress and the available documentation is currently limited. Both rllab and garage are fully compatible with OpenAI Gym and only support Python 3.5 and higher.

Other RL benchmarks for continuous control have also been proposed, but many are not in use anymore. Duan et al. (2016) provide a comprehensive list of benchmarks containing low-dimensional tasks as well as a wide range of tasks with high-dimensional continuous state and action spaces. They also discuss previously proposed benchmarks for high-dimensional control tasks do not include such a variety of tasks as in rllab. Where relevant, we mention some of these benchmarks in the next section that have additional interesting tasks.

3.3.1 Benchmark tasks

There are four categories for the rllab continuous control tasks: basic, locomotion, partially observable and hierarchical tasks.

Basic tasks: These five tasks are widely analysed in the reinforcement learning and control literature. Some of these tasks can also be found in the “Classic control” section of OpenAI Gym (Section 3.1). The tasks are cart-pole balancing, cart-pole swing up, mountain car, acrobot swing up and double inverted pendulum balancing (which can be found in OpenAI Gym Roboschool). A related benchmark involving a 20 link pole balancing task is proposed as part of the Tdlearn package (Dann et al., 2014).

Locomotion tasks: Six locomotion tasks of varying dynamics and difficulty are implemented with the goal to move forward as quickly as possible. These tasks are challenging due to high degrees of freedom as well as the need for a lot of exploration, since getting stuck at a local optima (such as staying at the origin or diving forward slowly) can happen easily when the agent acts greedily. These tasks are: Swimmer, Hopper, Walker, Half-Cheetah, Ant, Simple Humanoid and Full Humanoid.

Other environments with related locomotion tasks include dotRL (Papis & Wawrzyński, 2013) with a variable segment octopus arm (Woolley & Stanley, 2010), PyBrain (Schaul et al., 2010), and SkyAI (Yamaguchi & Ogasawara, 2010) with humanoid robot tasks like jumping, crawling and turning.

⁹<https://github.com/rll/rllab>

¹⁰<https://github.com/rlworkgroup/garage>

¹¹<https://garage.readthedocs.io/en/latest/>

Partially observable tasks: Realistic agents often do not have access to perfect state information due to limitations in sensory input. To address this, three variations of partially observable tasks are implemented for each of the five basic tasks mentioned above. This leads to 15 additional tasks. The three variations are *limited sensors* (only positional information is provided, no velocity), *noisy observations and delayed actions* (Gaussian noise is added to simulate sensor noise, and a time delay is added between taking an action and an action being executed) and *system identification* (the underlying physical model parameters vary across different episodes). These variations are not currently available in OpenAI Gym.

Hierarchical tasks: In many real-world situations higher level decisions can reuse lower level skills, for example a robot learning to navigate a maze can reuse learned locomotion skills. Here tasks are proposed where low-level motor controls and high-level decisions are needed, which operate on different time scales and a natural hierarchy exists in order to learn the task most efficiently. The tasks are as follows. *Locomotion and food collection:* where the swimmer or the ant robot operates in a finite region and the goal is to collect food and avoid bombs. *Locomotion and maze:* the swimmer or the ant robot has the objective to reach a specific goal location in a fixed maze environment. These tasks are not currently available in OpenAI Gym.

3.3.2 Published benchmark results

Duan *et al.* (Duan et al., 2016) provide performance results on the rllab tasks. The algorithms implemented are mainly gradient-based policy search methods, but two gradient-free methods are included for comparison. Almost all of the algorithms are batch algorithms and one algorithm is an online algorithm. The batch algorithms are REINFORCE (Williams, 1992), truncated natural policy gradient (TNPG) (Duan et al., 2016), reward-weighted regression (RWR) (Peters & Schaal, 2007), relative entropy policy search (REPS) (Peters et al., 2010), trust region policy optimization (TRPO) (Schulman et al., 2015), cross entropy method (CEM) (Rubinstein, 1999) and covariance matrix adaptation evolution strategy (CMA-ES) (Hansen & Ostermeier, 2001). The online algorithm used is deep deterministic policy gradient (DDPG) (Lillicrap et al., 2015). Direct applications of the batch-based algorithms to recurrent policies are implemented with minor modifications.

Of the implemented algorithms, TNPG, TRPO and DDPG were effective in training deep neural network policies. However, all algorithms performed poorly on the hierarchical tasks, which suggest that new algorithms should be developed for automatic discovery and exploitation of the tasks' hierarchical structure.

Recently a new class of reinforcement learning algorithms called proximal policy optimisation (PPO) (Schulman et al., 2017) was released by OpenAI. PPO's performance is comparable or better than state-of-the-art approaches to solving 3D locomotion, robotic tasks (similar to the tasks in the benchmark discussed above) and also Atari 2600, but it is simpler to implement and tune. OpenAI has adopted PPO as its go-to RL algorithm, since it strikes a balance between ease of implementation, sample complexity, and ease of tuning.

3.4 RoboCup Keepaway Soccer

RoboCup (Kitano et al., 1997) simulated soccer has been used as the basis for successful international competitions and research challenges since 1997. Keepaway is a subtask of RoboCup that was put forth as a testbed for machine learning in 2001 (Stone & Sutton, 2001). It has since been used for research on temporal difference reinforcement learning with function approximation (Stone, Sutton, & Kuhlmann, 2005), evolutionary learning (Pietro et al., 2002), relational reinforcement learning (Walker et al., 2004), behaviour transfer (Cheng et al., 2018; Didi & Nitschke, 2016a, 2016b, 2018; Nitschke & Didi, 2017; Schwab et al., 2018; Taylor & Stone, 2005), batch reinforcement learning (Riedmiller et al., 2009) and hierarchical reinforcement learning (Bai & Russell, 2017).

In Keepaway, one team (the keepers) tries to maintain possession of the ball within a limited region, while the opposing team (the takers) attempts to gain possession (Stone & Sutton, 2001). The episode ends whenever the takers take possession of the ball or the ball leaves the region. The players are then reset for another episode with the keepers being given possession of the ball again. Task parameters include the size of the region, the number of keepers, and the number of takers. Figure 3 shows an example episode with 3 keepers and 2 takers (called 3v2) playing in a $20m \times 20m$ region (Stone & Sutton, 2001).

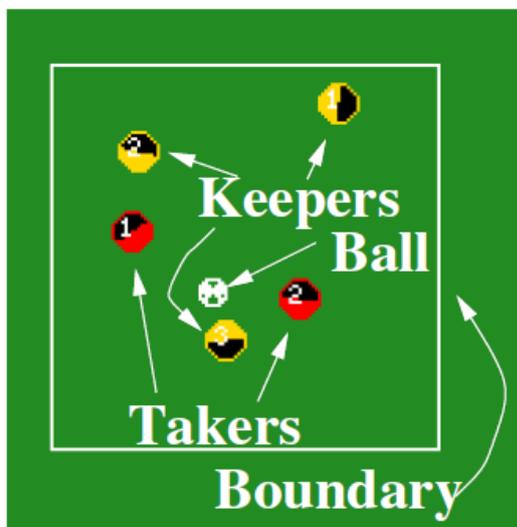


Figure 3: A screen shot from a 3v2 keepaway episode in a $20m \times 20m$ region from Stone and Sutton (2001).

In 2005 Stone, Kuhlmann, et al. (2005) elevated the Keepaway testbed to a benchmark problem for machine learning and provided infrastructure to easily implement the standardised task.

An advantage of the Keepaway subtask is that it allows for direct comparison of different machine learning algorithms. It is also good for benchmarking machine learning since the task is simple enough to be solved successfully, but complex enough that straightforward solutions are not sufficient.

3.4.1 Implementation

A standardized Keepaway player framework is implemented in C++ and the source code is available for public use at an online repository¹². The repository provides implementation for all aspects of the Keepaway problem except the learning algorithm itself. It also contains a step-by-step tutorial of how to use

¹²<http://www.cs.utexas.edu/~AustinVilla/sim/keepaway/>

the code, with the goal of allowing researchers who are not experts in the RoboCup simulated soccer domain to easily become familiar with the domain.

3.4.2 Standardised task

Robocup simulated soccer (and therefore also Keepaway) is a fully distributed, multiagent domain with both teammates and adversaries (Stone, 2000). The environment is partially observable for each agent and the agents also have noisy sensors and actuators. Therefore, the agents do not perceive the world exactly as it is, nor can they affect the world exactly as intended. The perception and action cycles of the agent are asynchronous, therefore perceptual input does not trigger actions as is traditional in AI. Communication opportunities are limited, and the agents must make their decisions in real-time. These domain characteristics all result in simulated robotic soccer being a realistic and challenging domain (Stone, 2000).

The size of the Keepaway region, the number of keepers, and the number of takers can easily be varied to change the task. Stone, Kuhlmann, et al. (2005) provide a framework with a standard interface to the learner in terms of macro-actions, states, and rewards.

3.4.3 Published benchmark results

Stone, Kuhlmann, et al. (2005) performed an empirical study for learning Keepaway by training the keepers using episodic SMDP SARSA(λ) (Stone, Sutton, & Kuhlmann, 2005; Sutton & Barto, 2018), with three different function approximators: CMAC function approximation (Albus, 1975, 1981), Radial Basis Function (RBF) (Sutton & Barto, 2018) networks (a novel extension to CMACs (Stone, Kuhlmann, et al., 2005)), and neural network function approximation. The RBF network performed comparably to the CMAC method. The Keepaway benchmark structure allows for these results to be quantitatively compared to other learning algorithms to test the relative benefits of different techniques.

3.4.4 Half Field Offense: An extension to Keepaway

Half Field Offense (HFO) (Hausknecht et al., 2016; Kalyanakrishnan et al., 2007) is an extension of Keepaway, which is played on half of the soccer field with more players on each team. The task was originally introduced in Kalyanakrishnan et al. (2007), but no code was made publicly available. In Hausknecht et al. (2016) the HFO environment was released publicly (open-source)¹³, however this repository is not currently being maintained.

Success in HFO means that the offensive players have to keep possession of the ball (the same as in Keepaway), learn to pass or dribble to get closer to the goal and shoot when possible. Agents can also play defence where they have to prevent goals from being scored. HFO also supports multi-agents which could be controlled manually or automatically.

In the same way as the Keepaway environment (Stone, Kuhlmann, et al., 2005), the HFO environment allows ease of use in developing and deploying agents in different game scenarios,

¹³<https://github.com/LARG/HFO>

with C++ and Python interfaces. The performance of three benchmark agents are compared in (Hausknecht et al., 2016), namely a random agent, a handcoded agent and a SARSA agent.

A similar platform to the Arcade Learning Environment (Section 3.2), the HFO environment places less emphasis on generality (the main goal of the ALE) and more emphasis on cooperation and multiagent learning.

3.5 Microsoft TextWorld

Recently, researchers from the Microsoft Research Montreal Lab released an open source project called TextWorld (Côté et al., 2018), which attempts to train reinforcement learning agents using text-based games.

In a time where AI agents are mastering complex multi-player games such as Dota 2 and StarCraft II, it might seem unusual to do research on text-based games. Text-based games can play a similar role to multi-player graphic environments which train agents to learn spatial and time-based planning, in advancing conversational skills such as affordance extraction (identifying which verbs are applicable to a given object), memory and planning, exploration etc. Another powerful motivation for the interest in text-based games is that language abstracts away complex physical processes, such as a robot trying not to fall over due to gravity. Text-based games require language understanding and successful play requires skills like long-term memory and planning, exploration (trial and error), common sense, and learning with these challenges.

TextWorld is a sandbox environment which enables users to handcraft or automatically generate new games. These games are complex and interactive simulations where text is used to describe the game state and players enter text commands to progress through the game. Natural language is used to describe the state of the world, to accept actions from the player, and to report subsequent changes in the environment. The games are played through a command line terminal and are turn-based, i.e. the simulator describes the state of the game through text and then a player enters a text command to change its state in some desirable way.

3.5.1 Implementation

In Figure 4 an example game is shown in order to illustrate the command structure of a typical text-based game generated by TextWorld.

TextWorld enables interactive play-through of text-based games and, unlike other text-based environments such as TextPlayer¹⁴ and PyFiction¹⁵, enables users to handcraft games or to construct games automatically. The TextWorld logic engine automatically builds game worlds, populates them with objects and obstacles, and generates quests that define a goal state and how to reach it (Côté et al., 2018). TextWorld requires Python 3 and currently only supports Linux and macOS systems. The code and documentation are available publicly¹⁶ and

¹⁴<https://github.com/danielricks/textplayer>

¹⁵<https://github.com/MikulasZelinka/pyfiction>

¹⁶<http://aka.ms/textworld>

the learning environment is described in full detail in Section 3 of (Côté et al., 2018), including descriptions of the two main components of the Python framework: a game generator and a game engine. To interact with TextWorld, the framework provides a simple application programming interface (API) which is inspired by OpenAI Gym.

```

You are now playing a life changing round of TextWorld! Here is how to play!
Your first objective is to venture west. And then, take the cuboid key from the
canteen. And then, take a trip east. After that, doublecheck that the cuboid box
within the cookery is closed. Then, assure the cuboid box is locked. And if you
do that, you're the winner!

-= Cookery -=
You are in a cookery. A typical one.

You make out an open cuboid box. Empty! What kind of nightmare TextWorld is
this?

There is a closed door leading east. You need an unguarded exit? You should try
going north. You need an unblocked exit? You should try going west.

> go west

-= Canteen -=
You arrive in a canteen. An usual kind of place.

You see a box. You make out a Henderson's limited edition locker.

There is an unguarded exit to the east.

There is a cuboid key and a Henderson's limited edition key on the floor.

> take cuboid key
You pick up the cuboid key from the ground.

Your score has just gone up by one point.

> go east

-= Cookery -=
You are in a cookery. A typical one.

You make out an open cuboid box. Empty! What kind of nightmare TextWorld is
this?

There is a closed door leading east. You need an unguarded exit? You should try
going north. You need an unblocked exit? You should try going west.

> look cuboid box

```

Figure 4: An example game generated by TextWorld to illustrate the command structure of a game.

The agent's observation(s) at any time in the game is the text information perceived by the agent. A probability function takes in the environment state and selects what information to show the agent based on the command entered. The agent receives points based on completion of (sub)quests and reaching new locations (exploring). This score could be used as the reward signal if it is available, otherwise positive reward signals can be assigned when the agent finishes the game. The agent's policy maps the state of the environment at any time and words generated in the command so far to the next word, which needs to be added to the command to maximise the reward received.

In an RL context, TextWorld games can be seen as partially observable Markov decision processes. The environment state at any turn t contains a complete description of the game state, but much of this is hidden from the agent. Once an agent has issued a command (of at least one word), the environment transitions to a next state with a certain probability. Since the interpreter in parser-based games can accept any sequence of characters (of any length), but only a fraction thereof is recognised, the resulting action space is very large. Therefore, two simplifying assumptions are made in Côté et al. (2018): the commands are sequences of at most L words taken from a fixed vocabulary V and the commands have to follow a specific structure: a verb, a noun phrase and an adverb phrase. The action space of the agent is therefore the set of all permissible commands from the fixed vocabulary V followed by a certain special token ("enter") that signifies the end of the command.

3.5.2 Benchmark tasks

TextWorld was introduced with two different sets of benchmark tasks (Côté et al., 2018) and a third task was added in the form of a competition that was available until 31 May 2019.

Task 1: A preliminary set of 50 hand-authored benchmark games are described in the original TextWorld paper (Côté et al., 2018). These games were manually analysed to ensure validity.

Task 2: This benchmark task is inspired by a treasure hunter task which takes place in a 3D environment (Parisotto & Salakhutdinov, 2017) and was adapted for TextWorld. The agent is randomly placed in a randomly generated map of rooms with two objects on the map. The goal object (the object which the agent should locate) is randomly selected and is mentioned in the welcome message. In order to navigate the map and locate the goal object, the agent may need to complete other tasks, for example finding a key to unlock a cabinet.

This task assesses the agent's skills of affordance extraction, efficient navigation and memory. There are different levels for the benchmark, ranging from level 1 to 30, with different difficulty modes, number of rooms and quest length.

Task 3: The TextWorld environment is still very new: TextWorld was only released to the public in July 2018. A competition – First TextWorld Problems: A Reinforcement and Language Learning Challenge¹⁶, which ran until 31 May 2019, was launched by Microsoft Research Montreal to challenge researchers to develop agents that can solve these text-based games. The challenge is gathering ingredients to cook a recipe.

Agents must determine the necessary ingredients from a recipe book, explore the house to gather ingredients, and return to the kitchen to cook up a delicious meal.

3.5.3 Published benchmark results

Côté et al. (2018) evaluate three baseline agents on the benchmark set in Task 1: BYU, Golovin and Simple. The BYU¹⁷ agent (Fulda et al., 2017) utilises a variant of Q-learning (Watkins & Dayan, 1992) where word embeddings are trained to be aware of verb-noun affordances. The agent won the IEEE CIG Text-based adventure AI Competition in 2016. The Golovin¹⁸ agent (Kostka et al., 2017) was developed specifically for classic text-based games and uses a language model pre-trained on fantasy books to extract important keywords from scene descriptions. The Simple¹⁹ agent uniformly samples a command from a predefined set at every step. Results indicated that all three baseline agents achieved low scores in the games. This indicates that there is significant scope for algorithms to improve on these results.

Côté et al. (2018) also provide average performance results of three agents (BYU, Golovin and a random agent) on 100 treasure hunter games (task 2) at different levels of difficulty. On difficulty level 1 the Golovin agents had the best average score, but the Random agent completed the game in the least number of steps. As the level of difficulty increase, the Random

¹⁷<https://github.com/danielricks/BYU-Agent-2016>

¹⁸https://github.com/Kostero/text_rpg_ai

¹⁹<https://github.com/Microsoft/TextWorld/tree/master/notebooks>

agent achieved the best score and also completed the game in the least number of steps. These results can be used as a baseline for evaluating improved algorithms.

It is evident that there is still enormous scope for research in the environment of text-based games, and that the generative functionality of the TextWorld sandbox environment is a significant contribution in the endeavour of researchers trying to solve these problems.

3.6 Summary

For the reader's convenience a summary of the discussed frameworks and algorithms that were shown to be effective are presented in Table 1. It should be noted that since the field moves at a rapid pace, the current state of the art will change (it may also be problem instance dependent within the benchmark class), however the listed algorithms can serve as a reasonable baseline for future research.

Framework	Benchmark class	Recent effective RL algorithm(s)
OpenAI Gym	Algorithmic	UREX (Nachum et al., 2017)
	Box2D	REINFORCE (Ha, 2019)
	Classic control	TNPG and TRPO (Duan et al., 2016)
	MuJoCo	PPO (Schulman et al., 2017)
	Roboschool	PPO (Schulman et al., 2017)
	Robotics	HER (Andrychowicz et al., 2017) ⁷
	Toy text	BIRL (Cundy & Filan, 2018)
The ALE	Atari 2600	A2C, ACER and PPO (Schulman et al., 2017); A3C (Mnih et al., 2016); Distribution DQN, Dueling DDQN, Prioritized DDQN and Rainbow (Hessel et al., 2018) ²⁰
Garage	Basic tasks	TNPG and TRPO (Duan et al., 2016)
	Locomotion tasks	PPO (Schulman et al., 2017)
	Partially observable tasks	TNPG and TRPO (Duan et al., 2016)
	Hierarchical tasks	HIRO (Nachum et al., 2018)
Keepaway soccer	Keepaway	Episodic SMDP SARSA(λ) (Stone, Sutton, & Kuhlmann, 2005; Sutton & Barto, 2018)
	Half-Field Offence	SARSA (Hausknecht et al., 2016)
TextWorld	Original tasks 1, 2 and 3	BYU and Golovin (Côté et al., 2018)
	Generalisation tasks	GATA (Adhikari et al., 2020)

Table 1: A summary of recent algorithms that performed well in different benchmark sets.

4 DISCUSSION

This section focuses on the ways that the different RL benchmarks discussed in Section 3 deal with or facilitate research in addressing the challenges for RL discussed in Section 2.2.

4.1 Partially observable environment

In many of the benchmark tasks, such as the classic control tasks in OpenAI Gym, the agent is provided with full information of the environment. The environment in TextWorld games, however, is partially observable since only local information and the player's inventory are available. The agent might also not be able to distinguish between some states based on observations if only the latest observation is taken into account, i.e. knowledge of past observations are important. In TextWorld games the environment might provide the same feedback for different commands and some important information about certain aspects of the environment might not be available by a single observation. Additionally, the agent might encounter observations that are time-sensitive, such as only being rewarded when it first examines a clue but not any other time. Controlling the partial observability of the state is also part of TextWorld's generative functionality. This is done by augmenting the agent's observations, where the agent can be provided with a list of present objects or even all game state information can be provided.

The partially observable tasks introduced in rllab (see Section 3.3.1), provide environments to investigate agents developed for dealing with environments where not all the information is known.

In RoboCup, a player can by default only observe objects in a 90-degree cone in front of them. In works from Kuhlmann and Stone (2003) and Stone, Sutton, and Kuhlmann (2005) it was shown that it is possible for learning to occur in this limited vision scenario, however players do not perform at an adequate level. For this reason, players in the standardised Keepaway task (Stone, Kuhlmann, et al., 2005) operate with 360-vision.

4.2 Delayed or sparse rewards

The tasks in the ALE and TextWorld are interesting when considering reward structure. In the ALE, reward or feedback may only be seen after thousands of actions. In TextWorld, the agent has to generate a sequence of actions before any change in the environment might occur or a reward is received. This results in sparse and delayed rewards in the games, in cases where an agent could receive a positive reward only after many steps when following an optimal strategy. In Keepaway, there is immediate reward, since the learners receive a positive reward after each action they execute.

²⁰A table summarising the best performance per game can be found at <https://github.com/cshenton/atari-leaderboard>.

4.3 Unspecified or multi-objective reward functions

In HFO (Section 3.4.4) success not only includes maintaining possession of the ball (the main objective in Keepaway), but the offense players also need to learn to pass or dribble to move towards the goal and shoot when an angle is open. Moreover, success is only evaluated based on a scored goal at the end of an episode, which is rare initially. This aspect of HFO could serve as an ideal environment for investigation into the challenge of problems with multi-objectives.

Due to the definition of a quest in TextWorld, i.e. a sequence of actions where each action depends on the outcomes of the previous action, quests in TextWorld are limited to simple quests. However, in text adventure games, quests are often more complicated, involving multiple sub-quests. Côté et al. (2018) remark that this limitation could be overcome by treating a quest as a directed graph of dependent actions rather than a linear chain. If this can be incorporated in TextWorld in the future, the platform can also be used to study problems with multi-objectives and rewards of varying difficulty.

4.4 Size of the state and action spaces

The benchmark tasks that are considered in this paper are ideal to investigate how the size of the state and/or action space challenge can be addressed. The tasks considered all have continuous or large discrete state spaces.

In the ALE the number of states in the games are very large and in TextWorld the state space is combinatorially enormous; since the number of possible states increases exponentially with the number of rooms and objects (Côté et al., 2018). In most of the tasks in OpenAI Gym, rllab, and in Keepaway, the state space is continuous. In Keepaway, the size of the Keepaway region can be varied along with the number of keepers and takers. This allows for investigation into a problem with various difficulties due to the size of the state space.

In TextWorld, the action space is large and sparse because the set of all possible word strings is much larger than the subset of valid commands. TextWorld's generative functionality also allows control over the size of the state space, i.e. the number of rooms, objects and commands. Different problem difficulties can therefore arise in terms of the size of the state space and this can aid in the investigation of algorithm behaviour with increasing state and action spaces.

4.5 The trade-off between exploration and exploitation

In the ALE the challenge of exploration vs. exploitation is difficult due to the large state spaces of games and delayed reward. Simple agents sometimes even learn that staying put is the best policy, since exploration can in some cases lead to negative rewards. Recently there has been some effort to address the exploration problem in the ALE, but these efforts are mostly successful only in individual games.

Exploration is fundamental to TextWorld games as solving them can not be done by learning a purely exploitative or reactive agent. The agent must use directed exploration as its

strategy, where it collects information about objects it encounters along the way. This information will provide knowledge about the goal of the game and provide insight into the environment and what might be useful later in the game. Due to this, exploration by curiosity driven agents might fair well in these types of problems.

Overall, there is still much work to be done to try and overcome this difficult challenge. Machado et al. (2018) suggest a few approaches for the ALE, such as agents capable of exploring in a more abstract manner (akin to humans) and agents not exploring joystick movements, but rather exploring object configurations and game levels. Agents with some form of intrinsic motivation might also be needed in order to continue playing even though achieving any reward might seem impossible.

4.6 Representation learning

The original goal of the ALE was to develop agents capable of generalising over many games making it desirable to automatically learn representations instead of hand crafting features. Deep Q-Networks (DQN) (Mnih et al., 2015) and DQN-like approaches are currently the best overall performing methods, despite high sample complexity. However, additional tuning is often required to obtain better performance (Islam et al., 2017), which suggest that there is still work to be done to improve performance by learning better representation in the ALE. Other different approaches and directions for representation learning that have been used in the literature are also mentioned in Machado et al. (2018) and should still be explored more in the ALE.

4.7 Transfer learning

Regarding the ALE, many of the Atari 2600 games have similar game dynamics and knowledge transfer should reduce the number of samples that are required to learn to play games that are similar. Even more challenging would be determining how to use general video game experience and share that knowledge across games that are not necessarily similar. Current approaches in the literature that apply transfer learning in the ALE are restricted to only a limited subset of games that share similarities and the approaches are based on using neural networks to perform transfer, combining representations and policy transfer. Machado et al. (2018) point out that it might be interesting to determine whether transferring each of these entities independently could be helpful. To help with the topic of transfer learning in the ALE, the new version includes different game modes and difficulty settings called flavours (see Section 3.2), which introduces many new environments that are very similar.

Some of the tasks in rllab and environments in OpenAI Gym have been used in studying the transferring of system dynamics from simulation to robots (Held et al., 2017; Peng et al., 2018; Wulfmeier et al., 2017). These simulation tasks are an ideal way to safely study the transferring of policies for robotic domains.

Transfer learning has also been studied in the Keepaway soccer domain (Taylor & Stone,

2005), which is a fitting setting since the number of players as well as the size of the action and state spaces can differ.

TextWorld's generative functionality (described in full in Côté et al. (2018)) allows for control of the size and the partial observability of the state space, and therefore a large number of games with shared characteristics can be generated. This could be used for studying transfer learning in text-based games, since agents can be trained on simpler tasks and behaviour transferred to harder problems.

4.8 Model learning

Planning and model learning in complex domains are challenging problems and little research has been conducted on this topic compared to traditional RL techniques to learn policies or value functions.

In the ALE, the Stella emulator provides a generative model that can be used in planning and the agent has an exact model of the environment. However, there has not been any success with planning using a learned generative model in the ALE, which is a challenging task since errors start to compound after only a few time steps. A few relatively successful approaches (Chiappa et al., 2017; Oh et al., 2015) are available, but the models are slower than the emulator. A challenging open problem is to learn a fast and accurate model for the ALE. On the other hand, related to this, is the problem of planning using an imperfect model.

On tasks in OpenAI Gym and rllab some research has also been conducted in model learning (Nagabandi et al., 2018; Wang et al., 2019), but the main focus in the literature is on model-free learning techniques. Therefore there is still scope for substantial research to address this problem.

Wang et al. (2019) attempted to address the lack of a standardised benchmarking framework for model-based RL. They benchmarked 11 model-based RL algorithms and four model-free RL algorithms across 18 environments from OpenAI Gym and have shared the code in an online repository²¹. They evaluated the efficiency, performance and robustness of three different categories of model-based RL algorithms (Dyna style algorithms, policy search with back-propagation through time and shooting algorithms) and four model-free algorithms (TRPO, PPO, TD3, and SAC – refer to Section 2.1.8 for these algorithms). They also propose three key research challenges for model-based methods, namely the dynamics bottleneck, the planning horizon dilemma, and the early termination dilemma and show that even with substantial benchmarking, there is no clear consistent best model-based RL algorithm. This again suggests that there is substantial scope and many opportunities for further research in model-based RL methods.

²¹<http://www.cs.toronto.edu/~tingwuwang/mbrl.html>

4.9 Off-policy learning

Deep neural networks have become extremely popular in modern RL literature, and the breakthrough work of Mnih et al. (2013; 2015) demonstrates DQN having human-level performance on Atari 2600 games. However, when using deep neural networks for function approximation for off-policy algorithms, new and complex challenges arise, such as instability and slow convergence. While discussing off-policy methods using function approximation, Sutton and Barto (2018) conclude the following: “The potential for off-policy learning remains tantalizing, the best way to achieve it still a mystery.” Nevertheless, off-policy learning has become an active research field in RL.

The use of off-policy learning algorithms in the ALE in current literature varies with most approaches using experience replay and target networks. This is an attempt at reducing divergence in off-policy learning, but these methods are very complex. New proposed algorithms such as GQ(λ) (Maei & Sutton, 2010) are theoretically sound, but there is still a need for a thorough empirical evaluation or demonstration of these theoretically sound off-policy learning RL algorithms. Other contributions of using off-policy learning in the ALE includes double Q-learning (van Hasselt et al., 2016) and Q(λ) with off-policy corrections (Harutyunyan et al., 2016).

Some of the tasks in rllab and OpenAI Gym have also been used in studying off-policy algorithms, for example introducing the soft actor-critic (SAC) algorithm (Haarnoja et al., 2018) and using the robotics environments from OpenAI Gym to learn grasping (Quillen et al., 2018). This area of research is still new and there is significant scope for further research in this domain.

4.10 Reinforcement learning in real-world settings

The robotics environments in the OpenAI Gym toolkit can be used to train models which work on physical robots. This can be used to develop agents to safely execute realistic tasks. A request for research from OpenAI⁷ indicates that work in this area is an active research field with promising results.

The Keepaway and HFO soccer tasks are ideal settings to study multi-agent RL (Buşoniu et al., 2008), an important research area for real-world problems since humans act in an environment where objectives are shared with others.

Challenges for RL that are unique to TextWorld games are related to natural language understanding: observation modality, understanding the parser feedback, common-sense reasoning and affordance extraction, and language acquisition. These challenges are explained in more detail in Côté et al. (2018). Natural language understanding is an important aspect of artificial intelligence, in order for communication to take place between humans and AI. TextWorld can be used to address many of the challenges described in Section 2.2 in simpler settings and to focus on testing and debugging agents on subsets of these challenges.

In addition to the frameworks covered in this survey, there are two further contributions that are focused on multi-agent and distributed RL. The MAgent research platform (Zheng et

al., 2018) facilitates research in many-agent RL, specifically in artificial collective intelligence. The platform aims at supporting RL research that scales up from hundreds to millions of agents and is maintained in an online repository²². MAgent also provides a visual interface presenting the state of the environment and agents.

A research team from Stanford has introduced the open-source framework SURREAL (Scalable Robotic REinforcementlearning ALgorithms) and the SURREAL Robotics Suite (Fan et al., 2018), to facilitate research in RL in robotics and distributed RL. SURREAL eliminates the need for global synchronization and improves scalability by decoupling a distributed RL algorithm into four components. The four-layer computing infrastructure can easily be deployed on commercial cloud providers or personal computers, and is also fully replicable from scratch, contributing to the reproducibility of results. The Robotics Suite is developed in the MuJoCo physics engine and provides OpenAI gym-style interfaces in Python. Detailed API documentation and tutorials on importing new robots and the creation of new environments and tasks are also provided, furthering the contribution to research in this field. The Robotics Suite is actively maintained in an online repository²³. The different robotics tasks include block lifting and stacking, bimanual peg-in-hole placing and bimanual lifting, bin picking, and nut-and-peg assembly. Variants of PPO and DDPG called SURREAL-PPO and SURREAL-DDPG were developed and examined on the Robotics Suite tasks, and experiments indicate that these SURREAL algorithms can achieve good results.

4.11 A standard methodology for benchmarking

The ALE consists of games with similar structure in terms of inputs, action movements, etc. This makes the ALE an ideal benchmark for comparative studies. A standard methodology is however needed and this is proposed by Machado et al. (2018):

- Episode termination can be standardised by using the game over signal than lives lost.
- Hyperparameter tuning needs to be consistently applied on the training set only.
- Training time should be consistently applied across different problems.
- There is a need for standard ways of reporting learning performance.

These same principles apply to groups of similar tasks in OpenAI Gym and rllab, and to TextWorld and Keepaway soccer.

4.12 Trends in benchmarking of RL

It is clear from Section 3 that the number of well thought-out frameworks designed for RL benchmarks has rapidly expanded in recent years, with a general move to fully open source

²²<https://github.com/geek-ai/MAgent>

²³<https://github.com/SurrealAI/surreal>

implementations being evident. A notable example is OpenAI Gym re-implementing, to an extent, open source variants of the benchmarks previously provided in the MuJoCo simulation environment. The move to fully open source implementations has had two primary benefits: reproducibility and accessibility.

The variety of RL frameworks and benchmark sets may present a challenge to a novice in the field, as there is no clear standard benchmark set or framework to use. This is not a surprising situation as the array of RL application areas has become relatively diverse and so different types of problems and their corresponding challenges will naturally be more interesting to certain sub-communities within the field.

One aspect of modern RL benchmarks that is relatively striking is the increase in problem complexity. While it is not immediately clear how to precisely define problem difficulty, it is clear that more and more problem features that are challenging for RL algorithms are being included in proposed benchmarks. Many established benchmark sets have been explicitly expanded to increase the challenge of a given problem instance. Some notable examples include the addition of sticky actions in the ALE and the addition of the partially observable variants of rllab's continuous control tasks.

It is also clear that the advancements made in the field of deep learning has allowed for certain types of RL tasks to be more readily solvable. Two notable examples are the use of convolution neural networks (Lecun et al., 1998) to assist in the vision problem present in Atari 2600 games of the ALE, and the use of modern neural network based approaches to natural language processing in Microsoft's TextWorld.

5 CONCLUSION

This paper provides a survey of some of the most used and recent contributions to RL benchmarking. A number of benchmarking frameworks are described in terms of their characteristics, technical implementation details and the tasks provided. A summary is also provided of published results on the performance of algorithms used to solve these benchmark tasks. Challenges that occur when solving RL problems are also discussed, including the various ways the different benchmarking tasks address or facilitate research in addressing these challenges.

The survey reveals that there has been substantial progress in the endeavour of standardising benchmarking tasks for RL. The research community has started to acknowledge the importance of reproducible results and research has been published to encourage the community to address this problem. However, there is still a lot to be done in ensuring the reproducibility of results for fair comparison.

There are many approaches when solving RL problems and proper benchmarks are important when comparing old and new approaches. This survey indicates that the tasks currently used for benchmarking RL encompass a wide range of problems and can even be used to develop algorithms for training agents in real-world systems such as robots.

References

- Adhikari, A., Yuan, X., Côté, M. A., Zelinka, M., Rondeau, M. A., Laroche, R., Poupart, P., Tang, J., Trischler, A., & Hamilton, W. L. (2020). *Learning dynamic knowledge graphs to generalize on text-based games* [arXiv preprint arXiv:2002.09127].
- Albus, J. S. (1975). A new approach to manipulator control: The cerebellar model articulation controller (CMAC). *Journal of Dynamic Systems, Measurement, and Control*, 97(3), 220–227. <https://doi.org/10.1115/1.3426922>
- Albus, J. S. (1981). *Brains, Behavior and Robotics*. McGraw-Hill, Inc.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., & Zaremba, W. (2017). Hindsight experience replay. *Advances in Neural Information Processing Systems* 30, 5048–5058.
- Arel, I., Liu, C., Urbanik, T., & Kohls, A. (2010). Reinforcement learning-based multi-agent system for network traffic signal control. *IET Intelligent Transport Systems*, 4(2), 128–135. <https://doi.org/10.1049/iet-its.2009.0070>
- Bai, A., & Russell, S. (2017). Efficient reinforcement learning with hierarchies of machines by leveraging internal transitions. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 1418–1424. <https://doi.org/10.24963/ijcai.2017/196>
- Bellemare, M. G., Dabney, W., & Munos, R. (2017). A distributional perspective on reinforcement learning. *Proceedings of the 34th International Conference on Machine Learning*, 449–458.
- Bellemare, M. G., Naddaf, Y., Veness, J., & Bowling, M. (2013). The Arcade Learning Environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47, 253–279. <https://doi.org/10.1613/jair.3912>
- Bellman, R. E. (1957). *Dynamic programming*. Princeton University Press.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). *OpenAI Gym* [arXiv:1606.01540].
- Buşoniu, L., Babuška, R., & De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2), 156–172. <https://doi.org/10.1109/TSMCC.2007.913919>
- Cheng, Q., Wang, X., Niu, Y., & Shen, L. (2018). Reusing source task knowledge via transfer approximator in reinforcement transfer learning. *Symmetry*, 11(1). <https://doi.org/10.3390/sym11010025>
- Chiappa, S., Racaniere, S., Wierstra, D., & Mohamed, S. (2017). *Recurrent environment simulators* [arXiv:1704.02254].
- Côté, M. A., Kádár, Á., Yuan, X., Kybartas, B., Barnes, T., Fine, E., Moore, J., Hausknecht, M. J., Asri, L. E., Adada, M., Tay, W., & Trischler, A. (2018). *TextWorld: A learning environment for text-based games* [arXiv:1806.11532].
- Cundy, C., & Filan, D. (2018). *Exploring hierarchy-aware inverse reinforcement learning* [arXiv:1807.05037].

- Dann, C., Neumann, G., & Peters, J. (2014). Policy evaluation with temporal differences: A survey and comparison. *Journal of Machine Learning Research*, 15, 809–883.
- Didi, S., & Nitschke, G. (2016a). Hybridizing novelty search for transfer learning. *IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–8. <https://doi.org/10.1109/SSCI.2016.7850180>
- Didi, S., & Nitschke, G. (2016b). Multi-agent behavior-based policy transfer. *European Conference on the Applications of Evolutionary Computation*, 181–197. https://doi.org/10.1007/978-3-319-31153-1_13
- Didi, S., & Nitschke, G. (2018). Policy transfer methods in RoboCup keep-away. *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 117–118. <https://doi.org/10.1145/3205651.3205710>
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., & Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, 1329–1338.
- Dulac-Arnold, G., Mankowitz, D., & Hester, T. (2019). Challenges of real-world reinforcement learning [arXiv:1904.12901]. *Reinforcement Learning for Real Life (RL4RealLife) Workshop in the 36th International Conference on Machine Learning*.
- Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2010). The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88(2), 303–338. <https://doi.org/10.1007/s11263-009-0275-4>
- Fan, L., Zhu, Y., Zhu, J., Liu, Z., Zeng, O., Gupta, A., Creus-Costa, J., Savarese, S., & Fei-Fei, L. (2018). SURREAL: Open-source reinforcement learning framework and robot manipulation benchmark. *Proceedings of The 2nd Conference on Robot Learning*, 87, 767–782.
- Fujimoto, S., van Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. *Proceedings of the 35th International Conference on Machine Learning*, 1587–1596.
- Fulda, N., Ricks, D., Murdoch, B., & Wingate, D. (2017). What can you do with a rock? Affordance extraction via word embeddings. *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 1039–1045. <https://doi.org/10.24963/ijcai.2017/144>
- Ha, D. (2019). Reinforcement learning for improving agent design. *Artificial Life*, 25(4), 352–365. https://doi.org/10.1162/artl_a_00301
- Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *Proceedings of the 35th International Conference on Machine Learning*, 1861–1870.
- Hansen, N., & Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2), 159–195. <https://doi.org/10.1162/106365601750190398>
- Harutyunyan, A., Bellemare, M. G., Stepleton, T., & Munos, R. (2016). $Q(\lambda)$ with off-policy corrections. *Algorithmic Learning Theory*, 305–320. https://doi.org/10.1007/978-3-319-46379-7_21

- Hausknecht, M., Mupparaju, P., Subramanian, S., Kalyanakrishnan, S., & Stone, P. (2016). Half field offense: An environment for multiagent learning and ad hoc teamwork. *AAMAS Adaptive Learning Agents (ALA) Workshop*.
- Held, D., McCarthy, Z., Zhang, M., Shentu, F., & Abbeel, P. (2017). Probabilistically safe policy transfer. *IEEE International Conference on Robotics and Automation (ICRA)*, 5798–5805. <https://doi.org/10.1109/ICRA.2017.7989680>
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., & Meger, D. (2018). Deep reinforcement learning that matters. *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 3207–3214.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., & Silver, D. (2018). Rainbow: Combining improvements in deep reinforcement learning. *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Islam, R., Henderson, P., Gomrokchi, M., & Precup, D. (2017). Reproducibility of benchmarked deep reinforcement learning tasks for continuous control [arXiv:1708.04133]. *ICML Workshop on Reproducibility in Machine Learning*.
- Kalyanakrishnan, S., Liu, Y., & Stone, P. (2007). Half Field Offense in RoboCup Soccer: A Multiagent Reinforcement Learning Case Study. *RoboCup 2006: Robot Soccer World Cup X*, 72–85. https://doi.org/10.1007/978-3-540-74024-7_7
- Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., & Osawa, E. (1997). RoboCup: The Robot World Cup Initiative. *Proceedings of the First International Conference on Autonomous Agents*, 340–347. <https://doi.org/10.1145/267658.267738>
- Kostka, B., Kwiecei, J., Kowalski, J., & Rychlikowski, P. (2017). Text-based adventures of the Golovin AI agent. *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, 181–188. <https://doi.org/10.1109/CIG.2017.8080433>
- Kuhlmann, G., & Stone, P. (2003). Progress in learning 3 vs. 2 Keepaway. *IEEE International Conference on Systems, Man and Cybernetics*, 52–59. <https://doi.org/10.1109/ICSMC.2003.1243791>
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324. <https://doi.org/10.1109/5.726791>
- Liang, Y., Machado, M. C., Talvitie, E., & Bowling, M. (2016). State of the art control of Atari games using shallow reinforcement learning. *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, 485–493.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2015). *Continuous control with deep reinforcement learning* [arXiv:1509.02971].
- Machado, M. C., Bellemare, M. G., Talvitie, E., Veness, J., Hausknecht, M., & Bowling, M. (2018). Revisiting the Arcade Learning Environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61, 523–562.
- Maei, H. R., & Sutton, R. S. (2010). GQ(λ): A general gradient algorithm for temporal-difference prediction learning with eligibility traces. *3rd Conference on Artificial General Intelligence (AGI-2010)*. <https://doi.org/10.2991/agi.2010.22>

- Marcus, M., Santorini, B., & Marcinkiewicz, M. A. (1993). Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2), 313–330.
- Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1), 8–30. <https://doi.org/10.1109/JRPROC.1961.287775>
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., & Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, 1928–1937.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, A., Wierstra, A., & Riedmiller, M. (2013). *Playing Atari with deep reinforcement learning* [arXiv:1312.5602].
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518, 529. <https://doi.org/10.1038/nature14236>
- Nachum, O., Gu, S., Lee, H., & Levine, S. (2018). Data-efficient hierarchical reinforcement learning. *Advances in neural information processing systems 31* (pp. 3303–3313). Curran Associates, Inc. <http://papers.nips.cc/paper/7591-data-efficient-hierarchical-reinforcement-learning.pdf>
- Nachum, O., Norouzi, M., & Schuurmans, D. (2017). Improving policy gradient by exploring under-appreciated rewards [arXiv:1611.09321]. *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Nagabandi, A., Kahn, G., Fearing, R. S., & Levine, S. (2018). Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. *IEEE International Conference on Robotics and Automation (ICRA)*, 7559–7566. <https://doi.org/10.1109/ICRA.2018.8463189>
- Nitschke, G., & Didi, S. (2017). Evolutionary policy transfer and search methods for boosting behavior quality: Robocup keep-away case study. *Frontiers in Robotics and AI*, 4, 62. <https://doi.org/10.3389/frobt.2017.00062>
- Oh, J., Guo, X., Lee, H., Lewis, R. L., & Singh, S. (2015). Action-conditional video prediction using deep networks in Atari games. *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, 2863–2871.
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>
- Papis, B., & Wawrzyński, P. (2013). DotRL: A platform for rapid reinforcement learning methods development and validation. *2013 Federated Conference on Computer Science and Information Systems*, 129–136.
- Parisotto, E., & Salakhutdinov, R. (2017). *Neural map: Structured memory for deep reinforcement learning* [arXiv:1702.08360].

- Peng, X. B., Andrychowicz, M., Zaremba, W., & Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. *IEEE International Conference on Robotics and Automation (ICRA)*, 1–8. <https://doi.org/10.1109/ICRA.2018.8460528>
- Peters, J., Mulling, K., & Altun, Y. (2010). Relative entropy policy search. *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 1607–1612.
- Peters, J., & Schaal, S. (2007). Reinforcement learning by reward-weighted regression for operational space control. *Proceedings of the 24th International Conference on Machine Learning*, 745–750. <https://doi.org/10.1145/1273496.1273590>
- Pietro, A. D., While, L., & Barone, L. (2002). Learning in RoboCup keepaway using evolutionary algorithms. *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation*, 1065–1072.
- Quillen, D., Jang, E., Nachum, O., Finn, C., Ibarz, J., & Levine, S. (2018). Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods. *IEEE International Conference on Robotics and Automation (ICRA)*, 6284–6291. <https://doi.org/10.1109/ICRA.2018.8461039>
- Riedmiller, M., Gabel, T., Hafner, R., & Lange, S. (2009). Reinforcement learning for robot soccer. *Autonomous Robots*, 27(1), 55–73. <https://doi.org/10.1007/s10514-009-9120-4>
- Rubinstein, R. (1999). The cross-entropy method for combinatorial and continuous optimization. *Methodology And Computing In Applied Probability*, 1(2), 127–190. <https://doi.org/10.1023/A:1010091220143>
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- Schaul, T., Bayer, J., Wierstra, D., Sun, Y., Felder, M., Sehnke, F., Rückstieß, T., & Schmidhuber, J. (2010). PyBrain. *Journal of Machine Learning Research*, 11(Feb), 743–746.
- Schulman, J., Levine, S., Moritz, P., Jordan, M., & Abbeel, P. (2015). Trust region policy optimization. *Proceedings of the 32nd International Conference on Machine Learning*, 1889–1897.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). *Proximal policy optimization algorithms* [arXiv:1707.06347].
- Schwab, D., Zhu, Y., & Veloso, M. (2018). Zero shot transfer learning for robot soccer. *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2070–2072.
- Silva, V. d. N., & Chaimowicz, L. (2017). *MOBA: A new arena for game AI* [arXiv:1705.10443].
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529, 484–489. <https://doi.org/10.1038/nature16961>

- Silver, D., Sutton, R. S., & Müller, M. (2012). Temporal-difference search in computer Go. *Machine Learning*, 87(2), 183–219. <https://doi.org/10.1007/s10994-012-5280-0>
- Stone, P. (2000). *Layered learning in multiagent systems: A winning approach to robotic soccer*. MIT Press.
- Stone, P., Kuhlmann, G., Taylor, M. E., & Liu, Y. (2005). Keepaway soccer: From machine learning testbed to benchmark. *Robot Soccer World Cup*, 93–105. https://doi.org/10.1007/11780519_9
- Stone, P., & Sutton, R. S. (2001). Keepaway soccer: A machine learning test bed. *Robot Soccer World Cup*, 214–223. https://doi.org/10.1007/11780519_9
- Stone, P., Sutton, R. S., & Kuhlmann, G. (2005). Reinforcement learning for RoboCup soccer Keepaway. *Adaptive Behavior*, 13(3), 165–188. <https://doi.org/10.1177/105971230501300301>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- Taylor, M. E., & Stone, P. (2005). Behavior transfer for value-function-based reinforcement learning. *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, 53–59. <https://doi.org/10.1145/1082473.1082482>
- Taylor, M. E., & Stone, P. (2009). Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul), 1633–1685.
- Todorov, E., Erez, T., & Tassa, Y. (2012). MuJoCo: A physics engine for model-based control. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033. <https://doi.org/10.1109/IROS.2012.6386109>
- van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double Q-learning. *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2094–2100.
- Vinyals, O., Ewalds, T., Bartunov, S., Georgiev, P., Vezhnevets, A. S., Yeo, M., Makhzani, A., Küttler, H., Agapiou, J., Schrittwieser, J., Quan, J., Gaffney, S., Petersen, S., Simonyan, K., Schaul, T., van Hasselt, H., Silver, D., Timothy Lillicrap, T., Kevin Calderone, K., ... Tsing, R. (2017). *StarCraft II: A new challenge for reinforcement learning* [arXiv:1708.04782], Collaboration between DeepMind and Blizzard.
- Walker, T., Shavlik, J., & Maclin, R. (2004). Relational reinforcement learning via sampling the space of first-order conjunctive features. *Proceedings of the ICML Workshop on Relational Reinforcement Learning*.
- Wang, T., Bao, X., Clavera, I., Hoang, J., Wen, Y., Langlois, E., Zhang, S., Zhang, G., Abbeel, P., & Ba, J. (2019). *Benchmarking model-based reinforcement learning* [arXiv:1907.02057].
- Watkins, J. C. H., & Dayan, P. (1992). Q-learning. *Machine learning*, 8(3-4), 279–292. <https://doi.org/10.1007/BF00992698>
- Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big data*, 3(9). <https://doi.org/10.1186/s40537-016-0043-6>
- Williams, R. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8, 229. <https://doi.org/10.1007/BF00992696>

- Woolley, B. G., & Stanley, K. O. (2010). Evolving a single scalable controller for an octopus arm with a variable number of segments. *Parallel Problem Solving from Nature, PPSN XI*, 270–279. https://doi.org/10.1007/978-3-642-15871-1_28
- Wulfmeier, M., Posner, I., & Abbeel, P. (2017). Mutual alignment transfer learning. *Proceedings of the 1st Annual Conference on Robot Learning*.
- Yamaguchi, A., & Ogasawara, T. (2010). SkyAI: Highly modularized reinforcement learning library. *2010 10th IEEE-RAS International Conference on Humanoid Robots*, 118–123. <https://doi.org/10.1109/ICHR.2010.5686285>
- Zheng, L., Yang, J., Cai, H., Zhou, M., Zhang, W., Wang, J., & Yu, Y. (2018). MAgent: A many-agent reinforcement learning platform for artificial collective intelligence. *Thirty-Second AAAI Conference on Artificial Intelligence*.