

---

# TOYBOX: Better Atari Environments for Testing Reinforcement Learning Agents

---

**John Foley\***

Department of Computer Science  
Smith College  
Northampton, MA 01060  
jfoley@smith.edu

**Emma Tosch\***

College of Information and Computer Science  
University of Massachusetts Amherst  
Amherst, MA 01060  
etosch@cs.umass.edu

**Kaleigh Clary**

College of Information and Computer Science  
University of Massachusetts Amherst  
Amherst, MA 01060  
kclary@cs.umass.edu

**David Jensen**

College of Information and Computer Science  
University of Massachusetts Amherst  
Amherst, MA 01060  
jensen@cs.umass.edu

## Abstract

It is a widely accepted principle that software without tests has bugs. Testing reinforcement learning agents is especially difficult because of the stochastic nature of both agents and environments, the complexity of state-of-the-art models, and the sequential nature of their predictions. Recently, the Arcade Learning Environment (ALE) has become one of the most widely used benchmark suites for deep learning research, and state-of-the-art Reinforcement Learning (RL) agents have been shown to routinely equal or exceed human performance on many ALE tasks. Since ALE is based on *emulation* of original Atari games, the environment does not provide semantically meaningful representations of internal game state. This means that ALE has limited utility as an environment for supporting testing or model introspection. We propose TOYBOX, a collection of reimplementations of these games that solves this critical problem and enables robust testing of RL agents.

## 1 Introduction

Advances in deep models have fueled impressive gains in Reinforcement Learning (RL). Agents trained using deep RL can learn control policies for a variety of complex tasks using only sensory input [1, 9]. One of the most impressive benchmarks is the set of Atari games served by the Arcade Learning Environment (ALE) [2, 13, 14] where deep models can often outperform human players. Consequently, Atari games have become the *de facto* academic and industrial benchmarks for RL. Frameworks such as OpenAI Gym have further lowered the barrier to entry for training RL agents, and have made popular models available off the shelf [4, 7].

However, these original Atari emulations are functionally black boxes, making it highly impractical to view and manipulate internal game state for testing purposes. Since RL agents are learned software, it is imperative that we can test these black-box software agents. This is difficult when the environment is also both complex and opaque.

ALE was originally designed strictly as a benchmark suite to enable comparison of the generalization properties of RL models [2]. In practice, however, it is now common for researchers to develop new

---

\* Authors contributed equally.

RL models using these Atari games. In a field where there are few instances of non-trivial separate training and testing environments, this trend presents challenges for how to address ALE’s original goal.

The central challenge for testing RL agent behavior is that typically both the agent and the environment are black boxes. When an RL agent achieves human or super-human performance on a complex task, researchers want to make claims about the reasoning power and utility of their trained agents. This requires the ability to form *and test* behavioral hypotheses in the environment. Therefore, the ideal RL test environment is a white-box environment with easily manipulable and composable parameters that permit the expression of high-level concepts. Such an environment ought to avoid design biases that would make the environment more tractable for learning.

Our core contribution is the development of TOYBOX, a suite of fully parameterizable, replicated Atari games. TOYBOX facilitates a host of analyses not previously possible in ALE, including, but not limited to:

- **Post-Training Acceptance Testing:** An agent that has learned a concept ought to be able to perform similarly in similar states. TOYBOX permits the programmatic generation of similar states. An agent can begin interacting with an intervened-upon TOYBOX implementation of a game at any point during gameplay. This means that we can accept or reject agents based on their ability to succeed in critical tasks under laboratory conditions.
- **Dynamic Analysis during Training:** RL training can take a considerable amount of wall-clock time. Therefore, the cost of getting stuck in a local minimum can be high [10]. Dynamic analysis of agent behavior could help diagnose problematic reward functions or random seeding earlier rather than later.
- **Test Set Generation:** Because TOYBOX games are parameterizable, users can specify a family of games by varying some of the parameter values. The specific features of what that entails will depend on the task. These families of games would be specifically useful for curriculum learning.

In this paper, we focus on the first novel analysis made possible via TOYBOX: that of determining the degree to which a specific, learned agent meets performance expectations expressed in human-understandable terms. To do this, we focus on *Breakout*, a specific game implemented within both ALE and TOYBOX. Testing whether a learned agent meets these performance expectations would be virtually impossible in ALE, but such testing is quite simple in TOYBOX.

**Example: Breakout.** With TOYBOX, we are able to convert speculation about RL agent behavior into testable *behavioral requirements*. We focus on the game of Breakout, as some of the most high profile instances of anthropomorphization have been expressed about RL agents playing this game [15]. We identify and test three behavioral requirements related to the playing of Breakout by a black-box RL agent:

- R1** An agent will be able to quickly eliminate each brick in isolation.
- R2** An agent will be able to hit the ball regardless of the angle of its initial trajectory.
- R3** An agent will quickly open tunnels to the roof, eagerly seeking the high reward of the titular "breaking-out" behavior.

These requirements are fairly straightforward, but it is not clear how to test them with ALE. Regardless of whether they correctly express what we mean when we say an agent has “learned” to play Breakout, TOYBOX makes it possible to express them and our core contribution is the ability to test that an agent meets requirements of this style.

In order to conduct our experiments, we trained an RL agent using the recommended PP02 algorithm, and the default parameter settings available in the OpenAI Baselines implementation [17, 7]. We find that evidence supports **R2**, but **R1** is complicated by certain bricks being more challenging than others and **R3** is completely unsupported: our tested agents are only good at making one specific tunnel. Before TOYBOX these hypotheses were completely untestable in the existing ALE environment without significant time invested into reverse-engineering RAM locations.

## 2 Background and Related Work

Reinforcement learning is a subdomain of machine learning wherein an *agent* interacts with an *environment* via sequential *actions*, seeking to maximize its expected cumulative *reward*. The agent operates in the environment by taking actions to transition between environment *states* [18].

Actions are typically selected according to a *policy*, which is a potentially stochastic function mapping states to actions. There are many methods for learning a policy, but essentially all RL agents learn by repeatedly acting in the environment and using the resulting experience to alter its policy. Deep reinforcement methods implement the policy function as a deep neural net.

Arcade Learning Environment exports Atari states either as 160x210 pixel images, or as 1024 bit RAM. Actions in these environments are a set of inputs that can be applied to the joysticks, which have been discretized by ALE wrappers. The ALE designers have been able to identify the current game score in RAM, and extract it. Typical implementations of learners define the reward to be the truncated normalized score, i.e., +1, 0, -1 if the score has increased, not changed, or decreased, respectively. The agent is a computer program that returns joystick and button actions in response to Atari states.

**Simulation and Testing Environments.** Historically, RL researchers have created their own environments to test the behavior of RL algorithms. Environments drawn from the classic GridWorld family consist of an agent moving through a grid of states from a fixed start state toward a fixed goal state, typically with one or more obstacles or other complicating factors (e.g., teleporting states) [18]. In this simplified world, the optimal policy is directly computable via dynamic programming.

Such simple environments are critically important for testing theoretical properties of RL agents' learned policies. Complex environments designed for human interaction do not typically have tractable ways of computing optimal policies. Among these more complicated environments, ALE is one of the most popular benchmarks, but real-time strategy games such as StarCraft, and simulations of real-world robotics problems using the MuJoCo physics engine, are also popular [1, 2, 19, 21, 22].

Despite these recent achievements, the gap between toy problems and human domains remains quite large. Some authors have focused on generating families of theoretically understood games, which can help bridge this gap [16]. Others have shown that, in sufficiently complex systems, the only feasible way to understand what an agent is learning is to build fast, intervenable systems [3, 12, 20].

The motivation for intervention has been largely about increasing variability during training. These systems all implemented either completely new games or simplified versions of existing games, having specialized parameterized implementations for the purposes of RL research. By re-implementing Atari games in TOYBOX we avoid any research-specific bias that might influence game design.

**Challenges to Testing in Reinforcement Learning.** In RL, it is often the case that the same environment is used for training, testing, and deployment. In fact, in the case of Atari games, testing and deployment are effectively the same: a researcher trains a new model using a particular game and then shows how well the agent learned by running the agent in that same environment.

While there has certainly been criticism of the field for testing in the training environment [23], there is a legitimate justification for doing so: if an agent is only to be deployed in the training environment, it is perfectly reasonable to overfit the agent's policy to that environment. However, claims about higher-level reasoning or generalization capabilities require testing, where testing means something more than simply running a trained agent in the environment.

Testing as construed in this paper is not the typical "testing" as referred to in supervised machine learning, where a classifier's ability to generalize is measured against a held-out evaluation set. Instead, we are arguing for a form of *acceptance testing* where we take a set of trained RL agents and select the successful ones based on tests written to ensure conceptual accuracy. While we believe that this kind of testing and monitoring of models is present in industrial applications, there is little published work describing these procedures.

Testing in the software engineering sense is notoriously difficult for machine learning. The lack of determinism, combined with at times incredibly byzantine software layers and few theoretical guarantees, makes any kind of testing non-trivial. The closest analogue to unit testing would be




Figure 1: (Left: ALE. Right: TOYBOX.) Images of near-start frames for both Atari and TOYBOX implementations of Breakout.

training agents on a variety of classical problems, e.g., GridWorld experiments to show the capabilities of new algorithms. There are no common practices equivalent to integration, nor regression testing.

### 3 System Design

TOYBOX is a growing suite of Atari reimplementations, written in Rust. All games are written to be highly parameterized, with swappable components. Game state can be exported at any time, altered, and restarted from the intervened-upon state.

We have implemented OpenAI Gym environments that are drop-in replacements for the ALE OpenAI Gym implementations through TOYBOX’s Python API. This means that, for example, a user can train on the TOYBOX version of Breakout simply by registering and then loading the TOYBOX Breakout OpenAI environment. With care taken to ensure the same preprocessing steps are used on TOYBOX as on ALE Atari games, the same agents and networks can be trained on the testable TOYBOX interface instead.

Our fidelity to the original game has reached a level where the visual output between environments is nearly indistinguishable (see Figure 1), but remains a work in progress. We can already ask meaningful semantic questions through the TOYBOX API. In the future, we hope to improve our fidelity to the original games further, but since so much of Atari games is implemented in hardware and lookup tables, it is nearly impossible to code game physics that match Atari games bug-for-bug. While certain videogame bugs, such as the infamous overflow error in PacMan [8], are worth reproducing, we’ve focused this current work on testing desired properties on an environment that is simply indistinguishable by humans.

### 4 Semantic Tests

As papers are published detailing new RL models for Atari playing, it is only natural to speculate as to the behaviors learned and understood by agents. TOYBOX provides us the ability to test these types of claims; here we test the three simple requirements (**R1**, **R2**, **R3** from §1) as a demonstration of TOYBOX’s versatility and analytic power.

**Shared Test Setup.** All of the tests presented were run using two agents trained using the default settings for the PP02 algorithm available in OpenAI Gym Baselines [7, 17]. One agent was trained for 10 million steps ( $1e7$ ) and the other for 50 million steps ( $5e7$ ) in the TOYBOX Breakout environment. The  $1e7$  model achieves an average score of 407 and standard deviation of 86 points. The  $5e7$  model achieves an average score of 777, and a standard deviation of 154 points. In order to win Breakout, an agent must clear two identical levels, where each level is worth 432 points.

We chose PP02 to use in our demonstration because this is one of OpenAI Gym’s recommended algorithms<sup>2</sup> but our contributions in this paper are model-agnostic.

<sup>2</sup><https://blog.openai.com/openai-baselines-ppo/>

Many RL agents are stochastic in nature: during training, randomness is used to balance exploration and exploitation, while during deployment, a learned policy may need to break ties. This results in agents whose policies are stochastic. Therefore, for each experimental condition, we run 30 trials of each for the test conditions and present the aggregated results.

For each requirement, we ran many trials and constructed a software test that exited under conditions of both success and failure. Each trial was allocated the equivalent of four minutes of human gameplay time for the agent to reach a success or failure state. We gave the agent a single life in order to achieve its objective. Dying or failing to pass the test within the time span was considered a failure for all hypotheses. We allowed the agent only a single level of breakout, ending when all the bricks disappeared, rather than see if it can clear both levels for each test.

**R1: Brick Elimination.** To win Breakout, an agent must eliminate every brick. A plausible state for such an agent might be after a “death” or lost ball, where there is one brick remaining. We might expect an agent that understands the physics of Breakout to remove the last brick quickly.

Figures 2a and 2d depict the reciprocal of the median number of steps taken to eliminate a given brick, with yellow indicating very few steps (median of 20) and red indicating very many steps (median of 400). For example, the agent typically targeted the paddle so that the ball would hit straight up when it was launched from a start state.

We suspect that the large difference we see between brick elimination scenarios is that improvements in paddle aim are not evenly distributed spatially. The same temporal range is visible on both charts, (maximums and minimums have not shifted) but there are many more bricks that are eliminated quickly for the agent with more training steps.

We had expected to see, for example, symmetry across the bricks. However, there is no obvious pattern here, suggesting that the agents are not particularly good at aiming at final, lingering bricks. In the future, we hope to explore whether these latencies are related to the frequency of their occurrence as the last brick in training data. Some alternative hypotheses for explaining the agent’s behavior are that it has memorized a sequence of actions, or that it has only learned to survive and is not targeting the ball at all.

**R2: Start Angle Invariance.** Our next behavioral requirement comes from observation of the game. Over time, the ball bounces through many different angles, and it is plausible (though maybe not strictly possible in the original version based on fixed-point math) to see all possible angles of movement of the ball. We therefore expect an agent to be able to catch the ball “thrown” at any start angle and to resume play.

Since Breakout operates over a finite set of angles, we did not expect as much success as this test indicated. Modifying a start state to change the initial launch angle of the ball led to a single life of gameplay, and a total score. Over 30 trials per agent per angle, we are able to visualize the mean (dark gray area), max (light gray area), median (red), 25th (blue) and 75th (green) percentiles on a polar plot in Figures 2b and 2e.

Both plots display failure to achieve any score with horizontal ball angles: since Breakout has no gravity, balls simply bounce horizontally forever, never hitting any bricks or threatening the paddle.

The agents under test were remarkably resilient to starting angles. While displaying a large amount of variance,<sup>3</sup> the maximum score achieved from each angle was much higher, suggesting that an agent can be successful even with balls traveling at angles it may never have observed in training. The agent trained for 50 million steps was much more robust – out of the 30 trials, there was at least one trial that completed the full level for each starting angle.

Looking closer at the mean, we can see that the agent trained for  $5e7$  steps had more difficulty with vertical angles. When we observed this behavior, the agent would sometimes keep the ball aligned perfectly in the center of the board, hitting it precisely in the center of the paddle, over and over, and therefore failing to make progress. This is a fascinating behavior that is entirely unlike the kind of behavior we would expect from human players.

---

<sup>3</sup>Although outside the scope of this work, there is a recent growing interest in the role of variability during evaluation in machine learning, as well as RL specifically [6, 5, 11]. We suspect the higher variance observed during evaluation of **R2** is an instance of this phenomenon.




Figure 2: Test Data Visualizations: Color of bricks in Figures 2a, 2c, 2d, and 2f indicate the median number of steps required to clear the particular brick in that test: Bright yellow represents fewer steps and dark red indicate many more steps. In Figures 2b and 2e, the black lines indicate the starting angles seen during training, the light gray area the maximum score achieved from this starting angle, and the dark gray area represents the mean score achieved across trials.

**R3: Tunnel Exploitation.** The highest reward in Breakout comes from “digging” a tunnel through the wall of bricks and then bouncing the ball through the hole and onto the ceiling. When the ball is trapped above the bricks, it will result in many bricks being removed very quickly, and a higher near-term reward than usual. Mnih et al., speculated that a high-performing agent was learning this behavior [15].

One way to test whether an agent “knows” to exploit a tunnel is to give it a board with a nearly built tunnel, save for a single brick, and test whether the agent can aim at that single brick within a given amount of time. Just as in the test presented for **R1**, we can generate a test for every brick.

In Fig. 2c and 2f, the value for each brick is the reciprocal of the median number of time steps before that brick was removed. The values range from 17 timesteps (bright-yellow) to 400 timesteps (dark red).

In order to satisfy our requirement, we would expect that an agent could exploit a tunnel quickly anywhere on the board, however we find that agents hit the ball to predictable locations regardless of the board configuration. We feel comfortable rejecting the hypothesis behind our requirement. We can see as well that specific agents learn to hit the ball to specific locations – the PPO2 model trained for 10 million steps prefers to send the ball to the right side of the screen, and the agent trained for 50 million steps prefers to hit the ball to the center column.

## 5 Conclusions

We present TOYBOX, a reimplementation of popular Atari benchmarks that allow for acceptance testing, dynamic analysis during training, and test-set generation. We present three requirement studies over deep RL models trained on Breakout and make evidence-based claims about whether particular agents satisfy our high-level behavioral requirements.

## Acknowledgements

This material is based upon work supported by the United States Air Force under Contract No, FA8750-17-C-0120. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

## References

- [1] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. A Brief Survey of Deep Reinforcement Learning. *arXiv preprint arXiv:1708.05866*, 2017.
- [2] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research*, 47:253–279, jun 2013.
- [3] Glen Berseth, Xue Bin Peng, and Michiel van de Panne. Terrain RL Simulator. *arXiv preprint arXiv:1804.06424*, 2018.
- [4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym, 2016.
- [5] Kaleigh Clary, Emma Tosch, John Foley, and David Jensen. Let’s Play Again: Variability of Deep Reinforcement Learning Agents in Atari Environments. In *NeurIPS 2018 Workshop on Critiquing and Correcting Trends in Machine Learning*, 2018.
- [6] Daniel Cohen, Scott M. Jordan, and W. Bruce Croft. Distributed evaluations: Ending neural point metrics. In *SIGIR; LND4IR*, 2018.
- [7] Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. OpenAI Baselines. <https://github.com/openai/baselines>, 2017.
- [8] Don Hodges. Why do Pinky and Inky have different behaviors when Pac-Man is facing up? [http://donhodges.com/pacman\\_pinky\\_explanation.htm](http://donhodges.com/pacman_pinky_explanation.htm), 2015.
- [9] Jemin Hwangbo, Inkyu Sa, Roland Siegwart, and Marco Hutter. Control of a Quadrotor with Reinforcement Learning. *IEEE Robotics and Automation Letters*, 2(4):2096–2103, 2017.
- [10] Alex Irpan. Deep reinforcement learning doesn’t work yet. <https://www.alexirpan.com/2018/02/14/r1-hard.html>, 2018.
- [11] Scott M. Jordan, Daniel Cohen, and Phillip S. Thomas. Using Cumulative Distribution Based Performance Analysis to Benchmark Models. In *NeurIPS 2018 Workshop on Critiquing and Correcting Trends in Machine Learning*, 2018.
- [12] Simon M Lucas. Game AI Research with Fast Planet Wars Variants. *arXiv preprint arXiv:1806.08544*, 2018.
- [13] Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew J. Hausknecht, and Michael Bowling. Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents. *CoRR*, abs/1709.06009, 2017.
- [14] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with Deep Reinforcement Learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level Control through Deep Reinforcement Learning. *Nature*, 518:529 EP –, 02 2015.
- [16] Maithra Raghu, Alex Irpan, Jacob Andreas, Robert Kleinberg, Quoc V Le, and Jon Kleinberg. Can Deep Reinforcement Learning solve Erdos-Selfridge-Spencer Games? *arXiv preprint arXiv:1711.02301*, 2017.
- [17] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [18] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*. MIT press, 1998.

- [19] Gabriel Synnaeve, Nantas Nardelli, Alex Auvolat, Soumith Chintala, Timothée Lacroix, Zeming Lin, Florian Richoux, and Nicolas Usunier. Torchcraft: a Library for Machine Learning Research on Real-time Strategy Games. *arXiv preprint arXiv:1611.00625*, 2016.
- [20] Yuandong Tian, Qucheng Gong, Wenling Shang, Yuxin Wu, and C. Lawrence Zitnick. ELF: An Extensive, Lightweight and Flexible Research Platform for Real-time Strategy Games. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2659–2669. Curran Associates, Inc., 2017.
- [21] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A Physics Engine for Model-based Control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 5026–5033. IEEE, 2012.
- [22] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. Starcraft II: A New Challenge for Reinforcement Learning. *arXiv preprint arXiv:1708.04782*, 2017.
- [23] Shimon Whiteson, Brian Tanner, Matthew E Taylor, and Peter Stone. Protecting against Evaluation Overfitting in Empirical Reinforcement Learning. In *Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), 2011 IEEE Symposium on*, pages 120–127. IEEE, 2011.