

# An Atari Model Zoo for Analyzing, Visualizing, and Comparing Deep Reinforcement Learning Agents

Felipe Petroski Such<sup>1</sup>, Vashisht Madhavan<sup>1</sup>, Rosanne Liu<sup>1</sup>, Rui Wang<sup>1</sup>  
 Pablo Samuel Castro<sup>2</sup>, Yulun Li<sup>1</sup>, Jiale Zhi<sup>1</sup>, Ludwig Schubert<sup>3</sup>  
 Marc G. Bellemare<sup>2</sup>, Jeff Clune<sup>1</sup> and Joel Lehman<sup>1\*</sup>

<sup>1</sup>Uber AI Labs

<sup>2</sup>Google Brain

<sup>3</sup>OpenAI

## Abstract

Much human and computational effort has aimed to improve how deep reinforcement learning (DRL) algorithms perform on benchmarks such as the Atari Learning Environment. Comparatively less effort has focused on understanding what has been learned by such methods, and investigating and comparing the representations learned by different families of DRL algorithms. Sources of friction include the onerous computational requirements, and general logistical and architectural complications for running DRL algorithms at scale. We lessen this friction, by (1) training several algorithms at scale and releasing trained models, (2) integrating with a previous DRL model release, and (3) releasing code that makes it easy for anyone to load, visualize, and analyze such models. This paper introduces the Atari Zoo framework, which contains models trained across benchmark Atari games, in an easy-to-use format, as well as code that implements common modes of analysis and connects such models to a popular neural network visualization library. Further, to demonstrate the potential of this dataset and software package, we show initial quantitative and qualitative comparisons between the performance and representations of several DRL algorithms, highlighting interesting and previously unknown distinctions between them.

## 1 Introduction

Since its introduction the Atari Learning Environment (ALE; [Bellemare *et al.*, 2013]) has been an important reinforcement learning (RL) testbed. It enables easily evaluating algorithms on over 50 emulated Atari games spanning diverse gameplay styles, providing a window on such algorithms’ generality. Indeed, surprisingly strong results in ALE with deep neural networks (DNNs), published in Nature [Mnih *et al.*, 2015], greatly contributed to the current popularity of deep reinforcement learning (DRL).

Like other machine learning benchmarks, much effort aims to quantitatively improve state-of-the-art (SOTA) scores. As

the DRL community grows, a paper pushing SOTA is likely to attract significant interest and accumulate citations. While improving performance is important, it is equally important to understand what DRL algorithms learn, how they process and represent information, and what are their properties, strengths, and weaknesses. These questions cannot be answered through simple quantitative measurements of performance across the ALE suite of games.

Compared to pushing SOTA, much less work has focused on understanding, interpreting, and visualizing *products* of DRL; in particular, little research compares DRL algorithms across dimensions other than performance. This paper thus aims to alleviate the considerable *friction* for those looking to rigorously understand the qualitative behavior of DRL agents. Three main sources of such friction are: (1) the significant computational resources required to run DRL at scale, (2) the logistical tedium of plumbing the products of different DRL algorithms into a common interface, and (3) the wasted effort in re-implementing standard analysis pipelines (like t-SNE embeddings of the state space [Mnih *et al.*, 2015], or activation maximization for visualizing what neurons in a model represent [Erhan *et al.*, 2009; Olah *et al.*, 2018; Nguyen *et al.*, 2017; Simonyan *et al.*, 2013; Yosinski *et al.*, 2015; Mahendran and Vedaldi, 2016]). To address these frictions, this paper introduces the *Atari Zoo*, a release of trained models spanning major families of DRL algorithms, and an accompanying open-source software package<sup>1</sup> that enables their easy analysis, comparison, and visualization (and similar analysis of future models). In particular, this package enables easily downloading particular frozen models of interest from the zoo on-demand, further evaluating them in their training environment or modified environments, generating visualizations of their neural activity, exploring compressed visual representations of their behavior, and creating synthetic input patterns that reveal what particular neurons most respond to.

To demonstrate the promise of this model zoo and software, this paper presents an initial analysis of the products of seven DRL algorithms spanning policy gradient, value-based, and evolutionary methods<sup>2</sup>: A2C (policy-gradient; [Mnih

<sup>1</sup><https://github.com/uber-research/atari-model-zoo>

<sup>2</sup>While evolutionary algorithms are excluded from some definitions of RL, their inclusion in the zoo can help investigate what dis-

\*Corresponding author: joel.lehman@uber.com

*et al.*, 2016]), IMPALA (policy-gradient; [Espeholt *et al.*, 2018]), DQN (value-based; [Mnih *et al.*, 2015]), Rainbow (value-based; [Hessel *et al.*, 2017]), Ape-X (value-based; [Horgan *et al.*, 2018]), ES (evolutionary; [Salimans *et al.*, 2017]), and Deep GA (evolutionary; [Such *et al.*, 2017]). The analysis illuminates differences in learned policies across methods that are independent of raw score performance, highlighting the benefit of going beyond simple quantitative measures and of having a unifying software framework that enables analyses with multiple, different, complementary techniques and applying them across many RL algorithms.

## 2 Background

### 2.1 Visualizing Deep Networks

One line of DNN research focuses on visualizing the internal dynamics of a DNN [Yosinski *et al.*, 2015] or examines what particular neurons detect or respond to [Erhan *et al.*, 2009; Zeiler and Fergus, 2014; Olah *et al.*, 2018; Nguyen *et al.*, 2017; Simonyan *et al.*, 2013; Yosinski *et al.*, 2015; Mahendran and Vedaldi, 2016]. The hope is to gain more insight into how DNNs are representing information, motivated both to enable more interpretable decisions from these models [Olah *et al.*, 2018] and to illuminate previously unknown properties about DNNs [Yosinski *et al.*, 2015]. For example, through live visualization of all activations in a vision network responding to different images, [Yosinski *et al.*, 2015] highlighted that representations were often surprisingly local (as opposed to distributed), e.g. one convolutional filter proved to be a reliable face detector. One practical value of such insights is that they can catalyze future research. The Atari Zoo enables animations in the spirit of [Yosinski *et al.*, 2015] that show an agent’s activations as it interacts with a game, and also enables creating synthetic inputs via activation maximization [Erhan *et al.*, 2009; Zeiler and Fergus, 2014; Olah *et al.*, 2018; Nguyen *et al.*, 2017; Simonyan *et al.*, 2013; Yosinski *et al.*, 2015; Mahendran and Vedaldi, 2016], specifically by connecting DRL agents to the Lucid visualization package<sup>3</sup>.

### 2.2 Understanding Deep RL

While much more visualization and understanding work has been done for vision models than for DRL, a few papers directly focus on understanding DRL agents [Greydanus *et al.*, 2017; Zahavy *et al.*, 2016], and many others feature some analysis of DRL agent behavior (often in the form of t-SNE diagrams of the state space; see [Mnih *et al.*, 2015]). One approach to understanding DRL agents is to investigate the learned features of models [Greydanus *et al.*, 2017; Zahavy *et al.*, 2016]. For example, [Zahavy *et al.*, 2016] visualize what pixels are most important to an agent’s decision by using gradients of decisions with respect to pixels. Another approach is to modify the DNN architecture or training procedure such that a trained model will have more interpretable features [Anasamy and Sycara, 2018]. For example, [Anasamy and Sycara, 2018] augment a model with an atten-

tinguishes such black-box optimization from more traditional RL.

<sup>3</sup><http://https://github.com/tensorflow/lucid>

tion mechanism and a reconstruction loss, hoping to produce interpretable explanations as a result.

The software package released here is in the spirit of the first paradigm. It facilitates understanding the most commonly applied architectures instead of changing them, although it is designed also to accommodate importing in new vision-based DRL models, and could thus be also used to analyze agents explicitly engineered to be interpretable. In particular, the package enables re-exploring many past DRL analysis techniques at scale, and across algorithms, which were previously applied only for one algorithm and across only a handful of hand-selected games.

### 2.3 Model Zoos

A useful mechanism for reducing friction for analyzing and building upon models is the idea of a *model zoo*, i.e. a repository of pre-trained models that can easily be further investigated, fine-tuned, and/or compared (e.g. by looking at how their high-level representations differ). For example, the Caffe website includes a model zoo with many popular vision models, as do Tensorflow, Keras, and PyTorch. The idea is that training large-scale vision networks (e.g. on the ImageNet dataset) can take weeks with powerful GPUs, and that there is little reason to constantly reduplicate the effort of training. Pre-trained word-embedding models are often released with similar motivation, e.g. for Word2Vec or GloVe. However, such a practice is much less common in the space of DRL; one reason is that so far, unlike with vision models and word-embedding models, there are few other downstream tasks from which Atari DRL agents provide obvious value. But, if the goal is to better understand these models and algorithm, both to improve them and to use them safely, then there is value in their release.

The recent Dopamine reproducible DRL package [Castro *et al.*, 2018] released trained ALE models; it includes final checkpoints of models trained by several DQN variants. However, in general it is non-trivial to extract TensorFlow models from their original context for visualization purposes, and to compare agent behavior across DRL algorithms in the same software framework (e.g. due to slight differences in image preprocessing), or to explore dynamics that take place *over* learning, i.e. from *intermediate* checkpoints. To remedy this, for this paper’s accompanying software release, the Dopamine checkpoints were distilled into frozen models that can be easily loaded into the Atari Zoo framework; and for algorithms trained specifically for the Atari Zoo, we distill intermediate checkpoints in addition to final ones.

## 3 Generating the Zoo

The approach is to run several validated implementations of DRL algorithms and to collect and standardize the models and results, such that they can then be easily used for downstream analysis and synthesis. There are many algorithms, implementations of them, and different ways that they could be run (e.g. different hyperparameters, architectures, input representations, stopping criteria, etc.). These choices influence the kind of post-hoc analysis that is possible. For example, Rainbow most often outperforms DQN, and if only

final models are released, it is impossible to explore scientific questions where it is important to *control* for performance.

We thus adopted the high level principles that the Atari Zoo should hold as many elements of architecture and experimental design constant across algorithms (e.g. DNN structure, input representation), should enable as many types of downstream analysis as possible (e.g. by releasing checkpoints across training time), and should make reasonable allowances for the particularities of each algorithm (e.g. ensuring hyperparameters are well-fit to the algorithm, and allowing for differences in how policies are encoded or sampled from). The next paragraphs describe specific design choices.

### 3.1 Frozen Model Selection Criteria

To enable the platform to facilitate a variety of explorations, we release multiple frozen models for each run, according to different criteria that may be useful to control for when comparing learned policies. The idea is that depending on the desired analysis, controlling for samples, or for wall-clock, or for performance (i.e. comparing similarly-performing policies) will be more appropriate. In particular, in addition to releasing the final model for each run, additional are models taken over training time (at one, two, four, six, and ten hours); over game frame samples (400 million, and 1 billion frames); over scores (if an algorithm reaches human level performance); and also a model before any training, to enable analysis of how weights change from their random initialization. The hope is that these frozen models will cover a wide spectrum of possible use cases.

### 3.2 Algorithm Choice

One important choice for the Atari Zoo is which DRL algorithms to run and include. The main families of DRL algorithms that have been applied to the ALE are policy gradient methods like A2C [Mnih *et al.*, 2016], value-based methods like DQN [Mnih *et al.*, 2015], and black-box optimization methods like ES [Salimans *et al.*, 2017] and Deep GA [Such *et al.*, 2017]. Based on representativeness and available trusted implementations, the particular algorithms chosen to train included two policy gradient algorithms (A2C [Mnih *et al.*, 2016] and IMPALA [Espeholt *et al.*, 2018]), two evolutionary algorithms (ES [Salimans *et al.*, 2017] and Deep GA [Such *et al.*, 2017]), and one value-function based algorithm (a high-performing DQN variant, Ape-X; [Horgan *et al.*, 2018]). Additionally, models are also imported from the Dopamine release [Castro *et al.*, 2018], which include DQN [Mnih *et al.*, 2015] and a sophisticated variant of it called Rainbow [Hessel *et al.*, 2017]. Note that from the Dopamine models, only final models are currently available. Hyperparameters and training details for all algorithms are available in supplemental material<sup>4</sup> section S3. We hope to include models from additional algorithms in future releases.

### 3.3 DNN Architecture and Input Representation

All algorithms are run with the DNN architecture from [Mnih *et al.*, 2015], which consists of three convolutional layers

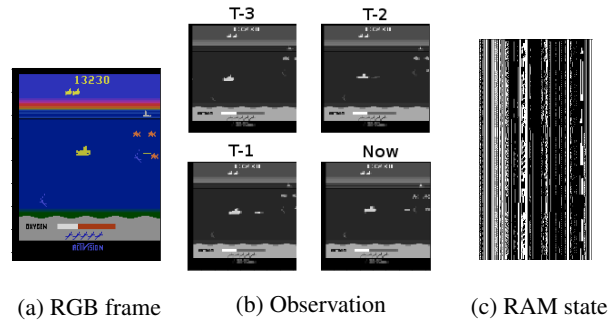


Figure 1: Input and RAM Representation. (a) One RGB frame of emulated Atari gameplay is shown, which is (b) preprocessed and concatenated with previous frames before being fed as an observation into the DNN agent. A compressed representation of a 2000-step ALE simulation is shown in (c), i.e. the 1024-bit RAM state (horizontal axis) unfurled over frames (vertical axis).

(with filter size 8x8, 4x4, and 3x3, followed by a fully-connected layer). For most of the explored algorithms, the fully-connected layer connects to an output layer with one neuron per valid action in the underlying Atari game. However, A2C and IMPALA have an additional output that approximates the state value function; Ape-X’s architecture features dueling DQN [Wang *et al.*, 2015], which has two separate fully-connected streams; and Rainbow’s architecture includes C51 [Bellemare *et al.*, 2017], which uses many outputs to approximate the distribution of expected Q-values.

Atari frames are 210x160 color images (see figure 1a); the canonical DRL representation is a tensor consisting of the four most recent observation frames, grayscaled and down-sampled to 84x84 (figure 1b). By including some previous frames, the aim is to make the game more fully-observable, to boost performance of the feed-forward architectures that are currently most common in Atari research (although recurrent architectures offer possible improvements [Mnih *et al.*, 2016; Espeholt *et al.*, 2018]). One useful Atari representation that is applied in post-training analysis in this paper, is the Atari RAM state, which is only 1024 bits long but encompasses the true underlying state (figure 1c).

### 3.4 Data Collection

All algorithms are run across 55 Atari games, for at least three independent random weight initializations. Regular checkpoints were taken during training; after training, the checkpoints that best fit each of the desired criteria (e.g. 400 million frames or human-level performance) were frozen and included in the zoo. The advantage of this post-hoc culling is that additional criteria can be added in the future, e.g. if Atari Zoo users introduce a new use case, because the original checkpoints are archived. Log files were stored and converted into a common format that are also released with the models, to aid future performance curve comparisons for other researchers. Each frozen model was run post-hoc in ALE for 2500 timesteps to generate cached behavior of policies in their training environment, which includes the raw game frames, the processed four-frame observations, RAM states, and high-level representations (e.g. neural representations at hidden layers). As a result, it is possible to do meaningful

<sup>4</sup>Supplemental material PDF: <http://bit.ly/2GJJstT>

analysis without ever running the models themselves.

## 4 Quantitative Analysis

The open-source software package released with the acceptance of this work provides an interface to the Atari Zoo dataset, and implements several common modes of analysis. Models can be downloaded with a single line of code; and other single-line invocations interface directly with ALE and return the behavioral outcome of executing a model’s policy, or create movies of agents superimposed with neural activation, or access convolutional weight tensors. In this section, we demonstrate analyses the Atari Zoo software can facilitate, and highlight some of its built-in features. For many of the analyses below, for computational simplicity we study results in a representative subset of 13 ALE games used by prior research [Such *et al.*, 2017], which we refer to here as the *analysis subset* of games.

### 4.1 Convolutional Filter Analysis

While understanding a DNN only by examining its weights is challenging, weights directly connected to the input can often be interpreted. For example, from visualizing the weights of the first convolutional layer in a vision model, Gabor-like edge detection filters are nearly always present. An interesting question is if Gabor-like filters also arise when DRL algorithms are trained from pixel input (as is done here). In visualizing filters across games and DRL algorithms, edge-detector-like features sometimes arise in the gradient-based methods, but they are seemingly never as crisp as in vision models; this may be because ALE lacks the visual complexity of natural images. In contrast, the filters in the evolutionary models are less regular. Representative examples across games and algorithms are shown in supplemental figure S1.

Learned filters commonly are tiled similarly across time (i.e. across the four DNN input frames), with past frames having lower-intensity weights. One explanation is that reward gradients are more strongly influenced by present observations. To explore this systematically, across games and algorithms we examined the absolute magnitude of filter weights connected to the present frame versus the past. In contrast to the gradient-based methods the evolutionary methods show no discernable preference across time (supplemental figure S2), again suggesting that their learning differs qualitatively from the gradient-based methods. Interestingly, a rigorous information-theoretic approximation of memory usage is explored by [Dann *et al.*, 2016] in the context of DQN; our measure well-correlates with theirs despite the relative simplicity of exploring only filter weight strength (supplemental section S1.1).

### 4.2 Robustness to Observation Noise

An important property is how agents perform in slightly out-of-distribution (OOD) situations; ideally they would not catastrophically fail in the face of nominal change. While it is difficult to freely alter the ALE game dynamics (without learning how to program in 6502 assembly code), it is possible to systematically distort observations. Here we explore one simple OOD change to observations by adding increasingly severe noise to the observations input to DNN-based

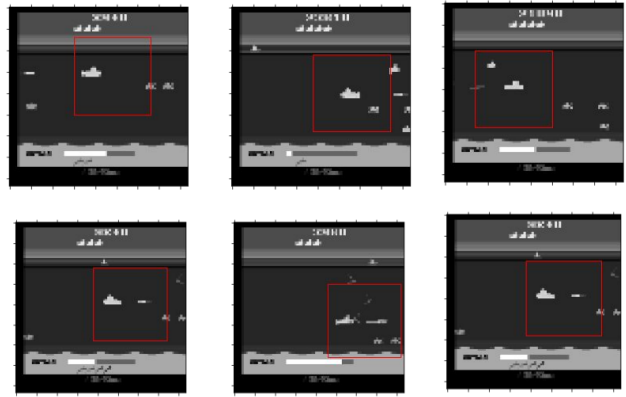


Figure 2: A sub-detecting neuron in Seaquest. Each image represents an observation from an Ape-X agent playing Seaquest. The red square indicates which image patch highly-activated the sub-detecting neuron on the third convolutional layer of the DNN. Having multiple tools (such as this image patch finder, or the activation movies which identified this neuron of interest) enables more easily triangulating and verifying hypotheses about the internals of a DRL agent’s neural network.

agents, and observe how their evaluated game score degrades. The motivation is to discover whether some learning algorithms are learning more robust policies than others. The results show that with some caveats, methods with a direct representation of the policy appear more robust to observation noise (supplemental figure S4). A similar study conducted for robustness to parameter noise (supplemental section S1.2) tentatively suggests that actor-critic methods are more robust to such noise.

### 4.3 Distinctiveness of Learned Policies

To explore the distinctive signature of solutions discovered by different DRL algorithms, we train image classifiers to identify the generating DRL algorithm given states sampled from independent runs of each algorithm (details are in supplemental section S1.3). Supplemental figure S6 shows the confusion matrix for Seaquest, wherein a cluster of policy search methods (A2C, ES, and GA) have the most inter-class confusion, reflecting (as confirmed in later sections) that these algorithms tend to converge to the same sub-optimal behavior in this game; results are qualitatively similar when tabulated across the analysis subset of games (supplemental figure S7).

## 5 Visualization

We next highlight the Atari Zoo’s capabilities to quickly and systematically visualize policies, which broadly can be divided into three categories: Direct policy visualization, dimensionality reduction, and neuron activation maximization.

### 5.1 Animations to Inspect Policy and Activations

To quickly survey the solutions being learned, our software generates grids of videos, where one grid axis spans different DRL algorithms, and the other axis covers independent runs

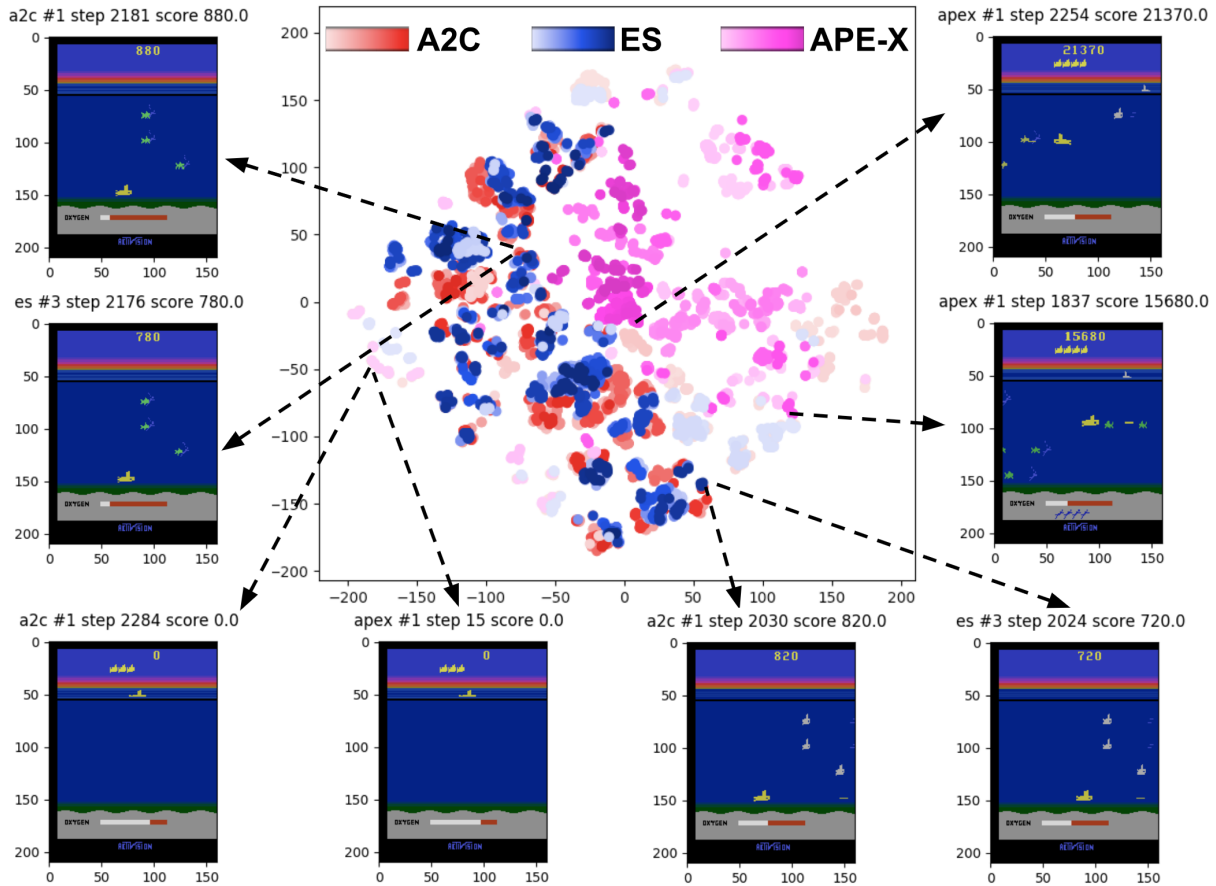


Figure 3: Multiple runs of algorithms and sharing the same RAM-space embedding in Seaquest. This plot shows one ALE evaluation per model for A2C, ES, and Ape-X, visualized in the same underlying RAM t-SNE embedding. Each dot represents a separate frame from each agent, colored by score (darker color indicates higher score). The plot highlights that in this game, A2C and ES visit similar distributions of states (corresponding to the same sub-optimal behavior), while Ape-X visits a distinct part of the state-space, i.e. matching what could manually be distilled from watching the policy movies shown in supplemental figure S9. The interface allows clicking on points to observe the corresponding RGB frame, and for toggling different runs of different algorithms for visualization.

of the algorithm. Such videos can highlight when different algorithms are converging to the same local optimum (e.g. supplemental figure S9 shows a situation where this is the case for A2C, ES, and the GA; video: <http://bit.ly/2Xpd5kO>).

To enable investigating the internal workings of the DNN, our software generates movies that display activations of all neurons alongside animated frames of the agent acting in game. This approach is inspired by the deep visualization toolbox [Yosinski *et al.*, 2015], but put into a DRL context. Supplemental figure S10 shows how this tool can lead to recognizing the functionality of particular high-level features (video: <http://bit.ly/2tFHiCU>); in particular, it helped to identify a submarine detecting neuron on the third convolution layer of an Ape-X agent. Note that for ES and GA, no such specialized neuron was found; activations seemed qualitatively more distributed for those methods.

### 5.2 Image Patches that Maximally Excite Filters

One automated technique for uncovering the functionality of a particular convolutional filter is to find which image patches evoke from it the highest magnitude activations. Given a

trained DRL agent and a target convolution filter to analyze, observations from the agent interacting with its ALE training environment are input to the agent’s DNN, and resulting maps of activations from the filter of interest are stored. These maps are sorted by the single maximum activation within them, and the geometric location within the map of that maximum activation is recorded. Then, for each of these top-most activations, the specific image patch from the observation that generated it is identified and displayed, by taking the receptive field of the filter into account (i.e. modulated by both the stride and size of the convolutional layers). As a sanity check, we validate that the neuron identified in the previous section does indeed maximally fire for submarines (figure 2).

### 5.3 Dimensionality Reduction

Dimensionality reduction provides another view on agent behavior; often DRL research includes t-SNE plots of agent DNN representations that summarize behavior in the domain [Mnih *et al.*, 2015]. Our software includes such an implementation (supplemental figure S12).



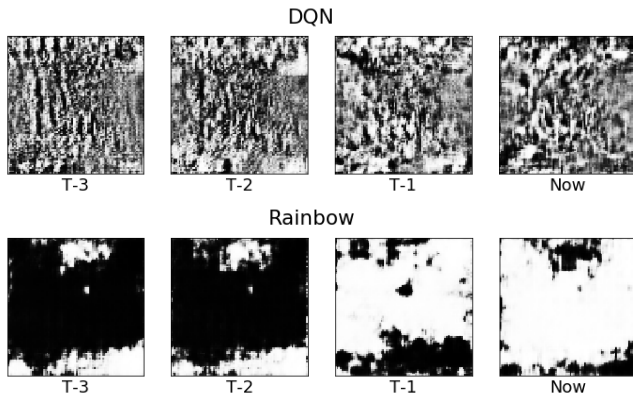


Figure 4: Synthesized inputs for output layer neurons in Seaquest. For a representative run of Rainbow and DQN, inputs are shown optimized to maximize the activation of the first neuron in the output layer of a Seaquest network. Because Rainbow includes C51, its image is in effect optimized to maximize the probability of a low-reward scenario; this neuron appears to be learning interpretable features such as submarine location and the seabed. When maximizing (or minimizing) DQN Q-value outputs (one example shown on left), this qualitative outcome of interpretability was not observed.

However, such an approach relies on embedding the high-level representation of *one* agent; it is unclear how to apply it to create an embedding appropriate for comparisons of different independent runs of the same algorithm, or runs from different DRL algorithms. As an initial approach, we implement an embedding based on the Atari RAM representation (which is the same across algorithms and runs, but distinct between games). Like the grid view of agent behaviors and the state-distinguishing classifier, this t-SNE tool provides high-level information from which to compare runs of or between different algorithms (figure 3); details of this approach are provided in supplemental section S2.1.

### 5.4 Synthesizing Inputs to Understand Neurons

While the previous sections explore DNN activations in the context of an agent’s training environment, another approach is to optimize synthetic input images that stimulate particular DNN neurons. Variations on this approach have yielded striking results in vision models [Nguyen *et al.*, 2017; Olah *et al.*, 2018; Simonyan *et al.*, 2013]; the hope is that these techniques could yield an additional view on DRL agents’ neural representations. To enable this analysis, we leverage the Lucid visualization library; in particular, we create wrapper classes that enable easy integration of Atari Zoo models into Lucid, and release Jupyter notebooks that generate synthetic inputs for different DRL models.

We now present a series of synthetic inputs generated by the Lucid library across a handful of games that highlight the potential of these kinds of visualizations for DRL understanding (further details of the technique used are described in supplemental section S2.2. We first explore the kinds of features learned across depth. Supplemental figure S13 supports what was learned by visualizing the first-layer filter weights for value-based networks (section 4.1; i.e. showing that first convolution layers in the value-based networks appear to be

learning edge-detector features). The activation videos of section 5.1 and the patch-based approach of section 5.2 help to provide grounding, showing that in the context of the game, some first-layer filters detect the edges of the screen, in effect to serve as location anchors, while others encode concepts like blinking objects (see figure S11). Supplemental figure S14 explores visualizing later-layer convolution filters, and figure 4 show inputs synthesized to maximize output neurons, which sometimes yields interpretable features.

Such visualizations can also reveal that critical features are being attended to (figure 5 and supplemental figure S15). Overall, these visualizations demonstrate the potential of this kind of technique, and we believe that many useful further insights may result from a more systematic application and investigation of this and many of the other interesting visualization techniques implemented by Lucid, which can now easily be applied to Atari Zoo models. Also promising would be to further explore regularization to constrain the space of synthetic inputs, e.g. a generative model of Atari frames in the spirit of [Nguyen *et al.*, 2017] or similar works.

## 6 Discussion and Conclusions

There are many follow-up extensions that the initial explorations of the zoo raise. One natural extension is to include more DRL algorithms (e.g. TRPO or PPO [Schulman *et al.*, 2017]). Beyond algorithms, there are many alternate architectures that might have interesting effects on representation and decision-making, for example recurrent architectures, or architectures that exploit attention. Also intriguing is examining the effect of the incentive driving search: Do auxiliary or substitute objectives qualitatively change DRL representations, e.g. as in UNREAL [Jaderberg *et al.*, 2016], curiosity-driven exploration [Pathak *et al.*, 2017], or novelty search [Conti *et al.*, 2017]? How do the representations and features of meta-learning agents such as MAML [Finn *et al.*, 2017] change as they learn a new task? Finally, there are other analysis tools that could be implemented, which might illuminate other interesting properties of DRL algorithms and learned representation, e.g. the image perturbation analysis of [Greydanus *et al.*, 2017] or a variety of sophisticated neuron visualization techniques [Nguyen *et al.*, 2017]. We welcome community contributions for these algorithms, models, architectures, incentives, and tools.

While the main motivation for the zoo was to reduce friction for research into understanding and visualizing the behavior of DRL agents, it can also serve as a platform for other research questions. For example, having a zoo of agents trained on individual games, for different amounts of data, also would reduce friction for exploring transfer learning within Atari, i.e. whether experience learned on one game can quickly benefit on another game. Also, by providing a huge library of cached rollouts for agents across algorithms, the zoo may be interesting in the context of learning from demonstrations, or for creating generative models of games. In conclusion, we look forward to seeing how this dataset will be used by the community at large.

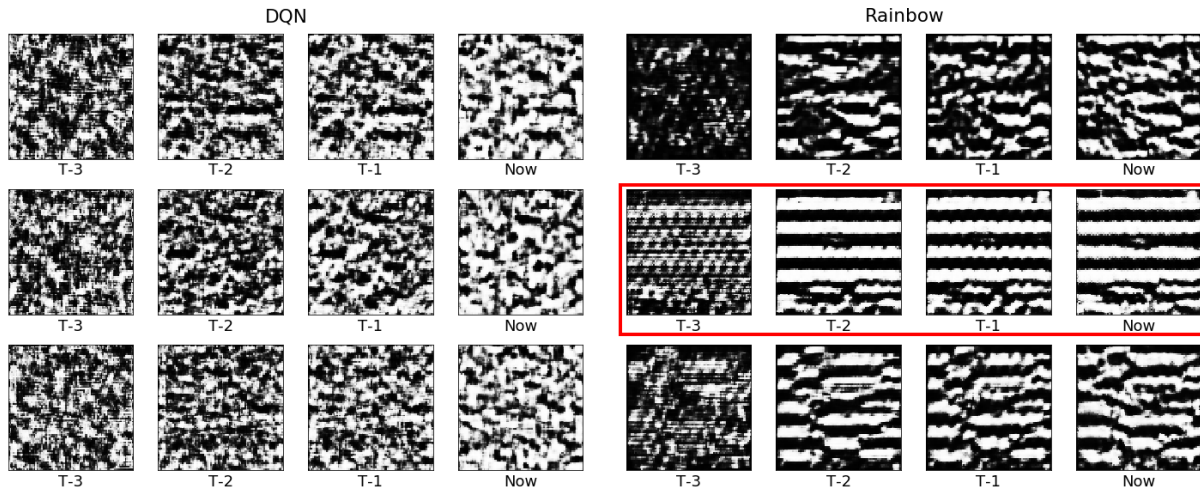


Figure 5: Synthesized inputs for fully-connected layer neurons in Freeway. Inputs synthesized to maximize activations of the first three neurons in the first fully connected layer are shown for a representative DQN and Rainbow DNN. One of the Rainbow neurons (in red rectangle) appears to be capturing lane features.

### Acknowledgements

We thank the members of Uber AI Labs, in particular, Jason Yosinski for helpful discussions; additionally, Chris Olah provided useful input during initial stages of the project.

### References

[Anasamy and Sycara, 2018] Raghuram Mandyam Anasamy and Katia Sycara. Towards better interpretability in deep q-networks. *arXiv preprint arXiv:1809.05630*, 2018.

[Bellemare *et al.*, 2013] Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.

[Bellemare *et al.*, 2017] Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. *arXiv preprint arXiv:1707.06887*, 2017.

[Castro *et al.*, 2018] Pablo Samuel Castro, Subhodeep Moitra, Carles Gelada, Saurabh Kumar, and Marc G Bellemare. Dopamine: A research framework for deep reinforcement learning. *arXiv preprint arXiv:1812.06110*, 2018.

[Conti *et al.*, 2017] Edoardo Conti, Vashisht Madhavan, Felipe Petroski Such, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. *arXiv preprint arXiv:1712.06560*, 2017.

[Dann *et al.*, 2016] Christoph Dann, Katja Hofmann, and Sebastian Nowozin. Memory lens: How much memory does an agent use? *arXiv preprint arXiv:1611.06928*, 2016.

[Erhan *et al.*, 2009] Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent. Visualizing higher-layer

features of a deep network. *University of Montreal*, 1341(3):1, 2009.

[Espeholt *et al.*, 2018] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.

[Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.

[Greydanus *et al.*, 2017] Sam Greydanus, Anurag Koul, Jonathan Dodge, and Alan Fern. Visualizing and understanding atari agents. *arXiv preprint arXiv:1711.00138*, 2017.

[Hessel *et al.*, 2017] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*, 2017.

[Horgan *et al.*, 2018] Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado Van Hasselt, and David Silver. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.

[Jaderberg *et al.*, 2016] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.

[Mahendran and Vedaldi, 2016] Aravindh Mahendran and Andrea Vedaldi. Visualizing deep convolutional neural networks using natural pre-images. *International Journal of Computer Vision*, 120(3):233–255, 2016.

- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Belle-mare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [Mnih *et al.*, 2016] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.
- [Nguyen *et al.*, 2017] Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. In *CVPR*, volume 2, page 7, 2017.
- [Olah *et al.*, 2018] Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander Mordvintsev. The building blocks of interpretability. *Distill*, 3(3):e10, 2018.
- [Pathak *et al.*, 2017] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, volume 2017, 2017.
- [Salimans *et al.*, 2017] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Simonyan *et al.*, 2013] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [Such *et al.*, 2017] Felipe Petroski Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv preprint arXiv:1712.06567*, 2017.
- [Wang *et al.*, 2015] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
- [Yosinski *et al.*, 2015] Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015.
- [Zahavy *et al.*, 2016] Tom Zahavy, Nir Ben-Zrihem, and Shie Mannor. Graying the black box: Understanding dqns. In *International Conference on Machine Learning*, pages 1899–1908, 2016.
- [Zeiler and Fergus, 2014] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833, 2014.