

# Towards Better Interpretability in Deep Q-Networks

**Raghuram Mandyam Annasamy**  
Carnegie Mellon University  
rannasam@cs.cmu.edu

**Katia Sycara**  
Carnegie Mellon University  
katia@cs.cmu.edu

## Abstract

Deep reinforcement learning techniques have demonstrated superior performance in a wide variety of environments. As improvements in training algorithms continue at a brisk pace, theoretical or empirical studies on understanding what these networks seem to learn, are far behind. In this paper we propose an interpretable neural network architecture for Q-learning which provides a global explanation of the model’s behavior using key-value memories, attention and reconstructible embeddings. With a directed exploration strategy, our model can reach training rewards comparable to the state-of-the-art deep Q-learning models. However, results suggest that the features extracted by the neural network are extremely shallow and subsequent testing using out-of-sample examples shows that the agent can easily overfit to trajectories seen during training.

## Introduction

The last few years have witnessed a rapid growth of research and interest in the domain of deep Reinforcement Learning (RL) due to the significant progress in solving RL problems (Arulkumaran et al. 2017). Deep RL has been applied to a wide variety of disciplines ranging from game playing, robotics, systems to natural language processing and even biological data (Silver et al. 2017; Mnih et al. 2015; Levine et al. 2016; Kraska et al. 2018; Williams, Asadi, and Zweig 2017; Choi et al. 2017). However, most applications treat neural networks as a black-box and the problem of understanding and interpreting deep learning models remains a hard problem. This is even more understudied in the context of deep reinforcement learning and only recently has started to receive attention. Commonly used visualization methods for deep learning such as saliency maps and t-SNE plots of embeddings have been applied to deep RL models (Greydanus et al. 2017; Zahavy, Ben-Zrihem, and Mannor 2016; Mnih et al. 2015). However, there are a few questions over the reliability of saliency methods including, as an example, sensitivity to simple transformations of the input (Kindermans et al. 2017). The problem of generalization and memorization with deep RL models is also important. Recent findings suggest that deep RL agents can easily memorize large amounts of training data with drastically varying test performance

and are vulnerable to adversarial attacks (Zhang et al. 2018; Zhang, Ballas, and Pineau 2018; Huang et al. 2017).

In this paper, we propose a neural network architecture for Q-learning using key-value stores, attention and constrained embeddings, that is easier to study than the traditional deep Q-network architectures. This is inspired by some of the recent work on Neural Episodic Control (NEC) (Pritzel et al. 2017) and distributional perspectives on RL (Bellemare, Dabney, and Munos 2017). We call this model i-DQN for Interpretable DQN and study latent representations learned by the model on standard Atari environments from Open AI gym (Brockman et al. 2016). Most current work around interpretability in deep learning is based on local explanations i.e. explaining network predictions for specific input examples (Lipton 2016). For example, saliency maps can highlight important regions of the input that influence the output of the neural network. In contrast, global explanations attempt to understand the mapping learned by a neural network regardless of the input. We achieve this by constraining the latent space to be reconstructible and inverting embeddings of representative elements in the latent space (keys). This helps us understand aspects of the input space (images) that are captured in the latent space across inputs. Our visualizations suggest that the features extracted by the convolutional layers are extremely shallow and can easily overfit to trajectories seen during training. This is in line with the results of (Zhang et al. 2018) and (Zhang, Ballas, and Pineau 2018). Although our main focus is to understand learned models, it is important that the models we analyze perform well on the task at hand. To this end, we show our model achieves training rewards comparable to Q-learning models like Distributional DQN (Bellemare, Dabney, and Munos 2017). Our contribution in this work is threefold:

- We explore a different neural network architecture with key-value stores, constrained embeddings and an explicit soft-assignment step that separates representation learning and Q-value learning (state aggregation).
- We show that such a model can improve interpretability in terms of visualizations of the learned keys (cluster), attention maps and saliency maps. Our method attempts to provide a global explanation of the model’s behavior instead of explaining specific input examples (local explanations). We also develop a few examples to test the

generalization behavior.

- We show that the model’s uncertainty can be used to drive exploration that reaches reasonably high rewards with reduced sample complexity (training examples) on some of the Atari environments.

## Related Work

Many attempts have been made to tackle the problem of interpretability with deep learning, largely in the supervised learning case. (Zhang and Zhu 2018) carry out an in-depth survey on interpretability with Convolutional Neural Networks (CNNs). Our approach to visualizing embeddings is in principle similar to the work of (Dosovitskiy and Brox 2016) on inverting visual representations. They train a neural network with deconvolution layers using HOG, SIFT and AlexNet embeddings as input and their corresponding real images as ground truth (the sole purpose of this network being visualization). Saliency maps are another popular type of method that generate local explanations which generally use gradient-like information to identify salient parts of the image. The different ways of computing saliency maps are covered exhaustively in (Zhang and Zhu 2018). Few of these have been applied in the context of deep reinforcement learning. (Zahavy, Ben-Zrihem, and Mannor 2016) use the Jacobian of the network to compute saliency maps on a Q-value network. Perturbation based saliency maps using a continuous mask across the image and also using object segmentation based masks have been studied in the context of deep-RL (Greydanus et al. 2017; Iyer et al. 2018; Li, Sycara, and Iyer 2017). In contrast to these approaches, our method is based on a global view of the network. Given a particular action and expected returns, we invert the corresponding key to try and understand visual aspects being captured by the embedding regardless of the input state. More recently, (Verma et al. 2018) introduce a new method that finds interpretable programs that can best explain the policy learned by a neural network- these programs can also be treated as global explanations for the policy networks.

Architecturally, our network is similar to the network first proposed by (Pritzel et al. 2017). The authors describe their motivation as speeding up the learning process using a semi-tabular representation with Q-value calculations similar to the tabular Q-learning case. This is to avoid the inherent slowness of gradient descent and reward propagation. Their model learns to attend over a subset of states that are similar to the current state by tracking all the states recently seen (up-to half-million states) using a k-d tree. However, their method does not have any notion of clustering or fixed Q-values. Our proposed method is also similar to Bellemare, Dabney, and Munos’s work on categorical/distributional DQN. The difference is that in our model the cluster embeddings (keys) for different Q-values are accessible freely (for analysis and visualization) because of the explicit soft-assignment step, whereas it is almost impossible to find such representations while having fully-connected layers like in (Bellemare, Dabney, and Munos 2017). Although we do not employ any iterative procedure (like refining keys; we train fully using backpropagation), works on combining deep embeddings with unsuper-

vised clustering methods (Xie, Girshick, and Farhadi 2016; Chang et al. 2017) (joint optimization/iterative refinement) have started to pick up pace and show better performance compared to traditional clustering methods.

Another important direction that is relevant to our work is that of generalizing behavior of neural networks in the reinforcement learning setting. (Henderson et al. 2017) discuss in detail about general problems of deep RL research and evaluation metrics used for reporting. (Zhang et al. 2018; Zhang, Ballas, and Pineau 2018) perform systematic experimental studies on various factors affecting generalization behavior such as diversity in training seeds and randomness in environment rewards. They conclude that deep RL models can easily overfit to random reward structures or when there is insufficient training diversity and careful evaluation techniques (such as isolated training and testing seeds) are needed.

## Proposed Method

We follow the usual RL setting and assume the environment can be modelled as a Markov Decision Process (MDP) represented by the 5-tuple  $(S, A, T, R, \gamma)$ , where  $S$  is the state space,  $A$  is the action space,  $T(s'|s, a)$  is the state transition probability function,  $R(s, a)$  is the reward function and  $\gamma \in [0, 1)$  is the discount factor. A policy  $\pi : S \rightarrow A$  maps every state to a distribution over actions. The value function  $V^\pi(s_t)$  is the expected discounted sum of rewards by following policy  $\pi$  from state  $s_t$  at time  $t$ ,  $V^\pi(s_t) = E[\sum_{i=0}^{\infty} \gamma^i r_{t+i}]$ . Similarly, the Q-value (action-value)  $Q^\pi(s_t, a)$  is the expected return starting from state  $s_t$ , taking action  $a$  and then following  $\pi$ . Q-value function can be recursively estimated using the Bellman equation  $Q^\pi(s_t, a) = E[r_t + \gamma \max_{a'} Q(s_{t+1}, a')]$  and  $\pi^*$  is the optimal policy which achieves the highest  $Q^\pi(s_t, a)$  over all policies  $\pi$ .

Similar to the traditional DQN architecture (Mnih et al. 2015), any state  $s_t$  (a concatenated set of input frames) is encoded using a series of convolutional layers each followed by a non-linearity and finally a fully-connected layer at the end  $h(s_t) = Conv(s_t)$ . This would usually be followed by some non-linearity and a fully-connected layer that outputs Q-values. Instead, we introduce a restricted key-value store over which the network learns to attend as shown in Figure 1.

Intuitively, the model is trying to learn two things. First, it learns a latent representation  $h(s_t)$  that captures important visual aspects of the input images. At the same time, the model also attempts to learn an association between embeddings of states  $h(s_t)$  and embeddings of keys in the key-value store. This would help in clustering the state (around the keys) based on the scale of expected returns (Q-values) from that state. We can think of the  $N$  different keys  $h^a$  (for a given action  $a$ ) weakly as the cluster centers for the corresponding Q-values, attention weights  $w^a(s_t)$  as a soft assignment between embeddings for current state  $h(s_t)$  and embeddings for different Q-values  $\{h_1^a, h_2^a, \dots, h_N^a\}$ . This explicit association step helps us in understanding the model in terms of attention maps and visualizations of the cluster centers (keys).

The key-value store is restricted in terms of size and values of the store. Each action  $a \in A$  has a fixed number of key-

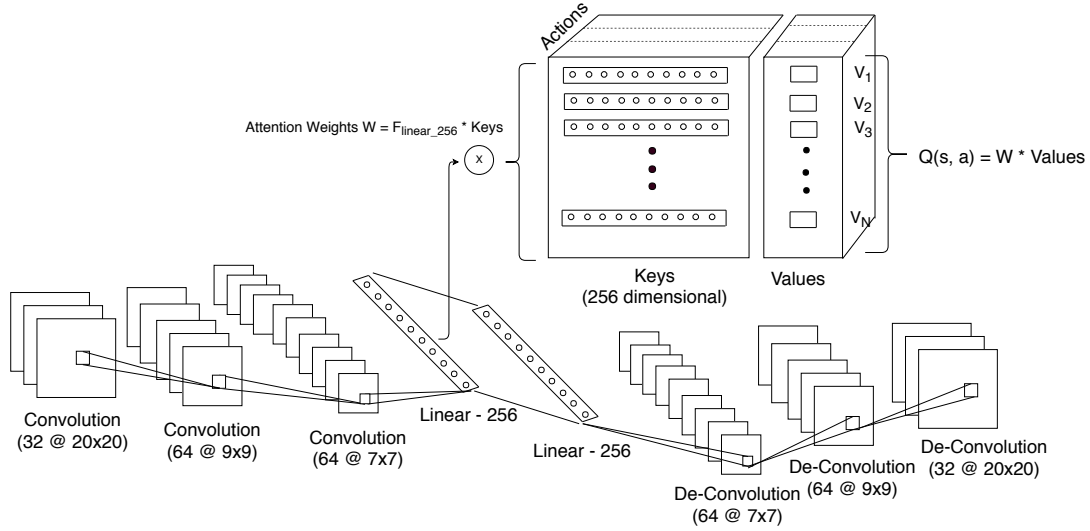


Figure 1: Model Architecture- Interpretable DQN (i-DQN)

value pairs (say  $N$ ) and the value associated with every key is also held constant.  $N$  values  $\{v_1, v_2, \dots, v_N\}$  are sampled uniformly at random from  $(V_{min}, V_{max})$  (usually  $(-25, 25)$ ) once and the same set of values are used for all the actions. All of the keys ( $N \times A$ ) are also initialized randomly. To compute attention, the embeddings  $h(s_t)$  for a state  $s_t$  are compared to all the keys  $\{h_1^a, h_2^a, \dots, h_N^a\}$  in the store for a particular action ( $a$ ) using a softmax over their dot products.

$$w(s_t)_i^a = \frac{\exp(h(s_t) \cdot h_i^a)}{\sum_j \exp(h(s_t) \cdot h_j^a)} \quad (1)$$

These attention weights over keys and their corresponding value terms are then used to calculate the Q-values.

$$Q(s_t, a) = \sum_i w_i^a(s_t) v_i$$

Now that we have Q-values, we can define the different losses that can be used to train the network,

- Bellman Error ( $\mathcal{L}_{bellman}$ ): The usual value function estimation error.

$$\mathcal{L}_{bellman}(\theta) = (Q(s_t, a, \theta) - Y_t)^2$$

where  $Y_t = R(s_t, a, s_{t+1}) + \gamma \max_a' Q(s_{t+1}, a', \theta)$

- Distributive Bellman Error ( $\mathcal{L}_{distrib}$ ): We force the distributive constraint on attention weights between current and next states similar to (Bellemare, Dabney, and Munos 2017) using values  $\{v_1, v_2, \dots, v_N\}$  as supports of the distribution. The distributive loss is defined as the KL divergence between  $\phi \mathcal{T} w(s_{t+1})^{a^*}$  and  $w(s_t)^a$  where  $\mathcal{T}$  is the distributional Bellman operator and  $\phi$  is the projection operator and  $a^*$  is best action at state  $s_{t+1}$  i.e.

$$a^* = \arg \max_a Q(s_{t+1}, a).$$

$$\mathcal{L}_{distrib}(\theta) = \mathcal{D}_{KL}(\phi \mathcal{T} w(s_{t+1})^{a^*}, w(s_t)^a) \quad (2)$$

$$= - \sum_i \phi \mathcal{T} w(s_{t+1})_i^{a^*} \cdot w(s_t)_i^a \quad (3)$$

Equation (3) is simply the cross entropy loss (assuming  $w(s_{t+1})^{a^*}$  to be constant with respect to  $\theta$ , similar to the assumption for  $Y_t$  in Bellman error).

- Reconstruction Error ( $\mathcal{L}_{reconstruct}$ ): We also constrain the embeddings  $h(s_t)$  for any state to be reconstructible. This is done by transforming  $h(s_t)$  using a fully-connected layer and then followed by a series of non-linearity and deconvolution layers.

$$h^{dec}(s_t) = W^{dec} h(s_t) \quad (4)$$

$$\hat{s}_t = Deconv(h^{dec}(s_t))$$

The mean squared error between reconstructed image  $\hat{s}_t$  and original image  $s_t$  is used,

$$\mathcal{L}_{reconstruct}(\theta) = \frac{1}{2} \|\hat{s}_t - s_t\|_2^2$$

- Diversity Error ( $\mathcal{L}_{diversity}$ ): The diversity error forces attention over different keys in a batch. This is important because training can collapse early with the network learning to focus on very few specific keys (because both the keys and attention weights are being learned together). We could use KL-divergence between the attention weights but (Lin et al. 2017) develop an elegant solution to this in their work.

$$\mathcal{L}_{diversity}(\theta) = \|(AA^T - I)\|^2$$

where  $A$  is a 2D matrix of size (batch size,  $N$ ) and each row of  $A$  is the attention weight vector  $w(s_t)^a$ . It drives

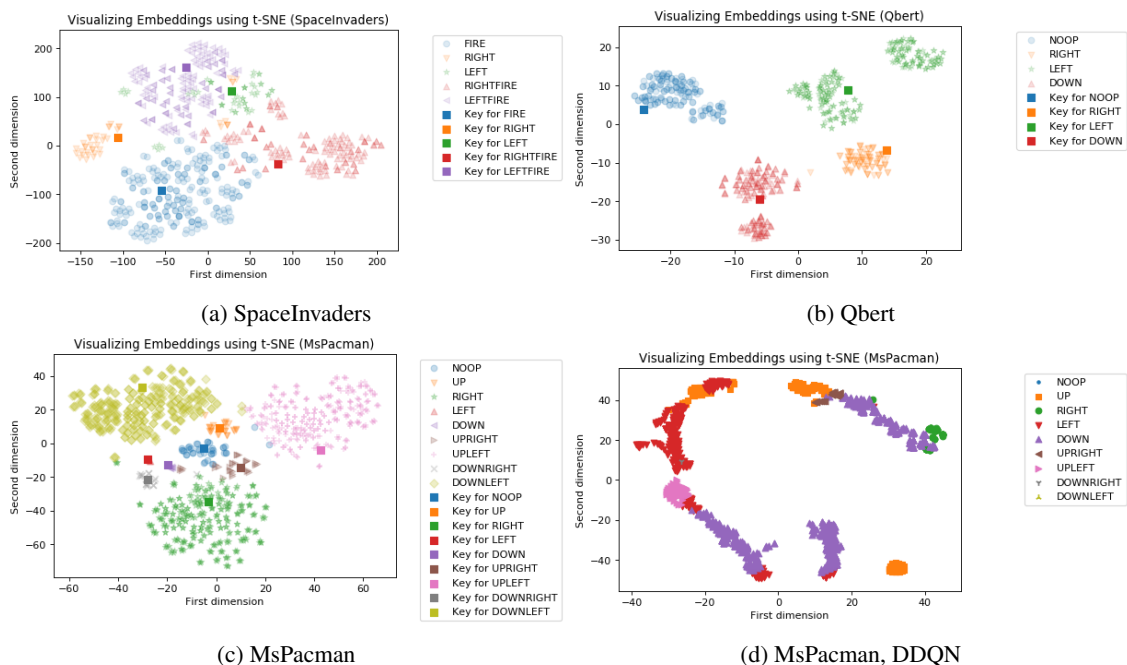


Figure 2: Visualizing keys, state embeddings using t-SNE: i-DQN, Q-value 25 (a)-(c); Double DQN(d)

$AA^T$  to be diagonal (no overlap between keys attended to within a batch) and  $l_2$  norm of  $w$  to be 1. Because of softmax, the  $l_1$  norm is also 1 and so ideally the attention must peak at exactly one key however in practice it spreads over as few keys as possible. Finally, the model is trained to minimize a weighted linear combination of all the four losses.

$$\mathcal{L}_{final}(\theta) = \lambda_1 \mathcal{L}_{bellman}(\theta) + \lambda_2 \mathcal{L}_{distrib}(\theta) + \lambda_3 \mathcal{L}_{reconstruct}(\theta) + \lambda_4 \mathcal{L}_{diversity}(\theta)$$

## Experiments and Discussions

We report the performance of our model on eight Atari environments (Brockman et al. 2016)- Alien, Freeway, Frostbite, Gravitar, MsPacman, Qbert, SpaceInvaders, and Venture, in Table 1.<sup>1</sup> Using the taxonomy of Atari games from (Belle-mare et al. 2016) and (Ostrovski et al. 2017), seven of the eight environments tested (all except SpaceInvaders) are considered hard exploration problems. Additionally, three of them (Freeway, Gravitar and Venture) are hard to explore because of sparse rewards. Since our focus is on interpretability, we do not carry out an exhaustive performance comparison. We simply show that training rewards achieved by i-DQN model are comparable to some of the state-of-the-art models. This is important because we would like our deep-learning models to be interpretable but also remain competitive at the same time. We look at scores against other exploration baselines for Q-learning that do not involve explicit reward shaping/exploration bonuses- Bootstrap DQN (Osband et al. 2016), Noisy DQN (Fortunato et al. 2017) and Q-ensembles (Chen et al. 2018).

<sup>1</sup>Code available at <https://github.com/maraghuram/I-DQN>

## Directed exploration

We use the uncertainty in attention weights to drive exploration during training.  $U(s_t, a)$  is an approximate upper confidence on the Q-values. Similar to (Chen et al. 2018) we select the action maximizing a UCB style confidence interval,

$$Q(s_t, a) = \sum_i w_i^a(s_t) v_i$$


$$U(s_t, a) = \sqrt{Q(s_t, a)^2 - \sum_i w_i^a(s_t) v_i^2}$$

$$a_t = \arg \max_{a \in A} Q(s_t, a) + \lambda_{exp} U(s_t, a)$$

Table 1 compares i-DQN’s performance (with directed exploration) against a baseline Double DQN implementation (which uses epsilon-greedy exploration) at 10M frames. Double DQN (DDQN), Distributional DQN (Distrib. DQN), Bootstrap DQN and Noisy DQN agents are trained for up to 50M steps which translates to 200M frames (Hessel et al. 2017; Fortunato et al. 2017; Osband et al. 2016). The Q-ensemble agent using UCB-style exploration is trained for up to 40M frames (Chen et al. 2018). We see that on some of the games, our model reaches higher training rewards within 10M frames compared to Double DQN, Distributional DQN models. Also, our model is competitive with the final scores reported by other exploration baselines like Bootstrap DQN, Noisy DQN and Q-ensembles, and even performs better on some environments (5 out of 8 games). The training time for i-DQN is roughly 2x slower because of the multiple loss functions compared to our implementation of Double DQN.


Environment	10M frames		Reported Scores (final)				
	DDQN	i-DQN	DDQN	Distrib. DQN	Q-ensemble	Bootstrap DQN	Noisy DQN
Alien	1,533.45	2,380.72	3,747.7	<b>4,055.8</b>	2,817.6	2,436.6	2,394.90
Freeway	22.5	28.79	33.3	33.6	<b>33.96</b>	33.9	32
Frostbite	754.48	<b>3,968.45</b>	1,683.3	3,938.2	1,903.0	2,181.4	583.6
Gravitar	279.89	517.33	412	<b>681</b>	318	286.1	443.5
MsPacman	2,251.43	<b>6,132.21</b>	2,711.4	3,769.2	3,425.4	2,983.3	2,501.60
Qbert	10,226.93	<b>19,137.6</b>	15,088.5	16,956.0	14,198.25	15,092.7	15,276.30
Space Invaders	563.2	979.45	2,525.5	<b>6,869.1</b>	2,626.55	2,893	2,145.5
Venture	70.87	985.11	98	<b>1,107.0</b>	67	212.5	0

Table 1: Training scores (averaged over 100 episodes, 3 seeds). Scores for Double DQN , Distributional DQN and Noisy DQN are from (Hessel et al. 2017); Scores for Bootstrap-DQN are as reported in the original paper (Osband et al. 2016); Scores for UCB style exploration with Q-ensembles are from (Chen et al. 2018)



(a) Down (b) Downleft (c) Upleft (d) Right

Figure 3: MsPacman, Inverting keys for Q-value 25



(a) Right (b) R-Fire (c) Left (d) L-Fire (e) Fire

Figure 4: SpaceInvaders, Inverting keys for Q-value 25

### What do the keys represent?

The keys are latent embeddings (randomly initialized) that behave like cluster centers for the particular action-return pairs (latent space being  $\mathbb{R}^{256}$ ). Instead of training using unsupervised methods like K-means or mixture models, we use the neural network itself to find these points using gradient descent. For example, the key for action right; Q-value 25 (Figure 2c) is a cluster center that represents the latent embeddings for all states where the agent expects a return of 25 by selecting action right. These keys partition the latent space into well formed clusters as shown in Figure 2, suggesting that embeddings also contain action-specific information crucial for an RL agent. On the other hand, Figure 2d shows embeddings for DDQN which are not easily separable (similar to the visualizations in (Mnih et al. 2015)). Since we use simple dot-product based distance for attention, keys and state embeddings must lie in a similar space and this can be seen in the t-SNE visualization i.e. keys (square boxes) lie within the state embeddings (Figure 2). The fact that the keys lie close to their state embeddings is essential to interpretability because state embeddings satisfy reconstructability constraints.

### Inversion of keys

Although keys act like cluster centers for action-return pairs, it is difficult to interpret them in the latent space. By inverting keys, we attempt to find important aspects of input space (images) that influence the agent to choose particular action-return pair ( $Deconv(h_t^a)$ ). These are ‘global explanations’ because inverting keys is independent of the input. For example, in MsPacman, reconstructing keys for different actions (fixing return of 25) indicates yellow blobs at many different places for each action (Figure 3). We hypothesize that these correspond to the Pacman object itself and that the model memorizes its different positions to make its decision i.e. the yellow blobs in Figure 3d correspond to different locations of Pacman and for any input state where Pacman is in one of those positions, the agent selects action right expecting a return of 25. Figure 5 shows such examples where the agent’s action-return selection agrees with reconstructed key (red boxes indicate Pacman’s location). Similarly, in SpaceInvaders, the agent seems to be looking at specific combinations of shooter and alien ship positions that were observed during training (Figure 4).

The keys have never been observed by the deconvolution network during training and so the reconstructions depend upon its generalizability. Interestingly, reconstructions for action-return pairs that are seen more often tend to be less noisy with less artifacts. This can be observed in Figure 3 for Q-value 25 where actions Right, Downleft and Upleft nearly 65% of all actions taken by the agent. We also look at the effect of different reconstruction techniques keeping the action-return pair fixed (Figure 6). Variational autoencoder with  $\beta$  set to 0 yields sharper looking images but increasing  $\beta$  which is supposed to bring out disentanglement in the embeddings yields reconstructions with almost no objects. Dense VAE with  $\beta = 0$  is a slightly deeper network similar to (Oh et al. 2015) and seems to reconstruct slightly clearer shapes of ghosts and pacman.

### Evaluating the reconstructions

To understand the effectiveness of these visualizations, we design a quantitative metric that measures the agreement between actions taken by the agent and the actions suggested using the reconstructed images. With a fully trained

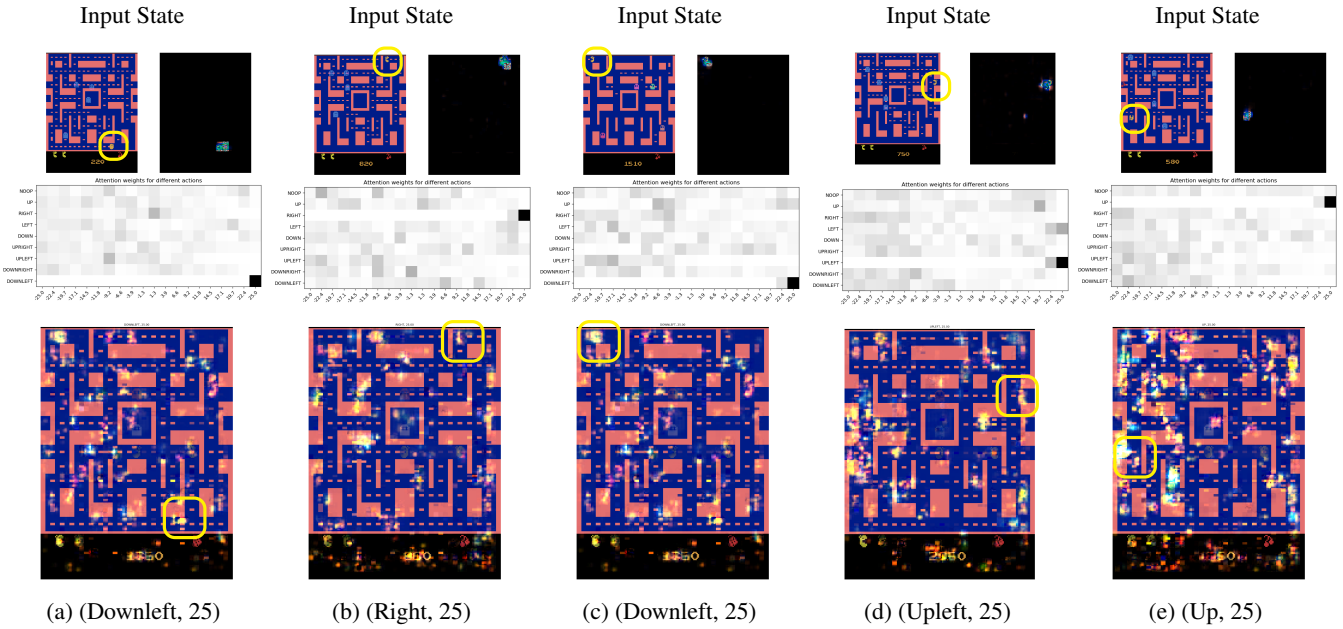


Figure 5: MsPacman, examples where agent’s decision agrees with the reconstructed image

Agreement	AE	VAE ( $\beta = 0$ )	VAE ( $\beta = 0.01$ )	Dense VAE ( $\beta = 0$ )
MsPacman (Color)	<b>30.76</b>	29.78	16.8	23.73
MsPacman (Gray, Rescaled)	19.87	18.14	10.97	14.56

Table 2: Evaluating visualizations: Agreement scores




Figure 6: Reconstruction: AE, VAE  $\beta = 0$ , VAE  $\beta = 0.01$ , Dense VAE

model, we reconstruct images  $\{s_1^a, s_2^a, \dots, s_N^a\}$  from the keys  $\{h_1^a, h_2^a, \dots, h_N^a\}$  for all actions  $a \in A$ . In every state  $s_t$ , we induce another distribution on the Q-values using the cosine similarity in the image space,

$$w'(s_t)_i^a = \text{Softmax}\left(\frac{s_t \cdot s_i^a}{\|s_t\|_2 \cdot \|s_i^a\|_2}\right)$$

similar to  $w(s_t)_i^a$  (which is also a distribution over Q-values but in the latent space). Using  $w'(s_t)_i^a$ , we can compute  $Q'(s_t, a)$  and  $U'(s_t, a)$  as before and select an action  $a'_t = \arg \max_{a \in A} Q'(s_t, a) + \lambda_{exp} U'(s_t, a)$ . Using  $a_t$  and  $a'_t$ , we

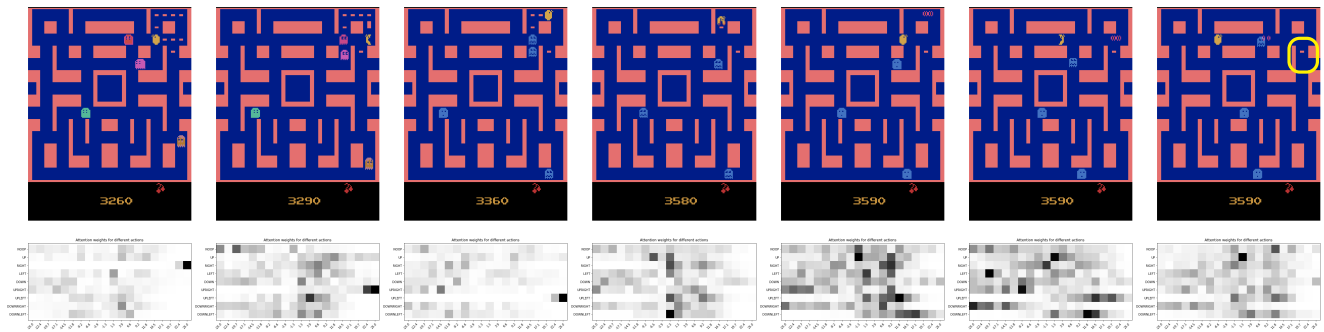
define our metric of agreeability as

$$\text{Agreement} = \frac{\mathbb{1}_{a_t = a'_t}}{\mathbb{1}_{a_t = a'_t} + \mathbb{1}_{a_t \neq a'_t}}$$


where  $\mathbb{1}$  is the indicator function. We measure this across multiple rollouts (5) using  $a_t$  and average them. In Table 2, we report Agreement as a percentage for different encoder-decoder models. Unfortunately, the best agreement between the actions selected using the distributions in the image space and latent space is around 31% for the unscaled color version of MsPacman. In MsPacman, the agent has 9 different actions and a random strategy would expect to have an Agreement of  $\sim 11\%$ . However, if the agreement scores were high (80-90%), that would suggest that the Q-network is indeed learning to memorize configurations of objects seen during training. One explanation for the gap is that reconstructions rely heavily on generalizability to unseen keys.

### Adversarial examples that show memorization


Looking at the visualizations and rollouts of a fully trained agent, we hand-craft a few out-of-sample environment states to examine the agent’s generalization behavior. For example, in MsPacman, since visualizations suggest that the agent may be memorizing pacman’s positions (also maybe ghosts and other objects), we simply add an extra pellet adjacent to a trajectory seen during training (Figure 7a). The agent does not clear the additional pellet and simply continues to execute actions performed during training (Figure 7b). Most importantly, the agent is extremely confident in taking actions initially (seen in attention maps Figure 7a) which suggest that the extra pellet was probably not even captured by the embeddings. Similarly, in case of SpaceInvaders, the agent has a strong bias towards shooting from the leftmost-end (seen in Figure 4). This helps in clearing the triangle



(a) Adversarial example



(b) Trajectory during training



(c) Adversarial example



(d) Trajectory during training

Figure 7: Adversarial examples for MsPacman (a)-(b) and SpaceInvaders (c)-(d)

like shape and moving to the next level (Figure 7d). However, when triangular positions of spaceships are inverted, the agent repeats the same strategy of trying to shoot from left and fails to clear ships (Figure 7c). These examples indicate that the features extracted by the convolutional channels seem to be shallow. The agent does not really model interactions between objects. For example in MsPacman, after observing 10M frames, it does not know general relationships between pacman and pellet or ghosts. Even if optimal Q-values were known, there is no incentive for the network to model these higher order dependencies when it can work with situational features extracted from finite training examples. (Zhang et al. 2018) also report similar results on simple mazes where an agent trained on insufficient environments tends to repeat training trajectories on unseen mazes.

### Sensitivity to hyperparameters

I-DQN’s objective function introduces four hyperparameters for weighting the different loss components ( $\lambda_1$ : for bellman error,  $\lambda_2$ : distributional error,  $\lambda_3$ : reconstruction error and  $\lambda_4$ : diversity error). The diversity error forces attention over multiple keys (examples for  $\lambda_4 = 0$  and  $\lambda_4 = 0.01$  are shown in the supplementary material). In general, we found the values  $\lambda_1 = 1.0$ ,  $\lambda_2 = 1.0$ ,  $\lambda_3 = 0.05$ ,  $\lambda_4 = 0.01$  to work well across games (detailed list of hyperparameters and their values is reported in the supplementary material). We ran experiments for different settings of  $\lambda_1$ ,  $\lambda_2$  and  $\lambda_3$  keeping  $\lambda_4 = 0.01$  constant on Ms Pacman (averaged over 3 trials). In general, increasing  $\lambda_3$  (coefficient on reconstruction error) to 0.5 and 5.0 yields visually better quality reconstructions but poorer scores- 3,245.3 and 3,013.1 respectively (drop by  $\sim 45\%$ ). Increasing  $\lambda_1 = \lambda_2 = 10.0$

also drops the score and yields poor reconstructions (sometimes without the objects of the game which loses interpretability) but converges quite quickly (5, 267.15, drop by  $\sim 12\%$ ). Out of  $\lambda_1$  (bellman loss) and  $\lambda_2$  (distributional loss),  $\lambda_2$  seems to play a more important in reaching higher scores compared to  $\lambda_1$  [ $\lambda_1 = 1.0, \lambda_2 = 10.0$  score: 4, 916.0 ;  $\lambda_1 = 10.0, \lambda_2 = 1.0$ , score: 4, 053.6]. So, the setting  $\lambda_1 = 1.0, \lambda_2 = 1.0, \lambda_3 = 0.05, \lambda_4 = 0.01$  seems to find the right balance between the q-value learning losses and regularizing losses. For the exploration factor, we tried a few different values  $\lambda_{\text{exp}} = \{0.1, 0.01, 0.001\}$  and it did not have a significant effect on the scores.

## Conclusion

In this paper, we propose an interpretable deep Q-network model (i-DQN) that can be studied using a variety of tools including the usual saliency maps, attention maps and reconstructions of key embeddings that attempt to provide global explanations of the model’s behavior. We also show that the uncertainty in soft cluster assignment can be used to drive exploration effectively and achieve high training rewards comparable to other models. Although the reconstructions do not explain the agent’s decisions perfectly, they provide a better insight into the kind of features extracted by convolutional layers. This can be used to design interesting adversarial examples with slight modifications to the state of the environment where the agent fails to adapt and instead repeats action sequences that were performed during training. This is the general problem of overfitting in machine learning but is more acute in the case of reinforcement learning because the process of collecting training examples depends largely on the agent’s biases (exploration). There are many interesting directions for future work. For example, we know that the reconstruction method largely affects the visualizations and other methods such as generative adversarial networks (GANs) (Goodfellow et al. 2014) can model latent spaces more smoothly and could generalize better to unseen embeddings. Another direction is to see if we can automatically detect the biases learned by the agent and design meaningful adversarial examples instead of manually crafting test cases.

## Acknowledgments

This work was funded by awards AFOSR FA9550-15-1-0442 and NSF IIS- 1724222.

## References

Arulkumaran, K.; Deisenroth, M. P.; Brundage, M.; and Bharath, A. A. 2017. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*.

Bellemare, M.; Srinivasan, S.; Ostrovski, G.; Schaul, T.; Sutton, D.; and Munos, R. 2016. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, 1471–1479.

Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A distributional perspective on reinforcement learning. *arXiv preprint arXiv:1707.06887*.

Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym.

Chang, J.; Wang, L.; Meng, G.; Xiang, S.; and Pan, C. 2017. Deep adaptive image clustering. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 5880–5888. IEEE.

Chen, R. Y.; Sidor, S.; Abbeel, P.; and Schulman, J. 2018. Ucb exploration via q-ensembles.

Choi, E.; Hewlett, D.; Uszkoreit, J.; Polosukhin, I.; Lacoste, A.; and Berant, J. 2017. Coarse-to-fine question answering for long documents. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 209–220.

Dosovitskiy, A., and Brox, T. 2016. Inverting visual representations with convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4829–4837.

Fortunato, M.; Azar, M. G.; Piot, B.; Menick, J.; Osband, I.; Graves, A.; Mnih, V.; Munos, R.; Hassabis, D.; Pietquin, O.; et al. 2017. Noisy networks for exploration. *arXiv preprint arXiv:1706.10295*.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.

Greydanus, S.; Koul, A.; Dodge, J.; and Fern, A. 2017. Visualizing and understanding atari agents. *arXiv preprint arXiv:1711.00138*.

Henderson, P.; Islam, R.; Bachman, P.; Pineau, J.; Precup, D.; and Meger, D. 2017. Deep reinforcement learning that matters. *arXiv preprint arXiv:1709.06560*.

Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2017. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*.

Huang, S.; Papernot, N.; Goodfellow, I.; Duan, Y.; and Abbeel, P. 2017. Adversarial attacks on neural network policies. *arXiv preprint arXiv:1702.02284*.

Iyer, R.; Li, Y.; Li, H.; Lewis, M.; Sundar, R.; and Sycara, K. 2018. Transparency and explanation in deep reinforcement learning neural networks.

Kindermans, P.-J.; Hooker, S.; Adebayo, J.; Alber, M.; Schütt, K. T.; Dähne, S.; Erhan, D.; and Kim, B. 2017. The (un) reliability of saliency methods. *arXiv preprint arXiv:1711.00867*.

Kraska, T.; Beutel, A.; Chi, E. H.; Dean, J.; and Polyzotis, N. 2018. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, 489–504. ACM.

Levine, S.; Finn, C.; Darrell, T.; and Abbeel, P. 2016. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research* 17(1):1334–1373.

Li, Y.; Sycara, K. P.; and Iyer, R. 2017. Object-sensitive deep reinforcement learning. In *GCAI 2017, 3rd Global Conference on Artificial Intelligence, Miami, FL, USA, 18-22 October 2017.*, 20–35.



Lin, Z.; Feng, M.; Santos, C. N. d.; Yu, M.; Xiang, B.; Zhou, B.; and Bengio, Y. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.

Lipton, Z. C. 2016. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*.

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.

Oh, J.; Guo, X.; Lee, H.; Lewis, R. L.; and Singh, S. 2015. Action-conditional video prediction using deep networks in atari games. In *Advances in Neural Information Processing Systems*, 2863–2871.

Osband, I.; Blundell, C.; Pritzel, A.; and Van Roy, B. 2016. Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, 4026–4034.

Ostrovski, G.; Bellemare, M. G.; Oord, A. v. d.; and Munos, R. 2017. Count-based exploration with neural density models. *arXiv preprint arXiv:1703.01310*.

Pritzel, A.; Uria, B.; Srinivasan, S.; Puigdomenech, A.; Vinyals, O.; Hassabis, D.; Wierstra, D.; and Blundell, C. 2017. Neural episodic control. *arXiv preprint arXiv:1703.01988*.

Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2017. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.

Verma, A.; Murali, V.; Singh, R.; Kohli, P.; and Chaudhuri, S. 2018. Programmatically interpretable reinforcement learning. *arXiv preprint arXiv:1804.02477*.

Williams, J. D.; Asadi, K.; and Zweig, G. 2017. Hybrid code networks: practical and efficient end-to-end dialog control with supervised and reinforcement learning. *arXiv preprint arXiv:1702.03274*.

Xie, J.; Girshick, R.; and Farhadi, A. 2016. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, 478–487.

Zahavy, T.; Ben-Zrihem, N.; and Mannor, S. 2016. Gray-ing the black box: Understanding dqns. In *International Conference on Machine Learning*, 1899–1908.

Zhang, Q.-s., and Zhu, S.-C. 2018. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering* 19(1):27–39.

Zhang, A.; Ballas, N.; and Pineau, J. 2018. A dissection of overfitting and generalization in continuous reinforcement learning. *arXiv preprint arXiv:1806.07937*.

Zhang, C.; Vinyals, O.; Munos, R.; and Bengio, S. 2018. A study on overfitting in deep reinforcement learning. *arXiv preprint arXiv:1804.06893*.