

Random Projection in Neural Episodic Control

Daichi Nishio

DNISHIO@CSL.EC.T.KANAZAWA-U.AC.JP

Satoshi Yamane

SYAMANE@IS.T.KANAZAWA-U.AC.JP

Kanazawa University, Ishikawa, Japan

Editors: Wee Sun Lee and Taiji Suzuki

Abstract

End-to-end deep reinforcement learning has enabled agents to learn with little preprocessing by humans. However, it is still difficult to learn stably and efficiently because the learning method usually uses a nonlinear function approximation. Neural Episodic Control (NEC), which has been proposed in order to improve sample efficiency, is able to learn stably by estimating action values using a non-parametric method. In this paper, we propose an architecture that incorporates random projection into NEC to train with more stability. In addition, we verify the effectiveness of our architecture by Atari’s five games. The main idea is to reduce the number of parameters that have to learn by replacing neural networks with random projection in order to reduce dimensions while keeping the learning end-to-end.

Keywords:

Neural Episodic Control, Random Projection, Deep Reinforcement Learning, Atari

1. Introduction

Recent advances in deep learning have allowed more complex function approximation and feature extraction. Reinforcement learning has also benefited from this, and Deep Q-Network (DQN) (Mnih et al. (2015)) has been proposed, which enables end-to-end learning from only image. However, there are still problems that deep reinforcement learning has to solve. One of them is poor sample efficiency. Neural Episodic Control (NEC) (Pritzel et al. (2017)) has been proposed to solve it. NEC has introduced a differentiable dictionary, called Differentiable Neural Dictionary (DND), into the neural network architecture. This makes it possible to learn action values from a feature of a state stably and to learn with a small number of learning steps. One of the reasons is that it has reduced the parameters which have to learn by changing parametric prediction to non-parametric one. Furthermore, NEC’s architecture allows end-to-end learning, although extraction of features from images is parametric and output action values from them is non-parametric.

In this research, we propose a method that incorporates random projection, which is one of the dimensionality reduction methods without training parameters, into a part of NEC’s network in order to learn more stably. We also verify its effectiveness in video game experiments.

2. Deep Reinforcement Learning

Reinforcement learning is an area of machine learning that an agent learns how to maximize the return of rewards obtained by interaction with the environment. The action value of a reinforcement learning agent taking the action a in the state s is defined by the following equation (1), where $G_t = \sum_t (\gamma^t r_t)$ is the sum of discounted rewards, and γ ($0 \leq \gamma < 1$) is the discount factor that represents how important future rewards are.

$$Q_\pi(s, a) = E_\pi [G_t | s, a] \quad (1)$$

The action value is called Q-value, and Q-learning (Watkins and Dayan (1992)) is a method to estimate the value. It is a bootstrap estimation using Bellman equation (Bellman (1952)) as in the equation (2).

$$Q(s, a) \leftarrow Q(s, a) + \alpha \left(r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right) \quad (2)$$

In value-based policy, the agent chooses an action that maximizes the estimated Q-value. However, it may not get better rewards for inexperienced states if it continues to choose the greedy action to maximize the value. In order to avoid this problem, there is a simple but powerful way called ε -greedy policy. It is written as the equation (3).

$$\pi(a|s) = \begin{cases} 1 - \varepsilon & (a = \operatorname{argmax}_a Q(s, a)) \\ \varepsilon & (\text{otherwise}) \end{cases} \quad (3)$$

Classically, Q-values are recorded in a Q-table, but in reality the state space is large and the action space may be continuous. Therefore, it is general to approximate an action value function.

Deep learning is also a way of function approximation. It is able to extract features of states and to estimate the function which it is difficult for other approximation methods to express. DQN has attracted much attention for end-to-end learning by extracting embeddings from Convolutional Neural Network (CNN) (Krizhevsky et al. (2012)) only based on images and outputting action values in a Fully Connected (FC) layer based on the embeddings.

However, Q-learning by nonlinear function approximation cannot generally guarantee the convergence, and it is difficult to learn stably. Therefore, DQN uses experience replay (Lin (1992)) which creates minibatches by taking experiences randomly, and has a network of the old parameters as the target value network. In particular, an efficient use of experience has been found to be important for fast and stable learning. Prioritized Experience Replay (PER) (Schaul et al. (2015)) has been proposed to weight experiences instead of conventional random sampling. The paper (Hessel et al. (2018a)) that has examined the importance and combination of various methods has also verified the effectiveness of PER, and it has shown that it has been very important to use experiences efficiently.

3. Neural Episodic Control

Neural Episodic Control (NEC) (Pritzel et al. (2017)) has a memory stored experience inside the architecture to enhance sample efficiency further. NEC is based on Model-Free

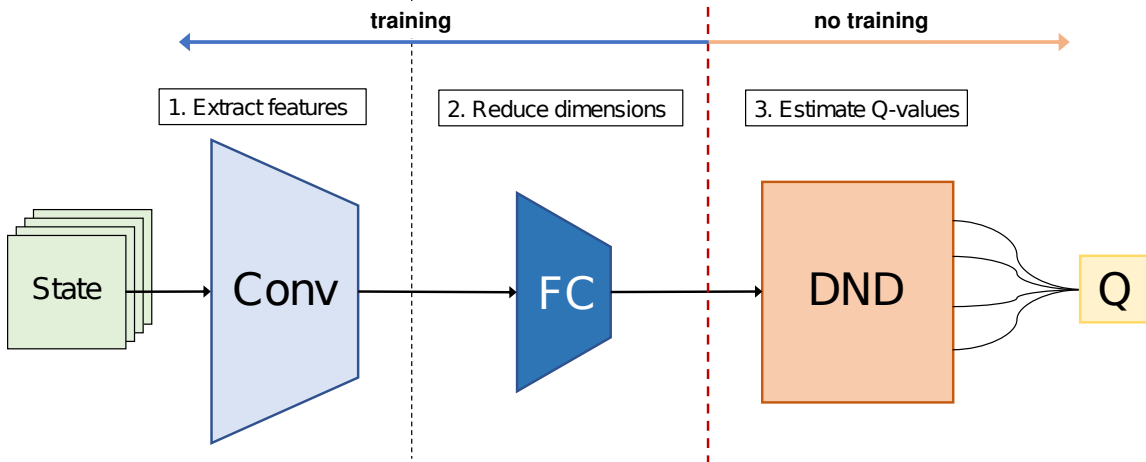


Figure 1: Neural Episodic Control

Episodic Control (MFEC) (Blundell et al. (2016)). The main idea of MFEC is to store many experiences in a Q-table, and it estimates action values by a non-parametric method for embeddings extracted by Variational AutoEncoder (VAE) (Kingma and Welling (2013)) or random projection. In contrast, NEC adopts this idea as a part of the architecture by Differentiable Neural Dictionary (DND), which is a dictionary that enables to update by gradients for a Q-table of embeddings. This makes it possible to learn end-to-end while keeping the Q-table inside the network.

Specifically, the architecture of NEC can be divided into the following three parts.

- (a) Obtain an embedding \mathbf{h} such as by convolutional layers.
- (b) Obtain an embedding \mathbf{h}' reduced dimensions of \mathbf{h} by fully connected layers.
- (c) Lookup to output action values from \mathbf{h}' using a k-nearest neighbor algorithm and a kernel function.

Here, we describe 2 and 3.

In 2, NEC reduces the embedding dimensions obtained in 1. There are two reasons for doing it, the first is to reduce the space complexity of DND. The second is the reduction of time complexity of the k-nearest neighbor algorithm in 3.

We explain an idea for non-parametric estimation of action values in 3. It is assumed that the $Q(\mathbf{h}', a)$ of each action a for the embedding \mathbf{h}'_i of a state s_i similar to the embedding \mathbf{h}' of a certain state s should be a similar value in many scenes.

When \mathbf{h}' obtained in 2 and the corresponding action a are input, p keys \mathbf{h}'_i resembling \mathbf{h}' among the keys existing in DND are searched by k-nearest neighbor algorithm using kd-trees (Bentley (1975)). Let Q_a be the weighted value v_i corresponding to the p keys.

$$w_i = \frac{k(\mathbf{h}', \mathbf{h}'_i)}{\sum_j k(\mathbf{h}', \mathbf{h}'_j)} \quad (4)$$

$$Q_a = \sum_{i=1}^p w_i v_i \quad (5)$$

The function k in the equation (4) is a kernel function. For example, it is written as the equation (6) by the inverse kernel function.

$$k(\mathbf{h}', \mathbf{h}'_i) = \frac{1}{\|\mathbf{h}' - \mathbf{h}'_i\|_2^2 + \delta} \quad (6)$$

Although δ is the parameter to prevent division by zero, we should make it little larger such as $\delta = 10^{-3}$ because each value of p neighbors is referred.

In this way, NEC allows stable learning by estimating action values non-parametrically. While it has been possible to learn faster and stabler than many algorithms including DQN, it has been found that long-term training makes it inferior to other methods. The cause is that it is necessary to store a large number of embedding-action pairs in DND, and the insufficient buffer size cannot estimate action values well.

NEC adopts multi-step Q-learning (Peng and Williams (1996)) as another technique for fast learning. One-step learning has the advantage that the variance of target value is low, but it also has the disadvantage that the propagation of rewards is slow. On the other hand, Monte Carlo Q-learning, which uses all the experiences of one episode, has the a rapid propagation of rewards, but unstable learning because of the high variance. Multi-step Q-learning has role as the trade-off the strengths and weaknesses of one-step Q-learning and Monte Carlo Q-learning.

In Rainbow (Hessel et al. (2018a)), it is better to set the value of this step number N to 3 or 5, but NEC sets N to 100 because of the stability of learning.

4. Random Projection

Random Projection (RP) is a linear projection with a random matrix and is used to reduce dimensions of high-dimensional data. As the properties of the projection matrix, it is necessary to consider the time to construct the matrix and the quality of embedding after dimensionality reduction. The difference between the main methods of random projection is as shown in Table 1.

From here, we describe Gaussian random projection (Hecht-Nielsen et al. (1994)) because it has the best **Embedding quality**. In this method, a d_x dimensional vector \mathbf{x} is multiplied by the Gaussian random matrix \mathbf{R} to convert it into a d_y dimensional vector \mathbf{y} .

$$\mathbf{y} = \mathbf{R}\mathbf{x} \quad (7)$$

If the elements of the random matrix \mathbf{R} are generated by random numbers in accordance with a Gaussian distribution (mean 0, variance $1/d_y$), the distance between the data is approximately maintained with high probability $1 - \mathcal{O}(e^{-d_x \varepsilon^2})$ when any n training data $\mathbf{x}^{(j)} (j = 1, \dots, n)$ are projected in the $d_{y=\mathcal{O}(\varepsilon^{-2} \log n)}$ dimensions [Johnson and Lindenstrauss (1984), Vempala (2004)] where ε is the distortion.

$$(1 - \varepsilon) \|\mathbf{x}_j - \mathbf{x}_i\|_2^2 \leq \|\mathbf{y}_j - \mathbf{y}_i\|_2^2 \leq (1 + \varepsilon) \|\mathbf{x}_j - \mathbf{x}_i\|_2^2 \quad (8)$$

Table 1: Comparison of random projection methods

RP method	Construction time	Projection time ^e	Embedding quality ^f
Gaussian	$\mathcal{O}(d_x d_y)$	$\mathcal{O}(n d_x d_y)$	$\mathcal{O}(\varepsilon^{-2} n)$
Achlioptas' ^a	$\mathcal{O}(d_x d_y)$	$\mathcal{O}(n d_x d_y)$	no proof
Li's ^b	$\mathcal{O}(\sqrt{d_x} d_y)$	$\mathcal{O}(n \sqrt{d_x} d_y)$	no proof
SRHT ^c	$\mathcal{O}(d_x d_y + d_x \log d_x)$	$\mathcal{O}(n d_x \log d_y)$	$\mathcal{O}(\varepsilon^{-2} (n + d_x) \log n)$
Count Sketch ^d	$\mathcal{O}(d_x)$	$\mathcal{O}(n d_x)$	$\mathcal{O}(\varepsilon^{-2} n^2)$

^a Achlioptas (2001) ^b Li et al. (2006) ^c Woodruff (2014) ^d Meng and Mahoney (2013)

^e The projection time is the case of dense inputs.

^f This quality is oblivious subspace embedding (OSE) lower bound.

Gaussian random projection is distinguished by its simplicity and high quality of embedding. The paper (Dasgupta (2013)) that actually experimented with random projection has shown that the relationship between the dimensions of the vector before projection and the dimensions of the vector after the projection is theoretically guaranteed. Moreover, it has been applied to the EM algorithm and the performance is improved. Another paper (Bingham and Mannila (2001)) applied to an image or text has reported to show good performance even if we have reduced dimensions under weaker conditions than the theoretically guaranteed the inequality (8). It has also shown that if the dimension is too small, the accuracy drops sharply.

Random projection has a good property that there is no restriction on the magnitude of each value of data. This means that the inequality (8) holds without normalization.

Also, random projection has good compatibility with the kd-trees because information other than Euclidean distance is redundant for the algorithm based on the closeness of the distance. In other words, we can use it for the data projected by random projection. In fact, RP-kd-Trees (Wu et al. (2011)), which uses random projection for dimension reduction, has been also proposed.

5. Related Work

There are several researches that have improved or combined NEC. A representative one is Ephemeral Value Adjustments (EVA) (Hansen et al. (2018)). It has improved the performance by combining a value buffer like NEC and other planning algorithms only in decision making. Also, there are combines parametric methods such as DQN with non-parametric methods such as NEC. Semiparametric Reinforcement Learning (SRL) (Jain and Lindsey (2018)) has proposed a method using action values that combine the values output by neural network and the values estimated by NEC architecture. NEC2DQN (Nishio and Yamane (2018)) has also made use of learning efficiency of NEC to assist to learn about a parametric network.

Deep reinforcement learning using random projection such as MFEC has been proposed. Episodic Memory Deep Q-Netwroks (EMDQN) (Lin et al. (2018)) has regularized action value with the values output using random projection. However, they differ from our

proposed algorithm in that they use random projection outside their neural networks, in other words, they are not end-to-end networks.

More recent research has proposed an architecture using random projection for deep neural networks (Wójcik (2018)). Although the networks train with embeddings projected by random projection, we do not use it as the input of neural networks, but utilize the relationship of the embedding distance.

6. Proposed Algorithm

As NEC has shown, non-parametric methods such as the k-nearest neighbor algorithm are good for increasing the learning speed of agents. One of the reasons is that there are no training parameters which have to learn. NEC reduces the dimensions of embeddings output by CNN with FC layers, calculates the Euclidean distance of nearby embeddings using k-nearest neighbor algorithm, then outputs the action value.

However, is it required to use fully connected layer for reducing dimensions? We think the layer must have the two important abilities.

- (a) It is able to keep the relationships of Euclidean distances for its inputs as possible.
- (b) It is differentiable in order to perform backpropagation.

It is important for the inverse kernel function to calculate the value is that the relationship of the distance is maintained as in the equation (6). In addition, the (b) is essential for end-to-end learning. Moreover, we would not like to train the layer because changing the parameters could cause unstable learning if it does not have to train.

Fortunately, random projection meets these conditions as stated later. Therefore, we perform stable learning by incorporating random projection which can keep the Euclidean distance relationship without training parameters. At the same time, we make the proposed architecture more general by end-to-end learning.

We propose NEC-RP, which incorporates random projection into the architecture of NEC. We show the algorithm in Algorithm 1.

6.1. Random Projection layers

Here we introduce Random Projection (RP) layer. It is a linear projection and differentiable as in the equation (7). The partial derivative of it is given by the equation (9).

$$\frac{\partial \mathbf{y}}{\partial \mathbf{x}} = \mathbf{R}^T \quad (9)$$

In addition, we can regard the equation (7) as in the following neural network's equation (10) by substituting the random matrix \mathbf{R} for \mathbf{W} and a zero vector for \mathbf{b} .

$$\mathbf{y} = \mathbf{W}\mathbf{x} + \mathbf{b} \quad (10)$$

In other words, we define our RP layer by Gaussian random projection as a layer of which the weight matrix is generated by Gaussian random (mean 0, variance $1/d_{\mathbf{y}}$) and the bias is a zero vector. Therefore, we regard our RP layer as a special case of the FC layer. We show the difference between these layers in Table 3.

Algorithm 1 NEC-RP

```

1: Initialize a replay memory  $D$  to capacity  $C_D$ .
2: Initialize DND memories  $M_a$  to capacity  $C_{M_a}$ .
3: Initialize an action value function  $Q$  with random weights.
4: Initialize the number of entire timesteps  $TS$  to zero.
5: Define the step  $CS$  to replace an RP layer with a FC layer. {if  $CS \leftarrow \infty$ , use only the
   RP layer forever.}
6: for each episode do
7:   for  $t = 1, 2, \dots, T$  do
8:     Receive an observation  $s_t$  from an environment.
9:     Convert  $s_t$  to an embedding  $h_t$ .
10:    if  $TS < CS$  then
11:      Reduce dimensions of  $h_t$  using an RP layer.
12:    else
13:      Reduce dimensions of  $h_t$  using a FC layer.
14:    end if
15:    Calculate  $Q(s_t, a)$  from  $M_a$ .
16:     $a_t \leftarrow \epsilon$ -greedy policy based on  $Q(s_t, a)$ .
17:    Take an action  $a_t$ , receive a reward  $r_t$ .
18:    Train on a random minibatch from  $D$ .
19:     $TS \leftarrow TS + 1$ 
20:  end for
21:  for  $t = 1, 2, \dots, T$  do
22:    Append  $(s_t, a_t, Q^{(N)}(s_t, a_t))$  to  $D$ .
23:    Append  $(h_t, Q^{(N)}(s_t, a_t))$  to  $M_{a_t}$ .
24:  end for
25: end for

```

Table 3: The differences between a FC layer and a RP layer

	FC	RP
Initializer (weight)	e.g. Gaussian (Mean=0, Variance=1)	e.g. Gaussian (Mean=0, Variance= $ units_{out} ^{-1}$)
Initializer (bias)	e.g. zeros	zeros
Parameters	update	fix

6.2. Neural Episodic Control with a Random Projection layer

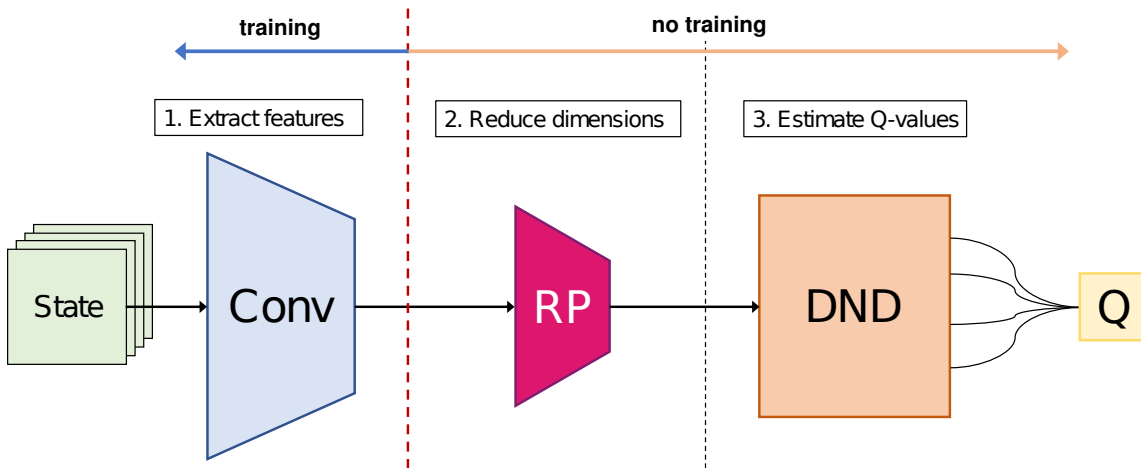


Figure 2: A random projection layer in Neural Episodic Control architecture

We show NEC-RP architecture in Fig. 2. Of the three parts of the NEC architecture mentioned in chapter 3, we use the RP layer in 2. In this way, the value of inverse kernel function is approximately maintained as in the inequality (11) from the inequality (8) with δ which is a constant that prevents division by zero.

$$\begin{aligned}
 (1 - \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 &\leq \|\mathbf{h}' - \mathbf{h}'_i\|_2^2 \leq (1 + \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 \\
 (1 - \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 + \delta &\leq \|\mathbf{h}' - \mathbf{h}'_i\|_2^2 + \delta \leq (1 + \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 + \delta \\
 \frac{1}{(1 + \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 + \delta} &\leq \frac{1}{\|\mathbf{h}' - \mathbf{h}'_i\|_2^2 + \delta} \leq \frac{1}{(1 - \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 + \delta} \\
 \frac{1}{(1 + \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 + \delta} &\leq k(\mathbf{h}', \mathbf{h}'_i) \leq \frac{1}{(1 - \varepsilon) \|\mathbf{h} - \mathbf{h}_i\|_2^2 + \delta} \quad (11)
 \end{aligned}$$

The value of the kernel function may be unstable if we use nonlinear function approximation methods such as deep neural network, but in NEC-RP it is guaranteed as in the inequality (11), hence it is possible to learn stably without training parameters.

6.3. Replacing a Random Projection layer with a Fully Connected layer

The RP layer does not learn its parameters as shown in Table 3. Although the methods which has no training parameters have the advantage of not requiring learning, they have the disadvantage that the performance is inferior to ones which have training parameters as the learning progresses. However, if we modify the parameters of the RP layer to learn, we can consider it as a FC layer because the RP layer is possible to learn from the middle of training. For example, the simplest method is to switch the RP layer to a FC layer according to the time step t and the hyperparameter CS .

$$\mathbf{h}' = \begin{cases} f_{RP}(\mathbf{h}) & (t < CS) \\ f_{FC}(\mathbf{h}) & (\text{otherwise}) \end{cases} \quad (12)$$

This switching is the same as pre-training the other parameters than the layer for dimensionality reduction (Fig. 3).

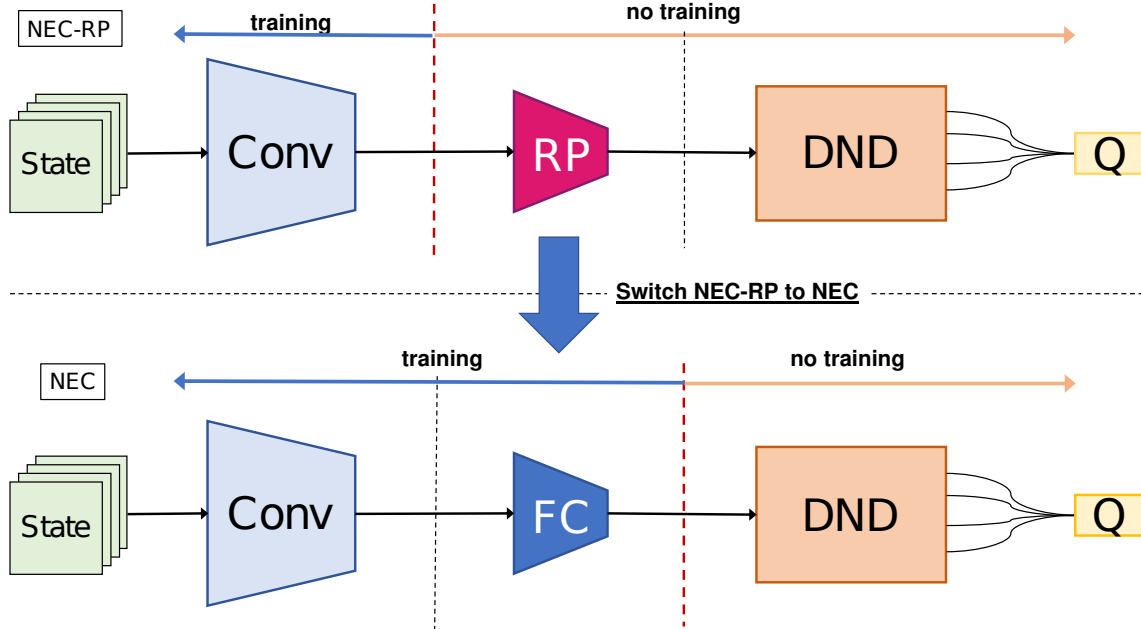


Figure 3: Changing an RP layer of NEC-RP to a FC layer

6.4. Initialization

The initialization of the convolution layers and the fully connected layer is the same as DQN and NEC. As for a random matrix in the RP layer, we should consider from the Table 1, but in NEC-RP we consider that **Embedding quality** is the most important factor. Therefore, we adopt Gaussian random projection. In other words, we use a random matrix \mathbf{R} generated by Gaussian random with mean 0, variance $1/d_{\mathbf{y}}$, where \mathbf{y} is the vector output by the previous layer. Naturally the bias is a zero vector.

7. Experiments

Alike DQN and NEC, we experiment with Atari2600, a video game environment provided by OpenAI Gym Brockman et al. (2016). As for baseline NEC, we use the one reproduced in reinforcement learning library Coach Caspi et al. (2017), and implement NEC-RP¹ with this framework. We also use Faiss Johnson et al. (2017) as an approximate nearest

1. Our implementation is available on <https://github.com/dnishio/NEC-RP>.

neighbor library. The target video games are $\{MsPacman, SpaceInvaders, Bowling, Boxing, DoubleDunk\}$. We experiment three times with different random seeds for each game, on the other hand, we use a fixed seed for random matrix values in our RP layer. In each experiment, our agent learns for 10M frames. In order to prevent the learning from becoming unstable due to the variation in the scale of the rewards of the games, DQN clips the rewards r to $-1 \leq r \leq 1$. However, clipping them affects the final sum of total rewards because high rewards are treated the same as small rewards [van Hasselt et al. (2016), Hessel et al. (2018b)]. NEC-RP also does not clip them because NEC achieved stable learning without it. The parameters such as CNN and DND are the same as NEC. We show the details of our experiment parameters and our architecture in Table ??, ?? in Appendix ??.

7.1. NEC vs. NEC-RP

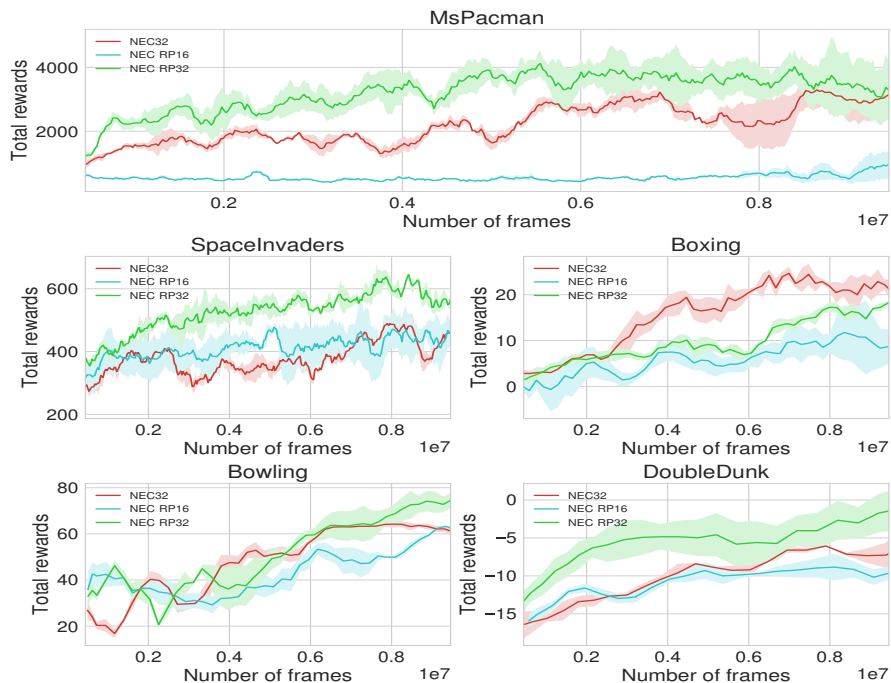


Figure 4: The result of NEC vs. NEC-RP

Here we compare NEC and NEC-RP. NEC reduces the embedding dimensions to 32 dimensions (NEC32), and NEC-RP reduces to 32 dimensions (NEC-RP32) and reduces to 16 dimensions (NEC-RP16). In the comparison of NEC32 and NEC-RP32, we verify whether our proposed architecture that introduced random projection outperforms the performances of NEC. Additionally, in the comparison between NEC-RP32 and NEC-RP16, we examine how the wide range of the inequality (11) affects their performances.

We have shown the result in Fig. 4. Comparing NEC32 and NEC-RP32, we have found that NEC-RP has obtained higher scores in the case of four out of five games. From this fact, we have considered that not only it has been possible to introduce random projection into NEC, but also NEC-RP has learned faster because the parameters required to update have been reduced. In particular, we have shown that it has been possible to stably learn even in the games with the rewards $r > 1$ such as *MsPacman* and *Bowling*. However, in *Boxing* performance of NEC has been superior to ours in the early stages of learning. This has implied that there have been some tasks have been difficult for NEC-RP.

We have also found that random projection has had bad performance if we reduce embedding dimensions too much. In fact, we have seen that the score has been worse in all five games when we have reduced the dimensions to 16 dimensions. Since NEC-RP uses k-nearest neighbor algorithm, the computation time can be reduced when we reduce the dimensions as small as possible, but the performance is sacrificed. Therefore, it is important to consider the tradeoff.

7.2. Switching NEC-RP to NEC

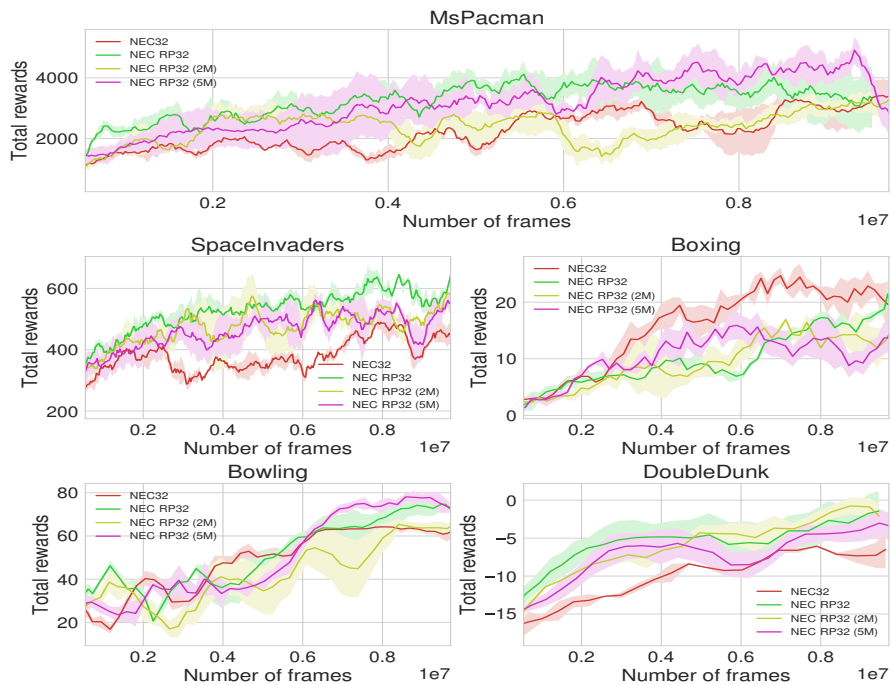


Figure 5: The result of switching NEC-RP to NEC

We compare the performance by switching NEC-RP to NEC, that is, we switch the RP layer to a FC layer with a heuristic step count CS . This comparison indicates experimentally

that it is possible to switch from NEC-RP to NEC, and we also verify whether the flexibility of learning about the neural networks can reduce the losses that random projection cannot minimize. We fix the embedding reduced dimensions to 32 dimensions, and we experiment in the case of switching in 2M frames (NEC-RP32 (2M)) and in 5M frames (NEC-RP32 (5M)).

We have shown the result in Fig. 5. Even when we have switched it in 2M frames or in 5M frames, it has learned without a sharp drop in their performances. From the result, we have found that it has been possible to switch the RP layer to a FC layer at any time. However, the performance has been about the same as NEC such as *MsPacman* and *Bowling* at 10M frames if the timing of switching is too early.

On the other hand, in the case of switching in 5M frames, we have shown that the result of exceeding the NEC-RP’s performance in *Bowling* even though it has been the architecture similar to NEC at 10M frames. We have also shown similar performance to NEC-RP in the other games, and we have found that it has been effective to use the RP layer only at the beginning of NEC’s training. However, in *MsPacman* we have observed that the score has dropped sharply around 10M frames. It is necessary to confirm whether we see this phenomenon in other games in long-term experiments in the future.

8. Conclusion

In this research, we have proposed NEC-RP as the more stable and efficient learning architecture than NEC. We have experimented with the Atari games and actually have outperformed the performance in four out of five games and have shown that our agent has learned efficiently. In addition, we have experimented to switch from NEC-RP to NEC, and we have found that it is possible to improve the performance by switching an RP layer to a FC layer.

As future work, we should experiment with long-term training. In NEC’s paper, there is a report that NEC the performance is inferior to DQN one as learning progresses.

We also need to verify the performance of NEC-RP by long-term training because it may be inferior to DQN like NEC.

Moreover, our architecture is easily available for other architectures that use NEC, and we can expect to improve the performance. We would like to research whether we can apply it to these methods in the future.

References

- Dimitris Achlioptas. Database-friendly random projections. In *Proceedings of the Twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 21-23, 2001, Santa Barbara, California, USA*, 2001. doi: 10.1145/375551.375608. URL <https://doi.org/10.1145/375551.375608>.
- Richard Bellman. On the theory of dynamic programming. *Proceedings of the National Academy of Sciences of the United States of America*, 38(8):716, 1952.

- Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975. doi: 10.1145/361002.361007. URL <http://doi.acm.org/10.1145/361002.361007>.
- Ella Bingham and Heikki Mannila. Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining, San Francisco, CA, USA, August 26-29, 2001*, pages 245–250, 2001. URL <http://portal.acm.org/citation.cfm?id=502512.502546>.
- Charles Blundell, Benigno Uria, Alexander Pritzel, Yazhe Li, Avraham Ruderman, Joel Z. Leibo, Jack W. Rae, Daan Wierstra, and Demis Hassabis. Model-free episodic control. *CoRR*, abs/1606.04460, 2016. URL <http://arxiv.org/abs/1606.04460>.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *CoRR*, abs/1606.01540, 2016. URL <http://arxiv.org/abs/1606.01540>.
- Itai Caspi, Gal Leibovich, Gal Novik, and Shadi Endrawis. Reinforcement learning coach, December 2017. URL <https://doi.org/10.5281/zenodo.1134899>.
- Sanjoy Dasgupta. Experiments with random projection. *CoRR*, abs/1301.3849, 2013. URL <http://arxiv.org/abs/1301.3849>.
- Steven Hansen, Alexander Pritzel, Pablo Sprechmann, André Barreto, and Charles Blundell. Fast deep reinforcement learning using online adjustments from the past. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 10590–10600, 2018.
- Robert Hecht-Nielsen et al. Context vectors: general purpose approximate meaning representations self-organized from raw data. *Computational intelligence: Imitating life*, pages 43–56, 1994.
- Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Gheshlaghi Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3215–3222, 2018a. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17204>.
- Matteo Hessel, Hubert Soyer, Lasse Espeholt, Wojciech Czarnecki, Simon Schmitt, and Hado van Hasselt. Multi-task deep reinforcement learning with popart. *CoRR*, abs/1809.04474, 2018b. URL <http://arxiv.org/abs/1809.04474>.
- Mika Sarkin Jain and Jack Lindsey. Semiparametric reinforcement learning. *ICLR 2018 Workshop*, 2018.

- Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017.
- William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. URL <http://arxiv.org/abs/1312.6114>.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States.*, pages 1106–1114, 2012.
- Ping Li, Trevor Hastie, and Kenneth Ward Church. Very sparse random projections. In *Proceedings of the Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Philadelphia, PA, USA, August 20-23, 2006*, pages 287–296, 2006. doi: 10.1145/1150402.1150436. URL <https://doi.org/10.1145/1150402.1150436>.
- Long Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8:293–321, 1992. doi: 10.1007/BF00992699. URL <https://doi.org/10.1007/BF00992699>.
- Zichuan Lin, Tianqi Zhao, Guangwen Yang, and Lintao Zhang. Episodic memory deep q-networks. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 2433–2439, 2018. doi: 10.24963/ijcai.2018/337. URL <https://doi.org/10.24963/ijcai.2018/337>.
- Xiangrui Meng and Michael W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 91–100, 2013. doi: 10.1145/2488608.2488621. URL <https://doi.org/10.1145/2488608.2488621>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas 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(7540):529–533, 2015. doi: 10.1038/nature14236. URL <https://doi.org/10.1038/nature14236>.
- Daichi Nishio and Satoshi Yamane. Faster deep q-learning using neural episodic control. In *2018 IEEE 42nd Annual Computer Software and Applications Conference, COMPSAC 2018, Tokyo, Japan, 23-27 July 2018, Volume 1*, pages 486–491, 2018. doi: 10.1109/COMPSAC.2018.00075. URL <https://doi.org/10.1109/COMPSAC.2018.00075>.

- Jing Peng and Ronald J. Williams. Incremental multi-step q-learning. *Machine Learning*, 22(1-3):283–290, 1996. doi: 10.1023/A:1018076709321. URL <https://doi.org/10.1023/A:1018076709321>.
- Alexander Pritzel, Benigno Uria, Sriram Srinivasan, Adrià Puigdomènech Badia, Oriol Vinyals, Demis Hassabis, Daan Wierstra, and Charles Blundell. Neural episodic control. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 2827–2836, 2017. URL <http://proceedings.mlr.press/v70/pritzel17a.html>.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *CoRR*, abs/1511.05952, 2015. URL <http://arxiv.org/abs/1511.05952>.
- Hado P. van Hasselt, Arthur Guez, Matteo Hessel, Volodymyr Mnih, and David Silver. Learning values across many orders of magnitude. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4287–4295, 2016.
- Santosh Srinivas Vempala. *The Random Projection Method*, volume 65 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. DIMACS/AMS, 2004. ISBN 0-8218-3793-1. URL <http://dimacs.rutgers.edu/Volumes/Vol65.html>.
- Christopher J. C. H. Watkins and Peter Dayan. Technical note q-learning. *Machine Learning*, 8:279–292, 1992. doi: 10.1007/BF00992698. URL <https://doi.org/10.1007/BF00992698>.
- Piotr Iwo Wójcik. Random projection in deep neural networks. *CoRR*, abs/1812.09489, 2018. URL <http://arxiv.org/abs/1812.09489>.
- David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2):1–157, 2014. doi: 10.1561/04000000060. URL <https://doi.org/10.1561/04000000060>.
- Pengcheng Wu, Steven C. H. Hoi, Duc Dung Nguyen, and Ying He. Randomly projected kd-trees with distance metric learning for image retrieval. In *Advances in Multimedia Modeling - 17th International Multimedia Modeling Conference, MMM 2011, Taipei, Taiwan, January 5-7, 2011, Proceedings, Part II*, pages 371–382, 2011. doi: 10.1007/978-3-642-17829-0_35. URL https://doi.org/10.1007/978-3-642-17829-0_35.