
Empowerment-driven Exploration using Mutual Information Estimation

Navneet Madhu Kumar*
Department of Computer Science
Technical University of Munich
Germany
navneet.nmk@gmail.com

Abstract

Exploration is a difficult challenge in reinforcement learning and is of prime importance in sparse reward environments. However, many of the state of the art deep reinforcement learning algorithms, that rely on epsilon-greedy exploration, fail on these environments. In such environments, empowerment can serve as an intrinsic reward signal to enable the agent to explore by maximizing the influence it has over the near future. The mutual information, required to calculate empowerment, is estimated using a Mutual Information Neural Estimator and a forward dynamics model. We demonstrate that an empowerment driven agent is able to find the rewards in Montezuma’s Revenge where a baseline DQN agent fails to find any.

1 Introduction

Reinforcement learning (RL) tackles sequential decision making problems by formulating them as tasks where an agent must learn how to act optimally through trial and error interactions with the environment. The goal is to maximize the sum of the numerical reward signal observed at each time step. Here we focus on the problem of exploration in RL, which aims to reduce the number of interactions an agent needs in order to learn to perform well.

The most common approach to exploration in absence of any knowledge about the environment is to perform random actions. As knowledge is gained, the agent can use it to attempt to increase its performance by taking greedy actions, while retaining some chance to choose random actions to further explore the environment (epsilon-greedy exploration). However, if rewards are sparse or are not sufficiently informative to allow performance improvements, epsilon-greedy fails to explore sufficiently far. Several methods have been described that bias the agent’s actions towards novelty mostly by using optimistic initialization or curiosity signals.

In particular, we opt for empowerment Klyubin et al. [2005], an information-theoretic formulation of the agent’s influence on the near future. The value of a state, its empowerment value, is given by the maximum mutual information between a control input and the successor state the agent could achieve.

Mutual Information is known to be very hard to calculate. Thankfully, recent advances in neural estimation Belghazi et al. [2018] enable effective computation of mutual information between high dimensional input / output pairs of deep neural networks, and in this work we leverage these techniques to calculate the empowerment of a state.

*Code is available at <https://github.com/navneet-nmk/pytorch-rl>

2 Empowerment-Driven Exploration

Our agent is composed of two networks: a reward generator that outputs a empowerment-driven intrinsic reward signal and a policy that outputs a sequence of actions to maximize that reward signal. In addition to intrinsic rewards, the agent optionally may also receive some extrinsic reward from the environment. Let the intrinsic curiosity reward generated by the agent at time t be r_i and the extrinsic reward be r_e . The policy sub-system is trained to maximize the sum of these two rewards $r=r_i+r_e$.

2.1 Empowerment

The Empowerment value for a state \mathbf{s} is defined as the channel capacity between the action \mathbf{a} and the following state \mathbf{s}' Klyubin et al. [2005],

$$\mathcal{E}(\mathbf{s}) = \max_{\pi} \mathcal{I}(\mathbf{s}', \mathbf{a} | \mathbf{s}). \quad (1)$$

Where \mathcal{I} is the mutual information. π is the empowerment maximizing policy. Empowerment, therefore, is the maximum information an agent can transfer to it's environment by changing the next states through it's actions.

This mutual information in eq. (1) can further be represented in the form of the KL Divergence,

$$\begin{aligned} \mathcal{I}(\mathbf{s}', \mathbf{a} | \mathbf{s}) &= \text{KL}(p(\mathbf{s}', \mathbf{a} | \mathbf{s}) || p(\mathbf{s}' | \mathbf{s}) \omega(\mathbf{a} | \mathbf{s})) \\ &= \iint p(\mathbf{s}', \mathbf{a} | \mathbf{s}) \ln \frac{p(\mathbf{s}', \mathbf{a} | \mathbf{s})}{p(\mathbf{s}' | \mathbf{s}) \omega(\mathbf{a} | \mathbf{s})} d\mathbf{s}' d\mathbf{a}, \end{aligned}$$

To compute the mutual information, we will be using the formulation as explained in Belghazi et al. [2018].

2.2 Mutual Information Neural Estimation

The Mutual Information Neural Estimation (MINE, Belghazi et al. [2018]) learns a neural estimate of the mutual information of continuous variables, is strongly consistent and can be used to learn the empowerment value of a state by using a forward dynamics model, $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ to get the samples from the marginal distribution of $p(\mathbf{s}_{t+1} | \mathbf{s}_t)$, and the policy, $\pi(\mathbf{a}_t | \mathbf{s}_t)$.

Following, Belghazi et al. [2018], we train a discriminator (a classifier) to distinguish between samples coming from the joint, $p(\mathbf{s}', \mathbf{a} | \mathbf{s})$, and the marginal distributions, $p(\mathbf{s}_{t+1} | \mathbf{s}_t)$ and $\pi(\mathbf{a}_t | \mathbf{s}_t)$. MINE relies on a lower-bound to the mutual information based on the Donsker and Varadhan [1983],

$$\begin{aligned} \mathcal{I}(\mathbf{s}', \mathbf{a} | \mathbf{s}) &= \text{KL}(p(\mathbf{s}', \mathbf{a} | \mathbf{s}) || p(\mathbf{s}' | \mathbf{s}) \omega(\mathbf{a} | \mathbf{s})) \\ &\geq \mathbb{E}_{p(\mathbf{s}', \mathbf{a} | \mathbf{s})} [T_{\omega}] - \log \mathbb{E}_{p(\mathbf{s}' | \mathbf{s}) \pi(\mathbf{a} | \mathbf{s})} [e^{T_{\omega}}] \end{aligned}$$

We use this estimate of the mutual information (or the empowerment for the state, \mathbf{s}) as the intrinsic reward to train the policy.

2.3 Forward dynamics model

The forward dynamics model, $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$, is used to sample from the marginal distribution, $p(\mathbf{s}_{t+1} | \mathbf{s}_t)$ by marginalizing out the actions,

$$p(\mathbf{s}' | \mathbf{s}) = \int \pi(\mathbf{a} | \mathbf{s}) p(\mathbf{s}', \mathbf{a} | \mathbf{s}) d\mathbf{a}.$$

The dynamics model is trained simultaneously with the statistics network, T and the policy, $\pi(\mathbf{a} | \mathbf{s})$.

3 Experimental Setup

3.1 Agent

The proposed empowerment driven DQN agent is composed of a policy network, the DQN, and the intrinsic reward generator that is composed of the statistics network T and the forward dynamics

Algorithm 1 Joint training of policy π , statistics T and forward dynamics model $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$

Require: cumulation horizon N , initializations for π , T , f and e

repeat

for $m = 1 : M$ **do**

 Sample a batch b of transitions from a replay buffer B

for each transition $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, r_e)$ **do**

 Update f to reduce $\sum_t [e(\mathbf{s}_{t+1}) - f(e(\mathbf{s}_t, \mathbf{a}_t))]^2$

 Update T to increase the mutual information

$$\mathbb{E}_{p(\mathbf{s}', \mathbf{a} | \mathbf{s})} [T_\omega] - \log \mathbb{E}_{p(\mathbf{s}' | \mathbf{s}) \pi(\mathbf{a} | \mathbf{s})} [e^{T_\omega}]$$

 Update the reward function for the transition $r = r_e + \mathcal{I}(\mathbf{s}', \mathbf{a} | \mathbf{s})$

end for

 Update the agent using the bellman update $y = r + \gamma * \max_a Q(\mathbf{s}_t, \mathbf{a}_t)$

end for

until convergence

model, $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$. The implementation is using Pytorch Paszke et al. [2017]. Inputs are a stack of four 84 x 84 gray scale frames. All the observations are then encoded to a 64 dimensional encoding using a random convolutional encoder. All other networks share the same network architecture but use separate artificial neural networks. We used double Q-learning with target network updates every 2000 steps and an experience replay of buffer capacity 1000000 steps. Training of the models starts at 1000 steps and follows Algorithm 1. All 3 networks are trained on batch size of 64, using Adam optimizers with learning rate of 1e-2, 1e-3 and 1e-4 for the forward dynamics, statistics and policy network respectively.

The DQN Agent predicts 18 Q-values corresponding to the 18 actions in ALE, uses relu activations and a discount factor of 0.99. The extrinsic reward is clipped between [-1, 1] and the gradients of the temporal difference loss is clamped between [-1, 1]. The value used for β is 0.1.




Figure 1: Rewards in Montezuma's Revenge

3.2 Results

The DQN Agent trained using the empowerment intrinsic motivation is able to consistently exit the room one and gather the rewards whereas the agent trained on the reward signal of the game fails to receive any reward (Figure 1).

Owing to computational limits, the size of the environment encoding was limited to 64 which could prove to be insufficient for the Atari Environment. This is a parameter which needs to be investigated further in future work.

4 Conclusion and future work

The experiments show that using empowerment, calculated using mutual information neural estimation, as an intrinsic motivator can help an agent to consistently achieve rewards.

Compared to an agent which just receives the external reward signal from the game, the empowerment driven agent is able to consistently achieve the rewards in the first level of Montezuma’s revenge and enter the second room.

Future work includes the testing of the method on the entire Atari game suite as well as increasing the model sizes and embedding sizes.

References

- Alexander S. Klyubin, Daniel Polani, and Chrystopher L. Nehaniv. All else being equal be empowered. In Mathieu S. Capcarrère, Alex A. Freitas, Peter J. Bentley, Colin G. Johnson, and Jon Timmis, editors, *Advances in Artificial Life*, pages 744–753, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. ISBN 978-3-540-31816-3.
- Ishmael Belghazi, Sai Rajeswar, Aristide Baratin, R. Devon Hjelm, and Aaron C. Courville. MINE: mutual information neural estimation. *CoRR*, abs/1801.04062, 2018. URL <http://arxiv.org/abs/1801.04062>.
- M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on Pure and Applied Mathematics*, 36(2):183–212, 1983. doi: 10.1002/cpa.3160360204. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.3160360204>.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- Shakir Mohamed and Danilo Jimenez Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 2125–2133, 2015.
- Maximilian Karl, Maximilian Soelch, Philip Becker-Ehmck, Djalel Benbouzid, Patrick van der Smagt, and Justin Bayer. Unsupervised real-time control through variational empowerment. *arXiv preprint arXiv:1710.05101*, 2017.
- Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. Large-scale study of curiosity-driven learning. In *arXiv:1808.04355*, 2018.
- Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Jul 2017. doi: 10.1109/cvprw.2017.70. URL <http://dx.doi.org/10.1109/CVPRW.2017.70>.
- Jurgen Schmidhuber. A possibility for implementing curiosity and boredom in model-building neural controllers. 1991.
- Manuel Lopes, Tobias Lang, Marc Toussaint, and Pierre-Yves Oudeyer. Exploration in model-based reinforcement learning by empirically estimating learning progress. In *NIPS*, 2012.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning, 2013.
- Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning, 2015.
- Marc G. Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Rémi Munos. Unifying count-based exploration and intrinsic motivation. In *NIPS*, 2016.
- Susanne Still and Doina Precup. An information-theoretic approach to curiosity-driven reinforcement learning. *Theory in Biosciences*, 131:139–148, 2011.
- A. PixelCNN Hyper-parameters. Count-based exploration with neural density models. 2017.
- Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. In *NIPS*, 2016.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *CoRR*, abs/1802.06070, 2018.
- Justin Fu, John D. Co-Reyes, and Sergey Levine. Ex2: Exploration with exemplar models for deep reinforcement learning. In *NIPS*, 2017.
- Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. exploration: A study of count-based exploration for deep reinforcement learning, 2016.
- Jürgen Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, Sep 2010. ISSN 1943-0612. doi: 10.1109/tamd.2010.2056368. URL <http://dx.doi.org/10.1109/TAMD.2010.2056368>.
- Bradly C. Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models, 2015.
- Sainbayar Sukhbaatar, Zeming Lin, Ilya Kostrikov, Gabriel Synnaeve, Arthur Szlam, and Rob Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play, 2017.
- Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, Charles Blundell, and Shane Legg. Noisy networks for exploration, 2017.
- Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y. Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration, 2017.