

Demystify the Gravity Well in the Optimization Landscape (Student Abstract)

Jason Xiaotian Dou¹, Runxue Bao¹, Susan Song², Shuran Yang³, Yanfu Zhang¹, Paul Pu Liang², Haiyi Harry Mao¹

¹ University of Pittsburgh

² Carnegie Mellon University

³ University of California, Berkeley

3700 O’Hara Street Benedum Hall of Engineering, Pittsburgh, PA 15261 USA

jasondpku@gmail.com

Abstract

We provide both empirical and theoretical insights to demystify the gravity well phenomenon in the optimization landscape. We start from describe the problem setup and theoretical results (an escape time lower bound) of the Softmax Gravity Well (SGW) in the literature. Then we move toward the understanding of a recent observation called ASR gravity well. We provide an explanation of why normal distribution with high variance can lead to suboptimal plateaus from an energy function point of view. We also contribute to the empirical insights of curriculum learning by comparison of policy initialization by different normal distributions. Furthermore, we provide the ASR escape time lower bound to understand the ASR gravity well theoretically. Future work includes more specific modeling of the reward as a function of time and quantitative evaluation of normal distribution’s influence on policy initialization.

Introduction

A gravity well is the result of the pull of gravity caused by a body in space. In the context of reinforcement learning (Huang et al. 2021), the policy gradient optimizes a parameterized policy to maximize the long-term expected reward. Softmax transform is normally used to produce conditional action distributions, which results in the softmax policy gradient (SPG). Softmax transform is as the following. For $\theta \in \mathbb{R}^K$, $\pi_\theta = \text{softmax}(\theta)$ is defined by $\pi_\theta(a) = \frac{\exp\{\theta(a)\}}{\sum_{a'} \exp\{\theta(a')\}}$ for all $a \in \{1, \dots, K\}$ (Mei et al. 2020).

The phenomenon “softmax gravity well (SGW)” is described as the following: gradient ascent trajectories are drawn toward suboptimal corners of the probability simplex and subsequently slow down in their progress toward the optimal vertex. The behavior of SPG is impacted by initialization (Mei et al. 2020). It’s similar to the concept of saddle point in the optimization landscape. An experiment on a Markov Decision Process (MDP) has been used to illustrate SGW in (Mei et al. 2020) very well. The reward r is defined as $r = (b + \Delta, b, \dots, b)^\top \in [0, 1]^K$ for some b , such that $\Delta > 0$ is the reward gap. And an escape time lower bound has been provided for the single state MDP setting:

Theorem 1 (SPG escape time lower bound (Mei et al. 2020)). *In a single-state MDP, for any learning rate $\eta_t \in (0, 1]$, there exists an initialization of the policy π_{θ_1} and a positive constant C , such that SPG with full gradients cannot escape a suboptimal corner before time $t_0 := \frac{C}{\Delta \cdot \pi_{\theta_1}(a^*)}$:*

$$(\pi^* - \pi_{\theta_t})^\top r \geq 0.9 \cdot \Delta \quad (1)$$

A single state MDP with 6 actions is a simplified version of real-world reinforcement learning settings. In a three-action, multiple-state MDP setting (Dou et al. 2022b), interestingly, a phenomenon similar to softmax gravity well is also observed, which is named ASR gravity well. The goal of this paper is to apply the theoretical and empirical insights of the softmax gravity well to demystify the ASR gravity well.




Figure 1: ASR Gravity Well (Dou et al. 2022b)

Main Contributions

The ASR gravity well is illustrated in Figure . Details of the experimental setup are in Dou et al. (2022b). When the policy is initialized with normal distribution with high variance, all the curves fall into a global minimum at around epoch 15. The “reverse peaks” are very significant within five (Recall@1, Recall@2, Recall@4, Recall@8, and NMI) out of six benchmarks. That means the phenomenon arises for both information retrieval and clustering applications (Chen et al. 2019; Mao et al. 2022b,a; Dou et al. 2022a,b; Liu et al. 2022). This empirical observation coincides elegantly with Theorem 1: a high standard deviation normal distribution can be the policy initialization of the policy π_{θ_1} which traps

the algorithm into the plateaus which cannot escape within a time lower bound. To understand this more clearly, we first recall that the probability density function of the normal distribution has μ as the mean and σ as the standard deviation. From a statistical physics point of view, a high variance in the normal distribution often leads to high energy, making molecules trap into energy wells (Zuckerman 2010; Liu and Tian 2022).

Then we dig into the details of the normal distribution of the policy initialization models. In the Adaptive Sampling with Reward (ASR) setting (Dou et al. 2022b), the policy initialization distribution guides a sampling distribution of $p(O_n | O_a)$, which builds on the distance between negative sample O_n and anchor sample O_a in triplet loss construction (Chen et al. 2019; Yan et al. 2022; Dou, Luo, and Yang 2022). A high variance normal distribution policy initialization means that at the beginning of the sampling stage, the model will sample negative samples from a long range of distances from the anchor sample, which includes hard, semi-hard, and easy negative samples. However, in Figure , the gravity well only appears phenomenally in the normal high case, while for the normal high initialization, the curves are relatively flat.

Curriculum learning (Bengio et al. 2009) offers two empirical insights: introducing gradually more difficult examples speeds up online training, and cleaner examples may yield better generation faster. Comparing the curves belonging to normal high and normal low distributions, we can see that from similar starting points, though the normal high curves fall into the gravity wells, they end up with higher benchmarks in all six benchmarks. So we can add a piece new insight to curriculum learning, learning broadly (from hard negatives to easy negatives) may confuse the model in the short term and compromise downstream tasks’ performance, but in the long run, learning broadly(normal high) will surpass learning within a narrow range (normal low).

Finally, we provide an ASR escape time lower bound which is adjusted from theorem 1. There are two key insights that distinguish the ASR escape time lower bound from the SPG escape time lower bound. First, $\Delta(s)$ becomes state-dependent. Second, the reward r becomes a bounded function of t , such that $r(t) = (b_1(t), b_2(t), \dots, b_k(t))^T$. As our reward is a weighted combination of recall and NMI evaluated on a validation dataset and it evolves over training. The proof of the theorem is in the supplementary material.

Theorem 2 (ASR escape time lower bound). *In the multi-state MDP of the ASR setting, we use s to denote the current state. The reward $r(t)$ at time t is defined as $r(t) = (b_1(t) + \Delta(s), b_2(t), \dots, b_k(t))^T \in [0, 1]^K$ for $b_i(t), 1 \leq i \leq k$, such that $\Delta(s) > 0$ is the reward gap, were $b_i \in \mathcal{B} : \mathbb{R} \rightarrow \mathbb{R}$ and $\exists \beta$, such that $b_i(t) \leq \beta \leq b_1(t) + \Delta(s)$. For any learning rate $\eta_t \in (0, 1]$, there exists an initialization of the policy π_{θ_1} and a positive constant C , such that the ASR framework cannot escape a suboptimal plateau before time $t_0 := \frac{C}{\Delta(s) \cdot \pi_{\theta_1}(a^*)}$, i.e., it will hold that*

$$(\pi^* - \pi_{\theta_t})^T r(t) \geq 0.9 \cdot \Delta(s) \quad (2)$$

Theorem 2 shows even ASR framework with relaxed conditions of reward, it still enjoys a similar lower bound.

Discussion

In the above section, we first demystify the ASR gravity well from the energy function’s viewpoint. Then we complement curriculum learning’s insights by comparison of different normal distribution initialization. We make a first step towards understanding the ASR gravity well theoretically by providing a new escape time lower bound. In future work, we plan to further enhance our understanding of the “wells” from two perspectives. First, we want to provide quantitative explanations of how normal distribution influences policy initialization. Second, we want to explore the approximation of the function $r(t)$ and $\Delta(s)$ to make the theory aligns more closely with empirical practice.

References

- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *ICML*, 41–48.
- Chen, S.; Luo, L.; Yang, J.; Gong, C.; Li, J.; and Huang, H. 2019. Curvilinear Distance Metric Learning. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Neurips*, volume 32. Curran Associates, Inc.
- Dou, J. X.; Jia, M.; Zaslavsky, N.; Bao, R.; Zhang, S.; Ni, K.; Liang, P. P.; Mao, H.; and Mao, Z. 2022a. Learning More Effective Cell Representations Efficiently. In *NeurIPS 2022 Workshop on Learning Meaningful Representations of Life*.
- Dou, J. X.; Luo, L.; and Yang, R. M. 2022. An optimal transport approach to deep metric learning (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 12935–12936.
- Dou, J. X.; Pan, A. Q.; Bao, R.; Mao, H. H.; Luo, L.; and Mao, Z. 2022b. Sampling through the lens of sequential decision making. *arXiv preprint arXiv:2208.08056*.
- Huang, B.; Feng, F.; Lu, C.; Magliacane, S.; and Zhang, K. 2021. Adarl: What, where, and how to adapt in transfer reinforcement learning. *arXiv preprint arXiv:2107.02729*.
- Liu, H.; and Tian, X. 2022. SGEM: stochastic gradient with energy and momentum. *arXiv preprint arXiv:2208.02208*.
- Liu, J.; Lian, J.; Sprott, J. C.; Liu, Q.; and Ma, Y. 2022. The Butterfly Effect in Primary Visual Cortex. *IEEE Transactions on Computers*.
- Mao, H.; Jia, M.; Dou, J. X.; Zhang, H.; and Benos, P. V. 2022a. COEM: cross-modal embedding for metacell identification. *ICML Workshop on Computational Biology*.
- Mao, H.; Liu, H.; Dou, J. X.; and Benos, P. V. 2022b. Towards Cross-Modal Causal Structure and Representation Learning. In *Machine Learning for Health*.
- Mei, J.; Xiao, C.; Dai, B.; Li, L.; Szepesvári, C.; and Schuurmans, D. 2020. Escaping the Gravitational Pull of Softmax. *Neurips*, 33.
- Yan, J.; Luo, L.; Xu, C.; Deng, C.; and Huang, H. 2022. Noise Is Also Useful: Negative Correlation-Steered Latent Contrastive Learning. In *CVPR*, 31–40.
- Zuckerman, D. M. 2010. *Statistical physics of biomolecules: an introduction*. CRC press.