Training Larger Networks for Deep Reinforcement Learning

Ziniu Li
ziniuli@link.cuhk.edu.cn

CUHKSZ

December 13, 2022

Introduction

Large neural network models have superior performance in many domains (CV, NLP, etc).

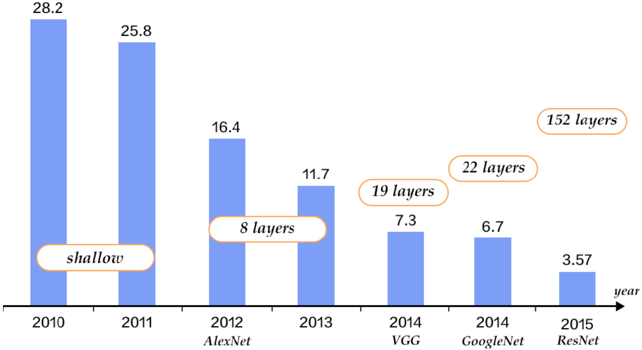

Figure 1: Classification error and model size on the Imagenet competition.

- Are large models helpful in RL?

Introduction

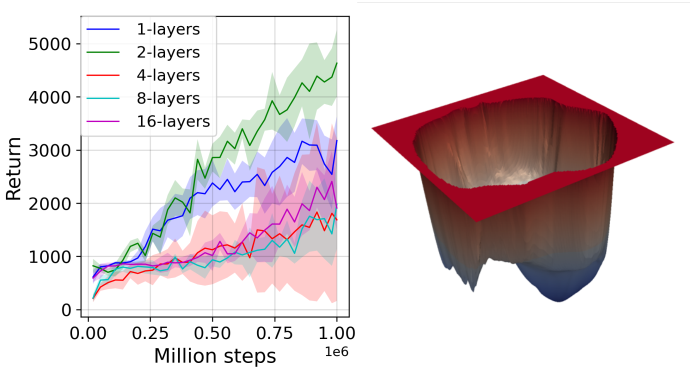

(a) Average return.

(b) Loss surface.

Figure 2: Results from [OJK21].

- Directly increasing the number of layers (with $N_{\text{unit}} = 256$) could even hurt the performance.

Introduction

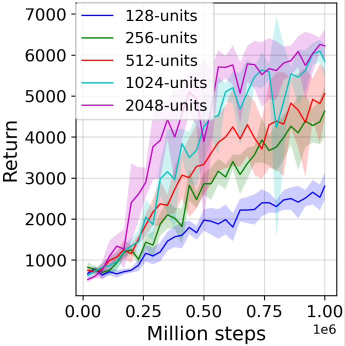

(a) Average return.

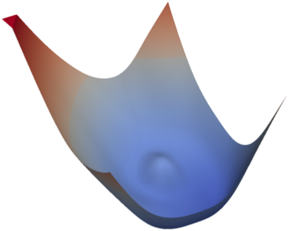

(b) Loss surface.

Figure 3: Results from [OJK21].

- Directly increase the width (with $N_{\text{layer}} = 2$) improve the performance by about 1.5 times.

Introduction

	SAC			TD3		
ENVIRONMENT	OURS	OFENET	ORIGINAL	OURS	OFENET	ORIGINAL
HOPPER-V2	3467.3	**3511.6**	3316.6	3206.7	3488.3	**3613.0**
WALKER2D-V2	**8802.4**	5237.0	3401.5	**7645.8**	4915.1	4515.6
HALFCHEETAH-V2	**19209.9**	16964.1	14116.1	**18147.5**	16259.5	13319.9
ANT-V2	**14021.0**	8086.2	5953.1	**12811.3**	8472.4	6148.6
HUMANOID-V2	**14858.2**	9560.5	6092.6	**13282.0**	120.6	340.5

Figure 4: Results from [OJK21].

- Training with large networks achieves the SOTA performance.

How to Train Larger Networks for Deep RL?

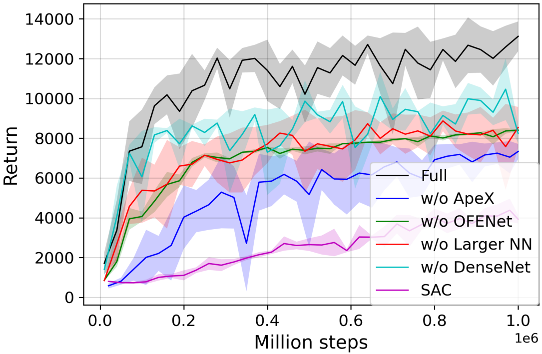

Figure 5: Results from [OJK21].

- OFENET: predicting the next state to learn presentation.
- Wider NN: $N_{\text{unit}} = 2048$.
- Dense Net: advanced architecture.
- Ape-X training: parallelization for on-policy samples.

References I

[OJK21] Kei Ota, Devesh K Jha, and Asako Kanezaki. Training larger networks for deep reinforcement learning. *arXiv preprint arXiv:2102.07920*, 2021.