




Figure 1: A comparison of OPEN and Adam in Craftax-Classic [1], a JAX-based reimplementaion of Crafter [2]. We examine 0-shot transfer for both *OPEN* and Adam [3], as well as the performance of Adam when finetuned directly in Craftax. OPEN generalizes significantly better than Adam, and performs only slightly worse than Adam (Finetuned). Plotted with standard error over 32 seeds.

Table 1: Hyperparameters for Craftax. Adam (Finetuned) takes most hyperparameters from Matthews et al. [1]. We set $W = 64$ and search over LR s in $\{0.00005, 0.0005\}$ with 8 validation seeds.

Hyperparameter	OPEN (0 shot)	Adam (0 shot)	Adam (Finetuned)
<i>Learning rate</i>	—	0.0001	0.0005
β_1	—	0.99	0.9
β_2	—	0.99	0.999
<i>Anneal LR</i>	—	✓	✓
<i>Number of Environments N_{envs}</i>	1024	1024	256
<i>Number of Environment Steps N_{steps}</i>	20	20	16
<i>Total Timesteps T</i>	3×7	3×10^7	3×7
<i>Number of Minibatches $N_{minibatch}$</i>	16	16	8
<i>Number of Epochs L</i>	2	2	4
<i>Discount Factor γ</i>	0.99	0.99	0.99
<i>GAE λ</i>	0.8	0.8	0.8
<i>PPO Clip ϵ</i>	0.2	0.2	0.2
<i>Value Function Coefficient c_1</i>	0.5	0.5	0.5
<i>Entropy Coefficient c_2</i>	0.01	0.01	0.01
<i>Max Gradient Norm</i>	0.5	0.5	0.5
<i>Layer Width W</i>	64	64	64
<i>Number of Hidden Layers H</i>	2	2	2
<i>Activation</i>	tanh	tanh	tanh

References

- [1] Michael Matthews, Michael Beukman, Benjamin Ellis, Mikayel Samvelyan, Matthew Jackson, Samuel Coward, and Jakob Foerster. Craftax: A lightning-fast benchmark for open-ended reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2024.
- [2] Danijar Hafner. Benchmarking the spectrum of agent capabilities. *arXiv preprint arXiv:2109.06780*, 2021.
- [3] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2017.