

Figure 1: (a) A simplified visualization of the crafter environment with many objects to interact with and a complex technology tree. Details can be found in Appendix G. (b) Learning curve of ELDEN (ours) compared to baseline approaches (mean \pm std dev of the number of stages completed across three random seeds), where ELDEN outperforms all baselines, including RND.




Figure 2: Learning curve of ELDEN (ours) compared to baseline approaches (mean \pm std dev of episode reward across three random seeds) in the Deepmind control suite Cheetah domain with sparse reward. As the environment is challenging in the aspect of precise low-level sensorimotor control, for both ELDEN and CAI, the reward is close to zero throughout the training.




Figure 3: Learning curve of ELDEN (ours) compared to baseline approaches (mean \pm std dev of the number of stages completed across three random seeds) in the Kitchen domain using low-level actions. None of the methods make meaningful progress due to the challenges of long-horizon goal reaching and precise low-level sensorimotor under sparse reward.