

We explore the capability of evolution strategies to train a transformer architecture in the reinforcement learning setting. We perform experiments using OpenAI’s highly parallelizable evolution strategy and its derivatives utilizing novelty and quality-diversity searches to train Decision Transformer in Humanoid locomotion environment, testing the ability of these black-box optimization techniques to train even such relatively large (compared to the previously tested in the literature) and complicated (using a self-attention in addition to fully connected layers) models. The tested algorithms proved to be, in general, capable of achieving strong results and managed to obtain high-performing agents both from scratch (randomly initialized model) and from a pretrained model.