GENETIC POLICY OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Genetic algorithms have been widely used in many practical optimization problems. Inspired by natural selection, operators, including mutation, crossover and selection, provide effective heuristics for search and black-box optimization. However, they have not been shown useful for deep reinforcement learning, possibly due to the catastrophic consequence of parameter crossovers of neural networks. Here, we present Genetic Policy Optimization (GPO), a new genetic algorithm for sample-efficient deep policy optimization. GPO uses imitation learning for policy crossover in the state space and applies policy gradient methods for mutation. Our experiments on Mujoco tasks show that GPO as a genetic algorithm is able to provide superior performance over the state-of-the-art policy gradient methods and achieves comparable or higher sample efficiency.

1 INTRODUCTION

Reinforcement learning (RL) has recently demonstrated significant progress and achieves state-of-the-art performance in games (Mnih et al., 2015; Silver et al., 2016), locomotion control (Lillicrap et al., 2015), visual-navigation (Zhu et al., 2017), and robotics (Levine et al., 2016). Among these successes, deep neural networks (DNNs) are widely used as powerful functional approximators to enable signal perception, feature extraction and complex decision making. For example, in continuous control tasks, the policy that determines which action to take is often parameterized by a deep neural network that takes the current state observation or sensor measurements as input. In order to optimize such policies, various policy gradient methods (Mnih et al., 2016; Schulman et al., 2015; 2017; Heess et al., 2017) have been proposed to estimate gradients approximately from rollout trajectories. The core idea of these policy gradient methods is to take advantage of the temporal structure in the rollout trajectories to construct a Monte Carlo estimator of the gradient of the expected return.

In addition to the popular policy gradient methods, other alternative solutions, such as those for black-box optimization or stochastic optimization, have been recently studied for policy optimization. Evolution strategies (ES) is a class of stochastic optimization techniques that can search the policy space without relying on the backpropagation of gradients. At each iteration, ES samples a candidate population of parameter vectors ("genotypes") from a probability distribution over the parameter space, evaluates the objective function ("fitness") on these candidates, and constructs a new probability distribution over the parameter space using the candidates with the high fitness. This process is repeated iteratively until the objective is maximized. Covariance matrix adaptation evolution strategy (CMA-ES; Hansen & Ostermeier (2001)) and recent work from Salimans et al. (2017) are examples of this procedure. These ES algorithms have also shown promising results on continuous control tasks and Atari games, but their sample efficiency is often not comparable to the advanced policy gradient methods, because ES is black-box and thus does not fully exploit the policy network architectures or the temporal structure of the RL problems.

Very similar to ES, genetic algorithms (GAs) are a heuristic search technique for search and optimization. Inspired by the process of natural selection, GAs evolve an initial population of genotypes by repeated applications of three genetic operators - mutation, crossover and selection. The major difference between GA and ES is that the crossover operator in GA is able to provide higher diversity of good candidates in the population. However, the crossover operator is often performed on the parameter representations of two parents, thus making it not suitable for nonlinear neural networks. The straightforward crossover of two neural networks by exchanging their parameters can often destroy the hierarchical relationship of the networks and thus cause a catastrophic drop in per-

formance. NeuroEvolution of Augmenting Topologies (NEAT; Stanley & Miikkulainen (2002b;a)), which evolves neural networks through evolutionary algorithms such as GA, provides a solution to exchange and augment neurons but has found limited success when used as a method of policy search in deep RL for high-dimensional tasks. A major challenge to making GAs work for policy optimization is to design a good crossover operator which efficiently combines two parent policies represented by neural networks and generates an offspring that takes advantage of both parents. In addition, a good mutation operator is needed as random perturbations are often inefficient for high-dimensional policies.

In this paper, we present Genetic Policy Optimization (GPO), a new genetic algorithm for sample-efficient deep policy optimization. There are two major technical advances in GPO. First, instead of using parameter crossover, GPO applies imitation learning for policy crossovers in the state space. The state-space crossover effectively combines two parent policies into an offspring or child policy that tries to mimic its best parent in generating similar state visitation distributions. Second, GPO applies advanced policy gradient methods for mutation. By randomly rolling out trajectories and performing gradient descent updates, this mutation operator is more efficient than random parameter perturbations and also maintains sufficient genetic diversity. Our experiments on several continuous control tasks show that GPO as a genetic algorithm is able to provide superior performance over the state-of-the-art policy gradient methods and achieves comparable or higher sample efficiency.

2 Background and Related Work

2.1 Reinforcement Learning

In the standard RL setting, an agent interacts with an environment $\mathcal{E}$ modeled as a Markov Decision Process (MDP). At each discrete time step t, the agent observes a state s_t and choose an action $a_t \in \mathcal{A}$ using a policy $\pi(a_t|s_t)$, which is a mapping from states to a distribution over possible actions. Here we consider high-dimensional, continuous state and action spaces. After performing the action a_t, the agent collects a scalar reward $r(s_t, a_t) \in \mathcal{R}$ at each time step. The goal in reinforcement learning is to learn a policy which maximizes the expected sum of (discounted) rewards starting from the initial state. Formally, the objective is

$$J(\pi) = \mathbb{E}_{\{(s_t,a_t)\}\sim\mathcal{E},\pi}\left[\sum_{t=0}^{\infty}\gamma^t r(s_t, a_t)\right]$$

where the states s_t are sampled from the environment $\mathcal{E}$ using an unknown system dynamics model $p(s_{t+1}|s_t, a_t)$ and an initial state distribution $p(s_0)$, the actions a_t are sampled from the policy $\pi(a_t|s_t)$ and $\gamma \in (0, 1]$ is the discount factor.

2.2 Policy Gradient Methods

Policy-based RL methods search for an optimum policy directly in the policy space. One popular approach is to parameterize the policy $\pi(a_t|s_t;\theta)$ with θ, express the objective $J(\pi(a_t|s_t;\theta))$ as a function of θ and perform gradient descent methods to optimize it. The REINFORCE algorithm (Williams, 1992) calculates an unbiased estimation of the gradient $\nabla_\theta J(\theta)$ using the likelihood ratio trick. Specifically, REINFORCE updates the policy parameters in the direction of the following approximation to policy gradient

$$\nabla_\theta J(\theta) \approx \sum_{t=0}^{\infty}\nabla_\theta \log\pi(a_t|s_t;\theta)R_t$$

based on a single rollout trajectory, where $R_t = \sum_{i=0}^{\infty}\gamma^i r(s_{t+i}, a_{t+i})$ is the discounted sum of rewards from time step t. The advantage actor-critic (A2C) algorithm (Sutton & Barto; Mnih et al., 2016) uses the state value function (or critic) to reduce the variance in the above gradient estimation. The contribution to the gradient at time step t is $\nabla_\theta \log\pi(a_t|s_t;\theta)(R_t - V^\pi(s_t))$. $R_t - V^\pi(s_t)$ is an estimate of the advantage function $A^\pi(s_t, a_t) = Q^\pi(s_t, a_t) - V^\pi(s_t)$. In practice, multiple rollouts are performed to get the policy gradient, and $V^\pi(s_t)$ is learned using a function approximator.

High variance in policy gradient estimates can sometimes lead to large, destructive updates to the policy parameters. Trust-region methods such as TRPO (Schulman et al., 2015) avoid this by restricting the amount by which an update is allowed to change the policy. TRPO is a second order

algorithm that solves an approximation to a constrained optimization problem using conjugate gradient. Proximal policy optimization (PPO) algorithm (Schulman et al., 2017) is an approximation to TRPO that relies only on first order gradients. The PPO objective penalizes the Kullback-Leibler (KL) divergence change between the policy before the update ($\pi_{\theta_{old}}$) and the policy at the current step (π_θ). The penalty weight β is adaptive and adjusted based on observed change in KL divergence after multiple policy update steps have been performed using the same batch of data.

$$J^{PPO}(\theta) = \hat{\mathbb{E}}_t \Big[\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)} \hat{A}_t - \beta KL\big[\pi_{\theta_{old}}(.|s_t), \pi_\theta(.|s_t)\big] \Big]$$

where $\hat{\mathbb{E}}_t[...]$ indicates the empirical average over a finite batch of samples, and $\hat{A}_t$ is the advantage estimation. Schulman et al. (2017) propose another objective based on clipping of the likelihood ratio, but we use the adaptive-KL objective due to its better empirical performance (Heess et al., 2017; Hafner et al., 2017).

2.3 Evolutionary Algorithms

There is growing interest in using evolutionary algorithms as a policy search procedure in RL. We provide a brief summary; a detailed survey is provided by Whiteson (2012). Recently, Salimans et al. (2017) proposed a version of Evolution Strategies (ES) for black-box policy optimization. At each iteration k, the algorithm samples candidate parameter vectors (policies) using a fixed covariance Gaussian $\mathcal{N}(0, \sigma^2 I)$ perturbation on the mean vector $m^{(k)}$. The mean vector is then updated in the direction of the weighted average of the perturbations, where weight is proportional to the fitness of the candidate. CMA-NeuroES (Heidrich-Meisner & Igel, 2009) uses CMA-ES to learn neural network policies for episodic reinforcement learning. CMA-ES samples candidate parameter vectors using a Gaussian $\mathcal{N}(0, C^{(k)})$ perturbation on the mean vector $m^{(k)}$. The covariance matrix and the mean vector for the next iteration are then calculated using the candidates with high fitness. Cross-Entropy methods use similar ideas and have been found to work reasonably well in simple environments (Szita & Lörincz, 2006).

Among genetic algorithm approaches to policy optimization, NEAT and its extension, HyperNEAT have enjoyed some success (Stanley & Miikkulainen, 2002a). NEAT can evolve connection weights as well as the network topology. The crossover between fixed-topology parents is done by copying the weights of each DNN edge randomly from one of the parents. The weights are mutated using random perturbations. This mode of mutation is also used by Salimans et al. (2017). A more principled take on mutating weights with perturbations can be found in (Hansen & Ostermeier, 2001; Sehnke et al., 2010). In this work, we use policy gradient algorithms for efficient mutation of high-dimensional policies, and also depart from prior work in implementing the crossover operator.

3 Genetic Policy Optimization

3.1 Overall Algorithm

Our procedure for policy optimization proceeds by evolving the policies (genotypes) through a series of selection, crossover and mutation operators (Algorithm 1). We start with an ensemble of policies initialized with random parameters. In line 3, we mutate each of the policies separately by performing a few iterations of updates on the policy parameters. Any standard policy gradient method, such as PPO or A2C, can be used for mutation. In line 4, we create a set of parents using a selection procedure guided by a fitness function. Each element of this set is a policy-pair (π_x, π_y) that is used in the reproduction (crossover) step to produce a new child policy (π_c). This is done in line 7 by mixing the policies of the parents. In line 10, we obtain the population for the next generation by collecting all the newly created children. The algorithm terminates after k rounds of optimization.

3.2 GPO Crossover and Mutation

We consider policies that are parameterized using deep neural networks of fixed architectures. If the policy is Gaussian, as is common for many robotics and locomotion tasks (Duan et al., 2016), then

Algorithm 1 Genetic Policy Optimization

population $\leftarrow \pi_1, \ldots, \pi_m$ ▷ Initial policies with random parameters
repeat
  population $\leftarrow$ MUTATE(*population*)
  parents_set $\leftarrow$ SELECT(*population*, FITNESS-FN)
  children $\leftarrow$ empty set
  for tuple$(\pi_x, \pi_y) \in$ *parents_set* **do**
    $\pi_c \leftarrow$ CROSSOVER(π_x, π_y)
    add π_c to *children*
  end for
  population $\leftarrow$ *children*
until k steps of genetic optimization

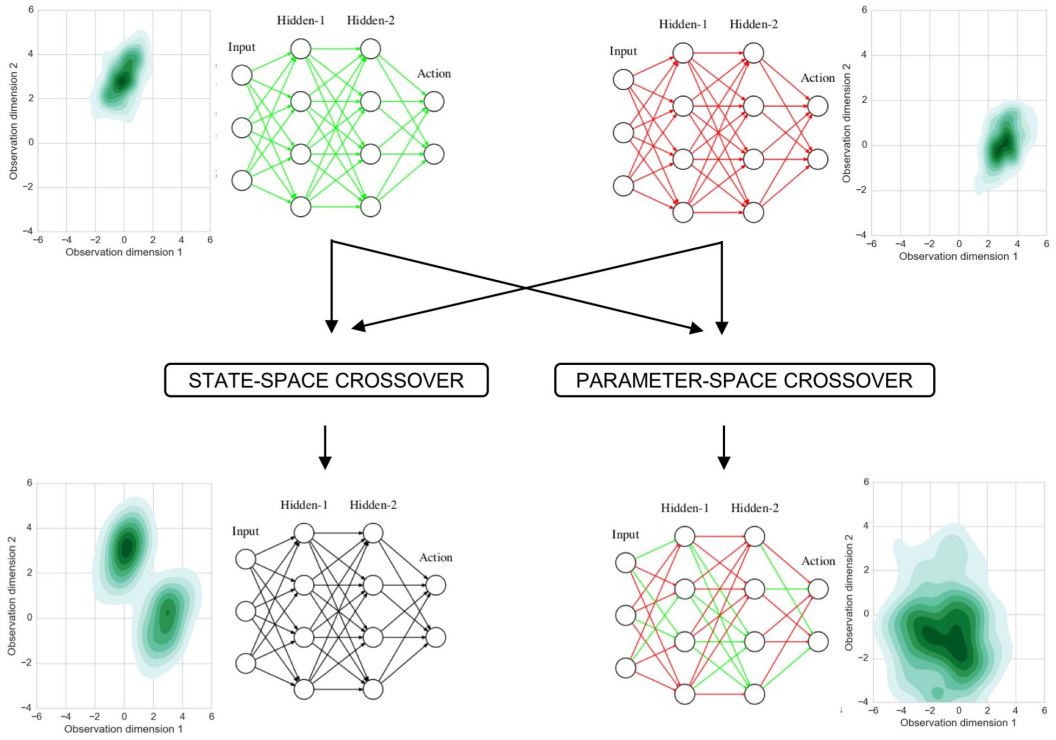

Figure 1: Crossover operator for neural network policies.

the network outputs the mean and the standard-deviation of each action in the action-space. Combining two DNN policies such that the final child policy possibly absorbs the *best traits* of both the parents is non-trivial. Figure 1 illustrates different crossover strategies. The figure includes neural network policies along with the state-visitation distribution plots (in a 2D space) corresponding to some high return rollouts using that policy. The two parent networks are shown in the top half of the figure. The state-visitation distributions are made non-overlapping to indicate that the parents policies have good state-to-action mapping for disparate local regions of the state-space.

A naïve approach is to do crossover in the parameter space (bottom-right in figure). In this approach, a DNN child policy is created by copying over certain edge weights from either of the parents. The crossover could be at the granularity of multiple DNN layers, a single layer of edges or even a single edge (e.g. NEAT(Stanley & Miikkulainen, 2002b)). However, this type of crossover is expected to yield a low-performance composition due to the complex non-linear interactions between policy-parameters and the expected policy return. For the same reason, the state-visitation distribution of the child doesn't hold any semblance to that of either of the parents. The bottom-left part of the figure shows the outcome of an ideal crossover in state-space. The state-visitation distribution of

the child includes regions from both the parents, leading to better performance (in expectation) than either of them. In this work, we propose a new crossover operator that utilizes imitation learning to combine the best traits from both parents and generate a high-performance child or offspring policy. So this crossover is not done directly in the parameter space but in the behavior or the state visitation space. We quantify the effect of these two types of crossovers in Section 4 by mixing several DNN pairs and measuring the policy performance in a simulated environment.

Our second contribution is in utilizing policy gradient algorithms for mutation of neural network weights in lieu of the Gaussian perturbations used in prior work on evolutionary algorithms for policy search. Because of the randomness in rollout samples, the policy-gradient mutation operator also maintains sufficient genetic diversity in the population. This helps our overall genetic algorithm achieve similar or higher sample efficiency compared to the state-of-the-art policy gradient methods.

3.3 Genetic Operators

This section details the three genetic operators. We use different subscripts for different policies. The corresponding parameters of the neural network are sub-scripted with the same letter (e.g. θ_x for π_x). We also use π_x and π_{θ_x} interchangeably. $\bigcup_{i=1}^{m} \pi_i$ represents an ensemble of m policies$\{\pi_1, \ldots, \pi_m\}$.

3.3.1 CROSSOVER(π_x, π_y)

This operator mixes two input policies π_x and π_y in state-space and produces a new child policy π_c. The three policies have identical network architecture. The child policy is learned using a two-step procedure. Firstly, we train a two-level policy $\pi_H(a|s) = \pi_S(h = x|s)\pi_x(a|s) + \pi_S(h = y|s)\pi_y(a|s)$ which, given an observation, first chooses between π_x and π_y, and then outputs the action of the chosen policy. π_S is a binary policy trained with the maximum likelihood objective (cross-entropy loss) on the recent rollout trajectories from both parents. This hierarchical reinforcement learning step acts a medium of knowledge transfer from the parents to the child. We use only high-reward trajectories from π_x and π_y as data samples for training π_S to avoid transfer of negative behavior. It is possible to further refine π_S by running a few iterations of any policy-gradient algorithm, but we find that the maximum likelihood approach works well in practice and can also avoid extra rollout samples. Next, to distill the information from π_H into a policy with the same architecture as the parents, we use imitation learning to train a child policy π_c. We use trajectories from π_H (expert) as supervised data and train π_c to predict the expert action under the state distribution induced by the expert. The surrogate loss for imitation learning is:

$$L^{IMIT}(\theta_c) = \mathbb{E}_{s\sim d^*}\Big[KL\big[\pi_c(.|s), \pi_H(.|s)\big]\Big] \tag{1}$$

where d^* is the state-visitation distribution induced by π_H. To avoid compounding errors due to state distribution mismatch between the expert and the student, we adopt the Dataset Aggregation (DAgger) algorithm (Ross et al., 2011). Our training dataset $\mathcal{D}$ is initialized with trajectories from the expert. After iteration i of training, we sample some trajectories from the current student ($\pi_c^{(i)}$), label the actions in these trajectories using the expert and form a dataset $\mathcal{D}_i$. Training for iteration $i+1$ then uses $\{\mathcal{D} \cup \mathcal{D}_1 \ldots \cup \mathcal{D}_i\}$ to minimize the loss. This helps to achieve a policy that performs well under its own induced state distribution. The direction of KL-divergence in Equation 1 encourages high entropy in π_c, and empirically, we found this to be marginally better than the reverse direction. For policies with Gaussian actions, the KL has a closed form and therefore the surrogate loss is easily optimized using a first order method. In experiments, we found that this crossover operator is very efficient in terms of sample complexity, which only requires a small size of rollout samples.

3.3.2 MUTATE($\bigcup_{i=1}^{m} \pi_i$)

This operator modifies (in parallel) each policy of the input policy ensemble by running some iterations of a policy gradient algorithm. The policies have different initial parameters and are updated with high-variance gradients estimated using rollout trajectories. This leads to sufficient genetic

diversity and good exploration of the state-space, especially in the initial rounds of GPO. For two popular policy gradient algorithms—PPO and A2C—the gradients for policy π_i are calculated as

$$\nabla_{\theta_i} L^{PPO}(\theta_i) = \hat{\mathbb{E}}_{i,t}\left[\frac{\nabla_{\theta_i}\pi_{\theta_i}(a_t|s_t)}{\pi_{\theta_i^{(old)}}(a_t|s_t)}\hat{A}_t - \nabla_{\theta_i}\beta KL\big[\pi_{\theta_i^{(old)}}(.|s_t), \pi_{\theta_i}(.|s_t)\big]\right] \tag{2}$$

$$\nabla_{\theta_i} L^{A2C}(\theta_i) = \hat{\mathbb{E}}_{i,t}\left[\nabla_{\theta_i}\log\pi_{\theta_i}(a_t|s_t)\hat{A}_t\right] \tag{3}$$

where $\hat{\mathbb{E}}_{i,t}[...]$ indicates the empirical average over a finite batch of samples from π_i, and $\hat{A}_t$ is the advantage. We use an MLP to model the critic baseline $V^{\pi}(s_t)$ for advantage estimation. PPO does multiple updates on the policy π_{θ_i} using the same batch of data collected using $\pi_{\theta_i^{(old)}}$, whereas A2C does only a single update.

During mutation, a policy π_i can also use data samples from other *similar* policies in the ensemble for off-policy learning. A larger data-batch (generally) leads to a better estimate of the gradient and stabilizes learning in policy gradient methods. When using data-sharing, the gradients for π_i are

$$\nabla_{\theta_i} L^{PPO}(\theta_i) = \left(\sum_{j\in\mathbb{S}_i}\hat{\mathbb{E}}_{j,t}\left[\frac{\nabla_{\theta_i}\pi_{\theta_i}(a_t|s_t)}{\pi_{\theta_j^{(old)}}(a_t|s_t)}\hat{A}_t\right]\right) - \nabla_{\theta_i}\hat{\mathbb{E}}_{i,t}\left[\beta KL\big[\pi_{\theta_i^{(old)}}(.|s_t), \pi_{\theta_i}(.|s_t)\big]\right] \tag{4}$$

$$\nabla_{\theta_i} L^{A2C}(\theta_i) = \sum_{j\in\mathbb{S}_i}\hat{\mathbb{E}}_{j,t}\left[\frac{\nabla_{\theta_i}\pi_{\theta_i}(a_t|s_t)}{\pi_{\theta_j^{(old)}}(a_t|s_t)}\hat{A}_t\right] \tag{5}$$

where $\mathbb{S}_i \equiv \{j \mid KL[\pi_i, \pi_j] < \epsilon$ before the start of current round of mutation$\}$ contains similar policies to π_i (including π_i).

3.3.3 SELECT($\bigcup_{i=1}^{m} \pi_i$, FITNESS-FN)

Given a set of m policies and a fitness function, this operator returns a set of policy-couples for use in the crossover step. From all possible $\binom{m}{2}$ couples, the ones with maximum fitness are selected. The fitness function $f(\pi_x, \pi_y)$ can be defined according two criteria, as below.

- Performance fitness as sum of expected returns of both policies, i.e. $f(\pi_x, \pi_y) \stackrel{\text{def}}{=} E_{\tau\sim\pi_x}[R(\tau)] + E_{\tau\sim\pi_y}[R(\tau)]$
- Diversity fitness as KL-divergence between policies, i.e. $f(\pi_x, \pi_y) \stackrel{\text{def}}{=} KL[\pi_x, \pi_y]$

While the first variant favors couples with high cumulative performance, the second variant explicitly encourages crossover between diverse (high KL divergence) parents. A linear combination provides a trade-off of these two measures of fitness that can vary during the genetic optimization process. In the early rounds, a relatively higher weight could be provided to KL-driven fitness to encourage exploration of the state-space. The weight could be annealed with rounds of Algorithm 1 for encouraging high-performance policies.

4 EXPERIMENTS

In this section, we conduct experiments to measure the efficacy and robustness of the proposed GPO algorithm on a set of continuous control benchmarks. We begin by describing the experimental setup and our policy representation. We then analyze the effect of our crossover operator. This is followed by learning curves for the simulated environments and comparison with baselines. We conclude the section with discussion on the quality of policies learned by GPO and scalability issues.

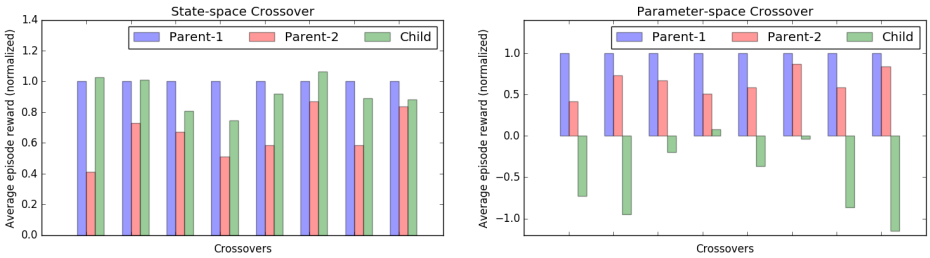

(a) Average episode reward for the child policies after state-space crossover (left) and parameter-space crossover (right), compared to the performance of the parents. All bars are normalized to the first parent in each crossover. Policies are trained on the HalfCheetah environment.

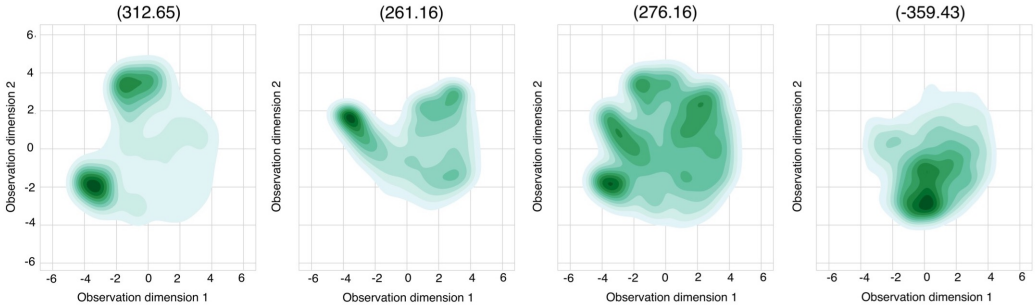

(b) State visitation distribution for high reward rollouts from policies trained on the HalfCheetah environment. From left to right - first parent, second parent, child policy using state-space crossover, child policy using parameter-space crossover. The number above each subplot is the average episode reward for 100 rollouts from the corresponding policy.

Figure 2: Comparison of Crossover operators.

4.1 Setup

All our experiments are done using the OpenAI rllab framework (Duan et al., 2016). We benchmark 9 continuous-control locomotion tasks based on the MuJoCo physics simulator [1]. All our control policies are Gaussian, with the mean parameterized by a neural network of two hidden layers (64 hidden units each), and linear units for the final output layer. The diagonal co-variance matrix is learnt as a parameter, independent of the input observation, similar to (Schulman et al., 2015; 2017). The binary policy (π_S) used for crossover has two hidden layers (32 hidden units each), followed by a softmax. The value-function baseline used for advantage estimation is also has two hidden layers (32 hidden units each). All neural networks use `tanh` as the non-linearity at the hidden units. We show results with PPO and A2C as policy gradient algorithms for mutation. PPO performs 10 steps of full-batch gradient descent on the policy parameters using the same collected batch of simulation data, while A2C does a single descent step. Other hyperparameters are in the Appendix.

4.2 Crossover Performance

To measure the efficacy of our crossover operator, we run GPO on the HalfCheetah environment, and plot the performance of all the policies involved in 8 different crossovers that occur in the first round of Algorithm 1. Figure 2a shows the average episode reward for the parent policies and their corresponding child. All bars are normalized to the first parent in each crossover. The left subplot depicts state-space crossover. We observe that in many cases, the child either maintains or improves on the better parent. This is in contrast to the right subplot where parameter-space crossover breaks the information structure contained in either of the parents to create a child with very low performance. To visualize the state-space crossover better, in Figure 2b we plot the state-visitation

[1] HalfCheetah, Walker2d, Hopper, InvertedDoublePendulum, Swimmer, Ant, HalfCheetah-Hilly, Walker2d-Hilly, Hopper-Hilly. The "Hilly" variants are more difficult versions of the original environments (`https://github.com/rll/rllab/pull/121`). We set difficulty to 1.0

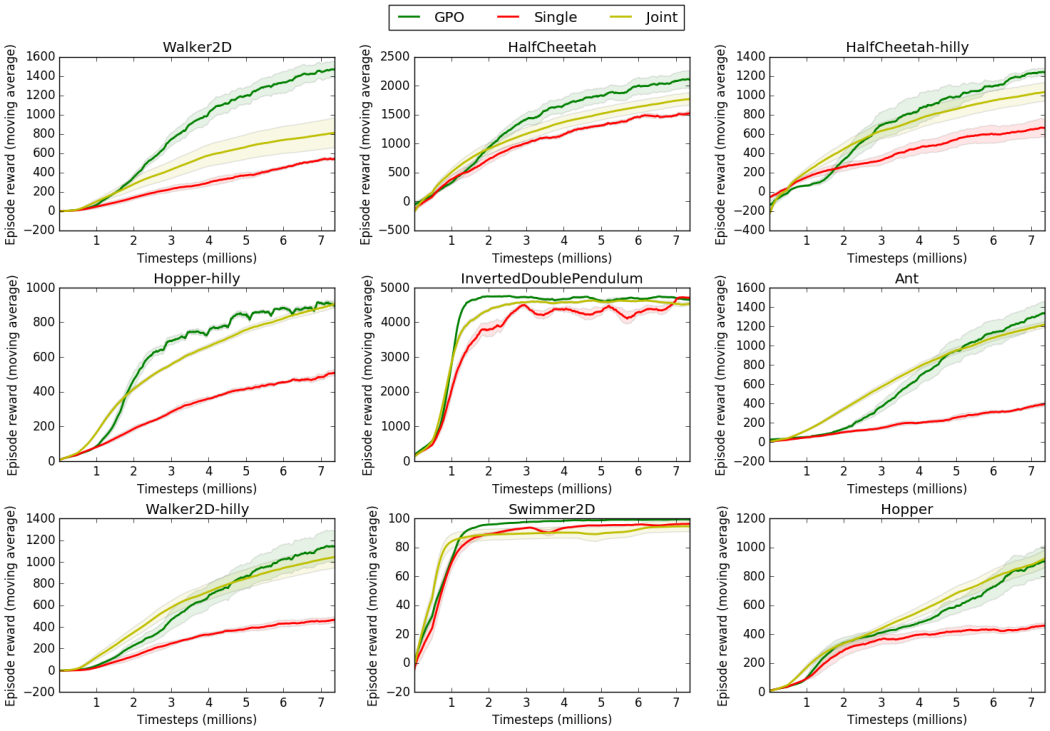

Figure 3: Performance of GPO and baselines on MuJoCo environments using PPO.

	GPO	Single	Joint
Walker2d	**1464.6** ± **93.42**	540.93 ± 13.54	809.8 ± 156.53
HalfCheetah	**2100.54** ± **151.58**	1523.52 ± 45.02	1766.11 ± 104.37
HalfCheetah-hilly	**1234.99** ± **38.72**	661.49 ± 86.58	1033.44 ± 99.15
Hopper-hilly	893.69 ± 13.81	508.62 ± 16.47	**904.87** ± **21.17**
InvertedDoublePendulum	4647.95 ± 39.69	**4705.72** ± **13.65**	4539.98 ± 37.49
Ant	**1337.75** ± **120.98**	393.74 ± 18.94	1215.16 ± 31.18
Walker2d-hilly	**1140.36** ± **146.69**	467.37 ± 24.9	1044.77 ± 98.34
Swimmer	**99.33** ± **0.14**	96.29 ± 0.14	94.55 ± 3.6
Hopper	903.16 ± 99.37	457.1 ± 16.01	**922.5** ± **61.01**

Table 1: Mean and standard-error for final performance of GPO and baselines using PPO.

distribution for high reward rollouts from all policies involved in one of the crossovers. All states are projected from a 20 dimensional space (for HalfCheetah) into a 2D space by t-SNE (Maaten & Hinton, 2008). Notwithstanding artifacts due to dimensionality reduction, we observe that high reward rollouts from the child policy obtained with state-space crossover visit regions frequented by both the parents, unlike the parameter-space crossover (rightmost subplot) where the policy mostly meanders in regions for which neither of the parents have strong supervision.

4.3 Comparison with Policy Gradient Methods

In this subsection, we compare the performance of policies trained using GPO with those trained with standard policy gradient algorithms. GPO is run for 12 rounds (Algorithm 1) with a population size of 8, and simulates 8 million timesteps in total for each environment (1 million steps per each candidate policy). We compare with two baselines which use the same amount of data. The first baseline algorithm, `Single`, trains 8 independent policies with policy gradient using 1 million timesteps each, and select the policy with the maximum performance at the end of training. Unlike GPO, these policies do not participate in state-space crossover or interact in any way. The second baseline algorithm, `Joint`, trains 1 independent policies with policy gradient using 8 million

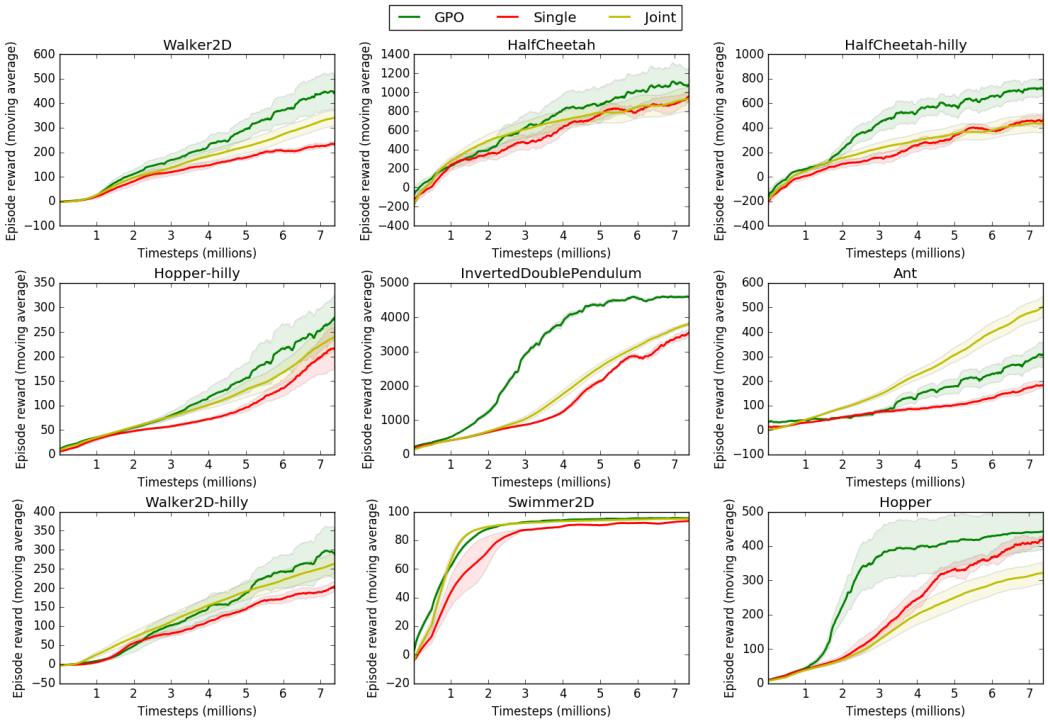

Figure 4: Performance of GPO and baselines on MuJoCo environments using A2C.

	GPO	Single	Joint
Walker2d	**444.7** $\pm$ **69.39**	233.98 $\pm$ 7.9	340.75 $\pm$ 33.59
HalfCheetah	**1071.49** $\pm$ **179.98**	956.84 $\pm$ 54.12	930.17 $\pm$ 123.83
HalfCheetah-hilly	**719.39** $\pm$ **63.74**	460.59 $\pm$ 43.2	434.48 $\pm$ 75.24
Hopper-hilly	**279.11** $\pm$ **40.28**	216.43 $\pm$ 39.78	240.26 $\pm$ 28.67
InvertedDoublePendulum	**4589.9** $\pm$ **45.37**	3545.43 $\pm$ 83.63	3802.8 $\pm$ 50.56
Ant	308.67 $\pm$ 49.63	182.61 $\pm$ 10.39	**503.64** $\pm$ **42.01**
Walker2d-hilly	**289.51** $\pm$ **56.9**	203.74 $\pm$ 12.77	263.49 $\pm$ 27.33
Swimmer	**95.43** $\pm$ **0.11**	93.43 $\pm$ 0.1	95.05 $\pm$ 0.05
Hopper	**441.79** $\pm$ **47.21**	421.06 $\pm$ 9.1	321.48 $\pm$ 30.92

Table 2: Mean and standard-error for final performance of GPO and baselines using A2C.

timesteps each, and averages the performance across policies. Both `Joint` and `Single` do the same number of gradient update steps on the policy parameters, but each gradient step in `Joint` uses 8 times the batch-size. For all methods, we replicate 8 runs with different seeds.

Figure 3 plots the moving average of per-episode reward when training with PPO as the policy gradient method for all algorithms. We observe that GPO achieves better performance than `Single` is almost all environments. `Joint` is a more challenging baseline since each gradient step uses a larger batch-size, possibly leading to well-informed, low-variance gradient estimates. Nonetheless, GPO reaches a much better score for environments such as Walker2D and HalfCheetah, and also their more difficult (hilly) versions. We think this is due to better exploration and exploitation by the nature of the genetic algorithm. The performance at the end of training is shown in Table 1. Results with A2C as the policy gradient method are in Figure 4 and Table 2. Here, GPO beats the baselines in all but one environments. In summary, these results indicate that, with the new crossover and mutation operators, genetic algorithms could be an alternative policy optimization approach that competes with the state-of-the-arts policy gradient methods.

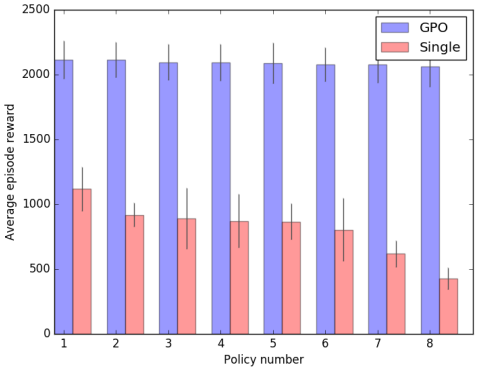

Figure 5: Final performance of the policy ensemble trained with GPO and `Single` on the HalfCheetah environment.

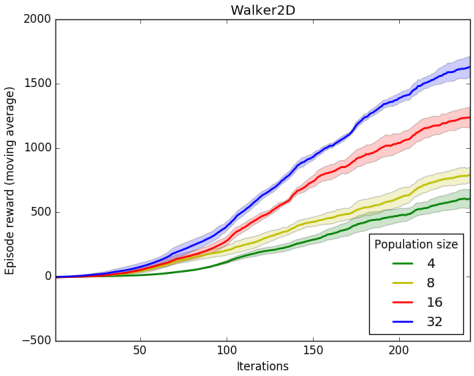

Figure 6: Scaling by increasing the GPO population size in Walker2d environment.

4.4 Robustness and Scalability

The SELECTION operator selects high-performing individuals for crossover in every round of Algorithm 1. Natural selection weeds out poorly-performing policies during the optimization process. In Figure 5, we measure the average episode reward for each of the policies in the ensemble at the final round of GPO. We compare this with the final performance of the 8 policies trained using the `Single` baseline. We conclude that the GPO policies are more robust. In Figure 6, we experiment with varying the population size for GPO. All the policies in this experiment use the same batch-size for the gradient steps and do the same number of gradient steps. Performance improves by increasing the population size suggesting that GPO is a scalable optimization procedure. Moreover, the MUTATE and CROSSOVER genetic operators lend themselves perfectly to multiprocessor parallelism.

5 Conclusion

We presented Genetic Policy Optimization (GPO), a new approach to deep policy optimization which combines ideas from evolutionary algorithms and reinforcement learning. First, GPO does efficient policy crossover in state space using imitation learning. Our experiments show the benefits of crossover in state-space over parameter-space for deep neural network policies. Second, GPO mutates the policy weights by using advanced policy gradient algorithms instead of random perturbations. We conjecture that the noisy gradient estimates used by policy gradient methods offer sufficient genetic diversity, while providing a strong learning signal. Our experiments on several MuJoCo locomotion tasks show that GPO has superior performance over the state-of-the-art policy gradient methods and achieves comparable or higher sample efficiency. Future advances in policy gradient methods and imitation learning will also likely improve the performance of GPO for challenging RL tasks.

References

Yan Duan, Xi Chen, Rein Houthooft, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pp. 1329–1338, 2016.

Danijar Hafner, James Davidson, and Vincent Vanhoucke. Tensorflow agents: Efficient batched reinforcement learning in tensorflow. *arXiv preprint arXiv:1709.02878*, 2017.

Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.

Nicolas Heess, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, Ali Eslami, Martin Riedmiller, et al. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.

Verena Heidrich-Meisner and Christian Igel. Neuroevolution strategies for episodic reinforcement learning. *Journal of Algorithms*, 64(4):152–168, 2009.

Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.

Stéphane Ross, Geoffrey J Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.

Tim Salimans, Jonathan Ho, Xi Chen, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 1889–1897, 2015.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559, 2010.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

Kenneth O Stanley and Risto Miikkulainen. Efficient reinforcement learning through evolving neural network topologies. In *Proceedings of the 4th Annual Conference on Genetic and Evolutionary Computation*, pp. 569–577. Morgan Kaufmann Publishers Inc., 2002a.

Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002b.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1.

István Szita and András Lörincz. Learning tetris using the noisy cross-entropy method. *Learning*, 18(12), 2006.

Shimon Whiteson. Evolutionary computation for reinforcement learning. In *Reinforcement learning*, pp. 325–355. Springer, 2012.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 3357–3364. IEEE, 2017.

6 APPENDIX

6.1 HYPERPARAMETERS

Hyperparameter	Value
Horizon (T)	512
Adam stepsize	5 x 10^{-4}
Discount (γ)	0.99
GAE parameter (λ)	1
PPO epochs	10
Batchsize (GPO, Single)	2048
Batchsize (Joint)	16384