PGPS : Coupling Policy Gradient with Population-based Search

Anonymous Author(s)
Affiliation
Address
email

Abstract

Gradient-based policy search algorithms (such as PPO, SAC, or TD3) in deep reinforcement learning (DRL) have shown successful results on a range of challenging control tasks. However, they often suffer from deceptive gradient problems in flat or gentle regions of the objective function. As an alternative to policy gradient methods, population-based evolutionary approaches have been applied to DRL. While population-based search algorithms show more robust learning in a broader range of tasks, they are usually inefficient in the use of samples. Recently, reported are a few attempts (such as CEMRL) to combine gradient with a population in searching optimal policy. This kind of hybrid algorithm takes advantage of both camps. In this paper, we propose yet another hybrid algorithm, which more tightly couples policy gradient with the population-based search. More specifically, we use Cross Entropy Method (CEM) for population-based search and Twin Delayed Deep Deterministic Policy Gradient (TD3) for policy gradient. In the proposed algorithm called Coupling Policy Gradient with Population-based Search (PGPS), a single TD3 agent, which learns by a gradient from all experiences generated by population, leads a population by providing its critic function Q as a surrogate to select better performing next generation population from candidates. On the other hand, if the TD3 agent falls behind the CEM population, then the TD3 agent is updated toward the elite member of CEM population using loss function augmented with the distance between the TD3 and the CEM elite. Experiments in five challenging control tasks in a MuJoCo environment show that PGPS is robust to deceptive gradient and also outperforms the state-of-the-art algorithms.

1 Introduction

In Reinforcement Learning (RL), an agent interacts with the environment, and its goal is to find the policy that maximizes the objective function, which is generally defined as a cumulative discounted reward. Recently, many researchers have worked on combining deep neural networks and a gradient-based RL algorithm, generally known as Deep Reinforcement Learning (DRL). This approach has achieved great success not only in the discrete action domain, such as in Go [32] and Atari games [23, 22], but also in the continuous action domain, such as in Robot control [7, 20, 29]. However, it is difficult to use the gradient-based method for the objective function (J), which includes "many wide flat regions" since the gradient ($\nabla_\theta J$) is near zero at a flat point. Figure 1 is an extreme case consisting of only flat regions, which is called a piece-wise constant

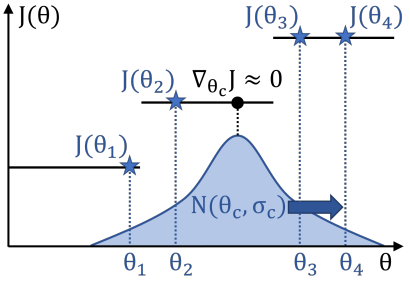

Figure 1: Flat gradient and population-based search on piece-wise constant function

Submitted to 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Do not distribute.

function. This problem remains an unsolved issue in gradient-based DRL with continuous control domains [4]. The Swimmer in a MuJoCo environment [36] has already been reported to be hard to use the gradient-based method [15, 17, 21, 26]. Our experiment shows that the objective function of Swimmer includes wide flat regions (Appendix A).

The population-based Evolutionary Approach (EA), which is an alternative to the gradient-based method, has also shown successful results in various control tasks [5, 21, 27, 34]. As a population-based search, the EA generates a population of agents to explore policy, and the population is regenerated with improvement in each generation. The EA is also known as a direct policy search [28] because it directly searches by perturbing the parameter of policy. In Figure 1, the Cross-Entropy Method (CEM) as a kind of population-based search is simply described, where the current population sampled from the target distribution is evaluated. Then the distribution is updated to the direction for generating a more promising population. Not depending on the gradient, these approaches are robust to flat or deceptive gradients [33, 21]. However, the EA is sample inefficient because it requires a Monte-Carlo evaluation, and the previous results and data generally cannot be reused.

The off-policy Policy Gradient (PG) algorithm uses the data from arbitrary policies to train its actor and critic functions. It generates exciting potential by combining the EA and PG, where the data which is discarded in a standard EA is directly used to train the PG's functions. Khadka and Tumer [17] and Pourchot and Sigaud [26] introduced a framework combining the EA and off-policy PG. However, the framework of [17] is less efficient to train the policy for general tasks than the PG algorithm alone, and the framework of [26] is unsuitable to train the policy for a task providing a deceptive gradient.

In this paper, we propose another hybrid algorithm, called Policy Gradient with Population-based Search (PGPS) in which the CEM and Twin Delayed Deep Deterministic Policy Gradient (TD3) [7] are combined. It is as robust to a deceptive gradient as the CEM and more efficient to train the policy for general tasks than TD3. To be robust to a deceptive gradient, the proposed algorithm is constructed in a way similar to the one in [18], where the TD3 is trained using data from the CEM and periodically participates in the CEM population as an individual (PG guides EA). However, in this basic framework, the TD3 sometimes falls into the inferior solution and inefficiently searches. To get the TD3 out of the inferior solution, we let the EA guide the TD3 by guided policy learning [16] (EA guides PG). Furthermore, the TD3 critic contributes to generating a more promising population by filtering the set of actors sampled from CEM (Q-critic filtering). Lastly, to control the trade-off between the frequency of search and stable estimation, we used evaluation step scheduling in the process of population evaluation (Increasing evaluation steps). It carries out frequent searches when searching far from the optimal, whereas it carries out stable estimation when searching close to the optimal. These approaches bring out more synergies between the CEM and the TD3 while maintaining both the population-based search and the gradient-based search. Consequently, the proposed algorithm is not only robust to a deceptive gradient, but also produces outstanding performances with a low additional computational cost.

2 Related Works

Recently, beyond the view of an alternative approach, few attempts have been proposed in the form of A supporting B. An attempt is to use EA to fill a replay buffer with diverse samples. In Colas et al. [4], a Goal Exploration Process (GEP), a kind of EA, is firstly applied to search the policy and to fill a replay buffer with the diverse samples, and then the off-policy PG algorithm is sequentially used for fine tuning the parameters of the policy. Another attempt is to combine a population-based approach and PG for efficiently searching a good policy or the good hyper-parameters of an algorithm in parallel multi-learners setting. These applications generally consist of periodically evaluating the population, followed by distributing good knowledge to the other learners. To find the best architecture and hyper-parameters, Jaderberg et al. [13] proposed a Population-Based Training (PBT) method in which the current best knowledge is periodically transferred to PG learners. Gangwani and Peng [8] developed the distilled crossover using imitation learning and mutation based on the PG. Proposed operators transfer the information on current good policies into the next population without destructive change to the neural network. Jung et al. [15] introduced a soft-manner guided policy learning to fuse the knowledge of the best policy with other identical multiple learners while maintaining a more extensive search area for the exploration.

The idea of combining the population-based EA and off-policy PG was recently introduced by Khadka and Tumer [17]. Their approach was called Evolutionary-Guided Reinforcement Learning (ERL) in which the Genetic Algorithm (GA) and the Deep Deterministic Policy Gradient (DDPG) [20] are combined. In ERL frameworks, the GA transfers the experience from evaluation into the DDPG through a replay buffer, and the DDPG transfers the knowledge learned from the policy gradient into the GA by periodically injecting a PG actor into the GA population. Khadka et al. [16] expanded the PG algorithm of ERL from a single DDPG learner to multiple TD3 learners with a resource manager. Bodnar et al. [2] revised the GA's crossover and mutation to the distilled crossover and proximal mutation inspired by [8] and [19] to prevent the destruction of neural networks. Pourchot and Sigaud [26] introduced another framework, which combines the CEM and the TD3. In this framework, the TD3 algorithm has only a critic function trained using the experience from the CEM. In order to propagate the knowledge learned by policy gradient to the CEM, half of the population is updated to the direction indicated by the TD3 critic for a fixed number of steps, followed by the evaluation. The policy gradient for half of the population not only enhances the gradient-based learning, but also deteriorate the CEM's robustness over a deceptive gradient.

In this paper, we introduce another hybrid algorithm, in which the CEM and the TD3 are combined as in CEMRL [26]. However, the TD3 has both actor and critic, which are trained by a gradient from experiences generated by CEM. And then, the TD3 actor periodically participates in CEM population as in ERL [17]. This structure is an effective way to maintain a direct policy search of CEM. To enhance the performance, we introduced new interactions processes between the CEM and TD3 instead of carrying out a policy gradient for numerous individual actors.

3 Backgrounds

Twin Delayed Deep Deterministic Policy Gradient (TD3) RL framework has an agent interacting with an environment generally defined by a Markov Decision Process (MDP). At each timestep t, an agent receives the state s_t, and takes an action a_t according to the policy π, and then receives a reward r_t and the next state s_{t+1} at next time step $t+1$. The goal of RL is to find the policy that maximizes the discounted cumulative return $R_t = \sum_{k=t}^{\infty} \gamma^{k-t} r_k$ where γ is a discounted factor. Off-policy RL can use the data from arbitrary policies to train its actor and critic functions repeatedly, which is a key point for improving recent gradient-based RL. Silver et al. [31] introduced the off-policy Deterministic Policy Gradient (DPG), which has an advantage for high-dimensional action spaces. The DDPG [20] was extended from the DPG to apply it to a deep neural network. TD3 [7] is an advanced version of the DDPG, which suffers from the overestimation bias of the critic. To correct this bias, two critics are introduced, and then the critic with the lowest state-action value is taken during the critic update as in the Double Deep Q-Network (DDQN) [37]. Figure 2(a) represents the architecture of the TD3.

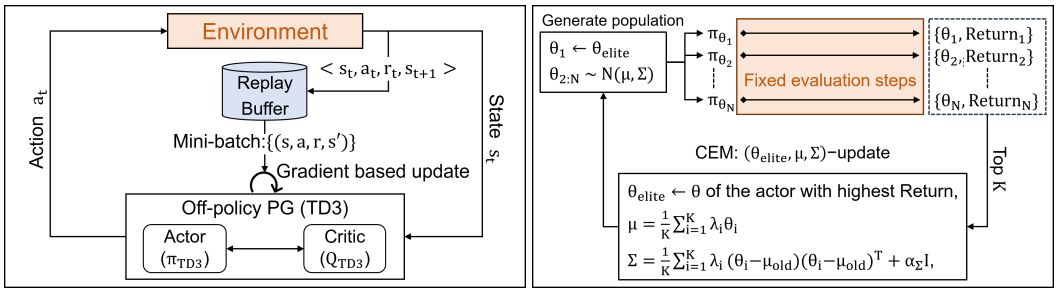

(a) TD3 (b) CEM

Figure 2: Architecture of TD3 and CEM. θ_i is sampled from $N(\mu, \Sigma)$, λ_i is a weight depending on the rank of the $Return_i$, and K is a fixed number about high performing actors.

Cross Entropy Method (CEM) The Evolutionary Approach (EA) is a heuristic search method inspired by nature, where the current population is evaluated, and the next population is regenerated using the current evaluation result in order to produce a higher Return, which is also known as Fitness and defined as a cumulative sum of immediate reward for a fixed number of steps. The Estimation of Distribution Algorithm (EDA) is a class of the EA: It updates the target distribution to generate a better population. Depending on the update method for the distribution, EDAs are classified as a CEM [6], a Covariance Matrix Adaptation Evolutionary Strategy [10], an Evolutionary Strategy [27],

and a Natural Evolutionary Strategy [38]. We used the CEM as one part of our proposed algorithm. As shown in Figure 2(b), the CEM procedures are as follows: The population is sampled from the multivariate Gaussian $N(\mu, \Sigma)$ and evaluated, and for the top K, which is smaller than the population size (N), high performing actors are used to update a new mean (μ) and covariance (Σ) of the target distribution. The weight can be given to each actor according to the rank of the Return [10, 26]. The elite actor can be passed to the next population, which is known as elitism. The more detailed procedure is reported in [6]. In this paper, we used a diagonal covariance to reduce the parameters.

Population-guided Parallel Policy Search (P3S) The Guided Policy Learning (GPL) is commonly used when the elite policy leads some sub-policies to a better solution in multiple policies setting. Teh et al. [35] introduced GPL for joint learning of numerous tasks in which a common policy encourages local policies to act better. Jung et al. [15] proposed a soft-manner GPL, called the Population-guided Parallel Policy Search, for multiple identical learners with the same objective, where a population is evaluated periodically. Then sub-policies are trained to maximize their critic value and to concurrently minimize the distance from the elite policy for the next period. For this purpose, Augmented Loss (2) is used to train the sub-policies instead of Original Loss (1).

$$Original\ Loss: \quad L^O(\pi) = E_{s\sim SS}[-Q_\pi(s, \pi(s))] \tag{1}$$

$$Augmented\ Loss: L^A(\pi, \pi_{elite}, \beta) = E_{s\sim SS}[-Q_\pi(s, \pi(s)) + \beta||\pi(s) - \pi_{elite}(s)||_2^2] \tag{2}$$

where π is a trained policy, Q_π is a critic function depending on π, π_{elite} is the elite policy, SS is the set of states, and $||\pi(s) - \pi_{elite}(s)||_2^2$ is the Euclidean distance measure between the trained policy and the elite policy. β is a distance weight and is controlled adaptively. In this paper, we used a revised GPL inspired by P3S so that the CEM elite actor guides the TD3 to better space.

4 Coupling Policy Gradient with Population based Search Algorithm

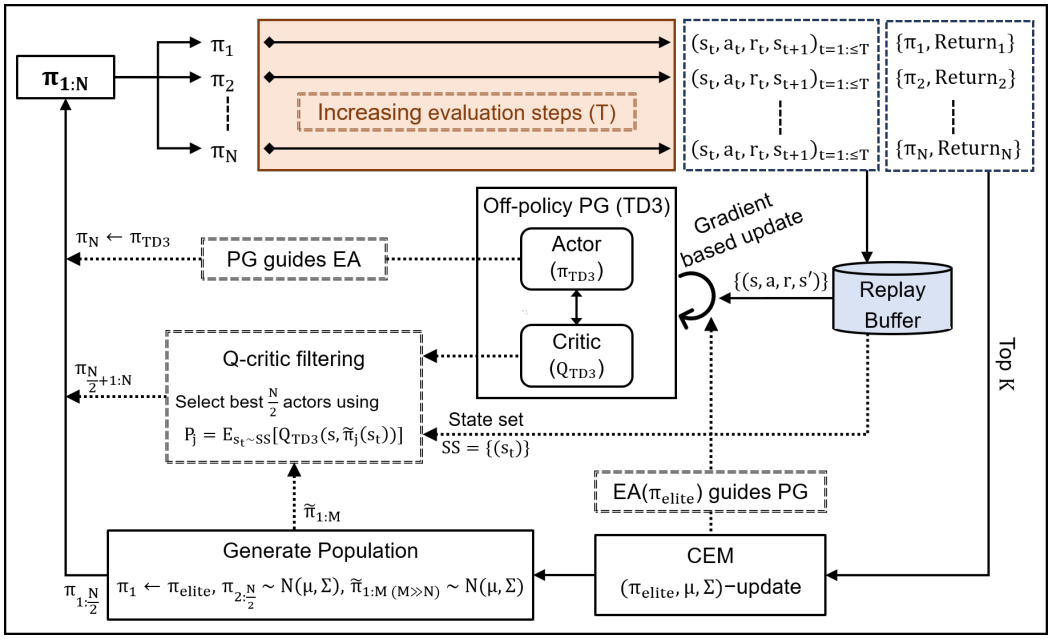

Figure 3: Architecture of PGPS. π is used instead of π_θ for the simplification.

As shown in Figure 3, the Policy Gradient with Population-based Search (PGPS) is a coupling framework between the CEM and the TD3. In this framework, two algorithms encourage each other to be efficient. The general flow is as follows. The generated actor population is evaluated by interacting with the environment, where each actor can fail to reach the max step (T). The experience (the set of state transitions) is saved in a replay buffer, and Returns go to the CEM. The CEM updates the parameters of the population using top K (K is set to half of the population) high performing actors, and then the TD3 trains its critic and actor using the mini-batches sampled from the replay buffer. The knowledge of TD3 is periodically transferred to the CEM by copying the TD3 actor

into the last actor of the CEM population. The knowledge of the CEM is transferred to the TD3 by the GPL when the TD3 is presumed to falls into the inferior solution. Lastly, the next population is regenerated, where the TD3 critic is used to selects promising actors among the set sampled from $N(\mu, \Sigma)$ so that the population will be better. A pseudocode of the PGPS is described in Appendix B.

Gradient-based Update In a standard CEM, the experience of state transitions is discarded immediately because they require only Returns to update the target distribution. However, the TD3 enables the discarded experience to be reused to train its actor and critic functions. Therefore, the experience is saved in the replay buffer and then is repeatedly used in a gradient-based update.

PG Guides EA In order to transfer the knowledge learned in TD3 to the CEM, the TD3 actor is periodically copied to the last actor of the population. If the last actor is included in the top-performing K actors, a multivariate Gaussian move to the direction indicated by TD3. On the other hand, if the TD3 actor is excluded from the top-performing K actors, the CEM ignores the knowledge from the TD3 and focuses on the direct policy search. High copying frequency benefits knowledge propagation from the TD3 to the CEM, but it can disturb the population-based search.

EA Guides PG When the TD3 falls into the inferior solution, mainly due to a deceptive gradient, it is difficult to get out by relying solely on the gradient-based method despite the experience from good behavior policy [4]. Therefore, we use the Guided Policy Learning (GPL), where the CEM elite actor leads the TD3 to escape the inferior solution. We judge the TD3 to have fallen into the inferior solution if its actor shows a lower Return than ***(mean − one standard deviation)*** of Returns of the current population. The TD3 actor that has fallen into the inferior solution is updated for the direction to minimize the Augmented Loss L^A, equation (2). Moreover, the TD3 critic is trained through the target actor indirectly guided by the elite actor, and it will fix the critic to be appropriate. The distance weight (β) in the in L^A is adapted several times during the GPL-based update by equation (3). It is a simplified version of P3S [15] and similar to Adaptive TRPO, which was introduced in [30].

$$\beta = \begin{cases} \beta \ \times \ 2 & \text{if } D(\pi_{TD3}, \pi_{elite}) > D_{target} \ \times \ 1.5 \\ \beta \ / \ 2 & \text{if } D(\pi_{TD3}, \pi_{elite}) < D_{target} \ / \ 1.5 \end{cases} \tag{3}$$

where the distance measure $D(\pi_{TD3}, \pi_{elite})$ is defined as $E_{s\sim SS}[||\pi_{TD3}(s) - \pi_{elite}(s)||_2^2]$, SS is the set of states in the replay buffer, and D_{target} as a hyper-parameter determines how close the TD3 actor and the elite actor should be. During the GPL-based update, the TD3 actor stays around the CEM elite actor while maximizing its critic value.

Q-critic Filtering It is important to generate a good population, since it not only leads $N(\mu, \Sigma)$ to be better but also encourages the TD3 to be trained well by filling the replay buffer with good experience. However, one cannot know which actor is better before the evaluation. To estimate the potential in advance, we use the Q-critic as a surrogate model [14]. It can sort out promising actors from the set of actors sampled from $N(\mu, \Sigma)$ before the evaluation.

Proposition1 If $E_{a\sim\pi_i(\cdot|s)}[Q_{TD3}(s,a)] \geq E_{a\sim\pi_{TD3}(\cdot|s)}[Q_{TD3}(s,a)] \ for \ all \ s,$
$\mathrm{E}_{a\sim\pi_i(\cdot|s)}[Q_{\pi_i}(s,a)] \geq E_{a\sim\pi_{TD3}(\cdot|s)}[Q_{TD3}(s,a)]$

Specific proof is described in Appendix C. We assume that a higher $E_{a\sim\pi_i(\cdot|s)}[Q_{TD3}(s,a)]$ indicates a higher $E_{a\sim\pi_i(\cdot|s)}[Q_{\pi_i}(s,a)]$. In this assumption, $E_{s\sim SS}[Q_{TD3}(s,\pi_i(s))]$ is used as the relative potential of the actor i (π_i) where SS is the set of states in the replay buffer. The overall procedure is as follows: M ($\gg$N) actors are sampled from $N(\mu, \Sigma)$, and then Q-critic fills half of the population with the actors with higher relative potential by filtering M actors. The remaining half consists of the elite actor and the actors sampled from $N(\mu, \Sigma)$ for the exploration.

Increasing Interaction Steps In order to efficiently control the trade-off between the frequency of searches and stable performance estimation of actors, we used evaluation steps scheduling, where the evaluation steps between actor and environment increase with cumulative evaluation step by equation (4).

$$T = min(T_{init} + 100 \times mod(cumulative\ evaluation\ steps, T_{inter}), T_{max}) \tag{4}$$

where T is the current evaluation step, which means that each actor can maximally interact with the environment as much as T. T_{init} is the initial evaluation step, T_{inter} is the interval for increasing evaluation step, and T_{max} is the maximum evaluation step depending on the task.

The evaluation step should be sufficiently long for stable performance estimation of the population, but it reduces the number of the CEM generation and delays the update of the TD3. Short evaluation steps make it possible to carry out more population-based searches and frequent TD3 updates but causes unstable performance estimation. As the arbitration approach, we used the increasing evaluation steps. This approach guarantees more searches and frequent updates at the beginning stage of learning and fine estimation at the later stage of learning.

5 Experiments

5.1 Comparison to Baselines

Environment The proposed algorithm (PGPS) was empirically evaluated on five games in Mujoco managed by OpenAI Gym [3], which is widely used as continuous control benchmarks [2, 7, 17, 26].

Baseline Algorithms We compared the performance of PGPS with various EA, PG, and EAPG algorithms, such as CEM, PPO[30], DDPG[20], TD3[7], SAC[9], CERL[16], PDERL[2], and CEMRL[26]. To reduce the reproductivity problem proposed by [11, 12], most algorithms were implemented using the code provided by its authors. However, the DDPG is implemented by the code provided by the authors of the TD3. OpenAI Spinningup [1] was used to implement PPO. PGPS and CEM were implemented using PyTorch [24]. Our code is available at http://github.com/NamKim88/PGPS.

Hyper-parameter Setting For stable learning, we performed tuning on the architecture of the neural network and the hyper-parameters of the learning rate, population size, the period of copying TD3 actors to the population, increasing evaluation steps, and Q-critic filtering. The detailed setting is described in Appendix D. Adam optimizer [18] was used to train the neural networks in the TD3. The hyper-parameters of baseline algorithms were set as the same as the reference code.

Performance Evaluation Metrics Each algorithm learns five times at different seeds for a task. Each learning runs for a million timesteps. That is, the total numbers of interacting all agents with the environment are a million. The evaluation tests are performed every 10,000 steps. Each evaluation test performed without any exploration behavior and reports the average reward over ten episodes, where the evaluation step of a episode is 1,000. For the evaluation test of the PGPS, the current mean (μ) of $N(\mu, \Sigma)$ was used as the parameters of the evaluation policy. If the learning is performed at seed 1, the evaluation test will proceed at random seed 1 + 100. This approach is the same as the original code of the TD3 [7] and applied to all baseline algorithms. The curves in Figure 4 reported over the average performance of policies trained at five random seeds from 1 to 5.

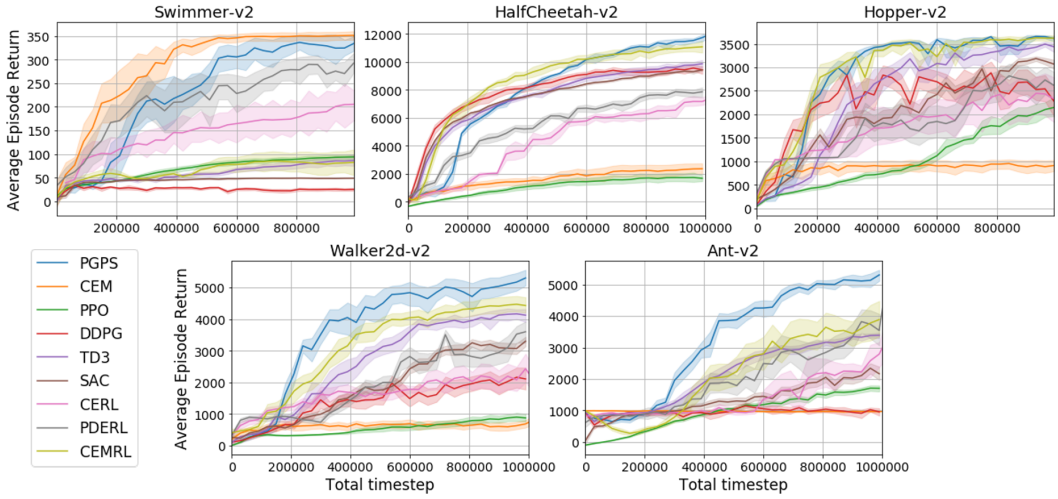

Figure 4: Learning curves obtained on 5 continuous control tasks.

Results In Figure 4 the performance of all baseline algorithms is similar to that reported in the original author's papers and reference papers. Some differences come from the training seed, evaluation metrics, the variance of algorithms, and the version of MuJoCo and Pytorch.

The results show that all PG algorithms suffer from a deceptive gradient in Swimmer-v2. In contrast, the CEM, which is a direct policy search method, yields the best result over a deceptive gradient in Swimmer-v2. The CERL and PDERL as a combining algorithm between GA and PG show better performance than PG algorithms. However, their performances are lower than the CEM because the knowledge propagated from the PG to the GA disturbs the direct policy search of the GA. Despite the CEMRL combining algorithm between the CEM and TD3, it shows similar performance to the TD3. This result comes from that the gradient-based update of half of the population ruins the population-based search of the CEM.

In the remaining four tasks which are advantageous to the gradient-based method, advanced PG algorithms show much better performance than the CEM. The performances of CERL and PDERL are located between an advanced PG algorithm and the CEM in most cases. Especially, CERL yields lower performance than the TD3, which is one part of CERL. This is due to the failure of GA and TD3 to combine effectively. CEMRL outperforms all baseline algorithms in four tasks since the gradient-based update for multiple actors amplifies the advantage of the gradient method.

PGPS carried out gradient-based update on the TD3, and then the TD3 actor is periodically copied to the last actor of the CEM population. It keeps the computational cost lower and also minimizes the disturbance from the gradient method on a direct policy search. As a result, the PGPS can achieve results comparable to the CEM in Swimmer-v2 which is advantageous to a direct policy search. Furthermore, the PGPS shows an outstanding performance in the remaining four tasks, which are advantageous to the gradient-based method. That performance is due to the additional interaction processes for coupling two algorithms efficiently, such as mutual guidance (PG-EA), Q-critic filtering, and increasing evaluation steps.

5.2 Ablations Studies

In this subsection, we performed ablations studies to investigate the effect on the final performance and computational time when a specific process is cumulatively added to the base model. The added sequence is as follows: PG guides EA (P to E), EA guides PG (E to P), Q-filtering, and increasing evaluation steps. HalfCheetah-v2, Walker2d-v2, and Ant-v2 are used for the ablation studies. Each task was learned five times at random seeds of 1 to 5. Each learning runs for a million timesteps. Table 1 reports the average over 15 (three games$\times$five learning) runs. Evaluation test is performed using the μ of the CEM except for the Base model.

Base model The experience from CEM is saved to a replay buffer. However, the TD3 and CEM are trained independently without any guidance. Two evaluation tests are performed using the TD3 and the μ of the CEM at the end of learning. A better one is selected for the final performance.

Table 1: Ablation studies on 3 continuous control tasks.

	Model	Average normalized score	Normalized computational time				
			Sum	TD3 update	Actors evaluation	Population generation	CEM update
(1)	Base model	1.00	1	0.934	0.065	0.001	≈ 0
(2)	(1) + P to E	3.73	1	0.934	0.065	0.001	≈ 0
(3)	(2) + E to P	4.26	1.107	1.041	0.065	0.001	≈ 0
(4)	(3) + Q-filtering	4.44	1.134	1.041	0.065	0.028	≈ 0
(5)	(4) + Increasing	5.26	1.134	1.041	0.065	0.028	≈ 0

P to E shows the most noteworthy performance improvement within the proposed algorithm. A good experience is essential to train TD3's actor and critic functions well. As the TD3 actor directly guides the population to be better, the population fills the replay buffer with good experience, and then it makes sure the TD3 will be trained well again. For the perspective of the CEM, the TD3 actor is a good exploration, which cannot be developed from a random sampling, and improves the population to be better. For the perspective of the TD3, the experience from the population is a richer set of exploration behavior, which is better than PG-based action space exploration [25]. The

other processes also additionally contribute to performance improvement of about $4 \sim 18\%$ The remarkable point is that these processes incur a low additional computational cost. The Q-filtering requires a relatively highest computational cost, but it is only about 11% of that of the TD3 alone.

The effect of EA guidance In contrast to the existing hybrid algorithm [2, 16, 17] in which the PG only guides the EA, the proposed algorithm lets the EA also guides the TD3. It is beneficial to pull the TD3 agent out of the inferior solution, especially if it is captured due to a deceptive gradient. Figure 5 represents the effect of EA guidance (EA guides PG) in swimmer-v2. As shown in Figure 5, the TD3 actor without EA guidance stays on the inferior solution most of the time, whereas the TD3 actor with EA guidance quickly gets out of the inferior solution and goes on searching for a better space. The proposed algorithm will also perform robust learning even if both deceptive and general gradient occurs in the environment since It lets CEM lead TD3 when a deceptive gradient occurs and lets TD3 lead CEM when the general gradient occurs.

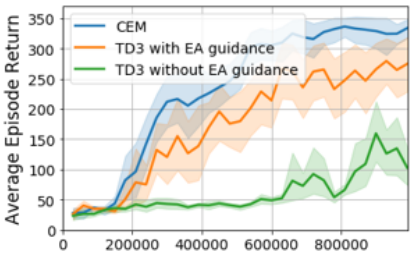

Figure 5: Learning curves of TD3 with(out) EA guidance on Swimmer-v2

Figure 6: Average Return of filtered actors and sampled actors during running

The effect of Q-critic filtering The proposed algorithm selects half of the population using Q-critic (Q-filter), and the remaining half is sampled from a multivariate Gaussian (Random). Figure 6 shows the average Return of Q-filter and Random in HalfCheetah-v2. Q-filter shows a 7.7% higher average performance than Random. (Because the evaluation steps were controlled, the average Return in Figure 6 is lower than the test episode return in Figure 4.) This result is additional evidence that Q-critic filtering is meaningful for sorting out promising actors. The remaining concern is the effect of Q-critic filtering on the exploration. If Q-critic selects similar actors, it might deteriorate the exploration behavior for both the EA and the PG. To further improve exploration, the distance-based criteria introduced in [2] can also be used with Q-critic filtering.

6 Conclusions

In this paper, we proposed another hybrid algorithm coupling the Cross-Entropy Method (CEM) and the Twin Delayed Deep Deterministic Policy Gradient (TD3), which is called Policy Gradient with Population-based Search (PGPS). Proposed algorithm is not only robust to a deceptive gradient, which is difficult to be learned by the TD3 alone, but it also achieves higher sample efficiency, which is deficient in the EA alone. To be robust to a deceptive gradient with low additional computational cost, we revised the existing hybrid algorithm framework to make it an improved structure. Also, to enhance the performance, we introduced new interaction processes such as mutual guidance (PG$\leftrightarrow$EA), Q-critic filtering, and increasing evaluation step. First, mutual guidance is the most crucial process, where the TD3 guides the CEM to make it better, and the CEM also guides the TD3 that has fallen into the inferior solution to make it better. Second, Q-critic helps the population to consist of more promising actors by filtering the set of actors sampled from a multivariate Gaussian. Lastly, the increasing evaluation step controls the trade-off between the frequency of searches and stable estimation. It takes the frequent searches and updates when searching on a coarse policy space far from the optimal at the beginning stage of learning and fine estimation close to the optimal at the later stage of learning. Our experiments on a MuJoCo confirmed that the proposed algorithm outperforms state-of-the-art PG and hybrid algorithms, while its computational cost is kept at about 13.5% higher than the TD3 algorithm.

Broader Impacts

Reinforcement learning (RL) is one of the areas that are currently in the spotlight in the field of artificial intelligence, but the degree to which it is being used in the real industry and public life is

relatively low compared to image processing and natural language processing. We think the reason lies in the fact that most RL algorithms do not directly handle the fitness value for the whole process and the unstable learning ability. For example, existing RL algorithms proposed learning through TD-error and most of them show different result values each time they learn. However, our proposed algorithm handles the fitness value for the whole process in the learning process and has a robust learning ability. Our final goal is to develop an RL algorithm applicable level in the real industry. The first target is learning a job shop problem in the manufacturing sector: not one step but the whole process throughput must be considered, and the stability of the algorithm is very important. Through this, we hope the positive impact of our work that we could contribute to the actual use of reinforcement learning in industrial sites and public life. Conversely, there can be a possible negative impact is that the automation of robots can cause conflicts of interest, such as "Robots vs Human Labor". However, if we develop RL technology in the direction of human and robot collaboration, we will be able to solve these conflicts.

References

[1] Achiam, J. (2018). Openai spinning up. *GitHub, GitHub repository*.

[2] Bodnar, C., Day, B., and Lió, P. (2019). Proximal distilled evolutionary reinforcement learning. *arXiv preprint arXiv:1906.09807*.

[3] Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym. *arXiv preprint arXiv:1606.01540*.

[4] Colas, C., Sigaud, O., and Oudeyer, P.-Y. (2018). Gep-pg: Decoupling exploration and exploitation in deep reinforcement learning algorithms. *arXiv preprint arXiv:1802.05054*.

[5] Conti, E., Madhavan, V., Such, F. P., Lehman, J., Stanley, K., and Clune, J. (2018). Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. In *Advances in neural information processing systems*, pages 5027–5038.

[6] De Boer, P.-T., Kroese, D. P., Mannor, S., and Rubinstein, R. Y. (2005). A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67.

[7] Fujimoto, S., Van Hoof, H., and Meger, D. (2018). Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*.

[8] Gangwani, T. and Peng, J. (2017). Policy optimization by genetic distillation. *arXiv preprint arXiv:1711.01012*.

[9] Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.

[10] Hansen, N. (2016). The cma evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*.

[11] Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. (2018). Deep reinforcement learning that matters. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

[12] Islam, R., Henderson, P., Gomrokchi, M., and Precup, D. (2017). Reproducibility of benchmarked deep reinforcement learning tasks for continuous control. *arXiv preprint arXiv:1708.04133*.

[13] Jaderberg, M., Dalibard, V., Osindero, S., Czarnecki, W. M., Donahue, J., Razavi, A., Vinyals, O., Green, T., Dunning, I., Simonyan, K., et al. (2017). Population based training of neural networks. *arXiv preprint arXiv:1711.09846*.

[14] Jin, Y. (2011). Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation*, 1(2):61–70.

[15] Jung, W., Park, G., and Sung, Y. (2020). Population-guided parallel policy search for reinforcement learning. *arXiv preprint arXiv:2001.02907*.

[16] Khadka, S., Majumdar, S., Miret, S., Tumer, E., Nassar, T., Dwiel, Z., Liu, Y., and Tumer, K. (2019). Collaborative evolutionary reinforcement learning. *arXiv preprint arXiv:1905.00976*.

[17] Khadka, S. and Tumer, K. (2018). Evolution-guided policy gradient in reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 1188–1200.

[18] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[19] Lehman, J., Chen, J., Clune, J., and Stanley, K. O. (2018). Safe mutations for deep and recurrent neural networks through output gradients. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 117–124.

[20] Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.

[21] Liu, G., Zhao, L., Yang, F., Bian, J., Qin, T., Yu, N., and Liu, T.-Y. (2019). Trust region evolution strategies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4352–4359.

[22] Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937.

[23] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.

[24] Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in pytorch.

[25] Plappert, M., Houthooft, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. (2017). Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*.

[26] Pourchot, A. and Sigaud, O. (2018). Cem-rl: Combining evolutionary and gradient-based methods for policy search. *arXiv preprint arXiv:1810.01222*.

[27] Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*.

[28] Schmidhuber, J. and Zhao, J. (1998). Direct policy search and uncertain policy evaluation. In *Aaai spring symposium on search under uncertain and incomplete information, stanford univ*, pages 119–124.

[29] Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897.

[30] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

[31] Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic policy gradient algorithms.

[32] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359.

[33] Staines, J. and Barber, D. (2013). Optimization by variational bounding. In *ESANN*.

[34] Such, F. P., Madhavan, V., Conti, E., Lehman, J., Stanley, K. O., and Clune, J. (2017). Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv preprint arXiv:1712.06567*.

[35] Teh, Y., Bapst, V., Czarnecki, W. M., Quan, J., Kirkpatrick, J., Hadsell, R., Heess, N., and Pascanu, R. (2017). Distral: Robust multitask reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 4496–4506.

[36] Todorov, E., Erez, T., and Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE.

[37] Van Hasselt, H., Guez, A., and Silver, D. (2016). Deep reinforcement learning with double q-learning. In *Thirtieth AAAI conference on artificial intelligence*.

[38] Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., Peters, J., and Schmidhuber, J. (2014). Natural evolution strategies. *The Journal of Machine Learning Research*, 15(1):949–980.