EFFICIENT EXPLORATION USING MODEL-BASED QUALITY-DIVERSITY WITH GRADIENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Exploration is a key challenge in Reinforcement Learning, especially in long-horizon, deceptive and sparse-reward environments. For such applications, population-based approaches have proven effective. Methods such as Quality-Diversity deals with this by encouraging novel solutions and producing a diversity of behaviours. However, these methods are driven by either undirected sampling (i.e. mutations) or use approximated gradients (i.e. Evolution Strategies) in the parameter space, which makes them highly sample-inefficient. In this paper, we propose a model-based Quality-Diversity approach, relying on gradients and learning in imagination. Our approach optimizes all members of a population simultaneously to maintain both performance and diversity efficiently by leveraging the effectiveness of QD algorithms as good data generators to train deep models. We demonstrate that it maintains the divergent search capabilities of population-based approaches while significantly improving their sample efficiency (5 times faster) and quality of solutions (2 times more performant).

1 INTRODUCTION

Reinforcement Learning (RL) has demonstrated tremendous abilities to learn challenging tasks across a range of applications (Mnih et al., 2015; Silver et al., 2016; Akkaya et al., 2019). However, they generally struggle with exploration as the agent can only gather data by interacting with the environment. On the other hand, population based learning methods have shown to be very effective approaches (Jaderberg et al., 2017; Vinyals et al., 2019; Ecoffet et al., 2021; Wang et al., 2020). In contrast to single agent learning, training a population of agents allow diverse behaviors and data to be collected. This results in exploration that can better handle sparse and deceptive rewards Ecoffet et al. (2021) as well as alleviate catastrophic forgetting (Conti et al., 2018).

An effective way to use the population of agents for exploration are novelty search methods (Lehman & Stanley, 2011; Conti et al., 2018) where the novelty of the behaviors of new agents is measured with respect to the population. This novelty measure is then used in place of the conventional task reward similar to curiosity and intrinsic motivation approaches (Oudeyer et al., 2007; Bellemare et al., 2016; Pathak et al., 2017). Quality-Diversity (QD) Pugh et al. (2016); Cully et al. (2015); Chatzilygeroudis et al. (2021) extends this but also optimizes all members of the population on the task reward while maintaining the diversity through novelty. Beyond exploration, the creativity involved in finding various ways to solve a problem/task (i.e. the QD problem) is an interesting aspect of general intelligence that is also associated with adaptability. For instance, discovering diverse walking gaits can enable rapid adaptation to damage (Cully et al., 2015).

However, a drawback of conventional population based approaches is the large amounts of samples and evaluations required, usually in the order of millions. Some methods that utilize Evolutionary Strategies (ES) and more recently MAP-Elites Mouret & Clune (2015) (a common QD algorithm), sidestep this issue as they can parallelize and scale better with compute (Salimans et al., 2017; Conti et al., 2018; Lim et al., 2022) than their Deep RL counterparts, resulting in faster wall-clock times. Despite this, they still come at a cost of many samples. One of the main reasons for this lies in the underlying optimization operators. QD methods generally rely on undirected search methods such as objective-agnostic random perturbations (Mouret & Clune, 2015; Vassiliades & Mouret, 2018) to favor creativity and exploration. More directed search such as ES has also been used (Colas et al.,

2020) but relies on a large number of such perturbations ($\sim$thousands) to approximate a single step of natural gradient to direct the improvement of solutions.

In this paper, we introduce an extended version of Dynamics-Aware QD (DA-QD-ext) as well as Gradient and Dynamics Aware QD (GDA-QD), a new model-based QD method to perform sample-efficient exploration in RL. GDA-QD optimizes an entire population of diverse policies through a QD process in imagination using a learned dynamics model. Additionally, GDA-QD augments the conventional QD optimization operators with policy gradient updates using a critic network to obtain a more performant population. Beyond the effective exploration capabilities of QD methods, they are also excellent data generators. We leverage this idea to harvest a diversity of transitions to train the dynamics model and the critic. Thus, GDA-QD combine the powerful function-approximation capabilities of deep neural networks with the directed-search abilities of gradient-based learning and the creativity of population-based approaches. We demonstrate that it successfully outperforms both Deep RL and QD baselines in a hard-exploration task. GDA-QD exceeds the performance of baseline QD algorithms by ~ 1.5 times, and can reach the same results in 5 times less samples.

2 Preliminaries

2.1 Reinforcement Learning

Reinforcement Learning (RL) is commonly formalised as a Markov Decision Process (MDP) (Sutton & Barto, 2018) represented by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$, where $\mathcal{S}$ and $\mathcal{A}$ are the set of states and actions. $\mathcal{P}(s_{t+1}|s_t, a_t)$ is the probability of transition from state s_t to s_{t+1} given an action a_t, where s_t, $s_{t+1} \in \mathcal{S}$ and $a_t \in \mathcal{A}$. The reward function defines the reward obtained at each timestep $r_t = r(s_t, a_t, s_{t+1})$ when transitioning from state s_t to s_{t+1} under action a_t. An agent acting in the environment selects its next action based on the current state s_t by following a policy $\pi_\theta(a_t|s_t)$. The conventional objective in RL is then to optimize the parameters θ of policy π_θ, such that it maximizes the expected cumulative reward $R(\tau) = \sum_{t=1}^{T} r_t$ over the entire episode trajectory τ:

$$J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\mathcal{R}(\tau) \right] \tag{1}$$

The transition probabilities $\mathcal{P}(s_{t+1}|s_t, a_t)$ of the environment are usually assumed to be unknown. Model-based RL (Wang et al., 2019b) methods learn a parametric model $p_\phi(s_{t+1}|s_t, a_t)$ typically using supervised learning, from data collected when interacting in the environment. Policies are then trained using transitions obtained by rolling out the model.

2.2 Quality Diversity

Quality-Diversity (QD) (Pugh et al., 2016; Cully & Demiris, 2017) are diversity-seeking population-based approaches to learning. QD methods maintain a diversity of policies in the population Θ while maximizing the performance of each policy $\theta \in \Theta$. The population usually contains thousands of policies. QD considers an objective function $F(\theta)$ acting on the parameters of the policy θ. Additionally, QD also considers a behavior descriptor $d(\theta)$ that characterizes the behavior induced by a policy. $d(\theta)$ is used to maintain solutions in their behavioral niche to guarantee the population diversity and that there is no two solutions in the population with similar behavior descriptor $d(\theta)$. When applying QD to a RL problem, they are defined as follows, where $d(\tau)$ is the behavior descriptor of a given trajectory τ:

$$F(\theta) = J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\mathcal{R}(\tau) \right] \quad \text{and} \quad d(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[d(\tau) \right] \tag{2}$$

Similar to Evolutionary Strategies (ES) (Salimans et al., 2017), QD methods operate on entire episodes and hence both the objective and the descriptor can be computed simultaneously for any parameter vector θ. QD methods then aim to maximize:

$$\max_{\Theta} \quad \text{QD-Score}(\Theta) = \sum_{\theta \in \Theta} F(\theta) \tag{3}$$

By maintaining a population Θ of both diverse and high-performing policies, QD uses the existing parameters in the population as stepping stones (Nguyen et al., 2016; Wang et al., 2019a) in

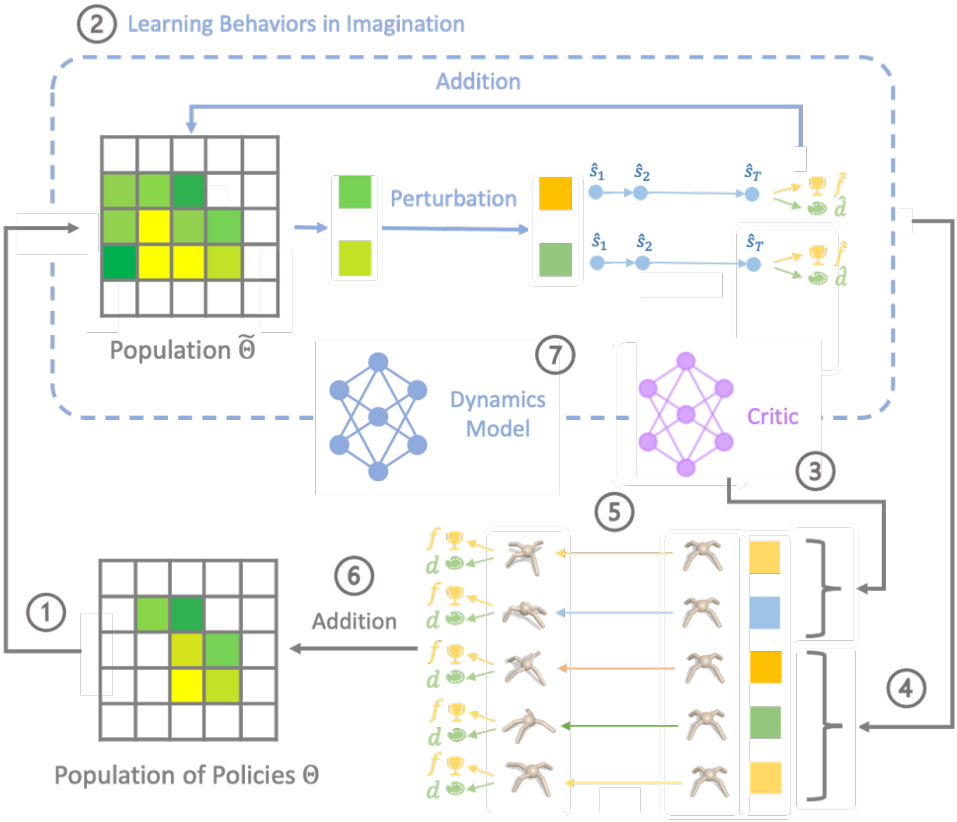

Figure 1: The GDA-QD algorithm can be summarized as follows: (1) the current population is copied, (2) this new population $\tilde{\Theta}$ is used to perform multiple steps of QD optimization fully in imagination using the dynamics model, (3) the critic is then used to update sampled policies learned in imagination with policy-gradient updates; they are concatenated with (4) policies sampled from the resulting population learned in imagination $\tilde{\Theta}$, (5) these concatenated batch of policies are then evaluated in the environment and (6) the real population of policies is updated before undergoing a new optimization loop, (7) the transitions collected in the environment are then used to train the dynamics model and the critic.

the optimization. At each iteration t, a random set of policies are sampled from the current population Θ_t, perturbed and evaluated in the environment. Based on the results of these evaluations, these perturbed policies might replace existing ones or fill in a new niche in the population. This incrementally improves the QD-Score(Θ_t) and encourages creativity and exploration.

3 METHOD

In this section, we introduce our two new model-based QD methods: DA-QD-ext and GDA-QD.

3.1 LEARNING AND SIEVING IN IMAGINATION

QD approaches are conventionally driven by random perturbations of the parameters of solutions. This undirected and divergent process gives QD its exploration capabilities but its major drawback is the number of samples it requires. This is especially evident for high-dimensional optimization problems such as optimizing deep neural networks where usually more directed gradient based methods are used. In this work, inspired by Lim et al. (2021), we perform the perturbation process in imagination, relying on a learned dynamics model of the environment to reduce the number of environment interactions when evaluating such perturbed policies.

Algorithm 1: GDA-QD

Inputs: T number of iterations, N number of imagined iterations, p_{model}
$\Theta \leftarrow$ initialised_population()
$q_\phi, Q_\psi \leftarrow$ initialised_models()
for $t = 1, ..., T$ **do**
  $\tilde{\Theta} \leftarrow \Theta$
  // Optimize population in imagination
  for $t_{imagination} = 1, ..., N$ **do**
    $\theta \leftarrow$ random_selection($\tilde{\Theta}$)
    $\tilde{\theta} \leftarrow$ perturb(θ)
    $F(\tilde{\theta}), d(\tilde{\theta}) \leftarrow$ evaluate_imgination($\tilde{\theta}, q_\phi$)
    $\tilde{\Theta} \leftarrow$ update_population($\tilde{\theta}, F(\tilde{\theta}), d(\tilde{\theta})$)
  end for
  // Get final policies
  $\theta_{model} \leftarrow$ last_added($p_{model}, \tilde{\Theta}$)
  $\theta_{gradient} \leftarrow$ apply_gradient($1 - p_{model}, \tilde{\Theta}, Q_\psi$)
  $\theta_{final} \leftarrow \theta_{model} + \theta_{gradient}$
  // Evaluate in environment, update all elements
  $F(\theta), d(\theta) \leftarrow$ evaluate(θ_{final})
  $\Theta \leftarrow$ update_population($\theta, F(\theta), d(\theta)$)
  $q_\phi, Q_\psi \leftarrow$ update_models()
end for

At each iteration t of the algorithm, the current population Θ_t is first "copied" into imagination $\tilde{\Theta}_t$. The policies of this provisional population $\pi_{\tilde{\theta}} \in \tilde{\Theta}_t$ are then perturbed, and evaluated in imagination using the rollouts of the dynamics model q_ϕ. Both the objective $F(\pi_{\tilde{\theta}})$ and the descriptor $d(\pi_{\tilde{\theta}})$ of the policies can be obtained from the state information present in the rollouts. Using this process, $\tilde{\Theta}_t$ undergoes multiple steps of QD optimization in imagination. The resulting policies $\pi_{\tilde{\theta}} \in \tilde{\Theta}_t$ that are added to the provisional population $\tilde{\Theta}_t$ during learning in imagination are then evaluated in the environment and used to update the population Θ_t if they improve the QD-Score of the population. The updated population Θ_t is then used as a start for the next iteration of the algorithm. This process of performing QD in imagination acts as a sieve and filters out perturbed solutions that are not likely to improve the quality and diversity of the population, hence increasing the sample efficiency.

Following Chua et al. (2018), we use a probabilistic bootstrap ensemble of models q_ϕ to capture uncertainties. Each model in the ensemble is a probabilistic model which predicts parameters of a Gaussian distribution $N(\mu_\phi(s_t, a_t), \Sigma_\phi(s_t, a_t))$ which we can then sample from, capturing the aleatoric uncertainty. We refer to this model-based QD method as DA-QD-ext as we extend DA-QD (Lim et al., 2021) to optimize closed loop neural network policies in complex exploration domains.

3.2 Incorporating Gradients in Quality-Diversity

As mentioned above, the random perturbations driving QD approaches prove inefficient when applied to high-dimensional search spaces such as the parameters of deep neural networks (Colas et al., 2020). To deal with this, we augment the usual mutation operator with policy gradient information as done by Nilsson & Cully (2021). The policy gradient can be more intuitively thought of as a more directed perturbation, hence being a more efficient optimization update procedure. To apply the policy gradient to the policies in the population, we maintain a critic network Q_ψ which approximates the action-value function $Q(s_t, a_t) = \mathbb{E}\left[\sum_{k=0}^{T-t} \gamma^k r_{t+k+1} \mid s_t, a_t\right]$ and gives the expected return from being in state s_t and following action a_t. The critic allows us to gradually improve any policy in the direction maximizing the expected return by computing policy-gradient in Equation 4, approximated over a batch of transitions. We train the critic using the same procedure as TD3 (Fujimoto et al., 2018).

	APPROACH	QD-SCORE	COVERAGE	MAX-TOTAL-REWARD
TD3	Policy Gradient	-	-	182.14
ES	Natural Gradient	302074.94	32.22	15.02
MAP-Elites	Mutation	977539.6	42.16	196.99
PGA-MAP-Elites	Policy Gradient	1016300.3	43.2	254.75
MAP-Elites-ES	Natural Gradient	316052.03	22.36	163.95
DA-QD-ext (Ours)	Mutation	1503323.6	63.16	267.66
GDA-QD (Ours)	Policy Gradient	**1724573.4**	**70.6**	**321.94**

Table 1: Final QD-Score, Coverage and Max-Total-Reward reached by all algorithms on the Ant-Trap task for 10^6 policies evaluations. Each experiment is replicated 3 times, we report in the table the median value across runs.

$$\nabla_{\theta_i} J(\theta_i) = \mathbb{E}_{\mathbf{s},\mathbf{a}\sim\pi_{\theta_i}} \left[\nabla_{\theta_i} \pi_{\theta_i}(\mathbf{s}) \; \nabla_{\mathbf{a}} Q_\psi(\mathbf{s},\mathbf{a}) \right] \tag{4}$$

As explained in Section 3.1, the population of policies Θ_t is copied in imagination $\tilde{\Theta}_t$ to undergo multiple steps of QD optimization. In imagination, only random perturbations are used to encourage exploration and optimize the population $\tilde{\Theta}_t$. Once the process in imagination finishes, B policies are sampled to be evaluated. A proportion p_{model} of B is sampled from the new $\tilde{\Theta}_t$ to be evaluated directly while another $p_{gradient} = 1 - p_{model}$ is sampled to be improved using the policy gradient operators before being evaluated. In the following, we use $p_{model} = 0.9$. We also study the impact of this value and illustrate the complementary properties of these two generation-procedures. We demonstrate that both of them are essential to guarantee performance. We refer to this new algorithm as GDA-QD as it leverages both gradient information and learns the dynamics of the environment. Figure 1 and Algorithm 1 provides a summary of this algorithm.

3.3 QUALITY-DIVERSITY AS DATA GENERATORS FOR DEEP MODELS

A key quality of QD algorithms is that they are excellent at generating diverse and high-quality data (Ecoffet et al., 2021; Grillotti & Cully, 2022). In our work, the search for a diversity of high-performing behaviors when optimizing the population results in a diverse dataset of transitions. We leverage this property to train deep models which require and often excel when provided with such data. As commonly done, the diverse dataset of transitions is stored in a replay buffer and used to train (1) the dynamics model q_ϕ and (2) the critic Q_ψ. Both these models are suitable candidates as the training of the critic Q_ψ is off-policy and can trained using transitions collected from any behavioral policy. Additionally, training the dynamics model q_ϕ is a supervised learning problem which would benefit from a large and diverse dataset of transitions.

We found this property to be especially important to our method. Specific to the training of the dynamics model, GDA-QD does not use transitions produce by gradient-optimized policies to train the model. This was found empirically in our studies as we observed the transitions produced by the gradient-optimized policies induced a shift in distribution that prove detrimental for the training of the dynamics models. Results of this study are detailed later in Section 4.3. It is important to recognize that the transitions used to train the dynamics model are not merely just from policies that have been randomly perturbed but a population of policies that have undergone multiple steps of QD optimization in imagination and are expected to improve the population. This results in a diverse and high performing dataset of transitions that GDA-QD uses to further train its models.

4 EXPERIMENTS

We aim to evaluate our method by answering four main questions: (1) Can we scale model-based QD approaches to RL domains and Neuroevolution? (2) Does GDA-QD results in more performant and sample-efficient learning than traditional QD approaches and simple model-based QD? (3) What is the importance of the policy-gradient perturbation in the performance of GDA-QD? (4) How does the data-generation capabilities of GDA-QD enforce efficient learning?

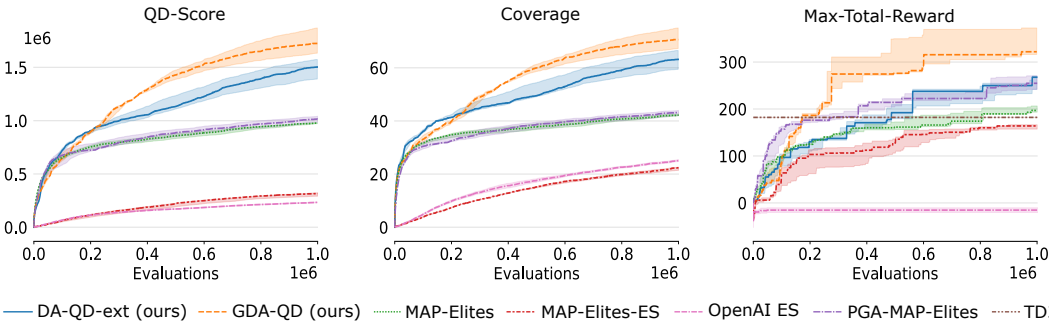

Figure 2: QD-Score (left), Coverage (middle) and Max-Total-Reward (right) of all algorithms on the AntTrap task for 10^6 policies evaluations. Each experiment is replicated 3 times, the solid line corresponds to the median over replications and the shaded area to the first and third quartiles.

4.1 Experimental setup

Tasks and Environments: We focus on the AntTrap domain (Conti et al., 2018; Colas et al., 2020; Cideron et al., 2020), considered in the literature as an hard exploration problems. In this task, an 8-DoF Ant robot learns how to walk, aiming to go beyond the trap (see Figure 3). It gets increasing rewards for going as fast as possible in a straight line while minimizing energy-usage. This reward definition makes this task deceptive as it drives the robot directly into the trap. The descriptor $d(\pi_\theta)$ is defined as the x-y position at the end of the trajectory.

Baselines: Across our experiments, we consider the following baselines:

- **TD3** as proposed in Fujimoto et al. (2018). For fairness, as this algorithm does not rely on a population, nor on complete-episode evaluations, we only reports its final value in the results.
- **OpenAI ES** the original ES algorithm (Salimans et al., 2017), relying on natural gradients approximation. To ensure comparability with other approaches, we consider every estimate-evaluation as one sample, making this algorithm highly sample-inefficient.
- **MAP-Elites** the most-commonly used QD algorithm (Mouret & Clune, 2015).
- **MAP-Elites-ES** a QD algorithm mixing MAP-Elites with OpenAI ES (Colas et al., 2020). This approach suffers from the same limitation as OpenAI ES.
- **PGA-MAP-Elites** from (Nilsson & Cully, 2021), a policy-gradient based QD method.

Metrics: We consider two metrics to assess the performance of GDA-QD:

- **QD-Score:** defined in Section 2.2. It quantifies the diversity and quality of the overall population and allows to compare population-based methods.
- **Coverage:** Coverage is defined as a percentage of the descriptor space filled by the population, used as a measure of the diversity of the population.
- **Max-Total-Reward:** the total reward of the best individual of the current population. This metrics allows comparison with single-policy methods such as RL baselines. .

Implementation and hyperparameters: Our source code is available at TO-BE-RELEASED.[1]. It includes a containerised environment to replicate our experiments. The methods presented in this paper as well as all our baselines are based on the implementation of MAP-Elites in the QDax open-source library (Lim et al., 2022), using the Brax simulator Freeman et al. (2021). We use a batch-size of 512 for MAP-Elites with a maximal population size of 2500 policies. All the hyperparameters used to train the critic networks are taken from the original TD3 paper Fujimoto et al. (2018). For the dynamics model we use an ensemble of 4 probabilistic models. The detailed hyperparameters can be found in Appendix A.

[1]The code will be made available upon acceptance to avoid revealing the authors' identities and breaking the double-blind review process.

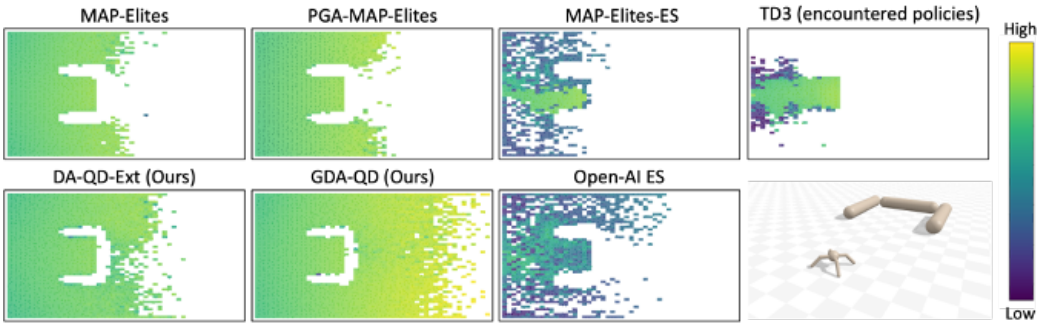

Figure 3: Final population of methods and the AntTrap environment (bottom right). Each policy in the population is represented with a dot in the final position it manages to reach within the episode. The color of the square around each policy indicates its total reward, the lighter the better. The obstacle clearly appears on this plot, as the empty area in the middle. Most of the baselines do not manage to go around the obstacle. Our approach GDA-QD clearly outperform all others approaches.

4.2 Results

The results of our experiments are summarized in Figure 2 and Table 1. We also display a visualization of the final populations of policies for each algorithm in Figure 3.

Figure 2 shows that both our proposed model-based versions, DA-QD-ext and GDA-QD, significantly outperform all baselines in terms of sample efficiency and final performance. This demonstrates that we can scale model-based QD methods to deep neuroevolution. The performance of DA-QD-ext suggests that learning diverse behaviors in imagination using the dynamics model is a simple but effective approach to save samples yet maintain the divergent search capabilities of the random perturbations. However, we can see that the maximum total reward obtained by DA-QD-ext seems to stagnate and increase slowly, especially when compared to GDA-QD and PGA-MAP-Elites. This can be explained by the absence of reward maximizing optimization updates such as the policy gradient updates which are present in GDA-QD and PGA-MAP-Elites. Hence, GDA-QD is shown to get the "best of both worlds" by benefitting from the QD in imagination as well as the policy gradient updates. This is evident in its performance across the QD-Score, coverage and max return metrics.

In terms of baselines, the poor performance of ES-based algorithms (ES and ME-ES) can be attributed to the number of samples ($\sim$ hundreds or thousands) required just to approximate a single gradient step. It is important to note that these methods commonly do not consider the notion of evaluations and generally evaluate the algorithms versus time or number of generations as they are suited to be heavily parallelized across clusters of CPUs. The baseline TD3 agent is plotted as a dotted line to represent the max performance obtained by the agent. This is done to allow comparison as TD3 does not rely on a population, and also performs policies-updates within episodes. It proves unable to solve this deceptive reward task and gets stuck inside the trap.

4.3 Importance of QD for training deep models

To investigate the data generation abilities of QD, we run an ablation where we filter out transitions collected by the different policies that were sampled from the imagined population $\tilde{\Theta}$ or perturbed using the policy gradient updates. This means that we investigate the performance of the algorithm when the dynamics model is trained on different data distributions based on the policies rolled out in the environment. Figure 4 shows the performance curves when running this ablation. We first test the effect of filtering out transitions collected from policy gradient updated policies when $p_{gradient} = 0.1$. We notice a minor difference in which experiments that filter such transitions out perform better than when not using any filter. To enable a fair comparison in terms of the number of transitions added to train the replay buffer when attempting to do the converse (i.e. removing transitions from policies output by QD in imagination), we use a $p_{model} = 0.5$ and $p_{gradient} = 0.5$. We observe better performance when the dynamics model is only trained with transitions obtained

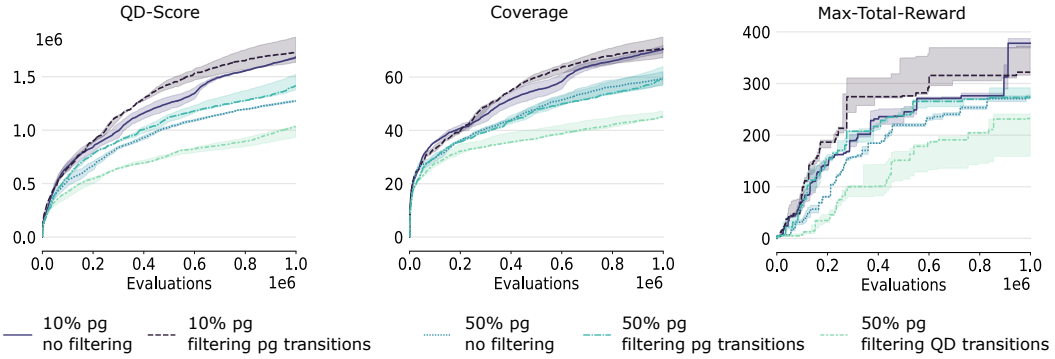

Figure 4: QD-Score (left), Coverage (middle) and Max-Total-Reward (right) on the AntTrap task of GDA-QD with different proportion of policy-gradients generated policies (pg) and different transitions filtering schemes. Each experiment is replicated 3 times, the solid line corresponds to the median over replications and the shaded area to the first and third quartiles.

from policies emitted and sieved through the QD process in imagination (filtering pg transitions). This is compared to training on the full mixture of transitions also given by the gradient-optimized policies (no filtering) where the difference is minimal as seen previously. However, performance significantly drops when training just on transitions given by the gradient optimized policies and ignoring transitions given by policies obtained from QD in imagination (filtering QD transitions). We hypothesize that this is due to bias in the transitions obtained through the policies perturbed by policy gradient resulting in a skewed dataset while the transitions given by policies obtained from QD in imagination provides a diverse and high-quality dataset.

4.4 Ablation of the proportion of policy-gradient

In this section, we aim to demonstrate the importance of both types of perturbations used in GDA-QD: policy-gradient and random model-filtered perturbations. We run an ablation of the different proportions p_{model} and $p_{gradient} = 1 - p_{model}$ of perturbations at each generation. In Figure 5, we can see that the configuration of $p_{gradient} = 0.1$ performs the best. This indicates that while the policy gradient operators are important, they are also generally wasteful evaluations. This is further backed up by the fact that the performance deteriorates as the proportion $p_{gradient}$ increases over p_{model}. However, it is also still key to maintain some proportion of policy gradient updates $p_{gradient}$, as not using them is detrimental to the max total reward which results in an overall lower performance of the population Θ.

5 Related Work

Searching for Diversity. Prior studies have shown the importance of maintaining a diverse set of solutions to solve a problem. Novelty-search approaches (Lehman & Stanley, 2011) inspired from evolutionary computation optimize for the novelty of solutions defined by a behavioral characterization with respect to a population instead of the optimization objective. QD methods (Cully & Demiris, 2017; Mouret & Clune, 2015) extend this approach by also considering the objective, aiming to find both diverse and high-performing population of solutions. Similarly, our work builds on QD approaches with the aim to maintain a diversity of high-performing solutions. Other approaches to searching for diversity also exist in the RL community. Unsupervised RL methods (Eysenbach et al., 2018; Sharma et al., 2020) commonly use a mutual information maximization objective to learn a diversity of behaviors in a skill-conditioned policies. Similar to QD, Kumar et al. (2020); Zahavy et al. (2022) have proposed to extend these unsupervised RL approaches by integrating objectives using constrained Markov decision processes. While these approaches typically focus on maintaining a dozen of different policies, our algorithm discovers thousands of independent diverse policies.

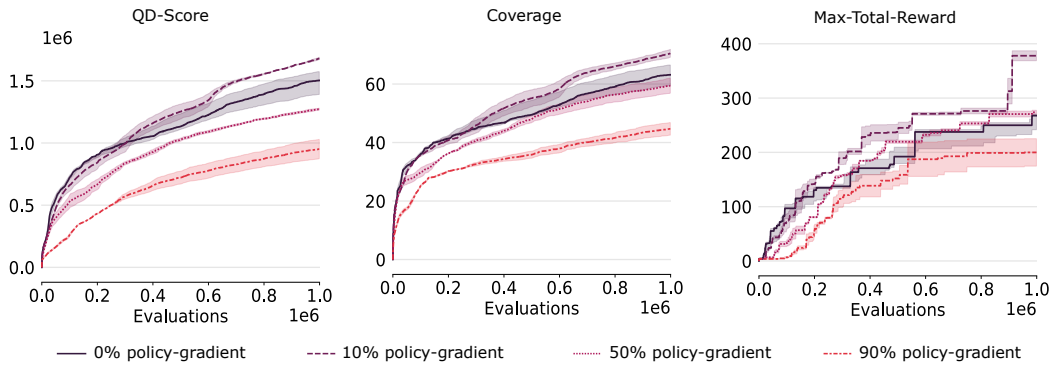

Figure 5: QD-Score (left), Coverage (middle) and Max-Total-Reward (right) on the AntTrap task of GDA-QD with different values of $p_{gradient}$. $p_{gradient}$ gives the proportion of policies improved using policy-gradients. Each experiment is replicated 3 times, the solid line corresponds to the median over replications and the shaded area to the first and third quartiles.

Neuroevolution seeks to evolve neural networks through biologically-inspired methods such as evolutionary algorithms and have interesting properties unavailable to common gradient-based methods (Stanley et al., 2019). However, a limitation in the neuroevolution domain is the dimensionality of the search-space, that quickly limits the effectiveness of random perturbations. Some methods have overcome this limitation by using indirect encoding methods (Clune et al., 2009) or through natural approximated gradients (Salimans et al., 2017). Our work hybridizes neuroevolution with deep reinforcement learning methods. Similar to Nilsson & Cully (2021), we use policy-gradients as directed perturbations to effectively maneuver the high-dimensional search space but, we significantly improve the sample efficiency and performance by also augmenting these operators with model-based methods.

Model-based Quality-Diversity. As the perturbations commonly used in QD are sample inefficient, prior work has sought the use of data-driven models to alleviate this. SAIL Gaier et al. (2018) introduced the use of a surrogate model in the form of a Gaussian Process model to predict the objective. As Gaussian processes generally only work well on low-dimensional data, more recent methods have explored the use of deep networks (Keller et al., 2020; Lim et al., 2021; Zhang et al., 2022) as forms surrogate models to predict both the objective and descriptors. We utilize the model-based QD framework from Lim et al. (2021) which first introduced the idea of maintaining an imagined population and also builds on model-based RL methods (Wang et al., 2019b) where a dynamics model is used. Critically, this work has only been applied to low-dimensional open-loop policies. To the best of our knowledge, our work is the first model-based QD algorithm that scales to the more complex deep neuroevolution domain where we optimize closed loop RL policies.

6 Conclusion and Future Work

In this paper, we introduce a novel model-based Quality-Diversity method, GDA-QD, which optimizes a population of diverse policies to explore more efficiently. To the best of our knowledge, this approach is the first model-based QD algorithm scaling to neuroevolution to optimize deep neural network policies. We leverage a key property of QD algorithms as effective data generators to train deep models in the form of a dynamics model and a critic. In turn, these models help to significantly improve the sample efficiency and final performance of the QD algorithm. The dynamics model is used to learn and sieve policies in imagination while the critic is used to apply policy gradient updates to sampled policies. Our experiments show that GDA-QD outperforms a range of Deep RL and QD baselines on a hard exploration task containing deceptive rewards. Overall, we demonstrate some of the powerful synergies that can arise between population-based learning and deep learning approaches. In future work, we hope to extend our work to more complex domains through the use of latent dynamics models (Ha & Schmidhuber, 2018; Hafner et al., 2019).

REFERENCES

Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik's cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.

Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29, 2016.

Konstantinos Chatzilygeroudis, Antoine Cully, Vassilis Vassiliades, and Jean-Baptiste Mouret. Quality-diversity optimization: a novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*, pp. 109–135. Springer, 2021.

Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *arXiv preprint arXiv:1805.12114*, 2018.

Geoffrey Cideron, Thomas Pierrot, Nicolas Perrin, Karim Beguir, and Olivier Sigaud. Qd-rl: Efficient mixing of quality and diversity in reinforcement learning. *arXiv preprint arXiv:2006.08505*, 2020.

Jeff Clune, Benjamin E Beckmann, Charles Ofria, and Robert T Pennock. Evolving coordinated quadruped gaits with the hyperneat generative encoding. In *2009 iEEE congress on evolutionary computation*, pp. 2764–2771. IEEE, 2009.

Cédric Colas, Vashisht Madhavan, Joost Huizinga, and Jeff Clune. Scaling map-elites to deep neuroevolution. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pp. 67–75, 2020.

Edoardo Conti, Vashisht Madhavan, Felipe Petroski Such, Joel Lehman, Kenneth Stanley, and Jeff Clune. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. *Advances in neural information processing systems*, 31, 2018.

Antoine Cully and Yiannis Demiris. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, 22(2):245–259, 2017.

Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.

Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. First return, then explore. *Nature*, 590(7847):580–586, 2021.

Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.

C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax - a differentiable physics engine for large scale rigid body simulation, 2021. URL http://github.com/google/brax.

Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.

Adam Gaier, Alexander Asteroth, and Jean-Baptiste Mouret. Data-efficient design exploration through surrogate-assisted illumination. *Evolutionary computation*, 26(3):381–410, 2018.

Luca Grillotti and Antoine Cully. Unsupervised behaviour discovery with quality-diversity optimisation. *IEEE Transactions on Evolutionary Computation*, 2022.

David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Advances in neural information processing systems*, 31, 2018.

Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.

Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.

Leon Keller, Daniel Tanneberg, Svenja Stark, and Jan Peters. Model-based quality-diversity search for efficient robot learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9675–9680. IEEE, 2020.

Saurabh Kumar, Aviral Kumar, Sergey Levine, and Chelsea Finn. One solution is not all you need: Few-shot extrapolation via structured maxent rl. *Advances in Neural Information Processing Systems*, 33:8198–8210, 2020.

Joel Lehman and Kenneth O Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pp. 211–218, 2011.

Bryan Lim, Luca Grillotti, Lorenzo Bernasconi, and Antoine Cully. Dynamics-aware quality-diversity for efficient learning of skill repertoires. *arXiv preprint arXiv:2109.08522*, 2021.

Bryan Lim, Maxime Allard, Luca Grillotti, and Antoine Cully. Accelerated quality-diversity for robotics through massive parallelism. *arXiv preprint arXiv:2202.01258*, 2022.

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.

Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

Anh Nguyen, Jason Yosinski, and Jeff Clune. Understanding innovation engines: Automated creativity and improved stochastic optimization via deep learning. *Evolutionary computation*, 24(3): 545–572, 2016.

Olle Nilsson and Antoine Cully. Policy gradient assisted map-elites. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 866–875, 2021.

Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.

Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pp. 2778–2787. PMLR, 2017.

Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.

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

Archit Sharma, Michael Ahn, Sergey Levine, Vikash Kumar, Karol Hausman, and Shixiang Gu. Emergent real-world robotic skills via unsupervised off-policy reinforcement learning. *arXiv preprint arXiv:2004.12974*, 2020.

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, Jeff Clune, Joel Lehman, and Risto Miikkulainen. Designing neural networks through neuroevolution. *Nature Machine Intelligence*, 1(1):24–35, 2019.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Vassilis Vassiliades and Jean-Baptiste Mouret. Discovering the Elite Hypervolume by Leveraging Interspecies Correlation. In *GECCO 2018 - Genetic and Evolutionary Computation Conference*, Kyoto, Japan, July 2018. doi: 10.1145/3205455.3205602. URL https://hal.inria.fr/hal-01764739.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Poet: open-ended coevolution of environments and their optimized solutions. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 142–151, 2019a.

Rui Wang, Joel Lehman, Aditya Rawal, Jiale Zhi, Yulun Li, Jeffrey Clune, and Kenneth Stanley. Enhanced poet: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions. In *International Conference on Machine Learning*, pp. 9940–9951. PMLR, 2020.

Tingwu Wang, Xuchan Bao, Ignasi Clavera, Jerrick Hoang, Yeming Wen, Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019b.

Tom Zahavy, Yannick Schroecker, Feryal Behbahani, Kate Baumli, Sebastian Flennerhag, Shaobo Hou, and Satinder Singh. Discovering policies with domino: Diversity optimization maintaining near optimality. *arXiv preprint arXiv:2205.13521*, 2022.

Yulun Zhang, Matthew C Fontaine, Amy K Hoover, and Stefanos Nikolaidis. Deep surrogate assisted map-elites for automated hearthstone deckbuilding. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 158–167, 2022.