
Parameter Space Noise for Exploration

Matthias Plappert[‡], Rein Houthoofd[†], Prafulla Dhariwal[†], Szymon Sidor[†],
Richard Y. Chen[†], Xi Chen[†], Tamim Asfour[‡], Pieter Abbeel[†], and Marcin Andrychowicz[†]

[†] OpenAI

[‡] Karlsruhe Institute of Technology (KIT)

Abstract

Deep reinforcement learning (RL) methods generally engage in exploratory behavior through noise injection in the action space. An alternative is to add noise directly to the agent’s parameters, which can lead to more consistent exploration and a richer set of behaviors. Methods such as evolutionary strategies use parameter perturbations, but discard all temporal structure in the process and require significantly more samples. Combining parameter noise with traditional RL methods allows to combine the best of both worlds. We demonstrate that both off- and on-policy methods benefit from this approach through experimental comparison of DQN, DDPG, and TRPO on high-dimensional discrete action environments as well as continuous control tasks. Our results show that RL with parameter noise learns more efficiently than traditional RL with action space noise and evolutionary strategies individually.

1 Introduction

Exploration remains a key challenge in contemporary deep reinforcement learning (RL). Its main purpose is to ensure that the agent’s behavior does not converge prematurely to a local optimum. Enabling efficient and effective exploration is, however, not trivial since it is not directed by the reward function of the underlying Markov decision process (MDP). Although a plethora of methods have been proposed to tackle this challenge in high-dimensional and/or continuous-action MDPs, they often rely on complex additional structures such as counting tables [1], density modeling of the state space [2], learned dynamics models [3–5], or self-supervised curiosity [6].

An orthogonal way of increasing the exploratory nature of these algorithms is through the addition of temporally-correlated noise, for example as done in bootstrapped DQN [7]. Along the same lines, it was shown that the addition of parameter noise leads to better exploration by obtaining a policy that exhibits a larger variety of behaviors [8, 9]. We discuss these related approaches in greater detail in Section 5. Their main limitation, however, is that they are either only proposed and evaluated for the on-policy setting with relatively small and shallow function approximators [10] or disregard all temporal structure and gradient information [9, 11, 12].

This paper investigates how parameter space noise can be effectively combined with off-the-shelf deep RL algorithms such as DQN [13], DDPG [14], and TRPO [15] to improve their exploratory behavior. Experiments show that this form of exploration is applicable to both high-dimensional discrete environments and continuous control tasks, using on- and off-policy methods. Our results indicate that parameter noise outperforms traditional action space noise-based baselines, especially in tasks where the reward signal is extremely sparse. This demonstrates that a fertile middle ground exists between evolutionary methods that discard temporal structure, and methods that rely entirely on unstructured noise injection.

Correspondence to matthias.plappert@partner.kit.edu.

2 Background

We consider the standard RL framework consisting of an agent interacting with an environment. To simplify the exposition we assume that the environment is fully observable. An environment is modeled as a Markov decision process (MDP) and is defined by a set of states \mathcal{S} , a set of actions \mathcal{A} , a distribution over initial states $p(s_0)$, a reward function $r : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, transition probabilities $p(s_{t+1}|s_t, a_t)$, a time horizon T , and a discount factor $\gamma \in [0, 1)$. We denote by π_θ a policy parametrized by θ , which can be either deterministic, $\pi : \mathcal{S} \mapsto \mathcal{A}$, or stochastic, $\pi : \mathcal{S} \mapsto (\mathcal{A} \mapsto \mathbb{R}_{\geq 0})$. The agent’s goal is to maximize the expected discounted return $\eta(\pi_\theta) = \mathbb{E}_\tau[\sum_{t=0}^T \gamma^t r(s_t, a_t)]$, where $\tau = (s_0, a_0, \dots, s_T)$ denotes a trajectory with $s_0 \sim p(s_0)$, $a_t \sim \pi_\theta(a_t|s_t)$, and $s_{t+1} \sim p(s_{t+1}|s_t, a_t)$. Experimental evaluation is based on the undiscounted return $\mathbb{E}_\tau[\sum_{t=0}^T r(s_t, a_t)]$.

2.1 Off-policy Methods

Off-policy RL methods allow learning based on data captured by arbitrary policies. This paper considers two popular off-policy algorithms, namely Deep Q-Networks (DQN) [13] and Deep Deterministic Policy Gradients (DDPG) [14].

Deep Q-Networks (DQN) DQN [13] uses a deep neural network as a function approximator to estimate the optimal Q -value function, which conforms to the Bellman optimality equation:

$$Q(s_t, a_t) = r(s_t, a_t) + \gamma \max_{a' \in \mathcal{A}} Q(s_{t+1}, a').$$

The policy is implicitly defined by Q as $\pi(s_t) = \operatorname{argmax}_{a' \in \mathcal{A}} Q(s_t, a')$. Typically, a stochastic ϵ -greedy or Boltzmann policy [16] is derived from the Q -value function to encourage exploration, which relies on sampling noise in the action space. The Q -network predicts a Q -value for each action, thus requiring that \mathcal{A} is discrete, and Q is updated using off-policy data from a replay buffer.

Deep Deterministic Policy Gradients (DDPG) DDPG [14] is an actor-critic algorithm, applicable to continuous action spaces. Similar to DQN, the critic estimates the Q -value function using off-policy data and the recursive Bellman equation:

$$Q(s_t, a_t) = r(s_t, a_t) + \gamma Q(s_{t+1}, \pi_\theta(s_{t+1})),$$

where π_θ is the actor or policy. The actor is trained to maximize the critic’s estimated Q -values by back-propagating through both networks. For exploration, DDPG uses a stochastic policy of the form $\widehat{\pi}_\theta(s_t) = \pi_\theta(s_t) + w$, where w is either $w \sim \mathcal{N}(0, \sigma^2 I)$ (uncorrelated) or $w \sim \text{OU}(0, \sigma^2)$ (correlated).¹ Again, exploration is realized through action space noise.

2.2 On-policy Methods

In contrast to off-policy algorithms, on-policy methods require updating function approximators according to the currently followed policy. In particular, we will consider Trust Region Policy Optimization (TRPO) [18], an extension of traditional policy gradient methods [19] using the natural gradient direction [20, 21].

Trust Region Policy Optimization (TRPO) The basis of TRPO [18] is REINFORCE [19], an algorithm that estimates the gradient of expected return $\nabla_\theta \eta(\pi_\theta)$ via the likelihood ratio trick:

$$\nabla_\theta \eta(\pi_\theta) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=0}^T \nabla_\theta \log \pi_\theta(a_t^i | s_t^i) (R_t^i - b_t^i),$$

where N is the number of episodes, $R_t^i = \sum_{t'=t}^T \gamma^{t'-t} r_{t'}^i$, and b_t^i a variance reducing baseline [22]. Hereafter, an ascent step is taken along the estimated gradient. TRPO improves upon REINFORCE by computing an ascent direction that ensures a small change in the policy distribution. This direction

¹OU(\cdot, \cdot) denotes the Ornstein-Uhlenbeck process [17].

is approximately $I(\theta)^{-1}\nabla_{\theta}\eta(\pi_{\theta})$, with $I(\theta)$ being the Fisher information matrix. More specifically, TRPO solves the following constrained optimization problem in each epoch:

$$\begin{aligned} \text{maximize}_{\theta} \quad & E_{s\sim\rho_{\theta'}, a\sim\pi_{\theta'}} \left[\frac{\pi_{\theta}(a|s)}{\pi_{\theta'}(a|s)} A_{\theta'}(s, a) \right] \\ \text{s.t.} \quad & E_{s\sim\rho_{\theta'}} [D_{\text{KL}}(\pi_{\theta'}(\cdot|s)\|\pi_{\theta}(\cdot|s))] \leq \delta_{\text{KL}} \end{aligned}$$

where $\rho_{\theta} = \rho_{\pi_{\theta}}$ is the discounted state-visitation frequencies induced by π_{θ} , $A_{\theta}(s, a)$ is estimated by the empirical return minus the baseline, and δ_{KL} is a step size parameter which controls how much the policy is allowed to change per iteration. We follow the procedure described in the original paper for solving this optimization problem [22].

3 Parameter Space Noise for Exploration

This work considers policies that are realized as parameterized functions, which we denote as π_{θ} , with θ being the parameter vector. We represent policies as neural networks but our technique can be applied to arbitrary parametric models. To achieve structured exploration, we sample from a set of policies by applying additive Gaussian noise to the parameter vector of the current policy: $\tilde{\theta} = \theta + \mathcal{N}(0, \sigma^2 I)$. Importantly, the perturbed policy is sampled at the beginning of each episode and kept fixed for the entire rollout. For convenience and readability, we denote this perturbed policy as $\tilde{\pi}_{\theta}$.

State-dependent exploration As pointed out in [10], there is a crucial difference between action space noise and parameter space noise. Consider the continuous action space case. When using Gaussian action noise, actions are sampled according to some stochastic policy, generating $a_t = \pi_{\theta}(s_t) + \mathcal{N}(0, \sigma^2 I)$. Therefore, even for a *fixed* state s , we will almost certainly obtain a different action whenever that state is sampled again in the rollout, since action space noise is completely *independent* of the current state s_t (notice that this is equally true for correlated action space noise). In contrast, if the parameters of the policy are perturbed at the beginning of each episode, we get $a_t = \tilde{\pi}_{\theta}(s_t)$. In this case, the same action will be taken every time the same state s_t is sampled in the rollout. This ensures consistency in actions, and directly introduces a dependence between the state and the exploratory action taken.

Perturbing deep neural networks It is not immediately obvious that deep neural networks, with potentially millions of parameters and complicated nonlinear interactions, can be perturbed in meaningful ways by applying spherical Gaussian noise. However, as recently shown in [9], a simple reparameterization of the network achieves exactly this. More concretely, we use layer normalization [23] between perturbed layers (in contrast to [9] that uses virtual batch normalization, which we found to perform less consistently). Due to this normalizing across activations within a layer, the same perturbation scale can be used across all layers, even though different layers may exhibit different sensitivities to noise.

Adaptive noise scaling Parameter space noise requires us to pick a suitable scale σ . This can be problematic since the scale will strongly depend on the specific network architecture, and is likely to vary over time as parameters become more sensitive to noise as learning progresses. Additionally, while it is easy to intuitively grasp the scale of action space noise, it is far harder to understand the scale in parameter space. We propose a simple solution that resolves all aforementioned limitations in an easy and straightforward way. This is achieved by adapting the scale of the parameter space noise over time and relating it to the variance in action space that it induces. More concretely, we can define a distance measure between perturbed and non-perturbed policy in action space and adaptively increase or decrease the parameter space noise depending on whether it is below or above a certain threshold. Since RL algorithms usually already have a hyperparameter for action space noise (e.g. ϵ for DQN), we can define this measure such that we can re-use this hyperparameter as the aforementioned threshold. The concrete realization depends on the policy representation and is described in detail for DQN, DDPG, and TRPO in Appendix C.

Parameter space noise for off-policy methods In the off-policy case, parameter space noise can be applied straightforwardly since, by definition, data that was collected off-policy can be used. More

concretely, we only perturb the policy for exploration and train the non-perturbed network on this data by replaying it.

Parameter space noise for on-policy methods Parameter noise can be incorporated in an on-policy setting, using an adapted policy gradient, as set forth in [24]. Policy gradient methods optimize $\mathbb{E}_{\tau \sim (\pi, p)}[R(\tau)]$. Given a stochastic policy $\pi_\theta(a|s)$ with $\theta \sim \mathcal{N}(\phi, \Sigma)$, the expected return can be expanded using likelihood ratios and the reparametrization trick [25] as

$$\nabla_{\phi, \Sigma} \mathbb{E}_\tau[R(\tau)] \approx \frac{1}{N} \sum_{\epsilon^i, \tau^i} \left[\sum_{t=0}^{T-1} \nabla_{\phi, \Sigma} \log \pi(a_t | s_t; \phi + \epsilon^i \Sigma) R_t(\tau^i) \right] \quad (1)$$

for N samples $\epsilon^i \sim \mathcal{N}(0, I)$ and $\tau^i \sim (\pi_{\phi + \epsilon^i \Sigma}, p)$ (see Appendix B for a full derivation). Rather than updating Σ according to the previously derived policy gradient, we fix its value to $\sigma^2 I$. This value may or may not be updated in a heuristic fashion. The reason is that experimentally, we noticed that learning noise variance through stochastic gradient descent takes a much longer time than learning the mean.

Comparison to other methods Traditional on-policy RL methods do not have an expectation over samples of the parameters θ , and thus do not obtain gradients from a diverse set of policies. On the other hand, while ES approaches like the one described in [9] sample parameter perturbations, they use the gradient estimator

$$\nabla_{\phi, \Sigma} \mathbb{E}_\tau[R(\tau)] \approx \frac{1}{N} \sum_{\epsilon^i, \tau^i} [\nabla_{\phi, \Sigma} \log p(\phi + \epsilon^i \Sigma) R(\tau^i)], \quad (2)$$

which completely ignores the temporal structure. A big advantage of ES is that they allow for effective exploration by sampling from a diverse set of behavioral policies. However, as the authors point out, when actions in the underlying MDP have no long-lasting effect, the discarding of temporal structure in ES can be highly detrimental to learning. In our case, by using the reparametrization trick, we get per-action gradient information, which allows for finer-grained credit assignment.

In this paper, we explore whether it is possible to combine both advantages, achieving effective exploration via parameter perturbations as in ES, while also having a high signal-to-noise ratio as in standard RL, which uses per-action gradient estimates.

4 Experiments

This section answers the following questions:

- (i) Do existing state-of-the-art RL algorithms benefit from incorporating parameter space noise?
- (ii) Does parameter space noise aid in exploring sparse reward environments more effectively?
- (iii) How does parameter space noise exploration compare against evolution strategies with respect to sample efficiency?

4.1 Comparing Parameter Space Noise to Action Space Noise

The added value of parameter space noise over action space noise is measured on both high-dimensional discrete-action environments and continuous control tasks. For the discrete environments, comparisons are made using DQN, while DDPG and TRPO are used on the continuous control tasks.

Discrete-action environments For discrete-action environments, we use the Arcade Learning Environment (ALE) [26] benchmark along with a standard DQN implementation. We compare a baseline DQN agent with ϵ -greedy action noise against a version of DQN with parameter noise. We linearly anneal ϵ from 1.0 to 0.1 over the first 1 million timesteps. For parameter noise, we adapt the scale using a simple heuristic that increases the scale if the KL divergence between perturbed and non-perturbed policy is less than the KL divergence between greedy and ϵ -greedy policy and decreases it otherwise (see subsection C.1 for details). By using this approach, we achieve a fair

comparison between action space noise and parameter space noise since the magnitude of the noise is similar and also avoid the introduction of an additional hyperparameter.

For parameter perturbation, we found it useful to reparametrize the network in terms of an explicit policy that represents the greedy policy π implied by the Q -values, rather than perturbing the Q -function directly. To represent the policy $\pi(a|s)$, we add a single fully connected layer after the convolutional part of the network, followed by a softmax output layer. Thus, π predicts a discrete probability distribution over actions, given a state. We find that perturbing π instead of Q results in more meaningful changes since we now define an explicit behavioral policy. In this setting, the Q -network is trained according to standard DQN practices. The policy π is trained by maximizing the probability of outputting the greedy action accordingly to the current Q -network. Essentially, the policy is trained to exhibit the same behavior as running greedy DQN. To rule out this double-headed version of DQN alone exhibits significantly different behavior, we always compare our parameter space noise approach against two baselines, regular DQN and two-headed DQN, both with ϵ -greedy exploration.

We furthermore randomly sample actions for the first 50 thousand timesteps in all cases to fill the replay buffer before starting training. Moreover, we found that parameter space noise performs better if it is combined with a bit of action space noise (we use a ϵ -greedy behavioral policy with $\epsilon = 0.01$ for the parameter space noise experiments). Full experimental details are described in subsection A.1.

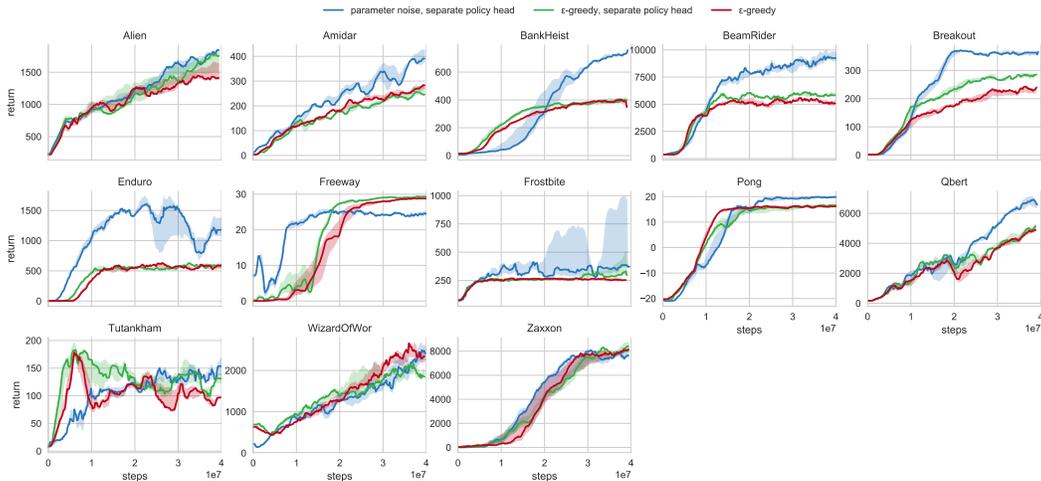


Figure 1: Median DQN returns for several ALE environment plotted over training steps.

We chose 21 games of varying complexity, according to the taxonomy presented by [27]. The learning curves are shown in Figure 1 for a selection of games (see Appendix D for full results). Each agent is trained for 40 M frames. The overall performance is estimated by running each configuration with three different random seeds, and we plot the median return (line) as well as the interquartile range (shaded area). Note that performance is evaluated on the exploratory policy since we are interested in its behavior especially.

On Enduro, which requires a player to continuously press a button in order to drive, parameter space noise finds this strategy significantly earlier than ϵ -greedy exploration. We also observe that parameter space noise achieves higher scores during training on Breakout, which can be explained by the fact that random action jittering, as induced by action space noise, is undesirable when the agent tries to hit the ball with precision. For Freeway, which requires an agent to cross a street before receiving any reward, parameter space exploration discovers this sparse reward much earlier than action space noise alone. Overall, our results show that using parameter space noise clearly outperforms the DQN baseline on 8 of the 21 games and performs comparably on the remaining ones for the final policy. However, learning usually starts much sooner when using parameter space noise. Finally, we also compare against a double-headed version of DQN with ϵ -greedy exploration to ensure that this change in architecture is not responsible for improved exploration, which our results confirm. Full results are available in Appendix D.

On a final note, proposed improvements to DQN like double DQN [28], prioritized experience replay [29], and dueling networks [30] are orthogonal to our improvements and would therefore likely improve results further. We leave the experimental validation of this theory to future work.

Continuous control environments We now compare parameter noise with action noise on the continuous control environments implemented in OpenAI Gym [31]. We use DDPG [14] as the RL algorithm for all environments with similar hyperparameters as outlined in the original paper except for the fact that layer normalization [23] is applied after each layer before the nonlinearity, which we found to be useful in either case and especially important for parameter space noise.

We compare the performance of the following configurations: (a) no noise at all, (b) uncorrelated additive Gaussian action space noise ($\sigma = 0.2$), (c) correlated additive Gaussian action space noise (Ornstein–Uhlenbeck process [17] with $\sigma = 0.2$), and (d) adaptive parameter space noise. In the case of parameter space noise, we adapt the scale so that the resulting change in action space is comparable to our baselines with uncorrelated Gaussian action space noise (see subsection C.2 for full details).

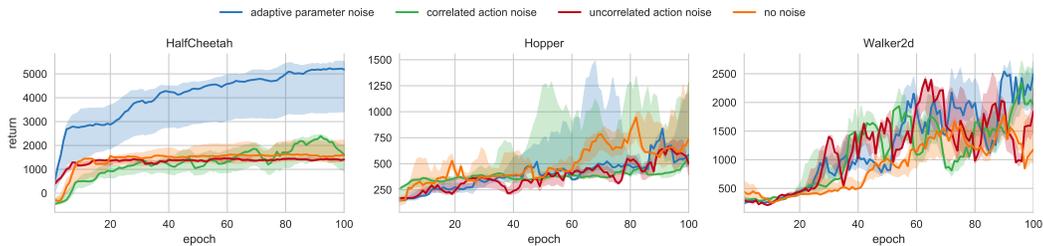


Figure 2: Median DDPG returns for continuous control environments plotted over epochs.

We evaluate the performance on several continuous control tasks. Figure 2 depicts the results for three exemplary environments. Each agent is trained for 1 M timesteps, where 1 epoch consists of 10 thousand timesteps. In order to make results comparable between configurations, we evaluate the performance of the agent every 10 thousand steps by using no noise for 20 episodes.

On *HalfCheetah*, parameter space noise achieves significantly higher returns than all other configurations. We find that, in this environment, all other exploration schemes quickly converge to a local optimum (in which the agent learns to flip on its back and then “wiggles” its way forward). Parameter space noise behaves similarly initially but still explores other options and quickly learns to break out of this sub-optimal behavior. Also notice that parameter space noise vastly outperforms correlated action space noise on this environment, clearly indicating that there is a significant difference between the two. On the remaining two environments, parameter space noise performs on par with other exploration strategies. Notice, however, that even if no noise is present, DDPG is capable of learning good policies. We find that this is representative for the remaining environments (see Appendix E for full results), which indicates that these environments do not require a lot of exploration to begin with due to their well-shaped reward function.

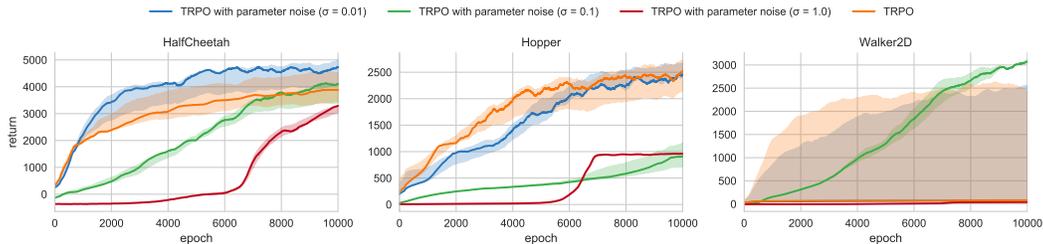


Figure 3: Median TRPO returns for continuous control environments plotted over epochs.

The results for TRPO are depicted in Figure 3. Interestingly, in the *Walker2D* environment, we see that adding parameter noise decreases the performance variance between seeds. This indicates that parameter noise aids in escaping local optima.

4.2 Does Parameter Space Noise Explore Efficiently?

The environments in the previous section required relatively little exploration. In this section, we evaluate whether parameter noise enables existing RL algorithms to learn on environments with very sparse rewards, where uncorrelated action noise generally fails [7, 4].

A scalable toy example We first evaluate parameter noise on a well-known toy problem, following the setup described in [7] as closely as possible. The environment consists of a chain of N states and the agent always starts in state s_2 , from where it can either move left or right. In state s_1 , the agent receives a small reward of $r = 0.001$ and a larger reward $r = 1$ in state s_N . Obviously, it is much easier to discover the small reward in s_1 than the large reward in s_N , with increasing difficulty as N grows. The environment is described in greater detail in subsection A.3.

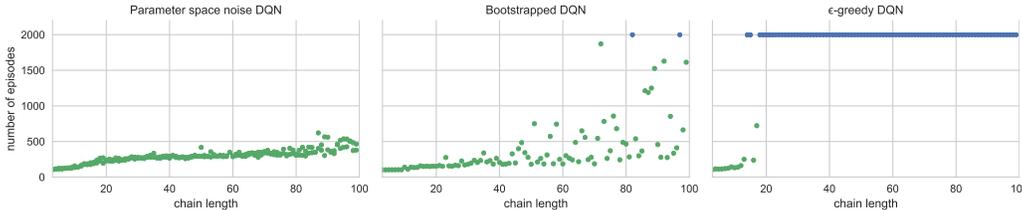


Figure 4: Median number of episodes before considered solved for DQN with different exploration strategies. Green indicates that the problem was solved whereas blue indicates that no solution was found within 2 K episodes. Note that less number of episodes before solved is better.

We compare adaptive parameter space noise DQN, bootstrapped DQN, and ϵ -greedy DQN. The chain length N is varied and for each N three different seeds are trained and evaluated. After each episode, we evaluate the performance of the current policy by performing a rollout with all noise disabled (in the case of bootstrapped DQN, we perform majority voting over all heads). The problem is considered solved if one hundred subsequent rollouts achieve the optimal return. We plot the median number of episodes before the problem is considered solved (we abort if the problem is still unsolved after 2 thousand episodes). Full experimental details are available in subsection A.3.

Figure 4 shows that parameter space noise clearly outperforms action space noise (which completely fails for moderately large N) and even outperforms the more computational expensive bootstrapped DQN.

Continuous control with sparse rewards We now make the continuous control environments more challenging for exploration. Instead of providing a reward at every timestep, we use environments that only yield a non-zero reward after significant progress towards a goal. More concretely, we consider the following environments from rllab² [22], modified according to [3]: (a) *SparseCartpoleSwingup*, which only yields a reward if the paddle is raised above a given threshold, (b) *SparseDoublePendulum*, which only yields a reward if the agent reaches the upright position, and (c) *SparseHalfCheetah*, which only yields a reward if the agent crosses a target distance, (d) *SparseMountainCar*, which only yields a reward if the agent drives up the hill, (e) *SwimmerGather*, yields a positive or negative reward upon reaching targets. For all tasks, we use a time horizon of $T = 500$ steps before resetting.

We consider both DDPG and TRPO to solve these environments (the exact experimental setup is described in subsection A.2). Figure 5 shows the performance of DDPG, while Figure 6 shows the performance of TRPO. The overall performance is estimated by running each configuration with five different random seeds, after which we plot the median return (line) as well as the interquartile range (shaded area).

For DDPG, *SparseDoublePendulum* seems to be easy to solve in general, with even no noise finding a successful policy relatively quickly. The results for *SparseCartpoleSwingup* and *SparseMountainCar* are more interesting: Here, only parameter space noise is capable of learning successful policies

²<https://github.com/openai/rllab>

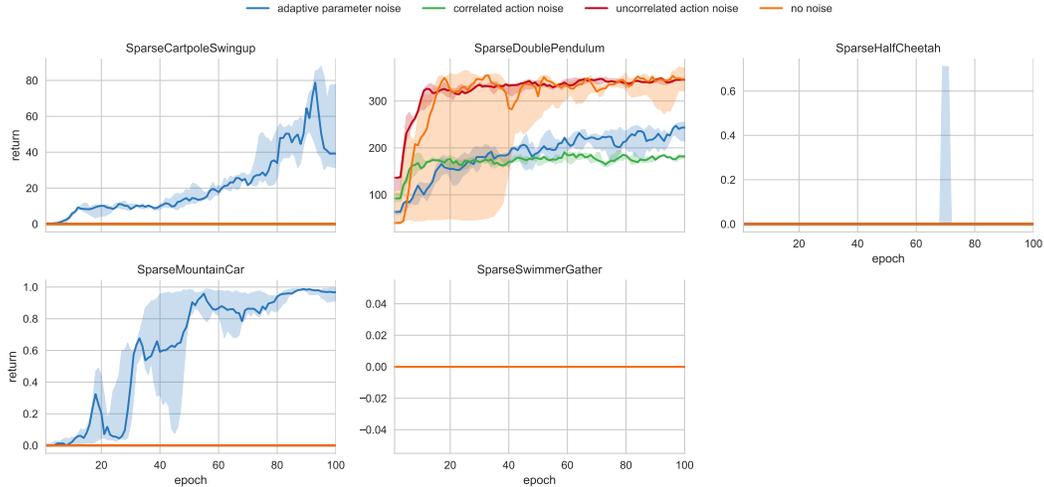


Figure 5: Median DDPG returns for environments with sparse rewards plotted over epochs.

since all other forms of noise, including correlated action space noise, never find states with non-zero rewards. For *SparseHalfCheetah*, DDPG at least finds the non-zero reward but never learns a successful policy from that signal. On the challenging *SwimmerGather* task, all configurations of DDPG fail.

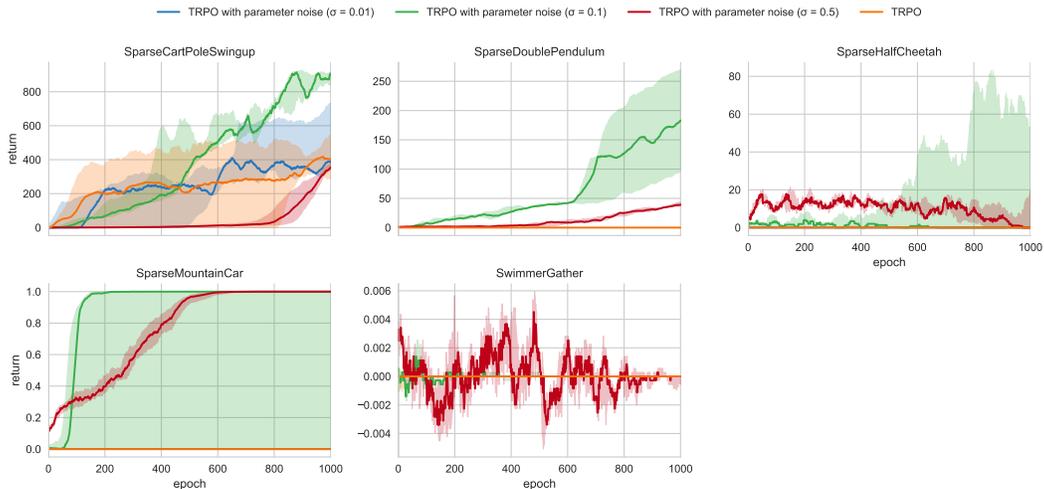


Figure 6: Median TRPO return for environments with sparse rewards plotted over epochs

The TRPO baseline uses only action noise by using a policy network that outputs the mean of a Gaussian distribution, while the variance is learned. These results show that adding parameter space noise aids in either learning much more consistently on these challenging environments (*SparseCartpoleSwingup*, *SparseDoublePendulum*, *SparseHalfCheetah*, and *SparseMountainCar*), or at least in reaching the sparse reward more often (*SwimmerGather*).

Our results clearly show that parameter space noise can be used to improve the exploration behavior of these off-the-shelf algorithms.

4.3 Is RL with Parameter Space Noise more Sample-efficient than ES?

Evolution strategies (ES) are closely related to our approach since both explore by introducing noise in the parameter space, which can lead to improved exploration behavior [9]. However, ES disregards

temporal information and uses black-box optimization to train the neural network. By combining parameter space noise with traditional RL algorithms, we can include temporal information as well rely on gradients computed by back-propagation for optimization while still benefiting from improved exploratory behavior. We now compare ES and traditional RL with parameter space noise directly.

We compare performance on the 21 ALE games that were used in subsection 4.1. The performance is estimated by running 10 episodes for each seed using the final policy with exploration disabled and computing the median returns. For ES, we use the numbers provided in [9], which were obtained after training on 1 000 M frames. For DQN, we use the same parameter space noise for exploration that was previously described and train on 40 M frames. Even though DQN with parameter space noise has been exposed to 25 times less data, it outperforms ES on 15 out of 21 Atari games (full results are available in Appendix D). Combined with the previously described results, this demonstrates that parameter space noise combines the desirable exploration properties of ES with the sample efficiency of traditional RL.

5 Related Work

The problem of exploration in reinforcement has been studied extensively. In general, one can differentiate between *undirected* and *directed* exploration. ϵ -greedy and Boltzmann exploration are both cases of undirected exploration. The same is obviously true for additive Gaussian action space noise for the continuous case. Directed exploration, on the other hand, can be count-based, error-based or recency-based. We refer the reader to [32, 16] for details on the respective technique.

A range of algorithms [33–35] have been proposed that guarantee near-optimal solutions after a number of steps that are polynomial in the number of states, number of actions, and the horizon time. However, in many real-world reinforcement learning problems both the state and action space are continuous and high dimensional so that, even with discretization, these algorithms become impractical. In the context of deep reinforcement learning, a large variety of techniques have been proposed to improve exploration [5, 3, 1, 7, 2, 36, 37]. However, all are non-trivial to implement and are often computational expensive.

The idea of perturbing the parameters of a policy has been proposed in [10] for policy gradient methods. The authors show that this form of perturbation generally outperforms random exploration and evaluate their exploration strategy with the REINFORCE [38] and Natural Actor-Critic [20] algorithms. However, their policies are relatively low-dimensional compared to modern deep architectures, they use environments with low-dimensional state spaces, and their contribution is strictly limited to the policy gradient case. In contrast, our method is applied and evaluated for both on and off-policy setting, we use high-dimensional policies, and environments with large state spaces.

Our work is also closely related to evolution strategies (ES) [39, 40], and especially neural evolution strategies (NES) [41, 8, 42–45]. In the context of policy optimization, our work is closely related to [11, 12]. More recently, ES was shown to even work for high-dimensional environments like Atari and OpenAI Gym continuous control problems [9]. However, ES generally disregards any temporal structure that may be present in trajectories and typically suffers from sample inefficiency.

In [46], the authors propose to use dropout to estimate model uncertainty. The authors show that their approach results in better exploration for a simple reinforcement learning problem. Their work can be related to ours since dropout can also be seen as a form of parameter space noise, where entire rows of the weight matrix are zeroed out. However, as shown in [7], this approach does not drive directed exploration (i.e. behavior that is consistent within episodes).

6 Conclusion

On the one hand, evolutionary methods discard temporal structure, which makes credit assignment more difficult and results in worse sample-efficiency. On the other hand, traditional RL methods often rely solely on unstructured action noise. This work shows that combining parameter perturbations with contemporary on- and off-policy deep RL algorithms such as DQN, DDPG, and TRPO allows for structured exploration while maintaining the properties of sample efficiency and exploitation of temporal structure that traditional RL approaches enjoy. We show that parameter noise can be applied to these off-the-shelf algorithms and often results in improved performance compared to action noise.

Experimental results further demonstrate that using parameter noise allows solving environments with very sparse rewards, in which action noise is unlikely to succeed. Results also indicate that RL with parameter noise exploration learns more efficiently than both RL and evolutionary strategies individually.

References

- [1] Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. #Exploration: A study of count-based exploration for deep reinforcement learning. *arXiv preprint arXiv:1611.04717*, 2016.
- [2] Georg Ostrovski, Marc G. Bellemare, Aäron van den Oord, and Rémi Munos. Count-based exploration with neural density models. *arXiv preprint arXiv:1703.01310*, 2017.
- [3] Rein Houthoofd, Xi Chen, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. VIME: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 1109–1117, 2016.
- [4] Joshua Achiam and Shankar Sastry. Surprise-based intrinsic motivation for deep reinforcement learning. *arXiv preprint arXiv:1703.01732*, 2017.
- [5] Bradley C. Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.
- [6] Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *ICML*, 2017.
- [7] Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped DQN. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 4026–4034, 2016.
- [8] Yi Sun, Daan Wierstra, Tom Schaul, and Jürgen Schmidhuber. Efficient natural evolution strategies. In *Genetic and Evolutionary Computation Conference, GECCO 2009, Proceedings, Montreal, Québec, Canada, July 8-12, 2009*, pages 539–546, 2009.
- [9] Tim Salimans, Jonathan Ho, Xi Chen, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- [10] Thomas Rückstieß, Martin Felder, and Jürgen Schmidhuber. State-dependent exploration for policy gradient methods. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases ECML/PKDD*, pages 234–249, 2008.
- [11] Jens Kober and Jan Peters. Policy search for motor primitives in robotics. In *Advances in Neural Information Processing Systems 21 (NIPS)*, pages 849–856, 2008.
- [12] Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559, 2010.
- [13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [14] 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. *CoRR*, abs/1509.02971, 2015.
- [15] John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1889–1897, 2015.
- [16] Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*, volume 135. MIT Press Cambridge, 1998.
- [17] George E Uhlenbeck and Leonard S Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930.
- [18] 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)*, pages 1889–1897, 2015.

- [19] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [20] Jan Peters and Stefan Schaal. Natural actor-critic. *Neurocomputing*, 71(7-9):1180–1190, 2008.
- [21] Sham Kakade. A natural policy gradient. *Advances in neural information processing systems*, 14:1531–1538, 2001.
- [22] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 1329–1338, 2016.
- [23] Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *CoRR*, abs/1607.06450, 2016.
- [24] Thomas Rückstieß, Martin Felder, and Jürgen Schmidhuber. State-dependent exploration for policy gradient methods. *Machine Learning and Knowledge Discovery in Databases*, pages 234–249, 2008.
- [25] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [26] Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- [27] Marc G Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems 29 (NIPS)*, pages 1471–1479, 2016.
- [28] Hado V Hasselt. Double Q-learning. In *Advances in Neural Information Processing Systems 23 (NIPS)*, pages 2613–2621, 2010.
- [29] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
- [30] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
- [31] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [32] Sebastian B. Thrun. Efficient exploration in reinforcement learning. Technical report, Pittsburgh, PA, USA, 1992.
- [33] Michael J. Kearns and Satinder P. Singh. Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2-3):209–232, 2002.
- [34] Ronen I. Brafman and Moshe Tennenholtz. R-MAX - A general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3:213–231, 2002.
- [35] Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in Neural Information Processing Systems 21 (NIPS)*, pages 89–96, 2008.
- [36] Sainbayar Sukhbaatar, Ilya Kostrikov, Arthur Szlam, and Rob Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. *arXiv preprint arXiv:1703.05407*, 2017.
- [37] Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. In *Proceedings of the 33rd International Conference on Machine Learning, ICML*, pages 2377–2386, 2016.
- [38] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.
- [39] Ingo Rechenberg and Manfred Eigen. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog Stuttgart, 1973.
- [40] Hans-Paul Schwefel. *Numerische Optimierung von Computermodellen mittels der Evolutionsstrategie*, volume 1. Birkhäuser, Basel Switzerland, 1977.
- [41] Yi Sun, Daan Wierstra, Tom Schaul, and Jürgen Schmidhuber. Stochastic search using the natural gradient. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, pages 1161–1168, 2009.

- [42] Tobias Glasmachers, Tom Schaul, and Jürgen Schmidhuber. A natural evolution strategy for multi-objective optimization. In *Parallel Problem Solving from Nature - PPSN XI, 11th International Conference, Kraków, Poland, September 11-15, 2010, Proceedings, Part I*, pages 627–636, 2010.
- [43] Tobias Glasmachers, Tom Schaul, Yi Sun, Daan Wierstra, and Jürgen Schmidhuber. Exponential natural evolution strategies. In *Genetic and Evolutionary Computation Conference, GECCO 2010, Proceedings, Portland, Oregon, USA, July 7-11, 2010*, pages 393–400, 2010.
- [44] Tom Schaul, Tobias Glasmachers, and Jürgen Schmidhuber. High dimensions and heavy tails for natural evolution strategies. In *13th Annual Genetic and Evolutionary Computation Conference, GECCO 2011, Proceedings, Dublin, Ireland, July 12-16, 2011*, pages 845–852, 2011.
- [45] Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, Jan Peters, and Jürgen Schmidhuber. Natural evolution strategies. *Journal of Machine Learning Research*, 15(1):949–980, 2014.
- [46] Yarın Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 1050–1059, 2016.
- [47] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [48] Ananth Ranganathan. The Levenberg-Marquardt algorithm. *Tutorial on LM algorithm*, pages 1–5, 2004.

A Experimental Setup

A.1 Arcade Learning Environment (ALE)

For ALE [26], the network architecture as described in [13] is used. This consists of 3 convolutional layers (32 filters of size 8×8 and stride 4, 64 filters of size 4×4 and stride 2, 64 filters of size 3×3 and stride 1) followed by 1 hidden layer with 512 units followed by a linear output layer with one unit for each action. ReLUs are used in each layer, while layer normalization [23] is used in the fully connected part of the network. For parameter space noise, we also include a second head after the convolutional stack of layers. This head determines a policy network with the same architecture as the Q -value network, except for a softmax output layer. The target networks are updated every 10 K timesteps. The Q -value network is trained using the Adam optimizer [47] with a learning rate of 10^{-4} and a batch size of 32. The replay buffer can hold 1 M state transitions. For the ϵ -greedy baseline, we linearly anneal ϵ from 1 to 0.1 over the first 1 M timesteps. For parameter space noise, we adaptively scale the noise to have a similar effect in action space (see subsection C.1 for details), effectively ensuring that the maximum KL divergence between perturbed and non-perturbed π is softly enforced. The policy is perturbed at the beginning of each episode and the standard deviation is adapted as described in Appendix C every 50 timesteps. Notice that we only perturb the policy head after the convolutional part of the network (i.e. the fully connected part, which is also why we only include layer normalization in this part of the network). To avoid getting stuck (which can potentially happen for a perturbed policy), we also use ϵ -greedy action selection with $\epsilon = 0.01$. In all cases, we perform 50 K random actions to collect initial data for the replay buffer before training starts. We set $\gamma = 0.99$, clip rewards to be in $[-1, 1]$, and clip gradients for the output layer of Q to be within $[-1, 1]$. For observations, each frame is down-sampled to 84×84 pixels, after which it is converted to grayscale. The actual observation to the network consists of a concatenation of 4 subsequent frames. Additionally, we use up to 30 noop actions at the beginning of the episode. This setup is identical to what is described in [13].

A.2 Continuous Control

For DDPG, we use a similar network architecture as described in [14]: both the actor and critic use 2 hidden layers with 64 ReLU units each. For the critic, actions are not included until the second hidden layer. Layer normalization [23] is applied to all layers. The target networks are soft-updated with $\tau = 0.001$. The critic is trained with a learning rate of 10^{-3} while the actor uses a learning rate of 10^{-4} . Both actor and critic are updated using the Adam optimizer [47] with batch sizes of 128. The critic is regularized using an $L2$ penalty with 10^{-2} . The replay buffer holds 100 K state transitions and $\gamma = 0.99$ is used. Each observation dimension is normalized by an online estimate of

the mean and variance. For parameter space noise with DDPG, we adaptively scale the noise to be comparable to the respective action space noise (see subsection C.2). For dense environments, we use action space noise with $\sigma = 0.2$ (and a comparable adaptive noise scale). Sparse environments use an action space noise with $\sigma = 0.6$ (and a comparable adaptive noise scale).

TRPO uses a step size of $\delta_{\text{KL}} = 0.01$, a policy network of 2 hidden layers with 32 tanh units for the nonlocomotion tasks, and 2 hidden layers of 64 tanh units for the locomotion tasks. The Hessian calculation is subsampled with a factor of 0.1, $\gamma = 0.99$, and the batch size per epoch is set to 5 K timesteps. The baseline is a learned linear transformation of the observations.

The following environments from OpenAI Gym³ [31] are used:

- *HalfCheetah* ($\mathcal{S} \subset \mathbb{R}^{17}, \mathcal{A} \subset \mathbb{R}^6$),
- *Hopper* ($\mathcal{S} \subset \mathbb{R}^{11}, \mathcal{A} \subset \mathbb{R}^3$),
- *InvertedDoublePendulum* ($\mathcal{S} \subset \mathbb{R}^{11}, \mathcal{A} \subset \mathbb{R}$),
- *InvertedPendulum* ($\mathcal{S} \subset \mathbb{R}^4, \mathcal{A} \subset \mathbb{R}$),
- *Reacher* ($\mathcal{S} \subset \mathbb{R}^{11}, \mathcal{A} \subset \mathbb{R}^2$),
- *Swimmer* ($\mathcal{S} \subset \mathbb{R}^8, \mathcal{A} \subset \mathbb{R}^2$), and
- *Walker2D* ($\mathcal{S} \subset \mathbb{R}^{17}, \mathcal{A} \subset \mathbb{R}^6$).

For the sparse tasks, we use the following environments from rllab⁴ [22], modified according to [3]:

- *SparseCartpoleSwingup* ($\mathcal{S} \subset \mathbb{R}^4, \mathcal{A} \subset \mathbb{R}$), which only yields a reward if the paddle is raised above a given threshold,
- *SparseHalfCheetah* ($\mathcal{S} \subset \mathbb{R}^{17}, \mathcal{A} \subset \mathbb{R}^6$), which only yields a reward if the agent crosses a distance threshold,
- *SparseMountainCar* ($\mathcal{S} \subset \mathbb{R}^2, \mathcal{A} \subset \mathbb{R}$), which only yields a reward if the agent drives up the hill,
- *SparseDoublePendulum* ($\mathcal{S} \subset \mathbb{R}^6, \mathcal{A} \subset \mathbb{R}$), which only yields a reward if the agent reaches the upright position, and
- *SwimmerGather* ($\mathcal{S} \subset \mathbb{R}^{33}, \mathcal{A} \subset \mathbb{R}^2$), which yields a positive or negative reward upon reaching targets.

A.3 Chain Environment

We follow the state encoding proposed in [7] and use $\phi(s_t) = (\mathbb{1}\{x \leq s_t\})$ as the observation, where $\mathbb{1}$ denotes the indicator function. DQN is used with a very simple network to approximate the Q -value function that consists of 2 hidden layers with 16 ReLU units. Layer normalization [23] is used for all hidden layers before applying the nonlinearity. Each agent is then trained for up to 2 K episodes. The chain length N is varied and for each N three different seeds are trained and evaluated. After each episode, the performance of the current policy is evaluated by sampling a trajectory with noise disabled (in the case of bootstrapped DQN, majority voting over all heads is performed). The problem is considered solved if one hundred subsequent trajectories achieve the optimal episode return. Figure 7 depicts the environment.

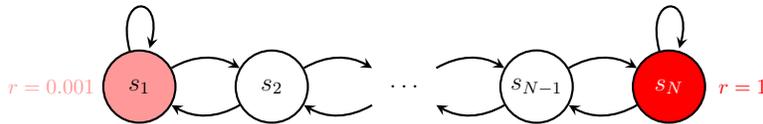


Figure 7: Simple and scalable environment to test for exploratory behavior [7].

We compare adaptive parameter space noise DQN, bootstrapped DQN [7] (with $K = 20$ heads and Bernoulli masking with $p = 0.5$), and ϵ -greedy DQN (with ϵ linearly annealed from 1.0 to 0.1 over

³<https://github.com/openai/gym>

⁴<https://github.com/openai/rllab>

the first one hundred episodes). For adaptive parameter space noise, we only use a single head and perturb Q directly, which works well in this setting. Parameter space noise is adaptively scaled so that $\delta \approx 0.05$. In all cases, $\gamma = 0.999$, the replay buffer holds 100 K state transitions, learning starts after 5 initial episodes, the target network is updated every 100 timesteps, and the network is trained using the Adam optimizer [47] with a learning rate of 10^{-3} and a batch size of 32.

B Parameter Space Noise for On-policy Methods

Policy gradient methods optimize $\mathbb{E}_{\tau \sim (\pi, p)}[R(\tau)]$. Given a stochastic policy $\pi_\theta(a|s)$ with $\theta \sim \mathcal{N}(\phi, \Sigma)$, the expected return can be expanded using likelihood ratios and the reparametrization trick [25] as

$$\nabla_{\phi, \Sigma} \mathbb{E}_\tau [R(\tau)] = \nabla_{\phi, \Sigma} \mathbb{E}_{\theta \sim \mathcal{N}(\phi, \Sigma)} \left[\sum_\tau p(\tau|\theta) R(\tau) \right] \quad (3)$$

$$= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} \nabla_{\phi, \Sigma} \left[\sum_\tau p(\tau|\phi + \epsilon \Sigma) R(\tau) \right] \quad (4)$$

$$= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I), \tau} \left[\sum_{t=0}^{T-1} \nabla_{\phi, \Sigma} \log \pi(a_t|s_t; \phi + \epsilon \Sigma) R_t(\tau) \right] \quad (5)$$

$$\approx \frac{1}{N} \sum_{\epsilon^i, \tau^i} \left[\sum_{t=0}^{T-1} \nabla_{\phi, \Sigma} \log \pi(a_t|s_t; \phi + \epsilon^i \Sigma) R_t(\tau^i) \right] \quad (6)$$

for N samples $\epsilon^i \sim \mathcal{N}(0, I)$ and $\tau^i \sim (\pi_{\phi + \epsilon^i \Sigma}, p)$, with $R_t(\tau^i) = \sum_{t'=t}^T \gamma^{t'-t} r_{t'}^i$. This also allows us to subtract a variance-reducing baseline b_t^i , leading to

$$\nabla_{\phi, \Sigma} \mathbb{E}_\tau [R(\tau)] \approx \frac{1}{N} \sum_{\epsilon^i, \tau^i} \left[\sum_{t=0}^{T-1} \nabla_{\phi, \Sigma} \log \pi(a_t|s_t; \phi + \epsilon^i \Sigma) (R_t(\tau^i) - b_t^i) \right]. \quad (7)$$

C Adaptive Scaling

Parameter space noise requires us to pick a suitable scale σ . This can be problematic since the scale will highly depend on the specific network architecture, and is likely to vary over time as parameters become more sensitive as learning progresses. Additionally, while it is easy to intuitively grasp the scale of action space noise, it is far harder to understand the scale in parameter space.

We propose a simple solution that resolves all aforementioned limitations in an easy and straightforward way. This is achieved by adapting the scale of the parameter space noise over time, thus using a time-varying scale σ_k . Furthermore, σ_k is related to the action space variance that it induces, and updated accordingly. Concretely, we use the following simple heuristic to update σ_k every K timesteps:

$$\sigma_{k+1} = \begin{cases} 1.01 \cdot \sigma_k, & \text{if } d(\pi, \tilde{\pi}) < \delta \\ \frac{1}{1.01} \cdot \sigma_k, & \text{otherwise,} \end{cases} \quad (8)$$

where $d(\cdot, \cdot)$ denotes some distance between the non-perturbed and perturbed policy (thus measuring in action space) and δ denotes some threshold value. This idea is based on the Levenberg-Marquardt heuristic [48]. The concrete distance measure and appropriate choice of δ depends on the policy representation. In the following sections, we outline our choice of $d(\cdot, \cdot)$ for methods that do (DDPG and TRPO) and do not (DQN) use behavioral policies.

C.1 A Distance Measure for DQN

For DQN, the policy is defined implicitly by the Q -value function. Unfortunately, this means that a naïve distance measure between Q and \tilde{Q} has pitfalls. For example, assume that the perturbed policy has only changed the bias of the final layer, thus adding a constant value to each action’s Q -value. In this case, a naïve distance measure like the norm $\|Q - \tilde{Q}\|_2$ would be nonzero, although the policies

π and $\tilde{\pi}$ (implied by Q and \tilde{Q} , respectively) are exactly equal. This equally applies to the case where DQN as two heads, one for Q and one for π .

We therefore use a probabilistic formulation⁵ for both the non-perturbed and perturbed policies: $\pi, \tilde{\pi} : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ by applying the softmax function over predicted Q values: $\pi(s) = \exp Q_i(s) / \sum_i \exp Q_i(s)$, where $Q_i(\cdot)$ denotes the Q -value of the i -th action. $\tilde{\pi}$ is defined analogously but uses the perturbed \tilde{Q} instead (or the perturbed head for π). Using this probabilistic formulation of the policies, we can now measure the distance in action space:

$$d(\pi, \tilde{\pi}) = D_{\text{KL}}(\pi \parallel \tilde{\pi}), \quad (9)$$

where $D_{\text{KL}}(\cdot \parallel \cdot)$ denotes the Kullback-Leibler (KL) divergence. This formulation effectively normalizes the Q -values and therefore does not suffer from the problem previously outlined.

We can further relate this distance measure to ϵ -greedy action space noise, which allows us to fairly compare the two approaches and also avoids the need to pick an additional hyperparameter δ . More concretely, the KL divergence between a greedy policy $\pi(s, a) = 1$ for $a = \operatorname{argmax}_{a'} Q(s, a')$ and $\pi(s, a) = 0$ otherwise and an ϵ -greedy policy $\tilde{\pi}(s, a) = 1 - \epsilon + \frac{\epsilon}{|\mathcal{A}|}$ for $a = \operatorname{argmax}_{a'} Q(s, a')$ and $\tilde{\pi}(s, a) = \frac{\epsilon}{|\mathcal{A}|}$ otherwise is $D_{\text{KL}}(\pi \parallel \tilde{\pi}) = -\log(1 - \epsilon + \frac{\epsilon}{|\mathcal{A}|})$, where $|\mathcal{A}|$ denotes the number of actions (this follows immediately from the definition of the KL divergence for discrete probability distributions). We can use this distance measure to relate action space noise and parameter space noise to have similar distances, by adaptively scaling σ so that it matches the KL divergence between greedy and ϵ -greedy policy, thus setting $\delta := -\log(1 - \epsilon + \frac{\epsilon}{|\mathcal{A}|})$.

C.2 A Distance Measure for DDPG

For DDPG, we relate noise induced by parameter space perturbations to noise induced by additive Gaussian noise. To do so, we use the following distance measure between the non-perturbed and perturbed policy:

$$d(\pi, \tilde{\pi}) = \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbb{E}_s \left[(\pi(s)_i - \tilde{\pi}(s)_i)^2 \right]}, \quad (10)$$

where $\mathbb{E}_s[\cdot]$ is estimated from a batch of states from the replay buffer and N denotes the dimension of the action space (i.e. $\mathcal{A} \subset \mathbb{R}^N$). It is easy to show that $d(\pi, \pi + \mathcal{N}(0, \sigma^2 I)) = \sigma$. Setting $\delta := \sigma$ as the adaptive parameter space threshold thus results in effective action space noise that has the same standard deviation as regular Gaussian action space noise.

C.3 A Distance Measure for TRPO

In order to scale the noise for TRPO, we adapt the sampled noise vectors $\epsilon\sigma$ by computing a natural step $H^{-1}\epsilon\sigma$. We essentially compute a trust region around the noise direction to ensure that the perturbed policy $\tilde{\pi}$ remains sufficiently close to the non-perturbed version via

$$E_{s \sim \rho_{\tilde{\theta}}} [D_{\text{KL}}(\pi_{\tilde{\theta}}(\cdot|s) \parallel \pi_{\theta}(\cdot|s))] \leq \delta_{\text{KL}}.$$

Concretely, this is computed through the conjugate gradient algorithm, combined with a line search along the noise direction to ensure constraint conformation, as described in Appendix C of [15].

D Additional Results on ALE

Figure 8 provide the learning curves for all 21 Atari games.

Table 1 compares the final performance of ES after 1 000 M frames to the final performance of DQN with ϵ -greedy exploration and parameter space noise exploration after 40 M frames. In all cases, the performance is estimated by running 10 episodes with exploration disabled. We use the numbers reported in [9] for ES and report the median return across three seeds for DQN.

⁵It is important to note that we use this probabilistic formulation only for the sake of defining a well-behaved distance measure. The actual policy used for rollouts is still deterministic.

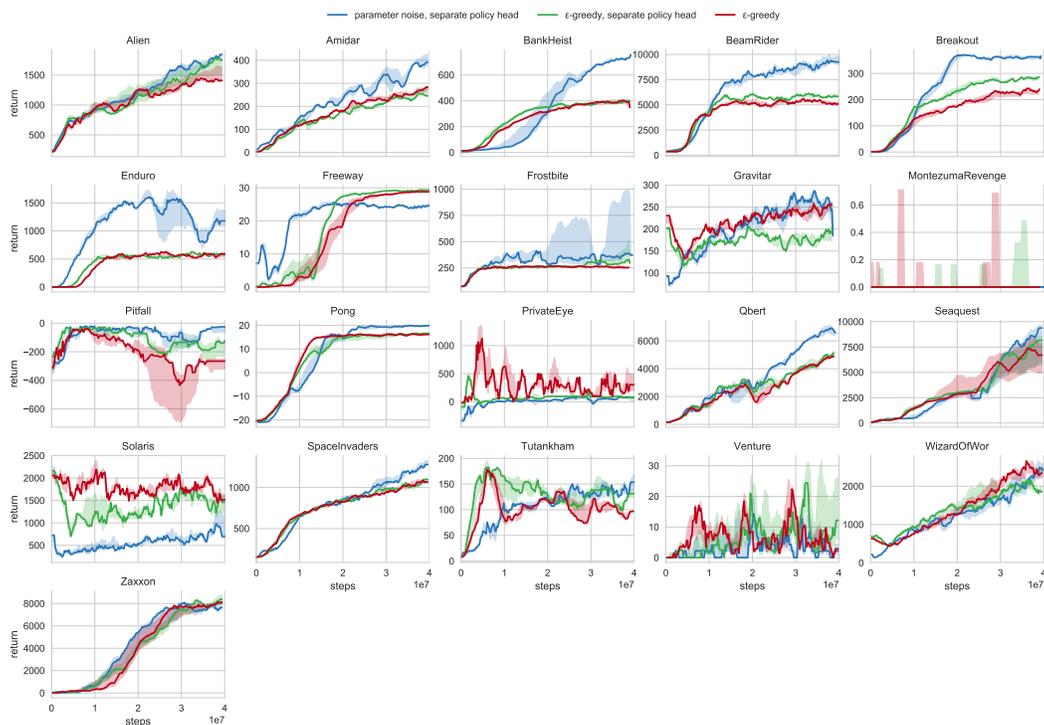


Figure 8: Median DQN returns for all ALE environment plotted over training steps.

Table 1: Performance comparison between Evolution Strategies (ES) as reported in [9], DQN with ϵ -greedy, and DQN with parameter space noise (this paper). ES was trained on 1 000 M, while DQN was trained on only 40 M frames.

Game	ES	DQN w/ ϵ -greedy	DQN w/ param noise
Alien	994.0	1535.0	2070.0
Amidar	112.0	281.0	403.5
BankHeist	225.0	510.0	805.0
BeamRider	744.0	8184.0	7884.0
Breakout	9.5	406.0	390.5
Enduro	95.0	1094	1672.5
Freeway	31.0	32.0	31.5
Frostbite	370.0	250.0	1310.0
Gravitar	805.0	300.0	250.0
MontezumaRevenge	0.0	0.0	0.0
Pitfall	0.0	-73.0	-100.0
Pong	21.0	21.0	20.0
PrivateEye	100.0	133.0	100.0
Qbert	147.5	7625.0	7525.0
Seaquest	1390.0	8335.0	8920.0
Solaris	2090.0	720.0	400.0
SpaceInvaders	678.5	1000.0	1205.0
Tutankham	130.3	109.5	181.0
Venture	760.0	0	0
WizardOfWor	3480.0	2350.0	1850.0
Zaxxon	6380.0	8100.0	8050.0

E Additional Results on Continuous Control with Shaped Rewards

For completeness, we provide the plots for all evaluated environments with dense rewards. The results are depicted in Figure 9.

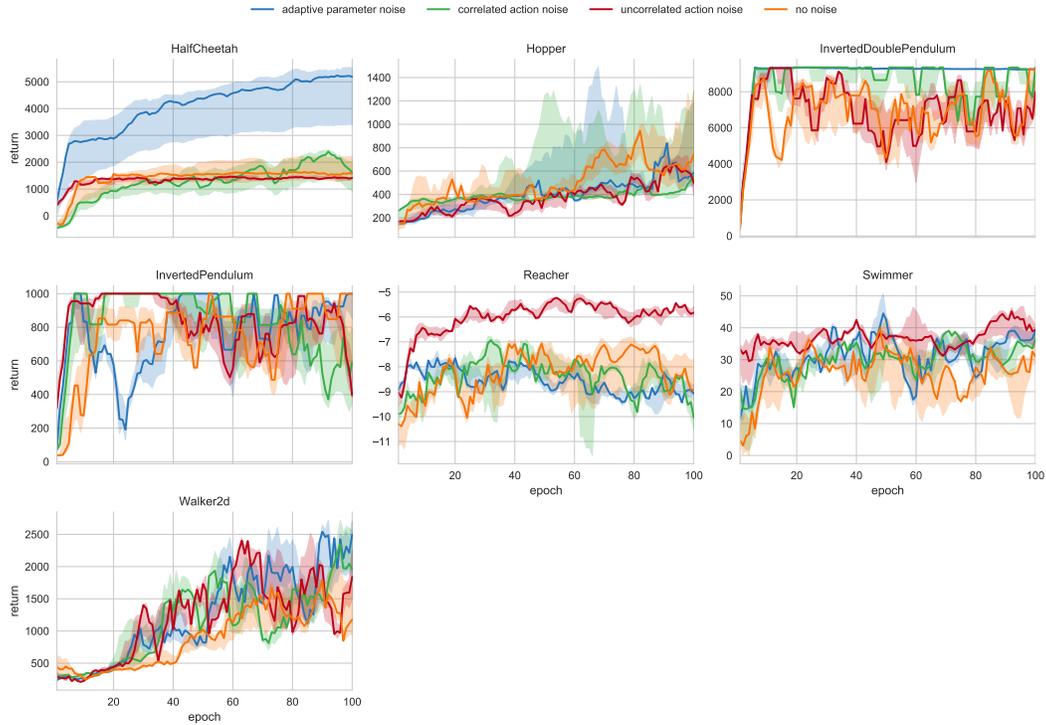


Figure 9: Median DDPG returns for all evaluated environments with dense rewards plotted over epochs.

The results for *InvertedPendulum* and *InvertedDoublePendulum* are very noisy due to the fact that a small change in policy can easily degrade performance significantly, and thus hard to read. Interestingly, adaptive parameter space noise achieves the most stable performance on *InvertedDoublePendulum*. Overall, performance is comparable to other exploration approaches. Again, no noise in either the action nor the parameter space achieves comparable results, indicating that these environments combined with DDPG are not well-suited to test for exploration.

F Additional Results on Continuous Control with Sparse Rewards

The performance of TRPO with noise scaled according to the parameter curvature, as defined in Section C.3 is shown in Figure 10.

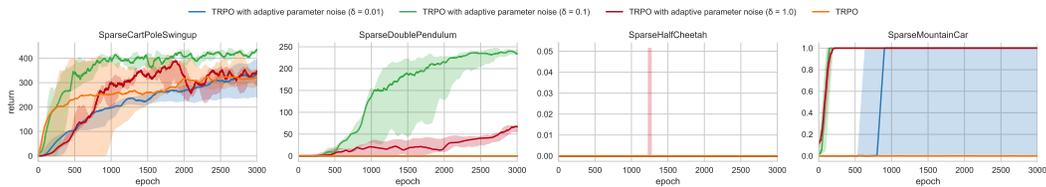


Figure 10: Median TRPO returns with three different environments with sparse rewards plotted over epochs.