

Colored Noise in PPO: Improved Exploration and Performance through Correlated Action Sampling

Jakob Hollenstein¹, Georg Martius^{3,4}, Justus Piater^{1,2}

¹Department of Computer Science, University of Innsbruck, Austria

²Digital Science Center, University of Innsbruck, Austria

³Max Planck Institute for Intelligent Systems, Tübingen, Germany

⁴Department of Computer Science, University of Tübingen, Germany

{jakob.hollenstein,justus.piater}@uibk.ac.at
georg.martius@uni-tuebingen.de

Abstract

Proximal Policy Optimization (PPO), a popular on-policy deep reinforcement learning method, employs a stochastic policy for exploration. In this paper, we propose a colored noise-based stochastic policy variant of PPO. Previous research highlighted the importance of temporal correlation in action noise for effective exploration in *off-policy* reinforcement learning. Building on this, we investigate whether correlated noise can also enhance exploration in *on-policy* methods like PPO. We discovered that correlated noise for action selection improves learning performance and outperforms the currently popular uncorrelated white noise approach in on-policy methods. Unlike off-policy learning, where pink noise was found to be highly effective, we found that a colored noise, intermediate between white and pink, performed best for on-policy learning in PPO. We examined the impact of varying the amount of data collected for each update by modifying the number of parallel simulation environments for data collection and observed that with a larger number of parallel environments, more strongly correlated noise is beneficial. Due to the significant impact and ease of implementation, we recommend switching to correlated noise as the default noise source in PPO.

1 Introduction

Exploration plays a crucial role in deep reinforcement learning, particularly in continuous action space applications like robotics, where the number of states and actions is infinite. In the continuous action space setting, exploration is typically achieved by introducing variability around the mean action proposed by the policy. For instance, deterministic off-policy algorithms such as TD3 (Fujimoto, van Hoof, and Meger 2018) and DDPG (Lillicrap et al. 2016) use additive action noise, while stochastic off-policy algorithms such as SAC (Haarnoja et al. 2019) and MPO (Abdolmaleki et al. 2018) employ a Gaussian distributed policy to sample actions and induce variation. On-policy algorithms such as TRPO (Schulman et al. 2015) and PPO (Schulman et al. 2017) follow a

similar approach by sampling variations from a Gaussian distribution.

While previous research (Hollenstein et al. 2022; Eberhard et al. 2023; Raffin, Kober, and Stulp 2021) has emphasized the importance of temporal correlation in action noise for exploration in *off-policy* algorithms, *on-policy* algorithms such as PPO do not incorporate correlated action noise and instead sample uncorrelated variations from a Gaussian distributed policy.

PPO, as an on-policy deep reinforcement learning algorithm, offers several advantages over off-policy algorithms: by not relying on replay buffers, on-policy algorithms exhibit fewer instabilities due to distributional shift between the distribution of previously collected data and the current policy-induced distribution. Moreover, on-policy algorithms mitigate Q-function divergence issues since the policy evaluation can leverage data collected under the current policy. These benefits make on-policy methods particularly advantageous when environment samples are inexpensive to obtain. However, exploration in on-policy methods needs to be induced by changes to the environment (exploration reward) or by changes to the policy itself, as, by definition, on-policy methods learn from data induced by the current policy. Consequently, altering the exploration technique without changing the environment implies a change in the policy distribution.

Contributions In this paper, we extend previous work on correlated action-noise in *off-policy* RL to *on-policy* RL methods: we introduce and empirically evaluate a modification to PPO incorporating temporally correlated colored noise into the stochastic policy’s distribution. Utilizing the re-parameterization trick, we maintain a Gaussian distributed behavior while injecting correlated noise instead of uncorrelated noise. This introduces a new hyperparameter, the *noise color* β parameterizing the correlated noise.

Our experiments demonstrate that adopting colored noise enhances performance in the majority of tested environments. In the off-policy setting Eberhard et al. (2023) found a common noise color, pink noise, to significantly outperform the previous default of uncorrelated white noise. Surprisingly we found the optimal common colored noise for improved performance — for the *on-policy* method PPO — to lie *between* white and pink noise ($\beta = 0.5$, see Sec. 2). Furthermore, we investigated the effect of update dataset size on the efficacy

of the noise colors: as the number of parallel data collection environments increases, and thus the update dataset size increases, we observe a trend towards more correlated noise. This provides a plausible explanation for the preference of more correlated noise in the off-policy setting. However, for the benchmarks considered in this work, our results indicate that utilizing about four parallel environments, resulting in 8192 samples per update step (see Sec. A11 for the relation of N_{steps} and N_{env}), together with noise in between white and pink noise, is the most efficient approach.

1.1 Related Work

Reinforcement learning is a promising technique for solving complex control problems. An important milestone for the development of deep reinforcement learning was the work of Mnih et al. (2015) showing human level performance of reinforcement learning on the Atari Games benchmark. In contrast to the Atari Games benchmarks which has complicated observation spaces but discrete action spaces, reinforcement learning for robotic applications has to deal with continuous action spaces, for example in the benchmarks provided by Tassa et al. (2018), where the task is to control simulated robots. Because the action selection process is more complicated in continuous spaces, policy search methods (Williams 1992) are favored. This includes off-policy methods such as DDPG (Lillicrap et al. 2016), TD3 (Fujimoto, van Hoof, and Meger 2018), MPO (Abdolmaleki et al. 2018) or SAC (Haarnoja et al. 2019) which can learn from arbitrary data and on-policy methods such as TRPO (Schulman et al. 2015) and PPO (Schulman et al. 2017) which iteratively improve the policy using data collected by the current iteration of the policy. On-policy methods exhibit better convergence properties and behave more stably but require larger amounts of training data.

Exploration, the problem of discovering better action sequences, is of pivotal importance in RL and is tackled in many ways. Amin et al. (2021), Ladosz et al. (2022) and Yang et al. (2022) provide recent surveys of the exploration research landscape in deep reinforcement learning. In this work, we turn to the simple yet effective method which forms the backbone of most RL algorithms: undirected noisy exploration.

In off-policy methods, exploration is often achieved by perturbing the action selection process, for example by perturbing the parameters of the policy (Plappert et al. 2018) or by additively perturbing the action, i.e., by adding *action noise*. This action perturbation can be done, for example, using temporally uncorrelated Gaussian noise, or temporally correlated noise generated by an Ornstein-Uhlenbeck process (Uhlenbeck and Ornstein 1930). Previous research has shown that increasing the temporal correlation of actions, i.e., switching from Gaussian noise to Ornstein-Uhlenbeck noise, tends to increase state space coverage, but an increase in state space coverage can be beneficial or harmful for policy learning; this depends on the task environment and its dynamics (Hollenstein et al. 2021, 2022). Another method for temporally correlated perturbations was introduced by Raffin, Kober, and Stulp (2021) where the “action noise” is generated deterministically by a function only dependent on

the state, but the function parameters are randomly changed after a number of steps. This approach is extended by Chiappa et al. (2023) to exploit the hidden layers of the policy, to induce correlation between the action dimensions.

While different environments respond differently to the temporal correlation of noise, Eberhard et al. (2023) found *pink noise* to act as a kind of middle ground between uncorrelated white Gaussian noise, and correlated Brownian-motion-like noise, e.g., red noise and Ornstein-Uhlenbeck noise. This middle ground, while not always the perfect choice, was found to be a much better default choice.

Petrazzini and Antonelo (2021) propose to improve exploration by changing the policy distribution of PPO from a Gaussian distribution to a Beta distribution. Similarly, we propose to change the policy distribution of PPO, but instead of a Beta distribution we propose to keep the Gaussian distribution and bias the sampling to be temporally correlated using a colored-noise process.

Algorithmic details and hyperparameter choices are known to be important for PPO and were empirically analyzed by Andrychowicz et al. (2021) and Engstrom et al. (2020). Similarly, we performed a comprehensive empirical evaluation of the impact of different noise correlation settings and the setting of parallel data collection. From this evaluation, we recommend switching the default noise process to a correlated noise process with $\beta = 0.5$.

2 Method

2.1 Colored Noise

Colored noise is a class of noise that exhibits temporal correlation and is characterized by a change of $\frac{1}{f^\beta}$ in its power spectral density (PSD) components, where f denotes the frequency and β determines the “color” of the noise. The power spectral density (PSD) of a sequence consists of the squared magnitude of the frequency components of the sequence’s Fourier transform. Given sequences of noise samples $\tau_\epsilon^{(i)} = \{\epsilon_1^{(i)}, \dots, \epsilon_T^{(i)}\}$, and the power spectral density of each sequence $|\mathcal{F}(\tau_\epsilon^{(i)})|^2$, the *expected PSD* is calculated by averaging over each trajectory’s PSD: $\mathbb{E}_i[|\mathcal{F}(\tau_\epsilon^{(i)})|^2]$.




Figure 1: Two-dimensional random walks caused by colored noise of different β . Lower β values cause more energy in high frequency parts of the power spectral density, causing the random walk to change direction more frequently and thus causing more local and less global exploration. Higher values of β result in more energy in the lower frequencies of the power spectral density. This translates to random walks that change direction less frequently, and thus explore more globally.

Uncorrelated noise is called *white noise* and has a constant, flat line, expected power spectral density ($\beta = 0$). The expected PSD of *pink noise* decreases with $1/f$ ($\beta = 1$) and for *red noise* the decrease happens at a rate of $1/f^2$ ($\beta = 2$).

The effect of sampling actions from colored noise processes of different colors can be exemplified by a velocity controlled robot that can take steps of various sizes in the x and y directions. The paths taken by the robot, when the steps are controlled by noise, depend on the correlation of the noise. The position (x_t, y_t) of the robot is the result of integrating over the noise sequence. Examples of such integrated sequences, *random walks*, for colored noises of different β coefficients are shown in Figure 1. Brownian motion, for example, exhibits red noise ($\beta = 2$) characteristics and is itself the result of integrating white noise ($\beta = 0$) over time.

Colored noise can be generated in the time domain by drawing white noise samples and applying an autoregressive filter or directly in the frequency domain by sampling the frequency components accordingly. In our work we use the latter approach, by building on an implementation by Patzelt (2019) and adapted by Eberhard et al. (2023) following an algorithm by Timmer and König (1995). We generate colored noise sequences of length τ : $\tau_\varepsilon = \{\varepsilon_1, \dots, \varepsilon_\tau\}$. Sampling action noise is then implemented as consuming the items from this sequence. Sec. A3 for further details.

2.2 PPO

PPO is a popular on-policy reinforcement learning algorithm that utilizes a policy gradient approach to optimize the policy. To avoid learning instabilities, PPO aims to prevent large policy updates. As an on-policy algorithm, PPO collects trajectories by interacting with the environment using the current policy. Advantages are then computed for these data and the policy is updated. For the next policy improvement, new data is collected. The number of data points collected for each update is a hyperparameter.

The action selection process uses a policy network that outputs the mean (μ_t) and standard deviation (σ_t) of a Gaussian distribution. An action is then sampled from this distribution:

$$a_t \sim \mathcal{N}(\mu_t, \sigma_t). \quad (1)$$

Using the re-parameterization trick, this is re-written:

$$a_t = \mu_t + \varepsilon_t \cdot \sigma_t; \quad \varepsilon_t \sim \mathcal{N}(0, 1). \quad (2)$$

Instead of sampling ε_t from a white noise Gaussian process, the default in PPO, we propose to use a colored noise Gaussian process, similar to the approach of Eberhard et al. (2023) for *off-policy* methods.

The generated noise samples $\varepsilon_t^{(i)}$ still show Gaussian distributions at each time step (Figure A1.1b). Because we modify the ε in $\mu + \sigma \cdot \varepsilon$ and ε_t remains Gaussian, the data collection, viewed at each individual step, remains asymptotically on-policy. However, while the marginal stays Gaussian, the two-step correlation of ε_t vs. ε_{t+1} changes with β (Figure A1.1a). In this paper, we empirically evaluate the impact of this correlation due to the use of colored noise.




Figure 2: Benchmark environments: (top) Mountain Car, Cartpole Balance, Cartpole Swingup, Ball in Cup (Catch), (middle) Hopper Hop, Cheetah Run, Walker Run, Reacher Hard, Pendulum Swingup, (bottom) Door, UMaze Point, UMaze Ant, 4 Rooms Point, UMaze Swimmer, 4 Rooms Swimmer, 4 Rooms Ant

3 Experiments

We perform an empirical evaluation of the impact of colored noise on the performance of PPO. To this end we perform training runs, repeated with 20 independent seeds, on the benchmarks 4 Rooms Ant, 4 Rooms Point, 4 Rooms Swimmer, Ball in Cup (Catch), Cartpole Balance, Cartpole Swingup, Cheetah Run, Door, Hopper Hop, Mountain Car, Pendulum Swingup, Reacher Hard, UMaze Ant, UMaze Point, UMaze Swimmer and Walker Run. We vary the noise color $\beta \in \{-1$ (blue), 0 (white), 0.2, 0.5, 0.75, 1 (pink), 1.25, 2 (red) $\}$ and test different numbers of parallel collection environments $N_{\text{envs}} \in \{1, 2, 4, 8, 16, 32, 64, 128\}$. This results in a total of 20480 experiments. See Sec. A5 for further details.

3.1 Evaluation Details

We train each agent for a total of 2 048 000 time steps and evaluate every 10240 steps, resulting in 200 evaluation points. At each evaluation point, 50 evaluation episodes are performed. We average the mean returns collected at each evaluation point, forming an estimate that captures the area under the learning curve, and refer to this as the *performance*. Evaluation results are grouped by environment and standardized to zero mean and unit variance to control for the impact of the environments. We combine these standardized results by averaging across different environments and seeds.

Does Colored-Noise Affect the Performance of PPO?

Adding action-noise to the action selection process creates a divergence between the action distribution of the policy and the actual action distribution, i.e., it induces a difference between the distribution of the data the undisturbed policy would collect, compared to the policy disturbed by additive action noise. Similarly, modifying the action selection process, by changing ε in (2), can potentially break the learning process, as on-policy algorithms operate under the assumption that the collected data match the distribution induced by the policy.

In this experiment, we compare how the learning performance (area under the learning curve, averaged across environments, Sec. 3.1) of PPO reacts to a modified action selection process: to the use of correlated noise ($\beta \neq 0$) as




Figure 3: Performance averaged across Environments: Correlated noise $\beta = 0.5$ significantly outperforms the default white noise ($\beta = 0$, Sec. A6) used by PPO. The bars indicate the 95% bootstrapped confidence intervals.

compared to the standard case of using uncorrelated noise ($\beta = 0$). Figure 3 shows the performance averaged across the different environments and across the different numbers of parallel data collection environments. Results for each environment are standardized to control for differences in reward scale. The mean performance, and 95% confidence intervals of the mean, are depicted. The confidence intervals are estimated using (bias-corrected and accelerated) bootstrapping. The results indicate a significant effect of the correlation of the noise. This is in line with previous findings highlighting an impact of action-noise correlation on learned performance (Hollenstein et al. 2022; Eberhard et al. 2023). The results depicted in Figure 3 show an overall preference for colored noise at $\beta = 0.5$, which lies between white noise ($\beta = 0$) and pink noise ($\beta = 1$). This contrasts with the results in off-policy learning, where $\beta = 1$ was found to be superior across many environments (Eberhard et al. 2023). On-policy methods assume matching state-visitation frequencies between the collected data and what the current policy would induce. More correlated noise induces larger state space coverage (Hollenstein et al. 2022) and thus a larger deviation from the states the deterministic-mean-action policy would visit. Thus, a potential reason for this difference in noise color preference is that increasing the state-space coverage, by increasing the noise correlation, increases the distributional shift between the data and policy-induced marginal state-visitation frequencies.

In summary, we found colored noise to have a significant impact on learned policy performance averaged across environments, with $\beta = 0.5$ performing best.

Is $\beta = 0.5$ a Better Default for PPO? The default stochastic policy for PPO relies on uncorrelated noise sampled from a Gaussian distribution. The results shown in Figure 3 indicate that across environments $\beta = 0.5$ performs best. Previous research found an environment specific response to action noise correlation for purely noisy policies (Hollenstein et al. 2021), as well as learning performance in off-policy methods (Hollenstein et al. 2022; Eberhard et al. 2023). In accordance with

Environment	β^*	Noise Color	β 0.5	$p_{0.5}$	β 0.0	$p_{0.0}$	β 1.0	$p_{1.0}$
4 Rooms Swimmer	0.00	White	✓	0.19	✓	–	✗	0.01
Cartpole Balance	0.00	White	✗	0.04	✓	–	✗	0.01
Cheetah Run	0.00	White	✓	0.14	✓	–	✗	0.01
Reacher Hard	0.00	White	✗	0.01	✓	–	✗	0.01
UMaze Ant	0.00	White	✓	0.24	✓	–	✗	0.01
UMaze Swimmer	0.00	White	✓	0.16	✓	–	✗	0.01
4 Rooms Ant	0.20		✓	0.21	✓	0.98	✗	0.01
Door	0.20		✓	0.06	✓	0.56	✓	0.25
Cartpole Swingup	0.50		✓	–	✗	0.03	✗	0.01
4 Rooms Point	0.75		✓	0.97	✗	0.03	✓	0.51
Ball in Cup (Catch)	0.75		✓	0.90	✗	0.01	✓	0.96
Hopper Hop	0.75		✓	0.68	✗	0.01	✓	0.09
Pendulum Swingup	0.75		✓	0.92	✗	0.04	✗	0.03
Walker Run	1.00	Pink	✗	0.01	✗	0.01	✓	–
Mountain Car	1.25		✗	0.01	✗	0.01	✓	0.45
UMaze Point	2.00	Red	✗	0.01	✗	0.01	✓	0.06

Table 1: Optimal noise color per environment. ✓ vs. ✗ indicate whether each fixed $\beta \in \{0.5, 0.0, 1.0\}$ performs comparable to β^* (✓) or is significantly (Welch t-test) outperformed by β^* (✗). The p-values are listed in $p_{0.5}$, $p_{0.0}$ and $p_{1.0}$. $\beta = 0.5$ performs favorable (✓11/16) and improves over the current default of white noise $\beta = 0$ (✓8/16)

these earlier results, we found the best β^* to vary depending on the environment. Table 1 lists the best performing β^* for each environment. In addition, we perform Welch t-tests to compare the performance of β^* to each $\beta \in \{0.5, 0.0, 1.0\}$. This shows that $\beta = 0.5$ performs comparable to β^* in 11/16 environments. In contrast, the current default $\beta = 0$ is significantly outperformed by β^* in eight environments and only performs comparable in 8/16 environments. The best default found for off-policy learning, pink noise ($\beta = 1.0$), is significantly outperformed in nine out of sixteen cases and performs comparable in 7/16 environments. *We therefore recommend switching the default from temporally uncorrelated noise $\beta = 0$ to temporally correlated noise $\beta = 0.5$.*

How Does the Number of Parallel Collection Environments Affect the Performance?

On-policy methods are less efficient with respect to environment interaction samples compared to off-policy methods. This makes on-policy methods particularly interesting when environment samples are cheap to collect and can be collected in parallel. We collected environment samples with different numbers of parallel environments $N_{\text{envs}} \in \{1, 2, 4, 8, 16, 32, 64, 128\}$ and kept the number of total environment interactions (total time steps) constant across experiments. With each of the parallel environments, 2048 samples are collected as the dataset for each update cycle. Thus, depending on the number of environments, more (or less) samples are used for each update. Because the total number of samples is limited (Sec. 3.1) this implies fewer (or more) updates in total. Figure 4 indicates that a larger number of parallel environments negatively affects performance. This is in line with findings by Andrychowicz et al. (2021), who also found that the most beneficial number of parallel environments is dependent on

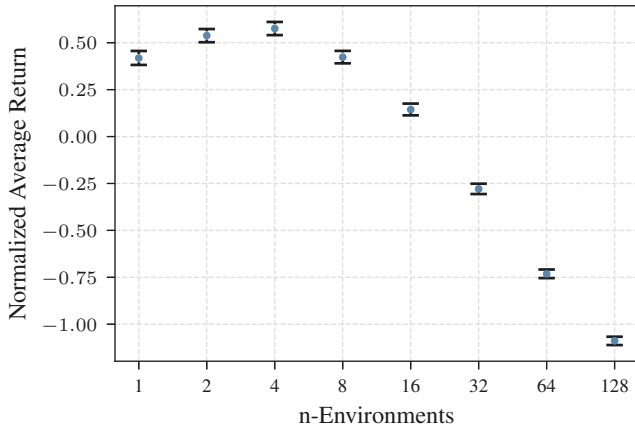
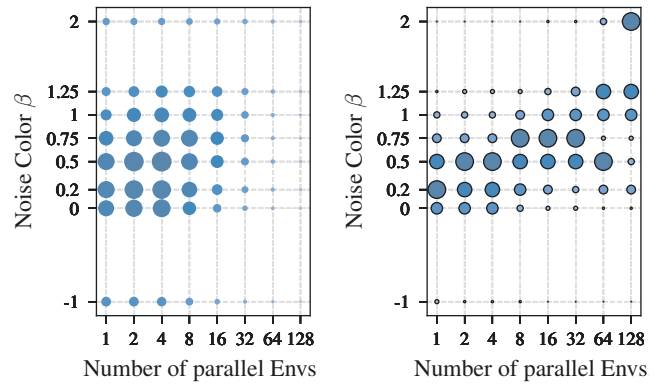


Figure 4: Performance averaged across environments and noise colors: the number of parallel data collection environments has a significant impact on the performance. Bootstrapped 95% confidence intervals for the mean are shown. With $N_{\text{env}} = 4$ achieving the highest performance, though not significantly outperforming $N_{\text{env}} = 2$

the environment type. In our study, we find that about four parallel environments are preferable. This translates to a preference for about $2048 \cdot 4 = 8192$ samples in each update if the episode lengths are small enough to fit the 2048 step limit. In summary, we found 8192 samples per update, collected by four environments in parallel, to perform most efficiently for the investigated class of tasks.


How Does the Number of Parallel Collection Environments and Noise Color Interact? The previous section indicated that the learning performance is impacted by the number of parallel collection environments and hence the size of the dataset used in each update cycle. We found a significant difference between the noise preference in off-policy methods ($\beta = 1.0$) and PPO as an on-policy method ($\beta = 0.5$). Since off-policy methods showed a preference for a larger β ($\beta = 1$) and off-policy methods employ a replay buffer, and thus have access to more samples in each update, this begs the question: does the number of samples used in each PPO update show any interaction with the preferred noise type β ?

Figure 5a shows the average performance across environments depending on the number of collection environments (x) and noise color (y). A positive association between larger β (more correlated noise) and larger number of environments N_{env} becomes visible. This trend is more discernible when the performances for each noise color are compared separately for each number of parallel collection environments. Figure 5b shows the rank of the average performance, separately ranked for each color: a larger number of environments leads to favorable performance of more correlated noise, indicating that with larger samples (i.e., larger N_{env}), more exploration is beneficial. However, in combination with the finding from Sec. 3.1, we find the best performing configuration as $\beta = 0.5, N_{\text{env}} = 4$.



(a) Marker size represents averaged performance (scaled to the sixth power to highlight differences) (b) Marker size represents rank of averaged performance within the same number of environments

Figure 5: Preferred noise color depends on number of environments: Average performance across environment, impact of noise color β combined with n-envs number of parallel environments: (a) A trend is visible: the averaged performance is larger for larger β when more collection environments are used, but the decline due to the number of environments outweighs this trend. (b) Ranks of average performance are indicated by circle size, ranks are calculated across noise-colors but within the same number of environments. The positive trend between number of environments and larger β is clearly visible.



(a) Distribution of the bias in ε_t , calculated by the mean in $\{\varepsilon_0, \dots, \varepsilon_T\}$, across sequences of perturbations $\{\varepsilon_0, \dots, \varepsilon_T\}$, as a function of the noise color β . (b) Standard deviation of the bias the standard deviations over the biases in (a) as a function of noise color β .

Figure 6: Increasing the noise color β , increases the correlation in the perturbations and the spread of the bias of each noise sequence.

Why Do Larger β Work Better for Larger N_{env} ? In the previous section we showed that, grouping for each N_{env} and comparing the different β within each group, there is a tendency for better performance in larger β values when N_{env}

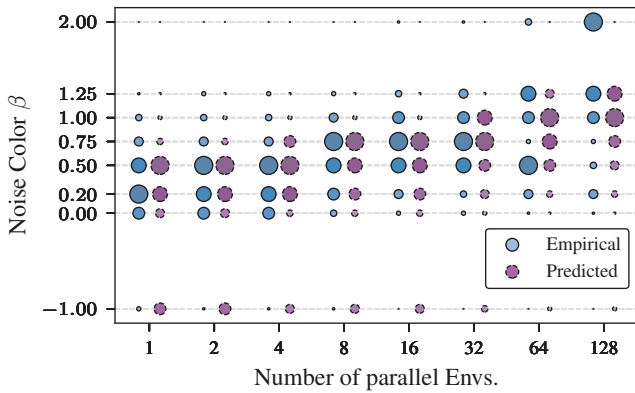


Figure 7: Performance rank, viewed separately for each number of parallel environments. Results from our empirical analysis are compared to a projected best noise color. Larger numbers of parallel environments decrease the variance of the bias, allowing for larger β .

was increased.

Larger β values lead to more correlated perturbation sequences $\{\varepsilon_1, \dots, \varepsilon_T\}$ and, since many environments feature integrative dynamics, larger state space coverage (Eberhard et al. 2023; Hollenstein et al. 2022). While this is beneficial for acquiring new information, it also implies that the samples are more spread out and that the density around the mean actions proposed by the policy is smaller. Similarly, a systematic non-zero bias would prevent the policy from collecting on-policy data and would lead to learning instability.

We measured the bias, as the mean across each sequence of $\{\varepsilon_1, \dots, \varepsilon_T\}$, for different β and found that, while there is no systematic bias ($= 0$) for all β , the spread of the bias increases with β (Figure 6a). The resulting standard deviation of the bias as a function of β is shown in Figure 6b.

When N_{env} is increased, more sequences of $\{\varepsilon_1, \dots, \varepsilon_T\}$ are pooled. This results in a decrease of the variance of the bias of the collected sample $\sigma_N = \frac{\sigma}{\sqrt{N_{\text{env}}}}$. We can estimate a standard deviation $\hat{\sigma}(\beta, N_{\text{env}})$ of the bias for a given β and N_{env} , by combining $\sigma(\beta)$, the standard deviation of the bias as a function of β (Figure 6b), and the reduction due to the sample size $\frac{1}{\sqrt{N_{\text{env}}}}$.

We assume that the variance of the bias was optimal for the best performing β for each N_{env} . We estimate this optimal variance of the bias σ^* as the average of $\hat{\sigma}(n, \beta)$ for the best performing β (Table A2.1). We calculate the difference between all combinations and the optimal combination: $E_{ij} = (\sigma^* - \hat{\sigma}(\beta_i, N_j))^2$ and rank these differences E_{ij} separately for each N_{env} (Figure 7). We find a trend closely matching the performance ranking: larger variance due to larger β is compensated by larger number of environments (and thus samples). The *theoretical* best β for each N_{env} (Figure 7) closely follows our *empirical* observation (Figure 5b).

This indicates, that the increased variance in the bias due to larger β is compensated by larger N_{env} and since generally more exploration (i.e., larger β) would be favorable to collect

more diverse samples, the set of best trade-off configurations of (β, N_{env}) shows a positive correlation between β and N_{env} .

4 Conclusion

Previous work has found temporally correlated noise to be an effective tool for noisy exploration in off-policy reinforcement learning. In this work, we investigated whether temporally correlated noise can also be applied to the on-policy method PPO. We propose to modify PPO by employing the re-parameterization trick, to sample ε from a Gaussian distributed, yet temporally correlated, noise source. We performed a large empirical analysis (in total 20480 experiments, in sixteen environments) and found (i) that colored noise indeed has a significant impact on the learning performance with colored noise ($\beta = 0.5$) intermediate of white noise ($\beta = 0$) and pink noise ($\beta = 1$) to perform best, when averaging across different environments. (ii) We found that, while the best noise color β depends on the environment, $\beta = 0.5$ is favorable in eleven out of the sixteen tested environments, and conclude that $\beta = 0.5$ is a *better default* noise color. (iii) We also investigated the effect of varying the update sample size by varying the number of parallel data collection environments and found $N_{\text{env}} = 4$ (8192 samples per update) to perform best when averaged across environments (and $N_{\text{env}} \in \{2, 4\}$ significantly outperforming the other options). (iv) Interestingly, we found that with larger N_{env} , and thus larger update sample size, the preference moves toward more correlated noise (i.e., larger β). (v) We hypothesize that larger β increase the uncertainty in the collected data: the variance of the effective bias over each action noise sequences increases with larger β and larger sample sizes counteract the variance. We observe that the best performing β as a function of the update size (by varying N_{env}) follows the trend predicted by the change of variance in the bias as a function of β .

In summary, the results from our large empirical evaluation indicate that $\beta = 0.5$ is a better default choice for the noise in PPO. We thus recommend switching the default noise source choice in PPO to colored noise with $\beta = 0.5$.

Acknowledgments

We want to thank Onno Eberhard and Samuele Tosatto for helpful feedback on earlier revisions of this text. Georg Martius is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

References

- Abdolmaleki, A.; Springenberg, J. T.; Tassa, Y.; Munos, R.; Heess, N.; and Riedmiller, M. A. 2018. Maximum a Posteriori Policy Optimisation. In *International Conference on Learning Representations*.
- Amin, S.; Gomrokchi, M.; Satija, H.; van Hoof, H.; and Precup, D. 2021. A Survey of Exploration Methods in Reinforcement Learning. *CoRR*, abs/2109.00157.
- Andrychowicz, M.; Raichuk, A.; Stańczyk, P.; Orsini, M.; Girgin, S.; Marinier, R.; Hussenot, L.; Geist, M.; Pietquin, O.; Michalski, M.; Gelly, S.; and Bachem, O. 2021. What Matters

- for On-Policy Deep Actor-Critic Methods? A Large-Scale Study. In *International Conference on Learning Representations*.
- Chiappa, A. S.; Vargas, A. M.; Huang, A. Z.; and Mathis, A. 2023. Latent Exploration for Reinforcement Learning. <http://arxiv.org/abs/2305.20065>. arxiv:2305.20065.
- Eberhard, O.; Hollenstein, J.; Pinneri, C.; and Martius, G. 2023. Pink Noise Is All You Need: Colored Noise Exploration in Deep Reinforcement Learning. In *International Conference on Learning Representations*.
- Engstrom, L.; Ilyas, A.; Santurkar, S.; Tsipras, D.; Janoos, F.; Rudolph, L.; and Madry, A. 2020. Implementation Matters in Deep RL: A Case Study on PPO and TRPO. In *International Conference on Learning Representations*.
- Fujimoto, S.; van Hoof, H.; and Meger, D. 2018. Addressing Function Approximation Error in Actor-Critic Methods. In *International Conference on Machine Learning*.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; and Levine, S. 2019. Soft Actor-Critic Algorithms and Applications. *arXiv:1812.05905 [cs, stat]*.
- Hollenstein, J.; Auddy, S.; Saveriano, M.; Renaudo, E.; and Piater, J. 2022. Action Noise in Off-Policy Deep Reinforcement Learning: Impact on Exploration and Performance. *Transactions on Machine Learning Research*.
- Hollenstein, J.; Saveriano, M.; Sayantan, A.; Renaudo, E.; and Piater, J. 2021. How Does the Type of Exploration-Noise Affect Returns and Exploration on Reinforcement Learning Benchmarks? In *Austrian Robotics Workshop*.
- Ladosz, P.; Weng, L.; Kim, M.; and Oh, H. 2022. Exploration in Deep Reinforcement Learning: A Survey. *Inf. Fusion*, 85: 1–22.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous Control with Deep Reinforcement Learning. In *International Conference on Learning Representations*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-Level Control through Deep Reinforcement Learning. *Nature*, 518(7540): 529–533.
- Patzelt, F. 2019. Colorednoise: Python Package to Generate Gaussian $(1/f)^{\beta}$ Noise. <https://github.com/felixpatzelt/colorednoise>.
- Petrazzini, I. G. B.; and Antonelo, E. A. 2021. Proximal Policy Optimization with Continuous Bounded Action Space via the Beta Distribution. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–8.
- Plappert, M.; Houthoofd, R.; Dhariwal, P.; Sidor, S.; Chen, R. Y.; Chen, X.; Asfour, T.; Abbeel, P.; and Andrychowicz, M. 2018. Parameter Space Noise for Exploration. In *International Conference on Learning Representations*.
- Raffin, A.; Kober, J.; and Stulp, F. 2021. Smooth Exploration for Robotic Reinforcement Learning. In Faust, A.; Hsu, D.; and Neumann, G., eds., *Conference on Robot Learning*.
- Schulman, J.; Levine, S.; Moritz, P.; Jordan, M.; and Abbeel, P. 2015. Trust Region Policy Optimization. In *International Conference on Machine Learning*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *CoRR*, abs/1707.06347.
- Tassa, Y.; Doron, Y.; Muldal, A.; Erez, T.; Li, Y.; Casas, D. d. L.; Budden, D.; Abdolmaleki, A.; Merel, J.; Lefrancq, A.; Lillicrap, T.; and Riedmiller, M. 2018. DeepMind Control Suite. *arXiv:1801.00690 [cs]*.
- Timmer, J.; and König, M. 1995. On Generating Power Law Noise. *Astron. Astrophys.*, 300: 707–710.
- Uhlenbeck, G. E.; and Ornstein, L. S. 1930. On the Theory of the Brownian Motion. *Physical review*, 36(5): 823.
- Williams, R. J. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine learning*, 8(3-4): 229–256.
- Yang, T.; Tang, H.; Bai, C.; Liu, J.; Hao, J.; Meng, Z.; Liu, P.; and Wang, Z. 2022. Exploration in Deep Reinforcement Learning: A Comprehensive Survey. *CoRR*, abs/2109.06668.