

Evolutionary Strategies at Scale lead to Catastrophic Forgetting

Immanuel Abdi*, Akshat Gupta*

Micah Mok, Alexander Lu, Nicholas Lee, Gopala Anumanchipalli

UC Berkeley

{immanuelazn, akshat.gupta, micahmok, alexander.lu,
nicholas.lee, gopala}@berkeley.edu

Abstract

One of the biggest missing capabilities in current AI systems is the ability to learn continuously after deployment. Implementing such continually learning systems have several challenges, one of which is the large memory requirement of gradient-based algorithms that are used to train state-of-the-art LLMs. Evolutionary Strategies (ES) have recently re-emerged as a gradient-free alternative to traditional learning algorithms and have shown encouraging performance on specific tasks in LLMs. In this paper, we perform a comprehensive analysis of ES and specifically evaluate its forgetting curves when training for an increasing number of update steps. We first find that ES is able to reach performance numbers close to GRPO for math and reasoning tasks with a comparable compute budget. However, and most importantly for continual learning, **the performance gains in ES is accompanied by significant forgetting of prior abilities, limiting its applicability for training models online.** We also explore the reason behind this behavior and show that the updates made using ES are much less sparse and have orders of magnitude larger ℓ_2 norm compared to corresponding GRPO updates, explaining the contrasting forgetting curves between the two algorithms. With this study, we aim to highlight the issue of forgetting in gradient-free algorithms like ES and hope to inspire future work to mitigate these issues.

1 Introduction

Despite rapid advances in AI with transformer-based LLMs (Vaswani et al., 2017; Brown et al., 2020; DeepSeek-AI et al., 2024), most state-of-the-art systems remain static after training and lack the ability to learn continually during deployment. In many real-world settings, models need to adapt to new tasks, user preferences, or data distributions

to perform optimally. While modern chatbots like ChatGPT do this by taking notes in the form of user memory (OpenAI, 2024) and use in-context learning (Brown et al., 2020) to incorporate this information, we currently lack solutions that can achieve this by modifying the model weights during deployment. One of the reasons that makes this challenging is that current post-training and adaptation methods for LLMs are exclusively gradient-based, including approaches such as SFT (Wei et al., 2022), RLHF (Ouyang et al., 2022), DPO (Rafailov et al., 2024), and GRPO (Shao et al., 2024). While effective, these methods require storing gradients, optimizer states, or intermediate activations, causing substantial memory overhead.

Evolutionary Strategies (ES) (Qiu et al., 2025; Korotyshova et al., 2025) have recently re-emerged as a gradient-free alternative for optimizing LLMs. By estimating updates through population-level perturbations rather than backpropagation, ES avoids explicit gradient storage and can significantly reduce memory requirements during deployment. Qiu et al. (2025) have shown that ES achieves comparable performance to GRPO on the Countdown task (Pan, 2026), presenting ES as a viable candidate for continual learning in LLMs. However, a more comprehensive analysis on task generalization was missing in their work. More importantly from the perspective of continual learning, Qiu et al. (2025) do not evaluate the extent to which ES preserves existing capabilities while learning new tasks.

In this work, we present a comprehensive empirical analysis of ES for fine-tuning LLMs, with a focus on continual learning and forgetting. We compare ES against GRPO on multiple math and reasoning benchmarks and evaluate forgetting curves over many update steps. Our results confirm that ES is able to reach performance levels comparable to GRPO on a large suite of tasks; however, contrary to results reported in Qiu et al. (2025), we find

*Equal contribution.

that GRPO still dominates ES marginally on almost all tasks. Additionally, we show that training LLMs using ES leads to significant model degradation and forgetting of existing abilities when compared to GRPO. To better understand this behavior, we analyze the structure of parameter updates produced by ES and compare them to those obtained using GRPO. We find that ES updates are significantly less sparse and exhibit much larger ℓ_2 norms, leading to more global parameter changes that interfere with previously learned capabilities.

Our findings highlight that although ES presents a tempting memory-efficient and gradient-free alternative to inference-time model adaptation, it is also accompanied by “catastrophic” forgetting (Kirkpatrick et al., 2017; Gupta et al., 2024) of prior abilities of the model. We hope these results can inspire future advancements in gradient-free algorithms with continual learning and catastrophic forgetting at the forefront of thought. We also release our codebase¹ and models² for reference.

To summarize, our work makes the following contributions:

1. We show that ES is able to reach comparable performance to GRPO on several math and reasoning benchmarks with similar number of update steps.
2. We show that training models using ES causes significant model degradation when compared to GRPO, leading to catastrophic forgetting of prior abilities.
3. Finally, we also explore the reason behind catastrophic forgetting in ES and show that this happens because model updates using ES are much less sparse when compared to GRPO with significantly larger ℓ_2 norms.

2 Related Work

Evolution Strategies are a class of algorithms that search for solutions to first-order optimization problems by randomly modifying population members to find better performing members (Rechenberg, 1989; Schwefel, 1977; Beyer, 1995). Although implementations such as CMA-ES (Hansen and Ostermeier, 2001) and natural ES (Wierstra et al., 2011;

¹Our codebase can be found here - <https://github.com/akshat57/es-catastrophic>

²Our models can be found here - <https://huggingface.co/collections/immanuelabdi/es-at-scale-lead-to-catastrophic-forgetting>

Sun et al., 2012) demonstrated success, initial implementations remained in the million-parameter scale (Such et al., 2018; Risi and Stanley, 2019; Zhang et al., 2017). However, recent updates have brought ES up to scale and in competition with GPRO, leveraging how it is highly parallelizable, (Salimans et al., 2017), memory efficient (Malladi et al., 2024; Korotyshova et al., 2025), faster (Sarkar et al., 2025), robust to sparse reward horizons (Salimans et al., 2017), and can be modified with LoRA adaptations (Jin et al., 2024; Korotyshova et al., 2025; Sarkar et al., 2025).

PathNet(Fernando et al., 2017) is closely related in spirit because it uses evolution in a continual-learning setting, but it is a hybrid method in which evolution selects/fixes pathways through a modular super-network while gradient descent/RL updates the weights on the active path. This differs from our setting, where we study modern ES for LLM fine-tuning in which the evolutionary procedure directly estimates parameter updates, and we directly analyze catastrophic forgetting through prior task retention during continued optimization.

Qiu et al. (2025) recently published a novel implementation of ES and showed that it outperforms GRPO. However, their study lacked a thorough analysis of model degradation during continued training. Additionally, a bulk of their study was focused on a single dataset. We extend their analysis to multiple datasets, evaluate model degradation during fine-tuning and also study the difference in weight updates in ES when compared to GRPO.

3 Experiments

3.1 ES vs GRPO Comparison

We use the ES implementation of Qiu et al. (2025) and compare it with the GRPO (Shao et al., 2024) implementation from the VERL library (Sheng et al., 2025). An algorithmic analogy between the ES and GRPO algorithms can be found in A.1 while implementations details can be found in A.2. We extend the analysis of ES and GRPO to three math and reasoning tasks – GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021) and Olympiad-Bench (He et al., 2024), in addition to the Countdown dataset which was extensively studied in prior work (Qiu et al., 2025). We perform this study for two models: Qwen2.5-1.5B-Instruct (Qwen et al., 2024) and Llama-3.2-1B-Instruct (Grattafiori et al., 2024). Following the experimental conditions of Qiu et al. (2025), we train our models on 200 ex-

Model	Task	ES	GRPO
Qwen-2.5-1.5B (Instruct)	Countdown	53.0	56.4
	GSM8K	77.4	80.4
	MATH	59.1	63.2
	OlympiadBench	15.2	18.2
Llama-3.2-1B (Instruct)	Countdown	15.2	37.6
	GSM8K	55.2	53.8
	MATH	32.2	35.6
	OlympiadBench	5.6	5.9

Table 1: Peak validation accuracy (%) across tasks and models using previously found optimal hyperparameters. The same hyperparameters were used across ES runs and across GRPO runs.

amples from each dataset with identical batch size and number of rollouts.

The results for comparison between ES and GRPO for fine-tuning LLMs can be found in Table 1. We see that for both models, ES is within 3-4 percentage points of GRPO in terms of task performance. These results are in contrast to prior work by Qiu et al. (2025), who claim that ES significantly outperforms GRPO on the Countdown task. In our experiments, we see that although ES performance numbers are close to GRPO, GRPO still outperforms ES for all but the GSM8K dataset with Llama-3.2-1B model. Therefore, we find different relative performance trends than those reported in prior work, which may stem from differences in GRPO implementations, hyperparameter choices, or evaluation protocols. We release our codebase and open source our trained models for reference.

The fact that the performance numbers of ES are comparable to a state-of-the-art post-training algorithm like GRPO is very encouraging and establishes ES as a potential gradient-free alternative to training LLMs. We also see that for all tasks except Countdown, ES is able to reach peak performance in similar number of update steps, which is shown in Figure A3. This makes the compute requirements of ES also comparable to GRPO.

3.2 ES and Catastrophic Forgetting

While Section 3.1 shows that ES performs competitively with GRPO on various downstream tasks, a defining factor in the viability of using a fine-tuning algorithm for continual learning is its relationship with catastrophic forgetting. We utilized Qwen2.5-1.5B-Instruct trained on the Countdown dataset with GRPO and ES to evaluate catastrophic forgetting. HellaSwag (Zellers et al., 2019) was used to evaluate LLMs on their prior capabilities.

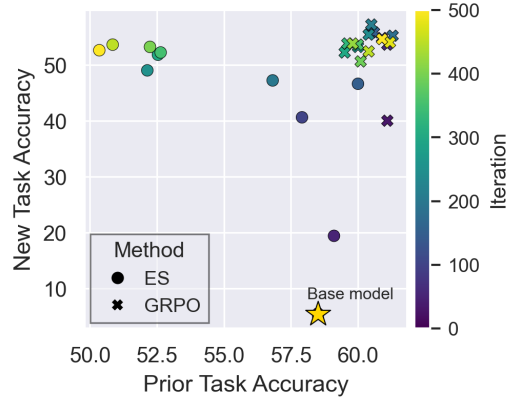


Figure 1: Pareto front between new task (Countdown) and old task (HellaSwag) performance across fine-tuning with ES and GRPO.




Figure 2: Prior task accuracy (%; HellaSwag) vs. training iteration for ES and GRPO fine-tuning. GRPO-trained models exhibit stable prior task accuracy across training iteration, while ES-trained models degrade with continued fine-tuning.

In an ideal scenario, performance on previous tasks should be preserved as new capability is gained. We thus evaluate task performance across each checkpoint of our trained models.

Figure 1 illustrates the relationship between new-task performance (Countdown) and prior-task performance (HellaSwag) across fine-tuning iterations. When training with ES, prior-task performance systematically deteriorates as fine-tuning proceeds, even after new-task performance has effectively converged. This can be observed in the convex Pareto front made by ES in Figure 1. The darker color dots, which depict early training iterations begin with a lower “New Task Accuracy”. With enough training iterations, the increase in “New Task Accuracy” for ES is accompanied by a gradual but evident decline in “Prior Task Accuracy”. Additionally, ES models reach near-maximum Countdown performance by approximately 200 iterations, after




Figure 3: Relationship between Frobenius norm of a model update and number of training iterations on a new task. ES-trained models drift several orders of magnitude more than GRPO-trained models.

which additional training yields negligible gains on the new task. As shown in Figure 2, despite this convergence, previous task performance continues to decline with further iterations, resulting in an approximately 10% drop relative to the best observed prior-task performance. **This pattern indicates that continued ES optimization disproportionately harms previously acquired capabilities, rather than trading off against improvements on the new task.**

In contrast, models fine-tuned with GRPO exhibit markedly different behavior. Across the full range of training iterations, GRPO maintains stable previous task performance while achieving strong new task accuracy. This can be seen by the cluster of crosses on the top-right corner of Figure 1. **This suggests that GRPO avoids the destructive interference observed with ES.** This property of GRPO has also been observed in prior work (Shenfeld et al., 2025).

Therefore, we see that although ES-trained models can be competitive to GRPO, they do so at the cost of severe catastrophic forgetting. Notably, this forgetting occurs within a single fine-tuning run rather than across sequential tasks, highlighting a fundamental instability in ES-based continual adaptation. These results suggest that ES is poorly suited for scenarios requiring task generalization or reuse of previously learned capabilities, whereas GRPO provides a substantially more stable fine-tuning regime.




Figure 4: Layerwise sparsity of parameter updates during fine-tuning (higher indicates more sparse updates). ES exhibits broadly distributed, dense updates across components, whereas GRPO updates are overall highly sparse across LLM parameter groups, consistent with more targeted parameter changes.

3.3 Dissecting ES Updates: Norm and Sparsity

In this section, we seek to determine the characteristics of fine-tuning with ES that cause catastrophic forgetting. We do this by analyzing two features: the update norm and sparsity.

Norm. Here we investigate norm growth of the updated matrix as a function of number of updates with ES and GRPO. We measure the Frobenius norm between model checkpoints within a training run. We do this for the Qwen2.5-1.5B-Instruct model trained on the Countdown dataset.

The results are shown in Figure 3. The Frobenius norm increases monotonically with the number of training iterations for ES-trained models. A similar trend is also present for GRPO-trained models (Figure A2); however, the key distinction lies in scale. After just 500 training iteration, the Frobenius norm of the ES-trained model relative to the base model is *three* orders of magnitude larger than the GRPO-trained model. When combined with what we learn from Figure 2, we see a clear association between the large increases in ES Frobenius norm and a decline in prior task accuracy. **Thus, ES updates have significantly higher ℓ_2 norm difference, causing orders or magnitude larger parameter-shifts compared to GRPO.**

Sparsity. Each update in ES is constructed from high-variance, global perturbations applied across all parameters, which may affect a large number

of stored parameters uniformly. In contrast, it is known that GRPO applies much sparser and targeted updates via backpropagation, limiting the extent of unintended parameter drift (Mukherjee et al., 2025). To check the difference in the number of parameters affected by these algorithms, we evaluate the update sparsity in ES when compared to GRPO.

We analyze Qwen2.5-1.5B-Instruct trained on the Countdown task with both GRPO and ES. We analyze the difference between a base model checkpoint and its corresponding fine-tuned checkpoint. For each shared parameter tensor, we compute the update $\Delta W = W_{new} - W_{base}$. Following prior work (Mukherjee et al., 2025), we define sparsity as the percentage of elements whose absolute magnitude is below a fixed threshold ($\tau = 10^{-4}$). Therefore, higher sparsity values mean a larger number of parameters are below this threshold, which means that the updates are more sparse. Parameters are grouped by architectural component, including attention projections (Q, K, V), the attention output projection (W_O), MLP layers and LayerNorms. Updates are further aggregated by transformer layer index to obtain layerwise sparsity profiles.

The results are shown in Figure 4. **ES updates are substantially less sparse across layers and parameter groups when compared to GRPO.** The sparsity levels for GRPO updates across all parameter types and layers are close to 95%, which means updates are concentrated around a very small number of parameters. However, the updates using ES have very low sparsity levels, showing that a much larger number of parameters are perturbed when fine-tuning with ES. The most sparse updates in ES appear for LayerNorm; however it also contains the least (and negligible) number of parameters compared to other parts of the model. Other layer updates, irrespective of model depth, are highly dense in ES-trained models.

Therefore, GRPO exhibits structured and comparatively sparse updates, aligning with the hypothesis that gradient-based optimization concentrates changes in task-relevant subspaces and mitigates interference with prior capabilities. When combined with KL regularization, these mechanisms provide a natural safeguard against large-scale parameter drift and, consequently, catastrophic forgetting. In contrast, we see that updates using ES have orders of magnitude larger norms and are much less sparse compared to GRPO. The lack of sparsity and large update norms in ES drifts the fine-tuned model fur-

ther away from the base model, potentially leading to the catastrophic forgetting behavior observed in previous sections.

4 Conclusion

We perform an empirical analysis of Evolutionary Strategies for fine-tuning LLMs based on recent work (Qiu et al., 2025) and show that it performs competitively with GRPO. Although, a critical roadblock still persists: **we observe that ES exhibits significant catastrophic forgetting and progressively deteriorates performance on prior skills of the model.** We show that this happens because ES updates have large norms and low sparsity levels (more dense), resulting in parameter drifts that are 1000x higher in magnitude than drifts observed with GRPO for the same number of update steps. These results imply that although recent progress in ES has bridged performance gap with state-of-the-art learning algorithms like GRPO, its intense model degradation still remains a challenge before its widespread adoption.

Limitations

ES has an inherent randomness associated with the updates due to the nature of the algorithm. Consequently, models trained with ES exhibit high variance in stochastic perturbations. Although we observed consistent qualitative trends across the settings we tested, where we worked with a population size of 30 as suggested in prior work (Qiu et al., 2025), increasing population size will decrease the variance and increase statistical stability of our performance numbers.

Our analysis has several limitations that constrain the generalizability of our findings. First, we assess catastrophic forgetting by tracking performance on a single prior task during continued fine-tuning on Countdown. While this provides a useful proxy for retention of a broad capability, it does not capture the full, multi-faceted nature of knowledge degradation across tasks. Second, due to computational constraints, our experiments are conducted on smaller models than those used in comparable studies. It is plausible that scaling effects alter the dynamics of forgetting, and thus our results may not fully extrapolate to larger models. Despite these limitations, the observed trends provide strong evidence for the presence of catastrophic forgetting under our training setup.

References

- H.-G. Beyer. 1995. [Toward a theory of evolution strategies: the \$\(\mu, \lambda\)\$ -theory](#). *Evolutionary Computation*, 2(4):381–407.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). *arXiv preprint arXiv:2005.14165*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- DeepSeek-AI, Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiu Shi Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong Guo, Guangbo Hao, and 57 others. 2024. [Deepseek llm: Scaling open-source language models with longtermism](#). *arXiv preprint arXiv:2401.02954*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#). *Preprint*, arXiv:2501.12948.
- Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. 2017. [Pathnet: Evolution channels gradient descent in super neural networks](#). *Preprint*, arXiv:1701.08734.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Akshat Gupta, Anurag Rao, and Gopala Anumanchipalli. 2024. [Model editing at scale leads to gradual and catastrophic forgetting](#). *arXiv preprint arXiv:2401.07453*.
- Nikolaus Hansen and Andreas Ostermeier. 2001. [Completely derandomized self-adaptation in evolution strategies](#). *Evolutionary Computation*, 9(2):159–195.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024. [Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems](#). *Preprint*, arXiv:2402.14008.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.
- Feihu Jin, Yin Liu, and Ying Tan. 2024. [Derivative-free optimization for low-rank adaptation in large language models](#). *Preprint*, arXiv:2403.01754.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, and 1 others. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Daria Korotyshova, Boris Shaposhnikov, Alexey Malakhov, Alexey Khokhulin, Nikita Surnachev, Kirill Ovcharenko, George Bredis, Alexey Gorbatsovski, Viacheslav Sinii, and Daniil Gavrilov. 2025. [Essa: Evolutionary strategies for scalable alignment](#). *Preprint*, arXiv:2507.04453.
- Sadhika Malladi, Tianyu Gao, Eshaan Nichani, Alex Damian, Jason D. Lee, Danqi Chen, and Sanjeev Arora. 2024. [Fine-tuning language models with just forward passes](#). *Preprint*, arXiv:2305.17333.
- Sagnik Mukherjee, Lifan Yuan, Dilek Hakkani-Tur, and Hao Peng. 2025. [Reinforcement learning finetunes small subnetworks in large language models](#). *arXiv preprint arXiv:2505.11711*.
- OpenAI. 2024. [Memory and new controls for ChatGPT](#). Accessed: 2026-01-24.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Preprint*, arXiv:2203.02155.
- Jiayi Pan. 2026. [Jiayi-Pan/TinyZero](#).
- Xin Qiu, Yulu Gan, Conor F. Hayes, Qiyao Liang, Elliot Meyerson, Babak Hodjat, and Risto Miikkulainen. 2025. [Evolution Strategies at Scale: LLM Fine-Tuning Beyond Reinforcement Learning](#). *Preprint*, arXiv:2509.24372.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin,

- Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2024. [Qwen2.5 Technical Report](#).
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2024. [Direct preference optimization: Your language model is secretly a reward model](#). *Preprint*, arXiv:2305.18290.
- Ingo Rechenberg. 1989. Evolution strategy: Nature’s way of optimization. In *Optimization: Methods and Applications, Possibilities and Limitations*, pages 106–126, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Sebastian Risi and Kenneth O. Stanley. 2019. [Deep neuroevolution of recurrent and discrete world models](#). *Preprint*, arXiv:1906.08857.
- Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. [Evolution Strategies as a Scalable Alternative to Reinforcement Learning](#). *Preprint*, arXiv:1703.03864.
- Bidipta Sarkar, Mattie Fellows, Juan Agustin Duque, Alistair Letcher, Antonio León Villares, Anya Sims, Dylan Cope, Jarek Liesen, Lukas Seier, Theo Wolf, Uljad Berdica, Alexander David Goldie, Aaron Courville, Karin Sevegnani, Shimon Whiteson, and Jakob Nicolaus Foerster. 2025. [Evolution strategies at the hyperscale](#). *Preprint*, arXiv:2511.16652.
- Hans-Paul Schwefel. 1977. *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie: mit einer vergleichenden Einführung in die Hill-Climbing- und Zufallsstrategie*. Birkhäuser Verlag, Basel, Stuttgart.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Idan Shenfeld, Jyothish Pari, and Pulkit Agrawal. 2025. [RL’s razor: Why online reinforcement learning forgets less](#). *Preprint*, arXiv:2509.04259.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. [Hybridflow: A flexible and efficient rlhf framework](#).
- Felipe Petroski Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth O. Stanley, and Jeff Clune. 2018. [Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning](#). *Preprint*, arXiv:1712.06567.
- Yi Sun, Daan Wierstra, Tom Schaul, and Juergen Schmidhuber. 2012. [Efficient natural evolution strategies](#). *Preprint*, arXiv:1209.5853.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *arXiv preprint arXiv:1706.03762*.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#). *Preprint*, arXiv:2109.01652.
- Daan Wierstra, Tom Schaul, Tobias Glasmachers, Yi Sun, and Jürgen Schmidhuber. 2011. [Natural evolution strategies](#). *Preprint*, arXiv:1106.4487.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) *Preprint*, arXiv:1905.07830.
- Xingwen Zhang, Jeff Clune, and Kenneth O. Stanley. 2017. [On the relationship between the openai evolution strategy and stochastic gradient descent](#). *Preprint*, arXiv:1712.06564.

A Appendix

A.1 Algorithmic Overview and Analogy Between ES and GRPO

A.1.1 ES Algorithm Overview

Qiu et al. (2025) implement a version of evolutionary strategies that features these techniques: weight adjustment in-place with noise generation from stored random seeds, ranked weight updates, and learning rate ingestion.

Each update step can be understood through the following equations.

Each population member at step t has a unique seed. With noise per iteration $\epsilon_{n,l} \sim \mathcal{N}(0, I)$, model parameters for timestep t θ_t , layer parameters for step t $\theta_{t,l}$, reward function R , reward score for the n th member R_n , z-score for n th member Z_n , and noise coefficient σ , and learning rate α .

Reset random seed generator. Sample noise $\epsilon_{n,l} \sim \mathcal{N}(0, I)$. For all layers, perturb in-place:

$$\theta_{t-1,l} \leftarrow \theta_{t-1,l} + \sigma \cdot \epsilon_{n,l}.$$

Reward for perturbed model is calculated:

$$R_n = R(\theta_{t-1}).$$

Reset random seed generator. Sample noise $\epsilon_{n,l} \sim \mathcal{N}(0, I)$. For all layers, restore in-place:

$$\theta_{t-1,l} \leftarrow \theta_{t-1,l} - \sigma \cdot \epsilon_{n,l}.$$

Z-score is calculated per population member:

$$Z_n = \frac{R_n - R_{\text{mean}}}{R_{\text{std}}},$$

Reset random seed generator. Sample noise $\epsilon_{n,l} \sim \mathcal{N}(0, I)$. For all layers, update with noise weighted by z-score and learning rate in-place:

$$\theta_{t,l} \leftarrow \theta_{t-1,l} + \alpha \cdot \frac{1}{N} Z_n \epsilon_{n,l}.$$

where R_{mean} and R_{std} are the mean and standard deviation of R_1, R_2, \dots, R_N .

A.1.2 GRPO Algorithm Overview

Shao et al. (2024) implement Group Relative Policy Optimization (GRPO), which eliminates the critic model by estimating advantages from group statistics.

For each prompt q , sample a group of G outputs $\{o_1, o_2, \dots, o_G\}$ from the current policy $\pi_{\theta_{\text{old}}}$.

Compute rewards $\{r_1, r_2, \dots, r_G\}$ for each output and calculate relative advantages via z-score normalization:

$$A_i = \frac{r_i - \text{mean}(\{r_1, \dots, r_G\})}{\text{std}(\{r_1, \dots, r_G\})}.$$

The policy π_{θ} is updated by maximizing the GRPO objective:

$$\rho_i(\theta) = \frac{\pi_{\theta}(o_i | q)}{\pi_{\theta_{\text{old}}}(o_i | q)}$$

$$\begin{aligned} \mathcal{J}_{\text{GRPO}}(\theta) = & \frac{1}{G} \sum_{i=1}^G \min \left(\rho_i(\theta) A_i, \right. \\ & \left. \text{clip}(\rho_i(\theta), 1 - \epsilon, 1 + \epsilon) A_i \right) \\ & - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}). \end{aligned}$$

To penalize divergence from the reference policy π_{ref} without additional sampling, the KL term is approximated:

$$\mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) = \frac{\pi_{\text{ref}}(o_i | q)}{\pi_{\theta}(o_i | q)} - \log \frac{\pi_{\text{ref}}(o_i | q)}{\pi_{\theta}(o_i | q)} - 1.$$

By replacing the value function with group-relative rewards, this implementation reduces computational overhead and memory usage compared to standard PPO.

A.1.3 Analogy Between ES Population Size and GRPO Rollout Count

Both GRPO and ES rely on creating different responses and then updating the model parameters via the fitness of those responses. The following section describes why the population size in ES

and number of rollouts in GRPO play an analogous role in controlling parameter updates.

Following the algorithm described by Qiu et al. (2025), an ES training update comprises of N different seeds used to generate perturbations to the baseline model, resulting in N different population members. Each population member is sampled with temperature at 0 to generate N different responses, which are evaluated by a reward function to determine their fitness, which is converted into a z-score to weight the contribution of that respective perturbation to the baseline model. Further explanation can be found in A.1.

Similarly, a GRPO training update samples N candidate outputs from the current policy, evaluates them to obtain relative reward signals, and updates the policy via a policy-gradient objective while constraining deviation from a fixed reference policy through KL regularization (DeepSeek-AI et al., 2025). Crucially, although ES simultaneously maintains multiple different versions of a model and GRPO maintains one, ES population size and GRPO number of rollouts both determine the number of samples used to estimate a stochastic update and to form a stochastic gradient or gradient-free estimator that drives the parameter update.

A.2 Implementation Details

Our implementations for both GRPO and ES model training and analysis is attached to this submission.

A.2.1 GRPO

The GRPO setup in this study is implemented on the VERL library, which employs the HybridFlow engine proposed by Sheng et al. (2025). Training was conducted on NVIDIA RTX A6000 GPUs and the Fully Sharded Data Parallel (FSDP) protocol was used to train across GPUs. Across all experiments, we maintained 30 rollouts for GRPO to mimic the 30 mutations generated by the original ES study by Qiu et al. (2025). To benchmark-finetuning, we used a batch-size of 200 examples along with a mini-batch size of 32 examples. A KL-Loss coefficient of $\beta = 0.001$ was used. The trainer was set to run for a total of 500 epochs, although once the validation accuracy appeared to plateau, we stopped training prematurely.

A.2.2 ES

We replicated the original author’s implementation of ES with two improvements: the authors found that using fp16 instead of bf16 improved validation

accuracy on certain tasks. Additionally, the application of the Qwen chat template to the original task prompts improved validation accuracy on the experiment replica Countdown task for Qwen2.5-1.5B, but left model performance on all other regimes virtually the same. Runs were performed both with and without the chat template to assess the effect.

A.2.3 Reward functions

For the countdown task, we employ the same reward function used by Qiu et al. (2025), adapted to fit the VERL API. An answer reward is calculated, which assigns a reward of 1.0 if the model’s answer uses all numbers once and evaluates to the provided target, and 0.0 otherwise. A separate format score is calculated, which serves to ensure that the model’s response obeys an XML-style format with `<think>...</think>` thinking tokens first followed by response tokens `<answer>...</answer>`. We take a weighted average of the two rewards to calculate the final reward to assign to the model: $\text{Reward} = 0.1 \cdot \text{Format Reward} + 0.9 \cdot \text{Answer Reward}$

For the GSM8K, MATH, and OlympiadBench benchmarks, we employ a rule-based reward function using a binary evaluation logic. An answer reward is calculated by extracting the model’s conclusion from the final 300 characters of the response using a regex pattern. The function first identifies the `#### [number]` format, falling back to `\boxed{...}` tags if necessary, and assigns a reward of 1.0 if the extraction matches the ground truth and 0.0 otherwise.

A.3 Hyperparameter Values

Hyperparameter	Value
Population size	30
Noise scale σ	0.001
Learning rate α	0.0005
Max tokens	1024

Table 2: Hyperparameters used for Evolution Strategies (ES) fine-tuning.

A.4 Additional Experiments

A.4.1 Catastrophic Forgetting and KL

Shenfeld et al. (2025) had previously established a negative correlation between KL-divergence and previous task score. We therefore searched whether

this trend also is reflected within ES-trained models. We first looked at KL-divergence between the trained and base models A1 on the newly trained task. While ES-trained models increase in KL-divergence with subsequent training steps, this behavior was not consistent when trained with GRPO. This can be attributed to the explicit KL-regularization factor in GRPO, preventing continuous drifts from the base model.

The trends for KL divergence and accuracy continue to diverge when evaluating previously known tasks A1. ES has a clear negative correlation between KL-divergence and old task performance. The KL-divergence between the new and base models also shows an increase over number of training iterations in ES. GRPO however continues to show no association between number of training steps and KL-divergence, as well as between KL-divergence and previous task accuracy. Therefore, KL-divergence is a less reliable indicator of catastrophic convergence across GRPO and ES.

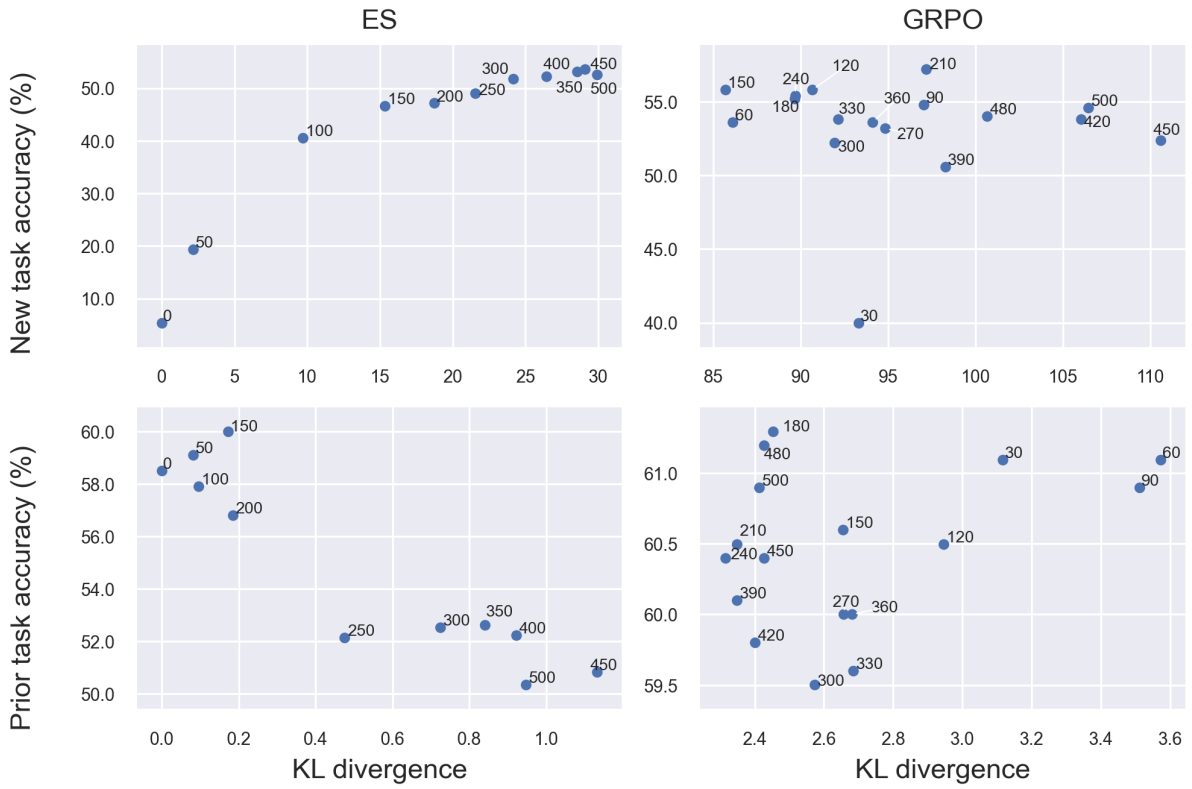


Figure A1: Relationship between KL divergence and task performance. Top row: new-task accuracy (Countdown). Bottom row: prior-task accuracy (HellaSwag). Training step indicated per sample. ES exhibits increasing KL accompanied by degradation on the prior task, whereas GRPO maintains stable performance across a broader KL range.






Figure A2: Log relationship between Frobenius norm of a model update and number of training iterations on a new task (Countdown). ES-trained models drift several orders of magnitude more than GRPO-trained models.



(a) Qwen-2.5-1.5B



(b) LLaMa-3.2-1B

Figure A3: Mean accuracy curves for ES and GRPO runs across across datasets: Countdown, GSM8K, Math, Olympiad.