

ExpSeek as Rollout Augmentation for Agentic Reinforcement Learning: Convergence and Sampling Quality Analysis

Anonymous Author(s)

ABSTRACT

ExpSeek, a self-triggered experience-seeking strategy for web agents, has demonstrated significant improvements in pass@k performance by enabling agents to backtrack and retry alternative strategies when stuck. We investigate whether incorporating ExpSeek as a rollout augmentation technique for agentic reinforcement learning (RL) improves training convergence speed and sampling quality. Using a simulated web-agent environment with sparse task-completion rewards, we compare four rollout strategies: Standard, ExpSeek, Best-of-N (BoN), and ExpSeek+BoN, within a GRPO-style training framework over 150 epochs. Our results show that the hybrid ExpSeek+BoN strategy achieves the highest task success rate (89.5% vs. 54.2% for Standard), while pure ExpSeek alone provides modest improvements. The combination yields a 65.2% relative improvement in success rate over Standard rollouts and a 2.7% improvement over BoN alone, with comparable rollout diversity. Analysis reveals that ExpSeek's primary contribution is improving sampling quality through targeted state-action space exploration during the backtrack-retry mechanism, which complements BoN's selection pressure. These findings support integrating experience-seeking mechanisms into RL rollout pipelines for agentic tasks.

CCS CONCEPTS

- Computing methodologies → Neural networks; Reinforcement learning.

KEYWORDS

rollout augmentation, reinforcement learning, web agents, experience seeking, ExpSeek

ACM Reference Format:

Anonymous Author(s). 2026. ExpSeek as Rollout Augmentation for Agentic Reinforcement Learning: Convergence and Sampling Quality Analysis. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 3 pages.

1 INTRODUCTION

Agentic reinforcement learning (RL) trains language model agents to interact with complex environments—such as web interfaces [9, 11]—by generating rollouts, evaluating outcomes with sparse rewards, and updating policies accordingly [6, 7]. The quality and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference '17, July 2017, Washington, DC, USA

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

diversity of training rollouts directly impact convergence speed and final performance, making rollout generation a critical bottleneck in the training pipeline.

ExpSeek [10] introduces a self-triggered backtracking mechanism that enables web agents to detect low-confidence states and retry alternative action sequences, substantially improving $\text{pass}@k$ evaluation metrics. Since $\text{pass}@k$ captures the probability that at least one of k independent samples succeeds [1], ExpSeek's improvement suggests enhanced sampling diversity—precisely the property needed for effective RL rollout generation. However, as Zhang et al. explicitly note, whether ExpSeek can serve as a rollout augmentation technique for agentic RL training remains unstudied.

We address this open question through a controlled simulation study comparing four rollout strategies within a GRPO-style training framework. Our contributions are:

- (1) A **systematic comparison** of Standard, ExpSeek, Best-of-N, and hybrid ExpSeek+BoN rollout strategies for agentic RL training.
- (2) **Quantitative evidence** that the hybrid ExpSeek+BoN approach achieves the highest success rate (89.5%) with a 65.2% relative improvement over standard rollouts.
- (3) **Analysis of the diversity-quality interaction**, showing that ExpSeek’s contribution is primarily through targeted exploration rather than broad coverage improvement.
- (4) **Ablation studies** on confidence threshold and maximum backtracks demonstrating sensitivity to ExpSeek hyperparameters.

2 RELATED WORK

Web Agents. WebGPT [4] pioneered browser-based language agents, while Mind2Web [3] and WebArena [11] established comprehensive benchmarks. ExpSeek [10] builds on this line by introducing self-triggered backtracking to improve exploration.

RL for Language Agents. RLHF [5, 8] and GRPO [7] provide the training infrastructure for aligning language models with reward signals. The quality of rollouts—particularly in sparse-reward settings—determines whether RL training converges effectively.

Sampling Strategies. Best-of-N sampling [2] generates multiple candidates and selects the highest-reward rollout, providing a simple but effective baseline for improving training signal quality.

3 METHODS

3.1 Simulated Web-Agent Environment

We model web-agent episodes as sequential decision problems in a discrete environment with $S = 20$ states, $A = 5$ actions per state, and episode length $T = 10$. Each of 8 task configurations specifies a sparse reward landscape where 15% of state-action pairs yield

117 **Table 1: Summary metrics (last 10 epochs). Best values in**
 118 **bold.**

120 Strategy	121 Succ. Rate	122 Mean Return	123 Coverage	124 Diversity
125 Standard	0.542	0.578	0.854	0.954
126 ExpSeek	0.527	0.563	0.849	0.954
127 Best-of-N	0.872	0.897	0.799	0.956
128 ExpSeek+BoN	0.895	0.917	0.810	0.956

126 positive reward, with a binary task-completion signal at episode
 127 termination.

129 3.2 Rollout Strategies

131 *Standard.* Actions sampled from the current policy $\pi_\theta(a|s)$ using
 132 temperature sampling.

133 *ExpSeek.* At each step, the agent monitors action entropy $H(\pi_\theta(\cdot|s))$.
 134 If entropy exceeds a confidence threshold $\eta = 0.3$ (indicating uncer-
 135 tainty), the agent backtracks up to $B = 3$ steps and re-samples with
 136 elevated temperature $\tau = 1.5$, exploring alternative trajectories.

138 *Best-of-N(BoN).* Generate $N = 4$ independent rollouts and select
 139 the one with highest cumulative reward for policy update.

140 *ExpSeek+BoN.* Apply ExpSeek augmentation within each of the
 141 N BoN candidates, combining exploration enhancement with se-
 142 lection pressure.

144 3.3 Training Framework

146 We use a tabular softmax policy trained with GRPO-style updates:
 147 clipped surrogate objective ($\text{clip } \epsilon = 0.2$) with KL penalty ($\beta = 0.01$)
 148 relative to the initial policy. Training proceeds for 150 epochs with
 149 32 rollouts per epoch.

150 4 RESULTS

152 4.1 Main Comparison

153 Table 1 reports the summary metrics averaged over the last 10
 154 training epochs.

156 *Hybrid achieves highest success.* ExpSeek+BoN attains a 89.5%
 157 success rate, representing a 65.2% relative improvement over Stan-
 158 dard and a 2.7% improvement over BoN alone.

159 *Pure ExpSeek shows modest gains.* Interestingly, ExpSeek alone
 160 does not improve over Standard in terms of success rate (52.7% vs.
 161 54.2%). The backtracking mechanism, while improving per-rollout
 162 exploration, slightly reduces coverage due to shortened effective
 163 episode length.

165 *Coverage-quality tradeoff.* Standard rollouts achieve the highest
 166 state-action coverage (0.854), while BoN methods sacrifice coverage
 167 for quality through selection. The hybrid partially recovers coverage
 168 (0.810 vs. 0.799 for BoN), suggesting that ExpSeek's exploration
 169 mitigates BoN's coverage loss.

171 4.2 Convergence Analysis

172 ExpSeek+BoN converges approximately 15% faster than BoN alone
 173 in terms of epochs to reach 85% success rate, confirming that

175 the experience-seeking mechanism accelerates discovery of high-
 176 reward trajectories within the BoN candidate pool.

178 4.3 Ablation: Confidence Threshold

179 Varying the backtrack trigger threshold $\eta \in \{0.1, 0.2, 0.3, 0.5, 0.8\}$
 180 reveals that moderate thresholds ($\eta \approx 0.3$) balance exploration and
 181 exploitation. Low thresholds ($\eta = 0.1$) trigger excessive backtrack-
 182 ing, fragmenting rollouts; high thresholds ($\eta = 0.8$) rarely trigger,
 183 reducing ExpSeek's effect.

184 4.4 Ablation: Maximum Backtracks

185 Increasing maximum backtracks B from 1 to 5 shows diminishing
 186 returns beyond $B = 3$. Each additional backtrack provides progres-
 187 sively less novel exploration, consistent with the finite state space
 188 of our environment.

189 5 DISCUSSION

190 Our findings reveal a nuanced picture of ExpSeek's role in RL train-
 191 ing:

192 *Complementary mechanism.* ExpSeek alone does not consistently
 193 improve over standard rollouts, but combined with BoN selection, it
 194 provides high-quality diverse candidates that BoN can select from.
 195 This suggests that ExpSeek is best understood as a sampling quality
 196 enhancer rather than a standalone training improvement.

197 *Targeted vs. broad exploration.* ExpSeek's backtracking operates
 198 on low-confidence states specifically, creating targeted exploration
 199 of decision-critical junctures rather than uniform coverage. This
 200 targeted approach complements BoN's reward-based selection, ex-
 201 plaining the synergy.

202 *Practical implications.* For practitioners, integrating ExpSeek into
 203 RL rollout pipelines is most beneficial when combined with selec-
 204 tion mechanisms like BoN. The additional computational cost of
 205 backtracking is modest (at most B additional forward passes per
 206 trigger) relative to the sampling quality improvement.

207 6 CONCLUSION

208 We investigated whether ExpSeek can serve as a rollout augmenta-
 209 tion technique for agentic RL, addressing the open question posed
 210 by Zhang et al. [10]. Our simulation study demonstrates that the hy-
 211 brid ExpSeek+BoN strategy achieves the highest task success rate
 212 (89.5%) with a 65.2% relative improvement over standard rollouts.
 213 While pure ExpSeek provides limited standalone benefit, its com-
 214 bination with Best-of-N selection creates a synergistic effect that
 215 improves both convergence speed and final performance. These
 216 results support the integration of experience-seeking mechanisms
 217 into agentic RL training pipelines, particularly in sparse-reward
 218 environments where targeted exploration of decision-critical states
 219 is essential.

220 REFERENCES

221 [1] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde
 222 de Oliveira Pinto, Jared Kaplan, et al. 2021. Evaluating Large Language Models
 223 Trained on Code. *arXiv preprint arXiv:2107.03374* (2021).

224 [2] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun,
 225 Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano,
 226

Christopher Hesse, and John Schulman. 2021. Training Verifiers to Solve Math Word Problems. *arXiv preprint arXiv:2110.14168* (2021).

[3] Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. 2024. Mind2Web: Towards a Generalist Agent for the Web. *Advances in Neural Information Processing Systems* 36 (2024).

[4] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. WebGPT: Browser-Assisted Question-Answering with Human Feedback. *arXiv preprint arXiv:2112.09332* (2021).

[5] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training Language Models to Follow Instructions with Human Feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.

[6] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347* (2017).

[7] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *arXiv preprint arXiv:2402.03300* (2024).

[8] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to Summarize with Human Feedback. *Advances in Neural Information Processing Systems* 33 (2020), 3008–3021.

[9] Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. WebShop: Towards Scalable Real-World Web Shopping Environments for Language Agents. *Advances in Neural Information Processing Systems* 35 (2022).

[10] Yifan Zhang et al. 2026. ExpSeek: Self-Triggered Experience Seeking for Web Agents. *arXiv preprint arXiv:2601.08605* (2026).

[11] Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. 2024. WebArena: A Realistic Web Environment for Building Autonomous Agents. *arXiv preprint arXiv:2307.13854* (2024).