BraintankAI ASIA

DISPATCH · AGENT REINFORCEMENT LEARNING · 3 MIN

BEIJING · Peking University

Reinforcement learning for AI agents that hands out credit turn by turn

An IQuest preprint trains search and tool-using agents by scoring each complete turn first, then sharing that score among its tokens, and tops five rival methods.

18 November 2026 · arXiv preprint, not yet peer reviewed

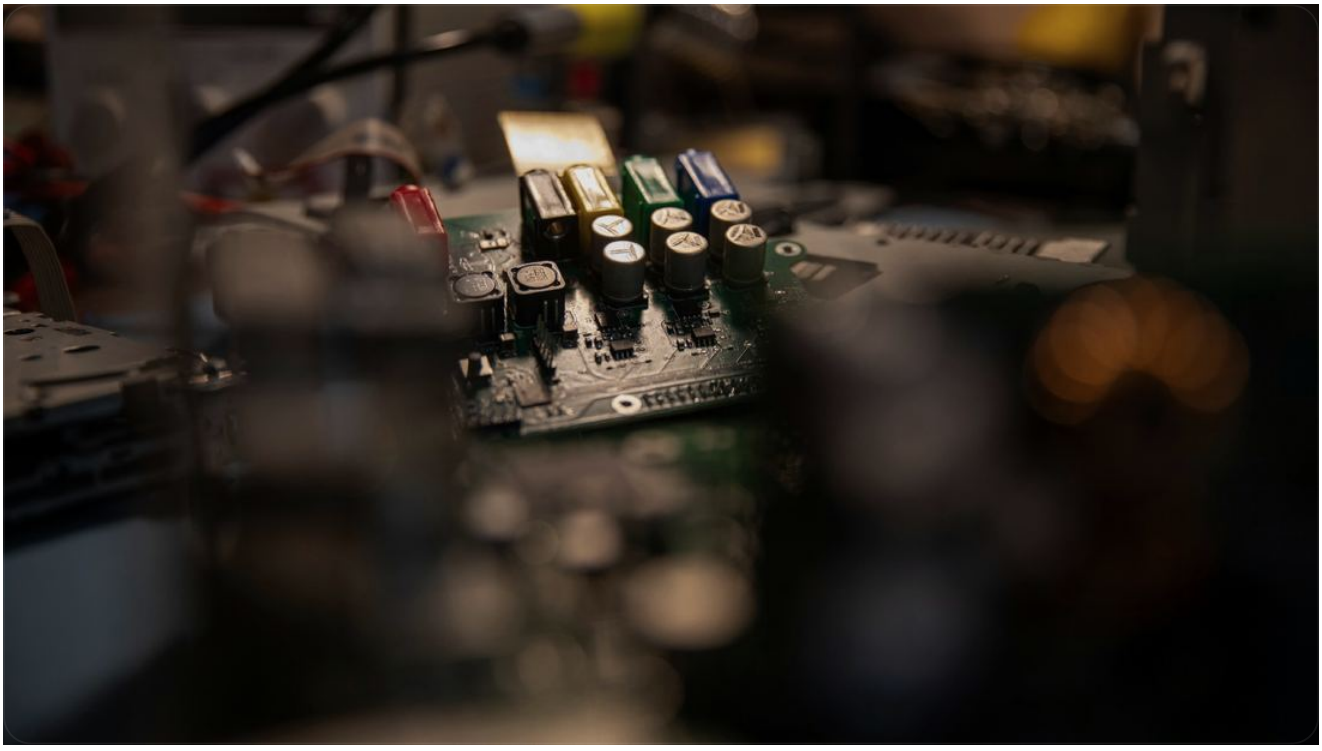


Photo: Laurens Van Der Drift

A language-model agent that searches the web or calls software tools acts one complete turn at a time: a query, a tool call, a final answer. Yet it writes each turn token by token, and it often learns from a single reward that arrives only at the end. Working out which turn, and which token within it, earned that reward is known as credit assignment. A 2026 preprint from IQuest, which has not yet been peer reviewed, describes Multiscale Policy Optimization (MPO), a reinforcement-learning method that assigns credit to whole turns first and only then to the tokens inside them.

The authors place MPO between two families of existing methods. Critic-free methods such as GRPO and GSPO give one shared advantage, a measure of how much better an outcome was than expected, to an entire sequence, so they cannot tell which turn mattered. Critic-based methods such as PPO, VAPO and SAO use a critic, a learned model that predicts the final reward to expect from a partial trajectory. They treat every token as a separate action, which blurs the line between

writing text and interacting with the environment. The motivating observation comes from one search-agent training run. Across 2,048 trajectories and 16,323 turns, differences between turn averages accounted for a median of 89.5% of the variation in the critic's predictions within a trajectory. In 93.8% of trajectories, variation between turns exceeded variation within them.

MPO therefore treats each turn as a single 'macro action', whose probability is the product of its token probabilities. Its estimator, Multiscale Advantage Estimation, applies generalised advantage estimation, a decaying sum of one-step surprise signals, along a timeline of turns rather than tokens. Credit therefore fades once per interaction instead of once per token. A second trace, confined to each turn, then redistributes the turn's credit among its tokens. It is constrained so that the tokens' average credit still equals the turn's, and a setting called β controls how strongly it applies. When the policy is updated, every token is weighted by the importance ratio of the whole turn, which compares the new and old model's probability of that turn. A stop-gradient, an operation that blocks learning signals through a value, keeps each token's update tied to its own credit.

The team trained Qwen3-4B-Thinking and Qwen3-30B-A3B-Instruct separately on search, Claw-Eval and Automation Bench tasks. Search agents were scored on BrowseComp-Plus, GAIA and HLE. The authors report that MPO scored highest on all nine metrics at both model sizes. On the 4B model, BrowseComp-Plus pass@1, the chance that a single attempt succeeds, reached 48.08, against 43.75 for SAO, the strongest rival on that measure. GAIA pass@1 reached 39.47, against 35.62. GAIA pass@3 was 52.67 against 47.56 on the smaller model and 56.24 against 50.34 on the larger. Gains in HLE pass@1 were slight, at 0.66 and 0.12 points, and Automation Bench scores stayed low, at most 15.00.

The internal comparisons point the same way. Swapping turn-level for token-level importance ratios lowered every reported score. Setting β to zero, so that all tokens in a turn share identical credit, cut 4B BrowseComp-Plus pass@1 from 48.08 to 41.80. Even that turn-only version beat PPO on all six search metrics.

The evidence has clear limits. The paper reports a single run for each ablation setting and for each method whose training it tracks, with no error bars or significance tests. The best checkpoint of each run was selected without saying on what data. The ablations cover only the 4B search setting, and the main turn-decay setting was tested at just three values. Only one model family was used. The test of larger sampling budgets compared MPO only with GRPO and PPO. The main theorems assume exact advantages and a fixed number of turns. MPO also needs a learned critic, unlike GRPO and GSPO, and the paper gives no comparison of computing cost.

The work tackles a problem braintankAI has covered from a different angle. One study graded the intermediate steps of a search agent directly, rather than only its final answer, and reported sturdier agents (earlier story). MPO keeps the reward at the end and instead relies on a critic to spread that reward across turns.

PAPER	*Multiscale Policy Optimization for Long-Horizon LLM Agents*
AUTHORS	Tao Ren, Hui Yang, Jia Deng, Tingyang Chen, Qiyao Wang, Shuo Tang et al.
VENUE	arXiv preprint — not yet peer reviewed · 26 October 2026
LINK	arxiv.org/abs/2610.35013
INDEX	stacks.braintank.ai
CODE	github.com/IQuestLab/Multiscale-Policy-Optimization