
Reward Is Enough: LLMs Are In-Context Reinforcement Learners

Kefan Song*
University of Virginia

Amir Moeini*
University of Virginia

Peng Wang
University of Virginia

Lei Gong
University of Virginia

Rohan Chandra
University of Virginia

Yanjun Qi
University of Virginia

Shangtong Zhang*
University of Virginia

Abstract

Reinforcement learning (RL) is a human-designed framework for solving sequential decision making problems. In this work, we demonstrate that, surprisingly, RL emerges in LLM's (Large Language Model) inference time – a phenomenon known as in-context RL (ICRL). Specifically, we propose a novel multi-round prompting framework called ICRL prompting. The goal is to prompt the LLM to complete a task. After the LLM generates a response at the current round, we give numerical scalar feedbacks for the response, called the rewards. At the next round, we prompt the LLM again with the same task and a context consisting of all previous responses and rewards. We observe that the quality of the LLM's response increases as the context grows. In other words, the LLM is able to maximize the scalar reward signal in the inference time, just like an RL algorithm. We evaluate ICRL prompting in three benchmarks (Game of 24, creative writing, and ScienceWorld) and demonstrate significant performance improvements over baseline methods such as Self-Refine and Reflexion. Surprisingly, in some experiments the reward signals are generated by the LLM itself, yet performance improvements are still observed from ICRL prompting, offering a promising paradigm for scaling test-time compute.

1 Introduction

For Large Language Models (LLMs) to act as effective agents to solve novel tasks, the capacity for LLMs to improve during the inference time is essential. Learning and search are the two general methods that can leverage scaling computation for performance improvement (Sutton, 2019). Search has been successfully applied to LLM's self-improvement during the inference time, starting from the simple Best-of-N (Stiennon et al., 2022) to Tree of Thoughts (Yao et al., 2023a) and Monte Carlo Tree Search (Ding et al., 2024). Learning, however, has yet to receive the same attention for LLM's self-improvement during the inference time. In-context (supervised) learning (ICL, Brown et al. (2020)), as a supervised learning paradigm during the inference time, cannot bring performance higher than the (expert) demonstrations in the context. Therefore, ICL does not allow the LLM to learn from its own experience (Silver and Sutton, 2024) to self-improve.

Reinforcement Learning (RL, Sutton and Barto (2018)) is perhaps the most successful human designed framework for sequential decision making. In this work, we demonstrate that, surprisingly,

*Correspondence to kefan@email.virginia.edu, {amoeini, shangtong}@virginia.edu.

RL also emerges in LLM’s inference time. Such inference-time reinforcement learning process is commonly called in-context reinforcement learning (ICRL, Laskin et al. (2023); Moeini et al. (2025)). In this paper, we propose a multi-round prompting framework called **ICRL prompting**, enabling LLM’s self-improvement during the inference time through (reinforcement) learning. The goal of ICRL prompting is to prompt the LLM to complete a task. Initially, the prompt is only the task description. After the LLM generates a response, we give numerical scalar feedbacks for the response, called the rewards. Then in the next round, we prompt the LLM again with the same task description and a context consisting of all previous responses and rewards. So on and so forth. We observe that the quality of the LLM’s response increases as the context grows. In other words, the LLM is able to maximize the scalar reward signal during the inference time, just like an RL algorithm. This demonstrates the potential of learning methods for LLM’s self-improvement during the inference time. Notably, the scalar reward signal is the only feedback we provide to the LLM. This complies with both the reward hypothesis (Sutton, 2004), “*that all of what we mean by goals and purposes can be well thought of as maximization of the expected value of the cumulative sum of a received scalar signal (reward)*”, and the “reward is enough” hypothesis (Silver et al., 2021), “*intelligence, and its associated abilities, can be understood as subserving the maximisation of reward*”.

To summarize, this paper makes three contributions:

- (1): We propose the ICRL prompting framework that demonstrates LLM’s inference-time self-improvement through RL. One key technical challenge is designing this framework to be as minimal as possible, ensuring that the observed improvements clearly stem from the LLM’s intrinsic capacity of ICRL, rather than from auxiliary external code or additional engineered LLM interactions. Specifically, our framework avoids the use of techniques such as alternating different sampling temperatures (Zhang et al., 2024), extra sampling steps (Yang et al., 2024), any form of prioritized experience replay (e.g., subsampling, selecting or re-ordering past trials based on rewards), textual descriptions of reward or textual gradients (Yuksekgonul et al., 2025) from additional LLM nodes. By adhering to this minimal design, we provide strong evidence that by merely providing instructions on exploration and exploitation, we can elicit the ICRL capabilities from the pretrained LLM itself.
- (2): We provide strong evidence suggesting the emergence of RL in LLM’s inference time when the ICRL prompting framework is used. Specifically, we demonstrate the maximisation of the scalar reward signal, the exploration-exploitation trade-off in LLM’s inference time, the performance improvements from the growth of the context, the performance drop with short context, and the performance drop when the reward is absent. All those observations are well expected for an RL algorithm. Essentially, this is a “duck test” (Heim, 2007)² for the inference process.
- (3): We demonstrate significant performance improvements of ICLR prompting over baseline methods like Self-Refine (Madaan et al., 2023) and Reflexion (Shinn et al., 2023), in three different benchmarks including Game of 24, creative writing, and ScienceWorld. In Game of 24 and creative writing, the scalar reward signal is generated by the LLM itself. Yet performance improvements are still observed.

2 Background

Reinforcement Learning. RL uses Markov Decision Processes (MDPs) to model a task, consisting of a state space \mathcal{S} , an action space \mathcal{A} , a reward function $r : \mathcal{S} \rightarrow \mathbb{R}$, an initial distribution $p_0 \in \Delta(\mathcal{S})$ with $\Delta(\mathcal{S})$ denoting the set of probability distributions over \mathcal{S} , and a transition function $p : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$. At time step 0, an initial state S_0 sampled from p_0 . At time t , an agent at S_t takes an action A_t according to its policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ with $\Delta(\mathcal{A})$ denoting the set of probability distributions over \mathcal{A} , i.e., $A_t \sim \pi(S_t)$. The action A_t is then executed, after which the agent transitions to a successor state $S_{t+1} \sim p(S_t, A_t)$ and receives a reward $R_{t+1} \doteq r(S_{t+1})$. This agent-environment interaction continues until a time T , which marks the end of an episode. The goal of the agent is to adapt its policy π such that the expected total rewards $J(\pi) \doteq \mathbb{E}[\sum_{t=1}^T R_t]$ is maximized. In modern deep RL (Mnih et al., 2015; Schulman et al., 2017), the policy π is usually parameterized by a neural network. We use θ to denote the network parameter and write the policy as π_θ . Typically, RL algorithms update θ to adapt its policy. For example, at time t , the action A_t is sampled from $\pi_{\theta_t}(S_t)$. The RL algorithm then updates θ_t to θ_{t+1} based on available information such as $S_0, A_0, R_1, \dots, S_t, A_t, R_{t+1}, S_{t+1}$. Then at time $t+1$, the action A_{t+1} is sampled from the updated policy $\pi_{\theta_{t+1}}(S_{t+1})$. Essentially, the RL process is reflected in the updates of $\{\theta_t\}$ and we call this in-weight RL.

²If it looks like a duck, swims like a duck, and quacks like a duck, then it probably is a duck.

In-Context Reinforcement Learning. ICRL (Moeini et al., 2025), first coined by Laskin et al. (2023), is an emerging inference-time compute paradigm where the RL process occurs in the inference time (i.e., the forward pass) of the network without any parameter update. In ICRL, the policy π_θ is additionally conditioned on a context called C_t , i.e., $A_t \sim \pi_\theta(S_t, C_t)$. The construction of C_t is an active research area but one example is to use all previous state-action-reward pairs obtained in the task. Notably, this usually includes state-action-reward pairs from all previous episodes, not just the current episode (Laskin et al., 2023). In ICRL, there is a pretraining stage where the network θ is pretrained in a wide range of tasks (MDPs). We use θ_* to denote the parameter after the pretraining. After the pretraining stage, the policy π_{θ_*} is evaluated in new tasks. In other words, in the new MDP, the action A_t is sampled from $\pi_{\theta_*}(S_t, C_t)$. Importantly, the θ_* is kept fixed. Nevertheless, it is observed that the quality of A_t increases as C_t grows in the new task. Since θ_* is fixed, this improvement can only come from the increase of the context. This is thus called in-context policy improvement. Notably, this in-context policy improvement is also observed even when the new task is out of the distribution of the pretraining tasks, e.g., Laskin et al. (2023) demonstrate in-context policy improvement in new bandit problems that have the opposite optimal arms to the pretraining bandit problems. Thus this in-context policy improvement cannot be attributed to the hypothesis that θ_* memorizes the pretraining tasks. The only plausible hypothesis seems to be that the forward pass of the network parameterized by θ_* implements some RL algorithm to process the information in the context C_t to generate the action A_t . This inference-time (forward pass) RL is called in-context RL.

LLMs. The token generation process of LLMs can be modeled via RL. In short, the state is all generated tokens and the action is the next token to generate. Namely, let \mathcal{V} be the set of all possible tokens. We consider a state space $\mathcal{S} \doteq \bigcup_{i=1}^{\infty} \mathcal{V}^i$ and an action space $\mathcal{A} \doteq \mathcal{V}$. At time step 0, an initial prompt is given, denoted as $S_0 \in \mathcal{S}$. In this work, S_0 contains a description of a task. We refer to the LLM with parameter θ as π_θ . At time t , given the current tokens S_t , a new token A_t is sampled from $\pi_\theta(S_t)$. The new state is then $S_{t+1} = [S_t A_t]$, i.e., the new state is obtained by concatenating current tokens and the new token. A reward signal $R_{t+1} \doteq r(S_{t+1})$ is then emitted according to a reward function r . This token generation process continues until a time T , where either T is the maximal allowed response length or A_{T-1} is a special end-of-sequence token. Either way, this marks the end of an episode and the final state S_T , called the terminal state, contains both the initial task description and the LLM's response. There are two types of reward functions. One is sparse (the outcome reward model, Ouyang et al. (2022)), where $r(s)$ is nonzero only when s is a terminal state. The other is dense (the progress reward model, Lightman et al. (2023)), where $r(s)$ can also be nonzero for non-terminal states.

3 In-Context Reinforcement Learning Prompting

We now present our main contribution, the ICRL prompting framework (Algorithm 1, Figure 1), consisting of the following ingredients.

Algorithm 1 ICRL Prompting

Require: An LLM π_θ . A reward function r . Number of episodes K . An experience buffer \mathcal{B} .
A task description $s_{\text{task}} \in \mathcal{S}$. The ICRL instruction $s_{\text{ICRL}} \in \mathcal{S}$.

- 1: **for** $k = 1$ to K **do**
- 2: Construct the initial prompt S_0 by concatenating all the tokens in \mathcal{B} , s_{task} , and s_{ICRL} .
- 3: $t \leftarrow 0$ // Execute the policy π_θ starting from S_0
- 4: **while** S_t is not terminal **do**
- 5: $A_t \sim \pi_\theta(S_t)$, $S_{t+1} \doteq [S_t A_t]$, $R_{t+1} \doteq r(S_{t+1})$, $t \leftarrow t + 1$
- 6: **end while**
- 7: // $[A_0 A_1, \dots, A_{T-1}]$ is the LLM's response to s_{task} at the current episode
- 8: Push $(A_0, R_1, A_1, R_2, A_2, R_3, \dots, A_{T-1}, R_T)$ into \mathcal{B} .
- 9: **end for**

LLM as the Policy. An LLM, denoted as π_θ , serves as the policy network. The goal is to prompt the LLM to solve a task. We assume a natural language description of the task is available and we denote it as $s_{\text{task}} \in \mathcal{S}$. At the beginning of each episode, we construct the initial prompt by concatenating the LLM's own previous attempts together with the corresponding rewards, the task description, and some meta instruction denoted as s_{ICRL} . The details of the concatenation of previous attempts and the

choice of the meta instruction will be discussed shortly. With this initial prompt, the LLM generates the response. Both the response and the rewards are stored in the buffer for future episodes.

Reward Function. A numerical scalar reward feedback is provided for each S_t in the episode. Notably, the reward can be either sparse (i.e., only R_T is nonzero) or dense. The reward function can be rule-based, learned separately, or instantiated via the same LLM for self-evaluation. The flexibility of using LLM’s self-evaluation as the reward function allows the ICRL prompting framework to be applied to a wide range of tasks. Notably, this scalar reward is the only feedback we provide to the LLM. But we do tell the LLM that this scalar is a reward. We do so by explicitly writing down the word “Reward” before this number when constructing the initial prompt. Notably, if the reward function is rule-based and learned separately, the reward signal constitutes an external feedback. But if the reward function is just the LLM’s own evaluation of the answer, there is no external feedback at all in the ICRL prompting framework. Yet we still expect the LLM’s response to improve over the episode. This is because of the widely believed hypothesis that evaluation is easier than generation. But we do hypothesize that the ceiling with self-evaluation is lower than that with external feedback.

Memory for Experience. We use an experience buffer \mathcal{B} to store the LLM’s responses and rewards for the task in previous episodes. Our underlying hypothesis is that pretrained LLM already has the ICRL ability. To use this innate ICRL ability to improve LLM’s response to the task, we concatenate its previous attempts and rewards as many as the context window allows. We expect that the LLM can reinforcement learn from the experiences in the context during the inference time.

ICRL Instructions. To facilitate LLM’s inference time RL, we additionally provide some instructions in initial prompt S_0 at each episode. The instruction is in natural language and is denoted as s_{ICRL} . We consider three types of instructions: (1) the exploration instruction (Figure 4 in App. A), (2) the exploitation instruction (Figure 5 in App. A), (3) the exploration or exploitation instruction (Figure 6 in App. A). For exploration instruction, the model is asked to provide a response that is different from all its previous responses. For exploitation instruction, the model is asked to generate the best response based on the the previous responses with the highest reward. We consider two strategies. **(1) ICRL Preset:** We alternate between the exploration and exploitation instructions. When the episode number K is even, we use the exploration instruction. When the episode number K is odd, we use the exploitation instruction. **(2) ICRL Autonomous:** We always provide the “exploration or exploitation” instruction at each episode and let the LLM itself to decide on which to use.

4 Related Works

4.1 In-Context Reinforcement Learning.

The study of inference time RL algorithms dates back to Duan et al. (2016); Wang et al. (2016) and recently Laskin et al. (2023) coined the word “in-context reinforcement learning”, after which the interest in ICRL grew quickly (Kirsch et al., 2023; Raparthy et al., 2023; Schmied et al., 2024; Lee et al., 2024; Zisman et al., 2023; Grigsby et al., 2024; Lu et al., 2023; Bauer et al., 2023; Wang et al., 2025; Cook et al., 2024; Xu et al., 2024; Shi et al., 2024; Huang et al., 2024; Liu and Abbeel, 2023; Dai et al., 2024) See Moeini et al. (2025) for a comprehensive survey. Most existing ICRL works are essentially a subarea of meta RL (Beck et al., 2023). They use games or robotics as benchmark and the network size is small (compared with pretrained LLMs like GPT-4.1). They focus on algorithms to induce this ICRL capability and train sequence models from scratch. By contrast, this work studies the ICRL capability of standard pretrained LLMs without any parameter updates. Nevertheless, some existing ICRL works do include pretrained LLM. For example, Brooks

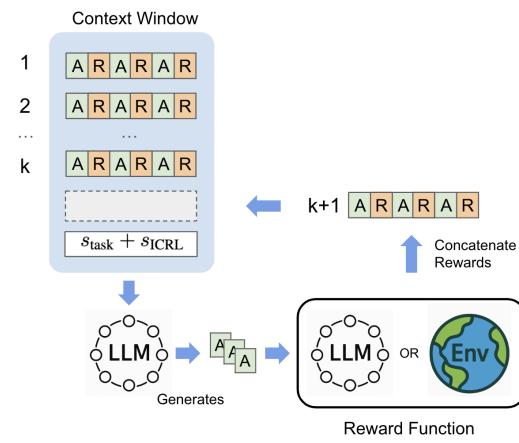


Figure 1: ICRL Prompting. At each episode $k + 1$, LLM generates action tokens based on previous experiences up to k , and receives numerical rewards either from itself as the evaluator or from the environment. At the end of the episode, the rewards are then concatenated with the action tokens and placed back into the context.

et al. (2024) use a pretrained LLM as a world model to simulate rollouts; Mirchandani et al. (2023) treated LLMs as general sequence predictors. Closest works to ours study LLMs in bandit tasks, specifically Multi-Arm Bandits (MAB) (Krishnamurthy et al., 2024; Nie et al., 2024; Park et al., 2024) and Contextual Bandits (CB) (Nie et al., 2024; Monea et al., 2025). They show that artificial interventions, such as including statistics commonly used in MAB algorithms in the context, can improve LLM’s performance on various MAB and CB tasks. In particular, Monea et al. (2025) reformulate a classification task as a CB problem by assigning a binary reward based on classification correctness. Their comprehensive empirical results indicate that LLMs perform poorly on this task without either prompts that encourage exploration or additional mechanisms to filter incorrect actions from the context. Overall, these works demonstrate that LLMs remain uncompetitive with algorithmic baselines for both CB and MAB, even after such interventions. We shift our focus away from tasks solvable by classic analytic bandit algorithms and instead tackle text tasks that require real-world knowledge about math, science, and linguistics, domains where LLMs uniquely excel. Our experiments show that interaction history alone, when paired with numerical rewards, outperforms baseline test-time self-improvement methods, including various prompting and sampling techniques. Importantly, we do not filter any “failed” action with low rewards since a key characterization of RL is to learn from failure as well.

Broadly speaking, ICRL can be viewed as a subarea of in-context learning (ICL, Brown et al. (2020)) if ICL is taken to encompass any inference-time learning paradigm. In practice, however, ICL often refers specifically to in-context supervised learning (ICSL) (Brown et al., 2020). The key difference between ICRL and ICSL is that in ICSL, the context consists of task-demonstration pairs, where demonstrations are typically generated by an expert. In contrast, ICRL contexts contain tasks, previous attempts, and associated rewards. Agarwal et al. (2024) show that scaling ICL with synthetic answers can outperform few-shot expert ICL, but their rule-based verification assumes access to many auxiliary problems drawn from the same distribution as the current task. In our setting we interact only with the current task and receive either a synthetic or a ground truth scalar reward, without access to expert or successful trajectories. Like any reinforcement learning method, ICRL exploits both rewards and trajectories, allowing the model to improve on the current task by learning from all of its prior failures.

4.2 Inference-Time LLM Self-Improvement

Chain-of-Though (CoT) (Wei et al., 2022) is the most commonly used method for inference-time improvement, where the LLM is prompted to reason step by step. **Long CoT** (OpenAI, 2024; DeepSeek-AI et al., 2025) extends this idea by training a LLM to generate extensive, multi-step reasoning processes to tackle complex problems. It has demonstrated that performance increases with the number of tokens generated at test time (OpenAI, 2024). **Best-of-N** is a simple yet effective way of optimizing LLM’s output at inference time. At each query, the model draws N independent samples, evaluates each candidate with a numerical score, and then selects the highest-scoring output. **ReAct** (Yao et al., 2023b) is a prompting method for long trajectory decision making. In ReAct, the LLM is prompted to complete a task through multiple steps. At each step, the LLM first generates a “thought”, reflecting on the problem, then decides on an “action” to gather new information or update state, and finally ingests the action’s result into its next thought. ReAct alternates between reasoning and action to complete a single trial (i.e., one episode in Algorithm 1) for the task. Our ICRL prompting is orthogonal to ReAct in that we study inference-time RL through multiple episodes. **Self-Refine** (Madaan et al., 2023) is the closest to our work and is a multi-round self-improvement method to improve the LLM’s response iteratively. First, the LLM generates an initial response to a prompt. Then, the same LLM provides verbal feedback on its own output, identifying areas for improvement. Finally, the LLM uses this feedback to revise and refine its previous response. This cycle is repeated until the output meets a desired quality or a stopping condition is reached. Our ICRL prompting is different from Self-Refine in two aspects. First, Self-Refine uses the LLM’s own natural language feedback. By contrast, ICRL prompting uses a numerical scalar feedback (from either the LLM itself or external if available). Second, our ICRL prompting additionally provides instructions to elicit exploration and exploitation behavior. One of our key contribution is to demonstrate the effectiveness of scalar feedback over verbal feedback. **Reflexion** (Shinn et al., 2023) is a multi-round self-improvement method for LLMs. At each iteration, the model generates an answer, receives an external feedback signal (e.g., a correctness score or error flag), and is then prompted to “reflect” on what went wrong and propose a plan for trying again. Those reflections are stored in a simple

memory and prepended to the next prompt. Ideally, in the next round, Reflexion proposes to provide all previous reflections in the new prompt. But in practice, since the reflection is long, usually only 3 most recent reflections are provided in the new prompt. Furthermore, the new prompt does not include the feedback signal (i.e., the rewards) nor the previous responses. By contrast, our ICRL prompting provides all previous responses and rewards in the new prompt. **Textual Gradient** (Yuksekgonul et al., 2025) prompts an LLM for an answer, then directs it to generate verbal feedback by comparing its output to a ground-truth evaluation. This verbal feedback is subsequently used to revise the prompt itself for future attempts. While Textual Gradient can use numerical scores as ground truth, their role is to inform the generation of this verbal feedback, and only the verbal feedback is used to refine the prompt. By contrast, our framework keeps the original task prompt unchanged. ICRL improves performance solely by growing the context by appending the complete history of previous responses along with their rewards.

A parallel line of research employs search-based methods for inference-time improvement in LLMs. Examples include Tree-of-Thoughts (ToT) (Yao et al., 2023a), Graph-of-Thoughts (GoT) (Besta et al., 2024), Monte Carlo Tree Search (MCTS) (Ding et al., 2024), and more recently, Intelligent Go-Explore (Lu et al., 2025). In general, their performance heavily relies on externally engineered components, such as search heuristics or memory management mechanisms, rather than directly harnessing the intrinsic learning capabilities within LLMs. Our work is also related to LLM as optimizers. Previous work on prompt optimization (Yang et al., 2024) shows that LLMs can leverage numerical scores to improve prompts, however, Yang et al. (2024) rely on selecting the top 20 prompts and filtering out error cases in context, thereby aligning it more with in-context supervised learning (e.g., filtered behavior cloning, Grigsby et al. (2024)) than ICRL. Again, one key characterization of our ICRL prompting is that we allow learning from failure.

5 Experiment

We evaluate ICRL prompting on three benchmarks: Game of 24, creative writing from Yao et al. (2023a), and ScienceWorld (Wang et al., 2022). We compare several baselines including **CoT-only**, **Long-CoT** style prompting, **Best-of-N**, **Self-Refine**, and **Reflexion**. Notably, in all the experiments, we allow the prompt of Self-Refine and Reflexion to grow as long as the LLM allows.

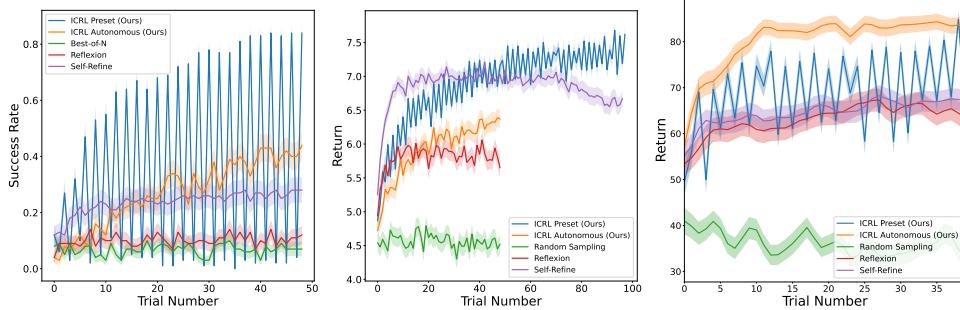


Figure 2: **Baseline Method Comparison.** **(Left)** Mean Success Rate on Game of 24. **(Middle)** Mean Coherence Reward on Creative Writing. Both ICRL Preset and Self-Refine went through an additional run of 50 episodes. **(Right)** Mean Return on Science World. A running max version of the plots is available in Figure 13 in App. B. This plot shows quality of the response at the current trial while the running max version shows the quality of the best response until now.

5.1 Game of 24

Task Setup. Given four input numbers, the model must use each number exactly once and apply only addition, subtraction, multiplication, or division in any order to reach 24. We choose the GPT-4.1 model for this experiment because of its excellent long-context capacity (OpenAI, 2025), accessed through api calls. Following Yao et al. (2023a), we use CoT prompting to elicit the model to provide a step-by-step solution, where each step the model picks two numbers from the remaining numbers and an operation to perform the calculation, and obtains one complete solution per LLM query, containing a total of 4 thinking steps. To ensure that the LLM generates response

with correct format, we additionally provide 5 in-context supervised learning demonstrations. The CoT instruction and the 5 demonstrations together form the task description s_{task} (Figure 7 in App. A).

Evaluation. To verify the correctness of the solutions, we leverage SymPy (Meurer et al., 2017), a library for symbolic mathematics, by extracting operands and operators and evaluating the reconstructed expression to confirm it equals 24, and report the mean success rate over the 100 problems. This rule-based success rate is referred to as r_* (i.e., the ground truth reward function) in the rest of the paper. We use r to denote the reward function that the algorithms actually have access to. In particular, we use GPT-4.1 as the r , the same LLM as the policy LLM but prompted differently (see Figure 12 in App. A). After the policy LLM generates the response, for each thinking step, GPT-4.1 scores the likelihood of reaching 24 with the remaining numbers on a 0-3 scale (0 = impossible, 3 = sure). The task is challenging in that no algorithm has access to r_* . Instead, they have to rely on their own (possibly imperfect) evaluation, generated by the same LLM, to improve the response.

Baselines. We compare our method with CoT-only, Long-CoT style prompting, Best-of-N, Reflexion, and Self-Refine. In CoT-only prompting, the model receives only the task description s_{task} and produces a single step-by-step solution. In Long-CoT style prompting, we explicitly ask the LLM to generate a long chain-of-thought, and keep retrying if the solution is incorrect in "<think>...</think>" tags, before finally providing the answer. Although GPT-4.1 is not specifically trained for long-form CoT reasoning, we find that Long-CoT style prompting can elicit significantly longer and self-correcting thought traces compared to zero-shot prompting, making it a strong baseline for the Game of 24. Both methods cannot make use of a reward signal and are run for one pass. For Best-of-N, to make it even stronger, we use the ground truth reward r_* to select the best response. Self-Refine does not require a reward function. It instead asks the LLM itself to provide textual verbal feedback. Reflexion generates reflection according to r . ICRL prompting is different from Self-Refine and Reflexion in that it uses the reward r directly, without any verbal feedback. So the comparison between ICRL prompting and Self-Refine / Reflexion is essentially the comparison between scalar feedback and verbal feedback.

ICRL prompting. As discussed before, both π_θ and r in Algorithm 1 are the same LLM GPT-4.1 (prompted differently). We now clarify how we compute S_0 at each episode. Since each action is a token, not all actions receive a reward. In fact, since we use CoT to prompt the LLM for 4 thinking steps, only 4 rewards are available for each episode. We thus only include those 4 rewards in S_0 . We add a "Reward: " tag before the actual scalar reward and then concatenate the tagged reward immediately after the corresponding action (i.e., token).

Results. The success rate (i.e., r_*) against the number of trials (i.e., the episodes in Algorithm 1) is reported in Figure 2. The ICRL Preset method achieves the highest performance, and the observed oscillations in success rate reflect the model’s alternating phases of exploration and exploitation. The mean of running max is also plotted in Figure 13 in App. B. For each problem, we compute its running maximum success rate up to each episode and then average these values across all problems at every episode. As summarized in Table 1, after 50 trials, our methods achieve a success rate of 90% which is significantly higher than 49% from Best-of-N sampling, 47% from Self-Refine, and 44% from Reflexion.

Table 1: Game of 24 Success Rate. The running max success rate of the last episode is reported.

Method	Success Rate
CoT-only	6%
Long-CoT	47%
Reflexion	44%
Best-of-N	49%
Self-Refine	47%
ICRL Preset (Ours)	90%
ICRL Autonomous (Ours)	84%

5.2 Creative Writing

Task Setup. We consider the creative writing task from Yao et al. (2023a), where four sentences are randomly sampled from a pool of sentences. The task for LLMs is to generate four paragraphs, each ending with a sentence, while ensuring that the generated passage is coherent. This is a difficult task, as it challenges the LLMs to craft a unified storyline that logically justifies each of the four sampled sentences by weaving them into a single narrative. A total of 100 problems are evaluated.

An example of s_{task} is in Figure 8 in App. A.

Evaluation. We evaluated model outputs using the Length-Controlled Alpaca-Eval 2 (Tatsu Lab, 2025) framework, an automated evaluation tool that is widely used as a proxy for human annotators (Hong et al., 2024; Ethayarajh et al., 2024; Meng et al., 2024) and has demonstrated a Pearson correlation of up to 0.98 with human judgments. For each of the 100 creative writing problems, we present the instruction alongside each method’s top response. For Reflexion and Best-of-N, the top response is selected as the one that receives the highest reward r among 50 trials. For ICRL prompting and Self-Refine, we simply use the response generated in the 50th episode. The pairwise length-controlled win rates are then computed by the Alpaca-eval framework and is denoted as the reward function r_* . We now introduce the reward function r that the algorithms have access to. We find prompting the LLM to directly provide an absolute numerical score to a response, even with explicit rubrics, induces a lot of variance in the reward values. Therefore, we follow the standard practice of pairwise comparison of LLM-as-a-judge (Zheng et al., 2023) for evaluating a numerical score from 1-10 with a reference answer. Since this particular creative writing task is especially challenging for coherence, we query GPT-4.1 to generate a coherent paragraph to create a single reference answer for all pairwise comparisons, serving as a reliable foundation for reward signals (see Figure 11 in App. A). This pairwise comparison reward function is denoted as r . Notably, although both r and r_* use an LLM as a judge, they are prompted very differently. The reward function r evaluates responses to 100 questions against a single, coherent base text. By contrast, r_* performs pairwise comparison between two responses generated by our method and a baseline method. In addition, r is prompted to specifically focus on coherence rating, which is the main challenge of this task. Empirically, we found that directly learning from r_* is difficult for the LLM, potentially due to the unstable quality of the different base texts.

Baselines. We compare our method with Best-of-N, Reflexion, and Self-Refine. We allow Best-of-N, Reflexion and ICRL prompting to use r . Self-Refine do not use r and instead asks GPT-4.1 to provide verbal feedback. Since it is hard to distinguish CoT and Long-CoT style prompting for this task, we include Long-CoT style prompting as the baseline.

ICRL prompting. GPT-4.1 is used as both the policy LLM π_θ and the reward model r in Algorithm 1. At each episode, the initial prompt S_0 is constructed by concatenating all of the previous generations along with their coherence scores from r . Notably, this reward is sparse and only R_T can be nonzero. We, therefore, only include R_T in constructing S_0 .

Results. Our method achieves a length-controlled win rate of 59.48% against Reflexion, 78.36 % against Long-CoT style prompting, 86.32 % against Self-Refine, and 93.81 % against Best-of-N as shown in Table 2. This shows the responses generated by our method outperform the ones by baselines in terms of following the instruction to write coherent paragraphs and achieving better human preference. The return curve from reward model r is plotted in Figure 2, and a running max of the return is plotted in Figure 13 in App. B. Although Self-Refine initially matches ICRL in terms of coherence reward, extending both methods by 50 additional episodes, our methods keep improving, whereas self-refine first plateaus, then declines, likely due to the significant growth of its context.

5.3 ScienceWorld

Task Setup. ScienceWorld (Wang et al., 2022) is an interactive, text-based benchmark consisting of 30 science-experiment tasks set in a multi-room environment populated with diverse objects. An agent is given a task such as “*Grow an apple by cross-pollination*” or “*Experiment and find the surface with the highest friction*”, and can perform actions like “*Move to workshop*” or “*Use lighter on wood*”. The environment is challenging due to sparse rewards, large action spaces, and the requirement for scientific knowledge and efficient exploration. At each step, the agent observes the result of its action and receives zero reward unless it completes a predefined subgoal. This reward signal is used both for evaluation and for inference-time self-improvement (i.e., r and r_* are identical in this task). Completing all subgoals yields a cumulative reward of 100 and marks the episode as successful. An episode ends in failure if the agent either reaches the maximum number of steps or executes an incorrect terminating action. The input s_{task} provided to the agent describes the environment, the task, and the template of all possible actions. An example of s_{task} is provided in App. A

Table 2: Length-Controlled Win Rate (LC) and Standard Error (SE) from Alpaca-Eval 2.0 on Creative Writing.

Comparison	LC \pm SE (%)
Ours vs Reflexion	59.48 \pm 3.47
Ours vs Long CoT	78.36 \pm 1.99
Ours vs Self-Refine	86.32 \pm 3.03
Ours vs Best-of-N	93.81 \pm 1.01

Evaluation. We use the environment-provided reward function for each task both to construct the trajectories used in the context (r), and to evaluate the model (r^*). We benchmark each method on all 30 tasks and aggregate the results. GPT-4.1 mini is used as the policy for all compared algorithms.

Baselines. In Reflexion, at the end of each episode, the agent is prompted to reflect on its attempt. The reflection is then sanitized and appended to a reflection buffer, which is formatted into the context for subsequent trials. Self-Refine similarly generates self-feedback, but appends it to a trajectory summary, which is then added to the buffer. To ensure a fair comparison, we allow these methods access to the reward signals of the current episode (unlike ICRL) before prompting for reflection.

ICRL Setup. Each trial corresponds to a single episode in the environment. After the trial, the new trajectory added to the buffer is constructed by concatenating the actions, observations, rewards, and the final outcome (success or failure). As each episode typically yields only a few rewards, we include only those. At the start of each trial, S_0 is constructed by concatenating the task description S_{task} , the collected trajectories, and then the instruction s_{ICRL} . An example of S_0 is provided in App. A.

Results. The mean return at each trial, is presented in Figure 2 Right. Steady improvements are observed for methods that make use of some form of history of interactions similar to ICRL prompting. However, ICRL prompting outperforms baseline methods by about 20% after enough iterations. To make the comparison fair for efficient baselines such as Best-of-N, in App. B, we compare baselines as we scale test-time compute budget and observe that ICRL also scales better than the baselines not only in terms of number of trials but also the test-time compute budget (in dollar amounts).

5.4 Ablation Study

To better understand ICRL prompting, we consider following ablations. **(1) Zero Rewards:** We set all rewards to 0. **(2) Short Context:** In Algorithm 1, the buffer \mathcal{B} is essentially a queue of infinite length. Instead, we now make it a deque of length 3. In other words, only the recent 3 episodes are used in constructing S_0 . **(3) Exploration Only.** We simply ask the LLM to provide a different response than the ones in context, using the exploration instruction as s_{ICRL} , without the reward signal. **(4) Exploitation Only.** We always use the exploitation instruction as s_{ICRL} , with the reward signal. **(5) No E&E.** We entirely remove s_{ICRL} .

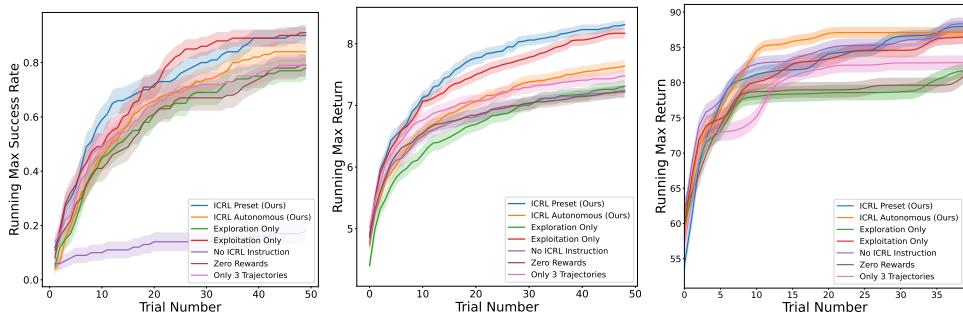


Figure 3: **Ablation Studies (Running Max).** **(Left)** The mean of running max success rate on Game of 24. **(Middle)** The mean of running max coherence reward on creative critiquing. **(Right).** The mean of running max return on ScienceWorld.

Results. The running max results of the ablation study are plotted in Figure 3. Both our two methods and the exploitation only with reward signals demonstrate the best-performing curves. This demonstrates our ICRL prompting framework is quite robust to the different prompts setup. We have also observed performance drop with short context and performance drop when the reward is absent. A key finding is that the “exploration only without reward signal” method (shown in green) performs significantly worse than our approach when comparing the maximum performance achieved over time (running max). This demonstrates that our method’s improvement is not just due to exploring various responses and then picking the best one previously seen as doing a Best-of-N. Instead, ICRL can genuinely generate novel responses that are better than the ones during the exploration phase.

6 Conclusion

This work proposes a novel ICRL prompting framework for inference-time LLM self-improvement, using LLM’s inherent ICRL ability to maximize a scalar reward signal. Most prior works (Section 4.2) instead rely on verbal textual reflections, which rely primarily on LLM’s pretrained parametric knowledge, essentially making them more like a knowledge-guided search instead of (reinforcement) learning (Liu et al., 2025). While it seems that textual descriptions would be the more natural option as feedback for LLMs, it is shown that such textual LLM self-verifications are filled with hallucinations and misleading feedbacks and lead to performance collapse over self-correcting iterations (Stechly et al., 2025). Given the proven success of numerical feedbacks in RL (Sutton, 2004; Silver et al., 2021) and the demonstrated performance improvement over baselines in our experiments, we conjecture that the numerical feedback might be a competing alternative for the verbal feedback.

Acknowledgments

This work is supported in part by the US National Science Foundation under grants III-2128019 and SLES-2331904 and by the Coastal Virginia Center for Cyber Innovation (COVA CCI) and the Commonwealth Cyber Initiative (CCI), an investment in the advancement of cyber research & development, innovation, and workforce development. For more information about CCI, visit www.covacci.org and www.cyberinitiative.org.

References

- Agarwal, R., Singh, A., Zhang, L. M., Bohnet, B., Rosias, L., Chan, S., Zhang, B., Anand, A., Abbas, Z., Nova, A., Co-Reyes, J. D., Chu, E., Behbahani, F., Faust, A., and Larochelle, H. (2024). Many-shot in-context learning. In *Advances in Neural Information Processing Systems*.
- Bauer, J., Baumli, K., Behbahani, F., Bhoopchand, A., Bradley-Schmiege, N., Chang, M., Clay, N., Collister, A., Dasagi, V., Gonzalez, L., et al. (2023). Human-timescale adaptation in an open-ended task space. In *Proceedings of the International Conference on Machine Learning*.
- Beck, J., Vuorio, R., Liu, E. Z., Xiong, Z., Zintgraf, L., Finn, C., and Whiteson, S. (2023). A survey of meta-reinforcement learning. *ArXiv Preprint*.
- Besta, M., Blach, N., Kubicek, A., Gerstenberger, R., Podstawski, M., Gianinazzi, L., Gajda, J., Lehmann, T., Niewiadomski, H., Nyczyk, P., and Hoefer, T. (2024). Graph of thoughts: Solving elaborate problems with large language models. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Brooks, E., Walls, L., Lewis, R. L., and Singh, S. (2024). Large language models can implement policy iteration. In *Advances in Neural Information Processing Systems*.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems*.
- Cook, J., Lu, C., Hughes, E., Leibo, J. Z., and Foerster, J. (2024). Artificial generational intelligence: Cultural accumulation in reinforcement learning. *ArXiv Preprint*.
- Dai, Z., Tomasi, F., and Ghiassian, S. (2024). In-context exploration-exploitation for reinforcement learning. *ArXiv Preprint*.
- DeepSeek-AI, Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., Zhang, X., Yu, X., Wu, Y., Wu, Z. F., Gou, Z., Shao, Z., Li, Z., Gao, Z., Liu, A., Xue, B., Wang, B., Wu, B., Feng, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., Dai, D., Chen, D., Ji, D., Li, E., Lin, F., Dai, F., Luo, F., Hao, G., Chen, G., Li, G., Zhang, H., Bao, H., Xu, H., Wang, H., Ding, H., Xin, H., Gao, H., Qu, H., Li, H., Guo, J., Li, J., Wang, J., Chen, J., Yuan, J., Qiu, J., Li, J., Cai, J. L., Ni, J., Liang, J., Chen, J., Dong, K., Hu, K., Gao, K., Guan, K., Huang, K., Yu,

K., Wang, L., Zhang, L., Zhao, L., Wang, L., Zhang, L., Xu, L., Xia, L., Zhang, M., Zhang, M., Tang, M., Li, M., Wang, M., Li, M., Tian, N., Huang, P., Zhang, P., Wang, Q., Chen, Q., Du, Q., Ge, R., Zhang, R., Pan, R., Wang, R., Chen, R. J., Jin, R. L., Chen, R., Lu, S., Zhou, S., Chen, S., Ye, S., Wang, S., Yu, S., Zhou, S., Pan, S., Li, S. S., Zhou, S., Wu, S., Ye, S., Yun, T., Pei, T., Sun, T., Wang, T., Zeng, W., Zhao, W., Liu, W., Liang, W., Gao, W., Yu, W., Zhang, W., Xiao, W. L., An, W., Liu, X., Wang, X., Chen, X., Nie, X., Cheng, X., Liu, X., Xie, X., Liu, X., Yang, X., Li, X., Su, X., Lin, X., Li, X. Q., Jin, X., Shen, X., Chen, X., Sun, X., Wang, X., Song, X., Zhou, X., Wang, X., Shan, X., Li, Y. K., Wang, Y. Q., Wei, Y. X., Zhang, Y., Xu, Y., Li, Y., Zhao, Y., Sun, Y., Wang, Y., Yu, Y., Zhang, Y., Shi, Y., Xiong, Y., He, Y., Piao, Y., Wang, Y., Tan, Y., Ma, Y., Liu, Y., Guo, Y., Ou, Y., Wang, Y., Gong, Y., Zou, Y., He, Y., Xiong, Y., Luo, Y., You, Y., Liu, Y., Zhou, Y., Zhu, Y. X., Xu, Y., Huang, Y., Li, Y., Zheng, Y., Zhu, Y., Ma, Y., Tang, Y., Zha, Y., Yan, Y., Ren, Z. Z., Ren, Z., Sha, Z., Fu, Z., Xu, Z., Xie, Z., Zhang, Z., Hao, Z., Ma, Z., Yan, Z., Wu, Z., Gu, Z., Zhu, Z., Liu, Z., Li, Z., Xie, Z., Song, Z., Pan, Z., Huang, Z., Xu, Z., Zhang, Z., and Zhang, Z. (2025). Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *ArXiv Preprint*.

Ding, R., Zhang, C., Wang, L., Xu, Y., Ma, M., Zhang, W., Qin, S., Rajmohan, S., Lin, Q., and Zhang, D. (2024). Everything of thoughts: Defying the law of penrose triangle for thought generation. In *Findings of the Association for Computational Linguistics: ACL 2024*.

Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., and Abbeel, P. (2016). RL2: fast reinforcement learning via slow reinforcement learning. *ArXiv Preprint*.

Ethayarajh, K., Xu, W., Muennighoff, N., Jurafsky, D., and Kiela, D. (2024). Kto: Model alignment as prospect theoretic optimization. *ArXiv Preprint*.

Grigsby, J., Fan, L., and Zhu, Y. (2024). Amago: Scalable in-context reinforcement learning for adaptive agents. In *Proceedings of the International Conference on Learning Representations*.

Heim, M. (2007). *Exploring Indiana Highways*.

Hong, J., Lee, N., and Thorne, J. (2024). ORPO: Monolithic preference optimization without reference model. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.

Huang, S., Hu, J., Chen, H., Sun, L., and Yang, B. (2024). In-context decision transformer: Reinforcement learning via hierarchical chain-of-thought. *ArXiv Preprint*.

Kirsch, L., Harrison, J., Freeman, C., Sohl-Dickstein, J., and Schmidhuber, J. (2023). Towards general-purpose in-context learning agents. In *NeurIPS Foundation Models for Decision Making Workshop*.

Krishnamurthy, A., Harris, K., Foster, D. J., Zhang, C., and Slivkins, A. (2024). Can large language models explore in-context? *ArXiv Preprint*.

Laskin, M., Wang, L., Oh, J., Parisotto, E., Spencer, S., Steigerwald, R., Strouse, D., Hansen, S. S., Filos, A., Brooks, E., maxime gazeau, Sahni, H., Singh, S., and Mnih, V. (2023). In-context reinforcement learning with algorithm distillation. In *Proceedings of the International Conference on Learning Representations*.

Lee, J., Xie, A., Pacchiano, A., Chandak, Y., Finn, C., Nachum, O., and Brunskill, E. (2024). Supervised pretraining can learn in-context reinforcement learning. In *Advances in Neural Information Processing Systems*.

Lightman, H., Kosaraju, V., Burda, Y., Edwards, H., Baker, B., Lee, T., Leike, J., Schulman, J., Sutskever, I., and Cobbe, K. (2023). Let's verify step by step. *ArXiv Preprint*.

Liu, H. and Abbeel, P. (2023). Emergent agentic transformer from chain of hindsight experience. In *Proceedings of the International Conference on Machine Learning*.

Liu, M., Yu, C.-H., Lee, W.-H., Hung, C.-W., Chen, Y.-C., and Sun, S.-H. (2025). Synthesizing programmatic reinforcement learning policies with large language model guided search. In *Proceedings of the International Conference on Learning Representations*.

- Lu, C., Hu, S., and Clune, J. (2025). Intelligent go-explore: Standing on the shoulders of giant foundation models. In *Proceedings of the International Conference on Learning Representations*.
- Lu, C., Schroecker, Y., Gu, A., Parisotto, E., Foerster, J., Singh, S., and Behbahani, F. (2023). Structured state space models for in-context reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegreffe, S., Alon, U., Dziri, N., Prabhumoye, S., Yang, Y., Gupta, S., Majumder, B. P., Hermann, K., Welleck, S., Yazdanbakhsh, A., and Clark, P. (2023). Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*.
- Meng, Y., Xia, M., and Chen, D. (2024). Simpo: Simple preference optimization with a reference-free reward. *ArXiv Preprint*.
- Meurer, A., Smith, C. P., Paprocki, M., Certik, O., Kirpichev, S. B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J. K., Singh, S., Rathnayake, T., Vig, S., Granger, B. E., Muller, R. P., Bonazzi, F., Gupta, H., Vats, S., Johansson, F., Pedregosa, F., Curry, M. J., Terrel, A. R., Roucka, S., Saboo, A., Fernando, I., Kulal, S., Cimrman, R., and Scopatz, A. (2017). Sympy: symbolic computing in python. *PeerJ Computer Science*.
- Mirchandani, S., Xia, F., Florence, P., Ichter, B., Driess, D., Arenas, M. G., Rao, K., Sadigh, D., and Zeng, A. (2023). Large Language Models as General Pattern Machines. *ArXiv Preprint*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M. A., Fidjeland, A., 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*.
- Moeini, A., Wang, J., Beck, J., Blaser, E., Whiteson, S., Chandra, R., and Zhang, S. (2025). A survey of in-context reinforcement learning. *ArXiv Preprint*.
- Monea, G., Bosselut, A., Brantley, K., and Artzi, Y. (2025). Llms are in-context bandit reinforcement learners. *ArXiv Preprint*.
- Nie, A., Su, Y., Chang, B., Lee, J. N., Chi, E. H., Le, Q. V., and Chen, M. (2024). EVOLvE: Evaluating and Optimizing LLMs For Exploration. *ArXiv Preprint*.
- OpenAI (2024). Learning to Reason with LLMs. <https://openai.com/index/learning-to-reason-with-llms/>.
- OpenAI (2024). Openai o1 system card. *ArXiv Preprint*.
- OpenAI (2025). Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1/>.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*.
- Park, C., Liu, X., Ozdaglar, A., and Zhang, K. (2024). Do llm agents have regret? a case study in online learning and games. *ArXiv Preprint*.
- Raparthy, S. C., Hambro, E., Kirk, R., Henaff, M., and Raileanu, R. (2023). Generalization to new sequential decision making tasks with in-context learning. *ArXiv Preprint*.
- Schmied, T., Adler, T., Patil, V., Beck, M., Pöppel, K., Brandstetter, J., Klambauer, G., Pascanu, R., and Hochreiter, S. (2024). A Large Recurrent Action Model: xLSTM enables Fast Inference for Robotics Tasks. *ArXiv Preprint*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms. *ArXiv Preprint*.
- Shi, L. X., Jiang, Y., Grigsby, J., Fan, L., and Zhu, Y. (2024). Cross-episodic curriculum for transformer agents. In *Advances in Neural Information Processing Systems*.

- Shinn, N., Cassano, F., Berman, E., Gopinath, A., Narasimhan, K., and Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Silver, D., Singh, S., Precup, D., and Sutton, R. S. (2021). Reward is enough. *Artificial Intelligence*.
- Silver, D. and Sutton, R. S. (2024). Welcome to the era of experience.
- Stechly, K., Valmeeckam, K., and Kambhampati, S. (2025). On the self-verification limitations of large language models on reasoning and planning tasks. In *Proceedings of the International Conference on Learning Representations*.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D. M., Lowe, R., Voss, C., Radford, A., Amodei, D., and Christiano, P. (2022). Learning to summarize from human feedback. *ArXiv Preprint*.
- Sutton, R. S. (2004). The reward hypothesis. <http://incompleteideas.net/rlai.cs.ulberta.ca/RLAI/rewardhypothesis.html>.
- Sutton, R. S. (2019). The bitter lesson. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd Edition). MIT press.
- Tatsu Lab (2025). alpaca_eval: An evaluation framework for llms. https://github.com/tatsu-lab/alpaca_eval.
- Wang, J., Blaser, E., Daneshmand, H., and Zhang, S. (2025). Transformers can learn temporal difference methods for in-context reinforcement learning. In *Proceedings of the International Conference on Learning Representations*.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. (2016). Learning to reinforcement learn. *ArXiv Preprint*.
- Wang, R., Jansen, P., Côté, M.-A., and Ammanabrolu, P. (2022). ScienceWorld: Is your Agent Smarter than a 5th Grader? *ArXiv Preprint*.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.
- Xu, T., Li, Z., and Ren, Q. (2024). Meta-reinforcement learning robust to distributional shift via performing lifelong in-context learning. In *Proceedings of the International Conference on Machine Learning*.
- Yang, C., Wang, X., Lu, Y., Liu, H., Le, Q. V., Zhou, D., and Chen, X. (2024). Large language models as optimizers. *ArXiv Preprint*.
- Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., and Narasimhan, K. (2023a). Tree of thoughts: Deliberate problem solving with large language models. *ArXiv Preprint*.
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. (2023b). React: Synergizing reasoning and acting in language models. *ArXiv Preprint*.
- Yuksekgonul, M., Bianchi, F., Boen, J., Liu, S., Lu, P., Huang, Z., Guestrin, C., and Zou, J. (2025). Optimizing generative ai by backpropagating language model feedback. *Nature*.
- Zhang, S., Bao, Y., and Huang, S. (2024). Edt: Improving large language models' generation by entropy-based dynamic temperature sampling. *ArXiv Preprint*.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., and Stoica, I. (2023). Judging llm-as-a-judge with mt-bench and chatbot arena. *ArXiv Preprint*.
- Zisman, I., Kurenkov, V., Nikulin, A., Sinii, V., and Kolesnikov, S. (2023). Emergence of in-context reinforcement learning from noise distillation. *ArXiv Preprint*.

A Prompt Examples

Instruction: Examine all the <attempt>...</attempt> examples, each showing a candidate Response, and the Rewards for each step of the Response. Provide a response that is completely different for any steps from every single one of the previous attempts demonstrated in the context.

Figure 4: The Exploration Instruction (s_{ICRL}).

Instruction: You will be given multiple <attempt>...</attempt> examples, each showing a candidate Response, and the Rewards for each step of the Response. Your task: Based on the previous attempts, try your best to produce a response that can achieve higher rewards.

Figure 5: The Exploitation Instruction (s_{ICRL}).

Instruction: Examine all the <attempt>...</attempt> examples, each showing a candidate Response and its Reward. You have two options: exploration or exploitation.

For exploration, provide a response that is completely different for any steps from every single one of the previous attempts demonstrated in the context, while making sure it correctly follows the task instruction.

For exploitation, based on the previous attempts, try your best to produce a response that can achieve higher rewards.

Pick one option to follow.

Figure 6: The Exploration or Exploitation Instruction (s_{ICRL}).

Prompt: Write a coherent passage of 4 short paragraphs. The end sentence of each paragraph must be: For some unfathomable reason, the response team didn't consider a lack of milk for my cereal as a proper emergency. You realize you're not alone as you sit in your bedroom massaging your calves after a long day of playing tug-of-war with Grandpa Joe in the hospital. He poured rocks in the dungeon of his mind. I'm a living furnace. Make a plan then write. Your output should be of the following format: Plan: Your plan here. Passage: Your passage here.

Figure 8: An example of s_{task} for creative writing.

```

<attempt>
Input: 4 4 6 8
Step1: 4 + 8 = 12 (left: 4 6 12)
Step2: 6 - 4 = 2 (left: 2 12)
Step3: 2 * 12 = 24 (left: 24)
Answer: (6 - 4) * (4 + 8) = 24
</attempt>

<attempt>
Input: 2 9 10 12
Step1: 12 * 2 = 24 (left: 9 10 24)
Step2: 10 - 9 = 1 (left: 1 24)
Step3: 24 * 1 = 24 (left: 24)
Answer: (12 * 2) * (10 - 9) = 24
</attempt>

<attempt>
Input: 4 9 10 13
Step1: 13 - 10 = 3 (left: 3 4 9)
Step2: 9 - 3 = 6 (left: 4 6)
Step3: 4 * 6 = 24 (left: 24)
Answer: 4 * (9 - (13 - 10)) = 24
</attempt>

<attempt>
Input: 1 4 8 8
Step1: 8 / 4 = 2 (left: 1 2 8)
Step2: 1 + 2 = 3 (left: 3 8)
Step3: 3 * 8 = 24 (left: 24)
Answer: (1 + 8 / 4) * 8 = 24
</attempt>

<attempt>
Input: 5 5 5 9
Step1: 5 + 5 = 10 (left: 5 9 10)
Step2: 10 + 5 = 15 (left: 9 15)
Step3: 15 + 9 = 24 (left: 24)
Answer: ((5 + 5) + 5) + 9 = 24
</attempt>

```

Task: Use numbers and basic arithmetic operations (+ - * /) to obtain 24. Put your answer in this format '<answer>**Response** Step1: ... (left: ...) Step2: ... (left: ...) Step3: ... (left: ...) **Answer**: <math operations of the 4 input numbers, even if it does not equal 24></answer>'. Whether it is correct or not, do not try again.

Prompt: Input: 1 8 10 11

Figure 7: An example of s_{task} for Game of 24 with few-shot CoT prompting.

You are a helpful assistant to do some scientific experiment in an environment.

<Environment description> # *s_{task}*

In the environment, there are several rooms: kitchen, foundry, workshop, bathroom, outside, living room, bedroom, greenhouse, art studio, hallway

You should explore the environment and find the items you need to complete the experiment.

The available actions are: Available Actions

FOCUS is a extremely critical action that can be only used the number of times 'focus' is mentioned in the task description. Using it more than that or inappropriately (such as on a wrong object) will terminate the session and the task WILL FAIL.

Task Description:

Your task is to change the state of matter of water. First, focus on the substance. Then, take actions that will cause it to change its state of matter.

</Environment description>

<Instruction> *s_{ICRL}* </Instruction>

<Attempts> # **Buffer**

:

Attempt *n*:

teleport to bathroom -> Observation: You teleport to the bathroom. (reward=3)

-> focus on substance in toilet -> Observation: You focus on the water. (reward=66)

-> use cup on substance in toilet -> Observation: I'm not sure how to do that. (reward=0)

-> activate sink -> Observation: The sink is now activated. (reward=0)

-> dunk cup into sink -> Observation: No known action matches that input. (reward=0)

-> move cup to sink -> Observation: You move the glass cup to the sink. (reward=0)

-> teleport to kitchen -> Observation: You teleport to the kitchen. (reward=0)

-> use cup on stove -> Observation: I'm not sure how to do that. (reward=0)

-> activate stove -> Observation: The stove is now activated. (reward=2)

-> move cup to table -> Observation: You move the glass cup to the chair. (reward=0)

-> examine cup -> Observation: a glass cup (containing nothing) (reward=0)

Task Failed. You have exceeded the maximum number of steps. (reward=0) Total reward: 71

Attempt *n* + 1:

teleport to bathroom -> Observation: You teleport to the bathroom. (reward=3)

-> focus on substance in toilet -> Observation: You focus on the water. (reward=66)

-> pick up cup -> Observation: You move the glass cup to the inventory. (reward=0)

-> move cup to sink -> Observation: You move the glass cup to the sink. (reward=0)

-> activate sink -> Observation: The sink is now activated. (reward=0)

-> dunk cup into sink -> Observation: No known action matches that input. (reward=0)

-> pour substance in cup into sink -> Observation: You pour the water into the sink. (reward=0)

-> activate stove -> Your generated action "activate stove" cannot be matched to a valid action. (reward=0)

-> teleport to kitchen -> Observation: You teleport to the kitchen. (reward=0)

-> pick up cup containing nothing in table -> Observation: You move the glass cup to the inventory. (reward=0)

-> activate sink -> Observation: The sink is now activated. (reward=0)

Task Failed. You have exceeded the maximum number of steps. (reward=0) Total reward: 69

:

</Attempts>

Figure 9: An example of *S₀* from ScienceWorld.

```

# Buffer:
<attempt>
Input: 4 9 10 13.
Response:
Step1: 10 - 4 = 6 (left: 6 9 13)  <Reward: 3.00>
Step2: 13 - 6 = 7 (left: 7 9)  <Reward: 0.00>
Step3: 9 * 7 = 63 (left: 63)  <Reward: 0.00>
*Answer**: (13 - (10 - 4)) * 9 = 63  <Reward: 3.00>
</attempt>

<attempt>
Input: 4 9 10 13.
Response:
Step1: 10 + 4 = 14 (left: 9 13 14)  <Reward: 0.00>
Step2: 14 + 9 = 23 (left: 13 23)  <Reward: 0.00>
Step3: 23 + 13 = 36 (left: 36)  <Reward: 0.00>
*Answer**: (10 + 4 + 9) + 13 = 36  <Reward: 0.00>
</attempt>

<attempt>
Input: 4 9 10 13.
Response:
Step1: 9 + 10 = 19 (left: 4 13 19)  <Reward: 3.00>
Step2: 19 - 13 = 6 (left: 4 6)  <Reward: 3.00>
Step3: 6 + 4 = 10 (left: 10)  <Reward: 0.00>
*Answer**: ((9 + 10) - 13) + 4 = 10  <Reward: 6.00>
</attempt>
SICRL
Task

```

Figure 10: An example of S_0 from Game of 24.

Instruction: You are a seasoned text coherence evaluator. Read the TEXT below and rate its overall coherence on a scale from 1 to 10, where 1 means significantly less coherent than the Base Answer, 5 means equally coherent, and 10 means significantly more coherent. Be a strict and conservative evaluator-only assign high scores when the TEXT is clearly better than the Base Answer.

Base Answer:

{At dawn, golden light slips past pale curtains, rousing the world in quiet celebration. A lone robin greets the morning with a clear, cheerful trill, its song drifting across dew-laden grass. A gentle breeze stirs the leaves, carrying the fresh, earthy scent of new growth. Nearby, rooftops and empty streets lie poised between night's calm and the city's stirring pulse, promising simple comforts like a warm cup of coffee. In this tranquil pause, one senses life's renewal and the gentle invitation to greet the day with hope and gratitude.}

TEXT: { model_answer}

Return your answer in exactly this format: Coherency score: <integer 1-10>.

Response:

Figure 11: Prompt for Pairwise Coherence Evaluation for Reward Model r .

Rule of the Game of 24: Use all four numbers provided in the input, without repetition, and only basic arithmetic operations ($+, -, \times, \div$) to obtain 24. Only three steps are allowed.

Given the following two remaining numbers from a previous step in the Game of 24, the current step is: {step}. Evaluate this step.

Examine the numbers shown in each “left: . . .” after the step and reason whether it is still possible to reach 24: • **Sure** → 3 • **Likely** → 1 • **Impossible** → 0

Return the score in the following format: ****Answer**:** <integer score>

Response:

Figure 12: Prompt for single-step evaluation used in the reward model r for Game of 24.

B Additional Experimental Results

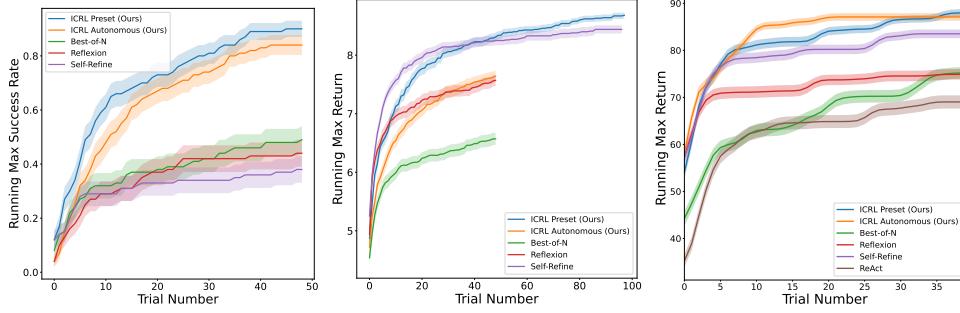


Figure 13: **Benchmark results: Mean of Running Max.** **(Left)** The mean of running max success rate on Game of 24. **(Middle)** The mean of running max coherence reward on creative writing. Both ICRL Preset and Self-Refine went through an additional run of 50 episodes. **(Right)**. The mean of running max return on ScienceWorld.

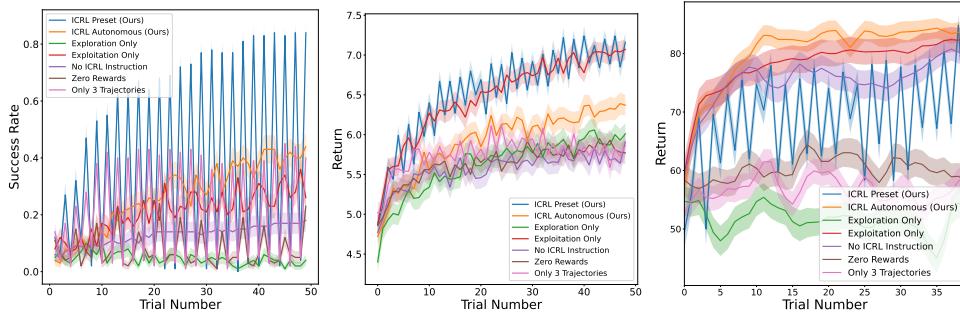


Figure 14: **Ablation study results: Original Curves.** **(Left)** The mean of success rate on Game of 24 ablation studies. **(Middle)** The mean of coherence reward on creative writing ablation studies. **(Right)**. The mean return on ScienceWorld ablation studies.

Table 3: Running max of return averaged over all the tasks in **ScienceWorld**.

Method	Return (max = 100)
ReAct	69 ± 1.4
Reflexion	74 ± 1.1
Best-of-N	75 ± 1.2
Self-Refine	83 ± 0.9
ICRL Preset (Ours)	88 ± 0.7
ICRL Autonomous (Ours)	87 ± 0.8

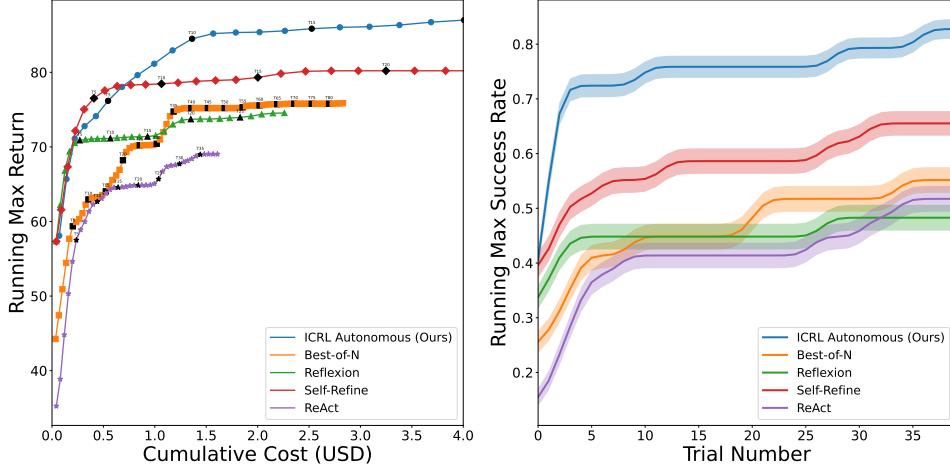


Figure 15: **Additional ScienceWorld Results.** **(Left)** Although ICRL’s context (comprising the experience buffer) is longer than that of random sampling methods, it still outperforms them and other experience-based approaches given additional compute budget. **(Right)** ICRL’s superior return improvement as seen in other results, also leads to a greater increase in success rate.