

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

TEACHING LANGUAGE MODEL TO ACT EFFICIENTLY

Anonymous authors

Paper under double-blind review

ABSTRACT

Tool-integrated reasoning (TIR) augments large language models (LLMs) with the ability to invoke external tools during long-form reasoning, such as search engines and code interpreters, to solve tasks beyond the capabilities of internal reasoning. While reinforcement learning (RL) has shown promise in training such agents, most of existing approaches typically optimize only for final correctness without considering the efficiency or necessity of external tool use. This often leads to excessive tool calling, leading to increased computational costs and additional latency, and may also shift reliance toward external tools rather than the model’s own reasoning – a phenomenon referred to as *cognitive offloading*. To this end, we propose Optimized Tool Call-controlled Policy Optimization (OTC-PO), a simple yet effective RL-based framework that encourages models to produce accurate answers with less tool calls. Our method introduces a tool-integrated reward that jointly considers answer correctness and corresponding tool use behavior of model to reach that answer. To validate the effectiveness, we introduce the additional metric of *tool productivity*, defined as the ratio between the number of correct answers and the total number of tool calls across all test cases. This metric reflects how efficiently and effectively tool usage contributes to successful task completion, with higher values indicating more productive external tool calls with the help of internal reasoning. We then instantiate this framework within both Proximal Policy Optimization (PPO) and Group Relative Preference Optimization (GRPO), resulting in OTC-PPO and OTC-GRPO. Experiments with Qwen-2.5 and Qwen-Math across multiple QA benchmarks show that our approach reduces tool calls by up to 68.3% and improves tool productivity by up to 215.4%, while maintaining comparable answer accuracy, especially for the larger models.

1 INTRODUCTION

Large language models (LLMs) have shown strong reasoning capabilities, particularly when trained with reinforcement learning (RL) using simple outcome rule-based rewards (Team, 2024; 2025). Building on this progress, recent work on tool-integrated reasoning (TIR) (Gou et al., 2024; Wang et al., 2024a) enables LLMs to interact with external tools such as search engines (Jin et al., 2025) and code interpreters (Li et al., 2025b) using similar reward design. By combining internal reasoning with access to various external tools, TIR substantially broadens the problem-solving scope of LLMs, allowing them to tackle tasks that exceed the limitations of parametric knowledge alone.

However, most existing RL approaches for TIR focus solely on final answer correctness, without explicitly accounting for the efficiency or necessity of tool use (Jin et al., 2025; Feng et al., 2025). As a result, models often rely excessively on external calls (Qian et al., 2025b). This behavior raises several concerns: (i) higher computational and infrastructure costs due to unnecessary tool invocations (Figure 2 left); (ii) longer inference latency due to lengthy observation returned by the tool (Figure 2 middle); and (iii) a tendency to rely on external tools instead of leveraging the model’s own reasoning, a phenomenon known as *cognitive offloading* (Risko & Gilbert, 2016). For example, recent baselines such as Search-R1 (Jin et al., 2025) frequently invoke tools even in cases where the question could be answered directly, missing the opportunity to leverage internal reasoning (§ 5.3).

Based on these observations, we argue that effective agentic LLMs should not only produce correct answers but also use tools judiciously. In practice, different tasks and models may require different levels of external assistance. For some questions, the best strategy involves no tool usage at all, while others may benefit from one or more targeted interactions. This motivates the need for training

054 methods that optimize both effectiveness (i.e., answer correctness) and efficiency (i.e., appropriate
 055 tool call). We aim to develop such model based on a foundational assumption: for *each question* and
 056 *each model*, there exists an optimal number of tool calls, defined as the minimal number required
 057 for the model to arrive at a correct answer (i.e., the minimum tool calls among correct trajectories
 058 practically). Crucially, this optimal number can vary depending on both the capabilities of models
 059 and the complexity of questions. Therefore, it is difficult to create supervised fine-tuning data or to
 060 design a single prompting strategy that generalizes across diverse models and tasks. In contrast, RL
 061 offers a promising alternative, as it allows models to adapt their tool-call behavior based on their
 062 own experience or interactions (Silver & Sutton, 2025).

063 We then propose Optimized Tool Calls controlled Policy Optimization (OTC-PO), a simple yet ef-
 064 fective RL-based method that enforces models to minimize the number of tool calls required to reach
 065 a correct solution. Specifically, we introduce a tool-integrated reward that modulates traditional out-
 066 come reward signals, such as correctness, with a scaling *coefficient* reflecting tool efficiency. This
 067 encourages the model to prefer trajectories that reach correct answers with fewer tool calls. There-
 068 fore, it shifts the optimization objective from correctness alone to tool productivity, defined as the
 069 ratio between task benefit (e.g., answer accuracy) and tool usage cost (e.g., number of tool calls).
 070 In addition, OTC-PO is lightweight and plug-and-play, requiring only minimal changes to standard
 071 RL pipelines, making it easy to adopt in existing systems. We conduct comprehensive experiments
 072 on two widely used tools: search and code based on several LLMs, Qwen-2.5-3B/7B-Base and
 073 Qwen2.5-Math-1.5B/7B-Base (Qwen et al., 2025), and derive models with a substantial tool calls
 074 reduction in trajectory while approximately maintaining accuracy with strong baselines. In sum-
 075 mary, the key contributions are as follows.

- 076 • We re-frame tool-integrated reasoning as jointly optimizing correctness and efficiency, and intro-
 077 duce tool productivity, the number of correct answers per tool call, as a metric to capture this
 078 trade-off. We additionally identify the cognitive offloading phenomenon in TIR.
- 079 • We propose a simple, faster, and generalizable OTC-PO algorithm to encourage the model to use
 080 fewer tool calls to solve the problem and therefore maximize tool productivity. We note that it is
 081 compatible with various RL algorithms and can be easily implemented.
- 082 • We implement OTC-PPO and OTC-GRPO as two typical methods without losing the adaptability
 083 and generalization based on PPO (Schulman et al., 2017) and GRPO (Shao et al., 2024) algo-
 084 rithms, and the experimental results on several benchmarks and baselines demonstrate significant
 085 reductions in tool call cost while preserving most of the accuracy in both in-domain and out-of-
 086 domain evaluation.

087 2 RELATED WORK

090 **Tool Utilization for LLMs.** Teaching LLMs to use tools enables them to interact with external
 091 environments while overcoming several inherent limitations such as restricted access to up-to-date
 092 or domain-specific knowledge and poor mathematical operation capabilities. There are three major
 093 methods which can achieve this goal: 1) prompting engineering (Qian et al., 2023; Wang et al.,
 094 2024b; Qian et al., 2024b; Yuan et al., 2024; Qian et al., 2024a; Wang et al., 2025a), which guides
 095 the model’s behavior through carefully designed input templates or few-shot examples without mod-
 096 ifying model weights; 2) supervised finetuning on tool-integrated reasoning datasets (Schick et al.,
 097 2023; Qin et al., 2023; Gou et al., 2024; Qian et al., 2024c; 2025b), where the model learns from
 098 annotated trajectories that demonstrate when and how to use tools in context; and 3) reinforcement
 099 learning (Jin et al., 2025; Li et al., 2025b; Feng et al., 2025; Qian et al., 2025a), which allows
 100 the model to directly learn tool-use strategies by interacting with an environment and optimizing
 101 long-term rewards, enabling more adaptive and goal-directed behaviors. Some previous methods
 102 combine SFT and RL with collected tool-use trajectories and preferences (Chen et al., 2025). While
 103 prompting and supervised fine-tuning have shown promising results, they rely heavily on expert-
 104 designed prompts and tool-integrated annotations, which limits their scalability and generalization.
 105 As a result, recent efforts have shifted toward reinforcement learning, which demonstrates improved
 106 performance and generalization through simple rule-based rewards.

107 **Tool-integrated Reasoning Efficiency.** Only a few studies from prompting engineering and su-
 108 pervised fine-tuning attention on tool-integrated reasoning efficiency issues in terms of the cost of

108 tool usages during the reasoning (Wang et al., 2025a; Shen et al., 2024; Qian et al., 2025b). In
 109 detail, (Wang et al., 2025a) first propose a prompting-based framework: Self Divide-and-Conquer
 110 (SelfDC) to leverage the self-aware confidence score of LLMs to decide whether or not need to
 111 call tools during reasoning, achieving better trade-off between effectiveness and efficiency in the
 112 context of RAG. Furthermore, several works follow this direction and explore more boarder appli-
 113 cations and larger tool spaces (Qian et al., 2025b; Shen et al., 2024; Li et al., 2025a). For example,
 114 SMART (Qian et al., 2025b) collect the well-designed dataset to finetune the model to only call tools
 115 when the knowledge is outside the inherent parametric knowledge of LLMs. Despite these advance-
 116 ments, most existing approaches still rely on complex prompt engineering or annotated datasets,
 117 which hinders their adaptability and scalability to new scenarios and benchmarks. In contrast, the
 118 efficiency of tool-integrated reasoning within reinforcement learning frameworks remains largely
 119 underexplored.

3 METHODOLOGY

120 In this section, we first provide a formal definition of task considering the both effectiveness and
 121 efficiency of tool-integrated reasoning, followed by general RL framework and then our proposed
 122 Optimized Tool Call-controlled Policy Optimization (OTC-PO).

3.1 TASK DEFINITION

123 Given a question q , and an environment \mathcal{E} that provides access to a set of tools $\mathcal{T} = \{t_0, t_1, \dots, t_n\}$,
 124 the language model \mathcal{M} can *optionally* interact with the environment by calling specific tools in \mathcal{T} ,
 125 obtaining the corresponding tool results from \mathcal{E} , and iteratively repeating this processing until the
 126 final answer is driven. Without losing generalization, the tool-integrated reasoning trajectory τ_k at
 127 step k is defined as follows:

$$\tau_k = (r_1, tc_1, o_1), (r_2, tc_2, o_2), \dots, (r_k, tc_k, o_k), \quad (1)$$

128 where r_i, tc_i, o_i denotes the reasoning, tool call and returned observation respectively. Importantly,
 129 we also account for reasoning steps that do not involve tool usage. Suppose the step p does not need
 130 to call tools, then the tc_p and o_p become empty string, the reasoning content r_p can either be merged
 131 with the subsequent reasoning step r_{p+1} to form the new r_{p+1}^* , or, if p is the last step¹, be used
 132 directly to derive the final answer. The objective of the task is to generate the correct answer \hat{a} with
 133 the minimal cost of the full trajectory τ as follows:

$$\arg \min_{\tau} \text{Cost}(\tau) \quad \text{subject to} \quad \mathcal{M}(q, \tau) = \hat{a}, \quad (2)$$

134 Here the cost is measured as the number of tool calls within the trajectory τ considering its simplicity
 135 and generalization. Thus the model is encouraged to not only generate correct answer but also
 136 minimize the cost. We emphasize that this revised objective highlights an overlooked dimension of
 137 TIR to balance correctness with the tool use behavior of models.

3.2 TOOL-INTEGRATED REINFORCEMENT LEARNING

138 Inspired by recent success to use RL for better reasoning in LLMs, several efforts try to extend RL
 139 to tool-integrated reasoning with the objective functions can be defined as follows:

$$\max_{\pi_{\theta}} \mathbb{E}_{q \sim \mathcal{D}, y \sim \pi_{\theta}(\cdot | q; \mathcal{E})} [r_{\phi}(q, y)] - \beta \mathbb{D}_{kl} [\pi_{\theta}(y | q; \mathcal{E}) \| \pi_{\text{ref}}(y | q; \mathcal{E})], \quad (3)$$

140 where $\pi_{\theta}, \pi_{\text{ref}}$ stand for the policy model and reference model respectively, r_{ϕ} is the reward function
 141 and \mathbb{D}_{kl} is the KL-divergence measure. q is the question drawn from the dataset \mathcal{D} and y is the
 142 generated output, consisting of the tool-integrated reasoning trajectory τ and the final answer a .
 143 To optimize this goal, there are two well-established policy-gradient RL methods: Proximal Policy
 144 Optimization (PPO) and Group Relative Policy Optimization (GRPO) in TIR. We provide detailed
 145 optimization objective in Appendix C to save space.

146¹The last step is practically determined by the predefined maximal tool calls.

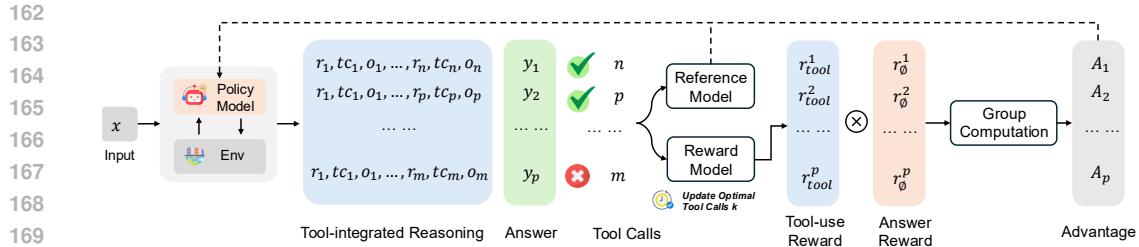


Figure 1: An overview of OTC-GRPO Algorithm.

Following the success of DeepSeek-R1 (Team, 2025), most prior predominantly focuses on rule-based outcome correctness rewards (Jin et al., 2025; Li et al., 2025b; Feng et al., 2025) to learn the policy model during training, as illustrated below:

$$r_\phi(q, y) = r_{correct} = 1 \text{ if } a = \hat{a} \text{ else } 0, \quad (4)$$

where a is the extracted predicted answer from the response y and \hat{a} is the ground truth answer.

3.3 OTC-PO: OPTIMIZED TOOL CALL-CONTROLLED POLICY OPTIMIZATION

In order to consider the tool call behavior, we propose Optimized Tool Call-controlled Policy Optimization (OTC-PO), which can be easily integrated with existing RL algorithms, aiming to preserve general accuracy while significantly reducing the cost of tool interactions. At the heart of OTC-PO is a reward mechanism that ranks the current number of tool calls m relative to the optimal number of tool calls n given the question and model. In most realistic settings, the true optimal number n is unknown in advance. To address this, OTC-PO approximates n by tracking the minimal number of tool calls observed across different trajectories to arrive the correct answer for the same question and model. This approach enables the model to learn efficient tool use without requiring prior knowledge of the optimal tool budget. Furthermore, the framework can be naturally extended to scenarios where the optimal tool call number is known a priori (e.g., $n = 0$ for language-only reasoning). We instantiate OTC-PO in two variants: OTC-PPO and OTC-GRPO, tailored to different underlying RL algorithms.

OTC-PPO. Since there is no multiple trajectories for same question q in single optimization step, we design the tool reward according to the number of tool calls m in the current trajectory as follows:

$$r_{tool} = \cos\left(\frac{m * \pi}{2m + c}\right) \quad (5)$$

The core idea is to map m to $[0, \pi/2]$ and c is the smooth constant that controls the reward decay rate. The smaller the value, the faster the punishment and the more it encourages less use of tools; the larger the value, the more tolerant it is. In this way, among trajectories, those requiring more tool calls will receive lower rewards than those achieved with fewer tool calls. We note that any function exhibiting such monotonicity is applicable, but we find sin/cos functions is more smooth and easy to train, as also evidence by Yeo et al. (2025). Although we can track the m here for approximation of n in later steps, we empirically find it requires much longer training steps due to poor sampling efficiency in PPO, which is undesirable.

OTC-GRPO. We first identify the trajectories $\{\tau^1, \tau^2, \dots, \tau^p\}$ leading to correct answer from the group for the question q , and get the corresponding number of tool calls for each trajectory: $\mathcal{C} = \{c^1, c^2, \dots, c^p\}$, and we can calculate the minimal tool calls $k = \min(\mathcal{C})$, serving as the *local* approximation of optimal tool calls for q . Furthermore, we can update k during multiple epochs to approximate the *global* optimal tool calls if the policy model finds the better solution with less than k calls in later iterations. We use n to indicate the approximation of optimal tool calls, and m the number of tool calls in the current trajectory for q . Therefore we design the reward as follows:

$$\begin{aligned}
r_{tool} &= \begin{cases} 1 & \text{if } f(m, n) = n = 0 \\ \cos\left(\frac{m*\pi}{2m+c}\right) & \text{if } n = 0 \\ \sin\left(\frac{f(m, n)*\pi}{2n}\right) & \text{otherwise} \end{cases} \\
f(m, n) &= \begin{cases} 0, & \text{if } m = 0 \text{ and } n = 0 \\ m, & \text{if } n = 0 \\ \frac{2nm}{m+n}, & \text{otherwise} \end{cases}
\end{aligned} \tag{6}$$

where $f(m, n)$ is the mapping function to re-map the m to the range from 0 to $2n$. Briefly, the key here is to assign the highest reward 1 (a.k.a, $\sin(\pi/2)$) when the policy model achieves optimal tool calls n , and when actual number of tool calls m deviates from n , either exceeding or falling short, the model receives a comparatively lower reward. Thus, we can assign different rewards dynamically according to the different n and m for the question q . We provide reward illustration to better motivate our reward design in Appendix E.1.

Tool-integrated Reward Design. Inspired by recent study (Arora & Zanette, 2025), we regard the tool-integrated reward r_{tool} as a *coefficient* of conventional reward function $r_\phi(q, y)$, leading to the final tool-integrated reward function:

$$r_\phi^{tool}(q, y) = \alpha * r_{tool} * r_\phi(q, y) \tag{7}$$

where α is a hyperparameter that controls the scaling of the tool reward as r_ϕ is usually same for correct or wrong answer. Notably, this multiplicative structure ensures that tool efficiency is mainly rewarded when the primary task objective (e.g., answer correctness) is satisfied. For example, when the final answer is incorrect, r_ϕ is 0, effectively disabling the influence of r_{tool} . This design helps mitigate reward hacking by preventing the policy model from being incentivized to use tools without ultimately producing a correct answer. If the answer is correct, r_ϕ is 1 and then r_{tool} begins to dominate the scale of the reward, thus encouraging less tool usage.

It is worth noting that this design offers several practical advantages: i) It is consistent with the task objective, which prioritizes achieving correct solutions while reducing the number of tool calls, with theoretical support provided in Appendix E.2. ii) It helps reduce the risk of reward hacking compared to simple additive reward formulations (e.g., $r_{tool} + r_\phi$), which in our experiments led to unstable behaviors. For instance, we observed cases where the model could achieve higher reward by minimizing tool usage without necessarily producing correct answers. iii) It remains flexible and can be applied to different reward formulations of r_ϕ , such as $r_\phi = r_{correct}$ or $r_\phi = r_{correct} + r_{format}$, depending on the task requirements, as discussed in Appendix E.3.

4 EXPERIMENTS

4.1 SET UP

Datasets and Baselines. We mainly follow the Search-R1 (Jin et al., 2025) and ToRL (Li et al., 2025b) setting and use same baselines and datasets for the fair and comprehensive evaluation. Specifically, we use NQ and HotpotQA as training dataset for search, and we use the dataset provided in ToRL for code. We also directly compare our methods with several baselines such as SFT, Base-RL, retrieval-augmented generation baselines (i.e., RAG, IRCoT) and more importantly the Search-R1 and ToRL themselves.

Evaluation Metrics. Besides the exact match (EM) and the average tool calls (TC), we additionally define a new metric, *tool productivity* (TP), which measures the effectiveness and efficiency of tool calls during inference. Specifically, TP is defined as the number of correctly answered questions per unit of tool call: $TP = \frac{\sum_{i=1}^N \mathbb{I}\{y_i = \hat{y}_i\}}{\sum_{i=1}^N tc_i}$ ², where \mathbb{I} is the indicator function that equals 1 if the predicted answer \hat{y}_i matches the ground truth y_i , and tc_i denotes the number of tool calls used in the i^{th} instance. This metric reflects how efficiently the model converts tool usage into correct answers, capturing both utility and cost in a single measure. While EM provides a standard measure of accuracy, it does not reflect the underlying reasoning cost. Therefore, we consider TP as more

²This can also be understood as the fraction between benefits and cost.

270 Table 1: The results of OTC-PO with different baselines in search. The results except Search-R1
 271 are directly copied from original paper (Jin et al., 2025). We highlight the relative improvement of
 272 OTC-PO and OTC-GRPO compared with the corresponding variants of Search-R1.
 273

| Models | NQ | | | HotpotQA | | |
|--------------------------|--------|-----------------|------------------|----------|-----------------|------------------|
| | EM (↑) | TC (↓) | TP (↑) | EM (↑) | TC (↓) | TP (↑) |
| Qwen2.5-3B(-Base) | | | | | | |
| R1-Base | 0.226 | - | - | 0.201 | - | - |
| SFT | 0.249 | - | - | 0.186 | - | - |
| RAG | 0.348 | 1.0 | 0.348 | 0.255 | 1.0 | 0.255 |
| IRCoT | 0.111 | 10.0 | 0.011 | 0.164 | 10.0 | 0.016 |
| Search-R1-PPO | 0.403 | 1.738 | 0.232 | 0.279 | 1.716 | 0.163 |
| Search-R1-GRPO | 0.404 | 1.426 | 0.283 | 0.312 | 1.802 | 0.173 |
| OTC-PPO | 0.355 | 1.010 (▼ 41.9%) | 0.351 (▲ 51.3%) | 0.260 | 1.026 (▼ 40.2%) | 0.253 (▲ 55.2%) |
| OTC-GRPO | 0.444 | 1.008 (▼ 29.3%) | 0.440 (▲ 55.5%) | 0.365 | 1.387 (▼ 23.0%) | 0.263 (▲ 52.0%) |
| Qwen2.5-7B(-Base) | | | | | | |
| R1-Base | 0.270 | - | - | 0.242 | - | - |
| SFT | 0.318 | - | - | 0.217 | - | - |
| RAG | 0.349 | 1.0 | 0.349 | 0.299 | 1.0 | 0.299 |
| IRCoT | 0.224 | 9.999 | 0.022 | 0.133 | 9.982 | 0.013 |
| Search-R1-PPO | 0.449 | 3.282 | 0.136 | 0.380 | 3.741 | 0.102 |
| Search-R1-GRPO | 0.399 | 1.697 | 0.235 | 0.341 | 2.109 | 0.162 |
| OTC-PPO | 0.446 | 1.040 (▼ 68.3%) | 0.429 (▲ 215.4%) | 0.383 | 1.464 (▼ 60.9%) | 0.262 (▲ 156.9%) |
| OTC-GRPO | 0.444 | 0.990 (▼ 41.7%) | 0.448 (▲ 90.6%) | 0.366 | 1.005 (▼ 52.3%) | 0.364 (▲ 124.7%) |

288 Table 2: The results of OTC-PO with different baselines in ToRL (Li et al., 2025b).
 289

| Models | AIME24 | | | AIME25 | | |
|---------------------------------|--------|---------------|-----------------|--------|---------------|-----------------|
| | EM (↑) | TC (↓) | TP (↑) | EM (↑) | TC (↓) | TP (↑) |
| Qwen2.5-Math-1.5B(-Base) | | | | | | |
| Qwen2.5-Math-1.5B-Ins | 10.0 | - | - | 10.0 | - | - |
| Qwen2.5-Math-1.5B-Ins-TIR | 13.3 | 1.1 | 12.1 | 13.3 | 1.4 | 9.5 |
| ToRL-GRPO | 23.3 | 2.2 | 10.6 | 23.3 | 2.3 | 10.1 |
| OTC-GRPO | 20.0 | 1.1 (▼ 50.0%) | 18.2 (▲ 71.7%) | 20.0 | 1.1 (▼ 41.2%) | 18.2 (▲ 80.2%) |
| Qwen2.5-Math-7B(-Base) | | | | | | |
| Qwen2.5-Math-7B-Ins | 10.0 | - | - | 16.7 | - | - |
| Qwen2.5-Math-7B-Ins-TIR | 26.7 | 1.6 | 16.4 | 16.7 | 1.4 | 12.2 |
| Base-RL | 33.3 | - | - | 6.7 | - | - |
| ToRL-GRPO | 36.7 | 2.1 | 17.5 | 26.7 | 2.1 | 12.7 |
| OTC-GRPO | 36.7 | 0.7 (▼ 66.7%) | 32.4 (▲ 199.4%) | 23.3 | 0.8 (▼ 61.9%) | 29.1 (▲ 129.1%) |

300 informative indicators of agentic reasoning efficiency – highlighting not only whether the model can
 301 provide correct answer, but how economically it arrives at that correctness.
 302

303 **Implementation Details.** We re-use the same parameter in Search-R1 (Jin et al., 2025) and
 304 ToRL (Li et al., 2025b) respectively. There are only minor modification we make to suit our method:
 305 i) We slightly change the template in Search-R1 to tell the model that it only need to call tools when
 306 necessary (Appendix D), and we do not change the template in ToRL; ii) We set the maximal num-
 307 ber of tool calls \mathcal{C} in ToRL to 3 to better study the effects of our methods when multiple tool calls
 308 are allowed and keep it as 4 as in original Search-R1. Moreover, we set c as corresponding max
 309 turns or maximal tool limits, α as 1. We conduct our experiments on 8 A100-80G GPU and re-
 310 produce the results of Search-R1 and ToRL independently. We implement OTC-GRPO using the
 311 global approximation of minimal tool calls since it leads to more stable and better optimization.
 312

313 4.2 MAIN RESULTS

314 **Search as Tool.** Table 1 shows the results for search-required tasks. There are several key insights
 315 can be drawn: i) OTC-PPO achieves significant improvement in terms of TC and TP compared
 316 with Search-R1-PPO, resulting in up to a 68.3% reduction in TC and a 215.4% increase in TP, and
 317 OTC-GRPO can further reduce the absolute tool calls due to more accurate approximation; ii) It
 318 is found that our method will not sacrifice the accuracy too much especially for larger LLMs, as
 319 evidenced by our OTC-PO achieves comparable EM score with Search-R1 on Qwen2.5-7B model;
 320 iii) Different models have different tool use behaviors on different datasets. If we look at the Search-
 321 R1, we can find that 3B model tends to use less tool calls compared with 7B model. This reveals a
 322 critical issue: as language models scale up, they may tend to over-rely on external tool calls unless
 323 explicitly penalized for such behavior. This not only exacerbates the problem with larger models
 but also leads to an underutilization of their inherent reasoning capabilities (See § 5.3). The number

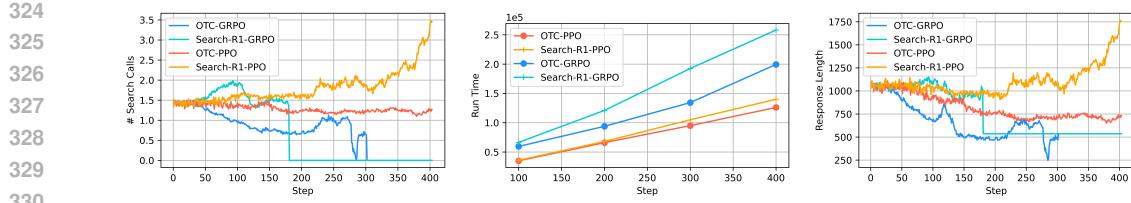


Figure 2: **Left:** Changes of number of search calls during the training; **Middle:** Running time analysis; and **Right:** Changes of response length during the training.

Table 3: Tool behavioral advantage analysis of OTC-PO against Search-R1. ME, LE, MA, LA, AE stands for more efficient, less efficient, more accurate, less accurate, and accurate and efficient, respectively. We found that OTC-GRPO and Search-R1-PPO achieves same results with same tool calls in 42.54% cases of HotpotQA on Qwen2.5-3B-Base model, leading to relatively lower ME.

| Method | NQ | | | | | HotpotQA | | | | |
|-------------------------|--------|--------|--------|--------|--------|----------|--------|--------|--------|--------|
| | ME (↑) | LE (↓) | MA (↑) | LA (↓) | AE (↑) | ME (↑) | LE (↓) | MA (↑) | LA (↓) | AE (↑) |
| Qwen2.5-3B-Base) | | | | | | | | | | |
| OTC-PPO | 63.55 | 0.02 | 4.60 | 7.67 | 3.74 | 61.78 | 0.53 | 4.30 | 6.34 | 3.35 |
| OTC-GRPO | 65.43 | 2.49 | 6.15 | 3.85 | 5.10 | 35.54 | 5.96 | 11.89 | 4.07 | 3.47 |
| Qwen2.5-7B-Base) | | | | | | | | | | |
| OTC-PPO | 86.2 | 0.03 | 6.45 | 6.92 | 6.45 | 81.49 | 0.08 | 8.60 | 8.68 | 8.53 |
| OTC-GRPO | 82.0 | 0.00 | 7.26 | 10.69 | 7.26 | 79.71 | 0.00 | 8.15 | 11.98 | 8.12 |

of TC is also various across the datasets, which relates to many factors such as inherent capabilities of different models (i.e., self-awareness) and complexity of different datasets, leading to varying minimal number of tool calls for each question and model; iv) As the model size increases, the TC and TP get bigger boost no matter in OTC-PPO or OTC-GRPO. We highlight the ii) and iv) are very important since they compose the great potential scalability for our methods as model scales.

Code as Tool. We mainly follow the same setting and only report the GRPO results to be consistent with original ToRL (Li et al., 2025b) paper. Table 2 shows the results. We observe several similar findings: i) our method does not sacrifice the accuracy too much and even brings some improvement when the model gets larger. We attribute this to the development of internal reasoning capability of models when it is enforced to minimize external tool calls, as evidenced in our case study; ii) our method can significantly boost the tool productivity, reducing the unnecessary tool calls; iii) As model size increases, our method is more effective to improve the tool productivity.

5 ANALYSIS

We mainly conduct our analysis using search as a tool in this section and leave more analysis in the Appendix F and G, respectively.

5.1 TOOL USE BEHAVIOR ANALYSIS

Training. Fig. 2 shows the training behaviors of different methods. It is clear that our method not only achieves comparable results with fewer tool calls and shorter responses, but also enables faster and more efficient training optimization. This is particularly important, as it significantly reduces the time and cost associated with real-time tool interactions during training, both in terms of financial expenses and computational overhead. We also find GRPO tends to achieve better performance than PPO due to more accurate approximation of minimal tool call required, but it is less stable than PPO, as also evidenced by several recent studies (Jin et al., 2025). It is encouraging to find that our method can delay the early collapse substantially with more training steps.

Inference. We identify several representative tool-use behavior types of our proposed method: i) the answer is same with *less* tool calls compared with baseline (*more efficient*); ii) the answer is same with *more* tool calls compared with baseline (*less efficient*); and iii) our method is able to produce the correct answer whereas the baseline fails to answer correctly (*more accurate*); iv) the baseline

378 Table 4: The results of Out-of-Domain (OOD) evaluation of OTC against Search-R1 in EM and TC.
379

| 380 381 Models | 382 TriviaQA | | 383 PopQA | | 384 2Wiki | | 385 Musique | | 386 Bamboogle | |
|--------------------------|-----------------|---------------|---------------|---------------|---------------|---------------|----------------|---------------|------------------|---------------|
| | 387 EM (↑) | 388 TC (↓) | 389 EM (↑) | 390 TC (↓) | 391 EM (↑) | 392 TC (↓) | 393 EM (↑) | 394 TC (↓) | 395 EM (↑) | 396 TC (↓) |
| Qwen2.5-3B(-Base) | | | | | | | | | | |
| Search-R1-PPO | 0.566 | 1.580 | 0.425 | 1.631 | 0.258 | 1.675 | 0.051 | 1.922 | 0.063 | 1.766 |
| Search-R1-GRPO | 0.587 | 1.455 | 0.345 | 1.542 | 0.257 | 1.991 | 0.084 | 2.263 | 0.203 | 1.859 |
| OTC-PPO | 0.551 | 1.008 | 0.409 | 1.009 | 0.235 | 1.050 | 0.045 | 1.051 | 0.063 | 1.016 |
| OTC-GRPO | 0.608 | 1.046 | 0.441 | 1.030 | 0.341 | 1.561 | 0.124 | 1.734 | 0.266 | 1.547 |
| Qwen2.5-7B(-Base) | | | | | | | | | | |
| Search-R1-PPO | 0.596 | 3.353 | 0.420 | 3.315 | 0.326 | 4.116 | 0.135 | 4.294 | 0.375 | 3.641 |
| Search-R1-GRPO | 0.578 | 1.704 | 0.411 | 1.754 | 0.340 | 2.521 | 0.130 | 2.616 | 0.203 | 1.859 |
| OTC-PPO | 0.623 | 1.066 | 0.425 | 1.083 | 0.363 | 1.868 | 0.152 | 1.942 | 0.391 | 1.828 |
| OTC-GRPO | 0.597 | 0.430 | 0.431 | 0.739 | 0.311 | 0.938 | 0.130 | 1.224 | 0.250 | 0.781 |

389 Table 5: The results of OTC-PO under Qwen2.5-7B-Base and Qwen2.5-7B-Instruct models.
390

| 391 Models | 392 NQ | | | 393 HotpotQA | | |
|------------------------------|---------------|---------------|---------------|-----------------|---------------|---------------|
| | 394 EM (↑) | 395 TC (↓) | 396 TP (↑) | 397 EM (↑) | 398 TC (↓) | 399 TP (↑) |
| Qwen2.5-7B(-Base) | | | | | | |
| OTC-PPO | 0.446 | 1.040 | 0.429 | 0.383 | 1.464 | 0.262 |
| OTC-GRPO | 0.444 | 0.990 | 0.448 | 0.366 | 1.005 | 0.364 |
| Qwen2.5-7B(-Instruct) | | | | | | |
| OTC-PPO | 0.389 | 1.404 | 0.277 | 0.381 | 1.880 | 0.203 |
| OTC-GRPO | 0.429 | 1.322 | 0.325 | 0.386 | 1.956 | 0.197 |

398 can produce the correct answer whereas the our method fails to answer correctly (*less accurate*); and
399 v) our method is able to produce the correct answer using fewer tool calls whereas the baseline fails
400 to answer correctly (*accurate and efficient*). Table 3 shows the results compared with the baseline
401 Search-R1-PPO³. On the efficiency side, it is clear to find that our method achieve same answers
402 with baseline using fewer tool calls in most cases, and more than 80% when the model becomes
403 larger. On the effectiveness side, we conclude that there is no significant loss when using same RL
404 algorithm, as evidenced by lower gap between MA and LA in OTC-PPO method. Moreover, among
405 all MA cases, it is observed that AE accounts for roughly 80% in 3B model (i.e., 3.74/4.60 =81%),
406 and exceeds 95% on 7B model. These results highlight the great strength of OTC-PO in promoting
407 efficiency while reserving effectiveness.

408 409 5.2 OUT-OF-DOMAIN EVALUATION

410 We run out-of-domain evaluation on more QA benchmarks. Fig. 6 shows the results of TP and Ta-
411 ble 4 shows the results of EM and TC. Generally, we find that our method achieves highest TP across
412 all the model and benchmark. More importantly, we observe that OTC-GRPO is better than OTC-
413 PPO on TP and as the model size increases, our method showcases bigger improvement, serving as
414 another evidence of great scalability of our method. Notably, we find that OTC-PPO outperforms
415 the Search-R1-PPO in terms of EM on all benchmarks with Qwen2.5-7B-Base model with much
416 less tool calls, demonstrating the better generalization capabilities of our methods.

417 418 5.3 MORE ANALYSIS

419 **Base v.s. Instruct Model.** We additionally compare the performance of OTC-PO using Qwen-
420 7B-Base and Qwen-7B-Instruct models. Table 5 shows the final results. It is observed that i) Base
421 models mostly lead to better performance, which is also consistent with results in the Search-R1
422 paper; ii) Instruct models tends to use more tools compared with base model while the TC still is
423 much lower than the Search-R1. We attribute this to the Instruct model tends to memorize specific
424 pattern during supervised fine-tuning and easily to be triggered to use tools even unnecessary.

425 **Case Study** Fig. 3 presents a case study to better illustrate the advantages of our proposed OTC-
426 PO. The analysis highlights several observations: (i) Search-R1 tends to issue multiple redundant
427 search calls, reflecting the absence of an explicit incentive to limit tool usage; (ii) both OTC-PPO
428 and OTC-GRPO substantially reduce unnecessary tool calls, but they do so in different ways. OTC-
429 PPO improves efficiency by generating more informative queries, while OTC-GRPO relies more on
430

³We use the same baseline to better distinguish OTC-PPO and OTC-GRPO.

432

433

Question: Are Giuseppe Verdi and Ambroise Thomas both Opera composers ?
 Answer: Yes

434

435

Search-R1-PPO: <think> I need to determine if Giuseppe Verdi and Ambroise Thomas are both opera composers. </think>
 <search> Giuseppe Verdi </search> *Undesirable tool use behavior*
 <observation> ... </observation>
 <search> Ambroise Thomas </search>
 <observation> ... </observation>
 <search> Are Giuseppe Verdi and Ambroise Thomas both Opera composers? </search>
 <observation> ... </observation>
 <think> Based on the observations, both Giuseppe Verdi and Ambroise Thomas are described as composers, specifically opera composers. </think>
 <answer> yes </answer> **(Tool Calls = 3)**

Cognitive Offloading fosters dependency on external tools, causes memory capabilities to diminish, and inhibits the development of internal reasoning..

OTC-PPO: <think> First, I need to determine if Giuseppe Verdi and Ambroise Thomas are both opera composers. </think>
 <search> Giuseppe Verdi and Ambroise Thomas both Opera composers </search>
 <observation> ... </observation>
 <think> Based on the search results, Giuseppe Verdi and Ambroise Thomas are both known as opera composers.
 </think>
 <answer> yes </answer> **(Tool Calls = 1)**

Search-R1-GRPO: <think> To answer this question, I need to check if both Giuseppe Verdi and Ambroise Thomas are composers of opera. I don't have immediate access to this knowledge, but I can conduct a search to find out. </think>
 <search> who are Giuseppe Verdi and Ambroise Thomas and their works </search>
 <observation>...</observation>
 <think>From the information gathered, ... This confirms Thomas as an opera composer. Next, I will search on Giuseppe Verdi. </think>
 <search> who is Giuseppe Verdi </search> *Undesirable tool use behavior*
 <observation>...</observation>
 <think>Giuseppe Verdi is named as an Italian opera composer of prominent stature. This confirms Verdi as an opera composer. Answering the initial question: Question: Are Giuseppe Verdi and Ambroise Thomas both Opera composers? </think>
 <answer> yes </answer> **(Tool Calls = 2)**

OTC-GRPO: To answer the question of whether Giuseppe Verdi and Ambroise Thomas are both Opera composers, let's first reason through the information:
 1. Giuseppe Verdi is a well-known Italian opera composer. He is famous for his numerous operas, including "Rigoletto," "La Traviata," and "Aida."
 2. Ambroise Thomas is a French opera composer. He is known for his operas such as "Mignon" and "Hamlet."
 Based on this reasoning, both Giuseppe Verdi and Ambroise Thomas are indeed Opera composers.
 <answer> yes </answer> **(Tool Calls = 0)**

451

452

Figure 3: A case study using *search* as a tool on Qwen-7B-Base model, where all four methods arrive at the correct answer but exhibit distinct behaviors in internal reasoning and external tool usage. OTC-GRPO does not rely on external tools while Search-R1-PPO uses extensive tool calls.

453

454

internal reasoning when the problem can be solved without external tools. Overall, the case study suggests that without guidance, models may overuse external tools, whereas encouraging minimal tool calls while preserving correctness promotes more balanced and deliberate tool use.

455

456

Reasoning v.s. Acting. We further examine the relationship between reasoning and tool use when models are encouraged to reduce external calls. Specifically, we compare OTC-PPO with Search-R1-PPO by recording, for each correctly answered sample, both the number of tool calls and the number of reasoning tokens (excluding search and observation tokens in response). As shown in Figure 4, OTC-PPO generally makes fewer tool calls while producing longer reasoning traces. Moreover, this difference becomes more pronounced as model size increases, suggesting that larger models are better able to leverage internal reasoning for external tool use under our framework⁴.

471

472

6 CONCLUSION AND FUTURE WORK

473

474

In this work, we redefine the objective of task for agentic RL not only provide the final correct answer, but also optimize the tool use behavior of model to achieve such goal. We then introduce OTC-PO, a simple yet effective RL framework that explicitly encourages LLMs to generate correct answers with fewer tool calls. Unlike prior work that primarily focuses on final answer correctness, our approach incorporates a tool-integrated reward that accounts for both effectiveness and efficiency of tool usage, thereby promoting tool productivity without sacrificing the accuracy a lot in several benchmarks with different tools. We also find that extensively rely on external tools hinder the development and utilization of internal reasoning capabilities of the model, and minimizing external tool calls alternatively foster the development of internal reasoning capabilities. Finally, in future work, we aim to extend our framework to more complex agentic tasks involving a broader set of tools and longer-horizon reasoning.

485

⁴More case studies, hyper-parameter analysis can be found in Appendix.

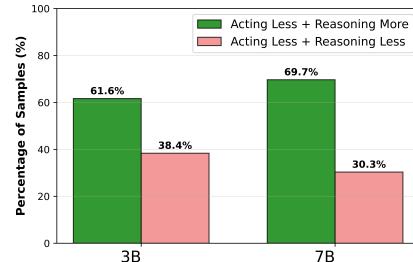


Figure 4: Comparison of reasoning (number of reasoning tokens) and acting (number of tool calls) between Search-R1-PPO and OTC-PPO on the NQ.

486 REPRODUCIBILITY STATEMENT
487

488 We strictly follow the same setting of two popular open-sourced codebases: Search-R1 (<https://github.com/PeterGriffinJin/Search-R1>) and ToRL (<https://github.com/GAIR-NLP/ToRL>), and use same hyper-parameters and data to conduct our main experiments. We
489 attach our reward design code in the supplementary for reproduction. The details for reproducing
490 results are also mentioned in Sections 4.1 of the main paper.

494 REFERENCES
495

496 Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with
497 reinforcement learning, 2025. URL <https://arxiv.org/abs/2503.04697>.

498 Daman Arora and Andrea Zanette. Training language models to reason efficiently, 2025. URL
499 <https://arxiv.org/abs/2502.04463>.

500 Sijia Chen, Yibo Wang, Yi-Feng Wu, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and
501 Lijun Zhang. Advancing tool-augmented large language models: Integrating insights from errors
502 in inference trees, 2025. URL <https://arxiv.org/abs/2406.07115>.

503 Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang,
504 Jinxin Chi, and Wanjun Zhong. Retool: Reinforcement learning for strategic tool use in llms,
505 2025. URL <https://arxiv.org/abs/2504.11536>.

506 Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujiu Yang, Minlie Huang, Nan Duan,
507 and Weizhu Chen. ToRA: A tool-integrated reasoning agent for mathematical problem solv-
508 ing. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Ep0TtjVoap>.

509 Bairu Hou, Yang Zhang, Jiabao Ji, Yujian Liu, Kaizhi Qian, Jacob Andreas, and Shiyu Chang.
510 Thinkprune: Pruning long chain-of-thought of llms via reinforcement learning, 2025. URL
511 <https://arxiv.org/abs/2504.01296>.

512 Bowen Jin, Hansi Zeng, Zhenrui Yue, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1:
513 Training llms to reason and leverage search engines with reinforcement learning, 2025. URL
514 <https://arxiv.org/abs/2503.09516>.

515 Adam Daniel Laud. *Theory and application of reward shaping in reinforcement learning*. University
516 of Illinois at Urbana-Champaign, 2004.

517 Wenjun Li, Dexun Li, Kuicai Dong, Cong Zhang, Hao Zhang, Weiwen Liu, Yasheng Wang, Ruiming
518 Tang, and Yong Liu. Adaptive tool use in large language models with meta-cognition trigger,
519 2025a. URL <https://arxiv.org/abs/2502.12961>.

520 Xuefeng Li, Haoyang Zou, and Pengfei Liu. Torl: Scaling tool-integrated rl, 2025b. URL <https://arxiv.org/abs/2503.23383>.

521 Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong
522 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kel-
523 ton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike,
524 and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.
525 URL <https://arxiv.org/abs/2203.02155>.

526 Cheng Qian, Chi Han, Yi Fung, Yujia Qin, Zhiyuan Liu, and Heng Ji. CREATOR: Tool creation for
527 disentangling abstract and concrete reasoning of large language models. In Houda Bouamor, Juan
528 Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP*
529 2023, pp. 6922–6939, Singapore, December 2023. Association for Computational Linguistics.
530 doi: 10.18653/v1/2023.findings-emnlp.462. URL <https://aclanthology.org/2023.findings-emnlp.462/>.

531 Cheng Qian, Peixuan Han, Qinyu Luo, Bingxiang He, Xiusi Chen, Yuji Zhang, Hongyi Du, Jiarui
532 Yao, Xiaocheng Yang, Denghui Zhang, et al. Escapebench: Pushing language models to think
533 outside the box. *arXiv preprint arXiv:2412.13549*, 2024a.

540 Cheng Qian, Shihao Liang, Yujia Qin, Yining Ye, Xin Cong, Yankai Lin, Yesai Wu, Zhiyuan Liu,
 541 and Maosong Sun. Investigate-consolidate-exploit: A general strategy for inter-task agent self-
 542 evolution. *arXiv preprint arXiv:2401.13996*, 2024b.

543

544 Cheng Qian, Chenyan Xiong, Zhenghao Liu, and Zhiyuan Liu. Toolink: Linking toolkit creation and
 545 using through chain-of-solving on open-source model. In *Proceedings of the 2024 Conference of*
 546 *the North American Chapter of the Association for Computational Linguistics: Human Language*
 547 *Technologies (Volume 1: Long Papers)*, pp. 831–854, 2024c.

548 Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan
 549 Tur, and Heng Ji. Toolrl: Reward is all tool learning needs. *arXiv preprint*, 2025a.

550

551 Cheng Qian, Emre Can Acikgoz, Hongru Wang, Xiusi Chen, Avirup Sil, Dilek Hakkani-Tür,
 552 Gokhan Tur, and Heng Ji. Smart: Self-aware agent for tool overuse mitigation. *arXiv preprint*
 553 *arXiv:2502.11435*, 2025b.

554 Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru
 555 Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein,
 556 Dahai Li, Zhiyuan Liu, and Maosong Sun. Toolllm: Facilitating large language models to master
 557 16000+ real-world apis, 2023. URL <https://arxiv.org/abs/2307.16789>.

558

559 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 560 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 561 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 562 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li,
 563 Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,
 564 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025.
 565 URL <https://arxiv.org/abs/2412.15115>.

566

567 Evan F Risko and Sam J Gilbert. Cognitive offloading. *Trends in cognitive sciences*, 20(9):676–688,
 568 2016.

569

570 Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambrø,
 571 Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can
 572 teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–
 573 68551, 2023.

574

575 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 576 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

577

578 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 579 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 580 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

581

582 Yuanhao Shen, Xiaodan Zhu, and Lei Chen. SMARTCAL: An approach to self-aware tool-
 583 use evaluation and calibration. In Franck Dernoncourt, Daniel Preoțiuc-Pietro, and Anastasia
 584 Shimorina (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural*
 585 *Language Processing: Industry Track*, pp. 774–789, Miami, Florida, US, November 2024.
 586 Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-industry.59. URL
 587 <https://aclanthology.org/2024.emnlp-industry.59/>.

588

589 David Silver and Richard S Sutton. Welcome to the era of experience. *Google AI*, 1, 2025.

590

591 DeepSeek-AI Team. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement
 592 learning, 2025. URL <https://arxiv.org/abs/2501.12948>.

593

594 OpenAI Team. Openai ol system card, 2024. URL <https://arxiv.org/abs/2412.16720>.

595

596 Hongru Wang, Huimin Wang, Zezhong Wang, and Kam-Fai Wong. Integrating pretrained language
 597 model for dialogue policy evaluation. In *ICASSP 2022 - 2022 IEEE International Conference*
 598 *on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6692–6696, 2022. doi: 10.1109/
 599 ICASSP43922.2022.9747593.

594 Hongru Wang, Yujia Qin, Yankai Lin, Jeff Z. Pan, and Kam-Fai Wong. Empowering large lan-
 595 guage models: Tool learning for real-world interaction. In *Proceedings of the 47th Interna-*
 596 *tional ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR*
 597 *'24*, pp. 2983–2986, New York, NY, USA, 2024a. Association for Computing Machinery. ISBN
 598 9798400704314. doi: 10.1145/3626772.3661381. URL <https://doi.org/10.1145/3626772.3661381>.

600 Hongru Wang, Rui Wang, Boyang Xue, Heming Xia, Jingtao Cao, Zeming Liu, Jeff Z. Pan,
 601 and Kam-Fai Wong. AppBench: Planning of multiple APIs from various APPs for com-
 602 plex user instruction. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Pro-*
 603 *ceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp.
 604 15322–15336, Miami, Florida, USA, November 2024b. Association for Computational Linguis-
 605 tics. doi: 10.18653/v1/2024.emnlp-main.856. URL <https://aclanthology.org/2024.emnlp-main.856/>.

607 Hongru Wang, Boyang Xue, Baohang Zhou, Tianhua Zhang, Cunxiang Wang, Huimin Wang, Guan-
 608 hua Chen, and Kam fai Wong. Self-dc: When to reason and when to act? self divide-and-
 609 conquer for compositional unknown questions, 2025a. URL <https://arxiv.org/abs/2402.13514>.

610 Rui Wang, Hongru Wang, Boyang Xue, Jianhui Pang, Shudong Liu, Yi Chen, Jiahao Qiu, Derek Fai
 611 Wong, Heng Ji, and Kam-Fai Wong. Harnessing the reasoning economy: A survey of efficient rea-
 612 soning for large language models, 2025b. URL <https://arxiv.org/abs/2503.24377>.

613 Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-
 614 of-thought reasoning in llms, 2025. URL <https://arxiv.org/abs/2502.03373>.

615 Lifan Yuan, Yangyi Chen, Xingyao Wang, Yi R. Fung, Hao Peng, and Heng Ji. Craft: Customizing
 616 llms by creating and retrieving from specialized toolsets, 2024. URL <https://arxiv.org/abs/2309.17428>.

617
 618
 619
 620
 621
 622
 623
 624
 625
 626
 627
 628
 629
 630
 631
 632
 633
 634
 635
 636
 637
 638
 639
 640
 641
 642
 643
 644
 645
 646
 647

648 A THE USE OF LARGE LANGUAGE MODELS (LLMs)
649650
651 LLMs did not play significant roles in this paper’s research ideation and/or writing to the extent
652 that they could be regarded as a contributor. In the experiments, LLMs are treated as the main
653 experimental object.
654655 B RELATED WORK
656657
658 **Reward Shaping in RL.** Reward shaping plays a critical role in reinforcement learning, as it di-
659 rectly shapes the behavior the model learns to optimize (Laud, 2004; Wang et al., 2022). Recent
660 advances have introduced several reward signals for LLMs to consider the correctness and the in-
661 ternal knowledge-only reasoning efficiency together such as the length of responses and difficulty
662 of the questions (Aggarwal & Welleck, 2025; Arora & Zanette, 2025; Wang et al., 2025b). For in-
663 stances, Length Controlled Policy Optimization (LCPO) (Aggarwal & Welleck, 2025) is proposed
664 to satisfy the length constraints while optimizing reasoning performance and some efforts try to dy-
665 namically allocate inference time compute based on task complexity (Arora & Zanette, 2025; Hou
666 et al., 2025).
667668 C TOOL-INTEGRATED REINFORCEMENT LEARNING
669670
671 **PPO in TRL.** Proximal Policy Optimization (PPO) is a widely used policy-gradient method in lots
672 of tasks (Ouyang et al., 2022). Given our formulation, PPO updates the policy, using trajectories
673 sampled from the previous policy, and maximizes the following objective:
674

675
676
$$\mathcal{J}_{\text{PPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, y \sim \pi_{\text{old}}} \left[\frac{1}{\sum_{t=1}^{|y|} \mathbb{I}(y_t)} \sum_{t=1}^{|y|} \mathbb{I}(y_t) \cdot \min \left(\rho_t A_t, \text{clip}(\rho_t, 1 - \epsilon, 1 + \epsilon) A_t \right) \right], \quad (8)$$

677

678
679 where π_θ and π_{old} are current and previous policy models, and $p_t = \frac{\pi_\theta(y_t | q, y_{<t}; \mathcal{E})}{\pi_{\text{old}}(y_t | q, y_{<t}; \mathcal{E})}$. Here, $\mathbb{I}(y_t)$
680 is an indicator function marking whether token y_t is generated by the model (i.e., r_i and $t c_i$) or
681 returned from the environment \mathcal{E} (i.e., o_i). The advantage estimate A_t is computed via Generalized
682 Advantage Estimation (GAE) and ϵ a PPO clipping threshold to constrain the policy update.
683684
685 **GRPO in TRL.** To improve the stability of policy optimization in language models and avoid re-
686 liance on an additional value function approximation, Group Relative Policy Optimization (GRPO)
687 is introduced which uses the relative ranking of multiple sampled outputs as a baseline for comput-
688 ing advantages, rather than fitting a value function. For each input question q , GRPO samples a
689 group of G response $\{y_1, y_2, \dots, y_G\}$ from the reference policy π_{ref} . The trained policy π_θ is then
690 updated by maximizing the following objective:
691

692
693
$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\text{old}}(\cdot | q; \mathcal{E})} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{\sum_{t=1}^{|y_i|} \mathbb{I}(y_{i,t})} \sum_{t=1}^{|y_i|} \mathbb{I}(y_{i,t}) \right. \\ \left. \cdot \min \left(p_t \hat{A}_{i,t}, \text{clip} (p_t, 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t} \right) \right] - \beta \mathbb{D}_{\text{KL}}[\pi_\theta \| \pi_{\text{ref}}], \quad (9)$$

694
695

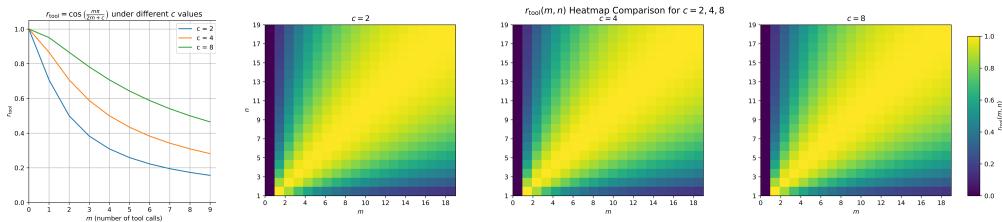
696
697 Here, $p_t = \frac{\pi_\theta(y_{i,t} | x, y_{i,<t}; \mathcal{R})}{\pi_{\text{old}}(y_{i,t} | x, y_{i,<t}; \mathcal{R})}$, $\hat{A}_{i,t}$ denotes the advantage at token t in response y_i , computed based
698 on the relative ranking of rewards within the group, β controlling the KL-regularization strength.
699 The clipping threshold ϵ ensures stable updates.
700

702 **D SEARCH TEMPLATE**
703

704
705 Answer the given question. You must conduct reasoning inside `<think>` and `</think>` first
706 every time you get new information. After reasoning, if you find you lack some knowledge,
707 you can call a search engine tool by `<search>` query `</search>`, and it will return the top
708 searched results between `<observation>` and `</observation>`. **You need to make every search**
709 **call count and gain helpful results.** If you find no further external knowledge needed, you can
710 directly provide the answer inside `<answer>` and `</answer>` without detailed illustrations.
711 For example, `<answer>` xxx `</answer>`. Question: question.

712 **E REWARD FUNCTION**
713714 **E.1 FIGURE ILLUSTRATION**
715

716 We draw the two types of r_{tool} defined in the main content for better understanding. Fig. 5 shows
717 the illustration of these reward functions for OTC-PPO and OTC-GRPO respectively. It is very clear
718 that: 1) Left: as the number of tool calls increases, the r_{tool} decreases accordingly. Thus when
719 multiple trajectories leads to correct answer, the one that use less tools will get higher reward and
720 the one that do not use tool will get the highest reward; 2) Right: It is obvious that the diagonal
721 achieves the highest reward as the number of tool calls is the less one, and as the m increases, the
722 color brightness becomes larger since multiple tool calls are involved.
723



732 Figure 5: **Left:** the illustration of tool-use reward functions in OTC-PPO; and **Right:** the illustration
733 of tool-use reward functions in OTC-GRPO.
734

735 **E.2 THEORETICAL JUSTIFICATION**
736

737 We mainly follow the justification (Section 4.2) in the paper (Arora & Zanette, 2025) to showcase
738 our proposed reward design leads to reducing the tool calls without compromising accuracy. We
739 strongly encourage readers to refer to the previous paper for complete details. All assumptions
740 made in that work still hold in our setting, with the only difference being that our response is defined
741 as $y = (\tau, a)$, which includes both the tool-integrated reasoning trajectory and the final answer, rather
742 than language-only reasoning.
743

744 Let θ_{eff}^* denote the population-level parameters of the policy models obtained by maximizing Equation
745 7, i.e.,
746

$$\theta_{eff}^* = \arg \max_{\theta} \{ \mathbb{E}_{x \sim p} \mathbb{E}_{y \sim p_{\theta}(x)} [(\alpha * r_{tool})] \} \quad (10)$$

747 as $r_{\phi}(q, y)$ is mainly the indicator function about the correctness and format. Therefore, the
748 population-level maximizer $p_{\theta_{eff}}^*$ is as accurate as the population-level maximizer p_{θ}^* and $\text{Acc}(p_{\theta_{eff}}^*) = 1$.
749

750 **E.3 GENERALIZATION OF REWARD**
751

752 We use two major forms of r_{ϕ} : i) $r_{\phi}^1 = r_{correct}$; and ii) $r_{\phi}^2 = r_{correct} + r_{format}$ to illustrate the
753 generalization of our proposed tool-integrated reward design $r_{\phi}^{tool} = \alpha * r_{tool} * r_{\phi}$.
754

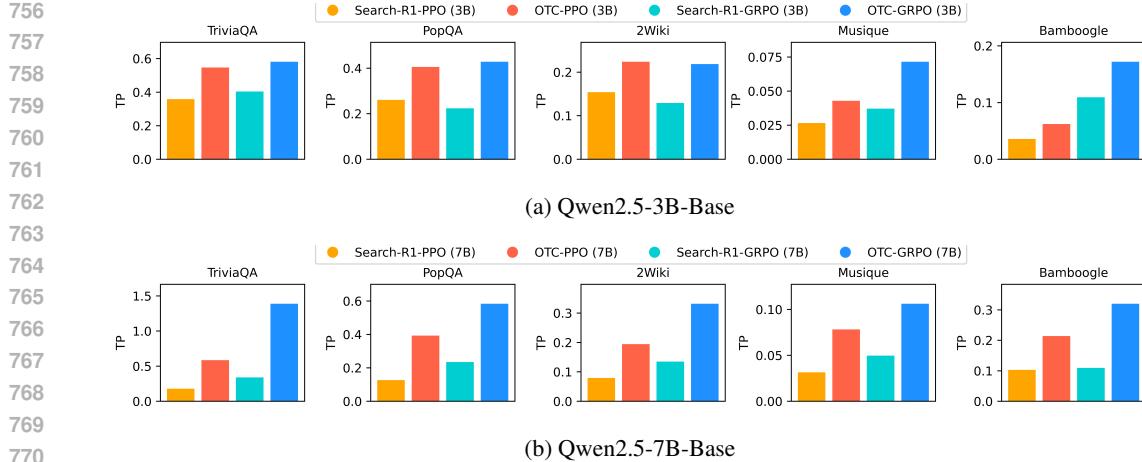


Figure 6: The Out-of-domain performance of OTC-PO and Search-R1 in TP.

Table 6: The results of OTC-PO with different α on Qwen2.5-7B-Base model

| Models | NQ | | | HotpotQA | | |
|--------------|-------------------|---------------------|-------------------|-------------------|---------------------|-------------------|
| | EM (\uparrow) | TC (\downarrow) | TP (\uparrow) | EM (\uparrow) | TC (\downarrow) | TP (\uparrow) |
| $\alpha = 1$ | 0.446 | 1.040 | 0.429 | 0.383 | 1.464 | 0.262 |
| $\alpha = 2$ | 0.354 | 1.571 | 0.225 | 0.320 | 2.062 | 0.155 |
| $\alpha = 3$ | 0.389 | 1.530 | 0.254 | 0.340 | 1.948 | 0.175 |

If the answer is correct, the r_ϕ^1 and r_ϕ^2 are both positive and fixed for all questions, therefore the r_ϕ^{tool} is only determined by the number of α and r_{tool} . Thus it can assign higher score for less tools and lower score for more tools by definition of r_{tool} .

If the answer is wrong, the r_ϕ^1 will be 0, and therefore disables the influence of r_{tool} , reducing the reward hacking issue. The r_ϕ^2 will be r_{format} and then r_ϕ^{tool} becomes $\alpha * r_{tool} * r_{format}$ which is acceptable. Considering two cases of r_{format} , if r_{format} is 0, then the final reward is also 0; and if r_{format} is a positive, the r_ϕ^{tool} still holds as less tools will be assigned more rewards. This is reasonable since less tool calls means less cost especially when the answer is wrong. We note that the reward gap here between wrong answer (i.e., r_{format}) and correct answer (i.e., $r_{correct} + r_{format}$) is significant for the policy model to learn the desirable behavior.

F ANALYSIS OF SEARCH AS TOOL

F.1 THE EFFECTS OF α

Table 6 shows the performance of OTC-PO with different α . It can be found that increasing α can not lead to better results. We further check the reward changes during the training and find that it becomes harder for the model to learn the desired behavior when using a larger α , because a single tool call can cause large fluctuations in the reward, especially considering the relatively lower C . We suspect the α will be more important in much longer tool-integrated reasoning trajectories.

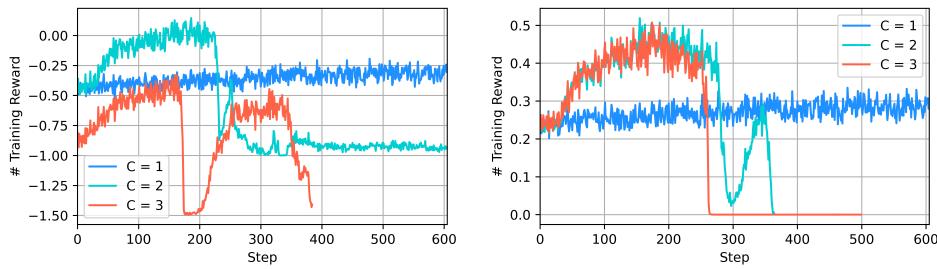
G ANALYSIS OF CODE AS TOOL

G.1 THE EFFECTS OF C

Table 7 shows the effects of C in ToRL. It is observed that i) the larger C always leads to more tool calls on ToRL-GRPO if we only consider the final correctness as the reward and do not penalize the tool use behaviors of LLMs; ii) OTC-GRPO achieves more stable tool calls which is more reasonable as the optimal number of tool calls should not be affected by C , and leads to bigger TP

Table 7: The effects of C on Qwen2.5-Math-7B Base model.

| Models | AIME24 | | | AIME25 | | |
|---------------------------|-------------------|---------------------|---------------------------|-------------------|---------------------|---------------------------|
| | EM (\uparrow) | TC (\downarrow) | TP (\uparrow) | EM (\uparrow) | TC (\downarrow) | TP (\uparrow) |
| $C = 1$ | | | | | | |
| ToRL-GRPO | 30.0 | 0.9 | 33.3 | 26.7 | 0.9 | 29.7 |
| OTC-GRPO | 35.8 | 0.8 | 44.8 ($\Delta 34.5\%$) | 26.7 | 0.8 | 33.4 ($\Delta 11.1\%$) |
| $C = 2$ | | | | | | |
| ToRL-GRPO | 33.3 | 1.4 | 23.8 | 25.8 | 1.4 | 18.4 |
| OTC-GRPO | 33.3 | 1.0 | 33.3 ($\Delta 39.9\%$) | 23.3 | 0.7 | 33.3 ($\Delta 81.0\%$) |
| $C = 3$ | | | | | | |
| ToRL-GRPO | 36.7 | 2.1 | 17.5 | 26.7 | 2.1 | 12.7 |
| OTC-GRPO | 36.7 | 0.7 | 52.4 ($\Delta 199.4\%$) | 23.3 | 0.8 | 29.1 ($\Delta 129.1\%$) |

Figure 7: **Left:** Training reward changes of ToRL-GRPO; and **Right:** Training reward changes of OTC-GRPO.Table 8: The results of OTC-PO on MATH when $C = 3$.

| Models | MATH | | |
|---------------------------------|-------------------|--------------------------|----------------------------|
| | EM (\uparrow) | TC (\downarrow) | TP (\uparrow) |
| Qwen2.5-Math-1.5B(-Base) | | | |
| ToRL-GRPO | 52.2 | 1.8 | 29.0 |
| OTC-GRPO | 53.7 | 1.1 ($\Delta 38.9\%$) | 48.8 ($\Delta 68.28\%$) |
| Qwen2.5-Math-7B(-Base) | | | |
| ToRL-GRPO | 59.7 | 1.49 | 40.07 |
| OTC-GRPO | 61.9 | 0.68 ($\Delta 54.4\%$) | 91.03 ($\Delta 127.2\%$) |

improvement as C increases. Further, Fig. 7 shows the rewards changes with different C . It can be found that GRPO tends to collapse earlier with bigger C and our OTC-PO can delay this, as also observed using search as tool.

G.2 ADDITIONAL RESULTS

We additionally run the evaluation on MATH dataset following ToRL (Li et al., 2025b), and Table 8 shows the final results. The conclusion also holds as analyzed in the main experiments, revealing the great generalization capability of our method.

864 **H CASE STUDY**
865866 **H.1 MORE EXAMPLES**
867868 Fig. 8 and Fig. 9 show more examples between Search-R1 and OTC-PO methods. Fig. 10 and
869 Fig. 11 show more examples between ToRL-GRPO and OTC-GRPO. It is clear that OTC-PO can
870 leverage more reasoning capabilities while minimizing external tool calls.
871872 **Case Study: No Tool Call Required**
873874 Question: The Bellagio resort owned by MGM Resorts international is located in which US
875 city?
876877 Answer: Las Vegas
878879 **Search-R1-PPO:** <think> I need to find the US city where the Bellagio resort owned by
880 MGM Resorts International is located. </think>
881 <search> Bellagio resort owned by MGM Resorts International </search>
882 <observation> ... </observation>
883 <search> Bellagio resort city </search>
884 <observation> ... </observation>
885 <think> Based on the observations, the Bellagio resort owned by MGM Resorts Interna-
886 tional is located in Las Vegas, Nevada. </think>
887 <answer> Las Vegas </answer> **(Tool Calls = 2)**
888889 **OTC-GRPO:** To answer this question, let's first reason through the information provided:
890 1. We need to find the US city where the Bellagio resort, owned by MGM Resorts Interna-
891 tional, is located.
892 2. The Bellagio resort is a well-known hotel and casino located in Las Vegas, Nevada.
893 3. Therefore, the US city where the Bellagio resort is located is Las Vegas.
894 <answer> Las Vegas </answer> **(Tool Calls = 0)**895 Figure 8: An example that OTC-GRPO provide correct answer solely rely on its own internal rea-
896 soning without any external tool calls while Search-R1 extensively call unnecessary tools.
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

918
919
920
921
922
923
924
925
926
927

928 Case Study: Less Tool Calls with Accurate Search 929

930 Question: What science fantasy young adult series, told in first person, has a set of
931 companion books narrating the stories of enslaved worlds and alien species?

932 Answer: Animorphs

933 **Search-R1-PPO:** <think> I need to find a science fantasy young adult series that is told in
934 first person and has a set of companion books narrating the stories of enslaved worlds and
935 alien species. </think>

936 <search> science fantasy young adult series told in first person </search>

937 <observation> ... <observation>

938 <search> companion books animorphs enslaved worlds alien species </search>

939 <observation> ... </observation>

940 <search> science fantasy young adult series told in first person enslaved worlds alien
941 species companion books </search>

942 <observation> ... </observation>

943 <think> Based on the observations, the science fantasy young adult series told in first person
944 and has a set of companion books narrating the stories of enslaved worlds and alien species
945 is "Animorphs". </think>

946 <answer> Animorphs </answer> **(Tool Calls = 3)**

947
948 **OTC-GRPO:** To answer this question, let's first reason through the information provided:

949 1. We need to find a science fantasy young adult series.

950 2. The series is told in first person.

951 3. It has a set of companion books narrating the stories of enslaved worlds and alien species.

952 Let's search for this information:

953 <search> What science fantasy young adult series, told in first person, has a set of companion
954 books narrating the stories of enslaved worlds and alien species? </search>

955 <observation> ... </observation>

956 Based on the search results, the science fantasy young adult series that fits the description is
957 "Animorphs".

958 <answer> Animorphs </answer> **(Tool Calls = 1)**

959
960 Figure 9: An example that OTC-GRPO uses less tool calls but with more accurate search to arrive a
961 correct answer.
962

963
964
965
966
967
968
969
970
971

972
973
974
975976 **Question:** Every morning Aya goes for a $9\frac{1}{2}$ -kilometer-long walk and stops at a coffee shop afterwards. When she walks at a constant speed
977 of s kilometers per hour, the walk takes her 4 hours, including 24 minutes spent in the coffee shop. When she walks $s+2$ kilometers per
978 hour, the walk takes her 2 hours and 24 minutes, including 24 minutes spent in the coffee shop. Suppose Aya walks at $s+\frac{1}{2}$ kilometers per hour. Find the number of minutes the walk takes her, including the 24 minutes spent in the coffee shop. **Answer:** 204979 **ToRL-GRPO:** To solve this problem, we need to determine ...980

```
```python
981 from sympy import symbols, Eq, solve
982 s = symbols('s')
983 equation = Eq(4 - 9/s, 2.4 - 9/(s+2))
984 s_solution = solve(equation, s)
985 print(s_solution)
986
987 ```output
988 [-4.50000000000000, 2.50000000000000]
```

989 The solutions for  $s$  are  $-4.5$  and  $2.5$ . Let's confirm  
990 the solution...991 

```
```python
992     ...
993     time_spent_walking = 9 / new_speed
994     total_time_minutes = (time_spent_walking * 60) + t_value
995     print(total_time_minutes)
996
997     ```output
998     204.0
```

999 The total time it takes for Aya to walk at a speed of 3 km/h,
1000 including the 24 minutes spent in the coffee shop, is
1001 $\boxed{204}$ (Tool Calls = 2)

1002

1003 **OTC-GRPO:** To solve this problem, we need to determine ...1004

```
```python
1005 import math
1006 a = 1.6, b = 3.2, c = -18
1007 discriminant = b**2 - 4*a*c
1008 s1 = (-b + math.sqrt(discriminant)) / (2*a)
1009 s2 = (-b - math.sqrt(discriminant)) / (2*a)
1010 print((s1, s2))
1011
1012 ```output
1013 (2.4999999999999996, -4.499999999999999)
```

1014 Since speed cannot be negative, we have  $s = 2.5$  km/h. Now we can substitute  $s = 2.5$  back into one of the original equations to  
1015 find  $t$ . Using the first equation:  $t = \frac{9}{s} = \frac{9}{2.5} = 3.6$ . So,  $t = 3.6 + 24 = 27.6$  minutes. Now we need to find the total time when Aya walks at  $s + \frac{1}{2}$  km/h. The time spent walking is  $\frac{9}{s + \frac{1}{2}} = 2.5 + 0.5 = 3$  km/h. The time spent walking is  $\frac{9}{3} = 3$  minutes. The total time including the time spent in the coffee shop is  $3 + 24 = 27$  minutes. Thus, the number of minutes the walk takes her, including the  $24$  minutes spent in the coffee shop, is  $27 + 24 = 51$ . The final answer is:  $\boxed{51}$ 

1016 (Tool Calls = 1)

Figure 10: An example of OTC-GRPO use *less tool calls* to get the correct answer.994  
995  
996  
997  
998  
999  
1000  
1001  
10021003  
10041005 **Question:** Find all values of  $x$  that satisfy the equation  $x = \sqrt{11-2x} + 4$ .  
1006 **Answer:** 51007 **ToRL-GRPO:** To solve the equation  $x = \sqrt{11-2x} + 4$ , we can follow these steps ....  
1008 

```
```python
1009     import sympy as sp
1010     x = sp.symbols('x')
1011     equation = x - (sp.sqrt(11 - 2*x) + 4)
1012     isolated_sqrt = sp.Eq(x - 4, sp.sqrt(11 - 2*x))
1013     squared_equation = sp.Eq((x - 4)**2, 11 - 2*x)
1014     simplified_equation = sp.simplify(squared_equation)
1015     solutions = sp.solve(simplified_equation, x)
1016     valid_solutions = [sol for sol in solutions if equation.subs(x, sol) == 0]
1017     print(valid_solutions)
1018
1019     ```output
1020     [5]
```

1021 The solution to the equation $x = \sqrt{11-2x} + 4$ is $x = 5$. Therefore, the final answer is $\boxed{5}$.
1022 (Tool Calls = 1)

1023

1024 **OTC-GRPO:** To solve the equation $x = \sqrt{11-2x} + 4$, we can follow these steps:1025 1. Isolate the square root term.
1026 2. Square both sides to eliminate the square root.
1027 3. Solve the resulting equation.
1028 4. Verify the solutions by substituting back into the original equation to ensure they satisfy it.

1029 Let's go through these steps systematically.

1030 1. Isolate the square root term:
1031 2. Square both sides:
1032 3. Expand and simplify the equation:
1033 4. Solve the quadratic equation:1034 So, the solutions are $x = 5$ and $x = 1$.1035 5. Verify the solutions by substituting back into the original equation.
1036 For $x = 5$: $5 = \sqrt{11 - 2 \cdot 5} + 4 \Rightarrow 5 = \sqrt{1} + 4 \Rightarrow 5 = 5$. This is true, so $x = 5$ is a valid solution.
1037 For $x = 1$: $1 = \sqrt{11 - 2 \cdot 1} + 4 \Rightarrow 1 = \sqrt{9} + 4 \Rightarrow 1 = 3 + 4 \Rightarrow 1 = 7$. This is false, so $x = 1$ is not a valid solution.
1038 The final answer is: $\boxed{5}$ Figure 11: An example of OTC-GRPO *does not need any external tool calls* to get the correct answer.1024
1025