# RLAX: Large-Scale, Distributed Reinforcement Learning for Large Language Models on TPUs

**Runlong Zhou**[w†] **Lefan Zhang**[†] **Shang-Chen Wu**[†] **Kelvin Zou**[†*]
**Hanzhi Zhou** **Ke Ye** **Yihao Feng** **Dong Yin** **Alex Guillen Garcia** **Dmytro Babych**
**Rohit Chatterjee** **Matthew Hopkins** **Xiang Kong** **Chang Lan** **Lezhi Li** **Yiping Ma**
**Daniele Molinari** **Senyu Tong** **Yanchao Sun** **Thomas Voice** **Jianyu Wang** **Chong Wang**
**Simon Wang** **Floris Weers** **Yechen Xu**[d] **Guolin Yin** **Muyang Yu** **Yi Zhang** **Zheng Zhou**
**Danyang Zhuo**[d] **Ruoming Pang** **Cheng Leong***
Apple [w]University of Washington [d]Duke University

## Abstract

Reinforcement learning (RL) has emerged as the de-facto paradigm for improving the reasoning capabilities of large language models (LLMs). We have developed **RLAX**, a scalable RL framework on TPUs. **RLAX** employs a parameter-server architecture. A master trainer periodically pushes updated model weights to the parameter server while a fleet of inference workers pull the latest weights and generates new rollouts. We introduce a suite of system techniques to enable scalable and preemptible RL for a diverse set of state-of-art RL algorithms. To accelerate convergence and improve model quality, we have devised new dataset curation and alignment techniques. Large-scale evaluations show that **RLAX** improves `QwQ-32B` 's pass@8 accuracy by 12.8% in just 12 hours 48 minutes on 1024 v5p TPUs, while remaining robust to preemptions during training.

## 1 Introduction

Reinforcement learning (RL) has become a critical step for large language model (LLM) post training. It significantly strengthens an LLM's capability at aligning with human preferences [28] and reasoning tasks (e.g., mathematics, coding). Today, almost all top performers in AI coding benchmarks have become RL-based reasoning models, including OpenAI o3 [27], Claude 4 [4], Grok 4 [41], Gemini 2.5 [11], DeepSeek R1 [14], and Qwen 3 [43].

We introduce **RLAX**, a reinforcement learning framework designed to execute state-of-the-art RL algorithms on large-scale distributed TPU clusters efficiently. At the same time, **RLAX** is fully preemptible, enabling higher-priority jobs (e.g., inference workloads) to immediately reclaim TPU resources when needed. **RLAX** has adopted a disaggregated architecture, where trainers, inference workers, and verifiers are logically separated. The logical separation enables **RLAX** to flexibly allocate compute resource to them independently.

**RLAX** uses an existing model training system, AXLearn [22], as its default training backend. AXLearn is built on top of JAX [7] and XLA [1]. It supports configurable model and data parallelism. On the inference worker side, we repurpose AXLearn as an inference engine by incorporating standard inference optimizations such as paged attention [21] and continuous batching [45]. Model weights are

[†]Runlong Zhou, Lefan Zhang, Shang-Chen Wu, and Kelvin Zou are core authors.
[w]Runlong Zhou, [d]Yechen Xu, and [d]Danyang Zhuo contributed to this work while visiting Apple.
*Kelvin Zou and Cheng Leong are the corresponding authors.
Kelvin Zou, Hanzhi Zhou, Ke Ye, Floris Weers, Chong Wang, Yi Zhang, and Ruoming Pang have left Apple. They contributed to this work while they were employed by Apple.

synchronized between inference workers and trainers via distributed parameter servers. The verifiers are language-specific code-execution runtimes optimized for each programming language in the training corpus. For example, to ensure efficient and deterministic verification of Python programs, we containerize standard Python dependencies and spawn identical containers for code execution.

**RLAX** focuses on addressing four key challenges in large-scale LLM post-training using RL: (1) how to build a training framework that flexibly supports a wide range of RL algorithms; (2) how to efficiently handle both on-policy and off-policy RL; (3) how to scale RL training to large distributed clusters while enabling seamless, anytime preemption; and (4) how to maintain numerical alignment between the trainer and the inference workers.

**RLAX** exposes programmable knobs for users to specify RL algorithms and different RL paradigms (i.e., on-policy, off-policy). RL algorithms are highly structured and can be expressed succinctly using polymorphism. The training paradigms are realized through parameter server versioning and runtime policies within **RLAX**. The key distinction between them lies in which rollouts the trainer consumes. In on-policy RL, all rollouts are generated using the most recent version of the model weights. In contrast, off-policy RL allows rollouts generated by older versions of the model to be reused. To manage this behavior, **RLAX** provides user-configurable knobs that enforce staleness bounds, specifying both how often inference workers pull new model weights and the maximum rollout staleness tolerated by the trainer. This enables users to flexibly choose between on-policy and off-policy RL, as well as to tune the degree of staleness permitted in off-policy training.

To support preemption at any point during execution, **RLAX** adopts a simple but powerful design philosophy: every component in the system is passive, with the controller serving as the only active entity. Moreover, all controller actions are idempotent. If **RLAX** is preempted, its current state is naturally captured using model checkpointing. The in-memory states can be lost, and the controller will reconstruct its state from the checkpoint and resume.

Finally, in large-scale RL for language models, small numerical mismatches between training and inference can distort log-probabilities and destabilize model training. **RLAX** recomputes log-probabilities using an inference-matched training graph. These measures significantly reduce numerical drift and provide the stable log-prob estimates needed for reliable off-policy RL.

We evaluate **RLAX** by training `QwQ-32B` using Codeforces problems from 2013-2024 as training data and testing on Codeforces problems in 2025. Our results show that **RLAX** improves the pass@8 accuracy of `QwQ-32B` by 12.8% in just 12 hours and 48 minutes on 1024 v5p TPUs, while maintaining robustness to preemptions during training.

## 2 Background

We first describe the general research area of LLM post-training using RL. We then describe several key problems **RLAX** has focused on.

**RL for post-training LLMs.** RL has become a critical process for post-training LLMs. Specifically, RL from human feedback [28, 31] enables better alignment with human preferences, such as following instructions and generating safe content. Scoring functions can also be mathematical functions or program compilers/interpreters, allowing RL to improve an LLM's mathematics and coding skills [14]. Today, all the AI coding benchmarks have been dominated by LLMs trained via RL [19, 23, 10].

A typical RL training iteration works as follows: an LLM samples token sequences based on a set user prompts using an LLM inference worker, e.g., "write a Python program that output the sums of all the integers from 1 to 100". The output token sequences is a Python program. These generated Python programs are then verified through a verifier (i.e., Python interpreter) to check whether the result is correct (in this case, the correct result is 5050). The trainer then updates the model weights in order to increase the probability of the correct program from being generated by the LLM.

**Problem #1: Support for Various RL Algorithms.** A large variety of RL algorithms exist for LLMs, making comprehensive support for these diverse and rapidly evolving algorithms a highly desirable property of an RL training system, while remaining easy to maintain and debug. To introduce these algorithms, we use $q$ to denote the question (user prompt), $\{o_i\}_{i=1}^{G}$ to denote a group of $G$ outputs generated for $q$, and $\{\widehat{A}_i\}_{i=1}^{G}$ to denote the advantage estimate for each output given $q$.

Here $o_i$ is generated using $\pi_{\theta_{\text{old}}}$, which can be either a previous checkpoint or a detached version of the current policy. To compensate for possible distribution shift, an importance sampling weight is typically used to make the loss function "on-policy":

$$r_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}. \tag{1}$$

We also denote $r_i(\theta) = \prod_{t=1}^{|o_i|} r_{i,t}(\theta)$ as the importance weight of the entire generation.

One type of algorithm uses $\log \pi_\theta$ as the source of gradients due to the log-derivative trick. These include REINFORCE [39, 36], which performs gradient ascent on the following surrogate objective:

$$\mathcal{J}_{\text{REINFORCE}}(\theta) = \mathop{\mathbb{E}}_{\substack{q\sim\mathcal{D} \\ \{o_i\}_{i=1}^{G}\sim\pi_{\theta_{\text{old}}}(\cdot|q)}} \left[ \frac{1}{G}\sum_{i=1}^{G} \mathsf{sg}[r_i(\theta)]\widehat{A}_i \log \pi_\theta(o_i|q) \right]. \tag{2}$$

Here $\mathsf{sg}[]$ is the stopping-gradient operator. Other notable instances of this type of algorithm include CISPO [9] and the single-sided clipping objective in LlamaRL [40].

Proximal policy optimization (PPO, [31]) uses the importance ratio $r_{i,t}(\theta)$ as the source of gradients. In the previous category, even though $r_{i,t}(\theta)$ may appear in the objective, its gradient is stopped. Here we use the clipped surrogate objective of PPO as an example:

$$\mathcal{J}_{\text{PPO}}(\theta) = \mathop{\mathbb{E}}_{\substack{q\sim\mathcal{D} \\ \{o_i\}_{i=1}^{G}\sim\pi_{\theta_{\text{old}}}(\cdot|q)}} \left[ \frac{1}{G}\sum_{i=1}^{G} \frac{1}{|o_i|}\sum_{t=1}^{|o_i|} \min\left\{ r_{i,t}(\theta)\widehat{A}_i,\ \mathsf{clip}(r_{i,t}(\theta), 1-\epsilon, 1+\epsilon)\widehat{A}_i \right\} \right].$$

$\widehat{A}_i$ is computed using a learned value network. The clipping mechanism in PPO constrains the importance ratio, preventing excessively large policy updates that could destabilize training. When the advantage $\widehat{A}_i$ is positive, the clipping discourages the new policy from deviating too far ($> 1+\epsilon$) from the old policy in the direction of increasing probability for that action. Conversely, when the advantage is negative, it limits how much the probability can be decreased. This conservative update strategy balances exploration with stable, incremental policy improvement. Other variants in this category include GRPO [32], DAPO [46], GSPO [47], and GFPO [34].

**Problem #2: Support for Different Training Paradigms.** RL has diverse distributed training paradigms. We would like to support all the paradigms to facilitate model engineers exploring different choices. The most obvious training paradigm is *on-policy* RL, where rollout generations and trainers do not overlap. Inference workers always use the most up-to-date model weights, and the controller must wait till all the inference workers finish their rollouts in every training iteration. On-policy RL provides our model engineers precise control and predictability during training process. However, this also means when the trainer trains the neural network, the inference workers are idle, and vice versa. To make training infrastructure achieve high utilization, *off-policy* RL allows the trainer to accept training data (i.e., the generated rollouts from the inference workers) from stale model weights. This means the trainer and inference workers can run concurrently. However, the staleness can adversely affect the training convergence. Several RL systems (e.g., StreamRL [48]) specifically optimize for training throughput given staleness constraints.

**Problem #3: Scaling and Supporting Preemption.** We leverage large-scale TPU infrastructure to support the demanding computational requirements of RL, driven by increasing model sizes, extended context lengths, and the integration of Chain-of-Thought (CoT). We want our RL system to be scalable, where we can easily configure the training cluster size and allocating resources to the inference workers and the trainer independent. Further, because we operate on large-scale shared infrastructure, we must support preemption when higher-priority jobs, such as inference, need to reclaim resources. However, RL systems have multiple components, and fully checkpointing and restoring the state of all the components in a black-box fashion is inefficient. For example, there is no need to checkpoint the kv cache states in the inference workers.

**Problem #4: Numerical Consistency between the Inference System and the Trainer.** A key technique in RL is importance sampling (IS), which uses a ratio $r(\theta)$ to correct for differences

between policies. This ratio is defined as: (c.f., Equation (1))

$$r(\theta) = \frac{\pi_\theta(a|s)}{\pi_{\theta_{\text{old}}}(a|s)}.$$

Here, $\pi_\theta$ represents the current policy being optimized, and $\pi_{\theta_{\text{old}}}$ is the behavior policy that was used to generate the training trajectories. In an off-policy setting, the policies are different by definition. The IS ratio $r(\theta)$ is essential to re-weight observed rewards and gradients, correcting for the data-policy mismatch. In a strictly on-policy setting, the behavior policy and the current policy are identical ($\pi_\theta = \pi_{\theta_{\text{old}}}$). Therefore, the importance sampling ratio should theoretically always be $r(\theta) = 1.0$. However, we have observed in practice that $r(\theta) \neq 1.0$ even in on-policy implementations. This phenomenon has also been reported in other work, such as [17] and [29].

The root cause of this discrepancy is the non-associative nature of floating-point arithmetic, especially in low-precision formats like bfloat16. Floating-point addition is not associative, meaning $(a + b) + c$ may not equal $a + (b + c)$ due to intermediate rounding errors. When calculations are performed in different orders, the final accumulated values can diverge.

This numerical instability can stem from two sources: compute non-determinism or implementation divergence. Compute non-determinism is a common issue on accelerators like GPUs. During concurrent operations like atomic-add, the accumulation order is not guaranteed, which can yield different results even with the exact same inputs. The fix on GPUs is to avoid concurrency in the reduction dimension by enforcing a sequential accumulation, which makes the result deterministic [17]. In contrast, a TPU's systolic array architecture inherently solves this problem by design, as it provides a fixed, deterministic accumulation order for numerical blocks.

Because compute non-determinism is not a problem on TPU, the remaining issue of **RLAX** is the implementation divergence between the trainer and inference service. This divergence is driven by two primary differences: (1) different parallelism strategies employed in the trainer and the inference service, and (2) variations in compiler-driven kernel fusion, both of which alter the final order of operations.

## 3 System Design

**RLAX** is a scalable, preemptible RL framework for LLM post-training on TPUs. Figure 1 shows the overall system architecture. **RLAX** adopts a parameter-server–based design with a single centralized controller. It comprises five key components: an inference service (which consists of a set of inference workers), a controller, a trainer, a parameter server, and a verification service (which includes coding and math verifiers). The controller orchestrates when the trainer, inference workers, and verifiers run based on the chosen RL algorithm.

Built on TensorStore [16], the **RLAX** parameter server extends TensorStore's native sharding schemes to support in-memory persistence and custom version management. We align parameter-server partitioning with the model's native weight sharding. For example, in a configuration with data parallelism of 4, expert parallelism of 8, and model parallelism of 8 (across 8-TPU-core hosts), **RLAX** deploys four parameter-server replicas, each spanning eight hosts to match the model-parallelism degree. Inference workers are evenly mapped to these replicas.

To support off-policy RL with bounded staleness, each parameter server maintains up to $N$ weight snapshots in host memory. Garbage collection of checkpoints prior to step $t$ occurs only once global progress to step $t$ is confirmed across all replicas. Furthermore, we extended the TensorStore client to support gRPC-based communication enriched with additional metadata. Loading times can be further accelerated by enabling RDMA for weight transfers.

**RLAX** uses AXLearn [22] as the trainer. For each RL iteration, the trainer updates the model weights in the parameter server using the training data from the controller. RL algorithms are implemented by extending AXLearn modules, which we collectively refer to as the RL plugin in AXLearn.

The inference service consists of multiple inference workers, each comprising a group of TPUs capable of independently running model inference. Each inference worker is an instance of AXLearn using its inference mode. Inference workers pull the latest model weights from the parameter server and generate rollouts based on prompts provided by the controller.

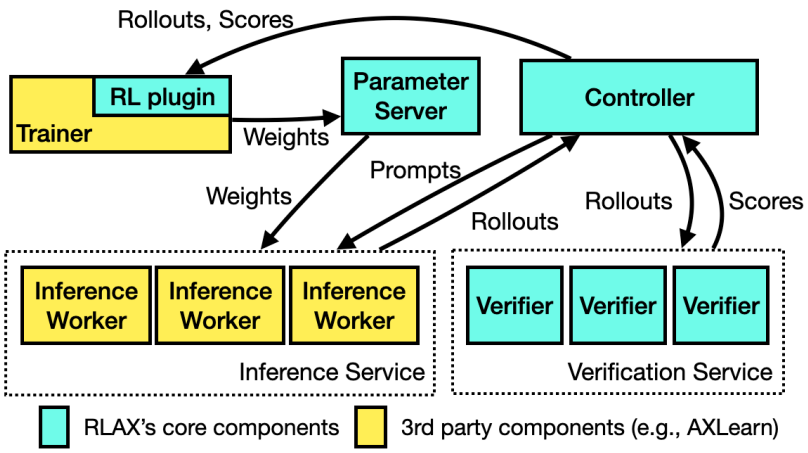


Figure 1: **RLAX**'s system diagram. Blue parts represent **RLAX**'s core software components. Yellow parts represent 3rd party components we use. **RLAX** uses AXLearn for trainer and inference workers.

The controller forwards these rollouts to the verification service. The verification service includes a pool of code-execution environments, one for each programming language present in the training distribution (e.g., Python, C++, Rust, Go). This design supports training batches that contain multiple programming languages code writing tasks. Both rollouts and model weights are periodically checkpointed to persistent storage. These checkpoints ensure fault tolerance and reproducibility in the event of preemption or system failures.

In the rest of the section, we describe how **RLAX** 's system design supports different training paradigms and handles preemption.

### 3.1 Support for Different Training Paradigms

**RLAX** is designed to flexibly support a wide range of training paradigms. Users can choose on-policy and off-policy RL by configuring two staleness parameters: the inference reload staleness and the trainer acceptance staleness. The inference reload staleness $j$ determines how frequently inference workers fetch updated model weights from the parameter server. In other words, $j$ means that workers reload only every $j$-th model version. The trainer acceptance staleness $k$ constrains how outdated a rollout may be when consumed by the trainer: each rollout is tagged with the model version used for its generation, and the trainer is allowed to use it only if the rollout's model version is at most $k$ steps behind the trainer's current model version.

This interface yields a unified formulation for several commonly used RL training modes:

- **On-policy RL** ($j = k = 1$). The trainer only consumes rollouts produced by the most recent model version, ensuring a strict lockstep schedule between inference, verification, and training. This corresponds to traditional synchronous PPO-style pipelines.
- **One-step off-policy RL** ($j = 1, k = 2$). The trainer may use rollouts generated by a model that is one step behind, allowing limited overlap between rollout generation and model updating. This mode aligns with the 1-step off-policy design used in StreamRL [48].
- **Off-policy RL** ($j \geqslant 1, k = \infty$). The trainer imposes no staleness constraint and may train from arbitrarily old rollouts. This maximizes throughput but sacrifices training stability, which may require additional safeguards such as off-policy corrections.

This unified abstraction allows **RLAX** to support a spectrum of RL algorithms and system-level trade-offs with minimal modification to user code or infrastructure. The trainer will always favor using more up-to-date rollouts and discard ones that are not used. We empirically found that bounded-staleness RL, with small $j$ and $k$ (e.g., $(j, k) = (16, 16)$ or $(16, 32)$, often provides a favorable trade-off: it significantly reduces idle time for inference workers while avoiding the generation of excessively stale rollouts that the trainer would later discard. In our experiments, we typically set

$(j, k) = (16, 32)$ unless otherwise noted, as this value provides strong empirical performance while enabling substantial hardware parallelism.

### 3.2 Support for Preemption

Preemption is a fundamental consideration when running large-scale RL workloads on shared or elastic clusters. **RLAX** is designed to tolerate preemption without loss of correctness or reproducibility. Our approach centers on a controller-driven, checkpoint-based recovery mechanism guided by three core design principles:

- **Principle #1: Controller-Centric, Passive Components.** All components outside the controller (i.e., inference workers, parameter-server replicas, verifiers, trainer) operate as passive executors. They perform work only in response to explicit instructions from the controller. This architecture ensures that system progress is fully serialized through the controller, eliminating hidden state transitions that would complicate recovery. Because no component autonomously advances its own state, a restarted controller can reliably reconstruct the global system state.
- **Principle #2: Persistent, Self-Contained Snapshots.** At each weight-reload boundary (typically once per training step), the controller writes a complete snapshot of its operational state to persistent storage such as Google Cloud Storage (GCS) or Amazon S3. These snapshots include the model checkpoint step, the progress of the prompt mixture, the state of RNG streams, and any additional metadata required for deterministic regeneration of rollouts. Importantly, these snapshots serve as the single source of truth for recovery, ensuring that the system can always resume from a well-defined, consistent point.
- **Principle #3: Idempotent Controller Execution.** The controller is designed to behave deterministically between any two snapshots. Starting from a saved snapshot, it will reload the same model weights, issue the same sequence of prompts, and drive the system through the same operations. This idempotence guarantees that failures between snapshots do not introduce divergence work in the training trace.

During normal operation, the trainer and inference workers read model weights from the in-memory parameter server for maximum throughput. In parallel, **RLAX** continuously checkpoints model weights to persistent object storage. This dual storage strategy ensures that training can always resume from persistent storage, even if the trainer process, the parameter server, or its host machine is entirely lost.

All randomness used in prompt mixture sampling is controlled by reproducible RNG streams whose states are included in every controller snapshot. The controller also logs the number of prompts consumed per training step. On recovery, **RLAX** reinstates both the RNG state and the consumption counters, enabling bitwise-identical prompt selection and ensuring that rollouts generated after recovery follow the exact trajectory of a preemption-free run.

When an ungraceful shutdown or preemption is detected, the controller initiates a deterministic recovery procedure: (1) The controller loads the most recent state snapshot from persistent storage, including model version, prompt-mixture progress, and RNG states. (2) The controller restores its internal bookkeeping structures and reattaches to the parameter server or reloads weights from persistent storage if needed. (3) The controller resumes rollout generation and training precisely from the recorded checkpoint and prompt position, ensuring continuity of both the training trajectory and logged metadata.

## 4 Support for RL Algorithms

**RLAX** adopts a modular design for computing the loss function given a batch of rollout results. This design provides substantial flexibility for algorithmic ablations. At each iteration, we first sample $B$ questions from the dataloader. For each question $q$, we sample a group of $G$ rollouts $\{o_i\}_{i=1}^{G}$ using $\pi_{\theta_{\text{old}}}$ ($\pi_{\theta_{\text{old}}}$ is detached from the current policy $\pi_\theta$ in the on-policy case). Finally, we receive rewards

$\{R_i\}_{i=1}^G$. In a unified view, the objective of modern RL algorithms can be designed as:

$$\mathcal{J}_{\text{unify}}(\theta) = \mathop{\mathbb{E}}_{\substack{q\sim\mathcal{D} \\ \{o_i\}_{i=1}^G\sim\pi_{\theta_{\text{old}}}(\cdot|q)}} \Big[ \sum_{i=1}^{G}\sum_{t=1}^{|o_i|} \mathsf{sg}[\text{Agg}_{i,t}^{\mathcal{X}} \cdot \text{IS}_{i,t}^{\mathcal{X}}] \cdot (\mathsf{sg}[\text{Adv}_{i,t}^{\mathcal{X}}] \cdot \text{GradTerm1}_{i,t}^{\mathcal{X}} + \text{GradTerm2}_{i,t}^{\mathcal{X}}) \Big], \tag{3}$$

where $\mathsf{sg}[]$ is the stopping-gradient operator, and $\mathcal{X} = (\pi_\theta, \pi_{\theta_{\text{old}}}, q, \{o_i, R_i\}_{i=1}^G)$ is the full information packet of the current batch. We explain each component as follows.

**Aggregation Weight.** $\text{Agg}$ is the aggregation weight for each token so that the objective is reasonably bounded. Common choices are individual trajectory average ($\frac{1}{G|o_i|}$), group-level average ($\frac{1}{\sum_{i=1}^G |o_i|}$) and max-length average ($\frac{1}{GL_{\max}}$, where $L_{\max}$ is the maximum generation length).

**Importance Sampling Weight.** IS accounts for the *distribution shift* between the current policy $\pi_\theta$ and the sampling policy $\pi_{\theta_{\text{old}}}$. From a loss function perspective, this term corrects an off-policy loss to an on-policy loss when $\pi_{\theta_{\text{old}}}$ is stale in asynchronous training. From a numerical perspective, a mismatch exists between the true $\pi_{\theta_{\text{old}}}$ and the actual sampled distribution due to approximations in the inference engine [44]. This issue appears even when training is fully synchronous ($\pi_{\theta_{\text{old}}} = \mathsf{sg}[\pi_\theta]$). For more details, refer to §5. For conceptual and engineering rigor, we require that IS does *not* contribute to the gradients. While PPO-like algorithms use the importance weight as the source of gradients, we attribute the gradients to $\text{GradTerm1}$.

**Advantage Estimator.** $\text{Adv}$ distinguishes responses in the "good" direction from those in the "bad" direction. In the **RLAX** design, different advantage estimators (the basic Monte-Carlo estimator, GAE [30], and others) can be easily plugged in.

**Gradient Terms.** $\text{GradTerm1}$ is the main source of gradients, as it is guided by $\text{Adv}$. $\text{GradTerm2}$ accommodates regularization. By explicitly masking out gradient terms with $0$, we can recover the clipping trick in PPO-like algorithms, which also provides a cleaner view of objective design.

### 4.1 Benefits of Modular Design

**Engineering.** We find that modern RL (policy optimization) algorithms for LLMs can be easily instantiated by this modular design. It enables plug-in implementation of each algorithmic component, making algorithm tweaking as easy as a few lines of code change. Any bug fixes or feature changes to any component (e.g., tricks to improve numerical stability for the importance weight) automatically take effect across all training configurations without re-inspecting its occurrences.

On the other hand, we note that while a modular design for advantage estimators is prevalent in modern RL libraries (verl [33], AReaL [15], slime [49], TorchRL [6], SkyRL [8]), other implementation details in the RL objectives are either *redundant* or suffer from low readability due to *inheritance hierarchy*: in these libraries, contributors either totally rewrite a class for a specific RL objective, resulting in redundancy, or inherit from an existing objective class and override its methods, resulting in low readability for both algorithm design and hyperparameter settings. A similar observation is made in AXLearn [22] that many training frameworks has redundancy in neural network layer implementation due to the use of subtyping.

**Perception.** Explicitly separating gradient terms and non-gradient terms provides a clearer view of each component's functionality. For example, although both GRPO (Equation (5)) and CISPO (Equation (7)) clip the importance sampling weights, their influences on the gradients are drastically different. For GRPO, the clipped terms do *not* have gradients, meaning that its clipping is in fact a gradient mask. For CISPO, the clipped terms still have gradients, but with weights confined to the neighborhood around $1$, meaning that its clipping is in fact a regularizer.

### 4.2 Instantiations of Modern RL Algorithms

We now present an example of instantiating modern RL algorithms using $\mathcal{J}_{\text{unify}}(\theta)$ (Equation (3)). For notational ease, we use the same definition as in Equation (7) in [9] to denote the gradient mask

due to PPO-like clipping:

$$M_{i,t} = \begin{cases} 0, & \text{if } \widehat{A}_{i,t} > 0 \text{ and } r_{i,t}(\theta) > 1 + \epsilon_{\text{high}}, \\ 0, & \text{if } \widehat{A}_{i,t} < 0 \text{ and } r_{i,t}(\theta) < 1 - \epsilon_{\text{low}}, \\ 1, & \text{otherwise.} \end{cases} \tag{4}$$

For algorithms with symmetric clipping, $\epsilon_{\text{high}} = \epsilon_{\text{low}} = \epsilon$.

**GRPO.** By Equation (3) in [32], the original GRPO objective is

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathop{\mathbb{E}}_{\substack{q \sim \mathcal{D} \\ \{o_i\}_{i=1}^{G} \sim \pi_{\theta_{\text{old}}}(\cdot|q)}} \left[ \frac{1}{G} \sum_{i=1}^{G} \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \Big( \min \left\{ r_{i,t}(\theta) \widehat{A}_i,\ \mathsf{clip}(r_{i,t}(\theta), 1-\epsilon, 1+\epsilon) \widehat{A}_i \right\} - \beta \widehat{\mathsf{KL}}_t(\pi_\theta || \pi_{\mathsf{ref}}) \Big) \right], \tag{5}$$

where

$$\widehat{A}_i = \frac{R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G})}{\mathsf{std}(\{R_{i'}\}_{i'=1}^{G})},$$

and $\widehat{\mathsf{KL}}_t(\pi_\theta || \pi_{\mathsf{ref}})$ is the K3 estimator for $\mathsf{KL}(\pi_\theta(\cdot|q, o_{i,<t}) || \pi_{\mathsf{ref}}(\cdot|q, o_{i,<t}))$ (c.f., Equation (4) of [32]):

$$\widehat{\mathsf{KL}}_t(\pi_\theta || \pi_{\mathsf{ref}}) = r_{i,t}(\theta) - \log r_{i,t}(\theta) - 1.$$

Thus, we can instantiate Equation (3) with

$$\begin{aligned} \text{Agg}_{i,t}^{\mathcal{X}} &= \frac{1}{G|o_i|}, \\ \text{IS}_{i,t}^{\mathcal{X}} &= 1, \\ \text{Adv}_{i,t}^{\mathcal{X}} &= \frac{R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G})}{\mathsf{std}(\{R_{i'}\}_{i'=1}^{G})}, \\ \text{GradTerm1}_{i,t}^{\mathcal{X}} &= M_{i,t} \cdot r_{i,t}(\theta), \\ \text{GradTerm2}_{i,t}^{\mathcal{X}} &= -\beta(r_{i,t}(\theta) - \log r_{i,t}(\theta) - 1). \end{aligned}$$

DAPO [46] can be instantiated with $\text{Agg}_{i,t}^{\mathcal{X}} = \frac{1}{\sum_{i'=1}^{G} |o_{i'}|}$, removing $\text{GradTerm2}$, and turning on the dynamic sampling switch in §6.2. Dr. GRPO [24] can be instantiated with $\text{Agg}_{i,t}^{\mathcal{X}} = \frac{1}{GL_{\max}}$, where $L_{\max}$ is the maximum generation length, disabling batch normalization, and removing $\text{GradTerm2}$. REINFORCE++ [18] can be instantiated by replacing $\text{Adv}$ with Equations (11) and (12) in [18] and replacing $\text{GradTerm2}$ with the K2 estimator.

We implement the original version of GRPO from [32] in Listing 1.

For other algorithms, refer to Appendix A.

# 5 Addressing Numerical Misalignment between the Inference Workers and the Trainer

There are two sources for divergent trainer and inference worker implementation that can lead to numerical misalignment.

**Source #1: Divergent Parallelism Strategies** The parallelism strategies used for training and inference are often different to optimize for their distinct objectives.

- **Inference:** Typically uses *Tensor Parallelism (TP)* to maximize effective memory bandwidth and reduce latency across multiple devices.

```
loss_func = RLObjective.default_config().set(
    agg=IndividualAggregation.default_config(),
        # DAPO: GroupAggregation.default_config(),
        # Dr. GRPO: MaxLengthAggregation.default_config()
    IS=None,
    adv=AdvantageEstimator.default_config().set(
        batch_mean=True,
        batch_norm=True, # Dr. GRPO: False
    ),
    grad_1=MaskedISGrad.default_config().set(
        mask_func=PPOClipMask.default_config().set(
            eps_high=0.2, # DAPO: 0.28
            eps_low=0.2,
        ),
    ),
    grad_2=K3Grad.default_config().set( # DAPO & Dr. GRPO: None
        beta=0.04,
    ),
).instantiate()
```

Listing 1: Configuration for the original GRPO objective.

- **Training:** Commonly employs more complex hybrid parallelism schemes as combinations of Data Parallelism, FSDP, and Context Parallelism to better overlap computation with communication and scale to larger models.

These different strategies inherently change the order of operations. For example, a parallel summation $\sum(x_1, ...x_n)$ might be computed as $\sum(x_1, ...x_{\frac{n}{2}}) + \sum(x_{\frac{n}{2}+1}, ...x_n)$ in one setup (e.g., TP) but as a single sequential sum in another. Because bfloat16 arithmetic is sensitive to the order of operations, such variations can lead to measurable numerical discrepancies.

**Source #2: Differences in JAX Kernel Fusion** The JAX compilation process also leads to different fused kernels for training and inference.

- **Inference:** A pure inference call triggers only the primal function. The compiler has full freedom to fuse operations aggressively for maximum speed.
- **Training:** A training step triggers both the forward and backward functions. The forward pass must emit intermediate values (residuals) required by the backward pass, which restricts the compiler's ability to fuse operations.

A concrete example is *RMSNorm*. In inference mode, JAX may fuse RMSNorm with subsequent operations (e.g., a residual add), computing the combined block in a highly optimized kernel. During training, however, the forward pass must explicitly output the output of *RMSNorm* for use in the backward pass. This requirement prevents fusion and results in a different kernel decomposition and thus a different computational order compared to the inference path. These subtle, low-level differences in computation order are sufficient to cause numerical divergence in bfloat16, ultimately leading to the unexpected observation $r(\theta) \neq 1.0$.

### 5.1 Mitigation Strategies

Obtaining highly accurate log-probabilities ($\pi_{\theta_{\text{old}}}(a|s)$) is critical for RL training convergence, particularly in off-policy settings where Importance Sampling (IS) is heavily relied upon. To mitigate numerical divergence, we employ log-prob recomputation with training-inference graph alignment, similar to approaches discussed in [44].

To ensure maximal alignment, we do not rely on cached log-probs from inference workers. Instead, we maintain a copy of the old policy that mirrors the exact model parallelism configurations of the trainer and re-run a prefill step to compute $\pi_{\theta_{\text{old}}}$. Crucially, we further force the training model's forward path to match the inference model's computational graph (e.g., kernel fusion patterns) by disabling most activation saving (except for essential layer outputs) in the trainer. This effectively

forces the training forward pass to behave like the inference worker's inference pass. While this heavy rematerialization ("nothing-saveable") combined with recomputation can incur up to some slowdown for the training process, it has negligible impact (less than 10%) for the end-to-end performance. This is because, in our RL workloads, the dominant performance bottleneck is the rollout generation on the inference workers rather than the trainer. Therefore we consider this performance tradeoff necessary to achieve the required numerical consistency. For Mixture of Experts (MoE) models, we have observed an additional source of divergence: the set of activated experts can differ significantly between the old and new policy models, a phenomenon also noted in [47]. We plan to address this in future work by enforcing routing alignment, ensuring the new policy adheres to the expert routing decisions of the old policy during specific training phases.

# 6 Other System Designs and Optimizations

Next, we discuss our verifier service implementation and data curation methods.

## 6.1 Verifiers

In contrast to answer-only math verifiers, which typically take tens of milliseconds and can be run using a single thread locally on the inference workers, code verifiers are more resource-intensive. For example, a high-quality competitive coding problem usually requires code to pass $10 - 100$ test cases, each within a time limit of $1 - 10$ seconds. As is often the case, sub-sampling from the large number of test cases easily leads to false positive results. For such problems, a local code verifier is too slow. To make matters worse, it can easily damage the system if the execution context is not safely isolated.

To make code verification scalable, we designed **OUBLIETTE**, a remote code execution service utilizing the AWS Lambda service. **OUBLIETTE** supports two fundamental operations: executing code (with compilation if necessary) and executing a binary executable directly, both supporting standard or file input/output.

**OUBLIETTE** plays a crucial role in verifying the official Codeforces[1] dataset. After de-duplication and removing problems without English statements or that are interactive [25], the official Codeforces dataset contains $9096$ problems and $449797$ test cases, with uncompressed total size $> 500$ GB. Apart from its extremely large scale, Codeforces further distinguishes itself from other competitive coding datasets by extensively using special checkers [3] for *each* problem. At a high level, a special checker handles problems with multiple feasible solutions and is itself an executable program.

We designed the following code verification pipeline:

**Before training.** Process all data (including compiling the checkers into executable binaries to reduce service-side compilation overhead) and upload them to a *persistent* **OUBLIETTE** repository. We obtain the *links* to each test case. This step only needs to be done *once*, as long as the data itself does not change.

**During training.** When code is being verified, the client sends only the code and *links* to inputs to **OUBLIETTE**. **OUBLIETTE** then executes the code with inputs drawn from the specified links, stores the outputs to a *transient* **OUBLIETTE** repository, and returns the *links* to the outputs. Next, the client sends *links* to the checker binaries, inputs, and answers (all stored in the same repository), as well as *links* to the outputs, to **OUBLIETTE**. Finally, **OUBLIETTE** executes the checker binaries and sends the verdicts (these are very short) generated by the checkers directly to the client.

We describe the efficiency and scalability of **OUBLIETTE** in our use case here. A typical request to **OUBLIETTE** contains $2000$ unique pairs of (code, input) to be executed in parallel. The time limit is set to $10$ seconds per pair. The average verification time, including two rounds of client-server interaction (code and checker), is less than 30 seconds per request. Moreover, there is no upper bound on the maximum concurrency of code runs **OUBLIETTE** can handle, as AWS Lambda is a serverless function service. Thus, there is no bottleneck even when there are large numbers of rollouts due to hitting max generation length in the first batch. The network transaction overhead is also drastically reduced thanks to the persistent repository storing test data. Assuming the average code size is 3 KB

[1] Apple acquired license. Website: `https://codeforces.com`

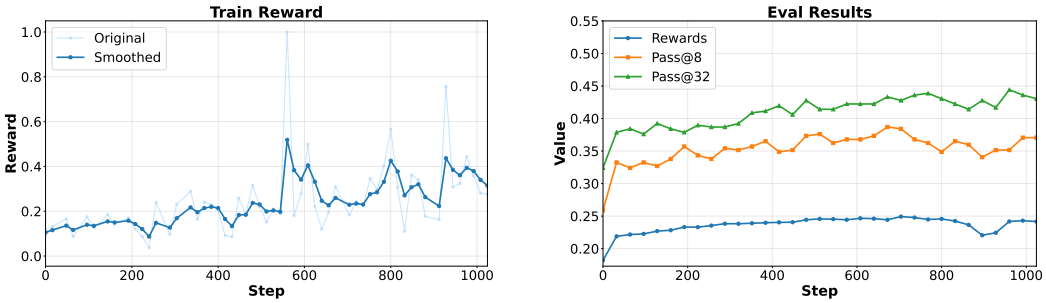


(a) Training rewards over steps

(b) Test set accuracy over steps

Figure 2: Training reward and test set accuracy over steps when training `QwQ-32B` on Codeforces dataset using **RLAX**.

and the average *link* size is 50 B, compared to full data transaction to any remote code execution server, one epoch on the entire Codeforces dataset is reduced from $> \frac{449797 \times 3}{1024 \times 1024} + 500 > 500$ GB to $\frac{449797 \times (3+3 \times 50/1024)}{1024 \times 1024} \approx 1.35$ GB.

### 6.2 Data Curation

Data quality represents a fundamental bottleneck in reinforcement learning, particularly for Group Relative Policy Optimization (GRPO) algorithms. data quality is extremely important for RL tasks, which usually have sparse 0/1 reward signals.

The core challenge lies in avoiding homogeneous rewards in a group that provide no learning signal. When coding problems are either trivially easy or prohibitively difficult, all samples generated from a single problem prompt tend to receive identical rewards—either all correct (reward = 1) or all incorrect (reward = 0). Such uniform distributions eliminate the reward variance necessary for any algorithms estimating advantages using a group of samples, including GRPO, etc. to compute meaningful policy gradients. The optimal training scenario requires a balanced reward distribution where, among $n$ samples generated from a single prompt, some solutions are correct while others contain errors, creating the reward diversity essential for effective learning.

To address this challenge, **RLAX** has an extention to optionally add a data filter to use one of the two data curation strategies as follows:

**Pre-filtering Approach:** This method conducts preliminary evaluation runs on the training dataset, retaining only problems where the model achieves partial success rates. While this maximizes dataset efficiency and enables curriculum learning by deploying problems of varying difficulty at appropriate stages, it introduces manual overhead that can slow the training pipeline.

**Dynamic Filtering Approach:** This strategy implements real-time filtering during training, forwarding trajectories only when batches contain both correct and incorrect samples from the same problem (similar to DAPO [46]). Although this eliminates manual evaluation overhead through automated filtering, it exacerbates computational bottlenecks by discarding trajectory portions.

## 7 Evaluation

We first present our main results demonstrating how **RLAX** improves the accuracy of `QwQ-32B` [37]. We then provide ablation studies on the choice of staleness bounds and the numerical alignment method. Finally, we evaluate **RLAX** 's ability to be preempted and resumed.

### 7.1 Main Results

**RLAX** is capable of training and delivering high-quality models, including both internal and open-source variants. Here, we show **RLAX**'s effectiveness through a case study involving RL training on `QwQ-32B` for competitive coding tasks. We use Codeforces 2013-2024 for training and 2025 for evaluation, with **OUBLIETTE** as the reward model for code execution against test cases. In this

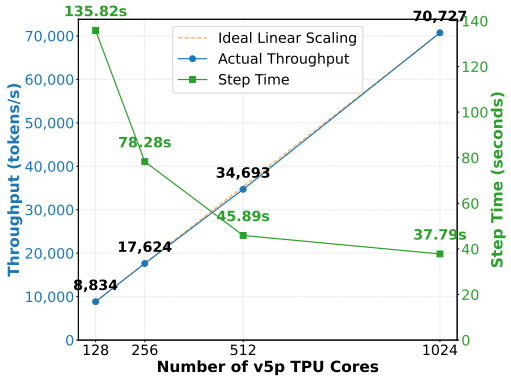


(a) Inference service throughput and training step time.

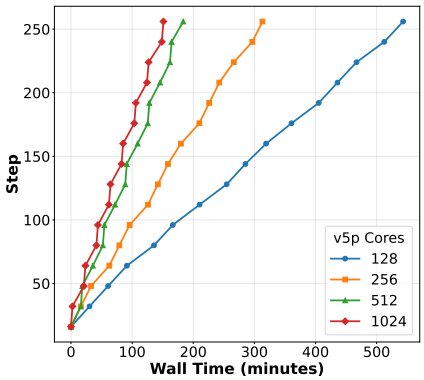


(b) Training steps at different wall time.

Figure 3: TPU Scalability: Rollout throughput and training step time vs core count

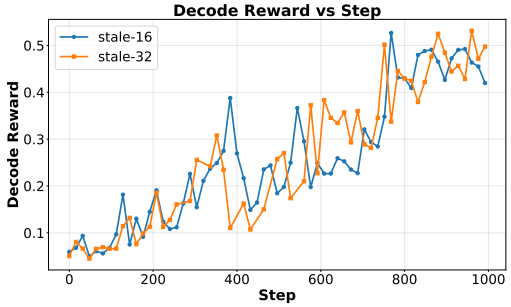


(a) Training rewards vs steps.

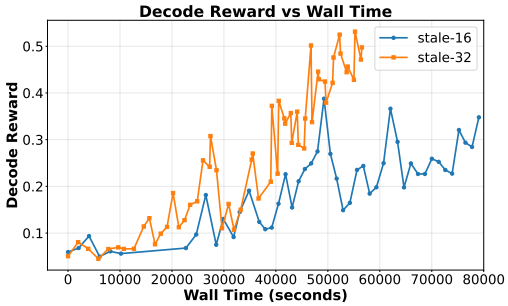


(b) Training rewards vs wall time.

Figure 4: **RLAX**'s training rewards under different staleness settings.

experiment, we deployed v5p-512 clusters for both training and rollout. Figure 2 shows the training rewards and test set accuracy over training steps. Using **RLAX**, we achieved test set accuracy improvements of $6.7\%$, $12.8\%$, and $12.0\%$ for reward, pass@8, and pass@32 respectively. Since `QwQ-32B`'s knowledge cutoff is November 28, 2024 [35], these gains on Codeforces 2025 represent genuine gains in coding capability.

A key objective of our system is to maintain *near-linear throughput scaling* as inference capacity increases. To assess this, we conduct identical RL workloads on TPU v5p inference clusters of size 128, 256, 512, and 1024, each hosting the same model and exposing the same serving interface. We used the same v5p-512 cluster as the trainer to compare generation throughput and training step times across different settings. Figure 3a reports the steady-state inference service throughput and training step latency for each configuration. As the blue curve in Figure 3a shows, the inference throughput grows roughly linearly with the size of the inference service, up to v5p-1024, with minimal deviations from ideal scaling. The system achieves an $8.0\times$ improvement from using v5p-128 to v5p-1024 as the size of the inference service, indicating that the controller introduces very low overhead even at large scale.

The overall training performance increases as we scale out the size of the inference service. The green curve in Figure 3a shows is the training step latency. The training step latency reduces by $3.6\times$, when the inference service's throughput improves by $8.0\times$. This is because, as inference throughput increased, the system became training-bound. Figure 3b shows step times vs wall time from the first 256 steps under each setting: **RLAX**'s training step latency is consistent across training steps.

### 7.2 Ablation Study

**Impact of staleness on convergence and training speed.** We study how different staleness settings affect convergence and training performance. In our main run, we had the inference reload staleness

to be 16 and trainer acceptance staleness to be 32. This means $(j, k) = (16, 32)$. We also tested a configuration where the trainer acceptance staleness was equal to the inference reload staleness $(j, k) = (16, 16)$, effectively causing the trainer and the inference workers to synchronize more frequently and thus wait for each other more often. We compare the training reward curves between these two settings.

Figure 4 shows the results for our experimental settings (Codeforces dataset, `QwQ-32B`, 16k context length). In the figure, stale-16 means $k = 16$, and stale-32 means $k = 32$. Both configurations achieve similar reward gains per step. However, the configuration with greater trainer acceptance staleness (trainer acceptance staleness > inference reload staleness) reduces average step time from 75 to 45 seconds, resulting in higher reward gains per unit of wall time.

However, high staleness can slow reward convergence per step, though not in our main experiments. In other experiments with a smaller model (3B) and fewer samples (10, overfitting scenario), high staleness both slowed convergence and increased time-to-plateau. How staleness bound changes accuracy and performance in RL is workload-dependent and requires validation runs before full experiments.

**Numerical alignment.** To quantify the discrepancy between the trainer policy and the behavior policy, we measure the absolute per-token log-probability difference between them. Let $\pi_\beta$ denote the behavior policy that generated the trajectories and $\pi_\theta$ the current trainer policy being optimized. For each action token, we compute the masked mean of $|\log \pi_\beta(a_t \mid s_t) - \log \pi_\theta(a_t \mid s_t)|$. Figure 5 reports results from two controlled experiments conducted with off-policy learning using the `QwQ-32B` model, comparing training runs with and without our numerical log-probability recomputation method (§5.1). Models trained without correction exhibit markedly more volatile log-probability differences, including sharp spikes that correlate with instability in reward learning. In contrast, enabling our method consistently suppresses these fluctuations and yields smoother, more stable policy updates. Quantitatively, we summarize stability using the 95th percentile of the per-step log-probability difference over 512 training steps: the baseline run without recomputation reaches a 95th percentile of $0.044348$, while our recomputation method reduces this to $0.019943$.

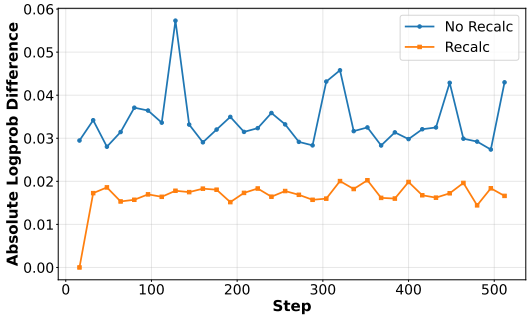


Figure 5: Absolute the Log Probability Difference Between Trainer and Rollout

### 7.3 Supporting Preemption

We evaluate **RLAX** 's robustness under preemption by examining its behavior during real preemptions. Several resource-related preemptions occurred naturally during our staleness ablation experiments (§7.2), providing representative stress tests without requiring synthetic fault injection. Figure 6 shows an example of **RLAX** being preempted. **RLAX** reliably resumed training from the most recent persisted checkpoint after the preemption, with no observable deviation in learning curves or system behavior. This result confirm that the checkpoint-and-restore mechanisms of **RLAX** are effective in practice and that the system can tolerate preemption without compromising training progress.

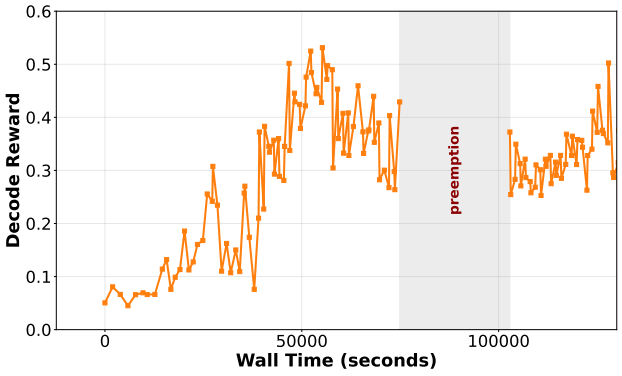


Figure 6: **RLAX**'s training rewards over time with a preemption event.

## 8 Related Work

**Off-policy RL.** There are many possible approaches to implementing off-policy reinforcement learning (RL) for large language models. **RLAX** currently adopts a straightforward, rollout-level off-policy design: each rollout is generated entirely using a single model version, and the trainer is allowed to consume any rollout whose originating model is at most $k$ steps stale. Under this scheme, inference workers must discard partially completed rollouts whenever they refresh their weights from the parameter server. Any continuation would produce rollouts whose model provenance cannot be cleanly attributed to a single version.

Recent work has begun to explore mechanisms that relax this constraint by enabling a single rollout to be generated using multiple model versions. Magistral [26], for example, continues generation using updated model weights while reusing the KV cache produced by older weights, thereby salvaging partial computations that would otherwise be discarded. LlamaRL [40] similarly improves data efficiency through mixed-version rollouts, wherein stale rollouts are prefixed with a fresh prefill computed using the latest weights. These approaches increase hardware utilization and reduce rollout waste, and can be extensions to **RLAX**.

**RL on TPUs.** **RLAX** is a full-stack RL post-training systems designed for TPU clusters, but it is not the only recent effort in this space. Google's Tunix library [5] provides a TPU-based RL fine-tuning framework focused on simplicity and extensibility. However, Tunix currently lacks multi-host distributed training capabilities and relies on standard attention implementations, which makes it less suitable for long-context LLMs. In contrast, **RLAX** is designed for large-scale, multi-host deployments and leverages optimized kernel implementations such as FlashAttention [13, 12] to support efficient training of long-context and high-capacity models.

**Checkpoint-and-restore in ML Systems.** Checkpoint-and-restore mechanisms have long been a standard practice of large-scale ML systems, enabling fault tolerance and preemption support in distributed training. Prior works typically focus on supervised training workloads with tightly synchronized iterations [42, 38]. In contrast, RL post-training introduces unique challenges: rollouts depend on prompt-mixture sampling, verifier outcomes, and model-version provenance, all of which must be kept consistent across restarts. Naively checkpointing all state, including large transient artifacts such as KV caches or intermediate verifier states, would result in unnecessarily large and slow-to-load checkpoints.

By tightly integrating deterministic controller logic, reproducible RNG streams, and dual in-memory/persistent storage of model weights, **RLAX** extends traditional checkpoint-and-restore techniques to the RL setting. This allows **RLAX** to guarantee correctness under frequent preemption, a property that is critical for large-scale RL workloads running on shared or elastic compute clusters.

## 9 Conclusion

We present **RLAX**, a scalable and distributed reinforcement learning framework for large language models (LLMs) that unifies a broad spectrum of RL algorithms and training paradigms within a single, flexible system architecture. Designed for practical large-scale deployment on TPUs, **RLAX** provides efficient rollout generation, high-throughput training, and built-in support for preemption and recovery. In large-scale evaluations, **RLAX** improves `QwQ-32B` 's pass@8 accuracy by 12.8% in just 12 hours and 48 minutes on 1024 v5p TPUs, and it does so while remaining robust to multiple unexpected preemptions throughout training.

## Acknowledgment

We thank Liang He, Meng Li, Jing Lu, Chang Liu, Mark Lee, Tom Gunter, Nan Du, Qibin Chen, John Peebles, Xianzhi Du and Bailin Wang for their earlier contributions to the system. We also thank Benoit Dupin and Daphne Luong for their leadership and our many colleagues for supporting the system development and providing valuable feedback.

## References


[1] Xla: Optimizing compiler for machine learning. `https://openxla.org/xla/tf2xla`, 2025.

[2] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98):1–76, 2021.

[3] Maxim Akhmedov. Checkers with testlib.h. `https://codeforces.com/blog/entry/18431`, 2025.

[4] Anthropic. Introducing Claude 4. `https://www.anthropic.com/news/claude-4`, 2025.

[5] Tianshu Bao, Lance Wang, Abheesht Sharma, Jiwon Shin, Ann Yan, Sizhi Tan, Haoyu Gao, Jen Ha, Lin Chai, Dangyi Liu, Rakesh Iyer, Mridul Sahu, et al. Tunix. `https://github.com/google/tunix`, 2025.

[6] Albert Bou, Matteo Bettini, Sebastian Dittert, Vikash Kumar, Shagun Sodhani, Xiaomeng Yang, Gianni De Fabritiis, and Vincent Moens. Torchrl: A data-driven decision-making library for pytorch, 2023.

[7] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[8] Shiyi Cao, Sumanth Hegde, Dacheng Li, Tyler Griggs, Shu Liu, Eric Tang, Jiayi Pan, Xingyao Wang, Akshay Malik, Graham Neubig, Kourosh Hakhamaneshi, Richard Liaw, Philipp Moritz, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. Skyrl-v0: Train real-world long-horizon agents via reinforcement learning, 2025.

[9] Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan, Changqing Yu, Chao Wang, Cheng Zhu, et al. Minimax-m1: Scaling test-time compute efficiently with lightning attention. *arXiv preprint arXiv:2506.13585*, 2025.

[10] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating llms by human preference, 2024.

[11] Gheorghe Comanici et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities, 2025.

[12] Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*, 2024.

[13] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

[14] DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025.

[15] Wei Fu, Jiaxuan Gao, Xujie Shen, Chen Zhu, Zhiyu Mei, Chuyi He, Shusheng Xu, Guo Wei, Jun Mei, Jiashu Wang, Tongkai Yang, Binhang Yuan, and Yi Wu. Areal: A large-scale asynchronous reinforcement learning system for language reasoning, 2025.

[16] Google. TensorStore. `https://google.github.io/tensorstore/`, 2024. Accessed: [Nov 7, 2025].

[17] Horace He and Thinking Machines Lab. Defeating nondeterminism in llm inference. *Thinking Machines Lab: Connectionism*, 2025. https://thinkingmachines.ai/blog/defeating-nondeterminism-in-llm-inference/.

[18] Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. Reinforce++: An efficient rlhf algorithm with robustness to both prompt and reward models. *arXiv preprint arXiv:2501.03262*, 2025.

[19] Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code, 2024.

[20] Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 REINFORCE samples, get a baseline for free!, 2019.

[21] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, SOSP '23, page 611–626, New York, NY, USA, 2023. Association for Computing Machinery.

[22] Mark Lee, Tom Gunter, Chang Lan, John Peebles, Hanzhi Zhou, Kelvin Zou, Sneha Bangalore, Chung-Cheng Chiu, Nan Du, Xianzhi Du, Philipp Dufter, Ruixuan Hou, Haoshuo Huang, Dongseong Hwang, Xiang Kong, Jinhao Lei, Tao Lei, Meng Li, Li Li, Jiarui Lu, Zhiyun Lu, Yiping Ma, David Qiu, Vivek Rathod, Senyu Tong, Zhucheng Tu, Jianyu Wang, Yongqiang Wang, Zirui Wang, Floris Weers, Sam Wiseman, Guoli Yin, Bowen Zhang, Xiyou Zhou, Danyang Zhuo, Cheng Leong, and Ruoming Pang. AXLearn: Modular Large Model Training on Heterogeneous Infrastructure, 2025.

[23] Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, et al. Spider 2.0: Evaluating language models on real-world enterprise text-to-sql workflows. *arXiv preprint arXiv:2411.07763*, 2024.

[24] Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*, 2025.

[25] Mike Mirzayanov. Interactive Problems: Guide for Participants. `https://codeforces.com/blog/entry/45307`, 2025.

[26] Mistral-AI, :, Abhinav Rastogi, Albert Q. Jiang, Andy Lo, Gabrielle Berrada, Guillaume Lample, Jason Rute, Joep Barmentlo, Karmesh Yadav, Kartik Khandelwal, Khyathi Raghavi Chandu, Léonard Blier, Lucile Saulnier, Matthieu Dinot, Maxime Darrin, Neha Gupta, Roman Soletskyi, Sagar Vaze, Teven Le Scao, Yihan Wang, Adam Yang, Alexander H. Liu, Alexandre Sablayrolles, Amélie Héliou, Amélie Martin, Andy Ehrenberg, Anmol Agarwal, Antoine Roux, Arthur Darcet, Arthur Mensch, Baptiste Bout, Baptiste Rozière, Baudouin De Monicault, Chris Bamford, Christian Wallenwein, Christophe Renaudin, Clémence Lanfranchi, Darius Dabert, Devon Mizelle, Diego de las Casas, Elliot Chane-Sane, Emilien Fugier, Emma Bou Hanna, Gauthier Delerce, Gauthier Guinet, Georgii Novikov, Guillaume Martin, Himanshu Jaju, Jan Ludziejewski, Jean-Hadrien Chabran, Jean-Malo Delignon, Joachim Studnia, Jonas Amar, Josselin Somerville Roberts, Julien Denize, Karan Saxena, Kush Jain, Lingxiao Zhao, Louis Martin, Luyu Gao, Lélio Renard Lavaud, Marie Pellat, Mathilde Guillaumin, Mathis Felardos, Maximilian Augustin, Mickaël Seznec, Nikhil Raghuraman, Olivier Duchenne, Patricia Wang, Patrick von Platen, Patryk Saffer, Paul Jacob, Paul Wambergue, Paula Kurylowicz, Pavankumar Reddy Muddireddy, Philomène Chagniot, Pierre Stock, Pravesh Agrawal, Romain Sauvestre, Rémi Delacourt, Sanchit Gandhi, Sandeep Subramanian, Shashwat Dalal, Siddharth Gandhi, Soham Ghosh, Srijan Mishra, Sumukh Aithal, Szymon Antoniak, Thibault Schueller, Thibaut Lavril, Thomas Robert, Thomas Wang, Timothée Lacroix, Valeriia Nemychnikova, Victor Paltz, Virgile Richard, Wen-Ding Li, William Marshall, Xuanyu Zhang, and Yunhao Tang. Magistral, 2025.

[27] OpenAI. Introducing OpenAI o3 and o4-mini. `https://openai.com/index/introducing-o3-and-o4-mini/`, 2025.

[28] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc.

[29] Penghui Qi, Zichen Liu, Xiangxin Zhou, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Defeating the training-inference mismatch via fp16, 2025.

[30] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.

[31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.

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

[33] Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv: 2409.19256*, 2024.

[34] Vaishnavi Shrivastava, Ahmed Awadallah, Vidhisha Balachandran, Shivam Garg, Harkirat Behl, and Dimitris Papailiopoulos. Sample more to think less: Group filtered policy optimization for concise reasoning. *arXiv preprint arXiv:2508.09726*, 2025.

[35] LLM Stats. Qwen3-30b-a3b vs qwq-32b model comparison. https://llm-stats.com/models/compare/qwen3-30b-a3b-vs-qwq-32b.

[36] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In S. Solla, T. Leen, and K. Müller, editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.

[37] Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025.

[38] John Thorpe, Pengzhan Zhao, Jonathan Eyolfson, Yifan Qiao, Zhihao Jia, Minjia Zhang, Ravi Netravali, and Guoqing Harry Xu. Bamboo: Making preemptible instances resilient for affordable training of large DNNs. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 497–513, Boston, MA, April 2023. USENIX Association.

[39] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

[40] Bo Wu, Sid Wang, Yunhao Tang, Jia Ding, Eryk Helenowski, Liang Tan, Tengyu Xu, Tushar Gowda, Zhengxing Chen, Chen Zhu, Xiaocheng Tang, Yundi Qian, Beibei Zhu, and Rui Hou. Llamarl: A distributed asynchronous reinforcement learning framework for efficient large-scale llm training, 2025.

[41] xAI. Grok 4. https://x.ai/news/grok-4, 2025.

[42] Wencong Xiao, Romil Bhardwaj, Ramachandran Ramjee, Muthian Sivathanu, Nipun Kwatra, Zhenhua Han, Pratyush Patel, Xuan Peng, Hanyu Zhao, Quanlu Zhang, Fan Yang, and Lidong Zhou. Gandiva: Introspective cluster scheduling for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 595–610, Carlsbad, CA, October 2018. USENIX Association.

[43] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025.

[44] Feng Yao, Liyuan Liu, Dinghuai Zhang, Chengyu Dong, Jingbo Shang, and Jianfeng Gao. Your efficient rl framework secretly brings you off-policy rl training, August 2025.

[45] Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. Orca: A distributed serving system for Transformer-Based generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 521–538, Carlsbad, CA, July 2022. USENIX Association.

[46] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.

[47] Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. Group sequence policy optimization, 2025.

[48] Yinmin Zhong, Zili Zhang, Xiaoniu Song, Hanpeng Hu, Chao Jin, Bingyang Wu, Nuo Chen, Yukun Chen, Yu Zhou, Changyi Wan, Hongyu Zhou, Yimin Jiang, Yibo Zhu, and Daxin Jiang. Streamrl: Scalable, heterogeneous, and elastic rl for llms with disaggregated stream generation, 2025.

[49] Zilin Zhu, Chengxing Xie, Xin Lv, and slime Contributors. slime: An llm post-training framework for rl scaling. https://github.com/THUDM/slime, 2025. GitHub repository. Corresponding author: Xin Lv.

# A More Instantiations of Modern RL Algorithms

**REINFORCE.** First introduced in [39], REINFORCE has many variations. An on-policy REINFORCE performs one-step gradient update using the objective in Equation (6), where $\widehat{A}_i = R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G})$ is the unbiased advantage estimator.

Thus, we can instantiate Equation (3) with

$$\begin{aligned}
\mathrm{Agg}_{i,t}^{\mathcal{X}} &= \mathbb{1}[t = |o_i|] \cdot \frac{1}{G}, \\
\mathrm{IS}_{i,t}^{\mathcal{X}} &= r_t(\theta), \\
\mathrm{Adv}_{i,t}^{\mathcal{X}} &= R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G}), \\
\mathrm{GradTerm1}_{i,t}^{\mathcal{X}} &= \mathbb{1}[t = |o_i|] \cdot \sum_{t'=1}^{|o_i|} \log \pi_\theta(o_{i,t'}|q, o_{i,<t'}), \\
\mathrm{GradTerm2}_{i,t}^{\mathcal{X}} &= 0,
\end{aligned}$$

A token-level REINFORCE

$$\mathcal{J}_{\text{token-REINFORCE}}(\theta) = \mathop{\mathbb{E}}_{\substack{q\sim\mathcal{D} \\ \{o_i\}_{i=1}^{G}\sim\pi_{\theta_{\text{old}}}(\cdot|q)}} \left[ \frac{1}{GL_{\max}} \sum_{i=1}^{G} \sum_{t=1}^{|o_i|} \mathsf{sg}[r_{i,t}(\theta)] \widehat{A}_i \log \pi_\theta(o_{i,t}|q, o_{i,<t}) \right]. \quad (6)$$

can be instantiated with

$$\begin{aligned}
\mathrm{Agg}_{i,t}^{\mathcal{X}} &= \frac{1}{GL_{\max}}, \\
\mathrm{IS}_{i,t}^{\mathcal{X}} &= r_{i,t}(\theta), \\
\mathrm{Adv}_{i,t}^{\mathcal{X}} &= R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G}), \\
\mathrm{GradTerm1}_{i,t}^{\mathcal{X}} &= \log \pi_\theta(o_{i,t}|q, o_{i,<t}), \\
\mathrm{GradTerm2}_{i,t}^{\mathcal{X}} &= 0,
\end{aligned}$$

where $L_{\max}$ is the maximum generation length. The leave-one-out advantage estimator [20] can be instantiated with $\mathrm{Adv}_{i,t}^{\mathcal{X}} = \frac{G-1}{G}(R_i - \frac{1}{G-1}\sum_{j\neq i} R_j)$. Regularization (such as Equation (12) in [2]) can be instantiated by setting $\mathrm{GradTerm2}$ accordingly. LlamaRL [40] can be instantiated with $\mathrm{IS}_{i,t}^{\mathcal{X}} = \min\{r_{i,t}(\theta), \rho\}$.

**CISPO.** Equation (4) in [9]:

$$\mathcal{J}_{\text{CISPO}}(\theta) = \mathop{\mathbb{E}}_{\substack{q\sim\mathcal{D} \\ \{o_i\}_{i=1}^{G}\sim\pi_{\theta_{\text{old}}}(\cdot|q)}} \left[ \frac{1}{\sum_{i'=1}^{G}|o_{i'}|} \sum_{i=1}^{G} \sum_{t=1}^{|o_i|} \mathsf{sg}[\widehat{r}_{i,t}(\theta)] \widehat{A}_i \log \pi_\theta(o_{i,t}|q, o_{i,<t}) \right], \quad (7)$$

where

$$\widehat{r}_{i,t}(\theta) = \mathsf{clip}(r_{i,t}(\theta), 1-\epsilon_{\text{low}}^{\text{IS}}, 1+\epsilon_{\text{high}}^{\text{IS}}), \quad \widehat{A}_i = \frac{R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G})}{\mathsf{std}(\{R_{i'}\}_{i'=1}^{G})}.$$

Thus, we can instantiate Equation (3) with

$$\begin{aligned}
\mathrm{Agg}_{i,t}^{\mathcal{X}} &= \frac{1}{\sum_{i'=1}^{G}|o_{i'}|}, \\
\mathrm{IS}_{i,t}^{\mathcal{X}} &= \mathsf{clip}(r_{i,t}(\theta), 1-\epsilon_{\text{low}}^{\text{IS}}, 1+\epsilon_{\text{high}}^{\text{IS}}), \\
\mathrm{Adv}_{i,t}^{\mathcal{X}} &= \frac{R_i - \mathsf{mean}(\{R_{i'}\}_{i'=1}^{G})}{\mathsf{std}(\{R_{i'}\}_{i'=1}^{G})}, \\
\mathrm{GradTerm1}_{i,t}^{\mathcal{X}} &= \log \pi_\theta(o_{i,t}|q, o_{i,<t}), \\
\mathrm{GradTerm2}_{i,t}^{\mathcal{X}} &= 0.
\end{aligned}$$

[9] also stated that $\mathrm{GradTerm1}$ could be masked with $M_{i,t}$.