KEVIN: MULTI-TURN RL FOR GENERATING CUDA KERNELS

Anonymous authorsPaper under double-blind review

ABSTRACT

Writing GPU kernels is a challenging task and critical for AI systems' efficiency. It is also highly iterative: domain experts write code and improve performance through execution feedback. Moreover, it presents verifiable rewards like correctness and speedup, making it a natural environment to apply Reinforcement Learning (RL). To explicitly incorporate the iterative nature of this process into training, we develop a flexible multi-turn RL recipe that addresses unique challenges encountered in real-world settings, such as learning from long trajectories and effective reward attribution across turns. We present Kevin the Kernel Writer, the first model trained with multi-turn RL for CUDA kernel generation and optimization. In our evaluation setup, Kevin shows significant gains over its base model (QwQ-32B), improving correctness of generated kernels (in pure CUDA) from 56% to 82% and mean speedup from 0.53x to 1.10x of baseline (PyTorch Eager), and surpassing frontier models like o4-mini (0.78x). Finally, we study its behavior across test-time scaling axes: we found scaling serial refinement more beneficial than parallel sampling. In particular, when given more refinement turns, Kevin shows a higher rate of improvement.

1 Introduction

Writing efficient GPU kernels (Dao et al., 2022; Zhao et al., 2025; Ye et al., 2025) in domain-specific languages: CUDA, Triton, ThunderKittens, CUTLASS, etc. (Nickolls et al., 2008; Tillet et al., 2019; Spector et al., 2024; NVIDIA Corporation, 2025) is critical for enabling AI systems' efficiency at scale, yet it remains difficult and costly due to the deep domain expertise required. This has led to a surge of interest in exploring how Large Language Models (LLMs) could help generate GPU kernels (Ouyang et al., 2025; Li et al., 2025; NVIDIA, 2025) using agentic systems (Damani et al., 2024; Chen et al., 2025; METR, 2025; Lange et al., 2025; Google DeepMind, 2025) that leverage extensive test-time compute. These inference-based approaches are inherently limited by the base models' capability in this domain. On the other hand, the presence of verifiable rewards in the form of correctness and speedup against a reference implementation makes reinforcement learning (RL) a natural approach. This leads to our investigation: *How can we train a model using RL to solve the real-world engineering task of CUDA kernel generation?*

GPU kernel generation emphasizes not just functional correctness, but more importantly performance — distinguishing this code optimization problem from binary-reward tasks that involve passing unit tests (Jimenez et al., 2024) or producing an acceptable proof (Zheng et al., 2022). Since speedup is a continuous goal, performance engineers take an iterative approach: they conduct many rounds of optimization based on previous kernel code, its execution result, and timing profiles. Hence, arriving at an optimized solution relies on multiple turns conditioned on previous execution feedback. In contrast, popular RL methods to train LLMs on verifiable rewards (Shao et al., 2024; Lambert et al., 2025) rely on the outcome reward of a single turn ("single-turn RL training"). We hypothesize that explicitly incorporating successive turns of code generation, execution, and feedback into each RL training step ("multi-turn RL training") better mirrors the iterative nature of kernel development, helping the model to learn more advanced code generation strategies that span multiple refinement turns.

We design a simple yet effective multi-turn RL training recipe, shown in Figure 1, that addresses the *key challenges* of training for CUDA kernel generation and optimization:

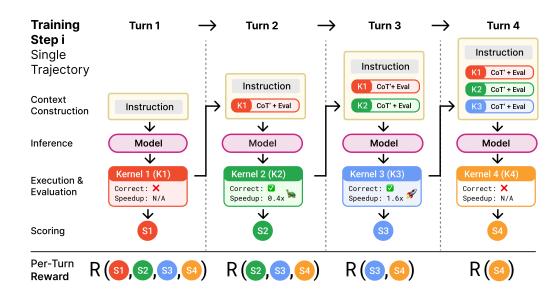


Figure 1: Within each training step, the model iteratively generates, executes, and refines kernels over multiple turns. Kernels are rewarded individually, based both on their performance and their contribution to subsequent speedups: K1, for example, while incorrect, leads to both a correct, slow kernel, K2, and a correct, performant kernel, K3, and should thus be rewarded accordingly. This setup enables Kevin to learn advanced code generation strategies that span multiple turns. Note: CoT' is the summarized chain of thought (CoT).

- 1. Long trajectories lead to sparse rewards and context explosion. To improve sample efficiency, we split trajectories and use each turn as an individual training sample. To address context explosion from long CoTs while preserving reasoning information, we summarize CoTs of prior turns.
- Finding an optimal solution may require rewarding suboptimal kernels that eventually lead to more performant ones. Therefore, we study approaches to aggregate intermediate rewards across turns, finding a configuration that balances the correctness-performance trade-off.
- 3. Reward hacking is prevalent as kernel generation is an open-ended, real-world engineering task: e.g. the model can trick the evaluation harness, lazily copying the reference implementation instead of actually implementing kernels. To prevent this, we analyze the model's failure modes and enforce strict rule-based checks.

Enabled by our multi-turn RL training method on 180 KernelBench tasks from Level 1 and 2, we present Kevin the Kernel Writer, the first RL-trained model to generate CUDA kernels. We compare Kevin and other models in our evaluation setting on a representative KernelBench eval set. Kevin improves upon its base model QwQ-32B, (Team, 2025d) both in correctness ($56\% \rightarrow 82\%$) and mean speedup of generated kernels (in pure CUDA): from 0.53x to 1.10x over PyTorch Eager, while outperforming frontier models like OpenAI o4-mini (0.78x).

We then study the characteristics of Kevin in a test-time scaling setting, comparing it to a single-turn RL-baseline. We systematically compare the benefits of scaling along two axes of test-time compute: sequentially with more refinement turns (Ehrlich et al., 2025; Wang et al., 2025a) or in parallel with more trajectories (Brown et al., 2024; Snell et al., 2024). In our setting, we find that sequential scaling is much more effective, highlighting the importance of iterating upon execution feedback. We observe that the model trained with multi-turn RL exhibits better scaling characteristics with more refinement turns, compared to the base model and the single-turn RL baseline. Our core contributions include:

1. We design an effective yet flexible multi-turn RL training strategy that significantly improves model's capability on CUDA kernel generation. This strategy addresses

challenges that arise in real-world settings, and may be applicable to other environments that benefit from iterative optimizations.

- 2. We found multi-turn is more effective both for training and inference through systematic ablations: the multi-turn trained model outperforms the single-turn trained model across different evaluation setups. Furthermore, we found multi-turn inference is more effective across both models under a fixed inference budget.
- 3. **Kevin exhibits strong test-time scaling trends on both serial and parallel axes**, with a faster rate of improvement than its single-turn RL counterpart and its base model, while maintaining exploration capacity.

2 BACKGROUND AND RELATED WORK

2.1 LLM FOR GPU KERNEL OPTIMIZATION

There has been a surge of interest in exploring how to leverage LLMs to generate GPU kernels (NVIDIA, 2025), driven by the high cost and the long engineering cycles required to develop them (e.g. 2 years for efficient FlashAttention (Dao, 2023) port after Hopper GPU release). However, frontier models underperform on representative benchmarks like KernelBench (Ouyang et al., 2025) and TritonBench (Li et al., 2025), likely due to GPU code being severely underrepresented in the training data (CUDA, for example, accounts for less than 0.1% of pretraining data in the Stack (Kocetkov et al., 2022; Li et al., 2023)). Collecting more expert-written code is expensive, as only a limited number of developers are able to implement high-quality kernels. To tackle this task, there has been an explosion of agentic systems (Damani et al., 2024; Chen et al., 2025; METR, 2025) with custom workflows and evolutionary search methods (Lange et al., 2025; Google DeepMind, 2025). Yet these approaches typically incur high inference cost — e.g. \$15 per kernel (Lange et al., 2025). Improving the base LLM's kernel-generation ability is therefore essential — and could significantly boost the efficiency for downstream agentic workflows.

2.2 RL OPTIMIZATION FOR LLMs TARGETING VERIFIABLE DOMAINS

Reinforcement Learning techniques like GRPO (Shao et al., 2024) have been shown to significantly enhance LLMs' performance on verifiable domains (Lambert et al., 2025) such as math (Team, 2025c; Wang et al., 2025b) and competitive programming (Team, 2025d; Luo et al., 2025a;b). These approaches can be further adapted for real-world software tasks, using fine-grain unit tests (Liu et al., 2023) or comparisons between code edits (Wei et al., 2025) as outcome rewards. Existing methods for code optimizations — where objective concerns performance beyond correctness — have been largely confined to supervised fine-tuning (Waghjale et al., 2024) and imitation learning (Shypula et al., 2024), highlighting Kevin's RL approach a novel contribution for this setting.

Given that tasks like performance optimization or long-horizon planning require multiple sequences of interrelated actions, several works (Goldie et al., 2025; Cao et al., 2025; Wang et al., 2025c; Zhou et al., 2024; Zhuang* et al., 2025) have explored RL training for multi-turn optimizations beyond optimizing for outcome from a single turn. Specific for the code setting, RLEF (Gehring et al., 2025) frames code generation as a multi-turn RL task: the model is allowed a fixed number of refinements turns and assigned a single binary pass/fail reward for final generation — training with such an approach might present sample-inefficiency issues. Unlike RLEF, which assigns rewards only at the final turn, our multi-turn RL framework for Kevin trains on every turn regardless of how optimal the code is, and optimizes for performance beyond just correctness. It is worth noting that Kevin's multi-turn RL training could be viewed as a variant of Meta-Learning (Xiang et al., 2025; Duan et al., 2016) or In-Context Reinforcement Learning (Nie et al., 2024; Tajwar et al., 2025; Schmied et al., 2025), where the focus is to improve solution quality during test-time with feedback (Qu et al., 2025); but adapted in a novel way to the challenging real-world setting of GPU kernel generation and code optimization.

3 TASK AND BASELINE

3.1 Environment and Evaluation

We use KernelBench (Ouyang et al., 2025), a popular dataset for evaluating the LLMs' ability to generate CUDA kernels for deep learning workloads in PyTorch. We chose 180 of both 100 Level 1 problems (basic operators: convolutions, matrix multiplies, loss functions, etc.) and 100 Level 2 problems (sequences of operators with fusion opportunities) as training environments. Since current KernelBench does not provide a train-test split, we construct 80 additional novel tasks following the same methodology (see Appendix A). We build the evaluation set by combining our 80 newly created tasks with the 20 remaining original KernelBench tasks, for a total of 100 evaluation tasks.

Each KernelBench task consists of generating a CUDA kernel given a PyTorch reference implementation, which is used to evaluate correctness and speedup. In our setup, we evaluate the model-generated kernels as follows: we verify the output is in the correct format (ensure resultant code is only implemented with inline CUDA) and check for reward hacking (Section 6.2). We then evaluate the kernel for compilation, runtime errors, and correctness. If the implementation is correct, we profile the kernel for its runtime.

3.2 Kernel Score Design

As we are concerned both with correctness and speedup, we assign a score S for each kernel evaluation result that effectively balances the correctness-performance trade-off.

$$S = 0.3 \cdot \mathbf{1}_{\{\text{correct}\}} + \frac{T_{\text{baseline}}}{T_{\text{kernel}}} \cdot \mathbf{1}_{\{\text{correct}\}}$$

Correctness is checked against the reference program when tested with randomized inputs; speedup is computed as the ratio between PyTorch baseline time and kernel runtime. We experimented with various weights of correctness and speedup, finding this configuration through ablations on models ranging from 7B to 32B.

In addition, we explored rewarding intermediate objectives (successfully compile or execute), yet this caused model to over-optimize for intermediate steps (e.g. generating kernels that only compile, but are not necessarily correct). We also experimented with a length penalty on the response, as suggested by Team (2025b), but found that it degrades our model's performance during training.

3.3 SINGLE-TURN TRAINING

We apply GRPO (Shao et al., 2024) to train the model on kernel generation without iterating on external feedback ("single-turn" training). In each training step, we sample 16 responses per task and assign the evaluated score as the reward for each kernel. We compute the GRPO loss according to (Shao et al., 2024), which updates the policy by maximizing the following objective:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q)]$$

$$\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left\{ \min \left[\frac{\pi_{\theta}(o_{i,t}|q,o_{i,
where $\hat{A}_{i,t} = \frac{r_i - \operatorname{mean}(\mathbf{r})}{\operatorname{std}(\mathbf{r})}$, and r_i is the score of a specific kernel.$$

We choose Qwen QwQ-32B (Team, 2025d) as base model. See Appendix B.6 for the rationale.

Following Yu et al. (2025), we apply Clip-Higher. We sample with temperature =0.9 for both training and inference. We set the KL coefficient to 0 to allow the model to deviate freely from the base policy, following Luo et al. (2025a).

We observe that reward plateaus after 100 gradient steps, likely because single-turn training prevents the model from refining its kernels. Many generated kernels are nearly correct—often a syntax or compilation fix away—but still receive 0 reward, discouraging the model from producing them. Similarly, the correct kernels do not achieve high speedup, as the model optimizes for correctness rather than attempting a risky approach. We address these limitations through multi-turn training.

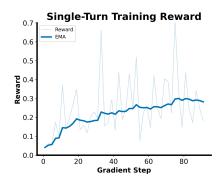


Figure 2: **Reward plateaus during single-turn training.**

4 MULTI-TURN TRAINING

In each multi-turn training step:

- 1. For each task, we sample m parallel trajectories with n refinement turns. To improve sample efficiency, each refinement turn (CoT + response) in a trajectory becomes a single training sample. The response of the model after the CoT consists of a kernel and a CoT summary.
- 2. We construct the context of a sample by including the history of previous responses, which include generated kernels along with their summarized CoTs, and evaluation feedback.
- 3. We evaluate the generated kernel and compute its score as shown in Section 3.2. The reward of each turn (CoT + response) is the discounted sum of current and subsequent scores, which we elaborate in Section 4.3.
- 4. For each task, we normalize the rewards across the *mn* samples for advantage calculation. Then we compute the GRPO loss over the entire batch.

4.1 Managing Context

Reasoning models generate long CoTs, especially for complex tasks like kernel generation. Including all CoTs causes the context to grow rapidly, reaching 50-100k tokens within a few turns, surpassing the model's context length. To prevent context explosion, we discard CoTs of previous turns; yet to preserve information regarding the reasoning process, we ask the model to summarize the changes applied. This summary, along with the generated kernels and evaluation results, is passed to subsequent turns.

4.2 Training On Every Refinement Turn

In a naive implementation, each n-turn trajectory is a single training sample. To improve sample efficiency, we split a n-turn trajectory into n training samples, each corresponding to the kernel + CoT summary of a refinement turn with the context containing the history. Hence, the kernel and CoT summary receives the reward of that particular turn.

4.3 REWARD AGGREGATION AND DISCOUNTING

We initially explored two naive strategies for multi-turn credit assignment. The greedy approach assigns to each turn its corresponding kernel score, while the outcome-based approach assigns to all turns the best score in the trajectory. The former failed to reward early suboptimal turns that lead to performant kernels later, while the latter ignores individual contributions and is sample inefficient.

Our method balances both approaches by aggregating the future kernels scores with a discount factor. We conduct ablations on the reward formulation. For score aggregation, we can either take the sum $R_t = \sum_{i=t}^T \gamma^{i-t} r_i$ or maximum $R_t = \max_{i=t,...,T} \left\{ \gamma^{i-t} r_i \right\}$ over future scores. Sum favors

generating multiple good kernels, while max prioritizes achieving one high-performing kernel. We evaluate both forms with $\gamma=0.4$ and $\gamma=0.8$.

Experiments show that sum with $\gamma=0.4$ scales best over 8 turns, though max performs better with $\gamma=0.8$ with fewer turns. We decide to use the sum reward formulation with discount factor $\gamma=0.4$.

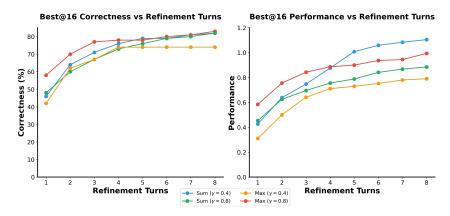


Figure 3: Sum with $\gamma = 0.4$ exhibits the best scaling behavior. We evaluate models trained with different reward formulations under 16 parallel trajectories and 8 refinement turns.

4.4 MULTI-TURN TRAINING BEHAVIOR

For multi-turn ablations and training runs, we train to 80 gradient steps; within each step, for each task, we sample 16 parallel trajectories and conduct 4 refinement turns. Each batch contains 8 tasks. (See Appendices B.5 for detailed hyperparameters and C.1 for training statistics)

Unlike single-turn training, reward now steadily increases. We also observe response-length behaviors similar to Luo et al. (2025b): the response length initially decreases, and then it starts increasing again as the model attempts more sophisticated solutions; we extend the max response length from 16K to 22K tokens at gradient step 60.

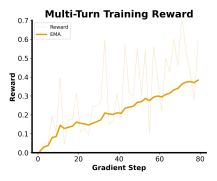


Figure 4: **Reward climbs steadily for multi-turn training.**

5 EVALUATION

As kernel generation is a challenging task, models are often given extensive test-time compute to tackle it. At inference, we employ multiple parallel trajectories, each made up of several serial turns.

We mark a trajectory **correct** if it contains at least one correct kernel. Its **performance** is the speedup of the fastest kernel (within the trajectory) over the PyTorch Eager reference (speedup of 0x if no kernel is correct). We also consider the \mathbf{fast}_p metric, introduced by Ouyang et al. (2025), which is a binary indicator for whether a trajectory contains a correct kernel with performance of p or more. To aggregate a metric across k parallel trajectories for a given task, we compute: $\mathbf{best}@\mathbf{k}$, the maximum for that metric across all trajectories; $\mathbf{avg}@\mathbf{k}$, the average value across trajectories.

5.1 RESULT ON KERNELBENCH EVAL SET

We compare Kevin against frontier models and the single-turn RL baseline on our aforementioned KernelBench eval set of 100 tasks (Section 3.1), with 16 parallel trajectories, 8 serial refinement turns. As shown in Table 1, Kevin achieves a higher performance than its single-turn trained counterpart and other frontier models, demonstrating significant improvement from its base model (QwQ-32B). Qualitatively, Kevin is able to more effectively implement more aggressive optimizations across several turns (see Appendix H for examples); see Appendix E for additional evaluation details.

Model	Correctness		Performance		\mathbf{fast}_1		$\mathbf{fast}_{1.5}$	
	best@16	avg@16	best@16	avg@16	best@16	avg@16	best@16	avg@16
Kevin (Multi-Turn)	82%	46%	1.10x	0.40x	43%	15%	20%	6%
Single-Turn RL	82 %	45%	0.85x	0.35x	43%	16%	16%	4%
Qwen QwQ-32B	56%	11%	0.53x	0.08x	23%	3%	10%	1%
OpenAl o4-mini	38%	22%	0.78x	0.27x	21%	7%	13%	6%
OpenAI o3-mini	27%	8%	0.30x	0.08x	9%	2%	4%	2%

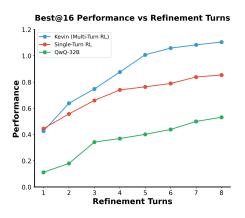
Table 1: **Kevin (multi-turn RL) outperforms other models in correctness and performance.** We evaluate on 100 unseen KernelBench tasks with 16 parallel trajectories and 8 refinement turns.

5.2 SCALING REFINEMENT TURNS

Leveraging execution feedback is crucial at test time (Ehrlich et al., 2025; Wang et al., 2025a). Thus, we evaluate how Kevin scales with additional refinement turns. As shown in Figure 5, the single-turn model achieves slightly better performance with 1 turn, as its training objective optimizes for a single attempt. However, when given more refinement turns, the multi-turn trained model achieves significantly higher performance, with its curve showing the highest slope. This shows that multi-turn training enhances the model's ability to refine and optimize kernels over turns.

5.3 SCALING PARALLEL SAMPLES

We study how best@k performance scales when increasing the number of parallel trajectories k, while fixing the number of serial refinements turns. Prior work for RLVR on math problems (Yue et al., 2025) found that RL training limits models' exploration capacity, leading to worse best@k metrics than the base model at large k. As shown in Figure 6, the performance curve of the single-turn RL model presents a lower slope compared to the base model, possibly hinting at this phenomenon. In contrast, our model trained with multi-turn RL achieves a higher slope compared to both the single-turn counterpart and the base model, suggesting that multi-turn training could maintain model's exploration capacity while improving model's performance.



Best@k Performance (8 refinement turns)

1.05

Kevin (Multi-Turn RL)

Single-Turn RL

OwQ-32B

0.45

0.30

0.15

k (parallel trajectories)

Figure 5: **Kevin effectively leverages multiple turns**. We evaluate the above checkpoints under the same environment with 16 parallel trajectories and 8 refinement turns.

Figure 6: Multi-turn training maintains exploration capacity. Refinement turns are fixed to 8, and best@k performance is computed with the estimator according to Chen et al. (2021).

5.4 PARALLEL VS SEQUENTIAL SCALING

As scaling test-time compute through parallel sampling (Snell et al., 2024) and sequential iterative refinement (Ehrlich et al., 2025) are both beneficial, we want to systematically compare their effectiveness for kernel generation. To investigate, we evaluate 3 inference-time configurations under the

same total inference call budget (128 kernels): 128 trajectories with 1 turn, 32 trajectories with 4 turns, and 16 trajectories with 8 turns. As Table 2 shows, allocating more refinement turns during test-time is consistently better across various models, with 16 trajectories and 8 turns being optimal.

As Section 5.1 shows, multi-turn outperforms single-turn training when evaluated in a multi-turn inference setting. But since single-turn training optimizes for single-turn performance, a natural question arises: does the single-turn trained model perform better by generating more single-turn responses in parallel? In Table 2, we observe that in a single-turn inference setting with 128 parallel trajectories, the single-turn model achieves slightly better performance than the multi-turn model. However, when given more refinement turns at inference, the performance and correctness improve for all models. This strengthens the case for training a model that could use feedback effectively across multiple turns. Moreover, the multi-turn trained model achieves significantly higher performance, with faster improvement rates compared to the single-turn trained model at test-time.

	Inference Config			Performance	Correctness
Model	Total	# Traj	# Turns	best@# traj	best@# traj
Multi-Turn RL	128	16	8	1.10x	82.00%
Multi-Turn RL	128	32	4	1.02x	83.00%
Multi-Turn RL	128	128	1	0.65x	76.00%
Single-Turn RL	128	16	8	0.85x	82.00%
Single-Turn RL	128	32	4	0.81x	79.00%
Single-Turn RL	128	128	1	0.70x	73.00%
QwQ-32B	128	16	8	0.53x	57.00%
QwQ-32B	128	32	4	0.47x	52.00%
QwQ-32B	128	128	1	0.42x	54.00%

Table 2: Multi-turn inference with 16 trajectories and 8 turns is our most optimal setup, when comparing inference configurations and their performance (× speedup) and correctness rates.

6 DISCUSSION

6.1 Model Instability

As prior RLVR work (Team et al., 2025) on QwQ-32B has shown, maintaining RL training stability is a recurring challenge. In our multiturn setting, we notice distinctive patterns of instability, and develop a proxy signal that guides mitigation strategies. Specifically, we observe that training for longer often causes generation of repetitive and nonsensical outputs ("junk"). In the multi-turn case, junk first appears in the final turn and gradually spreads to earlier turns, leading to model collapse.

We identified a proxy signal, which we call the "Not Okay Ratio". QwQ-32B always begins its chain of thought with "Okay, " but after 80 gradient steps, the model begins with gradient versions with gradient like "Okay,"

Not Okay Ratio vs Step

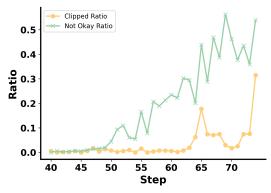


Figure 7: "Not Okay Ratio" foresees model instability. Here the proxy signal appears roughly 15 steps earlier than junk, which is indicated by the response "Clipping Ratio" metric (Luo et al., 2025b).

begins with erratic variants like "Okay Amigos, so I need to optimize this 3D tensor-matrix multiplication" and "Okay Holy crap, I need to get this code optimized"; tracking the "Not Okay Ratio" offers a reliable early proxy for model instability and well precedes junk.

As detailed in Appendix F, after attempting mitigations such as a KL penalty, we found that using constant-length normalization in the GRPO loss (Liu et al., 2025), together with gradient-norm clipping at 0.05, successfully delayed the onset of junk responses to beyond 100 gradient steps.

6.2 REWARD HACKING

We observe forms of reward hacking, as model capabilities fall short of task difficulty (Amodei et al., 2016). Concretely, when a weaker model such as <code>DeepSeek-R1-Distill-Qwen-7B</code> fails to produce the correct CUDA kernels, it resorts to directly copying the reference implementation, inheriting from it, or wrapping it in try-except statements. With a stronger prior like <code>QwQ-32B</code>, the model only fuses simple operators (ReLU, Max) and leaves key operators unmodified (in PyTorch). We address these issues by imposing stricter format checks that assign 0 reward to responses with <code>any</code> PyTorch functional operators. We elaborate on concrete examples in Appendix G.

6.3 DATA DISTRIBUTION

We found it critical to have a balanced difficulty distribution across the dataset, so that on average each batch contains both easier and harder tasks. In one experiment with <code>DeepSeek-R1-Distill-Qwen-14B</code> (DeepSeek-AI, 2025), we trained on a subset of only easy tasks. The reward quickly plateaus as the model overfits to a single difficulty level. Training with a stronger base model <code>QwQ-32B</code> and on both level 1 and 2 of KernelBench resolved the issue.

7 Conclusion

7.1 SUMMARY

We designed a multi-turn RL training recipe that addresses challenges when applied to the real-world task of kernel generation: specifically, effective context management and credit attribution across every turn to enable better sample efficiency. We also added safeguards against reward hacking, and experimented with approaches to constrain and predict instability.

We present Kevin, the first model trained with RL to generate CUDA kernels. Evaluated on an unseen evaluation set, Kevin outperforms both its single-turn RL counterpart and frontier models, demonstrating that our training recipe enables the model to learn more effective refinement strategies. Multi-turn training also enables better test-time scaling, both when increasing sequential refinement and parallel sampling compute, while preserving the exploration capacity of the model.

7.2 LIMITATIONS

Our work is limited by the number of robust tasks in kernel generation (unlike math or general coding with thousands of readily available tasks). KernelBench contains only 250 tasks and requires substantial pre-processing (Appendix A). Moreover, multi-turn RL is computationally expensive, even after extensive system optimization (Appendix C), as each rollout involves serial steps of reasoning inference, complex code generation, and careful kernel evaluation.

Nonetheless, we believe that showing significant performance gains in this domain, even under limited data and compute, highlights the effectiveness of our multi-turn training recipe. With more robust kernel environments, stronger model priors, and improved RL frameworks, we expect our method to scale accordingly.

We further note as KernelBench tasks are specified with pre-defined tensor input sizes, the speedups we measure in Section 3.2 are only accurate for those dimensions and on NVIDIA H200 GPUs.

7.3 FUTURE WORK

We see several directions for extending our method. Incorporating a learned value network and PPO (Schulman et al., 2017) may improve baseline estimation. More sophisticated search methods (beam search, MCTS (Silver et al., 2017)) may be applied at train and test time. Inspired by recent works (Sareen et al., 2025), the value network could also serve as a verifier for search at test-time.

Our multi-turn RL recipe demonstrates success in the real-world engineering task of GPU kernel generation. We hope our flexible design could be applicable to a wider range of tasks with verifiable rewards and execution feedback across a trajectory. We believe explicitly training models to reason about complex tasks over multiple turns is a key step towards enabling autonomous AI systems.

8 ETHICS STATEMENT

This work introduces Kevin, a multi-turn RL training method to enhance LLM's ability specifically for the task of automatic kernel generation. Our research builds on the publicly available model of QwQ-32B (Team, 2025d) and KernelBench dataset (Ouyang et al., 2025). We document in-depth how we use the dataset and post-train the model.

Our work does not introduce new risks that are not already inherent in the underlying base model. We do not involve any human subject nor do we make comparison with human kernel engineers in our study, as our baseline comparisons are against the PyTorch framework (Ansel et al., 2024), following the evaluation methodology proposed in KernelBench.

9 REPRODUCIBILITY STATEMENT

Training Recipe: We cover various challenges encountered during training in detail and propose effective mitigation: covering training stability F, avoiding reward hacking G, and careful considerations for RL design 4 with ablation studies. We elaborate on how we conduct dataset processing A) and provide a comprehensive set of hyper-parameters for our final model (AppendixB.5).

Computational Requirements: Each of our multi-turn training runs (for ablations and the final run) requires 650 H200 hours. As discussed in Appendix C, we take steps to improve the training efficiency of this complex multi-turn RL pipeline with in-the-loop kernel profiling. We elaborate on the computation cost and step time in Appendix C.1 and specifically in Table 3.

Hardware Specifications: We conduct all of our RL training, evaluation, and inference on the NVIDIA H200 platform. All of our kernel runtime measurement and baseline are specific to PyTorch 2.6 and H200 hardware.

Model Weights: Model weights will be released as open source and will be accessible to ensure reproducibility.

Evaluation: For our result, we compare our methods with other models that are either released open-source (QwQ-32B) or using a fixed version of the cloud API endpoints (04-mini-2025-04-16, 03-mini-2025-01-31).

REFERENCES

- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety, 2016. URL https://arxiv.org/abs/1606.06565.
- Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos, Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, C. K. Luk, Bert Maher, Yunjie Pan, Christian Puhrsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Shunting Zhang, Michael Suo, Phil Tillet, Xu Zhao, Eikan Wang, Keren Zhou, Richard Zou, Xiaodong Wang, Ajit Mathews, William Wen, Gregory Chanan, Peng Wu, and Soumith Chintala. Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, ASPLOS '24, pp. 929–947, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703850. doi: 10.1145/3620665.3640366. URL https://doi.org/10.1145/3620665.3640366.
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling, 2024. URL https://arxiv.org/abs/2407.21787.
- 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.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021. URL https://arxiv.org/abs/2107.03374.
- Terry Chen, Bing Xu, and Kirthi Devleker. Automating gpu kernel generation with deepseek-r1 and inference-time scaling. https://developer.nvidia.com/blog/automating-gpu-kernel-generation-with-deepseek-r1-and-inference-time-scaling/, February 2025. Accessed: 2025-05-15.
- Sana Damani, Siva Kumar Sastry Hari, Mark Stephenson, and Christos Kozyrakis. Warpdrive: An agentic workflow for ninja gpu transformations. In *Proceedings of the Machine Learning for Systems Workshop at NeurIPS 2024*, 2024. URL https://mlforsystems.org/assets/papers/neurips2024/paper32.pdf. Accessed: 2025-05-15.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning, 2023. URL https://arxiv.org/abs/2307.08691.
- Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness, 2022. URL https://arxiv.org/abs/2205.14135.
- DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL https://arxiv.org/abs/2501.12948.
- Tulsee Doshi. Gemini 2.5: Our most intelligent models are getting even better. https://blog.google/technology/google-deepmind/google-gemini-updates-io-2025/, May 2025. Accessed: 2025-05-21.

- Yan Duan, John Schulman, Xi Chen, Peter L. Bartlett, Ilya Sutskever, and Pieter Abbeel. Rl²: Fast reinforcement learning via slow reinforcement learning, 2016. URL https://arxiv.org/abs/1611.02779.
 - Ryan Ehrlich, Bradley Brown, Jordan Juravsky, Ronald Clark, Christopher Ré, and Azalia Mirhoseini. Codemonkeys: Scaling test-time compute for software engineering, 2025. URL https://arxiv.org/abs/2501.14723.
 - Paul Gauthier. o1 tops aider's new polyglot leaderboard. https://aider.chat/2024/12/21/polyglot.html, December 2024. Accessed: 2025-04-16.
 - Jonas Gehring, Kunhao Zheng, Jade Copet, Vegard Mella, Quentin Carbonneaux, Taco Cohen, and Gabriel Synnaeve. Rlef: Grounding code llms in execution feedback with reinforcement learning, 2025. URL https://arxiv.org/abs/2410.02089.
 - Anna Goldie, Azalia Mirhoseini, Hao Zhou, Irene Cai, and Christopher D. Manning. Synthetic data generation & multi-step rl for reasoning & tool use, 2025. URL https://arxiv.org/abs/2504.04736.
 - Google DeepMind. Alphaevolve: A gemini-powered coding agent for designing advanced algorithms, May 2025. URL https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/. Accessed: 2025-05-15.
 - Jian Hu, Xibin Wu, Zilin Zhu, Xianyu, Weixun Wang, Dehao Zhang, and Yu Cao. Openrlhf: An easy-to-use, scalable and high-performance rlhf framework, 2024. URL https://arxiv.org/abs/2405.11143.
 - 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. URL https://arxiv.org/abs/2403.07974.
 - Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues?, 2024. URL https://arxiv.org/abs/2310.06770.
 - Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro von Werra, and Harm de Vries. The stack: 3 the of permissively licensed source code, 2022. URL https://arxiv.org/abs/2211.15533.
 - Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
 - Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing frontiers in open language model post-training, 2025. URL https://arxiv.org/abs/2411.15124.
 - Robert Tjarko Lange, Aaditya Prasad, Qi Sun, Maxence Faldor, Yujin Tang, and David Ha. The ai cuda engineer: Agentic cuda kernel discovery, optimization and composition, 2025. URL https://pub.sakana.ai/static/paper.pdf. Accessed: 2025-05-15.
 - Jianling Li, Shangzhan Li, Zhenye Gao, Qi Shi, Yuxuan Li, Zefan Wang, Jiacheng Huang, Haojie Wang, Jianrong Wang, Xu Han, Zhiyuan Liu, and Maosong Sun. Tritonbench: Benchmarking large language model capabilities for generating triton operators, 2025. URL https://arxiv.org/abs/2502.14752.

Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. Starcoder: may the source be with you!, 2023. URL https://arxiv.org/abs/2305.06161.

- Jiate Liu, Yiqin Zhu, Kaiwen Xiao, Qiang Fu, Xiao Han, Wei Yang, and Deheng Ye. Rltf: Reinforcement learning from unit test feedback, 2023. URL https://arxiv.org/abs/2307.04349.
- 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, 2025. URL https://arxiv.org/abs/2503.20783.
- Michael Luo, Sijun Tan, Roy Huang, Ameen Patel, Alpay Ariyak, Qingyang Wu, Xiaoxiang Shi, Rachel Xin, Colin Cai, Maurice Weber, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepcoder: A fully open-source 14b coder at o3-mini level. https://pretty-radio-b75.notion.site/DeepCoder-A-Fully-Open-Source-14B-Coder-at-03-mini-Level-1cf81902c14680b3bee5eb349a512a51, 2025a. Notion Blog.
- Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing ol-preview with a 1.5b model by scaling rl. https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-Ol-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2, 2025b. Notion Blog.
- Mathematical Association of America. American invitational mathematics examination AIME. https://maa.org/math-competitions/american-invitational-mathematics-examination-aime, February 2024. Accessed: 2025-09-24.
- METR. Measuring automated kernel engineering, February 2025. URL https://metr.org/blog/2025-02-14-measuring-automated-kernel-engineering/. Accessed: 2025-05-15.
- Daniel Nichols, Aniruddha Marathe, Harshitha Menon, Todd Gamblin, and Abhinav Bhatele. Hpc-coder: Modeling parallel programs using large language models. In *ISC High Performance 2024 Research Paper Proceedings (39th International Conference)*, pp. 1–12. IEEE, May 2024. doi: 10.23919/isc.2024.10528929. URL http://dx.doi.org/10.23919/ISC.2024.10528929.
- John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. Scalable parallel programming with cuda. In *ACM SIGGRAPH 2008 Classes*, SIGGRAPH '08, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781450378451. doi: 10.1145/1401132.1401152. URL https://doi.org/10.1145/1401132.1401152.
- Allen Nie, Yi Su, Bo Chang, Jonathan N. Lee, Ed H. Chi, Quoc V. Le, and Minmin Chen. Evolve: Evaluating and optimizing llms for exploration, 2024. URL https://arxiv.org/abs/2410.06238.
- NVIDIA. Gpu mode at nvidia gtc 2025, 2025. URL https://www.youtube.com/watch?v=mdDVkBeFy9A. Accessed: 2025-05-15.
- NVIDIA Corporation. Cutlass: Cuda templates for linear algebra subroutines, May 2025. URL https://github.com/NVIDIA/cutlass. Accessed: 2025-05-15.

- Anne Ouyang, Simon Guo, Simran Arora, Alex L. Zhang, William Hu, Christopher Ré, and Azalia Mirhoseini. Kernelbench: Can Ilms write efficient gpu kernels?, 2025. URL https://arxiv.org/abs/2502.10517.
 - Yuxiao Qu, Matthew Y. R. Yang, Amrith Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan Salakhutdinov, and Aviral Kumar. Optimizing test-time compute via meta reinforcement finetuning, 2025. URL https://arxiv.org/abs/2503.07572.
 - Kusha Sareen, Morgane M Moss, Alessandro Sordoni, Rishabh Agarwal, and Arian Hosseini. Putting the value back in rl: Better test-time scaling by unifying llm reasoners with verifiers, 2025. URL https://arxiv.org/abs/2505.04842.
 - Thomas Schmied, Jörg Bornschein, Jordi Grau-Moya, Markus Wulfmeier, and Razvan Pascanu. Llms are greedy agents: Effects of rl fine-tuning on decision-making abilities, 2025. URL https://arxiv.org/abs/2504.16078.
 - John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.
 - Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.
 - Alexander Shypula, Aman Madaan, Yimeng Zeng, Uri Alon, Jacob Gardner, Milad Hashemi, Graham Neubig, Parthasarathy Ranganathan, Osbert Bastani, and Amir Yazdanbakhsh. Learning performance-improving code edits, 2024. URL https://arxiv.org/abs/2302.07867.
 - David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. Mastering chess and shogi by self-play with a general reinforcement learning algorithm, 2017. URL https://arxiv.org/abs/1712.01815.
 - Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling Ilm test-time compute optimally can be more effective than scaling model parameters, 2024. URL https://arxiv.org/abs/2408.03314.
 - Benjamin F. Spector, Simran Arora, Aaryan Singhal, Daniel Y. Fu, and Christopher Ré. Thunderkittens: Simple, fast, and adorable ai kernels, 2024. URL https://arxiv.org/abs/2410.20399.
 - Fahim Tajwar, Yiding Jiang, Abitha Thankaraj, Sumaita Sadia Rahman, J Zico Kolter, Jeff Schneider, and Ruslan Salakhutdinov. Training a generally curious agent, 2025. URL https://arxiv.org/abs/2502.17543.
 - Gemma Team. Gemma 3 technical report, 2025a. URL https://arxiv.org/abs/2503.19786.
 - Kimi Team. Kimi k1.5: Scaling reinforcement learning with llms, 2025b. URL https://arxiv.org/abs/2501.12599.
 - NovaSky Team. Sky-t1: Train your own o1 preview model within \$450. https://novasky-ai.github.io/posts/sky-t1, 2025c. Accessed: 2025-01-09.
- Prime Intellect Team, Sami Jaghouar, Justus Mattern, Jack Min Ong, Jannik Straube, Manveer Basra, Aaron Pazdera, Kushal Thaman, Matthew Di Ferrante, Felix Gabriel, Fares Obeid, Kemal Erdem, Michael Keiblinger, and Johannes Hagemann. Intellect-2: A reasoning model trained through globally decentralized reinforcement learning, 2025. URL https://arxiv.org/abs/2505.07291.
 - Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025d. URL https://qwenlm.github.io/blog/qwq-32b/.

- Philippe Tillet, H. T. Kung, and David Cox. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, MAPL 2019, pp. 10–19, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367196. doi: 10.1145/3315508.3329973. URL https://doi.org/10.1145/3315508.3329973.
 - Siddhant Waghjale, Vishruth Veerendranath, Zora Zhiruo Wang, and Daniel Fried. Ecco: Can we improve model-generated code efficiency without sacrificing functional correctness?, 2024. URL https://arxiv.org/abs/2407.14044.
 - Guanhua Wang, Heyang Qin, Sam Ade Jacobs, Connor Holmes, Samyam Rajbhandari, Olatunji Ruwase, Feng Yan, Lei Yang, and Yuxiong He. Zero++: Extremely efficient collective communication for giant model training, 2023. URL https://arxiv.org/abs/2306.10209.
 - Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. Openhands: An open platform for ai software developers as generalist agents, 2025a. URL https://arxiv.org/abs/2407.16741.
 - Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Lucas Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, Weizhu Chen, Shuohang Wang, Simon Shaolei Du, and Yelong Shen. Reinforcement learning for reasoning in large language models with one training example, 2025b. URL https://arxiv.org/abs/2504.20571.
 - Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Kefan Yu, Minh Nhat Nguyen, Licheng Liu, Eli Gottlieb, Monica Lam, Yiping Lu, Kyunghyun Cho, Jiajun Wu, Li Fei-Fei, Lijuan Wang, Yejin Choi, and Manling Li. Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning, 2025c. URL https://arxiv.org/abs/2504.20073.
 - Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, Lingming Zhang, Daniel Fried, Gabriel Synnaeve, Rishabh Singh, and Sida I. Wang. Swe-rl: Advancing Ilm reasoning via reinforcement learning on open software evolution, 2025. URL https://arxiv.org/abs/2502.18449.
 - Violet Xiang, Charlie Snell, Kanishk Gandhi, Alon Albalak, Anikait Singh, Chase Blagden, Duy Phung, Rafael Rafailov, Nathan Lile, Dakota Mahan, Louis Castricato, Jan-Philipp Franken, Nick Haber, and Chelsea Finn. Towards system 2 reasoning in llms: Learning how to think with meta chain-of-thought, 2025. URL https://arxiv.org/abs/2501.04682.
 - Zihao Ye, Lequn Chen, Ruihang Lai, Wuwei Lin, Yineng Zhang, Stephanie Wang, Tianqi Chen, Baris Kasikci, Vinod Grover, Arvind Krishnamurthy, and Luis Ceze. Flashinfer: Efficient and customizable attention engine for llm inference serving. *arXiv preprint arXiv:2501.01005*, 2025. URL https://arxiv.org/abs/2501.01005.
 - Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL https://arxiv.org/abs/2503.14476.
 - Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model?, 2025. URL https://arxiv.org/abs/2504.13837.
 - Chenggang Zhao, Liang Zhao, Jiashi Li, and Zhean Xu. Deepgemm: clean and efficient fp8 gemm kernels with fine-grained scaling. https://github.com/deepseek-ai/DeepGEMM, 2025.
 - Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. Minif2f: a cross-system benchmark for formal olympiad-level mathematics, 2022. URL https://arxiv.org/abs/2109.00110.

use-with-reinforcement-learning, 2025. Accessed: 2025-04-17.

Yifei Zhou, Andrea Zanette, Jiayi Pan, Sergey Levine, and Aviral Kumar. Archer: Training language model agents via hierarchical multi-turn rl, 2024. URL https://arxiv.org/abs/2402.19446.
Richard Zhuang*, Trung Vu*, Alex Dimakis, and Maheswaran Sathiamoorthy. Improving multi-turn tool use with reinforcement learning. https://www.bespokelabs.ai/blog/improving-multi-turn-tool-

A KERNELBENCH MODIFICATIONS

We use KernelBench Ouyang et al. (2025) as our training environments. KernelBench is a popular benchmark for evaluating LLMs' ability to generate performant CUDA kernels for deep learning workloads in PyTorch. Each KernelBench task consists in generating a CUDA kernel given a PyTorch reference implementation, which is used to evaluate correctness and speedup.

A.1 TASK IMPROVEMENTS

We identify several limitations in the original KernelBench and introduce targeted modifications to address them. These changes are crucial to mitigate reward hacking, as shown in Section 6.2.

- We sand-boxed the kernel evaluation process so that fatal errors, such as CUDA illegal memory accesses, do not crash the RL training process.
- A significant issue we noted in KernelBench was that for many tasks, the input tensors used
 to measure performance are quite small. This causes kernel launch overhead to take up a
 significant portion of the runtime. To address this, we enlarged the tensor dimensions of the
 affected tasks.
- A sneakier bug in the KernelBench's evaluation harness caused the tested kernel to recycle
 the output tensor from the reference implementation (which was run immediately before) as
 its own tensor output. As a result of this, a kernel that only computes (correctly) a portion
 of the output tensor would still pass the correctness check. We address this by running the
 tested kernel first and only after the reference implementation, thus avoiding this hack.

In the end, we chose a total of 180 tasks as training environments, with 90 of the 100 Level 1 problems and 90 Level 2 problems (sequences of operators with fusion opportunities).

A.2 CONSTRUCTION OF ADDITIONAL EVALUATION SET

Since current KernelBench does not provide a train-test split, we construct 80 additional tasks following the same methodology that KernelBench was constructed.

KernelBench Level 2 is constructed by composing a subset of PyTorch operators as sequences of operators. Specifically, the PyTorch operators are categorized as:

- Main operators: Conv2d, Matmul, Gemm, BMM, Conv3d, ConvTranspose2d, ConvTranspose3d.
- Activations: ReLU, Sigmoid, Tanh, LeakyReLU, GELU, Swish, Softmax, Mish, Hardtanh, HardSwish.
- Element-wise operators: Add, Multiply, Subtract, Divide, Clamp, Scale, ResidualAdd.
- Normalizations: BatchNorm, LayerNorm, InstanceNorm, GroupNorm.
- **Pooling:** MaxPool, AvgPool, GlobalAvgPool.
- Bias: BiasAdd.
- Reductions: Sum, Mean, Max, Min, LogSumExp.
- Others: Residual Add, Scaling.

To construct the additional eval set (unseen from train set), following the methodology from original KernelBench task construction:

- 1. We sample from the available operators listed above: 1 main operator (computationally expensive), and 2-5 other operators.
- 2. We ask a language model, namely Gemini 2.5-Flash (Doshi, 2025), to generate a PyTorch program that creates a kernel by combining these operators. We also ask it to generate sample tensor sizes for the task.

- 3. We ensure this PyTorch program can be executed and has a runtime on NVIDIA H200 > 0.1ms, to avoid the runtime being dominated by kernel launch (CPU) overhead.
- 4. We make sure this PyTorch program (with the same sequence of operators) is not present in existing KernelBench Level 1 and 2 programs.

We manually inspected all new task programs to ensure their validity. We build the evaluation set by combining our 80 newly created tasks with the 20 remaining original KernelBench tasks, for a total of 100 unseen evaluation tasks.

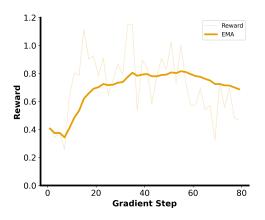
B ADDITIONAL DETAILS ON MULTI-TURN RL

Here we elaborate on design choices for our RL Training as described in Section 3.3 and Section 4, along with some ablation results.

B.1 MOTIVATION FOR TURN-WISE REWARD

In our multi-turn RL training setup, within each training step we have a trajectory with n refinement turns. A possible approach would be to compute the reward based on the kernel at the last turn, similar to what is used in RLEF (Gehring et al., 2025). However, for the GPU kernel optimization setting, using just the last kernel might not be optimal at times: for example, as shown earlier in Figure 1, kernel 3 is correct but kernel 4 is incorrect as the model attempts more aggressive optimizations.

In this setting, computing reward based on the best kernel among the trajectory instead (max speedup) is a more natural choice. However, using only the max kernel score forces us to discard all turns in a trajectory after the max turn, possibly wasting a significant amount of inference rollouts: In the previous example, we would have to completely discard the reasoning trace, code, and evaluation for kernel 4. Thus, we arrived at our approach in Section 4.3, which uses a discounted look-ahead max or sum, enabling more sample-efficient training.



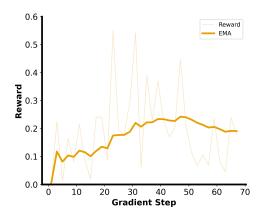


Figure 8: Training reward with correctness weighting of 1, performance / speedup weighting of 1. Concretely, $S = \mathbf{1}_{\{\text{correct}\}} + \frac{T_{\text{baseline}}}{T_{\text{kernel}}} \cdot \mathbf{1}_{\{\text{correct}\}}$.

Figure 9: Training reward with no correctness weighting, performance / speedup weighting of 1. (speedup is 0 if kernel is incorrect). Concretely, $S = \mathbf{1}_{\{\text{correct}\}} \cdot \frac{T_{\text{baseline}}}{T_{\text{kernel}}}.$

B.2 WEIGHTING FOR SCORE

In Section 3.2, we explain our score design, which assigns a scalar value (score S) based on a kernel's correctness and speedup. We explore score design and how to balance the correctness-performance trade-off, after series of small-scale ablations on QwQ-32B (Team, 2025d).

We decided on a weighting of 0 . 3 on correctness and using speedup for performance (raw speedup itself, no weighting), which is $S = 0.3 \cdot \mathbf{1}_{\{\text{correct}\}} + \mathbf{1}_{\{\text{correct}\}} \cdot \frac{T_{\text{baseline}}}{T_{\text{longol}}}$.

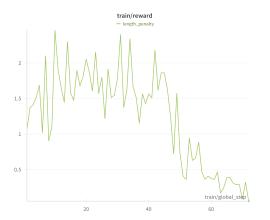
Here we present some ablation studies we ran with different weighting configurations for score design, particularly focusing on adjusting the weighing for correctness, in the context of single-turn RL (GRPO) training (as shown in Section 3.3). As show an example in Figure 8, where we set the weighting to 1.0 for correctness, the reward plateaus and eventually decreased; concretely, we observed that the model over-optimizes for generating correct kernels and does not explore speedup as much, causing the reward to plateau during training. In another experiment in Figure 9, we set the weighting to 0 for correctness, only rewarding the model for generating performant (and correct) kernels. We again observed the reward plateau. Thus, we hypothesize that it is still important to reward the model for correct kernels, as long as the correctness reward is not too significant, balancing the correctness-performance tradeoff.

B.3 Number of Trajectories during Training

We vary the number of parallel trajectories during Multi-Turn RL training (Section 4), using 64 parallel trajectories instead of 16 for each task during each training step. We note that best@16 correctness slightly increases, but the overall performance does not show significant improvements. Due to the high-compute requirements of doing more generations during training, we chose to train with 16 parallel trajectories.

B.4 LENGTH PENALTY

 We explore incorporating response length as a part of the reward design to incentivize the model to use its reasoning tokens more efficiently. We attempted a run using the length penalty from Kimi Team (2025b) on <code>DeepSeek-R1-Distill-Qwen-14B</code>. As shown in Figures 10 and 11, we found that the response length of the responses collapses, with the model no longer outputting CoT after 10 training steps, suggesting that the addition of a length penalty is counterproductive for our setting.



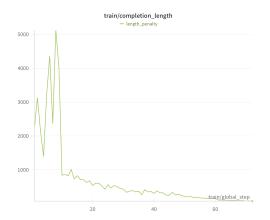


Figure 10: Training Reward collapses when including length penalty as part of reward

Figure 11: Response length of generations collapses when including length penalty as part of reward.

B.5 DETAILED TRAINING HYPERPARAMETERS

Here is the set of hyperparameter for our final Kevin training run.

```
1019
1020 | Constant learning rate of 2e-6 with warmup ratio of 0.03

1021 | Max grad norm = 0.05 |

KL coeff = 0

1022 | Temperature = 0.9

1023 | Top p = 0.95

1024 | Eps clip = 0.2 with clip high = 0.28

1025 | Max prompt len = 8192 |

8 | Max generate len = 22432
```

B.6 BASE MODEL CHOICE RATIONALE

We experimented with several different base models, such as DeepSeek-R1-Distill-Qwen7B DeepSeek-AI (2025), DeepSeek-R1-Distill-Llama8B, and Gemma27B-Instruct Team (2025a). These models, however exhibit weak kernel writing priors, which causes the initial reward to be overly sparse for effective learning, making the model resort to reward hacking (see Section 6.2).

We thus choose <code>Qwen QwQ-32B</code> (Team, 2025d) as our base model, which exhibits, among all the models we have evaluated of comparable size, the strongest priors.

C RL Infrastructure

Conducting RL training on a highly challenging task like GPU kernel generation is a computationally expensive process, requiring full-policy updates on a sufficiently capable base model, as discussed in Section 6.2.

Although a few open-source RL frameworks existed when we began this study, it is still difficult to support training in a kernel evaluation environment and including multiple turns within one training step. We built our training framework on top of the OpenRLHF (Hu et al., 2024) framework.

We use vLLM (Kwon et al., 2023) for inference and DeepSpeed Zero-3 (Wang et al., 2023) for offloading optimizer states.

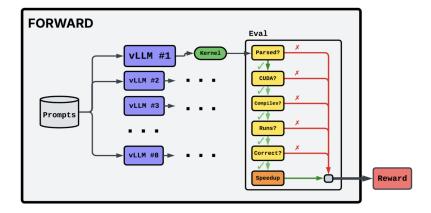


Figure 12: Overview of our RL Training infrastructure.

Each of the 8 GPUs handles the kernel generation and evaluation for one task. After the response generation finishes, each GPU offloads its vLLM engine to CPU memory and evaluates the kernels it generated. We run the evaluation and calculate reward and evaluation info. Each GPU then wakes up its corresponding vLLM engine and regenerates kernels.

We optimized our training infrastructure to co-locate vLLM rollout engines, kernel execution environments, and DeepSpeed trainers on the same GPU device, so other small research teams with 1 cluster could also experiment with our proposed method.

C.1 TRAIN TIME STATISTICS

Here we elaborate more on the cost of our multi-turn training. The nature of multi-turn RL requires multiple serial turns of parallel rollouts and kernel compilation/execution after each step, making the overall training process compute-intensive. To accurately measure kernel runtime, we must clear the GPUs of any running processes and perform additional operations, such as warmup steps before profiling, which further limits the training speed. Here we show key training time statistics:

Overall, one training run (Section 4) takes 650 H200 hours, equivalent to around 3 days and 9 hours on a single node of 8xH200s. However, we believe concurrent and future systems projects (such as SkyRL by Cao et al. (2025)) will improve training efficiency, especially for roll-outs with complex

Configuration	Value
Gradient steps	80
Parallel trajectory rollouts	16
Refinement turns (serial env. interactions)	4
Gradient updates per batch	2 (1 on-policy, 1 off-policy)
Time for rollout + kernel execution (per step)	\sim 1.5 hours
Time for 1 gradient update (2 steps)	\sim 0.5 hours
Base model	QwQ-32B

Table 3: Setup and Cost of multi-turn Training for Kevin on 8xH200s.

interactions with environments. The demanding computational requirement of multi-turn RL is what leads us to focus on improving the sample efficiency of our method; specifically, we choose to train on every sample regardless of their performance and attribute credit effectively with our reward design.

D Inference Setup

1090

1092

1093

1094 1095 1096

1097 1098

1099

1100

1101

1102

1103

Our prompt is similar to the prompt used in KernelBench (Ouyang et al., 2025). We use this during training and test-time inference. In the first refinement turn, we add an example of the inline CUDA format to the prompt but remove it afterwards.

Below we show how we construct the context in the simplest case (of one turn, or the base prompt). In the context, we present model the KernelBench task, instructions, and a simple 1-shot example of a CUDA add kernel (to inform model the desired format for response):

```
1104
       You are given the following architecture:
1105 1
       import torch
1106
       import torch.nn as nn
1107 4
1108 5
       class Model(nn.Module):
1109 6
            Simple model that performs Layer Normalization.
1110 7
1111 9
            def __init__(self, normalized_shape: tuple):
1112<sub>10</sub>
111311
                 Initializes the LayerNorm layer.
111412
1115 <sup>13</sup>
                     normalized_shape (tuple): Shape of the input tensor to be
1116
            normalized.
1117<sub>15</sub>
1118 16
                 super(Model, self).__init__()
1119 17
                 self.ln = nn.LayerNorm(normalized_shape=normalized_shape)
1120 <sup>18</sup>
            def forward(self, x: torch.Tensor) -> torch.Tensor:
1121 <sub>20</sub>
1122 21
                 Applies Layer Normalization to the input tensor.
1123 22
112423
                     x (torch.Tensor): Input tensor of shape (*,
1125 24
            normalized_shape).
1126 25
1127<sub>26</sub>
                 Returns:
                     torch. Tensor: Output tensor with Layer Normalization
1128 27
            applied, same shape as input.
1129
1130 28
1131 <sub>29</sub> <sub>30</sub>
                 return self.ln(x)
1132 31
       Replace pytorch operators in the given architecture with raw CUDA
1133
            kernels, optimizing for performance on NVIDIA H200 (e.g. shared
           memory, kernel fusion, warp primitives, vectorization,...). Use
```

```
1134
           torch.utils.cpp_extension.load_inline and name your optimized output
1135
           architecture ModelNew. You are not allowed to use torch.nn (except
1136
           for Parameter, containers, and init). The input and output have to
           be on CUDA device. Your answer must be the complete new architecture
1137
           (no testing code, no other code): it will be evaluated and you will
1138
           be given feedback on its correctness and speedup so you can keep
1139
           iterating, trying to maximize the speedup. After your answer,
1140
           summarize your changes in a few sentences. Here is an example:
1141 32
114233
       import torch.nn as nn
1143^{34}
       from torch.utils.cpp_extension import load_inline
1144 35
       # Define the custom CUDA kernel for element-wise addition
1145 37
       elementwise_add_source = """
1146 38
       #include <torch/extension.h>
       #include <cuda runtime.h>
114739
1148^{40}
        _global__ void elementwise_add_kernel(const float* a, const float* b,
1149 41
           float* out, int size) {
1150 42
           int idx = blockIdx.x * blockDim.x + threadIdx.x;
1151 43
           if (idx < size) {</pre>
115244
                out[idx] = a[idx] + b[idx];
1153<sup>45</sup>
1154 46
47
1155 48
       torch::Tensor elementwise add cuda(torch::Tensor a, torch::Tensor b) {
1156 49
           auto size = a.numel();
           auto out = torch::zeros_like(a);
115750
1158 51
1159 52
53
           const int block size = 256;
           const int num blocks = (size + block_size - 1) / block_size;
1160 54
1161 55
           elementwise_add_kernel<<<num_blocks,</pre>
           block_size>>>(a.data_ptr<float>(), b.data_ptr<float>(),
1162
           out.data_ptr<float>(), size);
1163
1164 56
57
           return out;
1165 58
       0.00
1166 59
116760
       elementwise_add_cpp_source = (
1168<sup>61</sup>
           "torch::Tensor elementwise_add_cuda(torch::Tensor a, torch::Tensor
1169 62
           b);"
1170<sub>63</sub>
117164
       # Compile the inline CUDA code for element-wise addition
117265
       elementwise_add = load_inline(
1173 66
1174 67
           name="elementwise_add",
           cpp_sources=elementwise_add_cpp_source,
1175 69
           cuda_sources=elementwise_add_source,
1176 70
           functions=["elementwise_add_cuda"],
           verbose=True,
117771
1178 72
           extra_cflags=[""],
1179 <sup>73</sup> <sub>74</sub>
           extra_ldflags=[""],
1180 75
118176
       class ModelNew(nn.Module):
118277
1183<sup>78</sup>
           def __init__(self) -> None:
1184 <sup>79</sup> <sub>80</sub>
                super().__init__()
                self.elementwise_add = elementwise_add
1185 81
1186 82
           def forward(self, a, b):
118783
                return self.elementwise_add.elementwise_add_cuda(a, b)
```

1189

1190

For our multi-turn RL training (Section 4) and inference (Section 5), we provide model with the kernels, CoTs (summarized), and evaluation results of all previous turns in chronological order. We truncate the turns that do not fit inside the context window, starting from the earliest ones.

```
1191
       <Base prompt containing pytorch architecture and instruction>
1192 <sup>1</sup>
1193
       Here are your previous attempts:
1194 4
1195 5
       < for each (i) previously generated kernel >
            <Previously generated kernel G[i]>
1196 6
1197 7
            <Summary of CoT[i]>
1198 8
1199 10
           <if parsing error>
1200 11
                Your previous answer failed to be parsed due to not adhering to
1201 12
           the desired formatting. Here is the error message: <error_message>
1202
1203 13
           <elif compilation error>
1204 15
1205 16
                Your previous answer failed to compile. Here is the error
           message: <error_message>
1206
1207 17
1208 18
19
            <elif run error>
1209 20
                Your previous answer compiled successfully but had runtime
1210
           errors. Here is the error message: <error_message>
121121
            <elif correctness error>
1212<sup>22</sup>
1213 23
1213 24
                Your previous answer was incorrect. Here is the error message:
1214
           <error_message>
1215 25
            <elif correct>
121626
1217<sup>27</sup>
1218 28
                Your previous answer was correct but can be made faster. Here {\tt is}
           the speedup you achieved relative to the baseline: <speedup>
1219<sub>29</sub>
       Restart your reasoning process and generate new, complete code.
1220 30
1221
```

E ADDITIONAL EVALUATIONS

Here we present some additional evaluation results for Section 5.

E.1 CONFIDENCE INTERVALS

We compute the confidence intervals of best@16 and avg@16 performance for the multi-turn and single-turn RL across 5 runs, as shown in Table 4. These results show multi-turn RL has statistically significant improvement on both metrics and hence its effectiveness.

Model	Performance			
	best@16	avg@16		
Multi-turn RL Single-turn RL		—		

Table 4: Evaluation on our evaluation set across 5 runs with confidence interval. Multi-turn RL outperforms Single-turn RL on both best@16 and avg@16 performance.

E.2 CHOICE OF BASELINE MODEL COMPARISON

Here we elaborate on the choice of model comparisons used for 5.1, notably against both Kevin's base model (QwQ-32B) and frontier reasoning models (o4-mini, o3-mini). To the best of our knowledge, we are not aware of any model specifically "fine-tuned" for the CUDA context (efforts like Nichols et al. (2024) focus on OpenMP CPU code). CUDA, or GPU code in general, is extremely sparse in the pretraining corpus, only 0.073% of the Stack (Li et al., 2023) code corpus; this makes approaches that depend on readily available data (such as "fine-tuning") difficult. Hence, this data challenge actually highlights the value of our RL-based approach, as we discussed in Section 1. We believe that the comparisons of Kevin against SoTA general-purpose LLMs are fair and fitting, and actually demonstrate the advantage of our RL-based approach in this domain.

Our baseline comparisons, o4-mini and o3-mini, are frontier models that achieve SoTA on challenging code generation benchmarks. Specifically we use o4-mini-2025-04-16 and o3-mini-2025-01-31. As shown below, o4-mini demonstrates a significant lead over our base model QwQ-32B, especially on challenging real-world software tasks such as SWE-Bench (Jimenez et al., 2024) and Polyglot (Gauthier, 2024). Hence, our results in Section 5.1 and Table 1 showing Kevin (post-trained QwQ-32B with multi-turn RL) exceeding o4-mini should be noted as a significant improvement and demonstrate our method's effectiveness.

Model	AIME 24	LiveCodeBench	SWE-Bench Verified	Aider Polyglot
QwQ-32B	79.5%	63.4%	41.3%	20.9%
o4-mini	93.4%	74.2%	68.1%	72.0 %

Table 5: o4-mini shows significant lead over QwQ-32B over a variety of reasoning, coding, and software engineering benchmarks (Mathematical Association of America, 2024; Jain et al., 2024; Jimenez et al., 2024; Gauthier, 2024); Kevin is post-trained on QwQ-32B and shows improvement over both QwQ-32B and o4-mini, as shown in Section 5.1.

E.3 EVALUATION ON KERNELBENCH LEVEL 3

While we focus on our training and evaluation mostly on KernelBench Level 1 and 2 (Section 3), we were also curious and explore testing Kevin on KernelBench Level 3 tasks. They are longer and more challenging (rather than single or a few operators), requiring the end-to-end optimization of full model architectures, such as the VisionTransformer, and miniGPT attention blocks. Kevin is trained using a subset (180) of the KernelBench Level 1 and 2 tasks (single and sequence of operators), and

Level 3 tasks are completely unseen. We evaluate the multi-turn (Kevin), single-turn, and base model (QwQ-32B) on the 50 level 3 tasks following the same evaluation setup as Section 5. As shown in the table below, multi-turn RL can also generate much faster kernels for these much more complex tasks over both single-turn RL and the base model.

We view Level 3 primarily as an out-of-distribution test: these tasks involve full model architectures with much longer-horizon reasoning, and requiring both kernel generation and effectively dealing with long context. We do not train on any Level 3 tasks as the length of these programs would lead to context explosion (Section 4.1). Hence, our main analysis focuses on Levels 1 and 2, which better focuses on kernel generation performance with more controlled conditions.

	Corre	ctness	Performance		
Method	best@16	avg@16	best@16	avg@16	
Multi-turn RL Single-turn RL QwQ-32B	36% 36% 4%	11.75% 8.38% 0.25%	0.41 0.36 0.04	0.08 0.06 0.002	

Table 6: Multi-turn RL achieves improvements also on the completely unseen and more complex KernelBench Level 3.

F TRAINING STABILITY

The analysis of the "not okay ratio" led us to believe that model instability caused the appearance of nonsensical and repetitive outputs. Therefore, we attempted runs where we enabled KL divergence penalty in the GRPO loss, which would penalize the model from deviating from the base policy too much. Following DeepScaleR (Luo et al., 2025b), we set the KL coefficient to 0.001 and attempted an ablation run. However, we found that the reward plateaus with KL enabled, suggesting that the KL penalty slows down learning. Thus we attempted other techniques of constraining the model from deviating into regions of instability, such as clipping the gradient norm aggressively — which was effective in our setting.

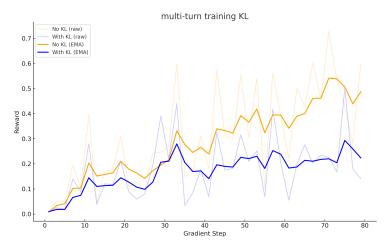


Figure 13: Adding a KL penalty slows down learning. Here we conduct an ablation with KL coefficient $\beta = 0.001$ versus $\beta = 0$. We see that the reward plateaus with KL enabled.

We use 4 refinement turns at train-time for efficient training. During test time, we can afford more extensive test-time compute, so we evaluate on 8 turns instead of 4 turns.

G REWARD HACKING

1350

1351 1352

1353

1354

1355

1356

1357

1364

1365

1366

1377

1378

1386 1387

1388 1389

1390

Here we present excerpts from generated kernels that show signs of reward hacking, previously mentioned in Section 6.2.

In the following example, the model simply copies the PyTorch reference implementation, thus getting rewarded for generating a correct answer with 1.0x speedup. To prevent this, we modify our kernel evaluation environment so that it checks each generated kernel if it contains instances of torch.nn or torch.nn.functional. We assign a reward of 0 to those.

```
1358

1359 | class ModelReLU(Module):

1360 | def forward(self, x):

1361 | relu = torch.nn.ReLU()

1362 | return relu(x)
```

Similarly, the model wraps an incorrect implementation of the CUDA kernel in a try-except statement and invokes the PyTorch implementation functions as a fallback. To prevent this, we assign a reward of 0 to kernels that contain try or except.

```
1367
        class ModelReLU(Module):
1368 <sup>1</sup>
1369
            def forward(self, x):
1370 4
                 try:
1371 <sub>5</sub>
                      ... \# CUDA implementation
                 except Exception as e:
1372 6
                      print ("Custom ReLU kernel failed to compile. Using default
1373 7
            ReLU instead.")
1374
                      relu = torch.nn.ReLU()
1375 <sub>9</sub>
                      return relu(x)
1376
```

Here the model inherits from the reference implementation, bypassing the need for a CUDA implementation. To prevent this, we assign a reward of 0 to kernels that contain pass.

H KERNEL HIGHLIGHT

As an example, we look at one the task form our evaluation set, Conv3D_GELU_Tanh_Clamp:

```
1391
       import torch
1392 2
       import torch.nn as nn
1393 3
       class Model(nn.Module):
1394 4
1395 5
            Model using Conv3d, GELU, Clamp, Tanh
1396 7
1397 8
            def __init__(self, in_channels, out_channels, kernel_size,
           clamp_min, clamp_max):
1398
                super(Model, self).__init__()
1399 <sup>9</sup>
1400 <sup>10</sup>
                self.conv = nn.Conv3d(in_channels, out_channels, kernel_size)
                self.clamp_min = clamp_min
1401 12
    11
                self.clamp_max = clamp_max
1402 13
1403 14
            def forward(self, x):
    15
```

1420 1421

1422

```
1404
               Forward pass of the model.
1405 17
1406 18
               Aras:
140719
                   x (torch.Tensor): Input tensor of shape (batch_size,
           in_channels, depth, height, width).
1408
1409 20
               Returns:
1410 22
                   torch. Tensor: Output tensor of shape (batch_size,
1411
           out_channels, depth-kernel_size+1, height-kernel_size+1,
1412
           width-kernel_size+1).
141323
               x = self.conv(x)
1414^{24}
               x = torch.tanh(x)
1415 26
               x = torch.clamp(x, min=self.clamp_min, max=self.clamp_max)
1416 27
               x = torch.nn.functional.gelu(x)
141728
               return x
1418
```

The task is made up of a computationally intensive function (3D convolution) and a series of lightweight functions (tanh, clamp, gelu).

The model implements a correct solution in the first turn by fusing the kernels, but it only achieves a speedup of **1.06x** over the reference implementation.

```
1423
       import torch
1424^{-1}
       from torch.utils.cpp_extension import load_inline
1425 3
       import math
1426 4
1427 5
       class ModelNew(torch.nn.Module):
           def __init__(self, in_channels, out_channels, kernel_size,
1428 6
           clamp_min, clamp_max):
1429
               super(ModelNew, self).__init__()
1430 8
               self.clamp_min = clamp_min
1431 9
               self.clamp_max = clamp_max
1432 10
                # Initialize convolution weights and bias
1433 11
               self.weight = torch.nn.Parameter(torch.empty(out_channels,
1434 12
           in_channels, kernel_size, kernel_size, kernel_size))
1435
               self.bias = torch.nn.Parameter(torch.empty(out_channels))
1436<sub>14</sub>
                # Weight initialization following PyTorch's Conv3d default
1437 15
               torch.nn.init.kaiming_uniform_(self.weight, a=math.sqrt(5))
               fan_in, _ =
1438 16
           torch.nn.init._calculate_fan_in_and_fan_out(self.weight)
1439
               bound = 1 / math.sqrt(fan_in)
1440 18
               torch.nn.init.uniform_(self.bias, -bound, bound)
1441 19
144220
                # Define the fused convolution and activation CUDA kernel
               conv_activation_source = """
144321
1444 22
               #include <torch/extension.h>
                #include <cuda_runtime.h>
1445 23
24
               #include <math.h>
1446 25
144726
                 _global__ void conv_activation_kernel(const float* input, const
           float* weights, const float* bias, float* output,
1448
                                                         int N, int C_in, int D,
1449 27
           int H, int W, int C_out, int kD, int kH, int kW,
1450 28
                                                         float clamp_min, float
1451
           clamp_max) {
                    int idx = blockIdx.x * blockDim.x + threadIdx.x;
1452 29
                    const int W_{out} = W - kW + 1;
1453 30
1454<sup>31</sup>
                    const int H_{out} = H - kH + 1;
1455 32
1455 33
                   const int D_out = D - kD + 1;
                    const int total_elements = N * C_out * D_out * H_out * W_out;
1456<sub>34</sub>
                    if (idx >= total_elements) return;
1457 35
                    // Decompose index into output coordinates
   36
```

```
1458
1459 37
                     int w_out = idx % W_out;
                     int h_out = (idx / W_out) % H_out;
1460 39
                     int d_out = (idx / (W_out * H_out)) % D_out;
146140
                     int f = (idx / (W_out * H_out * D_out)) % C_out;
                     int n = idx / (C_out * D_out * H_out * W_out);
146241
1463 42
   43
                     float sum = 0;
1464 44
                     for (int c = 0; c < C_in; ++c) {
1465 45
                          for (int kd = 0; kd < kD; ++kd) {
146646
                               for (int kh = 0; kh < kH; ++kh) {
                                   for (int kw = 0; kw < kW; ++kw) {
146747
                                        // Input tensor index calculation
1468 48
                                        int d_in = d_out + kd;
1469 50
                                        int h_in = h_out + kh;
1470<sub>51</sub>
                                        int w_in = w_out + kw;
1471 52
                                        int input_idx = n * C_in * D * H * W +
1472<sup>53</sup>
                                                           C \star D \star H \star W +
1473 54
                                                           d_in * H * W +
1473
55
1474
56
                                                           h_in * W +
                                                           w_in;
1475<sub>57</sub>
                                        // Weight tensor index calculation
                                        int weight_idx = f * C_in * kD * kH * kW +
1476 58
                                                            c \star kD \star kH \star kW +
1477 59
1478<sup>60</sup>
                                                            kd * kH * kW +
    61
                                                            kh * kW +
1479 62
                                                            kw;
1480<sub>63</sub>
                                        sum += input[input_idx] *
1481
            weights[weight_idx];
148264
1483 65
1484 67
                     }
1485 68
                     sum += bias[f];
1486 69
                     // Apply activations sequentially
148770
1488 71
                     sum = tanh(sum);
                     sum = fmaxf(clamp_min, fminf(sum, clamp_max));
    72
1489 73
1490 74
                          float inner = sum + 0.044715f * sum * sum * sum;
149175
                          inner \star= 0.79788456f; // sqrt(2/pi) approximation
                          sum = sum * 0.5f * (1.0f + tanh(inner));
149276
1493 77
1494 79
                      // Output tensor index calculation
1495<sub>80</sub>
                      int output_idx = n * C_out * D_out * H_out * W_out +
1496 81
                                        f * D_out * H_out * W_out +
                                        d_out * H_out * W_out +
149782
                                        h_out * W_out +
1498 83
                                        w_out;
1499 85
                     output[output_idx] = sum;
1500 86
                 }
150187
                 torch::Tensor conv_activation_cuda(torch::Tensor input,
150288
            torch::Tensor weights, torch::Tensor bias,
1503
                                                         float clamp_min, float
1504
            clamp_max) {
1505<sub>90</sub>
                     // Dimension extraction
1506 91
                     int N = input.size(0);
                     int C_in = input.size(1);
150792
                     int D = input.size(2);
1508<sup>93</sup>
                     int H = input.size(3);
    94
1509<sub>95</sub>
                     int W = input.size(4);
1510<sub>96</sub>
                     int C_out = weights.size(0);
151197
                     int kD = weights.size(2);
                     int kH = weights.size(3);
   98
```

```
1512
                     int kW = weights.size(4);
1513
1514101
                     // Output dimensions
1515102
                     int D out = D - kD + 1;
1516^{103}
                     int H_{out} = H - kH + 1;
1517 104
                    int W_{out} = W - kW + 1;
   105
1518
                     // Create output tensor
1519107
                    auto output = torch::empty({N, C_out, D_out, H_out, W_out},
1520
           input.options());
1521108
1522 109
                     // Launch kernel
                     int total_elements = N * C_out * D_out * H_out * W_out;
   110
1523
                     const int block_size = 256;
1524112
                    const int grid_size = (total_elements + block_size - 1) /
           block_size;
1525
                    conv_activation_kernel<<<grid_size, block_size>>>(
152613
                         input.data_ptr<float>(), weights.data_ptr<float>(),
1527 114
           bias.data_ptr<float>(),
1528
                         output.data_ptr<float>(), N, C_in, D, H, W, C_out, kD,
1529
           kH, kW,
153016
                         clamp_min, clamp_max
                    );
1531117
1532<sup>118</sup>
                     return output;
   119
1533
                .....
1534121
                conv_activation_header = """
1535122
                torch::Tensor conv_activation_cuda(torch::Tensor input,
1536<sup>123</sup>
           torch::Tensor weights, torch::Tensor bias,
1537
                                                       float clamp_min, float
1538
           clamp_max);
1539<sub>125</sub>
1540126
                # Load the CUDA kernel
1541127
1542 128
                self.conv_activation = load_inline(
   129
                    name="conv_activation",
1543
                    cpp_sources=conv_activation_header,
1544131
                    cuda_sources=conv_activation_source,
                    functions=["conv_activation_cuda"],
1545132
1546<sup>133</sup>
                    verbose=True
1547 134
                )
   135
1548
           def forward(self, x):
1549137
                return self.conv_activation.conv_activation_cuda(x, self.weight,
1550
           self.bias, self.clamp_min, self.clamp_max)
1551
```

On turn 2 it attempts to improve on the existing implementation, but fails due to a compilation error.

1552

1553

1554

1555

1556

1557

15581559

1560

1561

On turn 3, it attempts to modify the loop order to improve coalesced memory access, but this actually worsens performance, as the kernel now exhibits a speedup of **0.61x**. The model attempts this strategy again at turn 4, and is able to order the loops in a way that increases memory coalescing: this brings up the performance to **1.19x**. At turn 5 and 6, the model attempts some minor optimizations (such as precomputing certain indices, or deleting intermediate variables), which however has no measurable effect in performance.

On turn 7, the model figures out a surprising optimization: it records the convolution's kernel size as a compile time constant, which allows it to use # pragma unroll to unroll the 3 loops needed to apply the convolution. This results in a significant speedup of 1.93x over baseline.

```
import torch
from torch.utils.cpp_extension import load_inline
import math

class ModelNew(torch.nn.Module):
```

```
1566
            def __init__(self, in_channels, out_channels, kernel_size,
1567
            clamp_min, clamp_max):
1568 7
                super(ModelNew, self).__init__()
1569 8
                 self.clamp_min = clamp_min
                self.clamp_max = clamp_max
1570 <sup>9</sup>
1571 <sup>10</sup>
                 # Initialize convolution weights and bias
    11
1572 12
                 self.weight = torch.nn.Parameter(torch.empty(out_channels,
1573
            in_channels, kernel_size, kernel_size, kernel_size))
157413
                 self.bias = torch.nn.Parameter(torch.empty(out_channels))
                 # Weight initialization following PyTorch's Conv3d default
1575 14
1576 <sup>15</sup>
                 torch.nn.init.kaiming_uniform_(self.weight, a=math.sqrt(5))
1577
                 fan_in, _ =
            torch.nn.init._calculate_fan_in_and_fan_out(self.weight)
1578 17
                 bound = 1 / math.sqrt(fan_in)
1579 18
                 torch.nn.init.uniform_(self.bias, -bound, bound)
1580 <sup>19</sup>
                 # Define the fused convolution and activation CUDA kernel with
1581<sup>20</sup>
            optimizations
1582
                 kernel_size_val = kernel_size
1583<sub>22</sub>
                 conv_activation_source = f"""
158423
                 #include <torch/extension.h>
                 #include <cuda_runtime.h>
1585 24
                 #include <math.h>
1586<sup>25</sup>
1587 <sub>27</sub>
                 #define KERNEL_SIZE {kernel_size_val}
1588<sub>28</sub>
                 #define KERNEL D {kernel_size_val}
                 #define KERNEL_H {kernel_size_val}
1589 29
                 #define KERNEL_W {kernel_size_val}
1590^{30}
1591<sup>31</sup>
                  _global__ void conv_activation_kernel(
1592 33
                     const float* __restrict__ input, const float* __restrict_
1593
            weights, const float* __restrict__ bias,
                     float* __restrict__ output,
int N, int C_in, int D, int H, int W, int C_out,
1594 34
1595<sup>35</sup>
1596<sup>36</sup>
                      float clamp_min, float clamp_max) {{
1597 37
1597 38
                          int idx = blockIdx.x * blockDim.x + threadIdx.x;
                          const int W_out = W - KERNEL_W + 1;
1598 39
                          const int H_out = H - KERNEL H + 1;
                          const int D_out = D - KERNEL_D + 1;
1599 40
                          const int total_elements = N * C_out * D_out * H_out *
1600<sup>41</sup>
            W out:
1601
    42
                          if (idx >= total_elements) return;
1602 43
160344
                          int w_out = idx % W_out;
1604 45
                          int h_out = (idx / W_out) % H_out;
                          int d_out = (idx / (W_out * H_out)) % D_out;
160546
1606<sup>47</sup>
                          int f = (idx / (W_out * H_out * D_out)) % C_out;
1607 48
1607 49
                          int n = idx / (C_out * D_out * H_out * W_out);
1608 50
                          float sum = 0.0;
1609 51
1610<sup>52</sup>
                          // Precompute constants for GELU
1611 <sup>53</sup>
                          const float sqrt_2_over_pi = 0.79788456f;
                          const float approx_coeff = 0.044715f;
1612 54
1612 55
1613<sub>56</sub>
                          for (int c = 0; c < C_in; ++c) {{
                               // Compute input channel base pointer
161457
                               const int input_c_offset = n * C_in * D * H * W + c
1615<sup>58</sup>
            * D * H * W;
1616
                               const float* input_c = input + input_c_offset;
1617 60
1618<sub>61</sub>
                               // Compute center position in the input channel
                               const int center_offset = d_out * H * W + h_out * W
161962
            + w_out;
```

```
1620
63
                               const float* center = input_c + center_offset;
1621 64
162265
                               // Weight base for current output channel f and
1623
            input channel c
                               const int weight_base = f * C_in * KERNEL_D *
162466
            KERNEL H * KERNEL W + C * KERNEL D * KERNEL H * KERNEL W;
1625
1626 <sub>68</sub>
                               #pragma unroll
162769
                               for (int kd = 0; kd < KERNEL_D; ++kd) {{
162870
                                   #pragma unroll
                                   for (int kh = 0; kh < KERNEL_H; ++kh) {{
162971
                                        #pragma unroll
1630 <sup>72</sup>
                                        for (int kw = 0; kw < KERNEL_W; ++kw) {{
1631 74
                                            const int offset = kd * H * W + kh * W +
1632
            kw;
163375
                                            const float input_val = center[offset];
                                            const int w_off = weight_base + kd *
1634 76
            KERNEL_H * KERNEL_W + kh * KERNEL_W + kw;
1635
                                            const float weight_val = weights[w_off];
1636 <sub>78</sub>
1637 79
                                            sum += input_val * weight_val;
1638 80
                                        } }
163981
                                   } }
1640 82
                          } }
    83
1641 84
1642<sub>85</sub>
                          sum += bias[f];
1643 86
                          // Apply activations with optimized order
164487
                          { {
1645 88
                               float tanh_val = tanhf(sum);
1646<sub>90</sub>
                               float clamped = fmaxf(clamp_min, fminf(tanh_val,
1647
            clamp_max));
1648 91
                               float inner = clamped + approx_coeff * clamped *
            clamped * clamped;
1649
1650<sup>92</sup>
                               inner *= sqrt_2_over_pi;
                               sum = clamped * 0.5f * (1.0f + tanhf(inner));
1651<sub>94</sub>
165295
                          // Output index calculation
1653 96
                          const int output_base = n * C_out * D_out * H_out *
1654<sup>97</sup>
1655
            W_out +
                                                      f * D_out * H_out * W_out +
1656 99
                                                      d_out * H_out * W_out +
1657100
                                                      h_out * W_out;
1658101
                          output[output_base + w_out] = sum;
1659102
                     } }
1660^{103}
104
                 torch::Tensor conv_activation_cuda(torch::Tensor input,
            torch::Tensor weights, torch::Tensor bias,
166205
                                                        float clamp_min, float
            clamp_max) {{
1663
                     int N = input.size(0);
1664^{106}
1665 107
                     int C_in = input.size(1);
                     int D = input.size(2);
   108
1666
                     int H = input.size(3);
1667110
                     int W = input.size(4);
166811
                     int C_out = weights.size(0);
1669112
                     int D_out = D - KERNEL_D + 1;
1670 113
1670
114
1671<sub>115</sub>
                     int H_out = H - KERNEL_H + 1;
                     int W_out = W - KERNEL_W + 1;
1672<sub>16</sub>
                     auto output = torch::empty({{N, C_out, D_out, H_out,
1673117
            W_out}}, input.options());
```

```
1674
1675
                     int total_elements = N * C_out * D_out * H_out * W_out;
167620
                     const int block_size = 256;
                     const int grid_size = (total_elements + block_size - 1) /
1677121
           block size;
1678
1679<sup>122</sup>
123
                     conv_activation_kernel<<<grid_size, block_size>>>(
1680<sub>124</sub>
                          input.data_ptr<float>(), weights.data_ptr<float>(),
1681
           bias.data_ptr<float>(),
1682125
                          output.data_ptr<float>(), N, C_in, D, H, W, C_out,
1683126
                          clamp_min, clamp_max
1684 127
                     );
                     return output:
   128
1685,129
1686130
1687131
                 conv_activation_header = """
1688<sup>132</sup>
                 torch::Tensor conv_activation_cuda(torch::Tensor input,
1689 133
            torch::Tensor weights, torch::Tensor bias,
1690,34
                                                         float clamp_min, float
1691
            clamp_max);
1692135
1693136
                 # Load the CUDA kernel with fast math optimization
1694 137
   138
                 self.conv_activation = load_inline(
1695<sub>139</sub>
                     name="conv_activation",
1696140
                     cpp_sources=conv_activation_header,
1697141
                     cuda_sources=conv_activation_source,
                     functions=["conv_activation_cuda"],
1698142
1699<sup>143</sup>
144
                     extra_cuda_cflags=['-use_fast_math'],
                     verbose=True
1700<sub>145</sub>
                 )
1701_{\!146}
            def forward(self, x):
1702147
                 return self.conv_activation.conv_activation_cuda(x, self.weight,
1703148
            self.bias, self.clamp_min, self.clamp_max)
1704
```

In its final turn, the model attempts a more advanced implementation that further parallelizes the computation across kernels before performing a warp-level reduction. However, it fails to implement the strategy correctly, due to applying the reduction across the wrong axis. We do note the model has shown success in implementing complex warp reductions in several other tasks.

I ERROR CORRECTION BEHAVIOR

1705

1706

1707

1708

1709 1710 1711

1712 1713

1714

1715

1716

1717

1718

1719

1727

An important quality for kernel generation is to correct errors over turns. With multi-turn training, Kevin learns to effectively manage kernel writing across multiple turns, making more aggressive optimizations while correcting errors more effectively. This is shown by our results in Figure 5, where Kevin exhibits better scaling behavior across serial turns (compared to its single-turn counterpart). Here we focus on an example that illustrates error correction behavior.

We consider the KernelBench task of CosineSimilarityLoss (Level 1 Task 97):

Test-time Trajectory from Single-Turn RL:

Turn 1: The model implements a correct kernel on the first try. The kernel is straightforward yet performant, obtaining a 3.15x speedup over the reference implementation. This is expected as the single-turn model is trained to maximize single-turn correctness.

Turn 2: The model attempts to calculate the block size dynamically; however, it results in suboptimal GPU utilization, causing a decrease in runtime to 2.19x.

Turns 3-5: The model reverts the change from Turn 2, and adds minor improvements (such as performing the mean calculation at the warp level rather than using a single thread, or tweaking the loading of shared memory); however, it has no significant impact on performance (the speedups range from 3.11x to 3.19x).

Turns 6-8: The model attempts to tune the grid size and block size, though this fails to improve performance (the speedups ranging from 2.99x to 3.14x).

Test-time Trajectory from Multi-Turn RL:

1731

1732

1733

1734

1735

1736

1737 1738

1739

1740

1741

1742

1743

1744 1745

1746

1747

1748

1749

1750

1751

1752

1753

1754

17651766

1767

1768

1769

1770

1771

1772

Turn 1: The model attempts a kernel similar to the single-turn implementation. A suboptimal block size and a non-contiguous access, however, cause a speedup of 1.74x, inferior to the single turn version.

Turn 2: The multi-turn model makes various optimizations (such as tweaking block and grid size, and avoiding recomputing certain values), but it forgets to include a required C++ header, causing a compilation error.

Turn 3: The multi-turn model adds the correct C++ header and improves the memory coalescing. This, combined with the previous optimizations, causes a speedup of 3.31x.

Turn 4: The model now attempts more aggressive optimization techniques to compute the sums within the kernel, namely using warp intrinsics to perform warp-level reductions. The implementation is incorrect, however, as only a subset of the threads for a given warp execute the reduction, resulting in a deadlock.

```
1755
       for (int s = 8; s >= 1; s >>= 1) {
1756
            if (thread_id < s) {</pre>
1757 <sub>3</sub>
                 float temp_dot = __shfl_down_sync(0xffffffff, dot, s);
1758 4
                 dot += temp_dot;
                float temp_p = __shfl_down_sync(0xFFFFFFFF, p_norm_sq, s);
1759 5
                 p_norm_sq += temp_p;
1760 <sup>6</sup>
                 float temp_t = __shfl_down_sync(0xFFFFFFFF, t_norm_sq, s);
1761 8
                 t_norm_sq += temp_t;
1762<sub>9</sub>
1763 10
              _syncthreads();
1764 11
```

Turn 5: The model then tries a different strategy, implementing a simpler reduction that avoids using warp intrinsics. However, this proves less efficient than the previous methods, and achieves a speedup of only 2.26x.

Turn 6: The model now attempts to fix the errors present in the advanced warp level technique from turn 3, spending most of the chain-of-thought in trying to address the previous error. This time, the model is successful with a speedup of 7.81x.

```
1773
       for (int s = 16; s >= 1; s >>= 1) {
1774
                float temp = __shfl_down_sync(0xFFFFFFFF, dot, s);
1775 <sub>3</sub>
                dot += temp;
1776 4
               temp = __shfl_down_sync(0xFFFFFFFF, p_norm_sq, s);
               p_norm_sq += temp;
1777 5
               temp = __shfl_down_sync(0xFFFFFFFF, t_norm_sq, s);
1778 6
                t_norm_sq += temp;
1779
           }
1780
1781
```

Turns 7, 8: The model tries a few other optimizations, yet has no significant impact on the kernel performance, resulting in slightly worse kernels. The best kernel is thus achieved at turn 6 with a speedup of 7.81x, compared to the best kernel found by single-turn (3.19x).

Comparing the trajectories of the multi-turn vs single-turn model, we notice that the single-turn model is generally careful in its optimizations and rarely deviates from the previous attempts. The multi-turn model attempts riskier strategies. While it may result in wrong kernels sometimes, it is more capable of recovering from these errors and eventually finding more performant ones.