

CUDAFORGE: AN AGENT FRAMEWORK WITH HARDWARE FEEDBACK FOR CUDA KERNEL OPTIMIZATION

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Developing efficient CUDA kernels is increasingly critical for AI applications
012 such as large-scale LLM training. However, manual kernel design is both costly
013 and time-consuming, motivating automatic approaches that leverage LLMs for
014 code generation. Existing methods for automatic kernel generation, however, of-
015 ten produce low-efficiency kernels, incur high computational overhead, and fail to
016 generalize across settings.

017 In this work, we propose CudaForge, a training-free multi-agent workflow for
018 CUDA kernel generation and optimization. Our workflow is inspired by the iter-
019 ative workflow of human experts, which contains steps such as developing initial
020 kernels, testing correctness, analyzing hardware feedback, and iterative improve-
021 ment. More specifically, CudaForge employs two LLM agents – a Coder and
022 a Judge – that iteratively generate, correct, and optimize CUDA kernels, while
023 integrating hardware feedback such as Nsight Compute (NCU) metrics. In our
024 extensive evaluations, we show that CudaForge , by leveraging base models
025 like OpenAI-o3, achieves 97.6% correctness of generated kernels and an average
026 1.68 \times speedup over PyTorch baselines, substantially surpassing state-of-the-art
027 models including OpenAI-o3 and Kevin on KernelBench. Beyond accuracy and
028 speed, CudaForge demonstrates strong generalization across GPUs (A100, RTX
029 6000, 4090, 3090) and base models (OpenAI-o3, GPT5, gpt-oss-120B, Claude-
030 Sonnet-4, QwQ-32B), while maintaining high efficiency. In particular, generating
031 an optimized kernel takes about 25 minutes on one RTX 6000 and incurs \$0.30
032 API cost. Our results highlight that multi-agent, training-free workflows can en-
033 able cost-effective, generalizable, and high-performance CUDA kernel optimiza-
034 tion.

1 INTRODUCTION

035
036 **Motivation.** CUDA has become the *de facto* standard for deep learning training because modern
037 frameworks such as PyTorch and TensorFlow are deeply integrated with NVIDIA’s optimized GPU
038 libraries (NVIDIA, 2025b). Efficient CUDA kernels are crucial for accelerating deep learning work-
039 loads(Dao et al., 2022; Dao, 2024) .

040 However, developing high-efficiency cuda kernels has been known as challenging with very high
041 learning curve, requiring deep expertise in GPU architectures and parallel programming(Li et al.,
042 2024). For example, it took more than 2 years from the debut of the Hopper GPU architecture to the
043 release of FlashAttentionV3 (Shah et al., 2024), which is specially designed for Hopper GPUs.

044 This high development barrier has driven growing interest in finding automated ways of generating
045 highly efficient and customized CUDA kernels. For example, some work (Tillet et al., 2019) (Chen
046 et al., 2018) employs auto-tuning and evolutionary search to automatically explore kernel imple-
047 mentation spaces and optimize low-level parameters for specific hardware. More recently, there has
048 been a growing interest in leveraging large language models (LLMs) to perform such tasks. LLM is
049 believed to hold great promise in generating efficient and high-quality kernels, due to its capability
050 of code generation in other domains, such as Python, C++ and [Triton](#) (Dong et al., 2025; Jiang et al.,
051 2024; Anonymous, 2025; Li et al., 2025b; Woo et al., 2025; Li et al., 2025a).

052 **Existing Works and Key Challenges.** Generally, using LLMs for CUDA kernel generation is still
053 in an early stage. In KernelBench (Ouyang et al., 2025), the authors attempt to directly use state-of-

054 the-art (SOTA) models, such as OpenAI-o1 and Claude-3.5-Sonnet, to generate kernels. However,
 055 it has been observed that these SOTA models still struggle to produce correct or performant kernels
 056 out of the box, revealing fundamental limitations of existing LLMs in this domain.

057 To address this gap, recent studies have explored two main paradigms. The first approach is based
 058 on reinforcement learning (RL) (Schulman et al., 2017; Shao et al., 2024). CUDA-L1 (Team, 2025)
 059 and Kevin (Baronio et al., 2025) adopt RL to enhance LLMs’ ability to generate correct and per-
 060 formant CUDA code. The second approach is based on AI agents. In particular, in an independent and
 061 contemporaneous work (Lange et al., 2025)¹, researchers have explored agentic frameworks at in-
 062 ference time. Agents project PyTorch method into CUDA kernel design, then the CUDA kernels are
 063 further refined by sampling new kernels and verification filtering. This design effectively improves
 064 correctness in CUDA kernel generation without the high cost of RL training.

065 Despite these advances, several key challenges remain:

066 **(C1) Limited kernel efficiency.** While RL-based methods improve LLMs’ ability to generate
 067 CUDA kernels, their optimization capability remains insufficient. For example, Kevin-32B only
 068 achieves an average speedup of $1.10\times$ over KernelBench test cases, even after sampling 16 parallel
 069 trajectories with 8 refinement turns each per kernel (Baronio et al., 2025). Further, CUDA-L1 of-
 070 ten fails to directly optimize the CUDA kernels, but producing official implementation of PyTorch
 071 (Team, 2025) (see Appendix F for details).

072 **(C2) High training and inference cost.** RL-based approaches such as (Team, 2025; Baronio et al.,
 073 2025) require substantial computational resources and long training cycles, making them unsuitable
 074 for low-resource or rapid-prototyping settings. In addition, multi-stage agentic pipeline developed
 075 by (Lange et al., 2025) incurs high inference costs (about 6 H100 hours and \$5 API cost per kernel),
 076 which greatly limits its practical applicability of the approach.

077 **(C3) Lack of hardware feedback.** Human experts typically follow an iterative workflow to develop
 078 performant CUDA kernels through testing and refinement. They rely on hardware feedback like
 079 Nsight Compute (NCU)² to identify bottlenecks and optimize kernels accordingly (Wu et al., 2025;
 080 NVIDIA, 2025a; Hu et al., 2025). In contrast, RL-based approaches (Team, 2025; Baronio et al.,
 081 2025) train LLMs to directly generate or optimize kernels, but do not incorporate hardware feedback
 082 at all. As a result, they rely on blind exploration during generation, lacking the targeted guidance.
 083 This often leads to suboptimal kernel efficiency, limiting their practical applicability.

084 These challenges raise a natural question: *Can we design a simple but effective hardware-aware
 085 approach that reliably produces efficient CUDA kernels at low cost?*

086 **Our Contributions.** To address these challenges, we propose CudaForge, **a simple, effective and
 087 low-cost** multi-agent workflow for CUDA kernel generation and optimization, as shown in Figure 1.
 088 Our workflow is inspired by the iterative workflow of human experts (Wu et al., 2025; NVIDIA,
 089 2025a; Hu et al., 2025), which contains steps such as developing initial kernels, testing correctness,
 090 analyzing hardware feedback, and iterative improvement.

091 This workflow involves two specialized LLM agents that iteratively generate and optimize CUDA
 092 kernels: a Coder, which generates kernels given task instructions and Judge feedback, and a Judge,
 093 which analyzes kernels and hardware feedback to guide the Coder generation. One key novelty
 094 of CudaForge is its integration of external hardware feedback, including GPU specifications and
 095 Nsight Compute (NCU) metrics, enabling the Judge to identify performance bottlenecks like human
 096 experts and provide targeted optimization guidance to the Coder.

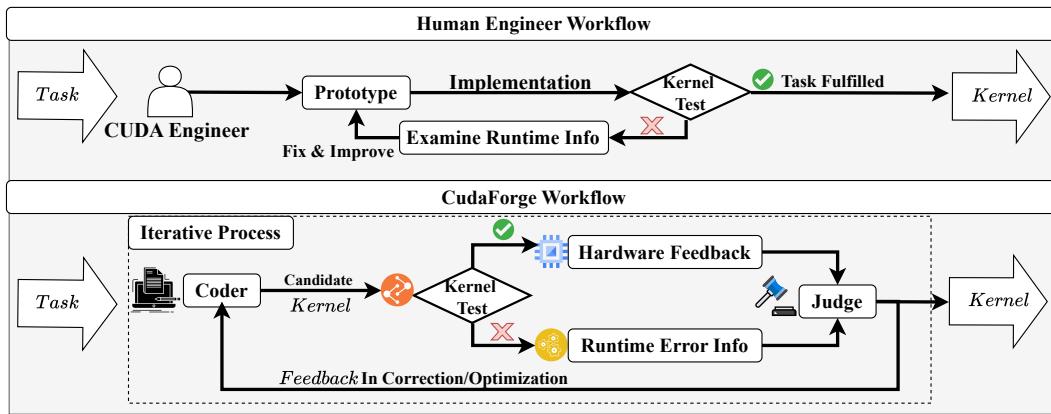
097 Compared to single-LLM approaches that generate and evaluate code using the same LLM, our
 098 framework separates these roles into an *independent* Coder and Judge, enabling more special-
 099 ized reasoning and more reliable iterative refinement. Unlike RL-based methods, CudaForge is
 100 training-free, avoiding the substantial cost of policy training. It is also hardware-aware, allowing
 101 it to tailor CUDA kernel optimizations to the underlying system, making the proposed framework
 102 easily generalizable across different GPUs. Finally, in contrast to existing multi-agent frameworks
 103 (Lange et al., 2025), CudaForge is lightweight and cost-efficient, running in just 25 minutes on a
 104 single RTX6000 GPU and \$0.3 per kernel in API costs, while still achieving better performance.

1published on arxiv Sept 16th, 2025

2Nsight Compute (NCU) is NVIDIA’s official kernel-level profiler for CUDA programs.

108 We evaluate CudaForge on 250 KernelBench tasks from Level 1 to Level 3. Though these tasks
 109 are challenging, CudaForge attains a 97.6% correctness rate and delivers an average speedup of
 110 1.68 \times over PyTorch baselines, which significantly outperforms advanced RL model like Kevin-32B
 111 and advanced frontier model like OpenAI-o3 (OpenAI, 2025). Further, we have conducted com-
 112 prehensive ablation studies of the features of CudaForge, such as its effectiveness across multiple
 113 GPU architectures, its inference-time scalability by increasing the number of generation, and the ef-
 114 fect of different base models. Overall, we observed that the proposed CudaForge achieves robust
 115 performance in all these settings.

116 These findings highlight the key contribution of this work: The proposed LLM agent workflow
 117 CudaForge is simple but effective: at very low cost, it develops performant CUDA kernels for
 118 many practical tasks, for a variety of GPU architectures and base models. It also exhibits strong
 119 test-time scaling capabilities where solution quality can improve substantially while increasing its
 120 iteration rounds. These results demonstrate CudaForge’s strong practical applicability.



135 Figure 1: Comparison between human and CudaForge workflows. Top: Human experts iter-
 136 atively refine kernels by writing a prototype, testing it, and analyzing runtime feedback. Bottom:
 137 CudaForge mimics human workflow with two specialized agents (Coder and Judge). The Coder
 138 generates candidate kernels, while the Judge analyzes runtime info and hardware feedback to pro-
 139 vide correction or optimization feedback. The process iterates until it reaches maximum round N .
 140

2 THE CUDAFORGE FRAMEWORK FOR CUDA KERNEL OPTIMIZATION

2.1 CUDAFORGE FRAMEWORK

145 Given a CUDA kernel generation task, the objective is to generate a kernel that is functionally
 146 equivalent to its PyTorch reference while achieving the lowest possible execution latency.
 147

148 Inspired by the iterative workflow of human experts (Wu et al., 2025; NVIDIA, 2025a; Hu et al.,
 149 2025), we design CudaForge as an iterative multi-agent framework, illustrated in Figure 1. The
 150 framework involves two independent agents: a **Coder** and a **Judge**. The Coder generates candidate
 151 kernels based on the task description and feedback from the Judge, while the Judge evaluates each
 152 candidate using the kernel itself, hardware feedback and runtime information.

153 More specifically, given a CUDA kernel generation task, the Coder first receives the task require-
 154 ments and PyTorch reference implementation, then produces an initial candidate kernel. This can-
 155 didate is compiled and executed on test cases to check correctness. If it fails, the Judge inspects
 156 *runtime information* (e.g., compilation errors, mismatched outputs with the PyTorch reference) and
 157 analyzes the faulty kernel. It then returns correction feedback (e.g., missing header file) to guide the
 158 next iteration. Once a kernel candidate passes the correctness test, the Judge profiles it with the NCU
 159 tool to obtain NCU metrics (e.g., memory throughput, occupancy, warp efficiency). Together with
 160 GPU specifications, these metrics form the *hardware feedback* that allows the Judge to identify the
 161 dominant bottleneck (e.g., compute-bound or memory-bound) and provide one specific optimization
 162 feedback (e.g. using shared memory) to the Coder.

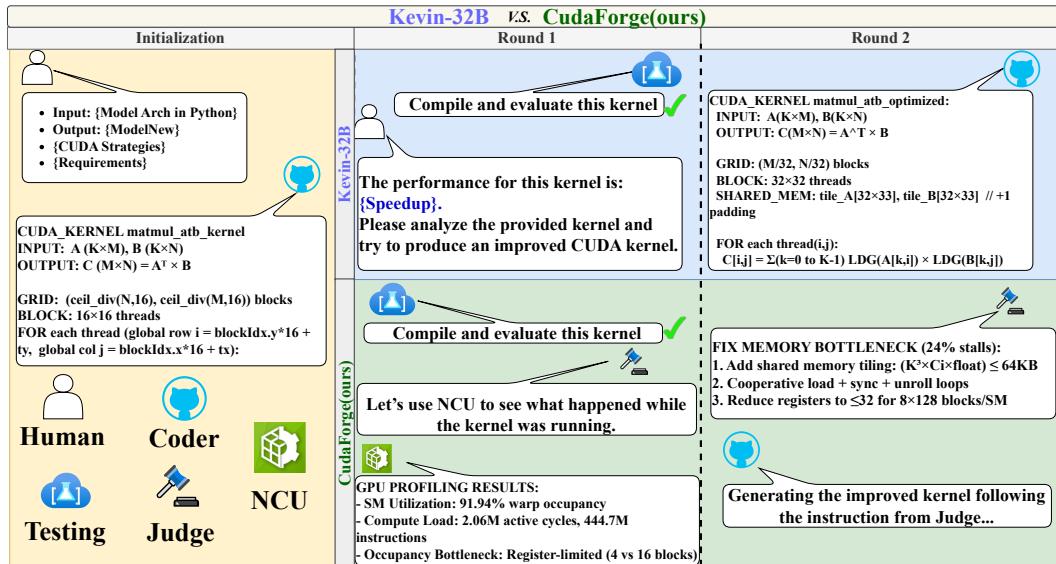


Figure 2: The overview of conversation between agents in Kevin-32B and CudaForge.

In the next iteration, the Coder is prompted with the previous kernel, Judge feedback, and the original task requirements, and generates a corrected or optimized kernel. This process repeats for up to N iterations, after which we select the most efficient correct kernel as the final solution.

CudaForge achieves reliability and efficiency through **three key design choices**. First, it adopts a two-agent system where the Coder focuses on generation and the Judge on evaluation, separating the “cognitive” load (See Section 3.4). The Coder receives only feedback from the Judge, while the Judge uses hardware and runtime information to guide generation and optimization. This division of labor mirrors human workflows and mitigate the risk of overlooking errors or inefficiencies. Second, the framework follows an iterative optimization process, progressively correcting errors and improving efficiency across rounds. This enables stable refinement, especially on hard tasks. Third, it explicitly incorporates hardware feedback, such as GPU specifications and NCU metrics, so the Judge can pinpoint bottlenecks and provide actionable guidance to the Coder. This targeted optimization avoids blind exploration and ensures directed performance gains.

2.2 HOW TO INTEGRATE HARDWARE FEEDBACK

In this subsection, we describe in detail a key design consideration, which enables CudaForge to utilize hardware feedback for kernel performance optimization. The hardware feedback module integrates static GPU specifications (e.g. architecture, memory bandwidth, per-thread register limits, per-SM shared-memory capacity) with performance metrics (e.g. memory throughput, occupancy, and warp efficiency) from Nsight Compute (NCU) collected during kernel execution. By cross-referencing GPU specifications and NCU metrics, the Judge infers the kernel’s primary performance-limiting cause and bottleneck mechanism. Figure 2 illustrates how Judge uses the hardware feedback to optimize kernels.

Just as CUDA engineers focus on key indicators, we do not pass the entire set of NCU metrics to the Judge. Feeding all metrics can overwhelm the decision process with excessive, partially redundant signals and lead to unstable judgments (See Appendix E.1 for detail). Instead, we design a novel protocol which profiles a subset of critical metrics provided by NCU and forward them to Judge so that we can improve the quality of the judge outputs. More specifically, the key subset of metrics are selected off-line (before the agent start to work), through the following steps:

(Step 1) Kernel sampling and Selection: We first profile key metrics on some preselected representative tasks (e.g., Conv2D, MatMul) to prepare a reliable metric set. Specifically, for each task we run 100 self-refine (repeating the cycle generating → execute/profile → evaluate → repair/optimize) with a single SOTA model (e.g., OpenAI-o3), collect the generated and correct kernels, and select 10 with the largest speed disparity (fastest vs. slowest). See Algorithm 1.

216 **Algorithm 1** Step 1: Kernel Sampling and Selection
 217 **Input:** Task set $Task = \{T_1, T_2, \dots, T_n\}$
 218 **Output:** Selected subsets K_i^* for each task T_i
 219 **for** $i \leftarrow 1$ **to** n **do**
 220 $K_i \leftarrow \emptyset$ **for** $j \leftarrow 1$ **to** 100 **do**
 221 $k_j \leftarrow \text{generate_kernel}(T_i)$ $K_i \leftarrow K_i \cup \{k_j\}$
 222 **end**
 223 Sort K_i in nondecreasing order according to kernel runtime $m \leftarrow |K_i|$; // Here $m = 100$
 224 $K_i^* \leftarrow \{K_i[1], K_i[2], K_i[3], K_i[4], K_i[5], K_i[m-4], K_i[m-3], K_i[m-2], K_i[m-1], K_i[m]\}$
 225 **end**
 226

227 **(Step 2)** Top-20 metrics within each task: We then refine the metrics within each task to identify
 228 the most relevant candidates. Specifically, for each task we consolidate the NCU metrics from the
 229 10 kernels selected from Step 1 into a single dataset. Since Nsight Compute reports a consistent
 230 full set of metrics across all kernels, the metric categories are aligned by default. We then remove
 231 aliases and strongly collinear indicators, and compute Pearson correlations between each metric and
 232 kernel runtime. We retain only the Top-20 metrics (by absolute correlation) as the candidate set for
 233 that task (see Appendix E.2 for examples).

234 **(Step 3)** Metrics selection across-tasks: Finally, we consolidate metrics across tasks to build a
 235 stable, task-agnostic set. We compare the Top-20 lists across tasks and keep metrics that consistently
 236 appear, show the same correlation direction, and achieve high global scores. This yields 24 metrics
 237 that are strongly correlated with kernel runtime across tasks. See Algorithm 2.

238 **Algorithm 2** Step 2-3: Profiling and Metrics Selection
 239 **Input:** $K^* = \{K_1^*, K_2^*, \dots, K_n^*\}$, where each $K_i^* = \{k_1^*, k_2^*, \dots, k_{10}^*\}$
 240 **Output:** Final metrics set $Final_Metrics$ containing 24 unique metrics
 241 $M^* \leftarrow \emptyset$
 242 **for** $i \leftarrow 1$ **to** n **do**
 243 $M_i^* \leftarrow \emptyset$ **foreach** $k \in K_i^*$ **do**
 244 $M \leftarrow \text{NCU_Profile}(k)$; // Run NCU profiling, $M = \{m_1, m_2, \dots, m_j\}$
 245 **foreach** $m \in M$ **do**
 246 | Compute Pearson correlation coefficient $r(m, \text{runtime}(k))$
 247 **end**
 248 $Top20(k) \leftarrow$ the 20 metrics in M with highest $|r(\cdot, \text{runtime}(k))|$ $M_i^* \leftarrow M_i^* \cup Top20(k)$
 249 **end**
 250 $M^* \leftarrow M^* \cup M_i^*$
 251 **end**
 252 // Final set contains 24 distinct metrics
 253 $Final_Metrics \leftarrow \bigcap_{i=1}^n M_i^*$ $|Final_Metrics| = 24$

254
 255 After the above steps are completed offline, during kernel optimization, the Judge profiles each generated
 256 kernel with NCU and uses only this 24 metrics as references (see Appendix E.3 for details).

257 Overall, at each iteration, the Judge collects hardware feedback, including static GPU specifications
 258 and the key subset of NCU metrics. Based on this information, the Judge identifies the dominant
 259 bottleneck by analyzing the 24 metrics and runtime log. To prevent AI agent searching without di-
 260 rection and generating random results, the Judge only captures 3-4 most important metrics in each
 261 round according to its own reasoning. For example, Judge can identify the current kernel is memory-
 262 bound when memory throughput is high but computing resources utilization is low, and then it will
 263 choose memory related metrics as critical metrics in this round. After this, Judge will generate sug-
 264 gestions on how to modify the kernel to address current critical bottleneck. The Coder incorporates
 265 this guidance in the next round generation accordingly. This mechanism enables our multi-agents
 266 system focus on addressing only one critical program bottleneck in each round, eventually optimize
 267 overall kernel performance step by step in iterative rounds, just like human expert's real workflow.

268
 269

270 3 EXPERIMENTS
271272 3.1 BENCHMARK AND EVALUATION
273

274 We evaluate our method on **KernelBench** (Ouyang et al., 2025), a popular benchmark designed to
275 assess the ability of LLMs to generate CUDA kernels. KernelBench consists of multiple difficulty
276 levels; we adopt all tasks from Level 1 to Level 3, resulting in a total of 250 tasks. Specifically, Level
277 1 contains relatively simple 100 tasks involving basic operators (e.g., matrix multiplication), Level
278 2 includes medium-difficulty 100 tasks composed of multi-step operator combinations, and Level 3
279 contains 50 challenging tasks involving full neural network architectures (e.g., AlexNet). Each task
280 is accompanied by a reference PyTorch implementation and predefined input/output specifications,
281 which enables fully automated and reliable evaluation of both correctness and performance.

282 We evaluate model performance on KernelBench using the following metrics:

283 (1) **Correctness**: the fraction of tasks for which the generated kernel compiles successfully and
284 produces outputs identical to the PyTorch reference on all test cases. (2) **Performance**: the ratio of
285 the execution speed (tested on a specific GPU), between a correct generated kernel and its PyTorch
286 reference. (3) **fast_p**: the proportion of correct kernels whose execution speed exceeds $p \times$ that of
287 the PyTorch reference (e.g., fast₁ indicates faster than PyTorch). (4) **Median speedup**: the median
288 of ‘Performance’ values across all tasks, reflecting typical rather than average behavior. (5) **75th
289 percentile speedup**: the 75th percentile of Performance values, capturing upper-quartile efficiency.

290 For methods that perform iterative refinement or generate multiple candidates (including
291 CudaForge), we report the best-performing correct kernel among all candidates for each task. [De-
292 tails of test cases, correctness evaluation and performance evaluation could be found in Appendix B.](#)
293

294 3.2 SETTINGS & BASELINES
295

296 In our main results, we instantiate CudaForge with OpenAI-o3 as both the Coder and the Judge
297 as our *default* setting. We set the maximum number of iteration rounds to $N=10$ to balance perfor-
298 mance improvements and inference cost. Unless otherwise stated, all methods are evaluated under
299 the same compilation/runtime environment in Quadro RTX 6000 and task-specific test suites.

300 To contextualize the performance of CudaForge and assess the effect of advanced foundation mod-
301 els, we include the following baselines for main results and ablation studies: (1) O3-S: OpenAI-o3
302 (single-shot), one-pass generation without iteration; (2) O3-10: OpenAI-o3-10-round (self-refine),
303 ten rounds of self-refinement without a Judge, where the model relies solely on itself to correct
304 and optimize kernels given hardware feedback; (3) O3-10-C: OpenAI-o3-10-round (correction-
305 only), a variant of CudaForge where the Judge provides only correctness feedback but no per-
306 formance optimization feedback; (4) O3-10-O: OpenAI-o3-10-round (optimization-only), a variant
307 of CudaForge where the Judge provides only optimization feedback but no correction feedback;
308 (5) Kevin-10: Kevin-32B-10-round(self-refine), the RL-based model run for ten iterative rounds
309 under the same protocol; (6) AgentBaseline: the agentic workflow from (Lange et al., 2025), a
310 strong multi-agent baseline. Due to the high computational cost of running Kevin-32B on the full
311 benchmark, we additionally construct a stratified random subset \mathcal{D}^* for fair comparison. Details of
312 KernelBench and \mathcal{D}^* are provided in Appendix G.

313 This suite enables a comprehensive comparison across (i) base model vs. corresponding agent-based
314 method, (ii) the presence/absence of Judge feedback, (iii) RL-based vs. training-free agent-based
315 approaches, and (iv) different agentic methods.

316 3.3 MAIN RESULTS
317

318 Table 1 reports the main results on KernelBench. CudaForge consistently outperforms all base-
319 lines across all metrics, both on the full benchmark \mathcal{D} and on the stratified subset \mathcal{D}^* .

320 On \mathcal{D} , CudaForge attains **97.6%** correctness with an average performance of **1.677 \times** , and **70.8%**
321 **Fast₁**, while achieving a median speedup of 1.107 \times with a 75th percentile speedup of 1.592 \times . This
322 is a clear improvement over its base model O3-S. On \mathcal{D}^* , which allows fair comparison with the
323 advanced RL model Kevin, CudaForge achieves 100% correctness, a median speedup of 1.322 \times ,

324 Table 1: Main results on KernelBench (Level 1-3, 250 tasks). Results of AgentBaseline is on Level
 325 1 and 2. All experiments here are run in RTX 6000. Methods evaluated on \mathcal{D}^* are marked with *.
 326

327 Method	328 Correct ↑	329 Median ↑	330 75% ↑	331 Perf ↑	332 Fast₁ ↑
O3-S	57.6%	0.390	1.014	0.680	31.60%
O3-10	90.8%	1.012	1.209	1.107	55.20%
O3-10-C	97.6%	1.031	1.238	1.222	59.60%
O3-10-O	88.4%	1.061	1.483	1.509	64.00%
AgentBaseline	95.0%	—	—	1.490	—
Kevin-10*	64.0%	0.472	1.047	0.608	36.00%
CudaForge	97.6%	1.107	1.592	1.677	70.80%
CudaForge*	100%	1.322	1.736	1.767	84.00%

337 Table 2: Main results on KernelBench (Level 1-3, 250 tasks) of CudaForge.
 338

339 Task	340 Correct ↑	341 Median ↑	342 75% ↑	343 Perf ↑	344 Fast₁ ↑
Level 1	96%	1.044	1.751	1.448	54.0%
Level 2	100%	1.124	1.427	2.104	89.0%
Level 3	96%	1.081	1.510	1.283	68.0%

345 a 75th percentile speedup of 1.736 \times , an average performance of 1.767 \times , and 84.0% Fast₁. This
 346 substantially surpasses Kevin-10, which reaches only 64.0% correctness, 0.472 \times median, 1.047 \times
 347 at the 75th percentile, 0.608 \times performance, and 36.0% Fast₁. This represents a +63.6% absolute
 348 gain in correctness and a +1.159 \times speedup, despite CudaForge being a training-free method while
 349 Kevin is a RL-trained model.

350 We also compare CudaForge with AgenticBaseline in KernelBench Level 1 and Level 2³. As
 351 shown in Table 2, CudaForge achieves 98% correctness and an average speedup of 1.776 \times , which
 352 outperforms AgenticBaseline (95.0%, 1.490 \times), especially in speedup. This result shows our advan-
 353 tage compared to existing agentic work.

354 Notably, on Level 3—the most challenging tier of KernelBench—CudaForge achieves **96%**
 355 correctness and an average **1.283** \times speedup. Given the complexity of Level 3 tasks, which in-
 356 volve full neural network architectures and multi-stage operations, these results demonstrate that
 357 CudaForge is capable of reliably generating and optimizing highly complex CUDA kernels, where
 358 prior approaches (Baronio et al., 2025; Lange et al., 2025) have not explored it.

359 We evaluate both API and time cost on KernelBench. On average, CudaForge requires only **25**
 360 **minutes** on a single RTX6000 GPU and incurs **\$0.3** API cost per kernel. This is highly cost-efficient
 361 compared with another agentic work (Lange et al., 2025), which reports about **6 GPU hours on**
 362 **H100** and **\$5** per kernel in their Appendix E. These results demonstrate that, by leveraging hardware
 363 feedback, our workflow can rapidly converge to high-quality solutions at low cost. [Details of where](#)
 364 [the 25 minutes is spent could be found in Appendix C.](#)

366 3.4 ABLATION STUDIES

368 **Comparison with O3-10 (self-refinement).** A key motivation behind CudaForge is to decou-
 369 ple the roles of generation and evaluation. In O3-10, the same model performs ten rounds of
 370 self-refinement, implicitly taking on both roles: it must both propose new kernels and evaluate
 371 its own outputs based on hardware feedback and runtime signals. While this strategy raises correct-
 372 ness 57.6% to 92.8%, performance remains limited (1.107 \times speedup, 55.2% Fast₁). In contrast,
 373 CudaForge explicitly separates responsibilities: the Coder focuses on code generation, while the
 374 Judge specializes in providing structured feedback. This division of labor proves critical—allowing
 375 each agent to concentrate on a distinct reasoning process—and results in significantly higher effi-
 376 ciency (1.677 \times speedup, 70.8% Fast₁) without sacrificing correctness.

377 ³Note that these works only report results in Level 1 and 2, and we directly take the results from their paper
 378 since the paper has not opened sourced the code.

Method	Correct \uparrow	Performance \uparrow	Fast $_1\uparrow$
CudaForge-Top 24 metrics(ours)	100%	1.767	84%
CudaForge-Full metrics	100%	1.414	80%
CudaForge-Random 24 metrics	100%	1.655	76%
CudaForge-Top 5 metrics(ours)	100%	1.641	76%
CudaForge-Top 10 metrics(ours)	100%	1.644	80%
CudaForge-Top 20 metrics(ours)	100%	1.827	88%

Table 3: Ablation study on NCU metric selection. Comparing full, random, and top- k subsets shows that concise and carefully chosen metrics (Top-24) provide strong overall performance, while Top-20 offers slightly higher speedup with similar behavior.

Comparison with O3-10-C (correction-only Judge). In O3-10-C, the Judge only provides correction feedback based on runtime signals, without optimization feedback. This setting achieves the same 97.6% correctness as CudaForge, confirming that iterative error correction is sufficient to ensure reliable kernel generation. However, efficiency remains much lower, with only $1.222\times$ performance and 58.8% Fast $_1$. The contrast with CudaForge(**1.677 \times , 70.8%**) highlights that while correctness feedback stabilizes generation, performance feedback—grounded in hardware profiling—is essential for driving substantial efficiency gains.

Comparison with O3-10-O (optimization-only Judge). We also evaluate the variant O3-10-O, where the Judge provides only optimization feedback, without correction feedback. In this setting, the Coder frequently generates kernels that fail to compile or run, since functional errors remain uncorrected. As a result, this setting achieves 88.4% correctness and a $1.509\times$ speedup, which are substantially lower than CudaForge($1.677\times$, 70.8%). The result demonstrates that correction feedback plays a significant role in CudaForge’s performance. The absence of it will lead to lower correctness. And without first ensuring functional validity, optimization feedback alone is ineffective and often wasted.

Ablation study on NCU metrics. A key design choice in CudaForge is to filter the full set of NCU metrics and retain a subset of 24 critical metrics for the Judge. This selective design enables the Judge to focus on the most informative performance indicators while avoiding redundancy and inconsistent feedback. To evaluate this choice, we conduct an ablation study comparing our top-24 metric subset against several variants, including using all NCU metrics, using a random subset of 24 metrics, and using smaller subsets of the top-5, top-10, and top-20 metrics.

As shown in Table 3, two conclusions emerge. First, using the complete set of NCU metrics degrades both correctness and speedup, as the Judge becomes overwhelmed by excessive and partially redundant profiler signals. Second, selecting too few metrics or selecting them randomly restricts the Judge’s ability to provide meaningful optimization feedback, resulting in inferior performance. Our 24-metric design consistently achieves the best overall results across all variants, while the top-20 subset yields slightly higher performance but remains very close in practice.

Furthermore, profiling with all NCU metrics significantly increases computational and API cost: each kernel requires approximately 40 minutes on an RTX 6000 GPU and incurs roughly \$1 in API usage. In contrast, our selective design reduces runtime to about 25 minutes and API cost to \$0.3 while achieving superior performance. These findings demonstrate that concise, carefully curated hardware feedback is both more effective and more efficient than exhaustive profiling. We further provide a case study illustrating this phenomenon in Appendix E.1.

3.5 GENERALIZATION CAPABILITY OF CUDAFORGE

In this section, we analyze CudaForge’s capabilities across various maximum iteration num N , GPU architectures and base models. Considering the high cost of full experiment, we use the stratified subset \mathcal{D}^* for this section.

Scaling up the maximum number of iteration rounds We investigate the effect of the maximum iteration number N on CudaForge’s performance.

432 Table 4: CudaForge’s performance on different GPUs. The system consistently achieves high
 433 correctness and strong performance across architectures by incorporating GPU specifications and
 434 *Nsight Compute* profiling signals during optimization.

436 GPU	437 Correct \uparrow	438 Median \uparrow	439 75% \uparrow	440 Perf \uparrow	441 Fast $_1\uparrow$
437 RTX 6000(Ada Arch-Data center level)	438 100%	439 1.322	440 1.736	441 1.767	442 84.0%
437 RTX 4090(Ada Arch-Desktop level)	438 100%	439 1.188	440 1.589	441 1.327	442 80.0%
437 A100(Ampere Arch-Data center level)	438 100%	439 1.371	440 1.762	441 1.841	442 84.0%
437 RTX 3090(Ampere Arch-Desktop level)	438 100%	439 1.155	440 1.706	441 1.320	442 72.0%

442 Table 5: Performance of CudaForge with different base model combinations. We fix one agent
 443 as OpenAI-o3 and replace the other with various models. All combinations achieve strong results,
 444 showing that the framework is not tied to a specific base model.

446 Models (Coder/Judge)	447 Correct \uparrow	448 Median \uparrow	449 75% \uparrow	450 Perf \uparrow	451 Fast $_1\uparrow$
O3 / O3	100%	1.322	1.736	1.767	84.0%
O3 / GPT-5	100%	1.131	1.561	2.114	96.0%
O3 / Claude	100%	1.265	1.456	1.829	84.0%
O3 / GPT-OSS-120B	100%	1.226	1.490	1.364	76.0%
GPT-5 / O3	100%	1.125	1.388	1.896	72.0%
Claude / O3	88%	1.052	1.207	1.398	56.0%
GPT-OSS-120B / O3	96%	1.080	1.477	1.653	68.0%
QwQ / O3	84%	0.965	1.153	0.790	44.0%

452 As shown in Figure 3, increasing N from 1 to 10 leads to substantial performance gains, indicating
 453 that CudaForge can rapidly improve kernel efficiency through iterative refinement. Further
 454 increasing N from 10 to 30 continues to improve performance, though with a slower growth rate,
 455 suggesting that the system gradually approaches its performance ceiling. These results demonstrate
 456 that CudaForge benefits from test-time scaling and has the potential to achieve even stronger per-
 457 formance given larger N with additional inference cost.

458 **Using CudaForge in different GPUs.** We also evaluate CudaForge on various GPU archi-
 459 tectures, including RTX 6000, RTX 4090, RTX 3090 and A100, to examine its effectiveness under
 460 different hardware conditions. As shown in Table 4, CudaForge consistently achieves high cor-
 461 rectness and strong performance on all tested GPUs. This is a direct consequence of its design:
 462 during the optimization phase, the Judge explicitly incorporates hardware feedback, including NCU
 463 metrics and GPU specifications when generating feedback to Coder. This allows the Coder to pro-
 464 duce kernels that are tailored to the target GPU at inference time, without training.

465 **Instantiate CudaForge with various LLM.** To examine whether CudaForge depends on a spe-
 466 cific base model, we conduct experiments by fixing one side (Coder or Judge) as *OpenAI-o3* and
 467 replacing the other with various advanced LLMs, including *QwQ-32B*, *GPT-5*, *Claude*, and *GPT-*
 468 *OSS-120B*. As shown in Table 5, all combinations achieve high correctness and strong performance,

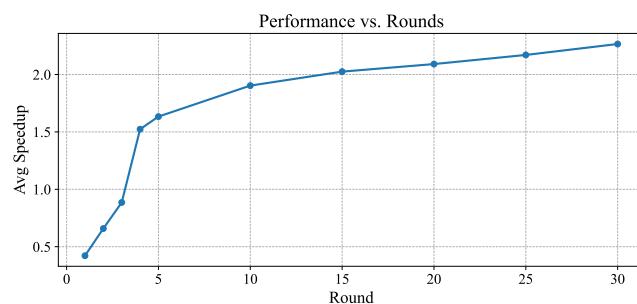


Figure 3: Scaling the number of iteration rounds to 30 on KernelBench (subset \mathcal{D}^*).

486 comparable to or even surpassing the original O3/O3 configuration. These results indicate that
487 CudaForge is not tied to a specific base model: its effectiveness stems from the workflow of
488 Coder and Judge, and it can readily benefit from stronger models as they emerge.
489

490 4 SUPPLEMENT EXPERIMENTS AND OBSERVATIONS 491

492 **Case study.** To comprehensively understand the details of CudaForge, we investigate to a specific
493 case to study its iterative workflow. As shown in Appendix A., it demonstrate a 10-round refine
494 process of KernelBench Level 1 task 95. Our workflow iteratively corrects and optimizes the kernel,
495 with the feedback of Judge model. More details could be found in A.
496

497 **Observations in CUDA-L1 results.** We carefully examined the kernel outputs reported by CUDA-
498 L1 (see Appendix F) and identified an interesting phenomenon that we term “*fake kernels*.” These
499 kernels, while reported as performant, often contain no actual CUDA code. Instead, they rely on
500 `try-except` constructs and fall back to PyTorch’s official implementations to solve the task. This
501 observation highlights a fundamental challenge in evaluating LLM-generated CUDA kernels. To
502 avoid this issue, we have manually checked all kernels in our experiments.
503

504 5 CONCLUSION 505

506 We presented CudaForge, a training-free multi-agent framework for CUDA kernel generation and
507 optimization. The framework mimics the iterative workflow of human experts, explicitly incor-
508 porating hardware feedback to guide targeted kernel refinement rather than blind exploration. On
509 the KernelBench benchmark, CudaForge achieves highest correctness rate and significant perfor-
510 mance gains compared with all existing method, while also demonstrating robustness across diverse
511 GPU architectures and base LLMs. Moreover, its performance scales effectively with the number of
512 refinement rounds. Finally, thanks to its low API and time cost, CudaForge provides a practical
513 and efficient solution for automated CUDA kernel development.
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540
541
ETHICS STATEMENT542
543
544
545
546
This paper proposes the CudaForge framework for automatically generating and optimizing
CUDA kernels, applied to diverse tasks in KernelBench. The design and experiments strictly adhere
to ethical guidelines, ensuring that no sensitive or personally identifiable information is involved.
All experiments rely solely on publicly available benchmarks and standard GPU hardware, and no
human data was collected or processed.547
548
549
550
551
552
We acknowledge the potential environmental concerns related to large-scale model training. While
our framework reduces inference-time cost compared to RL-based methods, more efficient kernels
could indirectly accelerate resource-intensive workloads. We therefore encourage responsible and
sustainable use of this technology. Our framework is intended strictly for research and scientific
purposes, and does not introduce additional risks beyond those already associated with standard
compiler or optimization tools.553
554
REPRODUCIBILITY STATEMENT
555556
557
558
559
560
We have taken extensive measures to ensure the reproducibility of our work. All experiments are
conducted on the publicly available KernelBench benchmark, which provides standardized tasks,
PyTorch references, and input/output specifications. We report detailed results across all difficulty
levels, including averaged metrics and stratified subsets, to ensure statistical robustness.561
562
563
564
To support replication, we provide a comprehensive description of our workflow in Section 2.1 and
include all prompts used for the Coder and Judge agents in Appendix D. Furthermore, we provide
full experimental details, including GPU hardware platforms, evaluation metrics, and iteration pro-
tocols. Since CudaForge is entirely training-free, no additional data collection or model training
is required, greatly simplifying reproducibility.565
566
567
We will release code and experiment scripts upon publication, ensuring that all results reported in
this paper can be faithfully reproduced.568
569
REFERENCES570
571
Anonymous. Tritongym: A benchmark for agentic LLM workflows in triton GPU code generation.
In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025.
572
URL <https://openreview.net/forum?id=oaKd1fVgWc>. under review.
573
574
Carlo Baronio, Pietro Marsella, Ben Pan, Simon Guo, and Silas Alberti. Kevin: Multi-turn rl for
575
generating cuda kernels, 2025. URL <https://arxiv.org/abs/2507.11948>.
576
577
Tianqi Chen, Lianmin Zheng, Eddie Yan, Ziheng Jiang, Thierry Moreau, Luis Ceze, Carlos Guestrin,
578
and Arvind Krishnamurthy. Learning to optimize tensor programs. *Advances in Neural Infor-
579
mation Processing Systems*, 31, 2018.
580
581
Tri Dao. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *Inter-
582
national Conference on Learning Representations (ICLR)*, 2024.
583
584
Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and
585
memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Process-
586
ing Systems (NeurIPS)*, 2022.
587
588
Yihong Dong, Xue Jiang, Jiaru Qian, Tian Wang, Kechi Zhang, Zhi Jin, and Ge Li. A survey on code
589
generation with llm-based agents, 2025. URL <https://arxiv.org/abs/2508.00083>.
590
591
Huangi Hu, Bowen Xiao, Shixuan Sun, Jianian Yin, Zhexi Zhang, Xiang Luo, Chengquan Jiang,
592
Weiwei Xu, Xiaoying Jia, Xin Liu, and Minyi Guo. Liquidgemm: Hardware-efficient w4a8 gemm
593
kernel for high-performance llm serving, 2025. URL <https://arxiv.org/abs/2509.01229>.
594
595
Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language
596
models for code generation, 2024. URL <https://arxiv.org/abs/2406.00515>.

594 Robert Tjarko Lange, Qi Sun, Aaditya Prasad, Maxence Faldor, Yujin Tang, and David Ha. Towards
 595 robust agentic cuda kernel benchmarking, verification, and optimization, 2025. URL <https://arxiv.org/abs/2509.14279>.

596

597 Jianling Li, Shangzhan Li, Zhenye Gao, Qi Shi, Yuxuan Li, Zefan Wang, Jiacheng Huang, Haojie
 598 Wang, Jianrong Wang, Xu Han, Zhiyuan Liu, and Maosong Sun. Tritonbench: Benchmarking
 599 large language model capabilities for generating triton operators, 2025a. URL <https://arxiv.org/abs/2502.14752>.

600

601 Shangzhan Li, Zefan Wang, Ye He, Yuxuan Li, Qi Shi, Jianling Li, Yonggang Hu, Wanxiang Che,
 602 Xu Han, Zhiyuan Liu, and Maosong Sun. Autotriton: Automatic triton programming with rein-
 603 force learning in llms, 2025b. URL <https://arxiv.org/abs/2507.05687>.

604

605 Shiyang Li, Jingyu Zhu, Jiaxun Han, Yuting Peng, Zhuoran Wang, Xiaoli Gong, Gang Wang, Jin
 606 Zhang, and Xuqiang Wang. Onegraph: a cross-architecture framework for large-scale graph
 607 computing on gpus based on oneapi. *CCF Transactions on High Performance Computing*, 6(2):
 608 179–191, 2024.

609

610 NVIDIA. Cuda c++ programming guide. <https://docs.nvidia.com/cuda/cuda-c-programming-guide/>,
 611 2025a. Accessed: 2025-09-21.

612 NVIDIA. Nvidia cudnn. <https://developer.nvidia.com/cudnn>, 2025b. Accessed:
 613 2025-09-21.

614

615 OpenAI. Openai o3 and o4-mini system card. <https://openai.com/index/o3-o4-mini-system-card/>, 2025. Accessed: 2025-09-24.

616

617 Anne Ouyang, Simon Guo, Simran Arora, Alex L. Zhang, William Hu, Christopher Ré, and Azalia
 618 Mirhoseini. Kernelbench: Can llms write efficient gpu kernels?, 2025. URL <https://arxiv.org/abs/2502.10517>.

619

620 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 621 optimization algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.

622

623 Jay Shah, Ganesh Bikshandi, Ying Zhang, Vijay Thakkar, Pradeep Ramani, and Tri Dao.
 624 Flashattention-3: Fast and accurate attention with asynchrony and low-precision, 2024. URL
 625 <https://arxiv.org/abs/2407.08608>.

626

627 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 628 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathe-
 629 matical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

630

631 DeepReinforce Team. Cuda-11: Improving cuda optimization via contrastive reinforcement learning.
 632 *arXiv preprint arXiv:2507.14111*, 2025.

633

634 Philippe Tillet, Hsiang-Tsung Kung, and David Cox. Triton: an intermediate language and compiler
 635 for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International
 636 Workshop on Machine Learning and Programming Languages*, pp. 10–19, 2019.

637

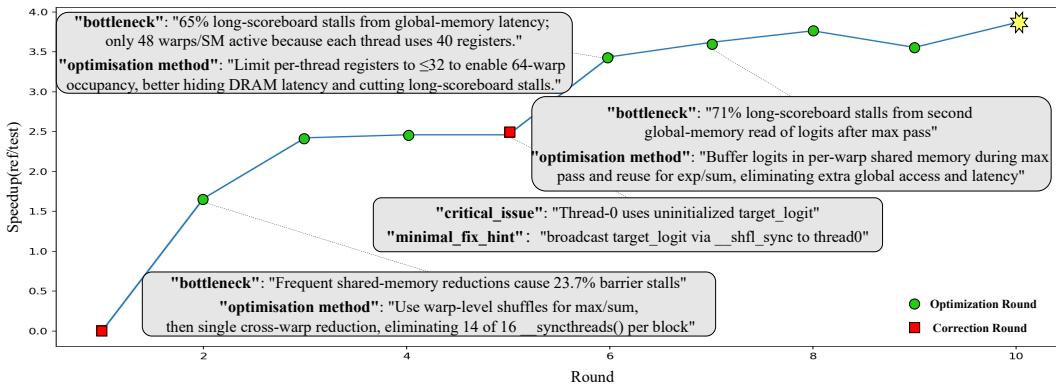
638 Jiin Woo, Shaowei Zhu, Allen Nie, Zhen Jia, Yida Wang, and Youngsuk Park. Tritonrl: Training llms
 639 to think and code triton without cheating, 2025. URL <https://arxiv.org/abs/2510.17891>.

640

641 Min Wu, Huizhang Luo, Fenfang Li, Yiran Zhang, Zhuo Tang, Kenli Li, Jeff Zhang, and Chubo Liu.
 642 Hsmu-spgemm: Achieving high shared memory utilization for parallel sparse general matrix-
 643 matrix multiplication on modern gpus. In *2025 IEEE International Symposium on High Perfor-
 644 mance Computer Architecture (HPCA)*, pp. 1452–1466, 2025. doi: 10.1109/HPCA61900.2025.
 645 00109.

646

647

648 A CASE STUDY
649
650651 A.1 A GOOD CASE
652653 KernelBench L1–95 • CrossEntropyLoss — Judge Outputs & Speedup
654

667 Figure 4: Illustration of the Judge’s outputs—bottleneck diagnoses and optimization suggestions—
668 on KernelBench Level-1 Task 95 (CrossEntropyLoss), as well as the corresponding speedup
669 across rounds (green = optimization, red = correction).

670 In this section, we present a case study on a single task to illustrate how the Judge diagnoses
671 issues and recommends optimizations. Figure 4 depicts the 10-round refinement process of
672 CudaForge on task 95_CrossEntropyLoss. We highlight four representative rounds—three
673 optimization rounds and one repair round—to demonstrate how the Judge leverages hardware feed-
674 back from NCU to provide targeted optimization or bug-fix suggestions.

675 In round 2, which is an optimization round, the Judge notices that 23.7% of active warps are stalled
676 due to barrier-type dependencies, which means roughly one quarter of potential issue opportunities
677 are blocked by synchronization. According to this, the Judge recommended replacing the original
678 shared-memory reduction that required multiple block-level synchronizations with a warp-level
679 shuffle reduction, giving below suggestion as prompt for coder: use warp-level shuffles in the max
680 and sum phases, then perform a single cross-warp combine, reducing `__syncthreads()` per block
681 from 16 to 2 (a reduction of 14). After applying this change, performance improved from $1.66\times$ to
682 $2.42\times$, with barrier stalls reduced and instruction-issue efficiency increased.

683 In round 5, it is a correction round. The previous round fails a numerical check with the following
684 error: “Outputs are not close, indicating a result mismatch”. The Judge diagnosed the root cause
685 as an uninitialized target_logit in thread 0 (“Thread-0 uses uninitialized target_logit”), which means
686 the variable target_logit is not updated to thread 0, leading wrong computing results. Accordingly,
687 the Judge gave the minimal fix suggestion, broadcast target_logit via `__shfl_sync` to thread
688 0. After applying the fix, the numerical issue disappeared.

690 In Rounds 6 & 7 (both optimization rounds), the Judge continues to track
691 `smsp_warp_issue_stalled_long_scoreboard_per_warp_active.pct`. In Round 6, this metric is
692 about 65%, primarily reflecting long-scoreboard stalls caused by global-memory latency. Per-thread
693 register usage is high, resulting in limited occupancy (only ~ 48 active warps/SM) and insufficient
694 latency hiding. The recommendation is to reduce per-thread registers to raise concurrency to ~ 64
695 warps/SM and thereby lower the long-scoreboard share. In Round 7, the metric rises to about 71%,
696 rooted in a second global read of logits after the max pass. The Judge therefore advises buffering
697 logits in per-warp shared memory during the max pass and reusing them in the expsum phase,
698 eliminating the redundant global memory access. Together, these strategies reduce global memory
699 access, significantly cut long-scoreboard stalls, improve issue efficiency and throughput; after these
700 two rounds, the speedup increases from $3.43\times$ to $3.76\times$.

701 This task demonstrates our CudaForge’s stability and expert-like workflow: first analyzing bottle-
necks from hardware feedback, then deriving the corresponding optimization strategy.

702 A.2 A BAD CASE
703704 We also explore some bad cases of our workflow. Take KernelBench Level 1 Task 5 as an example:
705

```

706 1 import torch
707 2 import torch.nn as nn
708 3
709 4 class Model(nn.Module):
710 5     """
711 6     Simple model that performs a matrix–scalar multiplication (C = A * s)
712 7     """
713 8     def __init__(self):
714 9         super(Model, self).__init__()
715 10
716 11     def forward(self, A: torch.Tensor, s: float) -> torch.Tensor:
717 12         """
718 13         Performs matrix–scalar multiplication.
719 14
720 15         Args:
721 16             A: Input matrix of shape (M, N)
722 17             s: Scalar value
723 18
724 19         Returns:
725 20             C: Resulting matrix of shape (M, N)
726 21         """
727 22     return A * s
728 23
729 24 M = 16384
730 25 N = 4096
731 26
732 27     def get_inputs():
733 28         A = torch.randn(M, N)
734 29         s = 3.14
735 30         return [A, s]
736 31
737 32     def get_init_inputs():
738 33         return [] # No special initialization inputs needed

```

738 In the first round, CudaForge generated a wrong kernel:
739

```

740 1 """python
741 2 # <complete ModelNew code>
742 3 import torch
743 4 import torch.nn as nn
744 5 from torch.utils.cpp_extension import load_inline
745 6
746 7 # -----
747 8 # CUDA source for fused ReLU + per-channel bias addition
748 9 # -----
749 10 cuda_src = r"""
750 11 #include <torch/extension.h>
751 12 #include <cuda.h>
752 13 #include <cuda_runtime.h>
753 14
754 15 template <typename scalar_t>
755 16     __global__ void relu_bias_add_kernel(const scalar_t* __restrict__ x,

```

```

756
757         const int spatial,
758         const int channels,
759         const int total_elems) {
760     int idx = blockIdx.x * blockDim.x + threadIdx.x;
761     if (idx >= total_elems) return;
762
763     int c = (idx / spatial) % channels; // channel index
764     scalar_t v = x[idx];
765     v = v > static_cast<scalar_t>(0) ? v : static_cast<scalar_t>(0);
766     v += bias[c];
767     y[idx] = v;
768 }
769
770 torch::Tensor relu_bias_add_cuda(torch::Tensor x, torch::Tensor bias) {
771     TORCH_CHECK(x.is_cuda(), "Input tensor x must reside on CUDA device");
772     TORCH_CHECK(bias.is_cuda(), "Bias tensor must reside on CUDA device");
773     TORCH_CHECK(x.scalar_type() == torch::kFloat32,
774                 "Only float32 data type is supported");
775
776     const int N = x.size(0);
777     const int C = x.size(1);
778     const int H = x.size(2);
779     const int W = x.size(3);
780     const int spatial = H * W;
781     const int total_elems = N * C * spatial;
782
783     auto y = torch::empty_like(x);
784
785     const int threads = 256;
786     const int blocks = (total_elems + threads - 1) / threads;
787
788     relu_bias_add_kernel<<<blocks, threads>>>(
789         x.data_ptr<float>(),
790         bias.view({C}).data_ptr<float>(),
791         y.data_ptr<float>(),
792         spatial,
793         C,
794         total_elems);
795
796     cudaError_t err = cudaGetLastError();
797     if (err != cudaSuccess)
798         throw std::runtime_error(cudaGetErrorString(err));
799
800     return y;
801 }
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
29
```

```

810
811    78
812    # -----
813    # Optimised model using the fused CUDA kernel
814    # -----
815    82 class ModelNew(nn.Module):
816        83     """
817        84     Same computation as original Model but with a fused
818        85     ReLU + bias-add CUDA kernel for higher performance.
819        86     """
820        87     def __init__(self, in_channels, out_channels, kernel_size, bias_shape):
821            88         super().__init__()
822            89         self.conv = nn.Conv2d(in_channels, out_channels, kernel_size)
823            90         # Bias after ReLU (per-channel, 11 spatial)
824            91         self.bias = nn.Parameter(torch.randn(bias_shape))
825            92         self.relu_bias_add = relu_bias_add
826
827        93     def forward(self, x):
828            94         x = self.conv(x)
829            95         # Fused kernel: y = relu(x) + bias
830            96         x = self.relu_bias_add.relu_bias_add_cuda(x, self.bias)
831            97         return x
832
833        99
834    100    # -----
835    101    # Helper functions (unchanged)
836    102    # -----
837    103    104    batch_size = 32
838    105    in_channels = 32
839    106    out_channels = 64
840    107    height = width = 64
841    108    kernel_size = 3
842    109    bias_shape = (out_channels, 1, 1)
843
844    110    def get_inputs():
845        111        return [torch.rand(batch_size, in_channels, height, width, device="cuda")]
846
847    112    def get_init_inputs():
848        113        return [in_channels, out_channels, kernel_size, bias_shape]
849        114
850        115    """
851
852
853    B DETAILS OF EVALUATION
854
855    Collecting Test Cases. To obtain reliable and representative test cases, we do not directly use
856    the single default input shape provided by KernelBench. Instead, for each task, we query the
857    GPT-4o API to generate ten diverse input shapes, ranging from moderately sized tensors (e.g.,
858    4096*4096) to shapes that utilize a substantial portion of the RTX 6000 GPU's memory capacity
859    (e.g., 16384*16384), as shown in the Table 6. This ensures that both correctness and performance
860    are evaluated across a broad spectrum of realistic workloads and prevents the evaluation from being
861    overly influenced by small-shape cases.
862
863    Correctness Evaluation. We then evaluate correctness through a two-stage procedure consisting of
864    compilation and execution. In the compilation stage, we verify that the generated kernel is syntacti-
865    cally valid and can be successfully compiled into executable CUDA code. In the execution stage, we

```

Task	Size in KernelBench	Perf	Max Size in Test	Perf	Change in Size
Level 1 Task 8	M = 8205 N = 2949 K = 5921	0.996 ×	M = 32820 N = 11796 K = 23684	0.993 ×	64 ×
Level 1 Task 15	M = 4096 N = 4096	3.063 ×	M = 16384 N = 16384	3.385 ×	16 ×
Level 2 Task 21	batch_size = 128 in_channels = 8 out_channels = 32 height = width = 256 kernel_size = 3 num_groups = 8	1.531 ×	batch_size = 128 in_channels = 8 out_channels = 32 height = width = 512 kernel_size = 3 num_groups = 8	1.449 ×	4 ×
Level 2 Task 75	batch_size = 1024 in_features = 8192 out_features = 8192 num_groups = 512	1.030 ×	batch_size = 1024 in_features = 19384 out_features = 19384 num_groups = 1024	1.017 ×	8 ×
Level 3 Task 18	batch_size = 64 input_channels = 3 height = 512 width = 512 num_classes = 1000	2.008 ×	batch_size = 64 input_channels = 3 height = 1024 width = 1024 num_classes = 1000	2.040 ×	4 ×

Table 6: Performance under KernelBench input sizes and maximum test sizes.

run the kernel on all ten input shapes and compare its outputs with those produced by the PyTorch reference implementation under the same inputs. A kernel is considered correct only if it successfully compiles and its numerical outputs are within a tolerance of 0.0001 for all test cases, which is a commonly adopted criterion (Ouyang et al., 2025; Lange et al., 2025; Baronio et al., 2025).

Performance Evaluation. Finally, we assess optimization performance using only the largest input shape in the generated test cases. The reason to select the largest is to aligns with the goal of CUDA kernel optimization, which is primarily motivated by large-scale workloads such as those found in LLM inference and training. For each task, we profile the candidate kernel on the largest shape. Then the Judge agent makes refinement decisions based on hardware feedback. After N refinement rounds, we select the most efficient correct kernel as the final result. When reporting speedup over the PyTorch baseline, we also use the largest input shape to ensure that GPU computation dominates runtime, instead of the PyTorch framework or OS overhead.

C ANALYSIS OF TIME COST IN CUDAFORGE

In this section, we provide an analysis of time cost in CudaForge. As shown in Table 7, for a typical KernelBench task with ten refinement rounds, the 25-minute end-to-end runtime is dominated by Nsight Compute profiling, which takes approximately 10–12 minutes in total. Kernel compilation accounts for 2–3 minutes, while LLM inference contributes about 9–11 minutes across all rounds. This breakdown shows that the overall runtime is determined primarily by profiling rather than model latency or compilation overhead. Since kernels are independent and can be processed concurrently, CudaForge scales effectively to larger codebases, with throughput largely governed by the degree of parallelism available for profiling rather than limitations of the framework itself.

Category	Time
Nsight Compute profiling	10–12 minutes
LLM inference	9–11 minutes
Kernel compilation	2–3 minutes

Table 7: Runtime breakdown of major components.

D PROMPT

D.1 SEED PROMPT FOR CODER(ONE-SHOT BASELINE PROMPT FROM KERNELBENCH)

We adopt the *One-shot Baseline Prompt* introduced in KERNELBENCH as our initial seed prompt for first round generation of all the baselines and our method. The full prompt is shown below.

```

1 You write custom CUDA kernels to replace the pytorch operators in the given architecture to
2 get speedups. You have complete freedom to choose the set of operators you want to replace.
3 You may make the decision to replace some operators with custom CUDA kernels and leave
4 others unchanged. You may replace multiple operators with custom implementations,
5 consider operator fusion opportunities (combining multiple operators into a single kernel, for
6 example, combining matmul+relu), or algorithmic changes (such as online softmax). You are
7 only limited by your imagination.
8
9 Here an example to show you the syntax of inline embedding custom CUDA operators in
10 torch:
11 The example given architecture is:
12 {few_base}
13
14 The example new arch with custom CUDA kernels looks like this:
15
16 {few_new}
17
18
19 You are given the following architecture:
20
21 ``python
22 {arch_src}
23
24 Optimize the architecture named Model with custom CUDA operators! Name your optimized
25 output architecture ModelNew. Output the new code in codeblocks. Please generate real
26 code, NOT pseudocode, make sure the code compiles and is fully functional. Just output
27 the new model code, no other text, and NO testing code!

```

D.2 PROMPT FOR JUDGE

In our prompt design for the Judge agent, we place the role specification and output schema in the system prompt. The input prompt only supplies per-round context(runtime information, NCU metrics, error_log). The system prompt is fixed; only the input prompt content changes each round.

The system prompt for cuda kernel optimization:

```

1 You are a senior CUDA performance engineer. Read the target GPU spec, the PyTorch
2 reference code, the current CUDA candidate, and the Nsight Compute metrics. Then identify
**exactly one** highest-impact speed bottleneck by 3–4 most important metrics, propose **
exactly one** optimization method and propose a modification plan. Be surgical and metrics-
driven.

```

```

972
973 3
974 4 Rules:
975 5 - Return **one and only one** optimization method the largest expected speedup.
976 6 - Prefer changes that directly address measured bottlenecks (occupancy limits,
977 7 memory coalescing, smem bank conflicts, register pressure, long/short scoreboard
978 8 stalls, tensor–core underutilisation, etc.).
979 9 - Keep fields brief; avoid lists of alternatives, disclaimers, or generic advice.
980
981 10
982 11 Output format (JSON):
983 12 {"json
984 13 {
985 14 "bottleneck": "<max 30 words>",
986 15 "optimization method": "<max 35 words>",
987 16 "modification plan": "<max 35 words>"}
988 17 }
989 18 "","",""
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

The input prompt for optimization:

```

990 1 # Target GPU
991 2 GPU Name: {gpu_name}
992 3 Architecture: {gpu_arch}
993 4 Details:
994 5 {gpu_items}
995 6
996 7
997 8 # Pytorch Reference
998 9 {python_code}
999 10
1000 11
1001 12 # CUDA candidate
1002 13 {"python
1003 14 {CUDA_CODE}
1004 15 "}
1005 16
1006 17 # Nsight Compute metrics (verbatim)
1007 18 {NCU_METRICS}
1008 19
1009 20 Read everything and follow the Rules exactly. Return the JSON in the specified format.
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

The system prompt for kernel correction:

```

1012 1 You are a senior CUDA + PyTorch correctness auditor. Your job is to read a PyTorch
1013 2 reference and a CUDA candidate and report exactly one most critical correctness issue in the
1014 3 CUDA code that would cause a behavioral mismatch vs. the PyTorch reference. Be terse and
1015 4 precise.
1016 5
1017 6 Rules:
1018 7
1019 8 5 Return one and only one issue the single highest-impact problem.
1020 9 6 Prefer semantic/correctness issues over micro-optimizations or style.
1021 10 7
1022 11 8 If multiple issues exist, pick the one that most changes outputs or gradients.
1023 12 9
1024 13 10 If nothing clearly wrong is found, say it explicitly.
1025 11 12
1026 13 13 Keep each field brief; avoid extra commentary, lists, or alternatives.
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1599

```

```

1026
1027 14
1028 15 Output format (JSON):
1029 16  ``"json
1030 17  {
1031 18  "critical_issue": "<max 20 words>",
1032 19  "why_it_matters": "<max 35 words>",
1033 20  "minimal_fix_hint": "<max 20 words>"
1034 21  }
1035 22  ``"

```

1036 The input prompt for kernel repair:

```

1037
1038 1 You are given:
1039 2
1040 3 ERROR_LOG:
1041 4 {ERROR_LOG}
1042 5
1043 6 PyTorch reference (ground truth):
1044 7
1045 8 {PYTORCH_CODE}
1046 9
1047 10 CUDA candidate (to audit):
1048 11
1049 12 {CUDA_CODE}
1050 13
1051 14
1052 15 Follow the Rules and produce the JSON exactly in the specified format.

```

D.3 PROMPT FOR CODER

1055 For the Coder, we use the default system prompt and put all task details in the input prompt. This
1056 keeps the agent simple and fully context-driven. .

1057 The prompt for kernel optimization:

```

1058
1059 1 # Target GPU
1060 2 GPU Name: {gpu_name}
1061 3 Architecture: {gpu_arch}
1062 4 Details:
1063 5 {gpu_items}
1064 6
1065 7 You are a CUDA–kernel optimization specialist.
1066 8
1067 9 Analyze the provided architecture and **strictly apply the following STRATEGY** to
1068 produce an improved CUDA kernel.
1069
1070 11 ``"python
1071 12 {CUDA_CODE}
1072 13 ``
1073
1074 15 [optimization instructions]
1075 16 {optimization_suggestion}
1076 17
1077 18 GOAL
1078 19
1079 20 - Improve latency and throughput on the target GPU.
21 - Maintain correctness within atol=1e-4 or rtol=1e-4.
22 - Preserve the public Python API (same inputs/outputs, shapes, dtypes).

```

```

1080
1081
1082 23
1083 24
1082 25 OUTPUT RULES (STRICT)
1083 26 1. Inside the block, follow **exactly** this order:
1084 27 1. Imports 'torch', 'torch.nn', 'load.inline'.
1085 28 2. 'source' triplequoted CUDA string(s) (kernel + host wrapper).
1086 29 3. 'cpp_src' prototypes for *all* kernels you expose.
1087 30 4. **One** 'load.inline' call per kernel group.
1088 31 5. 'class ModelNew(nn.Module)' mirrors original inputs/outputs but calls
1089 32     your CUDA kernels.
1090 33 2. **Do NOT include** testing code, 'if __name__ == "__main__"', or extra prose.
1091
1092 34
1092 35     """python
1093 36 # <your corrected code>
1093 37     """

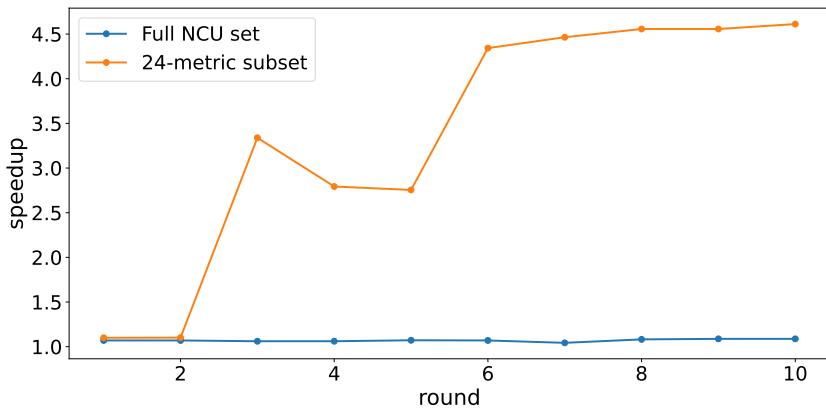
```

1094
1095
1096
1097
1098
1099 The prompt for kernel correction:

```

1100
1101 1 You are a senior CUDA-extension developer.
1102 2 Your job is to **FIX** the compilation or runtime errors in the Python script
1103 3 shown below.
1104
1105 4
1105 5 OUTPUT RULES (STRICT)
1106 6 1. Inside the block, follow **exactly** this order:
1107 7 1. Imports 'torch', 'torch.nn', 'load.inline'.
1108 8 2. 'source' triplequoted CUDA string(s) (kernel + host wrapper).
1109 9 3. 'cpp_src' prototypes for *all* kernels you expose.
1110 10 4. **One** 'load.inline' call per kernel group.
1111 11 5. 'class ModelNew(nn.Module)' mirrors original inputs/outputs but calls
1112 12     your CUDA kernels.
1113 13 2. **Do NOT include** testing code, 'if __name__ == "__main__"', or extra prose.
1114
1115 14
1115 15
1116 16 ERROR LOG
1117
1118 17 {ERROR_LOG}
1119
1120 18
1121 19
1122 20
1123 21 OLD CODE (read-only)
1124
1125 22 {CUDA_CODE}
1126
1127 23
1128 24
1129 25 Main Critical Problem
1130
1131 26 {Problem}
1132
1133 27
1133 28     """python
1134 29 # <your corrected code>
1135 30     """

```

1134 E DETAIL FOR THE NCU METRICS
1135
11361137 E.1 WHY CHOOSE NCU SUBSET METRICS?
11381139 We find that exposing large models to the full NCU metric set overwhelms them, reducing the
1140 accuracy and stability of their optimization suggestions and degrading Judge output quality. We
1141 illustrate this with following specific case study.
11421156 Figure 5: Generated kernel performance comparison between Full-metrics input(blue line) vs. Met-
1157 rics subset input(orange line) for Judge agent on Level 2 task 51.
11581159 As shown in Figure 5, we compare two configurations on a representative task. (i) The Judge agent
1160 receives all raw NCU metrics in each round and provides feedback to the Coder agent. (ii) The
1161 Judge agent receives only the 24 key metrics selected by our method and captures only 3-4 critical
1162 metrics in each round to generate optimization suggestion.
11631164 Within 10 iterative rounds, the Judge agent using full NCU metrics fails to guide the Coder agent
1165 toward better kernels. The performance of the generated kernels remains almost unchanged, despite
1166 repeated interactions. In contrast, when the Judge agent is restricted to use subset metrics, the per-
1167 formance of generated kernels varies across rounds and is often improved. This process eventually
1168 leads to a substantial speedup.
11691170 In addition, we carefully reviewed how the Judge identified bottlenecks and proposed optimizations
1171 in both configurations, full set vs. 24-subset. We select one Judge’s response from the optimization
1172 round to better illustrate the difference in judge output, with full NCU and with a subset of NCU as
1173 judge input.
11741175 Response of using the full NCU metrics set:
1176

- 1 {"bottleneck": "Long-scoreboard stalls from re-loading GEMM output out of DRAM (gemm_out round-trip dominates cycles).",
- 2 "optimization method": "Fuse the Linear GEMM and the current epilogue into one FP16-tensor-core CUTLASS kernel that performs subtract, row-mean, GELU and residual add in-register, avoiding global-memory re-reads.",
- 3 "modification plan": "Write CUTLASS custom epilogue functor; use cublasLt or cutlass::gemm_universal to compute GEMM+A, keep accumulators in registers, warp-reduce to mean, apply GELU, add x_orig, directly store final output; expose via new PyTorch extension."}

1184 Response of using the 24 key NCU subset:
1185

- 1 {"bottleneck": "Kernel is DRAM-bound (102.9 % peak); 57 % long-scoreboard stalls from second pass reading original_x.",

1188 2 "optimization method": "Fuse the second loop: cache each threads original_x elements in
 1189 registers during diff accumulation, then add gelu_val and store, eliminating the extra BF
 1190 global read.",
 1191 3 "modification plan": "Replace two for-loops with one: load linear_out, subtract, and
 1192 original_x together, keep original_x[16] in per-thread array, finish reduction, add gelu_val,
 1193 write to output; ~4 MB traffic per batch, >30 % speedup."}

1194
 1195 Based on these two responses, we find that judge with full set NCU metrics tends to misidentify the
 1196 true bottleneck. The judge with full set NCU metrics attributes the bottleneck to re-loading `gemm_out`
 1197 and recommends a monolithic CUTLASS epilogue that performs row-mean/GELU/residual in reg-
 1198 isters. This diagnosis is misaligned with our kernel's access pattern and is hard to realize for general
 1199 shapes due to cross-tile aggregation. In contrast, the judge with 24-key subset correctly identifies
 1200 a DRAM-bound kernel dominated by the second pass over `x_orig`, and proposes a one-pass rewrite
 1201 that caches `x_orig` in registers during the first traversal and writes back after GELU, eliminating
 1202 an entire $B \times F$ global memory read. This change is lightweight, architecture-agnostic, and yields
 1203 consistent speedups (e.g., about 4 MB less traffic per batch, more than 30% in our setting).

1204 E.2 TOP-20 NCU METRICS EXAMPLE

1205 This section reports, for several example tasks, the Top-20 Nsight Compute (NCU) metrics most
 1206 correlated with runtime, ranked by the absolute value of the Pearson correlation coefficient. Here,
 1207 runtime refers to the kernel's execution time. When the correlation coefficient is positive, larger met-
 1208 ric values typically imply longer execution time; when it is negative, larger metric values typically
 1209 imply shorter execution time. All metric names follow their original name in NCU.

1210 Table 8: Task-Conv2D: Pearson correlation with runtime (Top-20).

Metric Name	Correlation	Abs Correlation
<code>sm_cycles_active.avg</code>	1.000 000	1.000 000
<code>gpc_cycles_elapsed.max</code>	1.000 000	1.000 000
<code>launch_occupancy_limit_shared_mem</code>	0.945 507	0.945 507
<code>dram_bytes.sum.per_second</code>	-0.924 251	0.924 251
<code>gpu_dram_throughput.avg.pct_of_peak_sustained_elapsed</code>	-0.924 155	0.924 155
<code>smsp_inst_executed.avg</code>	0.916 287	0.916 287
<code>smsp_inst_executed.sum</code>	0.916 287	0.916 287
<code>smsp_inst_issued.avg</code>	0.916 262	0.916 262
<code>smsp_inst_issued.sum</code>	0.916 262	0.916 262
<code>lts_t_sector_hit_rate.pct</code>	0.839 237	0.839 237
<code>smsp_sass_average_branch_targets_threads_uniform.pct</code>	0.810 334	0.810 334
<code>lts_throughput.avg.pct_of_peak_sustained_elapsed</code>	-0.787 261	0.787 261
<code>smsp_inst_executed_op_branch.sum</code>	0.746 483	0.746 483
<code>launch_grid.size</code>	0.745 917	0.745 917
<code>l1tex_t_sector_hit_rate.pct</code>	0.728 356	0.728 356
<code>gpc_cycles_elapsed.avg.per_second</code>	0.728 053	0.728 053
<code>dram_cycles_elapsed.avg.per_second</code>	0.665 784	0.665 784
<code>launch_waves_per_multiprocessor</code>	0.627 478	0.627 478
<code>launch_thread_count</code>	0.627 478	0.627 478
<code>launch_shared_mem_per_block_static</code>	-0.610 501	0.610 501

1229 Table 9: Task-SpMM: Pearson correlation with runtime (Top-20).

Metric Name	Correlation	Abs Correlation
<code>gpc_cycles_elapsed.max</code>	0.999 993	0.999 993
<code>sm_cycles_active.avg</code>	0.998 432	0.998 432
<code>gpu_compute_memory_request_throughput.avg.pct....</code>	-0.967 284	0.967 284
<code>gpu_compute_memory_throughput.avg.pct.of_peak....</code>	-0.964 455	0.964 455
<code>lts_t_sector_hit_rate.pct</code>	0.951 201	0.951 201
<code>dram_bytes.sum.per_second</code>	-0.926 134	0.926 134
<code>gpu_dram_throughput.avg.pct.of_peak_sustained....</code>	-0.925 856	0.925 856
<code>l1tex_throughput.avg.pct.of_peak_sustained_active</code>	0.871 262	0.871 262
<code>sm_inst_executed.avg.per_cycle_elapsed</code>	-0.837 675	0.837 675
<code>smsp_issue_inst0.avg.pct.of_peak_sustained_active</code>	0.837 284	0.837 284
<code>smsp_issue_active.avg.pct.of_peak_sustained....</code>	-0.837 284	0.837 284

1241 Continued on next page

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

E.3 KEY SUBSET OF 24 NCU METRICS

1253

The table below lists the exact 24 metrics in our task-agnostic key subset.

1254

1255

Table 10: The 24-metric key subset.

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

F CUDA-L1

1277

In our replication efforts, we found that the authors of CUDA-L1 released only the final, generated kernels for each task. After carefully studying these cases, we identified several interesting findings.

1278

1279

First, We found that CUDA-L1 tends to emphasize PyTorch-level optimizations rather than generating and refining custom CUDA kernels. This pattern also emerged as the most frequent issue in their provided case. Although CUDA-L1 reports the top-10 cases with the largest speedups, our review shows that nine of these ten final solutions do not use custom CUDA kernels; instead, they rely heavily on official PyTorch implementations.

1280

1281

1282

1283

1284

1285

1286

1287

This is the top-ranked entry in their *KernelBench Tasks Ranked by RL-CUDA1 Acceleration (Top-10)*: Level-2 Task 83, with a reported $120.3 \times$ speedup

1288

1289

1290

1291

1292

1293

1294

1295

```

1 import torch
2 import torch.nn as nn
3
4 class ModelNew(nn.Module):
5     """
6     Optimized implementation of a model that performs a 3D convolution,
7     applies Group Normalization, minimum, clamp, and dropout.
8
9     Args:

```

```

1296     10     in_channels (int): Number of input channels
1297     11     out_channels (int): Number of output channels
1298     12     kernel_size (int or tuple): Size of the convolving kernel
1299     13     groups (int): Number of groups for GroupNorm
1300     14     min_value (float): Minimum value for clamp operation
1301     15     max_value (float): Maximum value for clamp operation
1302     16     dropout_p (float): Dropout probability
1303     17     """
1304     18     def __init__(self, in_channels, out_channels, kernel_size, groups, min_value, max_value,
1305      dropout_p):
1306         super(ModelNew, self).__init__()
1307         # Store the original layers for parameter compatibility
1308         self.conv = nn.Conv3d(in_channels, out_channels, kernel_size)
1309         self.norm = nn.GroupNorm(groups, out_channels)
1310         self.dropout = nn.Dropout(dropout_p)
1311         self.min_value = min_value
1312         self.max_value = max_value
1313         self.dropout_p = dropout_p
1314
1315         # Pre-compute convolution parameters for output shape calculation
1316         if isinstance(kernel_size, int):
1317             self.kernel_size = (kernel_size, kernel_size, kernel_size)
1318         else:
1319             self.kernel_size = kernel_size
1320
1321         self.stride = self.conv.stride
1322         self.padding = self.conv.padding
1323         self.dilation = self.conv.dilation
1324
1325         def forward(self, x):
1326             x = self.conv(x)
1327             x = self.norm(x)
1328             x = torch.minimum(x, torch.tensor(self.min_value, device=x.device))
1329             x = torch.clamp(x, min=self.min_value, max=self.max_value)
1330             x = self.dropout(x)
1331             return x
1332
1333         # Keep ALL hyperparameters EXACTLY as shown in the reference implementation
1334         batch_size = 128
1335         in_channels = 3
1336         out_channels = 16
1337         depth, height, width = 16, 32, 32
1338         kernel_size = 3
1339         groups = 8
1340         min_value = 0.0
1341         max_value = 1.0
1342         dropout_p = 0.2
1343
1344         def get_inputs():
1345             return [torch.randn(batch_size, in_channels, depth, height, width)]
1346
1347         def get_init_inputs():
1348             return [in_channels, out_channels, kernel_size, groups, min_value, max_value, dropout_p]
1349
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
2653
2654
2655
2656
2657
2658
2659
2660
2661
2662
2663
2664
2665
2666
2667
2668
2669
2670
2671
2672
2673
2674
2675
2676
2677
2678
2679
2680
2681
2682
2683
2684
2685
2686
2687
2688
2689
2690
2691
2692
2693
2694
2695
2696
2697
2698
2699
2700
2701
2702
2703
2704
2705
2706
2707
2708
2709
2710
2711
2712
2713
2714
2715
2716
2717
2718
2719
2720
2721
2722
2723
2724
2725
2726
2727
2728
2729
2730
2731
2732
2733
2734
2735
2736
2737
2738
2739
2740
2741
2742
2743
2744
2745
2746
2747
2748
2749
2750
2751
2752
2753
2754
2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788
2789
2790
2791
2792
2793
2794
2795
2796
2797
2798
2799
2800
2801
2802
2803
2804
2805
2806
2807
2808
2809
2810
2811
2812
2813
2814
2815
2816
2817
2818
2819
2820
2821
2822
2823
2824
2825
2826
2827
2828
2829
2830
2831
2832
2833
2834
2835
2836
2837
2838
2839
2840
2841
2842
2843
2844
2845
2846
2847
2848
2849
2850
2851
2852
2853
2854
2855
2856
2857
2858
2859
2860
2861
2862
2863
2864
2865
2866
2867
2868
2869
2870
2871
2872
2873
2874
2875
2876
2877
2878
2879
2880
2881
2882
2883
2884
2885
2886
2887
2888
2889
2890
2891
2892
2893
2894
2895
2896
2897
2898
2899
2900
2901
2902
2903
2904
2905
2906
2907
2908
2909
2910
2911
2912
2913
2914
2915
2916
2917
2918
2919
2920
2921
2922
2923
2924
2925
2926
2927
2928
2929
2930
2931
2932
2933
2934
2935
2936
2937
2938
2939
2940
2941
2942
2943
2944
2945
2946
2947
2948
2949
2950
2951
2952
2953
2954
2955
2956
2957
2958
2959
2960
2961
2962
2963
2964
2965
2966
2967
2968
2969
2970
2971
2972
2973
2974
2975
2976
2977
2978
2979
2980
2981
2982
2983
2984
2985
2986
2987
2988
2989
2990
2991
2992
2993
2994
2995
2996
2997
2998
2999
2999

```

1499

diag_mm_compare.py

||

```

1350
2  import time
3  import math
4  import torch
5  import torch.nn as nn
6  import torch.nn.functional as F
7
8  # -----
9  # Reference implementation
10 # -----
11 class Model(nn.Module):
12     """
13     Simple model that performs a matrix multiplication of a diagonal matrix with another
14     matrix.
15     C = diag(A) * B
16     """
17     def __init__(self):
18         super(Model, self).__init__()
19
20     def forward(self, A, B):
21         """
22         Args:
23             A (torch.Tensor): 1D tensor, diagonal entries. Shape: (N,)
24             B (torch.Tensor): 2D tensor. Shape: (N, M)
25         Returns:
26             torch.Tensor: (N, M)
27         """
28         return torch.diag(A) @ B
29
30     # -----
31     # Optimized implementation
32     # -----
33 class ModelNew(nn.Module):
34     """
35     Optimized model that performs a matrix multiplication of a diagonal matrix with another
36     matrix.
37     C = diag(A) * B
38     """
39     def __init__(self):
40         super(ModelNew, self).__init__()
41
42     def forward(self, A, B):
43         """
44         Args:
45             A (torch.Tensor): 1D tensor, diagonal entries. Shape: (N,)
46             B (torch.Tensor): 2D tensor. Shape: (N, M)
47         Returns:
48             torch.Tensor: (N, M)
49         """
50         # Equivalent to torch.diag(A) @ B, but avoids forming the full diagonal matrix
51         return B * A.unsqueeze(1)
52
53     # -----
54     # Hyperparameters & inputs
55     # -----
56     M = 4096
57     N = 4096
58

```

```
1404 59 def get_inputs(device=None, dtype=torch.float32):  
1405 60     A = torch.randn(N, device=device, dtype=dtype)  
1406 61     B = torch.randn(N, M, device=device, dtype=dtype)  
1407 62     return [A, B]  
1408 63  
1409 64 def get_init_inputs():  
1410 65     return [] # No special initialization inputs needed
```

1414 In addition, we observed many reported speedups that are effectively equal to one (clustered around
1415 1.00, typically within $\pm 5\%$). A closer inspection shows that, in these cases, the system falls back
1416 to the original PyTorch operator when the custom kernel fails to compile, which naturally yields no
1417 measurable speedup.

1418 For example, below is the forward method from the final solution for KernelBench Level-1 Task
1419 3 generated by CUDA-L1. This code get from the CUDA-L1's official Github. We observe that
1420 the method first attempts to call a *custom CUDA kernel*; however, upon any compilation failure
1421 or exception, it immediately falls back to `torch.bmm(A, B)`. Crucially, `torch.bmm(A, B)`
1422 is exactly the operator that this task asks to be replaced by a custom kernel, meaning the fallback
1423 undermines the task's objective. This explains why the reported speedup is only 1.006x.

```

1  def forward(self, A: torch.Tensor, B: torch.Tensor) -> torch.Tensor:
2      """
3          Performs batched matrix multiplication.
4
5          Args:
6              A: Input tensor of shape (batch_size, m, k).
7              B: Input tensor of shape (batch_size, k, n).
8
9          Returns:
10             C: Output tensor of shape (batch_size, m, n).
11             """
12
13         # Fall back to torch.bmm if CUDA module failed to load
14         if ModelNew._cuda_module is None:
15             return torch.bmm(A, B)
16
17         # Check if inputs are on CUDA
18         if not A.is_cuda or not B.is_cuda:
19             A = A.cuda() if not A.is_cuda else A
20             B = B.cuda() if not B.is_cuda else B
21
22         # Ensure inputs are contiguous and float32
23         A = A.contiguous().float()
24         B = B.contiguous().float()
25
26         # Use custom CUDA kernel
27         try:
28             result = ModelNew._cuda_module.batched_matmul(A, B)
29             if not A.is_cuda:
30                 result = result.cpu()
31             return result
32
33         except Exception as e:
34             print(f"Error in custom kernel: {e}, falling back to torch.bmm")
35             return torch.bmm(A, B)

```

1458 **G DETAILS OF BENCHMARK**
14591460 **G.1 KERNELBENCH**
14611462 **KernelBench** is a standardized benchmark designed to evaluate the capability of large language
1463 models (LLMs) in CUDA kernel generation and optimization. It consists of 270 tasks across four
1464 levels of increasing difficulty, of which Levels 1–3 (250 tasks in total) are commonly adopted for
1465 evaluation. Each task provides a PyTorch reference implementation f_{T_i} together with fixed in-
1466 put–output specifications, enabling automated correctness and performance validation.1467

- **Level 1 (Basic Operators):** Contains simple, low-level operators such as matrix multipli-
1468 cation, element-wise operations, and reductions. These tasks primarily test the ability to
1469 generate functionally correct CUDA kernels.
- **Level 2 (Composite Operations):** Involves multi-step operator combinations, requiring
1470 the model to compose multiple CUDA primitives and manage intermediate memory effi-
1471 ciently. These tasks test the capacity for more complex code synthesis.
- **Level 3 (End-to-End Models):** Includes challenging kernels derived from full neural net-
1472 work architectures such as AlexNet, VGG, and ResNet components. These tasks assess the
1473 ability to produce efficient, large-scale kernels under realistic deep learning workloads.
- **Level 4 (Optional):** The full benchmark also defines an advanced level with additional
1474 research-oriented tasks, but this is less frequently adopted due to its complexity and lack of
1475 standardized evaluation setups.

1476 KernelBench has become a widely used benchmark in recent work on LLM-based code generation
1477 (Team, 2025; Baronio et al., 2025; Lange et al., 2025), as it provides a controlled and reproducible
1478 environment to measure both *correctness* (functional equivalence to PyTorch) and *efficiency* (execu-
1479 tion speed relative to PyTorch). In our study, we adopt all Level 1–3 tasks, following prior work, to
1480 ensure fair comparison across baselines.1481 **G.2 OUR STRATIFIED RANDOM SUBSET \mathcal{D}^***
14821483 While our main evaluation is conducted on the full KernelBench Level 1–3 benchmark (250 tasks in
1484 total), we additionally construct a stratified subset \mathcal{D}^* to enable detailed analysis and fair comparison
1485 with prior work such as Kevin.1486 The construction of \mathcal{D}^* follows two principles: (1) **Coverage across difficulty levels.** Since Kernel-
1487 Bench is stratified by increasing task complexity (Level 1: single-operator tasks, Level 2: multi-step
1488 fused operators, Level 3: full network components), we ensure that the sampled subset preserves the
1489 relative distribution of difficulty. (2) **Diversity of task types.** Within each level, we sample tasks uni-
1490 formly across different operator categories (e.g., elementwise ops, reductions, convolutions, fused
1491 blocks) so that the subset remains representative of the overall benchmark.1492 Concretely, we perform stratified random sampling with a fixed 10% ratio for each level, resulting
1493 in a subset of 10 tasks from Level 1, 10 tasks from Level 2, and 5 tasks from Level 3, for a total of
1494 25 tasks. For reproducibility, the exact task IDs included in \mathcal{D}^* are:1495

- **Level 1 (10 tasks):** 13, 10, 16, 29, 35, 72, 7, 89, 93, 34
- **Level 2 (10 tasks):** 17, 19, 40, 3, 13, 21, 38, 28, 26, 34
- **Level 3 (5 tasks):** 5, 18, 32, 41, 21

1505 **USAGE OF LLM**
15061507 During the preparation of this paper, we employed large language models (LLMs) solely for **textual**
1508 **assistance**, including grammar correction, stylistic refinement, and clarity improvements. All core
1509 research contributions—including the design of CudaForge, implementation of experiments, and
1510 analysis of results—were conducted entirely by the authors. The LLM was not used to generate
1511 research ideas, experimental results, or any substantive content of the paper.