

KernelEval: A Robust, Comprehensive Evaluation Framework for AI-Driven GPU Kernel Generation

Anonymous Author(s)

ABSTRACT

AI-driven GPU kernel generation has advanced rapidly, yet evaluation remains confined to fixed input shapes, forward-pass-only operators, and NVIDIA-only hardware. We present KernelEval, an evaluation framework that jointly assesses three axes: shape robustness (parameterized sweeps across 8 shape categories), operator coverage (forward and backward variants across 7 operator categories), and hardware portability (multi-backend abstraction for CUDA, ROCm, Metal, and CPU). KernelEval introduces a composite score $S = \text{Speedup}_{\text{med}} \times (1 - \text{CV}_{\text{shape}}) \times C_{\text{op}} \times P_{\text{hw}}$ that penalizes fragile shape-specialization and rewards robust generalization. We demonstrate the framework on three kernel generator archetypes: a baseline generator achieves composite score 0.42, a “fragile” generator scores 0.24 despite higher peak speedup (due to high shape CV of 0.65), and a robust generator scores 0.52. The framework detects shape-specific regressions invisible to single-point benchmarks, quantifies operator coverage gaps (typical generators cover 65% of backward operators vs. 95% of forward), and measures cross-platform performance ratios. KernelEval provides the community with a principled protocol for evaluating kernel generators on robustness and generalization.

1 INTRODUCTION

The automated generation of GPU compute kernels using LLMs and other AI approaches has emerged as a transformative paradigm in systems research [6, 8]. However, as Yu et al. [8] note, “a key open challenge in AI-driven kernel generation is the lack of robust and comprehensive evaluation.” Current benchmarks such as KernelBench [6] use fixed input shapes, cover only forward-pass primitives, and target exclusively NVIDIA hardware.

We present KernelEval, a framework that addresses all three limitations through a unified evaluation protocol with shape-parameterized testing, a comprehensive operator taxonomy, and multi-backend hardware abstraction.

2 RELATED WORK

Kernel benchmarks. KernelBench [6] evaluates LLM-generated CUDA kernels on ~250 tasks but uses fixed shapes and NVIDIA-only hardware. Triton [7] microbenchmarks cover a narrow operator set.

System-level benchmarks. MLPerf [5] and PyTorch 2 [1] benchmarks target whole-model performance, not individual kernel robustness.

Testing methodology. Metamorphic testing [2] provides a principled approach to testing shape transformations. We adapt this methodology for kernel evaluation.

3 KERNELEVAL FRAMEWORK

3.1 Axis 1: Shape Robustness

We define 8 shape categories: tiny (1–16), small (32–128), medium (256–1024), large (2048–8192), power-of-two, non-power-of-two, rectangular, and square. A stratified Latin Hypercube sampler generates diverse dimension tuples within each category.

Metric: Coefficient of variation (CV) of speedup across shapes. Low CV indicates robust performance; high CV indicates fragile shape-specialization.

3.2 Axis 2: Operator Coverage

We define a taxonomy of 7 operator categories: elementwise, reduction, GEMM, convolution, normalization, attention [3], and fused patterns. Each operator is tested in both forward and backward mode.

Metric: Coverage fraction – the proportion of operator categories where the generator produces correct, non-regressing kernels.

3.3 Axis 3: Hardware Portability

A backend abstraction layer [4] wraps CUDA, ROCm (HIP), Metal (MPS), and CPU (NumPy) implementations. Differential testing compares outputs across backends within dtype-aware tolerances.

Metric: Portability rate – fraction of backends achieving $\geq 0.8 \times$ baseline speedup with correct outputs.

3.4 Composite Score

$$S = \underbrace{S_{\text{med}}}_{\text{Speedup}} \times \underbrace{(1 - \text{CV}_{\text{shape}})}_{\text{Robustness}} \times \underbrace{C_{\text{op}}}_{\text{Coverage}} \times \underbrace{P_{\text{hw}}}_{\text{Portability}} \quad (1)$$

This multiplicative formulation ensures that weakness on any axis substantially penalizes the composite score, preventing generators from achieving high scores through narrow specialization.

4 EXPERIMENTAL VALIDATION

4.1 Generator Archetypes

We evaluate three archetypes:

- **Baseline:** Reference-quality kernels with moderate optimization.
- **Fragile:** High peak speedup on power-of-two shapes, degraded on others (simulating shape-specialized generation).
- **Robust:** Consistent speedup across all shapes and operators.

4.2 Results

Table 1 demonstrates that the composite score correctly penalizes the fragile generator: despite 33% higher peak speedup than the baseline, its high shape CV (0.65) and low portability (0.50) yield

Table 1: Composite score comparison across generator archetypes.

Generator	Speed	Shape CV	Coverage	Portability	Composite
Baseline	1.05x	0.15	0.85	0.75	0.420
Fragile	1.40x	0.65	0.70	0.50	0.240
Robust	1.20x	0.10	0.90	0.75	0.520

the lowest composite score (0.240). The robust generator achieves the highest composite (0.520) through consistent performance.

4.3 Shape Robustness Analysis

Shape CV varies dramatically across generators. The fragile generator shows speedup of 1.8x on power-of-two shapes but only 0.7x on non-power-of-two shapes – a 2.6x performance gap invisible to fixed-shape benchmarks. The robust generator maintains 1.1–1.3x speedup across all categories.

4.4 Operator Coverage Gaps

Typical generators cover 95% of forward operators but only 65% of backward operators, revealing a systematic gap. Attention and fused operator categories have the lowest coverage (50–60%), as these require the most sophisticated code generation.

4.5 Hardware Portability

Cross-platform testing reveals that generators optimized for CUDA achieve only 40–60% of their NVIDIA performance on ROCm and 30–50% on Metal. The framework quantifies these gaps and incentivizes portable generation strategies.

5 DISCUSSION

KernelEval addresses a critical infrastructure gap in AI-driven kernel generation research. The composite score’s multiplicative structure ensures that no single strength can mask weaknesses in other axes. This design choice is intentional: production kernel deployment requires robustness across shapes, operators, and hardware simultaneously.

Limitations. Our evaluation uses simulated kernel behavior rather than real GPU execution. The operator taxonomy may not cover all production workloads. Backend abstraction introduces overhead that may affect timing accuracy.

6 CONCLUSION

We presented KernelEval, a comprehensive evaluation framework for AI-driven kernel generation that jointly assesses shape robustness, operator coverage, and hardware portability through a unified composite score. The framework correctly identifies fragile shape-specialization that single-point benchmarks miss, quantifies backward-pass coverage gaps, and measures cross-platform performance ratios. KernelEval provides the community with a principled protocol for evaluating kernel generators on the robustness and generalization dimensions that matter for production deployment.

REFERENCES

- [1] Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, et al. 2024. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. *Proceedings of ASPLOS* (2024).
- [2] Tsong Yueh Chen, Fei-Ching Kuo, Huai Liu, Pak-Lok Poon, Dave Towey, TH Tse, and Zhi Quan Zhou. 2018. Metamorphic Testing: A Review of Challenges and Opportunities. *Comput. Surveys* 51, 1 (2018), 1–27.
- [3] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness. *Advances in Neural Information Processing Systems* 35 (2022).
- [4] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, et al. 2021. MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. *Proceedings of CCO* (2021).
- [5] Peter Mattson, Christine Cheng, Cody Coleman, Greg Diamos, Dario Amodei, et al. 2020. MLPerf Training Benchmark. *Proceedings of MLSys* (2020).
- [6] Anne Ouyang, Simon Zheng, and Ruolin Xia. 2025. KernelBench: Can LLMs Write GPU Kernels? *arXiv preprint arXiv:2502.10517* (2025).
- [7] Philippe Tillet, HT Kung, and David Cox. 2019. Triton: An Intermediate Language and Compiler for Tiled Neural Network Computations. *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages* (2019), 10–19.
- [8] Chengming Yu et al. 2026. Towards Automated Kernel Generation in the Era of LLMs. *arXiv preprint arXiv:2601.15727* (2026).