

DiffBench Meets DiffAgent: End-to-End LLM-Driven Diffusion Acceleration Code Generation

Jiajun jiao^{1, 2*}, Haowei Zhu^{1, 3*}, Puyuan Yang¹, Jianghui Wang¹, Ji Liu¹, Ziqiong Liu¹, Dong Li¹, Yuejian Fang^{2†}, Junhai Yong³, Bin Wang^{3†}, Emad Barsoum¹

¹ Advanced Micro Devices, Inc.

² Peking University

³ Tsinghua University

{jiajjiao, haoweiz, puyuyang, jianghui.wang, liuji, ziqioliu, d.li, ebarsoum}@amd.com, fangyj@ss.pku.edu.cn, yongjh@tsinghua.edu.cn, wangbins@tsinghua.edu.cn

Abstract

Diffusion models have achieved remarkable success in image and video generation. However, their inherently multiple step inference process imposes substantial computational overhead, hindering real-world deployment. Accelerating diffusion models is therefore essential, yet determining how to combine multiple model acceleration techniques remains a significant challenge. To address this issue, we introduce a framework driven by large language models (LLMs) for automated acceleration code generation and evaluation. First, we present DiffBench, a comprehensive benchmark that implements a three stage automated evaluation pipeline across diverse diffusion architectures, optimization combinations and deployment scenarios. Second, we propose DiffAgent, an agent that generates optimal acceleration strategies and codes for arbitrary diffusion models. DiffAgent employs a closed-loop workflow in which a planning component and a debugging component iteratively refine the output of a code generation component, while a genetic algorithm extracts performance feedback from the execution environment to guide subsequent code refinements. We provide a detailed explanation of the DiffBench construction and the design principles underlying DiffAgent. Extensive experiments show that DiffBench offers a thorough evaluation of generated codes and that DiffAgent significantly outperforms existing LLMs in producing effective diffusion acceleration strategies.

Introduction

Diffusion models (Podell et al. 2023; Esser et al. 2024; Batifol et al. 2025) have rapidly become the preferred approach for high fidelity generative tasks in computer vision (Zhu et al. 2024b; Wan et al. 2025; Zhu et al. 2025). Their iterative sampling procedures, however, introduce considerable latency and computational overhead. Although a variety of acceleration strategies have been proposed to reduce inference cost (Lu et al. 2022; Ma, Fang, and Wang 2024; Liu et al. 2025; Bolya and Hoffman 2023), these methods require expert intervention and bespoke engineering for each model

architecture and deployment scenario. As diffusion architectures diversify to include U-Net variants (Ronneberger, Fischer, and Brox 2015) and transformer based models (Peebles and Xie 2023; Dosovitskiy et al. 2021) and deployment contexts become more varied, there is a pressing need for automated tools capable of navigating this complex design space and producing correct, efficient implementation code without manual effort.

Recent advances in large language models (LLMs) have demonstrated their potential for end to end code generation and optimization (Zhang et al. 2024a). Domain-aware LLM agents applied to specialized tasks, such as GPU kernel optimization benchmarks (KernelBench Research Team 2025; Li et al. 2025b), have achieved notable performance gains over naive implementations. Existing research, however, has not confronted the unique challenges of generating diffusion inference code, in particular the integration of diffusion architectures with tailored acceleration techniques while satisfying strict performance and accuracy constraints in target deployment scenarios. Developers currently rely on extensive documentation libraries and hardware specific evaluations to ensure that implementations meet both precision and efficiency requirements. Crafting such solutions demands deep expertise in diffusion modeling, acceleration methods and parameter tuning, capabilities that are beyond the reach of current LLMs.

Diffusion acceleration code generation differs from general code synthesis because it requires substantial domain knowledge and environmental feedback to evaluate and refine implementations. To address these needs, we first design an environment that emulates a human developer’s workflow and supports comprehensive evaluation of LLM outputs. This environment automates the selection and tuning of acceleration techniques across a wide range of diffusion algorithms and hardware platforms and assesses both functional correctness and performance. We introduce DiffBench, a benchmark consisting of 604 prompts and reference implementations drawn from real world deployment scenarios and covering diverse model architectures, control conditions and acceleration strategies. DiffBench defines a three stage evaluation protocol that enables rigorous assess-

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ment of LLM performance.

In addition, building on human programming practice, in which developers iteratively reflect on feedback to adjust their code, we propose DiffAgent, an LLM based agent framework for diffusion acceleration code generation. DiffAgent integrates planning (Wei et al. 2023; Jiang et al. 2024), coding and debugging (Li et al. 2022; Chen et al. 2023b) stages and employs a genetic algorithm based selector to generate code and guide successive refinements. We validate DiffAgent on DiffBench using several state-of-the-art LLMs including GPT-4.1, Claude Sonnet 4, Gemini 2.5 Flash (Comanici et al. 2025) and o3-mini. Compared to direct code generation by existing models, DiffAgent yields improvements of 54.30% to 81.59% across all benchmarks. A detailed case analysis shows that our framework significantly reduces common error modes in LLM generated code. Our main contributions are summarized as follows:

- We formalize the problem of diffusion acceleration code generation, providing a structured framework to evaluate the ability of large language models to autonomously produce optimized diffusion inference code for real-world deployment scenarios.
- We introduce DiffBench, a benchmark for diffusion acceleration code generation that comprises high-quality code repositories and spans a wide variety of real-world deployment scenarios and acceleration strategies.
- We propose DiffAgent, an agentic framework built on large language models that orchestrates planning, coding, debugging and evaluation stages and employs a genetic algorithm to tune implementation parameters in order to meet prescribed accuracy and efficiency targets.
- We conduct comprehensive experiments on four large language models to demonstrate the versatility and effectiveness of DiffAgent in generating diffusion acceleration code and to quantify its performance improvements under practical constraints.

Related Work

Diffusion Model Acceleration

Recent progress in diffusion model acceleration targets both algorithmic and system-level optimizations to enable real-time high-quality generation (Rombach et al. 2022). On the algorithmic front, learnable gating mechanisms (Zhang et al. 2024b) evaluate a scalar score at each denoising step and conditionally skip redundant updates, reducing network calls by nearly forty percent with minimal impact on visual fidelity. Complementing this, feature-level caching frameworks (Ma, Fang, and Wang 2024; Zhu et al. 2024a; Qin et al. 2025) store and reuse early UNet activations across adjacent timesteps, exploiting their slow variation to avoid redundant convolutions and achieve up to fifty percent speedups alongside lower memory peaks. From the transformer world, token merging techniques (Bolya and Hoffman 2023) iteratively fuse semantically similar feature tokens within self-attention blocks, halving spatial resolution while preserving content and doubling throughput on high-resolution outputs (Ryali et al. 2023) with only slight PSNR

degradation. Finally, just-in-time compilation APIs (e.g., torch compile) fuse elementwise operations, reorder memory accesses, and eliminate Python overhead in the denoising network, delivering twenty to thirty-five percent faster sampling across a wide range of diffusion checkpoints. By uniting adaptive skipping, activation reuse, token merging, and kernel-level compilation, these approaches jointly push the efficiency frontier for diffusion-based generative modeling.

Benchmarks and Evaluation

Evaluation methodologies for generative models have evolved to emphasize task-agnostic rigor and multi-dimensional assessment (Borji 2021; Maekawa, Iso, and Bhutani 2025). Low-level code-generation benchmarks (KernelBench Research Team 2025; Li et al. 2025b; Pan et al. 2025) measure a language model’s capacity to synthesize high-performance GPU kernels, providing insight into fine-grained computational reasoning. In multimodal generation, learned alignment metrics (Radford, Kim, and Sutskever 2021; Liu et al. 2024) quantify the correspondence between text prompts and synthesized images, while domain-specific criteria, such as transparency-based metrics for surface-defect detection (Zhang, Li, and Zhou 2023), offer targeted evaluations of visual fidelity. Three-dimensional reconstruction tasks further broaden the scope: event-based data augmentation protocols (Xu, Zhang, and Zhao 2023) stress neuromorphic vision pipelines with progressive stimuli, and point-level anomaly detection benchmarks assess performance on irregular 3D structures. Collectively, these benchmarks ensure that advances in speed and model architecture are balanced by comprehensive, multi-axis validation.

Agent-based Systems

The field of agent-based systems is undergoing a significant transformation, largely driven by the integration of Large Language Models (LLMs). This evolution is prominent in multi-agent coordination and LLM-driven decision-making (Xue, Zhang, and Cao 2024; Li et al. 2025a; Schmidgall et al. 2025; Dong et al. 2024). Within multi-agent systems, novel approaches such as structural-information objectives are enabling agents to achieve sophisticated role discovery in complex environments.

A promising frontier is the rise of specialized coding agents that exploit the high-level reasoning capabilities of large language models to autonomously plan, generate, and iteratively refine software solutions (Zhang et al. 2025; Islam, Ali, and Parvez 2024; Huang et al. 2024; Guo et al. 2024). These agents collaborate, decompose tasks, write tests, and optimize performance, consistently outperforming conventional baselines in both quality and speed. Reinforcement learning can play a supporting role, e.g., aligning low-level decision policies with latent textual rewards (Kim, Park, and Lee 2023), but the principal innovation lies in coupling LLM-driven natural-language reasoning with domain-specific tool use (Wölflein et al. 2025) and long-term memory (Xu et al. 2025; Wang et al. 2025) to enable end-to-end autonomous coding workflows.

Addressing the alignment and optimization of these systems remains a critical challenge. For alignment, methods like reward shaping and preference optimization are employed to mitigate bias and prevent catastrophic forgetting (Park, Kim, and Choi 2024). Concurrently, a new paradigm of LLMs-as-optimizers has emerged. These models use evolutionary search and self-reflection to automatically refine objectives, hyperparameters, and action plans, showcasing a powerful new direction for autonomous system improvement (Yang et al. 2024; Shinn et al. 2023a; Hemberg, Moskal, and O’Reilly 2024).

Methods

In this paper, we propose an end-to-end framework for the generation and evaluation of diffusion acceleration codes. We design two complementary components: *DiffBench*, a comprehensive benchmark suite, and *DiffAgent*, a unified LLM-based framework for code generation and optimization. Together, these components enable rigorous evaluation of generated code performance and fully automated generation of high-quality, optimized diffusion acceleration code.

DiffBench: A Framework for Diffusion Acceleration Code Generation

DiffBench is a new framework for evaluating the ability of language models to generate diffusion acceleration code that satisfies user requirements. In this section, we describe the task format, contents, and evaluation metric.

Task Format. DiffBench includes 604 tasks covering a range of diffusion acceleration scenarios and is easily extensible to new development contexts. The end-to-end specification of a task is shown in Figure 1 and described below.

- **Task Input:** The input to each task consists of a user prompt and reference code compiled from open source diffusion libraries. The user prompt may specify the diffusion model type, sampling steps, chosen acceleration methods, and the target speedup to be achieved on a specified hardware platform. The reference code implementations include the definition of the diffusion model instance, the code required to perform inference under specified conditions such as text or image inputs, and the associated preprocessing and postprocessing procedures.
- **Task Output:** Given the task input, the language model must produce a new diffusion inference implementation that meets the user requirements. For example, the model may combine multiple acceleration techniques in order to achieve the required speedup while keeping the relative quality loss within an acceptable threshold.

To succeed, the model must determine (1) which acceleration methods in the diffusion sampling process will most effectively contribute to the target speedup and (2) how to tune the parameters of those methods. The model therefore requires a deep understanding of diffusion architectures and acceleration strategies, as well as experience in parameter tuning, in order to fulfill these objectives.

Task Definition. The 604 tasks in DiffBench are organized into five levels according to deployment complexity. Each level presents distinct challenges:

- **Level 1 (41 tasks):** Basic pipeline generation. Tasks at this level require the model to construct a standard diffusion inference pipeline without any acceleration techniques. Defining the correct model type and architecture is challenging due to the diversity of existing diffusion frameworks.
- **Level 2 (116 tasks):** Single-method acceleration. These tasks require the addition of a single acceleration technique, chosen from a library of methods such as fast samplers, feature reuse or token merging. Although documentation for individual acceleration techniques is mature, correctly integrating a specified method into a particular pipeline remains difficult. An inappropriate choice can over-accelerate the model and produce noisy outputs.
- **Level 3 (261 tasks):** Compositional acceleration. This level introduces combinations of multiple acceleration methods, for example, pairing faster samplers with feature reuse, to achieve greater speedup potential. The LM model must reason about how to combine techniques effectively without compromising output quality.
- **Level 4 (93 tasks):** Acceleration with explicit speedup target. Tasks at this level specify a minimum speedup requirement while constraining relative quality loss to within a threshold δ (for example, $\delta = 5\%$). Models must iteratively adjust and debug their implementations to meet these quantitative criteria.
- **Level 5 (93 tasks):** Acceleration with latency constraint. In addition to maintaining relative quality loss below δ , tasks at this level impose a maximum allowable inference latency on a given hardware platform. Both Level 4 and Level 5 demand repeated verification and debugging in the target environment to satisfy user requirements.

Task Construction. We assemble a benchmark of 604 prompts paired with corresponding ground-truth code implementations to evaluate acceleration strategies across a broad spectrum of diffusion models. The benchmark covers U-Net backbones (SD1.5 (Rombach et al. 2022), SD2.1 (Rombach et al. 2022), SDXL (Podell et al. 2023)) and transformer variants (DiT (Peebles and Xie 2023), PixArt- α (Chen et al. 2023a), PixArt- Σ (Chen et al. 2024)). It supports text-to-image, class-to-image and image-to-image conditioning (Meng et al. 2021) and spans resolutions from 256×256 to 1024×1024 . Three popular samplers (DDIM (Song, Meng, and Ermon 2020), DPM-Solver (Lu et al. 2022), UniPC (Zhao et al. 2023)) are included along with four acceleration techniques: token merging (ToMe (Bolya and Hoffman 2023)), feature reuse (DeepCache (Ma, Fang, and Wang 2024)), gated activation (T-Gate (Zhang et al. 2024b)) and half-precision computation (FP16). We present the distribution of pipeline types and explicitly required acceleration methods in the benchmark prompts. The results are shown in Figure 2.

During benchmark construction, for Levels 1–3 we generate baseline pipeline code and verify correctness and output

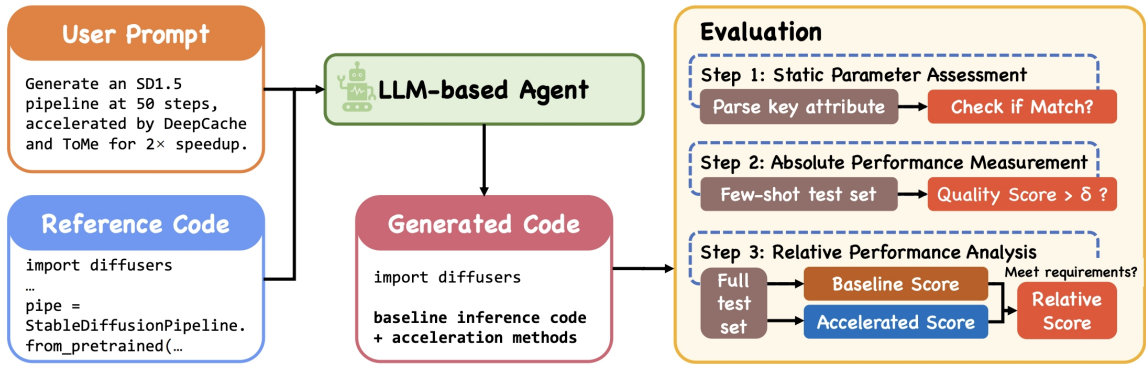


Figure 1: Overview of our proposed DiffBench. DiffBench tasks LLMs with generating diffusion acceleration code from a given user prompt and performs automated three-step evaluation.




Figure 2: Illustration of the pipeline and distribution of acceleration methods. Our benchmark covers diverse pipelines and acceleration strategies to enable comprehensive evaluation of LLM-driven acceleration code generation.

quality. For Levels 4 and 5, we define an acceleration search space for each baseline and perform a 50-iteration search on a 36-sample validation set, aiming to maximize speedup under a CLIP-Score degradation bound σ . Configurations meeting this bound are labeled as “medium” samples. We then scale the maximum speedup by factors Δ_1 and Δ_2 to generate “easy” and “hard” samples, noting that hard samples may not always have valid solutions.

Evaluation. To determine whether the generated code fulfills its intended purpose, we employ a three-stage automated evaluation pipeline comprising static parameter assessment, absolute performance measurement and relative performance analysis. These stages increase in rigor and are tailored to different prompt categories in the benchmark.

- **Static Parameter Assessment.** In the first stage, we verify that all critical parameters in the candidate code match those in the ground-truth code while permitting additional non-essential fields. We extract key attributes such as the pipeline class (e.g., `StableDiffusionPipeline`), model identifier (e.g., `stable-diffusion-v1-5`), scheduler class (e.g., `DDIMScheduler`), number of inference steps, applied acceleration methods, target resolution and the use of specific preprocessors (for example, Canny

edge detection). An automated parser extracts these fields from both the ground truth and the generated code, enforcing exact matches for ground-truth components and tolerating extraneous code. Samples that pass this static check proceed to the next stage.

- **Absolute Performance Measurement.** In the second stage, we evaluate the generated diffusion sampling code on a few-shot set of 10 examples using the CLIP-Score as a quality metric. Any sample scoring below a predefined threshold δ is marked as failing this evaluation. Evaluation samples are randomly selected from the COCO dataset (Lin et al. 2014).
- **Relative Performance Analysis.** The third stage conducts a quantitative analysis of the generated code’s performance relative to a reconstructed baseline implementation, and it is invoked when specific metrics are required, for example inference speedup in Level 4 or raw latency in Level 5. For Level 4 tasks, we define two metrics, L and U , as follows:

$$L = \frac{\frac{1}{N} \sum_{i=1}^N (S_{\text{base}}^{(i)} - S_{\text{acc}}^{(i)})}{\frac{1}{N} \sum_{i=1}^N S_{\text{base}}^{(i)}}, \quad (1)$$

$$U = \frac{\frac{1}{N} \sum_{i=1}^N T_{\text{base}}^{(i)}}{\frac{1}{N} \sum_{i=1}^N T_{\text{acc}}^{(i)}}, \quad (2)$$

where N denotes the number of validation samples. For each sample i , $S_{\text{base}}^{(i)}$ and $S_{\text{acc}}^{(i)}$ are the quality scores, i.e., CLIP scores, of the baseline and accelerated implementations, respectively. Likewise, $T_{\text{base}}^{(i)}$ and $T_{\text{acc}}^{(i)}$ denote the inference times under the baseline and accelerated implementations. A smaller value of L indicates stronger preservation of generative quality, while a larger value of U indicates greater inference efficiency. In Level 5 tasks, we substitute U with the raw latency measurement τ obtained on the specified hardware platform.

Only samples that satisfy all three stages are considered to pass. Following previous works (Zhang et al. 2024a; Zheng et al. 2023), we define the pass rate as the proportion of tasks for which the generated code is functionally correct

and meets the user’s requirements. We focus primarily on S_p , which represents the probability that the top-ranked submission is correct and reflects real-world usage, where only one suggestion is typically adopted. For tasks labeled as “hard,” we introduce the achievement rate s_a . Let $\rho = \frac{U}{U_{\text{req}}}$, where U is defined in Equation (2) and U_{req} denotes the required speedup. The **achievement rate** is then given by $S_a = \min\{\rho, 1\}$. This formulation provides a continuous scoring mechanism for evaluating performance on demanding optimization tasks.

DiffAgent: A Unified System for Code Generation and Optimization

We introduce DiffAgent, an LLM framework that unifies code generation and optimization to enhance the production of diffusion acceleration code. DiffAgent consists of four core components: planning agent, coding agent, debugging agent and genetic algorithm based selector. The overview pipeline is illustrated in Figure 3.

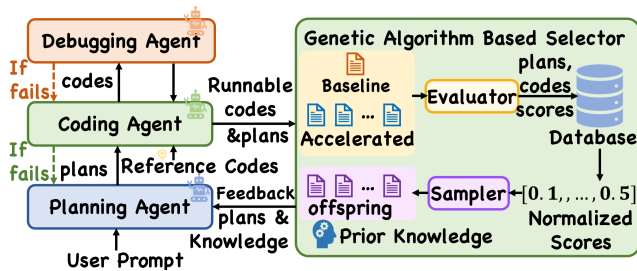


Figure 3: Overview of our proposed DiffAgent. Through multi-agent collaboration and genetic algorithm based optimization, DiffAgent generates high-quality diffusion acceleration code and iteratively refines it.

Planning Agent. The planning agent generates detailed code generation plans based on the user prompt. For code generation tasks (Levels 1–3 in DiffBench), the planning agent formulates a coding plan and passes it to the coding agent directly. For optimization tasks (Levels 4 and 5 in DiffBench), the planning agent begins by creating a baseline plan that outlines the implementation of a reference version without any acceleration techniques. It then generates prompt-conditioned acceleration plans, which define the selection, ordering, and parameterization of optimization strategies needed to achieve the desired performance target. As shown in Figure 3, during the optimization loop, the planning agent receives the generation plans of the M promising offspring. For each of these, the generation plan is augmented with a feedback report and prior tuning insights (empirical knowledge for optimizing diffusion acceleration parameters). The planning agent uses this augmented context to produce up to M refined, next-generation plans. Concurrently, to maintain diversity and avoid converging to a local optimum, it generates $P - M$ entirely new plans based on the original user prompt. This produces a total of P plans for the subsequent iteration. By default, $P = 7$ and $M = 4$.

In scenarios where no optimization is required, both P and M are set to zero.

Coding Agent. The coding agent is used to generate diffusion inference code conditioned on each specified plan. Given the baseline code plan and P acceleration code plans, it produces the baseline implementation and P accelerated variants respectively. To improve the model’s understanding of diffusion inference code structure, we compile templates of acceleration code as reference codes from established diffusion libraries. These reference codes are concatenated with each code plan before being provided to the coding agent. With this prior information, the agent can generate accurate and well-structured diffusion acceleration code.

Debugging Agent. The debugging agent verifies the correctness of the generated code by leveraging the Reflexion (Shinn et al. 2023b) architecture to identify errors and propose corrections. It coordinates with the coding agent to iteratively produce revised implementations for up to T_{debug} iterations. If no runnable version is obtained after T_{debug} attempts, the coding agent restarts code generation conditioned on the same plan (the orange arrow in Figure 3), and this regeneration process may be repeated for up to T_{code} restart cycles. If the code remains nonfunctional after these cycles, the framework backtracks to the planning stage (the green arrow in Figure 3), prompting the planning agent to generate a new plan. In the worst case, the system will invoke the LLM $T_{\text{code}} \times T_{\text{debug}}$ times per user prompt. By default, we set $T_{\text{code}} = 5$ and $T_{\text{debug}} = 3$.

Genetic Algorithm Based Selector. After obtaining runnable code implementation, we employ a genetic algorithm-based selector to evaluate the performance of each implementation and guide the agent toward producing high-performance implementations. Specifically, we assess each implementation along two dimensions: quality and efficiency. We compute the relative quality loss as defined in Equation 1 and the speedup as defined in Equation 2. These two metrics are combined using a weighted sum to produce a fitness score for each variant. Each code implementation, along with its corresponding generation plan and computed fitness score, is recorded in a database. If any implementation satisfies the user’s specified quality and efficiency requirements, the process terminates and the corresponding code is returned. Otherwise, we normalize the fitness scores to derive sampling probabilities and select M implementations as promising offspring.

For each of these M offspring, we use its associated generation plan to guide the next iteration. A feedback report is created to describe the current code’s quality and efficiency, explicitly comparing it against the performance targets. This feedback, along with the original generation plan and prior tuning insights, is passed back to the planning agent to generate the next round of plans. The genetic optimization process is repeated for up to T_{sel} iterations, where T_{sel} is set to 5 by default, until the user’s requirements are satisfied.

Method	Level 1	Level 2	Level 3	Level 4	Level 5	Avg.
o3-mini	41.46	24.14	4.60	9.68	6.45	11.92
Claude Sonnet 4	78.04	72.41	76.25	5.38	8.60	54.30
GPT-4.1	56.10	18.97	7.28	10.75	12.90	14.24
Gemini 2.5 Flash	39.02	29.31	7.66	2.15	1.08	12.09
o3-mini w/ DiffAgent	73.17 _{+31.71}	70.69 _{+46.55}	69.73 _{+65.13}	22.58 _{+12.90}	27.96 _{+21.51}	56.46 _{+44.54}
Claude Sonnet 4 w/ DiffAgent	90.24 _{+12.20}	91.38 _{+18.97}	99.23 _{+22.98}	33.33 _{+27.95}	63.44 _{+54.84}	81.59 _{+27.29}

Table 1: Comparison of pass rate S_p on the DiffBench benchmark across various LLMs, reported per difficulty level. All values are percentages, gains relative to the corresponding baseline are shown as subscripts in **bold**.

Experiments

We conducted extensive experiments to address three research questions: (1) How well do existing code-generation LLMs perform when evaluated on DiffBench, and what challenges do they face? (2) To what extent does our DiffAgent outperform these baseline LLMs in generating diffusion-model acceleration code? (3) How effectively do the individual modules we’ve integrated into the agent system support the code-writing process?

Experimental Setup

Benchmark. To evaluate our method on diffusion acceleration code generation, we conducted experiments on the DiffBench benchmark. We set $\sigma = 5\%$, $\Delta_1 = 0.8$ and $\Delta_2 = 1.2$.

Base LLMs. We evaluated the performance of several recent large language models on our benchmark: GPT-4.1, Claude Sonnet 4, Gemini 2.5 Flash, and o3-mini. These represent some of the most advanced models available at the time of our evaluation.

DiffBench Baseline Evaluation

In this section, we evaluate a variety of language models on DiffBench without further fine-tuning to assess their out-of-the-box performance and identify failure modes.

Baseline. We applied our custom diffusion-model acceleration benchmark, DiffBench, to evaluate four leading LLMs, including both foundation models and inference-optimized models. DiffBench proves more challenging than existing suites: as shown in Table 1, despite their strong performance on standard coding tasks, these models achieve pass rates below 35% on our acceleration challenges.

Error Analysis. Figure 4 provides a comprehensive breakdown of the failure modes in diffusion acceleration code generated by Claude Sonnet 4, both with and without DiffAgent integration. We classify these failures into five categories: Compile Error and Key Attributes Error, corresponding to evaluation stage 1; Absolute Quality Error for stage 2; and Relative Quality Error and Relative Speed Error for stage 3. Note that Relative Quality Error and Relative Speed Error are not mutually exclusive, so the sum of error rates may exceed 100%.




Figure 4: Error mode analysis on the DiffBench benchmark across various LLMs, reported per difficulty level. We categorize failure modes of generated code into five failure models. All values are percentages.

Our results indicate that DiffAgent yields substantial improvements in addressing fundamental code generation issues. Across all difficulty levels, the rate of Compile Error is markedly reduced. For example, at Level 5, the Compile Error rate decreases from 31.18% for the baseline Claude Sonnet 4 model to 2.15% when DiffAgent is employed. In addition, the incidence of Key Attributes Error is significantly lower, and Low Quality Error is completely eliminated in the scenarios tested, dropping to 0% in several cases.

When examining tasks that involve complex performance constraints (Levels 4 and 5), DiffAgent demonstrates a pronounced ability to balance competing objectives effectively. While the baseline model exhibits high rates of Relative Quality Error and Relative Speed Error, DiffAgent achieves a dramatic reduction in both error types at these higher difficulty levels. This finding underscores the capacity of DiffAgent to optimize for quality and speed simultaneously without compromising one objective for the other.

DiffAgent Coding Performance

In these experiments, we employed our diffusion acceleration code generation agent, DiffAgent, to augment the capabilities of leading LLMs. The results, summarized in Table 1, show that DiffAgent consistently delivers substantial performance improvements across all base models and scales. For the Claude Sonnet 4 model, we observe a maximum gain of 27.29%. Overall improvements range from 54.30% to 81.59%, validating the effectiveness of our proposed approach. These findings demonstrate that the inte-

Method	S_p						S_a (Hard-task)		
	Level 1	Level 2	Level 3	Level 4	Level 5	Avg.	Level 4	Level 5	Avg.
DiffAgent	90.24	91.38	99.23	33.33	63.44	81.59	56.61	79.92	68.27
w/o Knowledge Base	82.93	74.14	77.78	26.88	47.31	64.90	39.54	52.33	45.94
w/o GA	90.24	91.38	99.23	4.30	4.30	67.88	6.28	10.03	8.16
w/o Debugging Agent	87.80	77.59	81.99	31.18	33.33	66.23	56.99	67.40	62.02

Table 2: Ablation study of key components in DiffAgent. We report pass rate S_p across five task levels and hard-task achievement rates S_a on Level 4 and Level 5.

grated modules and domain knowledge within DiffAgent provide the necessary guidance for LLMs to generate accurate diffusion-model acceleration code and to successfully address the challenges of diffusion acceleration tasks.

For the Claude Sonnet 4 model, we illustrate the performance gains achieved by integrating DiffAgent across five levels of benchmark code, as shown in Figure 3. The most significant improvements appear in the third and fourth categories, because DiffAgent more accurately incorporates acceleration methods and better understands their parameter semantics, enabling correct usage that balances speed and accuracy to meet the prompt requirements.

Ablation Study

Effectiveness of Each Component. To assess the individual impact of DiffAgent’s three modules: the Knowledge Base, Genetic Algorithm and Debugging Agent. We conducted a leave-one-out ablation on the Claude Sonnet 4 backbone and evaluated performance on the full benchmark (Table 2). Removing any module degrades performance: the average pass rate S_p decreases from 81.59% to a range of 64.90%–67.88%, and the hard-task achievement rate S_a falls from 68.27% to between 8.16% and 62.02%. (1) The Genetic Algorithm module is essential for complex tasks. Without it, pass rates on Levels 4 and 5, with S_p dropping to 4.30% and S_a to 8.16%. This result underscores the role of GA in exploring and refining candidate programs for challenging problems. (2) The Knowledge Base provides broad benefits across all levels. Its removal causes the largest absolute decline in overall pass rate, a 16.69% reduction, demonstrating that domain knowledge accelerates search and improves solution quality. (3) The Debugging Agent primarily enhances robustness on high-difficulty tasks: without it, Level 5 pass rate declines by approximately 30% and S_a by 6.25%, showing that iterative error analysis prevents cascading mistakes. Collectively, these findings confirm that the three modules operate synergistically to deliver DiffAgent’s state-of-the-art performance.

P	$T_{sel} = 2$		$T_{sel} = 4$		$T_{sel} = 6$	
	S_p	S_a	S_p	S_a	S_p	S_a
4	18.28	32.41	26.88	45.42	31.18	59.91
7	30.11	50.08	33.33	63.44	37.63	64.80
10	26.88	59.78	37.63	66.48	38.71	66.28

Table 3: Comparison of P and T_{sel} .

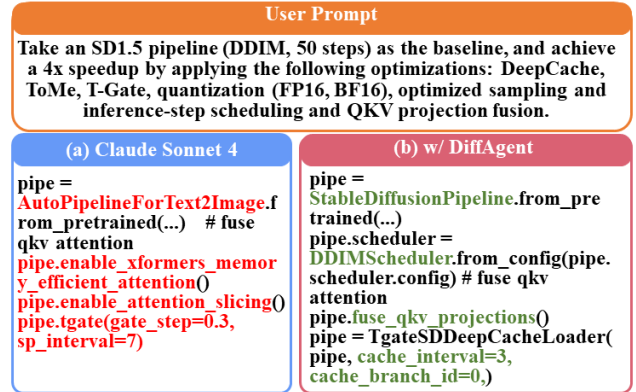


Figure 5: Sample visualization of code generation results. Incorrect code segments originally generated by the LLM are highlighted in red, while corrected segments produced by DiffAgent are shown in green.

Comparison of GA Hyperparameters. To investigate the impact of two key genetic algorithm hyperparameters: population size P and number of iterations T_{sel} on code generation, we conducted experiments on the fourth and fifth levels of DiffBench. The results in Table 3 indicate that performance plateaus once the population size exceeds 7 and the generation count exceeds 4. Beyond these thresholds, further increases yield negligible gains. To balance efficacy and runtime, we therefore set P to 7 and T_{sel} to 4.

Qualitative Analysis

Figure 5 contrasts code generated directly by the LLM with the version refined by DiffAgent. Our approach not only produces code that more faithfully meets the prompt, exhibiting consistent style and correct key attributes (e.g., pipeline, sampler), but also integrates acceleration techniques like DeepCache and T-Gate, making the resulting diffusion-acceleration code immediately applicable to real-world development scenarios.

Conclusion

We introduce DiffBench, the first benchmark for LLM-generated diffusion pipelines, and DiffAgent, an LLM-driven automation framework for their acceleration. Evaluations show DiffAgent delivers substantial speedups with minimal quality loss across diverse scenarios, surpassing leading LLMs on complex and latency-sensitive tasks.

References

- Batifol, S.; Blattmann, A.; Boesel, F.; Consul, S.; Diagne, C.; Dockhorn, T.; English, J.; English, Z.; Esser, P.; Kulal, S.; et al. 2025. FLUX. 1 Kontext: Flow Matching for In-Context Image Generation and Editing in Latent Space. *arXiv e-prints*, arXiv:2506.
- Bolya, D.; and Hoffman, J. 2023. Token merging for fast stable diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4599–4603.
- Borji, A. 2021. Pros and Cons of GAN Evaluation Measures: New Developments. arXiv:2103.09396.
- Chen, J.; Ge, C.; Xie, E.; Wu, Y.; Yao, L.; Ren, X.; Wang, Z.; Luo, P.; Lu, H.; and Li, Z. 2024. PixArt- Σ : Weak-to-Strong Training of Diffusion Transformer for 4K Text-to-Image Generation. arXiv:2403.04692.
- Chen, J.; Yu, J.; Ge, C.; Yao, L.; Xie, E.; Wu, Y.; Wang, Z.; Kwok, J.; Luo, P.; Lu, H.; and Li, Z. 2023a. PixArt- α : Fast Training of Diffusion Transformer for Photorealistic Text-to-Image Synthesis. arXiv:2310.00426.
- Chen, X.; Lin, M.; Schärli, N.; and Zhou, D. 2023b. Teaching Large Language Models to Self-Debug. arXiv:2304.05128.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Dong, Y.; Jiang, X.; Jin, Z.; and Li, G. 2024. Self-collaboration Code Generation via ChatGPT. arXiv:2304.07590.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Guo, D.; Zhu, Q.; Yang, D.; Xie, Z.; Dong, K.; Zhang, W.; Chen, G.; Bi, X.; Wu, Y.; Li, Y. K.; Luo, F.; Xiong, Y.; and Liang, W. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming – The Rise of Code Intelligence. arXiv:2401.14196.
- Hemberg, E.; Moskal, S.; and O’Reilly, U.-M. 2024. Evolving Code with A Large Language Model. arXiv:2401.07102.
- Huang, D.; Zhang, J. M.; Luck, M.; Bu, Q.; Qing, Y.; and Cui, H. 2024. AgentCoder: Multi-Agent-based Code Generation with Iterative Testing and Optimisation. arXiv:2312.13010.
- Islam, M. A.; Ali, M. E.; and Parvez, M. R. 2024. MapCoder: Multi-Agent Code Generation for Competitive Problem Solving. arXiv:2405.11403.
- Jiang, X.; Dong, Y.; Wang, L.; Fang, Z.; Shang, Q.; Li, G.; Jin, Z.; and Jiao, W. 2024. Self-planning Code Generation with Large Language Models. arXiv:2303.06689.
- KernelBench Research Team. 2025. KernelBench: Can LLMs Write Efficient GPU Kernels? arXiv:2502.10517.
- Kim, H.; Park, S.; and Lee, J. 2023. LaRe: Latent Reward Learning for Episodic Reinforcement. In *Proc. NeurIPS*.
- Li, A.; Xie, Y.; Li, S.; Tsung, F.; Ding, B.; and Li, Y. 2025a. Agent-Oriented Planning in Multi-Agent Systems. In *The Thirteenth International Conference on Learning Representations*.
- Li, J.; Li, S.; Gao, Z.; Shi, Q.; Li, Y.; Wang, Z.; Huang, J.; Wang, H.; Wang, J.; Han, X.; et al. 2025b. TritonBench: Benchmarking Large Language Model Capabilities for Generating Triton Operators. *arXiv preprint arXiv:2502.14752*.
- Li, Y.; Choi, D.; Chung, J.; Kushman, N.; Schrittwieser, J.; Leblond, R.; Eccles, T.; Keeling, J.; Gimeno, F.; Dal Lago, A.; Hubert, T.; Choy, P.; de Masson d’Autume, C.; Babuschkin, I.; Chen, X.; Huang, P.-S.; Welbl, J.; Goyal, S.; Cherepanov, A.; Molloy, J.; Mankowitz, D. J.; Sutherland Robson, E.; Kohli, P.; de Freitas, N.; Kavukcuoglu, K.; and Vinyals, O. 2022. Competition-level code generation with AlphaCode. *Science*, 378(6624): 1092–1097.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, H.; Zhang, W.; Xie, J.; Faccio, F.; Xu, M.; Xiang, T.; Shou, M. Z.; Perez-Rua, J.-M.; and Schmidhuber, J. 2025. Faster Diffusion via Temporal Attention Decomposition. arXiv:2404.02747.
- Liu, X.; Lei, X.; Wang, S.; Huang, Y.; Feng, Z.; Wen, B.; Cheng, J.; Ke, P.; Xu, Y.; Tam, W. L.; Zhang, X.; Sun, L.; Gu, X.; Wang, H.; Zhang, J.; Huang, M.; Dong, Y.; and Tang, J. 2024. AlignBench: Benchmarking Chinese Alignment of Large Language Models. arXiv:2311.18743.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35: 5775–5787.
- Ma, X.; Fang, G.; and Wang, X. 2024. Deepcache: Accelerating diffusion models for free. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15762–15772.
- Maekawa, S.; Iso, H.; and Bhutani, N. 2025. Holistic Reasoning with Long-Context LMs: A Benchmark for Database Operations on Massive Textual Data. In *The Thirteenth International Conference on Learning Representations*.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.
- Pan, Z.; Cao, R.; yongchang Cao; MA, Y.; Li, B.; Huang, F.; Liu, H.; and Li, Y. 2025. Codev-Bench: How Do LLMs Understand Developer-Centric Code Completion?

- Park, S.; Kim, N.; and Choi, J. 2024. SPO: Sequential Preference Optimization for Multi-Dimensional Alignment. In *Proc. ICLR*.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Qin, R.; Wang, Q.; Sun, M.; Zhu, H.; Zhou, C.; and Wang, B. 2025. Accelerating Diffusion-based Super-Resolution with Dynamic Time-Spatial Sampling. *arXiv preprint arXiv:2505.12048*.
- Radford, A.; Kim, J. W.; and Sutskever, I. 2021. CLIP: Connecting Text and Images. In *Proc. ICML*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv:1505.04597*.
- Ryali, C.; Hu, Y.-T.; Bolya, D.; Wei, C.; Fan, H.; Huang, P.-Y.; Aggarwal, V.; Chowdhury, A.; Poursaeed, O.; Hoffman, J.; Malik, J.; Li, Y.; and Feichtenhofer, C. 2023. Hiera: A Hierarchical Vision Transformer without the Bells-and-Whistles. *arXiv:2306.00989*.
- Schmidgall, S.; Su, Y.; Wang, Z.; Sun, X.; Wu, J.; Yu, X.; Liu, J.; Moor, M.; Liu, Z.; and Barsoum, E. 2025. Agent Laboratory: Using LLM Agents as Research Assistants. *arXiv:2501.04227*.
- Shinn, N.; Cassano, F.; Berman, E.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023a. Reflexion: Language Agents with Verbal Reinforcement Learning. *arXiv:2303.11366*.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023b. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36: 8634–8652.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Wan, T.; Wang, A.; Ai, B.; Wen, B.; Mao, C.; Xie, C.-W.; Chen, D.; Yu, F.; Zhao, H.; Yang, J.; et al. 2025. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*.
- Wang, Z.; Mao, J.; Fried, D.; and Neubig, G. 2025. Agent Workflow Memory.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q.; and Zhou, D. 2023. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *arXiv:2201.11903*.
- Wölflein, G.; Ferber, D.; Truhn, D.; Arandjelović, O.; and Kather, J. N. 2025. LLM Agents Making Agent Tools. *arXiv:2502.11705*.
- Xu, M.; Zhang, Y.; and Zhao, Q. 2023. EventZoom: Progressive Augmentation for Neuromorphic Vision. In *Proc. ICML*.
- Xu, W.; Mei, K.; Gao, H.; Tan, J.; Liang, Z.; and Zhang, Y. 2025. A-MEM: Agentic Memory for LLM Agents. *arXiv:2502.12110*.
- Xue, C.; Zhang, H.; and Cao, H. 2024. Multi-agent modelling and analysis of the knowledge learning of a human-machine hybrid intelligent organization with human-machine trust. *Systems Science & Control Engineering*, 12(1): 2343301.
- Yang, C.; Wang, X.; Lu, Y.; Liu, H.; Le, Q. V.; Zhou, D.; and Chen, X. 2024. Large Language Models as Optimizers. *arXiv:2309.03409*.
- Zhang, C.; Li, W.; and Zhou, H. 2023. TransFusion: Transparency-based Evaluation for Surface Defect Detection. In *Proc. ICCV*.
- Zhang, K.; Li, J.; Li, G.; Shi, X.; and Jin, Z. 2024a. Codeagent: Enhancing code generation with tool-integrated agent systems for real-world repo-level coding challenges. *arXiv preprint arXiv:2401.07339*.
- Zhang, W.; Liu, H.; Xie, J.; Faccio, F.; Shou, M. Z.; and Schmidhuber, J. 2024b. Cross-Attention Makes Inference Cumbersome in Text-to-Image Diffusion Models. *arXiv preprint arXiv:2404.02747v1*.
- Zhang, Y.; Yang, J.; Yuan, Y.; and Yao, A. C.-C. 2025. Cumulative Reasoning with Large Language Models. *arXiv:2308.04371*.
- Zhao, W.; Bai, L.; Rao, Y.; Zhou, J.; and Lu, J. 2023. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36: 49842–49869.
- Zheng, Q.; Xia, X.; Zou, X.; Dong, Y.; Wang, S.; Xue, Y.; Shen, L.; Wang, Z.; Wang, A.; Li, Y.; et al. 2023. Codegeex: A pre-trained model for code generation with multilingual benchmarking on humaneval-x. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5673–5684.
- Zhu, H.; Pan, T.; Qin, R.; Yong, J.-H.; and Wang, B. 2025. ReCon: Region-Controllable Data Augmentation with Rectification and Alignment for Object Detection. *arXiv preprint arXiv:2510.15783*.
- Zhu, H.; Tang, D.; Liu, J.; Lu, M.; Zheng, J.; Peng, J.; Li, D.; Wang, Y.; Jiang, F.; Tian, L.; Tiwari, S.; Sirasao, A.; Yong, J.; Wang, B.; and Barsoum, E. 2024a. DiP-GO: A Diffusion Pruner via Few-step Gradient Optimization. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 92581–92604. Curran Associates, Inc.
- Zhu, H.; Yang, L.; Yong, J.-H.; Yin, H.; Jiang, J.; Xiao, M.; Zhang, W.; and Wang, B. 2024b. Distribution-aware data expansion with diffusion models. *Advances in Neural Information Processing Systems*, 37: 102768–102795.