# QiMeng-SALV: Signal-Aware Learning for Verilog Code Generation

<sup>1</sup>State Key Lab of Processors, Institute of Computing Technology, CAS

<sup>2</sup>University of Chinese Academy of Sciences

<sup>3</sup>University of Science and Technology of China

https://zy1xxx.github.io/SALV

## **Abstract**

The remarkable progress of Large Language Models (LLMs) presents promising opportunities for Verilog code generation which is significantly important for automated circuit design. The lacking of meaningful functional rewards hinders the preference optimization based on Reinforcement Learning (RL) for producing functionally correct Verilog code. In this paper, we propose Signal-Aware Learning for Verilog code generation (QiMeng-SALV) by leveraging code segments of functionally correct output signal to optimize RL training. Considering Verilog code specifies the structural interconnection of hardware gates and wires so that different output signals are independent, the key insight of QiMeng-SALV is to extract verified signal-aware implementations in partially incorrect modules, so as to enhance the extraction of meaningful functional rewards. Roughly, we verify the functional correctness of signals in generated module by comparing with that of reference module in the training data. Then abstract syntax tree (AST) is employed to identify signal-aware code segments which can provide meaningful functional rewards from erroneous modules. Finally, we introduce signal-aware DPO which is optimized on the correct signal-level code segments, thereby preventing noise and interference from incorrect signals. The proposed QiMeng-SALV underscores the paradigm shift from conventional module-level to fine-grained signal-level optimization in Verilog code generation, addressing the issue of insufficient functional rewards. Experiments demonstrate that our method achieves state-of-the-art performance on VerilogEval and RTLLM, with a 7B parameter model matching the performance of the DeepSeek v3 671B model and significantly outperforming the leading open-source model CodeV trained on the same dataset. Our code is available at https://github.com/zy1xxx/SALV.

# 1 Introduction

Circuit design is inherently complex and time-consuming, particularly with respect to Hardware Description Language (HDL) such as Verilog. Automating the HDL coding process is of paramount importance as it significantly enhances the efficiency of circuit design. The remarkable progress of Large Language Models (LLMs) in program code generation, such as Python, presents promising

<sup>\*</sup>Corresponding author. Contact: {zhangyang22s2,cyj}@ict.ac.cn




Figure 1: Due to Verilog's characteristics of relatively independent and parallel signals, we can extract the code implementation of a specific signal through AST. Here there are two output signals, we can separately extract their related code implementations. Among them, the implementation of signal a is incorrect while signal d is correctly implemented. We can utilize the code implementation of the correct signal d to provide functional correctness rewards for RL.

opportunities for Verilog code generation. To equip LLMs with superior code generation capabilities, post-training procedures are essential, primarily comprising Supervised Fine-Tuning (SFT) and preference-based Reinforcement Learning (RL). SFT provides foundational programming language syntax comprehension and domain-specific knowledge, while RL optimization techniques, like Proximal Policy Optimization (PPO) [1], Group Relative Policy Optimization (GRPO) [2], and Direct Preference Optimization (DPO) [3], enhance the model's ability to produce functionally correct code.

Currently, most Verilog code generation methods primarily focus on the SFT stage, particularly through dataset preparation [4, 5, 6]. Few methods [7, 8] employ RL through using code structure similarity with reference code as the reward. However, relying solely on code structure similarity as a reward has inherent limitations, since a module may has many different Verilog implementations which exhibit structural differences but are functionally correct. Thus, it is more reasonable to use functional correctness as the reward in RL, which requires the model to sample functionally correct module implementations to provide the reward. However, the lack of high-quality SFT training data frequently results in inadequate initial code generation capability, obstructing the acquisition of meaningful functional rewards and consequently impairing RL optimization performance. Thus, enhancing the availability of actionable functional rewards to facilitate effective RL optimization remains an open research challenge.

In contrast to high-level program language that defines sequential execution behavior, Verilog code specifies the structural interconnection of hardware gates and wires, granting inherent independence between different output signals. This characteristic enables scenarios where individual signal implementations remain functionally correct despite errors in the overall module implementation. As shown in Fig.1, the faulty implementation of the signal a results in total malfunction of the top module, whereas the signal d maintains proper functionality. Due to Verilog's inherent signal independence and parallel processing nature, we can easily extract a certain signal implementation from the top module through abstract syntax tree (AST). Such partially valid signal-level implementations can be leveraged within RL frameworks to derive meaningful functional rewards, thereby expanding the effective sample space for RL optimization.

Based on the above analysis, in this work, we propose Signal-Aware Learning for Verilog code generation (QiMeng-SALV) by leveraging code segments of functionally correct output signal to optimize RL training. The key insight of QiMeng-SALV is to extract verified signal-aware implementations in partially incorrect modules, so as to enhance the extraction of meaningful

functional rewards. Specifically, to verify the functional correctness of signals, some random test inputs are generated to compare signal-level outputs between generated module and reference module in the training data. Additionally, abstract syntax tree (AST) analysis is employed to precisely identify preferred and dispreferred code segments, thus these partially correct signal-level implementation can be utilized to generate meaningful functional rewards. Finally, to enable the model to learn correct signal implementations from erroneous modules, we introduce signal-aware DPO which is based on DPO, a method renowned for its computational efficiency, rapid convergence, and strong performance. Unlike standard DPO that computes probabilities for all module code tokens, signalaware DPO exclusively calculates token probabilities for the correct signal-level code segments from the preferred and dispreferred samples, thereby preventing noise and interference from incorrect signals during RL training. Crucially, the proposed QiMeng-SALV expands the effective training dataset by incorporating all modules containing any valid signal implementations, regardless of overall module correctness. This represents a fundamental shift from conventional module-level to fine-grained signal-level optimization in RL-based Verilog code generation, addressing the issue of insufficient functional rewards in RL training. To the best of our knowledge, this is the first fine-grained signal-level reinforcement learning algorithm designed for Verilog code generation.

Comprehensive evaluations demonstrate that QiMeng-SALV establishes new state-of-the-art results on both VerilogEval [9, 10] and RTLLM benchmarks [11, 12]. The 7B-parameter model achieves remarkable 62.6% pass@1 and 75.1% pass@5 accuracy on RTLLM v1.1, not only matching the performance of the 671B-parameter DeepSeek-V3 [13], but also delivering substantial improvements over CodeV [4], the leading open-source model of comparable scale. Ablation studies further validate the substantial impact of filtering incorrect signal segments, confirming the efficacy of QiMeng-SALV's design.

## 2 Related Work

#### 2.1 RTL code generation

Recent years have witnessed remarkable achievements of Large Language Models (LLMs) in software code generation tasks, prompting both academia and industry to actively investigate their potential applications in Hardware Description Language (HDL) code generation [14]. This novel research frontier has achieved significant progress, unveiling promising paradigms for hardware design automation.

Instruction tuning has emerged as a pivotal methodology for enhancing LLMs' task adaptability, owing to its methodological simplicity and demonstrated performance gains [4]. This technique shows broad application prospects for automated HDL generation. However, compared to software programming languages (such as Python, C++), the domain-specific expertise of hardware design languages and the scarcity of high-quality training data result in current LLMs' performance in HDL generation tasks being significantly inferior to their performance in software programming languages [15]. Therefore, many works focus on the construction of hardware language datasets [6, 5, 12, 4] and employ Supervised Fine-Tuning (SFT) strategies to enhance LLMs' HDL generation capability. However, the heterogeneous quality of training data, whether scraped from public sources or synthesized by advanced LLMs, remains a critical bottleneck, compromising both validity and reliability. This has spurred innovative approaches incorporating verification frameworks and feedback mechanisms during model training/inference to overcome performance bottlenecks and improve RTL generation quality [16, 14, 17].

# 2.2 RL Training for LLM

The effectiveness of supervised fine-tuning (SFT) in hardware description languages is fundamentally constrained by data quality and diversity, particularly in specialized domains where conventional approaches often fail to meet professional requirements[18]. This limitation has driven interest in reinforcement learning (RL) techniques that can enhance model capabilities beyond what SFT alone can achieve.

Reinforcement Learning from Human Feedback (RLHF)[19] is among the earliest and most impactful approaches for aligning language models with human intent with Proximal Policy Optimization[1]. Subsequent work [20, 21, 3, 22, 23] has focused on developing more efficient alternatives, with

Direct Preference Optimization (DPO)[3] emerging as particularly impactful due to its simplicity and stable training process. The recent Group Relative Policy Optimization (GRPO) approach[2] further advanced this direction by eliminating the need for explicit reward modeling, demonstrating particular success in RL-based post-training optimization for higher-performance LLMs.

In the domain of Verilog generation, current RL-based training approaches face two key challenges: (1) highly rely on structural similarity metrics [7, 8] that cannot properly evaluate functional equivalence between different implementations, and (2) the scarcity of high-quality SFT data limits the initial model capability needed for effective RL optimization. While recent work has explored various reward formulations [24, 25, 7] and feedback mechanisms [26, 27, 17], these approaches often introduce evaluation biases, limiting their reliability in assessing functional correctness. VeriPrefer [28] mitigates such biases by employing module-level functional rewards during reinforcement learning. Our approach differs fundamentally: we adopt signal-level functional rewards, allowing partially correct modules to still provide valuable feedback during training when they contain correctly implemented signals. Consequently, our method delivers denser and more informative reward signals, leading to superior reinforcement learning performance.

## 3 Method

#### 3.1 Framework Overview

In this work, we propose an innovative training methodology named QiMeng-SALV, which is designed to leverage correct signal-aware implementations within erroneous module codes for Verilog code generation. We start from a naive code generator, which is typically fine-tuned using a Verilog training dataset  $D = \{(x,y)\}$  based on a general-purpose LLM, where (x,y) is the description-code pair. The naive code generator samples k candidate module code implementations  $E = \{y_1, y_2, ..., y_k\}$  given each module design prompt x. Our primary objective is to extract and learn correct signal-aware implementations from these candidates for improving RL optimization. Here we employ DPO as our RL optimization approach, a technique widely recognized for its computational efficiency, fast convergence, and robust performance.

Our training methodology comprises three stages: 1) Signal-aware Verification: By generating random input signals and comparing the output signals between the generated module and the reference module in the training set, we identify correct output signals to obtain preference dataset  $P = \{(y_w, y_l, c)\}$ , where  $y_w$  and  $y_l$  represent preferred and dispreferred module code respectively, c represent the contrast signal which is correct in  $y_w$  and incorrect in  $y_l$ . 2) Signal-aware Code Extraction: Leveraging abstract syntax tree (AST) analysis, we establish signal dependency graphs and isolate relevant preferred and dispreferred code segments  $(S_w^c, S_l^c)$  corresponding to the contrast signals c from  $y_w$  and  $y_l$ , respectively. 3) Signal-aware DPO training: By computing token probabilities only for the preferred and dispreferred code segments related to contrast signals, we employ DPO to enable the model to learn correct signal implementations, even when the entire module implementation is erroneous. The following sections provide detailed explanations of each stage.

# 3.2 Signal-aware Verification

Enabling RL to train models that generate functionally correct code fundamentally requires functional reward derived from verified implementations. This makes functional correctness verification a critical challenge when applying RL to LLM-based code generation, especially for hardware description languages such as Verilog. Conventional verification methods depend on testbenches to assess module functional correctness, but the lack of such testbenches in training datasets substantially restricts the capacity to provide meaningful correctness feedback during RL training.

Our approach addresses this limitation through an automated signal verification process for generating functional reward. Considering the functional reward is provided by a preference dataset in DPO, we therefore establish the preference dataset based on automated functional verification by three steps: 1) generating comprehensive random input stimuli, 2) verifying generated signals by comparison with reference modules, 3) collecting preference dataset from the verifying results of generated signals, as depicted in Fig.2 (b).

First, we use Yosys [29] to analyze the reference module head, automatically generating N sets of input signals at equal time intervals. Second, the input signals are simultaneously applied to the




Figure 2: Overview of the QiMeng-SALV framework. a) The proposed QiMeng-SALV comprises three stage: signal-aware verification, signal-aware code extraction, and signal-aware DPO training. b) In signal-aware verification stage, verification is performed by analyzing output signal discrepancies between generated modules and their reference counterparts, allowing precise identification of correctly functioning output signals. c) In signal-aware code extraction stage, AST parsing reveals critical dependencies between output signals and intermediate signals to obtain relevant preferred and dispreferred code segments pertinent to the contrastive signals.

reference module y and generated modules  $E = \{y_1, y_2, ..., y_k\}$ . By comparing the output signals produced by the generated modules with the corresponding output signals from the reference module, we can verify the correctness of generated modules, and obtain the pairs  $\{(y_1, c_1), (y_2, c_2), ..., (y_k, c_k)\}$ , where  $c_i$  is the correct output signals for generated module  $y_i$ . Notably, only output signals exhibiting complete matching across all N sets of input signals are verified as correct, while any discrepancies result in classification as erroneous signals. Third, to select preference samples suitable for DPO training, we choose sample pairs that there exists a set of contrast signals that is correct in the preferred sample but incorrect in the dispreferred sample. This can be formally expressed as:

$$P = \{(y_w, y_l, c)\}, \exists c \in c_w \land c \notin c_l, \tag{1}$$

where P denotes the preference dataset, c represents the contrast signals,  $c_w$  and  $c_l$  represent the correct signals in the preferred sample  $y_w$  and dispreferred sample  $y_l$  respectively.

# 3.3 Signal-aware Code Extraction

In RL-based code generation, there exist two typical feedback mechanisms: outcome feedback and process feedback. Compared to outcome feedback which only indicates whether the final output is correct or not, process feedback can precisely identify which parts of the code implementation are correct, thereby providing a more dense reward for more effective learning. Therefore, by extracting the specific code implementation corresponding to the output signal and adopt the RL method based on process feedback is of great significance for code generation.

In contrast to program code, which sequentially describes execution logic, Verilog code characterizes the interconnection between wires and gates. This indicates that signal-describing code blocks are largely self-contained, allowing individual signal implementations to maintain functional accuracy even when the broader module contains errors. These partially correct signal-level implementations can be utilized in RL frameworks to generate meaningful functional rewards, thus increasing the viable sample space for RL optimization.

Based on the above analysis, in the signal-aware code extraction stage, for given target signals, we parse the Abstract Syntax Tree (AST) to establish dependency relationships between signals, and then extract the complete code segment corresponding to the target signals, as depicted in Fig.2 (c). Specifically, given the module code of the target signal, we use Yosys [29] to parse the module

code and obtain the AST. Subsequently, we analyze the AST to derive the dependency relationships between all the output signals and intermediate signals defined in the module code, forming a signal topology graph. By analyzing this topology graph and performing backward traversal from the target output signal leaf nodes, we obtain the dependent signals of the target signal. Based on the locations of these dependent signals in the AST, we retain their corresponding code segments, thereby obtaining the complete code segment of the target signal. To facilitate the model's learning by using code segment of contrast signal during DPO training, given the contrast signals c, we isolates the relevant code segment  $S_w^c$  from the preferred sample  $y_w$ . Meanwhile, to enable comparison with the preferred sample, we also extract the code segment  $S_l^c$  of contrast signal c from the dispreferred sample  $y_l$ .

#### 3.4 Signal-aware DPO Training

In this stage, we employ the preferred and dispreferred code segments related to contrast signals in DPO training to strengthen the model's ability of generating functional correct signal implementations. Standard DPO approaches uniformly increase the likelihood of all tokens in preferred code samples while suppressing all tokens in dispreferred samples. This paradigm inherently assumes the absolute correctness of the preferred samples, which is often violated in practice due to suboptimal SFT models. Because of the limited training data of Verilog code implementations during the SFT stage, the SFT models frequently fail to generate completely accurate code samples. Our key insight recognizes that while generating completely accurate modules remains challenging, we can still leverage partially code segments of functionally correct output signal from incorrect modules to optimize DPO training. Therefore, we introduce signal-aware DPO, an improved DPO algorithm for Verilog code generation, which increases the probability of correct signal-related code in preferred samples while decreasing the probability of erroneous signals in dispreferred samples. To achieve this, during loss computation, we only calculate the token probabilities for the contrasting signal-related code segments in both preferred and dispreferred samples, while ignoring the probabilities of other code segment tokens. Formally expressed as:

$$\mathcal{L}(\pi_{\theta}; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l, S_w^c, S_l^c) \sim \mathcal{D}} \left[ \log \sigma \left( \beta \sum_{y_t \in S_w^c} \log \frac{\pi_{\theta}(y_t | y_{w, < t}, x)}{\pi_{\text{ref}}(y_t | y_{w, < t}, x)} - \beta \sum_{y_t \in S_c^c} \log \frac{\pi_{\theta}(y_t | y_{l, < t}, x)}{\pi_{\text{ref}}(y_t | y_{l, < t}, x)} \right) \right], \quad (2)$$

where  $y_w$  and  $y_l$  represents the preferred and dispreferred sample respectively.  $S_w^c$  denotes code segments related to contrast signals c in preferred samples, while  $S_l^c$  represents the same signal related segments in dispreferred samples.  $\pi_\theta$  is the policy model to be optimized and  $\pi_{\rm ref}$  is the reference model used for regularizing  $\pi_\theta$  with Kullback-Leibler divergence and  $\beta$  is a constant to control the degree of regularization.

Through the objective function in Eq. (2), signal-aware DPO shift DPO's learning focus from entire modules to individual signals, refining the learning granularity and increasing the quantity of functional correctness rewards in DPO, thereby enhancing optimization effectiveness.

## 4 Experiment

# 4.1 Experiment Setup

**Datasets.** Our training dataset is sourced from CodeV [4], consisting of 165K samples obtained by crawling all publicly available Verilog module code on GitHub. However, we observed that a subset of these modules are not syntactically valid, prompting us to filter out such cases and retain a cleaned dataset of 135K samples. The 135K dataset is initially employed to fine-tune the general-purpose code generation model Qwen2.5 Coder Instruct [30], yielding a naive code generator. For every prompt in the 135K dataset, this generator produces 5 candidate module codes, while the corresponding module codes in the dataset serve as reference codes to validate signal correctness during Signal-aware Verification.

**Training Settings.** In training, we perform full-parameter fine-tuning for 2 epochs during the SFT, followed by about 1 epoch (7000 steps) of LoRA-based[31] fine-tuning in the Signal-aware DPO. Optimization is carried out using the Adam [32] optimizer with a cosine annealing learning rate schedule, where the initial learning rates are set to 1.0e-5 for SFT and 5.0e-6 for Signal-aware DPO. The training configuration employs a global batch size of 64 and a maximum sequence length of 2048 tokens.

Table 1: Main Results on VerilogEval1.0 and VerilogEval2.0 Benchmarks

_		Size	VerilogEval1.0 (%)						VerilogEval2.0 (%)			
Type	Model			Machine			Human		Speci	fication	Com	pletion
			Pass@1	Pass@5	Pass@10	Pass@1	Pass@5	Pass@10	T=0	T=0.8	T=0	T=0.8
Foundation General Model	GPT-3.5 GPT-4 GPT-4o Deepseek v3	- - - 671B	46.7 60.0 67.7 <b>77.6</b>	69.1 70.6 75.5 <b>86.2</b>	74.1 73.5 77.2 <b>87.4</b>	26.7 43.5 60.1 <b>70.7</b>	45.8 55.8 71.4 <b>77.4</b>	51.7 58.9 74.5 <b>78.8</b>	32.0 62.5 <b>68.8</b>	31.7 61.4 <b>66.9</b>	42.3 59.0 <b>68.0</b>	41.6 56.1 <b>66.1</b>
General Code Model	Deepseek Coder CodeQwen1.5 Qwen2.5 Coder Instruct	6.7B 7B 7B	52.2 46.5 50.1	55.4 54.9 <b>66.5</b>	56.8 56.4 <b>70.9</b>	30.2 22.5 22.9	33.9 26.1 <b>36.0</b>	34.9 28.0 <b>39.5</b>	21.7 1.9 23.0	19.5 15.0 <b>22.0</b>	25.0 21.8 <b>30.1</b>	<b>29.3</b> 17.9 25.6
Verilog-Specific Model	RTLCoder CodeV(Qwen1.5) CodeV(Qwen2.5) Origen VeriSeek VeriPrefer QiMeng-SALV (Ours)	6.7B 7B 7B 7B 6.7B 7B 7B	61.2 77.6 77.3 74.1 61.6 72.7 <b>81.4</b>	76.5 88.2 87.9 82.4 76.9 85.8 <b>88.6</b>	81.8 90.7 90.1 85.7 81.7	41.6 52.7 57.9 54.4 30.5 49.7 <b>60.4</b>	50.1 62.5 66.7 60.1 43.4 62.3 <b>68.6</b>	53.4 67.3 69.7 64.2 49.2	36.8 7.7 44.8 49.3 28.8 - <b>57.1</b>	30.9 8.4 37.4 46.8 15.9	35.9 50.0 58.3 49.3 49.3 55.7 <b>62.2</b>	31.5 45.7 49.9 47.2 43.1 51.3 <b>58.8</b>

**Metric** Following previous work [11], the pass@k metric is employed to evaluate model performance, estimating the probability that at least one correct solution is generated within k independent attempts for each problem:

$$pass@k := \mathbb{E}_{problems} \left[ \frac{1 - \binom{n-c}{k}}{\binom{n}{k}} \right]. \tag{3}$$

Here,  $n \ge k$  denotes the total number of independent solution attempts per problem instance, while c corresponds to the count of functionally correct solutions among these trials.

**Benchmark** We conduct comprehensive evaluations on the VerilogEval (including VerilogEval1.0 [9] and VerilogEval2.0 [10]) and RTLLM benchmarks (including RTLLM v1.1 [11] and RTLLM v2.0 [12]).

VerilogEval1.0 is divided into two subtasks: Machine and Human. The Machine subset contains 143 Verilog design problems, with each prompt generated by LLMs; the Human subset comprises 156 Verilog design problems, with each prompt manually crafted. VerilogEval2.0 revisit some limitaions of VerilogEval1.0 [9] and support specification-to-RTL tasks in addition to the original code completion task. Since VerilogEval-machine can be overly descriptive compared to real-world code generation problem, VerilogEval2.0 only evaluate models against VerilogEval-human to highlight the most useful LLM evaluation.

RTLLM v1.1 [11] incorporates 29 distinct Verilog code generation tasks categorized into Arithmetic and Logic domains. Each task specification provides complete interface definitions (module ports) along with detailed functional requirements. RTLLM v2.0 expands from the 29 Verilog designs in RTLLM v1.1 to 50, covering four categories of tasks: Arithmetic, Control, Memory, and Miscellaneous.

Our experimental setup employs a sampling configuration of n=20 generations per prompt. We evaluate the model generated module code using pass@1, pass@5, and pass@10 metrics. The pass@1 metric specifically quantifies solution accuracy on demonstrably solvable problems, reflecting the model's consistency and stability in generating correct implementations. In contrast, pass@10 assesses the model's overall problem-solving capacity by measuring its ability to produce at least one valid solution within twenty attempts, thereby characterizing the breadth of its code generation capability.

Following previous practices [16], we conduct tests at temperature settings of 0.2, 0.5, and 0.8, and report the highest results on VerilogEval1.0 and RTLLM benchmark. According to the VerilogEval2.0 paper [10], VerilogEval2.0 only reports pass@1 results under low-temperature settings (T=0.0, top\_p=0.01, n=1) and high-temperature settings (T=0.8, top\_p=0.95, n=20) using nucleus sampling [33]. Notably, pass@5 and pass@10 scores are excluded from the evaluation.

**Baseline Methods** In our experimental evaluation, we conduct a comprehensive comparison between our proposed method and several baseline approaches, categorized into three groups: 1) General-purpose Foundation Models: GPT-3.5, GPT-40, GPT-4 [34], and DeepSeek-v3 [13], which demonstrate broad capabilities across diverse domains. 2) General Code Models: CodeQwen1.5 [35], Qwen2.5 Coder Instruct [30] and Deepseek Coder [36], possessing strong coding proficiency but

Table 2: Main Results on RTLLM v1.1 and RTLLM v2.0 Benchmarks

_		Size	RTLLM v1.1 (%)			RTLLM v2.0 (%)			
Type	Model		Pass@1	Pass@5	Pass@10	Pass@1	Pass@5	Pass@10	
Foundation	GPT-3.5	-	28.3	36.9	41.4	34.4	49.8	52.1	
General Model	GPT-4o	-	41.7	65.9	-	56.5	70.3	75.2	
General Model	Deepseek v3	671B	62.0	72.0	72.4	59.1	71.5	73.3	
General Code	Deepseek Coder	6.7B	23.1	29.3	34.5	26.5	36.3	42.7	
Model	CodeQwen1.5	7B	28.8	38.8	43.3	25.8	29.0	-	
Model	Qwen2.5 Coder Instruct	7B	30.1	49.2	55.9	33.2	52.5	57.7	
	RTLCoder	6.7B	35.8	40.3	43.1	43.5	48.1	-	
	CodeV(Qwen1.5)	7B	36.6	53.3	61.3	48.1	56.9	-	
Verilog-Specific	CodeV(Qwen2.5)	7B	39.3	63.5	74.2	41.0	60.1	68.1	
Model	Origen	7B	50.6	68.3	74.3	50.9	60.9	64.0	
	VeriSeek	6.7B	29.3	47.1	53.1	31.9	54.2	52.0	
	VeriPrefer	7B	53.2	67.7	-	52.4	66.4	-	
	QiMeng-SALV (Ours)	7B	62.6	75.1	81.1	62.0	71.7	76.0	

lacking Verilog-specific optimization. 3) Domain-Specialized Verilog Models: Including RTL Coder [16], CodeV [4], Origen [17], VeriSeek [7] and VeriPrefer [28] which employ module-level rewards in reinforcement learning.

#### 4.2 Main Results

Table1 and Table2 comprehensively compare the performance of our QiMeng-SALV method against baseline models across the VerilogEval [9, 10] and RTLLM [11, 12] benchmarks. To establish a more rigorous evaluation framework, we re-implemented CodeV using Qwen2.5 Coder Instruct as its foundation model, replacing its original CodeQwen1.5 which demonstrated inferior performance. This architectural alignment with our base model eliminates potential confounding factors arising from fundamental model discrepancies.

Comparison results in Table 1 and Table 2 show that QiMeng-SALV establishes new state-of-the-art results across both benchmarks in the open-source domain. As shown in Table 1, in VerilogEval, QiMeng-SALV achieves leading performance among open-source solutions on both specification understanding and code completion tasks—attaining 81.4 pass@1 on Machine and 60.4 pass@1 on Human in VerilogEval 1.0, as well as 57.1 (T = 0) on Specification and 62.2 (T = 0) on Completion in VerilogEval2.0. Its performance is comparable to GPT-4o on the VerilogEval1.0 Human benchmark and even surpasses Deepseek v3 and GPT-4o on the VerilogEval1.0 Machine.

As shown in Table 2, QiMeng-SALV achieves a remarkable 62.6% functional pass@1 accuracy on the RTLLM v1.1 benchmark and 62.0% on RTLLM v2.0 with merely 7B parameters, significantly exceeding all existing open-source alternatives and rivaling the performance of DeepSeek-v3, a 671B parameter model. Impressively, its functional pass@10 accuracy reaches 81.1% on RTLLM v1.1, surpassing DeepSeek-v3's 72.4%.

Remarkably, when compared to CodeV (Qwen2.5), trained on identical datasets with the same base model, QiMeng-SALV shows substantial improvements: about 59.7% improvement pass@1 accuracy on RTLLM v1.1, and 27.4%(T=0)/50.2%(T=0.8) improvement on VerilogEval2.0's specification subtasks. These findings underscore that leveraging RL to learn from correct signal-level implementations yields substantially stronger performance compared to models trained exclusively via supervised fine-tuning.

Furthermore, when compared with VeriPrefer, which employs module-level rewards in reinforcement learning, our method achieves substantial improvements across all evaluation metrics. This highlights the advantage of fine-grained supervision: signal-level rewards allow the model to effectively learn from partially correct samples that are often disregarded under coarser reward schemes. Consequently, signal-level reinforcement learning enables the generation of more robust and correct Verilog code.

# 4.3 Ablation Study

In this section, we present a systematic ablation study to evaluate our proposed QiMeng-SALV. Given that RTLLM contains more complex problems than VerilogEval, making it more suitable for revealing




Figure 3: (a) Pass rate performance comparison between different training stages across different tasks on RTLLM v1.1 benchmark. (b) Signal-aware DPO training datasets scaling on RTLLM benchmark.

nuanced performance differences among models, we focus our ablation analysis on the RTLLM benchmark. Our investigation spans two key dimensions: 1) configuration variations of Signal-aware DPO training, and 2) different training stage implementations of QiMeng-SALV.

**Ablation of configuration variations of signal-aware DPO** Unlike standard DPO, which is limited to processing complete correct modules and must consider all response tokens, our signal-aware DPO can not only handle complete correct modules but also extract functional reward in partial correct modules containing correct signal implementations. Therefore, we design three different options to evaluate the efficiency of our method: 1) Complete correct datasets (263k): all of preferred sample in DPO training data are complete correct modules. 2). Partial correct datasets (110k): all of preferred sample in DPO training data are partial correct modules containing both correct and incorrect signals. 3). Filter incorrect signals: determines whether erroneous signals in preferred samples participate in DPO computation. As evidenced by the results in Table 3, standard DPO trained exclusively on complete correct datasets achieves superior performance compared to training on partially correct datasets. Interestingly, combining both complete and partially correct datasets during training leads to degraded performance relative to using only complete correct datasets. This degradation likely stems from the presence of incorrect signal implementations in partially correct module, which introduce noise and hinder the model's ability to discern meaningful reward signals. By contrast, our proposed signal-aware DPO method, which actively filters out incorrect signal implementations, substantially improves performance in both scenarios: training on partially correct datasets alone and training on mixed datasets. These findings highlight two key advantages of signal-aware DPO: (1) its ability to leverage a broader spectrum of training data, and (2) its robustness against noisy signal implementations through effective filtering.

Ablation of different training stage implementations To systematically analyze the performance gains of the model across different training phases, we conducted comprehensive evaluations at each stage. The results, as presented in the Table 4, indicate that SFT yields a substantial improvement of approximately 15 percentage points in all three accuracy dimensions. Subsequent signal-aware DPO training delivers additional gains of 10-14 percentage points. Figure 3 (a) illustrates the pass rate of the model on the RTLLM v1.1 benchmark tasks across different training stages. The base code model not only answered fewer questions correctly but also exhibited low accuracy for the questions it could answer. After the SFT training stage, the model could correctly answer a broader range of questions. Following signal-aware DPO training, the model further improved its accuracy on questions it could answer while also correctly answering previously unsolvable questions. These findings substantiate that signal-aware DPO effectively utilizes functional correctness feedback signals to drive additional performance improvements beyond what is achieved through SFT alone.

**Scaling on Signal-aware DPO training datasets** To examine the influence of training dataset size on Signal-aware DPO's efficacy, we conducted experiments by training models with preference datasets of different scales. Figure 3 (b) reveals a consistent upward trend in pass@1, pass@5, and pass@10 metrics as the dataset size grows. These results indicate that incorporating more meaningful

Table 3: Ablation Different Settings on RTLLM Benchmark

Complete	Partial	Filter	RT	TLLM v1.1	(%)	RTLLM v2.0 (%)			
Correct Correct Datasets Datasets		Incorrect Signals	Pass@1	Pass@5	Pass@10	Pass@1	Pass@5	Pass@10	
$\checkmark$			57.4	74.0	78.3	57.1	70.4	73.4	
	$\checkmark$		56.0	67.9	71.6	54.7	66.0	69.2	
$\checkmark$	$\checkmark$		55.3	73.1	77.1	55.9	69.8	73.3	
	$\checkmark$	$\checkmark$	56.6	71.0	76.4	54.1	68.4	73.1	
$\checkmark$	$\checkmark$	$\checkmark$	62.6	75.1	81.1	62.0	71.7	76.0	

Table 4: Ablation Different Training Stage on **RTLLM** Benchmark

Model	RT	TLLM v1.1	(%)	RTLLM v2.0 (%)			
1,10001	Pass@1	Pass@5	Pass@10	Pass@1	Pass@5	Pass@10	
Base Model(Qwen2.5)	30.1	49.2	55.9	33.2	52.5	57.7	
+ SFT	48.3	65.2	69.8	50.4	64.3	68.7	
+ Signal-aware DPO	62.6	75.1	81.1	62.0	71.7	76.0	

functional rewards during reinforcement learning training significantly contributes to improved model performance.

# **5** Runtime Analysis

We measured the average runtime of simulation and AST parsing over all Verilog samples in training data. On average, simulation takes 0.0391 seconds and AST parsing takes 0.0957 seconds per sample. By utilizing 60 CPU cores for parallel processing, we reduce the per-sample average simulation time to 0.65 milliseconds and AST parsing time to 1.59 milliseconds. These results show that the overhead of simulation and AST parsing is negligible in practice.

# 6 Conclusion

In this paper, we propose QiMeng-SALV, a novel reinforcement learning method for training Verilog code generation, which learns correct signal implementations from erroneous modules, thereby increasing functional correctness rewards and improving capability. Experiments show that our model achieves state-of-the-art performance on VerilogEval and RTLLM, significantly outperforming other open-source models. On RTLLM v1.1, it reaches pass@1 62.6% and pass@5 75.1%, matching the performance of DeepSeek-v3 (671B parameters) with only 7B parameters. QiMeng-SALV shifts the training paradigm for Verilog code generation from the module level to the signal level.

## References

- [1] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [2] Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. Li, Y. Wu, *et al.*, "Deepseekmath: Pushing the limits of mathematical reasoning in open language models," *arXiv* preprint arXiv:2402.03300, 2024.
- [3] R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, and C. Finn, "Direct preference optimization: Your language model is secretly a reward model," *Advances in Neural Information Processing Systems*, vol. 36, pp. 53728–53741, 2023.
- [4] Y. Zhao, D. Huang, C. Li, P. Jin, Z. Nan, T. Ma, L. Qi, Y. Pan, Z. Zhang, R. Zhang, *et al.*, "Codev: Empowering llms for verilog generation through multi-level summarization," *arXiv* preprint arXiv:2407.10424, 2024.
- [5] S. Thakur, B. Ahmad, H. Pearce, B. Tan, B. Dolan-Gavitt, R. Karri, and S. Garg, "Verigen: A large language model for verilog code generation," *ACM Transactions on Design Automation of Electronic Systems*, vol. 29, no. 3, pp. 1–31, 2024.
- [6] H. Pearce, B. Tan, and R. Karri, "Dave: Deriving automatically verilog from english," in *Proceedings of the 2020 ACM/IEEE Workshop on Machine Learning for CAD*, pp. 27–32, 2020.
- [7] N. Wang, B. Yao, J. Zhou, X. Wang, Z. Jiang, and N. Guan, "Large language model for verilog generation with golden code feedback," *arXiv preprint arXiv:2407.18271*, 2024.
- [8] P. Wu, N. Guo, X. Xiao, W. Li, X. Ye, and D. Fan, "Itertl: An iterative framework for fine-tuning llms for rtl code generation," *arXiv preprint arXiv:2407.12022*, 2024.
- [9] M. Liu, N. Pinckney, B. Khailany, and H. Ren, "Verilogeval: Evaluating large language models for verilog code generation," in 2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD), pp. 1–8, IEEE, 2023.
- [10] N. Pinckney, C. Batten, M. Liu, H. Ren, and B. Khailany, "Revisiting verilogeval: Newer llms, in-context learning, and specification-to-rtl tasks," *arXiv preprint arXiv:2408.11053*, 2024.
- [11] Y. Lu, S. Liu, Q. Zhang, and Z. Xie, "Rtllm: An open-source benchmark for design rtl generation with large language model," in 2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC), pp. 722–727, IEEE, 2024.
- [12] S. Liu, Y. Lu, W. Fang, M. Li, and Z. Xie, "OpenIlm-rtl: Open dataset and benchmark for Ilm-aided design rtl generation," in *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, pp. 1–9, 2024.
- [13] A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, C. Zhao, C. Deng, C. Zhang, C. Ruan, et al., "Deepseek-v3 technical report," arXiv preprint arXiv:2412.19437, 2024.
- [14] Z. Pei, H. Zhen, M. Yuan, Y. Huang, and B. Yu, "Betterv: Controlled verilog generation with discriminative guidance," in *ICML*, OpenReview.net, 2024.
- [15] M. Gao, J. Zhao, Z. Lin, W. Ding, X. Hou, Y. Feng, C. Li, and M. Guo, "Autovcoder: A systematic framework for automated verilog code generation using llms," in 2024 IEEE 42nd International Conference on Computer Design (ICCD), pp. 162–169, IEEE, 2024.
- [16] S. Liu, W. Fang, Y. Lu, Q. Zhang, H. Zhang, and Z. Xie, "Rtlcoder: Outperforming gpt-3.5 in design rtl generation with our open-source dataset and lightweight solution," in 2024 IEEE LLM Aided Design Workshop (LAD), pp. 1–5, IEEE, 2024.
- [17] F. Cui, C. Yin, K. Zhou, Y. Xiao, G. Sun, Q. Xu, Q. Guo, Y. Liang, X. Zhang, D. Song, *et al.*, "Origen: Enhancing rtl code generation with code-to-code augmentation and self-reflection," in *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, pp. 1–9, 2024.

- [18] G. Tie, Z. Zhao, D. Song, F. Wei, R. Zhou, Y. Dai, W. Yin, Z. Yang, J. Yan, Y. Su, *et al.*, "A survey on post-training of large language models," *arXiv preprint arXiv:2503.06072*, 2025.
- [19] Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, N. DasSarma, D. Drain, S. Fort, D. Ganguli, T. Henighan, et al., "Training a helpful and harmless assistant with reinforcement learning from human feedback," arXiv preprint arXiv:2204.05862, 2022.
- [20] H. Dong, W. Xiong, D. Goyal, Y. Zhang, W. Chow, R. Pan, S. Diao, J. Zhang, K. Shum, and T. Zhang, "Raft: Reward ranked finetuning for generative foundation model alignment," *arXiv* preprint arXiv:2304.06767, 2023.
- [21] H. Yuan, Z. Yuan, C. Tan, W. Wang, S. Huang, and F. Huang, "Rrhf: Rank responses to align language models with human feedback," *Advances in Neural Information Processing Systems*, vol. 36, pp. 10935–10950, 2023.
- [22] Y. Zhao, R. Joshi, T. Liu, M. Khalman, M. Saleh, and P. J. Liu, "Slic-hf: Sequence likelihood calibration with human feedback," *arXiv preprint arXiv:2305.10425*, 2023.
- [23] T. Liu, Y. Zhao, R. Joshi, M. Khalman, M. Saleh, P. J. Liu, and J. Liu, "Statistical rejection sampling improves preference optimization," *arXiv preprint arXiv:2309.06657*, 2023.
- [24] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. D. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, *et al.*, "Evaluating large language models trained on code," *arXiv* preprint arXiv:2107.03374, 2021.
- [25] H. Le, Y. Wang, A. D. Gotmare, S. Savarese, and S. C. H. Hoi, "Coderl: Mastering code generation through pretrained models and deep reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 21314–21328, 2022.
- [26] J. Liu, Y. Zhu, K. Xiao, Q. Fu, X. Han, W. Yang, and D. Ye, "Rltf: Reinforcement learning from unit test feedback," *arXiv preprint arXiv:2307.04349*, 2023.
- [27] S. Dou, Y. Liu, H. Jia, L. Xiong, E. Zhou, W. Shen, J. Shan, C. Huang, X. Wang, X. Fan, et al., "Stepcoder: Improve code generation with reinforcement learning from compiler feedback," arXiv preprint arXiv:2402.01391, 2024.
- [28] N. Wang, B. Yao, J. Zhou, Y. Hu, X. Wang, N. Guan, and Z. Jiang, "Insights from verification: Training a verilog generation llm with reinforcement learning with testbench feedback," *arXiv* preprint arXiv:2504.15804, 2025.
- [29] C. Wolf, J. Glaser, and J. Kepler, "Yosys-a free verilog synthesis suite," in *Proceedings of the 21st Austrian Workshop on Microelectronics (Austrochip)*, vol. 97, 2013.
- [30] B. Hui, J. Yang, Z. Cui, J. Yang, D. Liu, L. Zhang, T. Liu, J. Zhang, B. Yu, K. Lu, et al., "Qwen2. 5-coder technical report," arXiv preprint arXiv:2409.12186, 2024.
- [31] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, *et al.*, "Lora: Low-rank adaptation of large language models.," *ICLR*, vol. 1, no. 2, p. 3, 2022.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [33] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi, "The curious case of neural text degeneration," *arXiv preprint arXiv:1904.09751*, 2019.
- [34] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al., "Gpt-4 technical report," arXiv preprint arXiv:2303.08774, 2023.
- [35] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang, B. Hui, L. Ji, M. Li, J. Lin, R. Lin, D. Liu, G. Liu, C. Lu, K. Lu, J. Ma, R. Men, X. Ren, X. Ren, C. Tan, S. Tan, J. Tu, P. Wang, S. Wang, W. Wang, S. Wu, B. Xu, J. Xu, A. Yang, H. Yang, J. Yang, S. Yang, Y. Yao, B. Yu, H. Yuan, Z. Yuan, J. Zhang, X. Zhang, Y. Zhang, Z. Zhang, C. Zhou, J. Zhou, X. Zhou, and T. Zhu, "Qwen technical report," arXiv preprint arXiv:2309.16609, 2023.

[36] D. Guo, Q. Zhu, D. Yang, Z. Xie, K. Dong, W. Zhang, G. Chen, X. Bi, Y. Wu, Y. Li, *et al.*, "Deepseek-coder: When the large language model meets programming—the rise of code intelligence," *arXiv preprint arXiv:2401.14196*, 2024.

# **Supplementary Material**

# **A** Experiment Compute resources

For the SFT phase, we execute full-parameter optimization across 2 training epochs utilizing a cluster of 4 NVIDIA A100-80GB SMX GPUs, with the complete training procedure consuming approximately 20 hours. Subsequently, during the Signal-aware DPO stage, we implement LoRA-based fine-tuning over 7000 steps on an array of 8 NVIDIA A100-40GB GPUs, resulting in a total training duration of roughly 15 hours.

# **B** Societal Impacts

Our research delivers significant positive societal impact by substantially improving the functional correctness of automatically generated Verilog code. This advancement enhances productivity in circuit design workflows, accelerates development cycles, and provides industry practitioners with a reliable assistive tool. However, we acknowledge potential negative implications in academic settings. The model's capabilities could be misused by students to complete Verilog programming assignments or examinations, potentially facilitating academic dishonesty. We emphasize the importance of developing appropriate usage guidelines and detection mechanisms to mitigate such risks while preserving the technology's beneficial applications.

## **C** Limitation

QiMeng-SALV determines the functional correctness of the generated module by producing random input signals and comparing the output signals between the generated module and the reference module. If the reference module in the training dataset itself is incorrect, it may lead to erroneous judgments of functional correctness, necessitating a relatively high-quality training dataset. In this paper, we do not discuss how to obtain a high-quality dataset, as this is beyond our scope. In our experiment, we use the CodeV dataset, a high-quality dataset obtained by crawling all publicly available Verilog module code on GitHub. We also perform data cleaning to ensure the reference modules are syntactically correct and compilable.

# **NeurIPS Paper Checklist**

#### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

#### Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

#### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: In Appendix.

#### Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

# 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

#### Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

# 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We will open-source our model on huggingface.

#### Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will open-source our datasets and code on github.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/ public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https: //nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- · The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: In 4.1.

#### Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

# 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Our experiments rely heavily on large language model (LLM) training, and repeating each setting multiple times would significantly increase the cost.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error
  of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In Appendix.

## Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The research strictly adheres to the NeurIPS Code of Ethics. All experiments were conducted with integrity, transparency, and respect for human and environmental welfare.

#### Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In Appendix.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

# 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

#### Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

# 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original paper.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

• If this information is not available online, the authors are encouraged to reach out to the asset's creators.

#### 13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

#### Guidelines:

- The answer NA means that the paper does not release new assets.
- · Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

# 14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

# 15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- · For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

# 16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This research does not involve LLMs as any important, original, or non-standard components.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.