

AgenticTCAD: A LLM-based Multi-Agent Framework for Automated TCAD Code Generation and Device Optimization

Guangxi Fan^{1,2}, Tianliang Ma^{1,2}, Xuguang Sun^{1,2}, Xun Wang^{1,2}, Kain Lu Low^{3*}, and Leilai Shao^{1,2*}

¹State Key Laboratory of Micro-nano Engineering Science

²Micro-nano Engineering Sciences Research Center

School of Mechanical Engineering, Shanghai Jiao Tong University, Shanghai, China

³Xi'an Jiaotong-Liverpool University, Suzhou, China

Abstract— With the continued scaling of advanced technology nodes, the design–technology co-optimization (DTCO) paradigm has become increasingly critical, rendering efficient device design and optimization essential. In the domain of TCAD simulation, however, the scarcity of open-source resources hinders language models from generating valid TCAD code. To overcome this limitation, we construct an open-source TCAD dataset curated by experts and fine-tune a domain-specific model for TCAD code generation. Building on this foundation, we propose AgenticTCAD, a natural language–driven multi-agent framework that enables end-to-end automated device design and optimization. Validation on a 2 nm nanosheet FET (NS-FET) design shows that AgenticTCAD achieves the International Roadmap for Devices and Systems (IRDS)-2024 device specifications within 4.2 hours, whereas human experts required 7.1 days with commercial tools.

Index Terms—Code Generation, TCAD Simulation, Multi-agent Flow, Device Optimization

I. INTRODUCTION

Over the past decades, the semiconductor industry has advanced by following Moore’s law, which has consistently increased transistor density, performance, and energy efficiency [1]. As CMOS technology approaches its physical and economic limits, however, further scaling is hindered by challenges such as rising leakage currents, variability, and manufacturing costs, making it difficult to sustain historical improvements. To address these issues, the concept of design technology co-optimization (DTCO) has emerged, highlighting the need for joint optimization of process technology and circuit design to continue driving progress in integrated circuits and systems [2].

Within the DTCO paradigm, the technology-oriented dimension plays a central role, encompassing material exploration [3], device architecture innovation [4], and standard-cell layout optimization [5]. Among these tasks, device design and optimization are especially critical, as transistor characteristics directly impact overall system performance. Consequently, Technology Computer-Aided Design (TCAD) simulation [6] is indispensable for guiding technology scaling and supporting effective DTCO.

TCAD simulation involves highly complex physical mechanisms, such like carrier transport, quantum confinement,

and trap effects [7]. Manual device optimization under these conditions is extremely challenging, as the large design space coupled with intricate physics requires significant expertise and effort. In addition, TCAD tools require domain-specific coding languages to describe device structures, meshing, and simulation setups. Since semiconductor device research is often proprietary and lacks open-source references, writing such codes poses a significant barrier for researchers and further limits efficient adoption.

Large language models (LLMs) have recently demonstrated considerable potential in the design domain of DTCO [8]. They have been employed for code generation, enabling the automatic synthesis of design scripts and hardware descriptions directly from natural language inputs [9]. Furthermore, LLMs have been investigated for automated circuit design, offering a pathway from high-level specifications to transistor-level design with reduced human involvement [10].

The technology dimension of DTCO critically depends on device design and optimization. Sentaurus Structure Editor (SDE) and Sentaurus Device (SDevice) are widely used for device construction and electrical simulation, but their domain-specific and proprietary scripts leave the community without publicly accessible datasets for TCAD code generation. Current research efforts have so far focused mainly on SDE script generation. Early attempts explored nanowire device structure generation [11], but these did not include mesh construction, making direct electrical simulation infeasible. Subsequent work has investigated end-to-end code generation through fine-tuned LLM [12], yet this required translating SDE scripts into Python descriptions beforehand, which introduces additional complexity and has so far been limited to two-dimensional structures. Other approaches have employed agent-based methods to map layout descriptions into SDE scripts [13], but these often depend on case-specific templates for in-context learning, reflecting the difficulty commercial LLMs still face in handling TCAD-specific syntax.

These limitations highlight the absence of a unified framework that can transform natural-language specifications into both SDE and SDevice codes for end-to-end TCAD simulation, while also enabling iterative device parameter optimization with high physical interpretability. To address this gap, we propose **AgenticTCAD**, a domain-specific multi-agent LLM-based framework for automated device simulation and optimization, as illustrate in Fig. 1. Our contributions are as follows:

- We release the first open-source comprehensive TCAD dataset containing expert-crafted SDE and SDevice scripts, further expanded through an LLM-based augmentation pipeline, to support community research in this field. The dataset includes a wide spectrum of device

*Corresponding author: Leilai Shao (leilaishao@sjtu.edu.cn); Kain Lu Low (kainlu.low@xjtu.edu.cn). This work was supported by the National Key Research and Development Program of China: Design Technology Co-Optimization Methodology (2023YFB4402700), Shanghai Jiao Tong University AI for Engineering Initiative (26X010100040), Shanghai Municipal Science and Technology Major Project in collaboration with Shanghai Artificial Intelligence Laboratory and Xi’an Jiaotong-Liverpool University Research Development Fund under Grant RDF-24-01-107. The authors would like to thank Prof. Libin Liu for the valuable discussions.

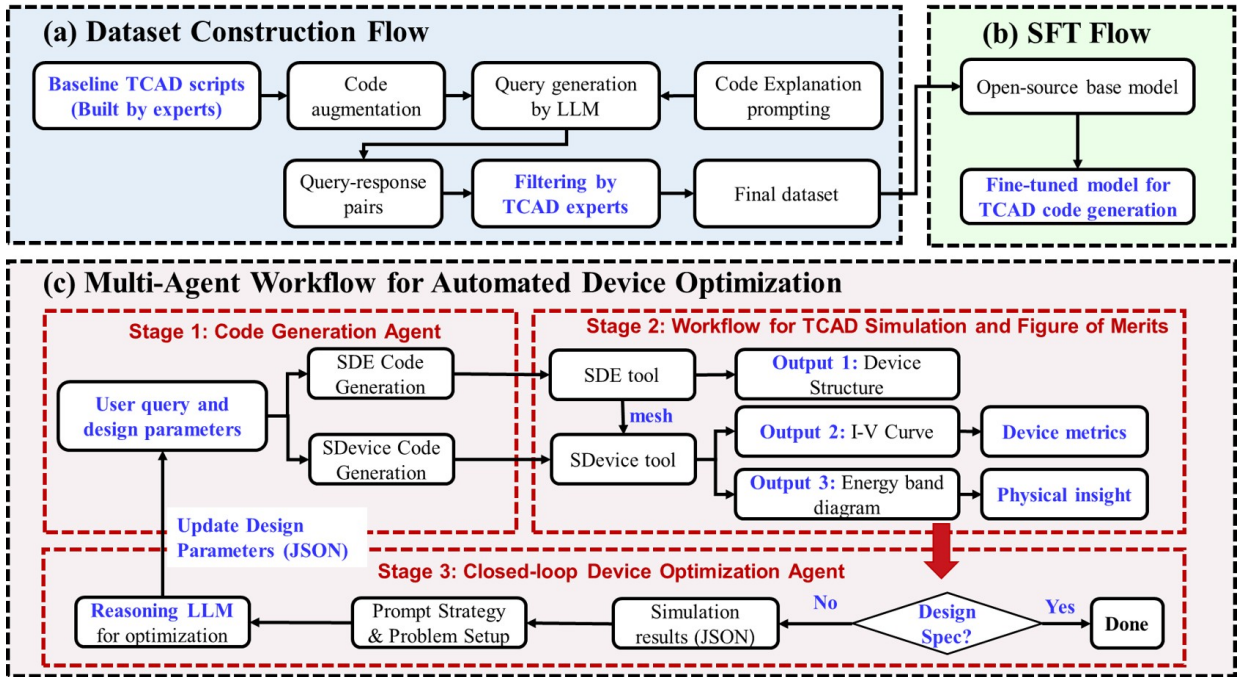


Fig. 1. Overall flow of TCADAgent. (a) Dataset construction flow; (b) Supervised fine-tuning (SFT) flow; (c) Multi-Agent flow for End-to-end device TCAD simulation and optimization based on natural language.

structures in both two- and three-dimensional forms, various operating mechanisms including conventional MOSFETs as well as emerging junctionless and tunneling transistors, and multiple material systems including silicon and III-V semiconductors.

- We fine-tune an open-source LLM to generate complete SDE and SDevice codes directly from natural language. The SDE codes specify geometry, doping profiles, contact definitions, and meshing strategies, while the SDevice codes define boundary conditions, physical models, and solver configurations, enabling direct end-to-end simulation.
- We introduce **AgenticTCAD**, a multi-agent framework that integrates our fine-tuned TCAD code generation LLM with a commercial reasoning model, with each LLM agent using a tailored prompting strategy to realize automated device simulation and optimization. To demonstrate its effectiveness, we apply the framework to the optimization of design parameters for nanosheet field-effect transistors (NS-FETs) at advanced technology nodes. The framework also automatically extracts physical information from simulation results, such as energy band diagrams, to enhance the interpretability of the optimization flow.

The remainder of this paper is organized as follows: Section II reviews related work on TCAD simulation, LLM-based code generation, and agent-based methods in EDA. Section III presents the proposed **AgenticTCAD** framework and its design details. Section IV reports experimental results and Section V concludes the paper. The code and dataset of **AgenticTCAD** are available at github.

II. PRELIMINARIES AND RELATED WORKS

A. TCAD Simulation

TCAD simulation has been widely applied in semiconductor device research and development to predict device behavior

before fabrication. It supports a variety of material systems and device architectures, including advanced silicon-based technologies [4], III-V semiconductors [14], and emerging devices such as carbon nanotube (CNT) [15]. In addition, TCAD enables comprehensive multiphysics analysis, capturing effects such as defects [7] and self-heating effects [16]. Owing to these capabilities, TCAD plays a crucial role in accelerating early-stage device design and optimization.

B. EDA Code Generation with LLM

Code generation is a key application of large language models (LLMs) and has been widely explored in mainstream programming language [17]. Within DTCO's design stage, LLMs have been applied to realize RTL code generation [9], [18], [19], debugging [20], [21] and optimization [22]. In the back-end stage, researchers are investigating how LLMs can automatically construct physical design scripts, including flows such as placement and routing, clock-tree synthesis [23]. In the technology domain, there have been initial attempts to use LLMs for generating TCAD SDE input code [11]–[13] and SDevice decks to support subsequent device simulations, though this line of research remains relatively limited. Our work aims to build a complete TCAD SDE/SDevice code generation flow and release an open-source dataset to foster further investigation in this direction. In summary, LLM-based code generation holds great potential to accelerate the EDA workflow across the entire DTCO iteration process.

C. Agent Flow in EDA

Agents built upon LLMs provide a flexible paradigm for decomposing complex design tasks into modular, cooperative subtasks. Recent studies have applied LLM-based agents to analog design, such as topology exploration of operational amplifiers [10] and hybrid frameworks that combine LLMs with Bayesian optimization for circuit design [24]. LLM-based agents have further been investigated in digital design,

including recent demonstrations on CPU development with successful tape-out results [25]. In addition, LLM agents have been explored for generating and optimizing circuit layouts [26], [27], aiming to improve downstream manufacturing quality and design efficiency. Overall, LLM-driven agent flows have shown strong potential to enhance automation and efficiency in EDA.

III. METHODOLOGY

A. Dataset Construction

To address the absence of publicly available datasets tailored for LLM applications in the design–technology co-optimization (DTCO) domain, particularly in device design, we develop a semi-automated dataset construction flow, as illustrated in Fig. 1(a). In the proposed dataset construction flow, TCAD experts first identify a diverse set of representative device architectures. The selection covers both conventional planar transistors and advanced three-dimensional structures such as FinFETs, nanosheets, and nanowires, as well as their array configurations. To enhance material diversity, the dataset also includes devices based on compound semiconductors such as SiGe and InGaAs in addition to traditional silicon. For each device type, experts develop baseline SDE scripts that specify geometry, doping profiles, and contact configurations. Doping distributions are modeled using both uniform concentrations and physically realistic Gaussian profiles derived from ion implantation processes. To ensure simulation robustness and scalability, the meshing strategy is parameterized to adapt to varying device dimensions, with both axis-aligned and offset meshes incorporated for accurate representation of complex structures. Complementary SDevice scripts are further generated to define boundary conditions, physical models, and biasing schemes. These configurations span traditional drift–diffusion models as well as advanced tunneling-based models, enabling simulations to produce a wide range of electrical characteristics, including transfer characteristics ($I_d - V_g$), output characteristics ($I_d - V_d$), and capacitance–voltage ($C - V$) curves. To further scale the corpus, the baseline scripts are systematically expanded through parameter variation in Sentaurus Workbench, resulting in thousands of executable simulation scripts.

However, these scripts alone are insufficient for direct use in LLM training, as effective SFT requires paired natural-language queries with corresponding code responses, and manually generating such queries at scale is impractical. To transform these scripts into natural language–to–code training data, we design specialized prompting strategies that leverage a commercial LLM to generate diverse user queries consistent with TCAD semantics and design intent, as illustrated in Fig. 2(a). To support this process, domain experts modularize the baseline codes, annotate them with descriptive comments on structural and parametric details, and package each script together with its metadata into a structured JSON format. In the prompting strategy, the LLM is explicitly designated as an expert in TCAD simulation and prompt engineering, and is provided with precise instructions regarding the information it receives. The prompt description specifies the global context of the target device or study, the type of simulation such as SDE for structure creation or SDevice for $I_d - V_g$, $I_d - V_d$ and $C - V$ analyses, as well as the form of metadata, which may be given as default parameters, symbolic templates, or explicit

numerical values. The task is formulated to require the generation of exactly one concise and natural-sounding query that reflects typical engineering practice, while adhering to strict output format constraints to ensure clarity and consistency. In addition, representative examples are included, and controlled variation in phrasing style is encouraged, thereby capturing the diversity with which TCAD engineers naturally articulate their requests.

This representation enables the LLM to generate high-quality queries that are consistent with the underlying code. Finally, domain experts conduct manual verification to filter and refine the augmented data, ensuring both physical plausibility and the feasibility of subsequent simulations, resulting in the finalized dataset.

B. Domain-Specific LLM Fine-Tuning

After constructing the TCAD-specific dataset, we further adapt a general-purpose LLM to the task of code generation through SFT [28]. In this stage, the pretrained model is exposed to a large number of carefully designed input–output pairs so that it can gradually adjust its parameters toward the characteristics of TCAD-related tasks. The fine-tuning objective is formulated as a next-token prediction problem, where the model is optimized to maximize the likelihood of generating the correct target sequence conditioned on the given input. Concretely, each training example is represented as a pair: the input corresponds to a natural-language query derived from the prompt templates in Fig. 2(a), while the output is the associated SDE or SDevice script that implements the requested functionality. This design explicitly bridges the gap between high-level textual descriptions and low-level simulation code.

During training, we adopt a causal language modeling objective in which all tokens belonging to the prompt are masked out during loss computation. By doing so, the model does not simply learn to copy the input, but instead concentrates on capturing the correct dependencies required to produce the desired output sequence. This setup encourages the model to internalize domain-specific conventions such as the syntax, parameter formatting, and hierarchical structure of TCAD scripts. As a result, the fine-tuned model becomes better aligned with the practical requirements of semiconductor device simulation tasks, enabling agile and efficient transistor design.

C. Multi-Agent Workflow for Automated Device Optimization

Building on our domain-specific TCAD code generation LLM, we design **AgenticTCAD**, a multi-agent workflow that integrates the fine-tuned model with a commercial reasoning LLM, as illustrated in Fig. 1(c). Each agent is equipped with a tailored prompting strategy to handle specialized subtasks, thereby enabling a fully automated loop that covers device code generation, TCAD simulation execution, post-processing, and iterative optimization. To validate the proposed framework, we apply **AgenticTCAD** to the optimization of design parameters for nanosheet field-effect transistors (NS-FETs) at advanced technology nodes. This workflow effectively connects natural language design specifications to the automated optimization of device performance, providing a practical framework for accelerating TCAD-driven device exploration.

In the stage 1, as shown in Fig. 1(c), **AgenticTCAD** employs the fine-tuned code generation LLM to convert user-provided design parameters, expressed in JSON format, into

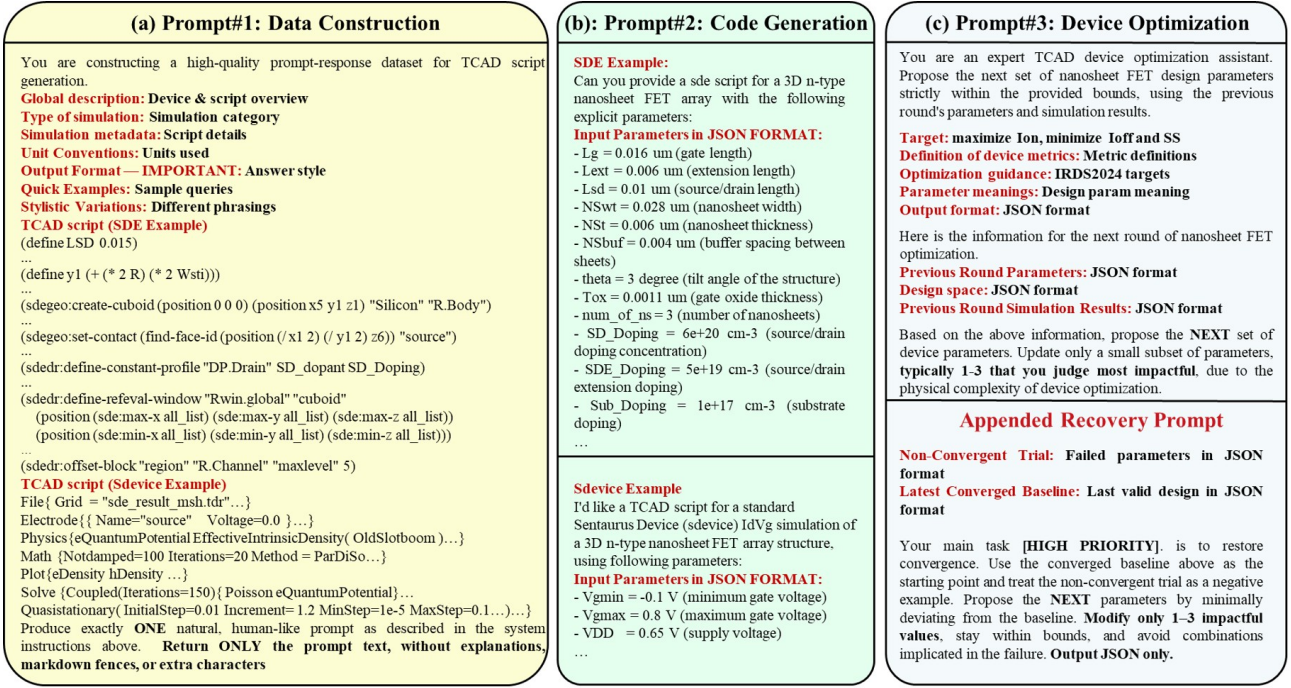


Fig. 2. Prompt templates and examples used in our framework. (a) Data Construction, (b) Code Generation and (c) Device Optimization.

simulation-ready inputs. As illustrated in Fig. 2(b), these parameters are incorporated into tailored prompting strategies, enabling the TCAD code generation LLM to produce the NS-FET structure description through SDE, which specifies device geometry, doping profiles, contact definition, and meshing. In parallel, the LLM generates the corresponding SDevice script, which defines the physical models, boundary conditions, and solver configurations required for electrical simulations. This stage enables the agent to autonomously design NS-FET architectures with user-specified parameters and generate their corresponding simulation scripts, thereby establishing a consistent and automated foundation for subsequent simulation and optimization.

The workflow in stage 2 executes the TCAD simulation and subsequent post-processing. The SDE script is first executed to construct the NS-FET geometry and corresponding device meshing. The resulting structure is then passed to the SDevice solver, which performs electrical simulations based on the associated physical models, boundary conditions, and solver settings. According to the International Roadmap for Devices and Systems (IRDS)-2024 logic device requirements [29], three key performance metrics are selected, and the post-processing module automatically extracts them from the simulated $I - V$ characteristics. The on-state current (I_{on}) is the drive current that determines both the switching speed of logic gates and the drive strength of the device. The off-state current (I_{off}) corresponds to the leakage current in the closed state and directly impacts static power consumption and overall energy efficiency. The subthreshold swing (SS) quantifies the sharpness of the switching transition and governs the trade-off between performance and power efficiency. These metrics, together with the complete $I - V$ characteristics, are stored in a structured JSON format to enable seamless integration with the subsequent reasoning stage. In addition, the post-processing module also extracts the energy band diagrams in the channel region under different bias conditions. These

diagrams include the conduction band (E_C), the valence band (E_V), the electron quasi-Fermi level (E_{Fn}), and the hole quasi-Fermi level (E_{Fp}). Such information provides physical insight into carrier transport and enhances the interpretability of the optimization process.

The results of the first two stages complete the forward process of device simulation, while Stage 3 employs a commercial reasoning LLM as a device optimization agent to realize feedback optimization. The agent receives the simulation results, the corresponding design parameter space, and the previous round of design parameters as references. These inputs are integrated into carefully engineered prompts, as illustrated in Fig. 2(c), where the performance metrics are supplemented with detailed physical interpretations and the optimization task is explicitly specified. Based on this information, the device optimization agent is instructed to generate updated design parameters for the next round of code generation and simulation. At the same time, it evaluates whether the simulated performance meets the specifications defined by the IRDS-2024 or whether the maximum number of iterations has been reached. Since TCAD inherently involves solving nonlinear semiconductor device equations, convergence issues may arise for certain parameter sets. To address this challenge, **AgentTCAD** incorporates a recovery mechanism, as shown in Fig. 2(c). In cases where non-convergence is encountered, the system provides the agent with the most recent convergent design, the failed configuration, and an additional prompt specifically designed for handling non-convergence. This mechanism enables the reasoning agent to revise the design parameters and ensures that the optimization loop continues without interruption.

IV. EXPERIMENTS AND RESULTS

A. Experimental Setup

All experiments were conducted on a Linux-based workstation equipped with two Intel Xeon Platinum 8375C CPUs,

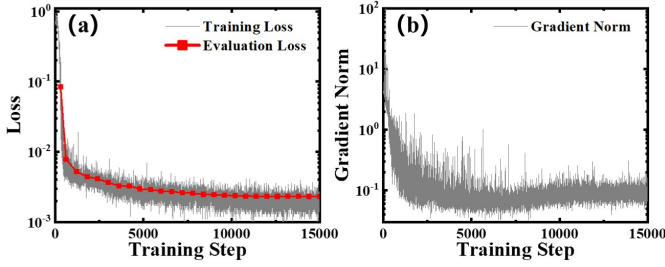


Fig. 3. SFT process: (a) training and evaluation loss, and (b) gradient norm.

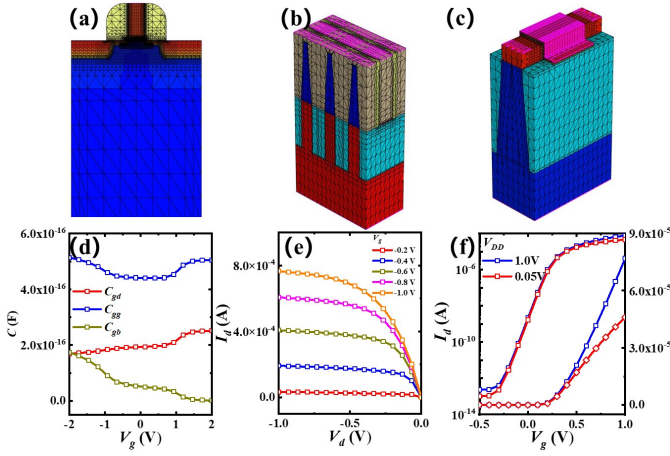


Fig. 4. Case study results with our TCAD code generation LLM with tailored meshing strategies. (a) 2D n-type MOSFET with Gaussian doping profile and (d) the corresponding $C - V$ simulation. (b) P-type FinFET array and (e) the corresponding $I_d - V_d$ simulation. (c) InGaAs-based n-type FinFET with high- k /metal gate stack and (f) the corresponding $I_d - V_g$ simulation.

each operating at 2.9 GHz with 64 cores, and 4 NVIDIA A6000 GPUs with 48 GB of memory. For SFT, we employed the open-source Qwen2.5-14B-Instruct model [30], implemented with the Transformers library. The fine-tuning was performed on 30,000 training samples for 5 epochs with a batch size of 8 in bfloat16 precision. The AdamW optimizer was used with momentum parameters $\beta_1 = 0.9$ and $\beta_2 = 0.95$, a weight decay of 1×10^{-4} , an initial learning rate of 1×10^{-5} , a cosine learning-rate decay schedule, and a warmup ratio of 0.05. The training procedure required approximately 107 hours on four A6000 GPUs. For TCAD execution, we utilized Synopsys Sentaurus W-2024.09-SP1. The code generation model was deployed on a single A6000 GPU, and the device optimization was carried out using DeepSeek V3.1 reasoning version [31].

B. Open TCAD Dataset and Model Evaluation

Fig. 3 summarizes the SFT process for TCAD code generation. The training loss in Fig. 3(a) decrease consistently, with periodic evaluations confirming the absence of overfitting. The gradient norm shown in Fig. 3(b) gradually declines as the parameters are updated, indicating stable convergence and effective acquisition of domain-specific knowledge. Representative case studies are presented in Fig. 4, where the fine-tuned model generates both device structures and corresponding simulations. The 2D n-type MOSFET in Fig. 4(a) includes a Gaussian doping distribution to emulate ion implantation, and the generated script produces the expected $C - V$ characteristics in Fig. 4(d). The p-type FinFET array shown in Fig.

TABLE I
COMPARISON OF TCAD CODE GENERATION APPROACHES.

Functionality	SDE Code		Sdevice Code		Device Agent
	Structure	Mesh	I-V	C-V	Optimization
[11]	✓	×	×	×	×
MALTS [13]	✓	✓	×	×	×
ChatTCAD [12]	✓	✓	×	×	×
AgenticTCAD	✓	✓	✓	✓	✓

4(b) yields valid $I_d - V_d$ results at different V_g in Fig. 4(e). Moreover, The InGaAs-based n-type FinFET with a high- k metal gate stack in Fig. 4(c) generates both linear and saturated $I_d - V_g$ curves in Fig. 4(f). These results demonstrate both the diversity of our dataset and the ability of SFT to enable reliable domain-specific TCAD code generation.

Table I compares our framework with the state-of-the-art approaches. Existing studies on TCAD code generation remain largely focused on the SDE stage with preliminary explorations, and the community still faces the bottleneck of lacking diverse datasets. In contrast, we provide an expert-crafted TCAD dataset and establish a complete natural language-driven framework that employs a domain-specific TCAD LLM as the code generator to cover device structure and electrical simulation, and further integrates a closed-loop device optimization stage, thereby pushing this research area forward.

C. Multi-Agent Workflow Evaluation

To evaluate the effectiveness of the proposed multi-agent workflow, we apply **AgenticTCAD** to the design and optimization of a NS-FET at advanced technology nodes, which represents a key direction for next-generation transistor architectures. Guided by the high-performance specifications defined in IRDS-2024 for the 2 nm technology node with a supply voltage of 0.65 V, the framework aims to adjust device parameters such that the resulting performance meets design requirements.

Fig. 5(a) illustrates the three-dimensional device structure and cross-section automatically generated through the code generation process. Fig. 5(b) shows the corresponding $I_d - V_g$ characteristics before and after optimization. The blue curve corresponds to an initial design with randomly assigned parameters, while the red curve represents the outcome of our agent-based optimization, which achieves a significant improvement in device performance. Table II summarizes the quantitative results of **AgenticTCAD** in comparison with the IRDS-2024 specifications. Although the randomly initialized design parameters yield an I_{on} that already satisfies the requirement, the SS reaches 286.72 mV/dec and the I_{off} exceeds the specification by more than 3 orders of magnitude, indicating unacceptable leakage power and poor switching characteristics. After optimization, **AgenticTCAD** successfully reduces I_{off} below the target while achieving an SS of 60.38 mV/dec, surpassing the IRDS requirement of 72 mV/dec for 2 nm technology node and approaching the theoretical limit of silicon-based devices. At the same time, I_{on} is further enhanced to 2.31×10^{-3} A/ μm , which is 2.94 times higher than the specified requirement, leading to a significant improvement in the ON-OFF ratio.

We further examine how different prompting strategies influence the reasoning ability of the device optimization agent. The results are presented in Fig. 6, which compares optimization trajectories under qualitative and quantitative




Fig. 5. TCADAgent for NS-FET design and optimization: (a) device structure and cross-section, (b) $I_d - V_g$ before (blue) and after (red) optimization.

TABLE II
OPTIMIZATION RESULTS COMPARED WITH IRDS SPECIFICATIONS

Metric	IRDS Spec	Before opt.	After opt.	Meets Spec
SS (mV/dec)	72 (↓)	286.72	60.38	✓
I_{off} (A/μm)	1.0×10^{-8} (↓)	3.40×10^{-5}	8.26×10^{-9}	✓
I_{on} (A/μm)	7.87×10^{-4} (↑)	4.23×10^{-3}	2.31×10^{-3}	✓
ON-OFF ratio	4.90 (↑)	2.10	5.45	✓

guidance, shown in blue and red curves respectively. In the case of qualitative guidance, the agent is only informed of the general optimization directions, namely to increase I_{on} , to decrease I_{off} , and to reduce the SS. As illustrated in Fig. 6(a), the optimization of SS under this setting is unstable. Although the agent is able to approach the target value at certain iterations, the performance often deteriorates in later steps even when the results of previous simulations are provided. A similar trend can be observed in Fig. 6(b), where I_{off} cannot be effectively improved when the initial value is far from the specification. In this case, the best result obtained through qualitative guidance still exceeds the design requirement by nearly an order of magnitude, which indicates the limitation of providing only directional hints.

When quantitative guidance is adopted, the agent is instead provided with explicit numerical targets relative to the design specifications. This enables stable convergence of the optimization process. As shown in Fig. 6(a), the SS target is rapidly achieved, and the optimization continues toward values close to the physical limit of silicon-based devices. Since I_{on} and I_{off} are typically in a trade-off relationship, the quantitative setting allows the agent to recognize that I_{on} already exceeds the requirement by a large margin and to strategically sacrifice a small portion of drive current in exchange for significant leakage reduction. This behavior can be clearly observed in Figs. 6(b) and 6(c), where I_{off} is successfully optimized even when the initial parameters are highly unfavorable. Furthermore, Fig. 6(d) shows that the ON-OFF ratio consistently improves during the entire optimization trajectory. This outcome confirms that although some I_{on} is traded off, the overall device performance is substantially enhanced. These results demonstrate that quantitative guidance enables stable and convergent optimization, underscoring the crucial role of prompt design in multi-agent workflows.

In terms of optimization efficiency, Fig. 6 demonstrates that **AgenticTCAD** requires only 25 iterations to achieve the device performance targets specified in IRDS-2024 for the 2 nm NS-FET, with a total runtime of 4.2 hours. By contrast, a TCAD expert conducted 5,120 simulations using five commercial licenses and required 7.1 days to identify the same design point. These results confirm that **AgenticTCAD**




Fig. 6. Optimization trajectories of device performance under qualitative guidance (blue) and quantitative guidance (red) for (a) SS, (b) I_{off} , (c) I_{on} , and (d) ON-OFF ratio. Missing iteration points correspond to non-convergent design parameters




Fig. 7. Physical insight of NS-FET with energy band diagrams along source-channel-drain in (a) OFF state and (b) ON state.

can substantially accelerate device design and optimization. Furthermore, during the optimization process, the workflow automatically extracts energy band diagrams, as shown in Fig. 7. Under different operating conditions, the diagrams provide direct evidence of carrier transport mechanisms, thereby offering a physically grounded insight for interpreting the optimization trajectory. In this way, **AgenticTCAD** not only expedites convergence toward target specifications but also enhances the interpretability and transparency of the LLM-based TCAD workflow by explicitly linking quantitative performance improvements to the corresponding underlying device physics.

V. CONCLUSION

We have presented **AgenticTCAD**, a natural language-driven framework for automated device simulation and optimization. Alongside this framework, we constructed an open-source TCAD dataset and fine-tuned a domain-specific language model for TCAD code generation. The effectiveness of **AgenticTCAD** is demonstrated through the design and optimization of a 2 nm NS-FET, where the framework successfully meets the IRDS-2024 specifications within an end-to-end flow. This contribution marks a significant advance for the TCAD community by establishing a scalable, data-driven pathway to automated device design and optimization and lays the groundwork for future research in interpretable, AI-assisted semiconductor technology development.

REFERENCES

- [1] A. Kelleher, "Celebrating 75 years of the transistor a look at the evolution of moore's law innovation," in *2022 International Electron Devices Meeting (IEDM)*, 2022, pp. 1.1.1–1.1.5.
- [2] Z. Ji, Y. Xue, P. Ren, J. Ye, Y. Li, Y. Wu, D. Wang, S. Wang, J. Wu, Z. Wang, Y. Wen, S. Xia, L. Zhang, J. Zhang, J. Liu, J. Luo, H. Deng, R. Wang, L. Yang, and R. Huang, "Toward reliability- and variability-aware design-technology co-optimization in advanced nodes: Defect characterization, industry-friendly modeling, and ml-assisted prediction," *IEEE Transactions on Electron Devices*, vol. 71, no. 1, pp. 138–150, 2024.
- [3] L. Silvestri, M. Palsgaard, R. Rhyner, M. Frey, J. Wellendorff, S. Smidstrup, R. Gull, and K. E. Sayed, "Hierarchical modeling for tcad simulation of short-channel 2d material-based fets," *Solid-State Electronics*, vol. 200, p. 108533, 2023.
- [4] S. Kim, K. Lee, S. Kim, M. Kim, J.-H. Lee, S. Kim, and B.-G. Park, "Investigation of device performance for fin angle optimization in finfet and gate-all-around fets for 3 nm-node and beyond," *IEEE Transactions on Electron Devices*, vol. 69, no. 4, pp. 2088–2093, 2022.
- [5] H. Kükner, G. Mirabelli, S. Yang, L. Verschuere, J. Bömmels, J. Y. Lin, D. Abdi, A. Farokhnejad, A. Zografos, N. Horiguchi, M. G. Bardón, G. Hellings, and J. Ryckaert, "Double-row cfet: Design technology co-optimization for area efficient a7 technology node," in *2024 IEEE International Electron Devices Meeting (IEDM)*, 2024, pp. 1–4.
- [6] Synopsys, "Tcad — technology computer-aided design," <https://www.synopsys.com/manufacturing/tcad.html>, accessed: 2025-09-07.
- [7] D. M. Sathaiya, T. Y. Hung, E. Chen, W.-C. Wu, A. Wei, C.-P. Chuu, S.-K. Su, A. S. Chou, C.-T. Chung, C.-H. Chien, H. Wang, J. Cai, C.-C. Wu, I. P. Radu, and J. Wu, "Comprehensive physics based tcad model for 2d mx2 channel transistors," in *2022 International Electron Devices Meeting (IEDM)*, 2022, pp. 28.4.1–28.4.4.
- [8] W. Fang, J. Wang, Y. Lu, S. Liu, Y. Wu, Y. Ma, and Z. Xie, "A survey of circuit foundation model: Foundation ai models for vlsi circuit design and eda," 2025. [Online]. Available: <https://arxiv.org/abs/2504.03711>
- [9] F. Cui, C. Yin, K. Zhou, Y. Xiao, G. Sun, Q. Xu, Q. Guo, Y. Liang, X. Zhang, and D. Lin, "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 (ICCAD)*. New York, NY, USA: Association for Computing Machinery, 2025, pp. Article 99, 1–9.
- [10] Z. Chen, J. Huang, Y. Liu, F. Yang, L. Shang, D. Zhou, and X. Zeng, "Artisan: Automated operational amplifier design via domain-specific large language model," in *Proceedings of the 61st ACM/IEEE Design Automation Conference*, ser. DAC '24. New York, NY, USA: Association for Computing Machinery, 2024.
- [11] L. M. Long Nguyen, A. Lu, and H. Y. Wong, "Tcad structure input file generation using large language model," in *2024 International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*, 2024, pp. 1–4.
- [12] Z. Zhang, S. Nassif, and M. Tahoori, "ChatTCAD: Leveraging Large Language Models for Automated TCAD Simulation File Generation," in *2025 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*. Los Alamitos, CA, USA: IEEE Computer Society, Jul. 2025, pp. 1–6.
- [13] J. Dou, J. Zhang, B. Ye, Y. Shen, X. Li, Y. Shi, Y. Zhang, and Y. Sun, "Malts: A multi-agent large language model-based layout-to-tcad structure synthesizer," in *2025 International Symposium of Electronics Design Automation (ISED)*, 2025, pp. 795–800.
- [14] S. Carapezzi, E. Caruso, A. Gnudi, P. Palestri, S. Reggiani, and E. Gnani, "Tcad mobility model of iii-v short-channel double-gate fets including ballistic corrections," *IEEE Transactions on Electron Devices*, vol. 64, no. 12, pp. 4882–4888, 2017.
- [15] T. Srimani, G. Hills, X. Zhao, D. Antoniadis, J. A. del Alamo, and M. M. Shulaker, "Asymmetric gating for reducing leakage current in carbon nanotube field-effect transistors," *Applied Physics Letters*, vol. 115, no. 6, p. 063107, 08 2019.
- [16] S. Chen, B. Peng, Y. Jiao, Y. Li, L. Zhang, Z. Sun, Y. Xue, M. Li, Z. Ji, R. Wang, and R. Huang, "Self-heating effects of vertical gate-all-around transistors: Analysis and modeling," *IEEE Transactions on Electron Devices*, vol. 71, no. 11, pp. 6478–6485, 2024.
- [17] J. Jiang, F. Wang, J. Shen, S. Kim, and S. Kim, "A survey on large language models for code generation," 2024. [Online]. Available: <https://arxiv.org/abs/2406.00515>
- [18] K. Min, S. Park, H. Park, J. Cho, and S. Kang, "Improving llm-based verilog code generation with data augmentation and rl," in *2025 Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–7.
- [19] M. ul Islam, H. Sami, P.-E. Gaillardon, and V. Tenace, "Eda-aware rtl generation with large language models," in *2025 Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–6.
- [20] Y. Tsai, M. Liu, and H. Ren, "Rtlfixer: Automatically fixing rtl syntax errors with large language model," in *Proceedings of the 61st ACM/IEEE Design Automation Conference*, ser. DAC '24. New York, NY, USA: Association for Computing Machinery, 2024.
- [21] K. Xu, J. Sun, Y. Hu, X. Fang, W. Shan, X. Wang, and Z. Jiang, "Meic: Re-thinking rtl debug automation using llms," in *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*. New York, NY, USA: Association for Computing Machinery, 2025.
- [22] X. Yao, Y. Wang, X. Li, Y. Lian, R. Chen, L. Chen, M. Yuan, H. Xu, and B. Yu, "Rtlrewriter: Methodologies for large models aided rtl code optimization," in *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, ser. ICCAD '24. New York, NY, USA: Association for Computing Machinery, 2025.
- [23] H. Wu, Z. He, X. Zhang, X. Yao, S. Zheng, H. Zheng, and B. Yu, "Chateda: A large language model powered autonomous agent for eda," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 43, no. 10, pp. 3184–3197, 2024.
- [24] Y. Yin, Y. Wang, B. Xu, and P. Li, "Ado-llm: Analog design bayesian optimization with in-context learning of large language models," in *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*. New York, NY, USA: Association for Computing Machinery, 2025.
- [25] X. Wang, G.-W. Wan, S.-Z. Wong, L. Zhang, T. Liu, Q. Tian, and J. Ye, "Chatcpu: An agile cpu design and verification platform with llm," in *Proceedings of the 61st ACM/IEEE Design Automation Conference*, ser. DAC '24. New York, NY, USA: Association for Computing Machinery, 2024.
- [26] C.-T. Ho and H. Ren, "Large language model (llm) for standard cell layout design optimization," in *2024 IEEE LLM Aided Design Workshop (LAD)*, 2024, pp. 1–6.
- [27] B. Liu, H. Zhang, X. Gao, Z. Kong, X. Tang, Y. Lin, R. Wang, and R. Huang, "Layoutcopilot: An llm-powered multiagent collaborative framework for interactive analog layout design," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 44, no. 8, pp. 3126–3139, 2025.
- [28] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. de Oliveira Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, and W. Zaremba, "Evaluating large language models trained on code," 2021. [Online]. Available: <https://arxiv.org/abs/2107.03374>
- [29] IEEE International Roadmap for Devices and Systems (IRDS), "2024 irds international focus teams (ifts) roadmap reports," <https://irds.ieee.org/editions/2024>, 2024.
- [30] Q. Team, "Qwen2.5 technical report," 2025. [Online]. Available: <https://arxiv.org/abs/2412.15115>
- [31] DeepSeek-AI Team, "Deepseek v3.1," <https://www.deepseek.com/>, 2025.