

STEP-LLM: Generating CAD STEP Models from Natural Language with Large Language Models

Xiangyu Shi, Junyang Ding, Xu Zhao, Sinong Zhan, Payal Mohapatra, Daniel Quispe, Kojo Welbeck, Jian Cao, Wei Chen, Ping Guo, Qi Zhu
Northwestern University, Evanston, USA

Abstract—Computer-aided design (CAD) is vital to modern manufacturing, yet model creation remains labor-intensive and expertise-heavy. To enable non-experts to translate intuitive design intent into manufacturable artifacts, recent large language models (LLM)-based text-to-CAD efforts focus on command sequences or script-based formats like CadQuery. However, these formats are kernel-dependent and lack universality for manufacturing. In contrast, the Standard for the Exchange of Product Data (STEP, ISO 10303) file is a widely adopted, neutral boundary representation (B-rep) format directly compatible with manufacturing, but its graph-structured, cross-referenced nature poses unique challenges for auto-regressive LLMs. To address this, we curate a dataset of $\sim 40\text{K}$ STEP-caption pairs and introduce novel preprocessing tailored for the graph-structured format of STEP, including a depth-first search (DFS)-based reserialization that linearizes cross-references while preserving locality and chain-of-thought (CoT)-style structural annotations that explicitly guide global coherence. We integrate retrieval-augmented generation (RAG) to ground predictions in relevant examples for supervised fine-tuning (SFT), and further refine generation quality through reinforcement learning (RL) with a specific Chamfer Distance-based geometric reward. Experiments demonstrate consistent gains of our STEP-LLM in geometric fidelity over the Text2CAD baseline, with improvements arising from multiple stages of our framework: the RAG module substantially enhances completeness and renderability, the DFS-based reserialization strategy strengthens overall accuracy, and the RL refinement further reduces geometric discrepancy. Both metrics and visual comparisons confirm that STEP-LLM generates shapes with higher fidelity than Text2CAD. These results demonstrate the feasibility of LLM-driven STEP model generation from natural language, showing its potential to democratize CAD design for manufacturing.

Index Terms—Computer-aided design, STEP file, large language models, design automation.

I. INTRODUCTION

Computer-aided design (CAD) plays a foundational role in modern design and manufacturing. However, creating CAD models remains a specialized and labor-intensive process that requires extensive expertise and time, limiting accessibility for non-expert users and slowing the pace of prototyping. More broadly, the entire pipeline from design to manufacturing still relies heavily on manual intervention, with limited automation in translating design intent into directly manufacture models (Fig. 1). Recently, the rapid progress of Generative AI has opened new possibilities for CAD workflow, raising the prospect that intuitive design intent expressed in natural




Fig. 1. The typical workflow of CAD from design to manufacturing.

language could be transformed into manufacture artifacts. Realizing this vision would allow non-experts to translate high-level ideas into manufacturable artifacts, lowering barriers for CAD design and manufacturing.

Recent attempts at text-to-CAD generation have largely relied on **command sequences** [1] or **script-based formats** such as CadQuery [2]. Command sequences represent the design history as a series of modeling operations (e.g., sketch, extrude) that can be replayed by a CAD kernel, while script-based formats express CAD models as executable code (e.g., Python or macros) that calls modeling APIs. These representations align well with large language models (LLMs)' code generation abilities and allow validation through external engines [3]. However, tied to specific kernels, these formats cannot directly support manufacturing and are constrained to simple objects by omitting complex operations such as fillets or free-form surfaces [4]. In contrast, **boundary representation (B-rep)** encodes complete topology and geometry, providing the expressiveness needed to model complex designs. Recent works such as SolidGen [5] and BrepGen [4] demonstrate its learnability, yet remain focused on geometric validity rather than manufacturability. Within this family, the Standard for the Exchange of Product Data (STEP) file is a widely adopted, an engine-agnostic standard that **directly interfaces with industrial pipelines**, ensuring models can be used in downstream CAD/CAM. Although a few studies have analyzed STEP for machining feature recognition or entity parsing [6], [7], direct generation of STEP files from natural language has not been explored. This gap motivates our work: leveraging LLMs to directly produce STEP models from natural language, combining LLMs' generative strengths with STEP's universality for manufacture-ready design.

However, direct STEP generation presents unique challenges. A STEP file is intrinsically **graph-structured**, with **cross-references and non-sequential dependencies** that conflict with the left-to-right auto-regressive paradigm of LLMs [8]. In addition, small errors in entity ordering or identifier usage can render the entire file invalid.

To address this, we propose STEP-LLM, a complete fine tun-

Source code is available at <https://github.com/JasonShiii/STEP-LLM>. Corresponding authors: Xiangyu Shi (xiangyushi2029@u.northwestern.edu), Qi Zhu (qzhu@northwestern.edu). The work was supported in part by the National Science Foundation grants 2328032, 2324936, and 2328973.

ing framework for natural language to STEP generation based on state-of-the-art small models. We first construct a dataset consisting around 40K caption-STEP pairs, and introduce a depth-first search (DFS)-style reserialization strategy that linearizes STEP’s graph structure to preserve **local sequential ordering**. Moreover, CoT-style annotations are used to guide **global coherence**. We further integrate retrieval-augmented generation (RAG) into supervised fine-tuning (SFT), grounding outputs in relevant examples. In addition, we generalize model’s reasoning capability through reinforcement learning (RL) with a specific scaled Chamfer Distance-based geometric reward, which is robust to translation, orientation, and scale differences, providing a reliable objective for optimizing shape fidelity.

In the absence of broadly established benchmarks, we evaluate our method against the Text2CAD [1] baseline on renderability, geometric fidelity, and the alignment of generated STEP file complexity with the distribution of ground-truth models. In addition, we conduct a comprehensive ablation study to verify the contribution of each design choice, including RAG, DFS-based reserialization, RL, and base model selection. Experiment results demonstrate that our framework achieves high completion and renderability rates, while significantly improving geometric accuracy compared to baselines. Our main contributions are as follows:

- We propose **STEP-LLM**, the first unified framework for direct STEP file generation from natural language, bridging LLMs with a universal CAD standard. The framework integrates RAG to enrich training context and augment SFT, and RL with geometry-aware rewards to further enhance geometric fidelity and robustness.
- We introduce a **novel DFS-based reserialization strategy** for STEP file preprocessing, which linearizes the graph-like structure of STEP files to better align with the autoregressive nature of LLMs. In addition, we incorporate **CoT-style structural annotations** to guide global coherence, enabling the model to capture both local sequential ordering and long-range dependencies.
- We release a **curated dataset** of caption-STEP pairs and a set of **evaluation metrics** for benchmarking text-to-STEP generation, facilitating future research in this area.

II. BACKGROUND AND RELATED WORK

A. Large Language Models for CAD Generation

Recent works have explored applying LLMs to CAD generation, primarily focusing on command sequences [1], [9]–[15] or script-based formats such as CadQuery [2], [16]–[18] and CAD Macros [19]. Command sequences, popularized by datasets such as DeepCAD [20], represent CAD design history as procedural instructions recording the modeling operations. FlexCAD [10] converts a CAD model into a brief, structured text and employs hierarchy-aware masking to fine-tune LLMs for controllable CAD generation tasks. CAD-Llama [12] follows an adaptive pretraining paradigm combined with instruction tuning on a multitask instructional dataset for multi-task modeling. Several works [9], [11] further expand the input to multi-modality such as image, point cloud. Script-based

formats express CAD models as executable code or scripts that call modeling APIs in a human-readable style. Representative examples include Query2CAD [19], which uses FreeCAD macros with self-refinement loops, and CAD-Coder [17], which translates natural language into CadQuery programs. Due to LLMs’ strong capability in code generation, these approaches show clear advantages in producing syntactically valid and easily verifiable CAD scripts.

Despite these advances, both command-sequence and script-based approaches rely heavily on external CAD engines for execution, limiting universality across design platforms. Moreover, the expressiveness of these formats is constrained by the APIs of the target engine or the vocabulary of the sequence, restricting their ability to capture complex operations such as fillet, chamfer, or B-spline surfaces. Therefore, while effective for controlled tasks, these representations are not directly suitable for downstream manufacturing, motivating the need for more universal CAD formats.

B. Learning-based B-rep Generation and STEP File Analysis

Unlike command sequences or script-based formats, which abstract away from geometric detail, B-rep encodes the complete topology and geometry of a solid model, including vertices, edges, faces, and their connectivity, making it inherently capable of describing complex operations. Several works in geometric deep learning addressed B-rep generation directly. SolidGen [5] employed an auto-regressive strategy to synthesize vertices, edges, and faces in sequence. BrepGen [4] advanced this line with a diffusion-based model over a structured latent geometry tree, enabling generation of free-form surfaces. ComplexGen [21] proposed a chain-complex view of B-reps, reconstructing corners, curves, and surface patches jointly, while Point2CAD [22] extended this paradigm to reverse-engineer B-reps from point clouds. These studies confirm that the graph-structured nature of B-reps is learnable and expressive, yet they remain primarily theoretical explorations focused on geometric validity rather than manufacturability and cannot be directly deployed in industrial workflows.

Along with these advances, the STEP file provides a standardized textual realization of B-rep and has become one of the most universal, engine-independent formats for exchanging product manufacturing information. A STEP file is organized into two main sections: the HEADER SECTION, which specifies metadata such as schema, author, and units, and the DATA SECTION, which encodes the actual geometry and topology. In the DATA SECTION, models are represented as a collection of entities (e.g., CARTESIAN_POINT, EDGE, FACE) linked together by unique identifiers, enabling cross-references that capture the full structure of the design. Limited prior work has treated STEP file as a structured language [6], using recursive neural networks for tasks such as machining feature recognition and entity analysis [7]. These efforts demonstrate the feasibility of parsing STEP files directly, but to the best of our knowledge, none has attempted LLM-based generation of STEP. This gap motivates our work: we investigate whether an LLM can be taught to “speak” the language of STEP,

combining the generative strengths of modern language models with the universality and manufacturability of a CAD standard.

III. THE STEP-LLM FRAMEWORK

Our framework (Fig. 2) consists of three main stages. First, we construct a paired caption–STEP dataset through rendering, captioning, and reserialization (Section III-A). Second, we perform supervised fine-tuning (SFT) with RAG (Section III-B), allowing the model to learn STEP grammar and leverage relevant prior examples for more consistent outputs. Finally, we refine the model with RL alignment (Section III-C), where a geometric reward based on scaled Chamfer Distance further improves geometric fidelity.

A. Dataset Construction and Preprocessing

We build our dataset upon the ABC dataset [23], a large-scale collection of CAD models released for geometric deep learning. The ABC dataset contains over one million B-rep models exported to STEP format, alongside corresponding parametric curves and surface annotations. To control training efficiency and account for model complexity, we adopt the entity number of a STEP file as a proxy for shape complexity, preparing a filtered dataset for proceeding supervised fine-tuning (see Section III-B). This selection strategy ensures that the data scale and difficulty are appropriate for LLM training.

To construct text–STEP paired data, we render each STEP file into multi-view images (9 views, resolution 1200×1200) following the common practice in prior work such as Text2CAD, and as illustrated in our system framework (Fig. 2). These include both orthographic and perspective views, providing comprehensive geometric coverage. We then employ GPT-4o to assess the renderings and generate captions for those deemed describable. In practice, however, we observe that current foundation models, including GPT-4o, face challenges in describing multi-view renderings: they often lack **perspective consistency** and may incorrectly interpret different views of the same object as multiple objects [24]–[26].

To overcome this limitation, we apply tailored prompt engineering to guide GPT-4o’s captioning process. Our prompts encourage concise yet informative textual descriptions that emphasize salient geometric features (e.g., symmetry, through holes, fillets) while also linking objects to real-world categories whenever appropriate. For example, a model might be captioned as “A flat circular lid with two rectangular mounting tabs and a central recessed feature.” This strategy produces captions that are human-interpretable and sufficiently structured to supervise STEP file generation in subsequent model training.

From a structural perspective, STEP files are challenging targets for LLMs. Each file encodes a model as a collection of entities connected by cross-references, forming a **directed acyclic graph (DAG)** that captures the hierarchical relationships between primitives and higher-level geometry. However, in their raw textual form this DAG structure becomes highly non-sequential: related entities may be scattered far apart, and long-range identifier dependencies must be recalled precisely. For an auto-regressive model, this easily leads to incoherent or invalid outputs.

To mitigate this issue, we introduce a DFS-based reserialization strategy. Specifically, we first parse the STEP file into a **hierarchical tree structure** where each node represents an entity and its child nodes correspond to directly referenced entities. We then serialize the tree using a depth-first traversal. This approach allows each branch of the tree to be expressed as a relatively local and coherent sequence, reducing the burden of long-range dependency tracking. To prevent entity explosion when traversing large graphs, we employ **strategic pruning**, ensuring only structurally relevant branches are expanded and each reference relationship appears once. Furthermore, we **renumber entity identifiers sequentially**, eliminating irregular gaps in the raw file and simplifying reference tracking. We also normalize floating-point precision, reducing unnecessary digits while preserving geometric validity, thereby lowering textual complexity without altering topology.

Nevertheless, DFS traversal alone cannot fully resolve discontinuities when switching between parallel branches. For example, once one branch of an entity is fully expanded, the model must switch context to a different branch of the same parent entity. Such discontinuities can cause the model to lose coherence. Inspired by recent progress in code generation with CoT-style annotations [27], we augment the reserialized STEP files with lightweight statistical annotations. These annotations summarize branch-level statistics (e.g., number of child entities, branch depth) and act as guidance tokens, helping the LLM to reason about the **global structure** while maintaining consistency across branch transitions.

Overall, this DFS-based reserialization, combined with structural annotations, preserves locality while enhancing global comprehensibility. This strategy also mitigates the **inherent heterogeneity** of STEP files, where different entity orders may in fact correspond to the same model, thereby providing a more consistent representation for downstream learning.

B. Supervised Fine-Tuning with RAG

Supervised fine-tuning is an essential step in adapting LLMs for domain-specific applications [3], [28], [29]. In our setting, SFT allows the model to internalize the grammar of STEP files and learn the mapping between natural language captions and precise geometric representations.

To enhance this process, we integrate RAG into the fine-tuning pipeline, which enriches the model’s context by grounding generation in semantically relevant prior examples. This approach reduces the reliance on parameter memorization and promotes output faithfulness. Such grounding is particularly critical for STEP generation, where long-range dependencies and structural consistency challenges LLMs.

Our retrieval module is implemented as follows. We construct an external database of paired STEP files and captions. For a given input caption, the system retrieves the most semantically similar caption and its associated STEP file from the database. Specifically, we use SentenceTransformer [30] to embed captions into dense vectors and FAISS [31] to index them for efficient nearest-neighbor search based on cosine similarity, ensuring that structurally relevant cases are retrieved. During SFT, the input prompt is augmented with both the

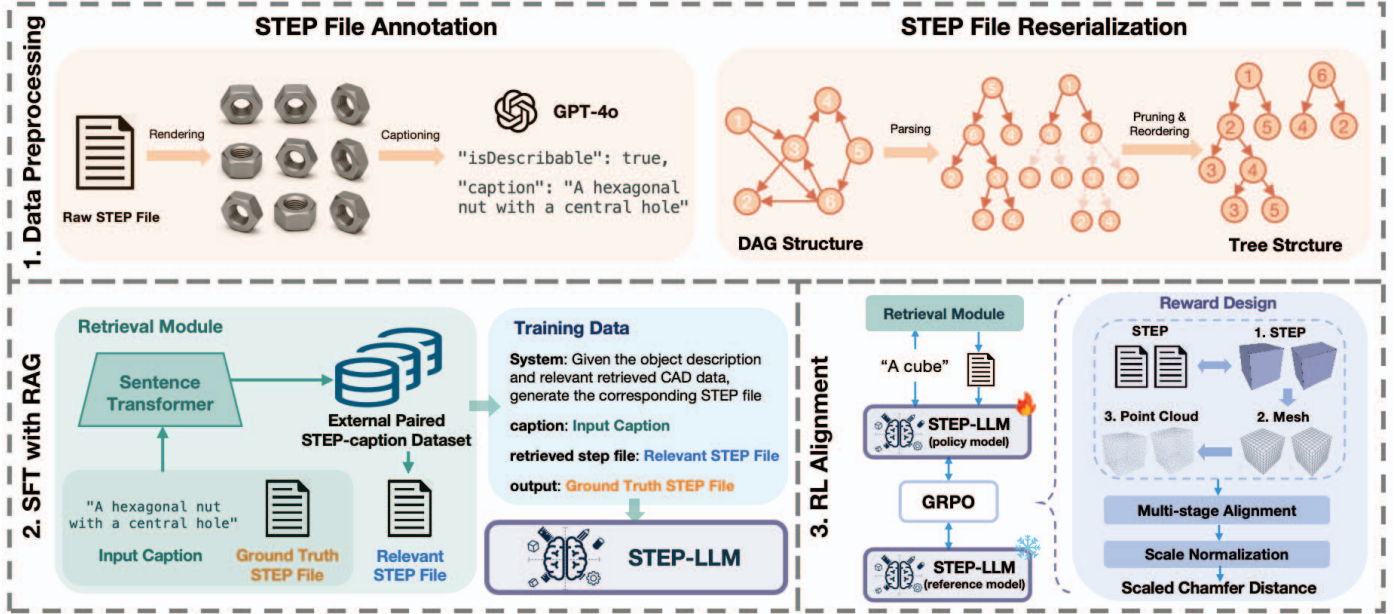


Fig. 2. Framework of STEP-LLM. The top panel illustrates data preprocessing: raw STEP files are rendered into nine views and captioned with GPT-4o (left), and their internal DAG structures (the numbers represent the identifier of the entity in a STEP file) are reserialized into locality-preserving trees (right). These captions and DFS-style STEP files together form the paired STEP-caption dataset for SFT (bottom-left). A specific geometric reward based on scaled Chamfer Distance is designed for further RL training (bottom-right).

original caption and the retrieved STEP file, while the target output remains the ground-truth STEP file of the original caption. This design provides the model with structural cues while still requiring faithful reproduction of the target object.

We adopt SFT on a curated subset of STEP files with fewer than 500 entities. Complex STEP files with higher entity counts (500–1000) often contain long sequences of repetitive entities, such as CARTESIAN_POINT or DIRECTION, which can cause LLMs to fall into repetitive loops and fail to produce complete files. By focusing training on simpler files, the model effectively learns fundamental STEP syntax and compositional rules, which improves stability and robustness when generalizing to more complex cases. This strategy enhances both performance and reliability in generating valid STEP files.

C. Reinforcement Learning Alignment

Recent advances in natural language processing have shown that RL can further improve model performance beyond supervised fine-tuning [32], [33]. While SFT equips the model with the grammar of STEP files and the ability to map captions to geometry, it does not guarantee that generated outputs are geometrically faithful. In RL, one of the most critical factors is the **reward design** [3], [34], [35]. To explicitly encode geometric fidelity into the training objective, we design a reward function based on the **Scaled Chamfer Distance (SCD)**, which is robust to translation, orientation, and scale differences. We implement this reward within a standard policy optimization framework using Group Relative Policy Optimization (GRPO) [36], and observe clear improvements in geometric accuracy.

Scaled Chamfer Distance (SCD). To evaluate geometric fidelity, we first compute the Chamfer Distance (CD) between the generated model and the ground truth. Both STEP files

are converted into STL meshes, from which point clouds are sampled. Given point sets P (prediction) and Q (ground truth), the bidirectional Chamfer Distance is defined as:

$$CD(P, Q) = \frac{1}{|P|} \sum_{p \in P} \min_{q \in Q} \|p - q\|_2 + \frac{1}{|Q|} \sum_{q \in Q} \min_{p \in P} \|q - p\|_2 \quad (1)$$

However, raw CD is sensitive to translation, orientation, and scale. To ensure fair comparison, we apply a multi-stage alignment strategy:

- **Center alignment:** shift both point clouds so that their centroids coincide, removing large translations.
- **Global registration:** compute a coarse rigid transform via feature matching (e.g., FPFH + RANSAC).
- **Iterative Closest Point (ICP):** iteratively refine the alignment for precise point-to-point correspondence.

We then normalize the distance by dividing it with the squared scale factor:

$$SCD = \frac{CD}{(\text{Scale Factor})^2}, \quad (2)$$

where the scale factor is defined as the root mean square distance of ground-truth points from its centroid.

Reward design. Inspired by CAD-Coder [17], we adopt a piecewise linear reward function with two thresholds. If the SCD is below the lower bound, the reward is 1; if above the upper bound, the reward is 0; for intermediate values, the reward is interpolated linearly to avoid sparse rewards:

$$R^{\text{geo}}(S) = \begin{cases} 1, & \text{if } SCD(S, S_{\text{gt}}) \leq \delta_{\text{low}}, \\ 0, & \text{if } SCD(S, S_{\text{gt}}) \geq \delta_{\text{high}}, \\ \frac{\delta_{\text{high}} - SCD(S, S_{\text{gt}})}{\delta_{\text{high}} - \delta_{\text{low}}}, & \text{otherwise.} \end{cases} \quad (3)$$

This design encourages the model to generate STEP files that not only remain syntactically valid but also yield geometrically accurate shapes, with robustness to common variations in scale, position, and orientation.

IV. EXPERIMENTS AND EVALUATION

A. Experimental Setup

We conduct our SFT on two representative open-source models: Llama-3.2-3B-Instruct and Qwen-2.5-3B. All training is performed using the Unsloth [37] framework for parameter-efficient fine-tuning. Our training corpus consists of 14,396 STEP-caption pairs, restricted to models with fewer than 500 entities. Additionally, we adopt LoRA [38] to reduce GPU memory consumption and improve efficiency. The training configuration includes an effective batch size of 16, the AdamW (8-bit) optimizer, and a learning rate of $2e-4$ with a linear decay scheduler. A warm-up of 5% of the total training steps is applied to stabilize early training. The models are trained for 10 epochs on a single NVIDIA A100 GPU. To facilitate evaluation, we save a checkpoint after each epoch and perform asynchronous validation.

For evaluation, we adopt a set of metrics that together capture syntactic validity, structural correctness, and geometric fidelity. Three primary metrics are used: Completion Rate (CR), Renderability Rate (RR), and Median Scaled Chamfer Distance (MSCD). CR verifies whether a generated STEP file terminates correctly, RR checks whether it can be reconstructed into a valid 3D object by OpenCASCADE [39], and MSCD quantifies geometric fidelity after alignment and scale normalization. Since baselines such as Text2CAD rely on command sequences and export STEP files through external engines, CR is not directly applicable in baseline comparisons. To address this, we additionally report Average Entity Count (AEC) as a proxy for design complexity, where a closer match between generated and ground-truth entity counts indicates stronger alignment with the expected difficulty level. More specifically, the metrics are as follows:

- **Completion Rate (CR):** percentage of generated STEP files that terminate correctly with the standardized "END-ISO-10303-21;" line.
- **Renderability Rate (RR):** proportion of files successfully parsed by OpenCASCADE into non-null shapes that can be meshed into valid STL files.
- **Median Scaled Chamfer Distance (MSCD):** median of scale-normalized Chamfer Distances across the test set, robust to translation, rotation, and scale differences.
- **Average Entity Count (AEC):** mean number of entities in generated STEP files compared to ground truth.

B. Advantages of STEP-LLM over Baseline

Evaluating text-to-CAD generation remains challenging as the field is still in its early stage, and there are no widely adopted benchmarks or standardized baselines [1], [40]. Among the available open-sourced approaches, we select Text2CAD [1] as the primary baseline. Published recently in NeurIPS 2024 as a spotlight paper, Text2CAD represents the strongest openly




Fig. 3. Entity count distribution of generated STEP files vs. ground truth.

available method for text-to-parametric CAD generation, and has quickly become the standard point of comparison in this emerging area [12], [40]. While several recent works report improvements under their own closed settings, their models or code are not publicly accessible, making Text2CAD the only suitable open-source baseline for fair and reproducible evaluation. In addition, it is capable of directly outputting STEP files during inference, which allows fair comparison and seamless integration into our evaluation pipeline without requiring additional kernel-dependent conversions.

We evaluate both methods on a held-out set of 2,056 samples. The results are summarized in Table I. **Our STEP-LLM delivers a renderability ratio comparable to Text2CAD, while achieving a substantially lower median scaled Chamfer distance, indicating higher geometric fidelity.** It is worth noting that Text2CAD benefits from relying on a CAD-kernel-based reconstruction and export pipeline, which enforces syntactic validity of STEP files and thus naturally yields a slightly higher renderability rate. In contrast, our framework generates STEP files directly, which makes the renderability criterion more stringent. Nevertheless, the obtained RR of over 95% shows that our approach remains robust and practically reliable, especially considering our clear advantage in geometric fidelity.

Beyond renderability and geometric distance, we further analyze the distribution of entity counts in generated STEP files (Table II and Fig. 3). Our results show that **STEP-LLM produces entity counts whose average and distribution are much closer to the ground truth**, whereas Text2CAD often generates excessively large structures, with some files exceeding 3,000 entities. Such over-generation highlights instability in Text2CAD, while our method better preserves the realistic complexity range observed in human-designed CAD models.

Finally, Fig. 4 shows qualitative results across different complexity levels. For both simple and complex prompts, STEP-LLM produces outputs that are visually more faithful to the described objects than Text2CAD, with notable improvements on

TABLE I
COMPARISON WITH TEXT2CAD BASELINE. RR: RENDERABILITY RATE; MSCD: MEDIAN SCALED CHAMFER DISTANCE; AEC: AVG. ENTITY COUNT (GROUND TRUTH: 265.64).

Method	RR (%) \uparrow	MSCD \downarrow	AEC
Text2CAD	98.38	3.99	390.41
STEP-LLM	95.18	0.53	240.99

TABLE II
ENTITY COUNT STATISTICS ON THE TEST SET.

Method	Avg.	Min	Max
Ground Truth	265.64	47	477
Text2CAD	390.41	50	3665
STEP-LLM	240.99	47	989

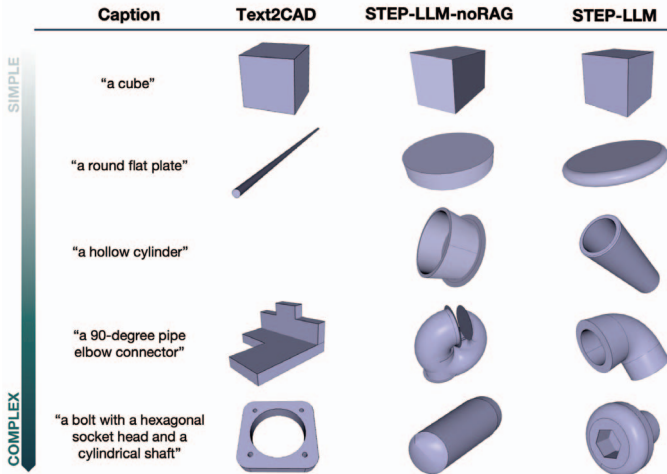


Fig. 4. The visual comparison between Text2CAD, STEP-LLM-noRAG and STEP-LLM.

curved geometries and detailed features. These results further demonstrate the effectiveness of our framework in advancing STEP file generation beyond current baseline.

C. Ablation Studies on Module Importance

To better understand the contributions of individual components in our framework, we conduct an ablation study comparing different configurations of model architecture and data processing. Specifically, we examine (i) the effect of DFS reserialization versus training on raw STEP files, (ii) the benefit of RAG versus non-RAG conditions, and (iii) differences across backbone models (Llama-3.2-3B-Instruct and Qwen-2.5-3B).

The results are summarized in Table III. Our full configuration, which combines DFS-based reserialization and RAG with a Llama base model (denoted as Llama-dfs-RAG), achieves the best trade-off among completion rate, renderability rate, and geometric fidelity, and is therefore adopted as the backbone of STEP-LLM. Removing either dfs (Llama-nodfs-RAG) or RAG (Llama-dfs-noRAG) leads to significantly worse performance, **showing the importance of both modules in our approach**. Note that when comparing RAG with no-RAG conditions, to ensure fairness, we trained all models for the same number of steps. Since RAG prompts are longer, this means the RAG condition processed more tokens overall. We regard step-matching as the fairer comparison, as it keeps the optimizer schedule identical across conditions. Finally, Qwen backbone model provides similar performance as Llama, and we chose the latter due to its better results in geometric fidelity (MSCD).

TABLE III
ABLATION STUDIES ON STEP RESERIALIZATION, RAG, AND BACKBONE MODELS. BOLD INDICATES BEST RESULT PER COLUMN.

Model Setting	CR (%) ↑	RR (%) ↑	MSCD ↓
Llama-dfs-RAG	0.97	0.95	0.53
Llama-nodfs-RAG	0.91	0.76	0.95
Llama-dfs-noRAG	0.84	0.13	0.61
Qwen-dfs-RAG	0.99	0.94	0.59

TABLE IV
RESULTS OF RL REFINEMENT ON LLAMA3.2-3B-INSTRUCT WITH CHAMFER DISTANCE REWARD.

Model	CR (%) ↑	RR (%) ↑	MSCD ↓
STEP-LLM	0.97	0.95	0.53
STEP-LLM-GRPO	0.99	0.92	0.098

D. Effectiveness of RL

We further explore the effectiveness of RL for STEP file generation using GRPO. For the threshold in reward design, we set the upper bound to 0.5 and the lower bound to 0.01 to ensure that the model is encouraged toward geometrically accurate generations while avoiding sparse reward issues. We cold start RL training from the previous SFT checkpoint of Llama3.2-3B-Instruct, and conduct GRPO optimization with the following hyperparameters: batch size of 8, 8 sampled responses per prompt, KL penalty coefficient of 0.02, entropy coefficient of 0.005, and a learning rate of 3e-6. Training was conducted on 4 NVIDIA H100 GPUs for a total of 80 optimization steps.

The results are summarized in Table IV. Compared to the supervised baseline, RL refinement leads to further improvement in completion ratio and a significant reduction in median Chamfer Distance, **demonstrating the effectiveness of even a small number of RL updates in enhancing geometric fidelity**. While the renderability rate shows a minor decline, this reflects the inherent trade-off introduced by optimizing for fine-grained geometric accuracy, as the model is encouraged to produce richer details. Importantly, the overall reliability remains high, confirming the effectiveness of using RL.

CONCLUSION

In this work, we introduced STEP-LLM, the first unified framework for direct LLM-based generation of STEP files. It incorporates four key components: DFS-based reserialization to linearize STEP’s graph structure, CoT-style annotation to ensure global coherence, retrieval-augmented fine-tuning to improve renderability, and reinforcement learning with geometry-aware rewards to refine fidelity and robustness. Together, these components enable the generation of syntactically valid, geometrically accurate, and manufacture-ready STEP representations.

Looking ahead in future work, scaling to larger models, refining reward functions with manufacturability or constraint checks, and expanding the dataset with richer captions and diverse CAD domains may further improve the model’s performance.

REFERENCES

- [1] M. S. Khan, S. Sinha, T. Uddin, D. Stricker, S. A. Ali, and M. Z. Afzal, "Text2cad: Generating sequential cad designs from beginner-to-expert level text prompts," *Advances in Neural Information Processing Systems*, vol. 37, pp. 7552–7579, 2024.
- [2] H. Xie and F. Ju, "Text-to-cadquery: A new paradigm for cad generation with scalable large model capabilities," *arXiv preprint arXiv:2505.06507*, 2025.
- [3] Y. Zhang, T. Wang, J. Gesi, Z. Wang, Y. Lu, J. Lin, S. Zhan, V. Gao, R. Jiao, J. Liu *et al.*, "Shop-r1: Rewarding llms to simulate human behavior in online shopping via reinforcement learning," *arXiv preprint arXiv:2507.17842*, 2025.
- [4] X. Xu, J. Lambourne, P. Jayaraman, Z. Wang, K. Willis, and Y. Furukawa, "Brepgen: A b-rep generative diffusion model with structured latent geometry," *ACM Transactions on Graphics (TOG)*, vol. 43, no. 4, pp. 1–14, 2024.
- [5] P. K. Jayaraman, J. G. Lambourne, N. Desai, K. D. Willis, A. Sanghi, and N. J. Morris, "Solidgen: An autoregressive model for direct b-rep synthesis," *arXiv preprint arXiv:2203.13944*, 2022.
- [6] V. Miles, S. Giani, and O. Vogt, "Approaching step file analysis as a language processing task: A robust and scale-invariant solution for machining feature recognition," *Journal of Computational and Applied Mathematics*, vol. 427, p. 115166, 2023.
- [7] M. Victoria, G. Stefano, and V. Oliver, "Recursive encoder network for the automatic analysis of step files," *Journal of Intelligent Manufacturing*, vol. 34, no. 1, pp. 181–196, 2023.
- [8] H. Wang, P. Wang, M. Li, S. Liu, S. Miao, Z. Wang, and P. Li, "Graphkv: Breaking sequence via injecting structural biases into large language models," *arXiv preprint arXiv:2506.07334*, 2025.
- [9] S. Wang, C. Chen, X. Le, Q. Xu, L. Xu, Y. Zhang, and J. Yang, "Cad-gpt: Synthesising cad construction sequence with spatial reasoning-enhanced multimodal llms," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 8, p. 7880–7888, Apr. 2025. [Online]. Available: <http://dx.doi.org/10.1609/aaai.v39i8.32849>
- [10] Z. Zhang, S. Sun, W. Wang, D. Cai, and J. Bian, "Flexcad: Unified and versatile controllable cad generation with fine-tuned large language models," *arXiv preprint arXiv:2411.05823*, 2024.
- [11] J. Xu, C. Wang, Z. Zhao, W. Liu, Y. Ma, and S. Gao, "Cad-mllm: Unifying multimodality-conditioned cad generation with mllm," *arXiv preprint arXiv:2411.04954*, 2024.
- [12] J. Li, W. Ma, X. Li, Y. Lou, G. Zhou, and X. Zhou, "Cad-llama: leveraging large language models for computer-aided design parametric 3d model generation," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 18 563–18 573.
- [13] Y. You, M. A. Uy, J. Han, R. Thomas, H. Zhang, S. You, and L. Guibas, "Img2cad: Reverse engineering 3d cad models from images through vlm-assisted conditional factorization," *arXiv preprint arXiv:2408.01437*, 2024.
- [14] S. Wu, A. H. Khasahmadi, M. Katz, P. K. Jayaraman, Y. Pu, K. Willis, and B. Liu, "Cadvlm: Bridging language and vision in the generation of parametric cad sketches," in *European Conference on Computer Vision*. Springer, 2024, pp. 368–384.
- [15] Z. Yuan, J. Shi, and Y. Huang, "Openecad: An efficient visual language model for editable 3d-cad design," *Computers & Graphics*, vol. 124, p. 104048, 2024.
- [16] A. C. Doris, M. F. Alam, A. H. Nobari, and F. Ahmed, "Cad-coder: An open-source vision-language model for computer-aided design code generation," *arXiv preprint arXiv:2505.14646*, 2025.
- [17] Y. Guan, X. Wang, X. Ming, J. Zhang, D. Xu, and Q. Yu, "Cad-coder: Text-to-cad generation with chain-of-thought and geometric reward," *arXiv preprint arXiv:2505.19713*, 2025.
- [18] X. Li, Y. Sun, and Z. Sha, "Llm4cad: Multimodal large language models for three-dimensional computer-aided design generation," *Journal of Computing and Information Science in Engineering*, vol. 25, no. 2, p. 021005, 2025.
- [19] A. Badagabettu, S. S. Yarlagadda, and A. B. Farimani, "Query2cad: Generating cad models using natural language queries," 2024. [Online]. Available: <https://arxiv.org/abs/2406.00144>
- [20] R. Wu, C. Xiao, and C. Zheng, "Deepcad: A deep generative network for computer-aided design models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6772–6782.
- [21] H. Guo, S. Liu, H. Pan, Y. Liu, X. Tong, and B. Guo, "Complexgen: Cad reconstruction by b-rep chain complex generation," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, pp. 1–18, 2022.
- [22] Y. Liu, A. Obukhov, J. D. Wegner, and K. Schindler, "Point2cad: Reverse engineering cad models from 3d point clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 3763–3772.
- [23] S. Koch, A. Matveev, Z. Jiang, F. Williams, A. Artemov, E. Burnaev, M. Alexa, D. Zorin, and D. Panozzo, "Abc: A big cad model dataset for geometric deep learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9601–9611.
- [24] Y. Wang, R. Jiao, S. S. Zhan, C. Lang, C. Huang, Z. Wang, Z. Yang, and Q. Zhu, "Empowering autonomous driving with large language models: A safety perspective," *arXiv preprint arXiv:2312.00812*, 2023.
- [25] C. Li, W. Tang, Y. Huang, S. S. Zhan, M. Hu, X. Jia, and Y. Liu, "Shedding light on vln robustness: A black-box framework for indoor lighting-based adversarial attack," *arXiv preprint arXiv:2511.13132*, 2025.
- [26] S. S. Zhan, Y. Liu, P. Wang, Z. Wang, Q. Wang, Z. Ruan, X. Shi, X. Cao, F. Yang, K. Wang *et al.*, "Sentinel: A multi-level formal framework for safety evaluation of llm-based embodied agents," *arXiv preprint arXiv:2510.12985*, 2025.
- [27] J. Jiang, F. Wang, J. Shen, S. Kim, and S. Kim, "A survey on large language models for code generation," *arXiv preprint arXiv:2406.00515*, 2024.
- [28] W. Lu, R. K. Luu, and M. J. Buehler, "Fine-tuning large language models for domain adaptation: Exploration of training strategies, scaling, model merging and synergistic capabilities," *npj Computational Materials*, vol. 11, no. 1, p. 84, 2025.
- [29] P. Mohapatra, A. Pandey, X. Zhang, and Q. Zhu, "Can LLMs understand unvoiced speech? exploring EMG-to-text conversion with LLMs," in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, 2025, pp. 703–712.
- [30] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," *arXiv preprint arXiv:1908.10084*, 2019.
- [31] M. Douze, A. Guzhva, C. Deng, J. Johnson, G. Szilvasy, P.-E. Mazaré, M. Lomeli, L. Hosseini, and H. Jégou, "The faiss library," *arXiv preprint arXiv:2401.08281*, 2024.
- [32] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, "Training language models to follow instructions with human feedback," *Advances in neural information processing systems*, vol. 35, pp. 27 730–27 744, 2022.
- [33] W. Feng, L. Wang, T. Wei, J. Zhang, C. Gao, S. Zhan, P. Lv, and W. Dong, "Token buncher: Shielding llms from harmful reinforcement learning fine-tuning," *arXiv preprint arXiv:2508.20697*, 2025.
- [34] Z. Zhang, C. Zheng, Y. Wu, B. Zhang, R. Lin, B. Yu, D. Liu, J. Zhou, and J. Lin, "The lessons of developing process reward models in mathematical reasoning," *arXiv preprint arXiv:2501.07301*, 2025.
- [35] S. S. Zhan, P. Wang, Q. Wu, Y. Wang, R. Jiao, C. Huang, and Q. Zhu, "Model-based reward shaping for adversarial inverse reinforcement learning in stochastic environments," *arXiv preprint arXiv:2410.03847*, 2024.
- [36] 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.
- [37] M. H. Daniel Han and U. team, "Unslloth," 2023. [Online]. Available: <http://github.com/unslothai/unslloth>
- [38] 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.
- [39] T. Laughlin, "pyocct – python bindings for opencascade via pybind11," 2020, <https://github.com/trelau/pyOCCT>.
- [40] Z. Zhou, J. Han, L. Du, N. Fang, L. Qiu, and S. Zhang, "Cad-judge: Toward efficient morphological grading and verification for text-to-cad generation," *arXiv preprint arXiv:2508.04002*, 2025.