

Idea23D: Collaborative LMM Agents Enable 3D Model Generation from Interleaved Multimodal Inputs

Junhao Chen^{1,2*}, Xiang Li^{3*}, Xiaojun Ye⁴, Chao Li⁵, Zhaoxin Fan^{6†}, Hao Zhao^{1†}

¹Institute for AI Industry Research (AIR), Tsinghua University,

²Tsinghua Shenzhen International Graduate School, Tsinghua University,

³School of Software and Microelectronics, Peking University,

⁴College of Computer Science, Zhejiang University,

⁵College of Computer Science and Technology, Harbin Engineering University,

⁶Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing,
School of Artificial Intelligence, Beihang University

Abstract

With the success of 2D diffusion models, 2D AIGC content has already transformed our lives. Recently, this success has been extended to 3D AIGC, with state-of-the-art methods generating textured 3D models from single images or text. However, we argue that current 3D AIGC methods still don't fully unleash human creativity. We often imagine 3D content made from multimodal inputs, such as what it would look like if my pet bunny were eating a doughnut on the table. In this paper, we explore a novel 3D AIGC approach: generating 3D content from IDEAs. An IDEA is a multimodal input composed of text, image, and 3D models. To our knowledge, this challenging and exciting 3D AIGC setting has not been studied before. We propose the new framework Idea23D, which combines three agents based on large multimodal models (LMMs) and existing algorithmic tools. These three LMM-based agents are tasked with prompt generation, model selection, and feedback reflection. They collaborate and critique each other in a fully automated loop, without human intervention. The framework then generates a text prompt to create 3D models that align closely with the input IDEAs. We demonstrate impressive 3D AIGC results that surpass previous methods. To comprehensively assess the 3D AIGC capabilities of Idea23D, we introduce the Eval3DAIGC-198 dataset, containing 198 multimodal inputs for 3D generation tasks. This dataset evaluates the alignment between generated 3D content and input IDEAs. Our user study and quantitative results show that Idea23D significantly improves the success rate and accuracy of 3D generation, with excellent compatibility across various LMM, Text-to-Image, and Image-to-3D models. Code and dataset are available at <https://idea23d.github.io/>.

* Indicates Equal Contribution. † Indicates Corresponding Author, email to zhaohao@air.tsinghua.edu.cn and zhaoxinf@buaa.edu.cn

1 Introduction

Recently the success of 2D AIGC foundation models (Rombach et al., 2022; Podell et al., 2023; Zhang et al., 2023c; Shi et al., 2023; Liu et al., 2023f; OpenAI, 2023b,c; Yang et al., 2023b; Chen et al., 2023b; Labs, 2024) has been translated to the 3D domain (Liu et al., 2023d; Long et al., 2023; Poole et al., 2022; Lin et al., 2023; Voleti et al., 2024; Xu et al., 2024b; Tochilkin et al., 2024; Yang et al., 2024). However, state-of-the-art models take RGB images or text prompts as inputs, which still fails to match the (wildest) creativity of humanity. Arguably, the nature of creativity is connecting (seemingly unrelated) dots that share intrinsic harmony. So we propose a novel 3D AIGC setting in which all prior arts fail: generating textured 3D models from IDEAs. The formal definition of an IDEA is an interleaved sequence of multi-modal inputs, covering modalities like text, images and 3D models. We show some typical IDEAs in Fig. 1, which contain a text prompt and images or 3D models. We use rendered 4 images to represent the 3D model in Fig. 1. This kind of IDEAs come into our minds now and then in the daily life: Someone takes a doughnut into the room and looks at her pet rabbit, imaging what it may look like if the rabbit is eating the doughnut using front paws. This is the moment that creativity happens but as far as we know, no existing 3D AIGC foundation models can take this kind of IDEAs as input. More different types of IDEAs are shown in Fig. 3 and Fig. 8.

Existing methods in text-based 3D model generation (Wang et al., 2024b; Poole et al., 2022; Lin et al., 2023; Liu et al., 2023d; Qian et al., 2023; Yang et al., 2024), known as Text-to-3D (T-2-3D), have made progress in certain aspects such as fidelity, but they still face substantial challenges, particularly when dealing with complex and abstract interleaved multimodal inputs (IDEAs). A poten-

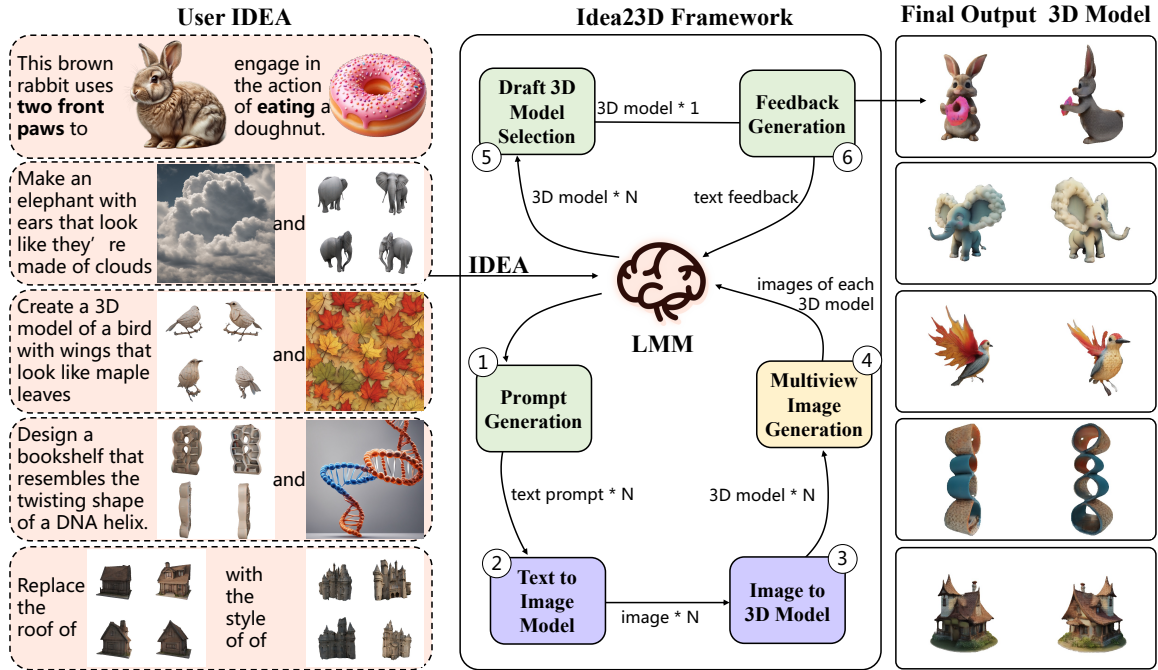


Figure 1: The *Idea23D* framework synergizes the capabilities of the Large Multimodal Model (LMM), Text-to-Image (T-2-I), and Image-to-3D (I-2-3D) models to transform complex multimodal input IDEAs into tangible 3D models. This process begins with the user articulating high-level 3D design requirements (**IDEA**). Following this, the LMM generates textual prompts (**Prompt Generation**) that are then converted into 3D models. These models are evaluated through a **Multiview Image Generation and Evaluation** process, leading to the **Selection of an Optimal 3D Model**. Subsequently, the T-2-I prompt is refined (**Feedback Generation**) using insights from the LMM. Additionally, an integrated memory module (see Sec. 3.7), meticulously records each iteration, facilitating a multimodal, iterative self-refinement cycle within the framework. Note that this procedure is fully automatic **without any human intervention**.

tial solution for adapting existing T-2-3D methods to handle IDEA inputs is converting the images and 3D models in the IDEAs into natural language descriptions. However, this approach is time-consuming and requires a certain level of expertise from the user.

Our proposal is to use a multi-agent collaboration framework. LLM (Large Language Model) agent systems (Xi et al., 2023; Reworkd, 2023; Team, 2023; Richards, 2023; Gong et al., 2023; Liu et al., 2023h; Gu et al., 2023b; Liu et al., 2023a; Chen et al., 2023c; Gou et al., 2023) have already demonstrated remarkable effectiveness in solving complex natural language processing tasks, suggesting their potential application in the T-2-3D domain. There are already some recent successful methods that leverage LLM agents for computer vision applications (Gupta and Kembhavi, 2023; Surís et al., 2023; Wei et al., 2024; Yang et al., 2023a; Huang et al., 2023; Yang et al., 2023b; Ye et al., 2024). They exploit the generic methodology of prompting LLM agents to write codes and invoke existing computer vision functions. We in-

herit this methodology but exploit LMMs (Large Multimodal Models) as agents because the visual inputs are critical in understanding IDEAs.

However, designing a LMM agent system to generate 3D models from IDEAs is not straightforward and presents its own set of challenges, especially effective integration and understanding of multimodal inputs. To tackle these challenges, we propose *Idea23D*, a framework that employs three different agents based upon the powerful LMM, for iterative self-improvement in automated 3D design and generation. Specifically, as shown in Fig. 1, *Idea23D* consists of three LMM agents (green boxes indexed by 1,5,6) acting in the roles of prompt generation, model selection and feedback reflection.

Idea23D combines the capabilities of LMM agents and other multimodal algorithmic modules (purple boxes indexed by 2,3 and yellow box indexed by 4 in Fig. 1) to generate textual prompts from interleaved user inputs (IDEAs), which are then converted into 3D models. This process involves iterative refinement, utilizing a memory

module to record each iteration and support continuous improvement. As shown in Fig. 3 and Fig. 8, *Idea23D* can generate high-quality 3D models that well align input **IDEAs in a fully automatic manner** while caption-based baselines constructed from prior T-2-3D models can hardly generate meaningful results.

Our qualitative comparisons and quantitative experiments demonstrate the effectiveness of *Idea23D*, especially in handling complex and challenging IDEA inputs. The contributions of this paper are as follows.

- (1) *Idea23D* is the first work to achieve the transformation of high-level, abstract user IDEAs (multimodal interleaved inputs) into concrete 3D models, realizing a fully automated 3D AIGC task.
- (2) Surpassing the capabilities of existing LLM agent systems in 3D AIGC, *Idea23D* demonstrates the effectiveness of LMM-based agents in improving, evaluating, and validating multimodal content for 3D model generation.
- (3) Proposes a challenging evaluation dataset **Eval3DAIGC-198** with multimodal inputs, and proves the effectiveness of *Idea23D* through comprehensive user preference studies and 3D visual caption experiments.

2 Related Works

2.1 Self-refining Agents

Our research builds on the self-refinement capability of large language models (LLMs). Recent studies show that LLMs are effective at self-refinement, when exploring unknown environments and tasks (Xi et al., 2023; Madaan et al., 2024; Pan et al., 2023; Shinn et al., 2023; Lee et al., 2023). For example, projects such as *Self-refine* (Madaan et al., 2024) and *Reflexion* (Shinn et al., 2023) utilize LLMs to iteratively critique their outputs and use feedback to improve predictions, resulting in significant performance improvements in natural language processing tasks. However, these approaches mainly excel in tasks dealing with natural language descriptions (Shridhar et al., 2020). In contrast, our *Idea23D* project employs an iterative self-refinement system based on LMM in a multimodal environment, especially for interleaved inputs of text, images, and 3D models (IDEAs), in a different way other than the traditional approach focusing solely on natural language inputs.

2.2 Large Multimodal Model

Building on the Large Language Model (LLM), the development of the Large Multimodal Model (LMM) marks an important evolution from unimodal to multimodal processing capabilities. Initial LLMs, such as the GPT family, focused on the generation of textual data, demonstrating superior capabilities in understanding and creating natural language (Brown et al., 2020; Radford et al., 2019; OpenAI, 2023d; Bai et al., 2023a,b; Zhang et al., 2023b; He et al., 2024). As the field evolves, the processing capabilities of the models expanded from pure text to include multimodal data including images, audio, and video (Ramesh et al., 2021; Radford et al., 2021). For example, CLIP (Radford et al., 2021) was the first to achieve cross-modal alignment between images and text, enabling cross-modal understanding between text and image. Some models (Liu et al., 2024; Xu et al., 2024a; Chen et al., 2024b; Bai et al., 2023c; Wang et al., 2024a), demonstrates multimodal understanding and generation of mixed text, image, and video inputs. Further, projects such as *Uni-3D* (Zhang et al., 2023d) and *SDFusion* (Cheng et al., 2023) extend this concept to the 3D design domain, enabling good 3D model understanding, generation and reconstruction. LMM has also enabled open-set scene understanding in various settings (Tian et al., 2023; Li et al., 2022, 2023; Liu et al., 2023e; Jin et al., 2023). Despite the progress made in the field of 3D understanding and generation, current LMMs can still only handle input content at the text and image level, and struggle to handle multimodal high-level inputs containing text, images, and 3D models (IDEAs in our case). In contrast, the *Idea23D* project takes a much larger step forward in multimodal 3D model generation. Our system is capable of processing not only single-modal inputs, but also composite multimodal inputs containing text, images, and 3D models at the same time.

2.3 Extensions of T-2-3D Models

There is already a large literature extending seminal T-2-3D models, including variants enabling T-2-3D models to better follow user prompts (Black et al., 2023; Chefer et al., 2023; Feng et al., 2022), refine keywords in T-2-3D prompts for better visual quality (Gu et al., 2023a), support for additional image inputs for image processing (Brooks et al., 2023; Kwar et al., 2023), support for additional 3D model inputs for 3D model processing (Cheng

et al., 2023), 3D style migration (Pan and Ke, 2023; Ma et al., 2014; Segu et al., 2020), visual concept customization (Wei et al., 2023; Kumari et al., 2023), using image generation models to generate 3D textures (Chen et al., 2023a; Zeng et al., 2024; Chen et al., 2024a; Yu et al., 2024), and more. Going even further, *Idea23D* offers users a more natural way to design and create the 3D content they want. Similar to *Visual ChatGPT* (Wu et al., 2023a), which extends *ChatGPT*'s (OpenAI, 2023e) ability to understand and generate 2D images, *Idea23D* is designed to provide a more unified and broadly applicable framework for automated 3D model design and generation. *Idea23D* extends the multimodal input and 3D model understanding and generation capabilities of the LMM (OpenAI, 2023d) using the T-2-I model (Shi et al., 2020) and I-2-3D model (Ramesh et al., 2021), as well as the multimodal input capabilities of the T-2-3D model (Wang et al., 2024b; Poole et al., 2022; Lin et al., 2023; Wang et al., 2023; Metzger et al., 2023; Tsalicoglou et al., 2023; Qian et al., 2023; Liu et al., 2023d; Long et al., 2023) and other standard algorithmic modules like multiview rendering.

3 *Idea23D* Framework

The *Idea23D* framework represents a novel approach to generate detailed 3D models from high-level, abstract multimodal inputs (IDEAs), shown in Fig. 2. It integrates three LMM-based agents and several off-the-shelf tools for agents to invoke. Specifically, three agents are responsible for prompt generation (green box indexed by 1 in Fig. 2), model selection (green box indexed by 5 in Fig. 2) and feedback generation (green box indexed by 6 in Fig. 2). Two foundation models for T-2-I (purple box indexed by 2 in Fig. 2) and I-2-3D (purple box indexed by 3 in Fig. 2) are exploited together to turn natural language prompts into textured 3D models. As shown in Tab. 1, various foundation model variants are evaluated for comprehensiveness. A unique memory module enhances the system, retaining insights from previous iterations to optimize future outputs.

The process begins with the LMM (**agent 1**) converting multimodal IDEAs into T-2-I prompts, which facilitate the creation of preliminary 3D drafts. These drafts undergo a selection process (**agent 5**) where the best 3D model is either considered finalized or subjected to further refinement based on LMM feedback (**agent 6**). The cycle con-

tinues until the model satisfies the user's IDEA (as judged by **agent 5**) or reaches a pre-set iteration limit. This framework, illustrated in our system diagram (Fig. 2), enables continuous improvement through a loop with automatically generated feedbacks.

3.1 Multimodal IDEA Input

The user-provided IDEA X encapsulates the overarching 3D modeling requirement, represented by a multimodal input set composed of text, images, and 3D models. $X = \{T, I, M\}$, where T is a suite of textual directives encompassing descriptive phrases, keywords, and design specifications that **may refer to both** 2D and 3D information. I is an assortment of images such as reference shots, diagrams, or related illustrations. M is a compilation of 3D models, including pre-existing constructions or particular design ingredients furnished by the user. Each element of X reflects a facet of the user's design intent, and their aggregation forms a comprehensive IDEA. This input specification is designed to capture the user's intent in a multi-faceted and multimodal manner, laying the ground for the subsequent procedural stages, including initial prompt formulation and 3D model synthesis. As mentioned before, designing from this kind of abstract IDEAs is of great needs and our method is the first to fulfill this need.

3.2 Initial Prompt Generation

Recall that, technically, the *Idea23D* framework is tailored to convert complex multimodal user inputs X into textual prompts for 3D model generation. Specifically, it employs the Large Multimodal Model (LMM) to comprehend and articulate these inputs into a format digestible by the Text-to-3D (T-2-3D) model. Addressing the high dimensionality of 3D models, *Idea23D* leverages LMM's text and image processing strengths by representing 3D models M as multi-view image sets I' via a conversion function $\text{CM2I}(\ast)$:

$$I' = \text{CM2I}(M) \quad (1)$$

Specifically, Function $\text{CM2I}(\ast)$ render each 3D model into six images, depicting the model from various perspectives: front, back, left, right, top, and bottom. The aggregation of these images I' with the original image set I and textual components yields an augmented IDEA X' :

$$X' = \{T, I' \cup I\} \quad (2)$$

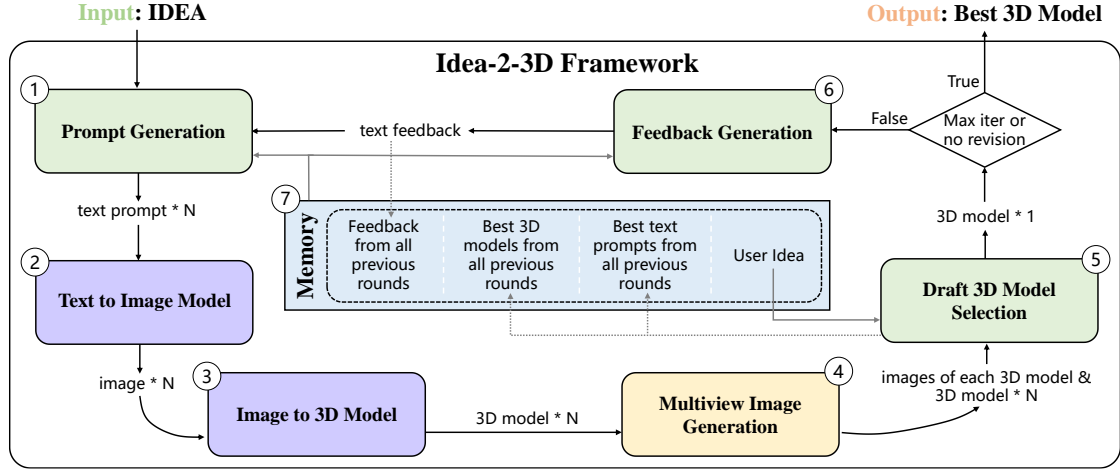


Figure 2: Overview of the *Idea23D* framework: LMM agents iteratively refine T-2-3D prompts to maximize the T-2-3D model’s potential for user input IDEA. Green rounded rectangles represent steps completed by LMM agents. Purple ones denote T-2-3D modules, including T-2-I and I-2-3D models. The yellow rounded rectangle represents the algorithm for rendering 3D models into multi-view images. Blue represents the memory module, storing feedback, the best 3D model, and the best text prompt. This cycle is **fully automated** by LMM agents, requiring no human intervention.

The **LMM agent 1 for prompt generation** receives the IDEA X' and generates N specific descriptive instructions $\{P_0, P_1, \dots, P_{N-1}\}$:

$$P = f_{\text{LMM}}(X', p_{\text{gen}}) \quad (3)$$

Here, p_{gen} represents a prompt facilitating the generation of T-2-3D prompts¹. This enables *Idea23D* to extract and interpret not only direct descriptions (e.g. X' with only T is also acceptable) but also high-level concepts and intermixed modalities within the IDEA, such as images and 3D models.

In subsequent iterations as denoted by iter , each T-2-3D prompt in the set of P^{iter} is used as input to the T-2-3D model to produce draft 3D models $D_i^{\text{iter}} = T23D(P_i^{\text{iter}})$, iteratively refining until the output aligns with the user’s intents.

3.3 3D Model Generation

Idea23D transforms the set of N text prompts $\{P_0, P_1, \dots, P_{N-1}\}$ into an equivalent set of N 3D models $\{D_0, D_1, \dots, D_{N-1}\}$ utilizing T-2-3D models. As outlined in Sec. 3.2, T-2-3D encompasses a two-step conversion process: initial Text-to-Image (T-2-I) generation followed by Image-to-3D (I-2-3D) generation. To improve 3D model creation quality, we employ a background removal module² on T-2-I outputs before I-2-3D processing.

¹Note here p_{gen} and P are prompts for LMMs and T-2-3D models, respectively. All prompts are in our code.

²<https://github.com/danielgatis/rembg>

The T-2-3D function is:

$$T23D(P_i) = I23D \circ \text{rembg} \circ T2I(P_i) \quad (4)$$

In detail: **(1) Text-to-Image model $T2I(*)$** : Each prompt P_i generates a 2D image $G_i = T2I(P_i)$. **(2) Background Removal $\text{rembg}(*)$** : The generated image G_i undergoes background removal $G'_i = \text{rembg}(G_i)$, enhancing the focus on foreground (i.e., the primary subject). **(3) Image-to-3D model $I23D(*)$** : The refined image G'_i is input into I-2-3D, producing the 3D model $D_i = I23D(G'_i)$. This methodology from text prompt to 3D model via an intermediary image phase, particularly with background removal, ensures a more accurate and intent-aligned 3D reconstruction, elevating the overall quality of the generated models. These functions are tools invoked by **LMM agent 1**.

3.4 Draft 3D Model Selection

Then the **LMM agent 5 for model selection** in *Idea23D* selects the superior draft 3D model D_{best} from the generated set $\{D_0, D_1, \dots, D_{N-1}\}$ based on the fidelity and relevance to the user’s IDEA. This critical step filters out subpar models, ensuring high-quality iterative generations.

$$D_{\text{best}} = f_{\text{select}}(D_i, X', p_{\text{select}}) \quad (5)$$

Here, p_{select} is the prompt for the Large Multimodal Model (LMM), guiding the selection of

the best draft 3D model. f_{select} uses specific few-shot prompts for the LMM. It renders six views of each 3D model, combines them into a single image, and then inputs this image into the LMM. The LMM then selects the 3D model that most closely matches the user’s IDEA input to serve as the draft model for the current iteration. This mechanism compensates for the discrepancy observed when high-quality T-2-I outputs do not necessarily translate into satisfactory I-2-3D models. By assessing the semantic coherence and visual quality of N similar draft models, the LMM identifies the best one. This comparative analysis, akin to a *find the difference* task, is impossible to achieve by conventional techniques, yet state-of-the-art LMMs like GPT-4V (OpenAI, 2023a), GPT-4o (OpenAI, 2024) and InternVL (Chen et al., 2024b) have demonstrated reliable performance in this selection process.

3.5 Feedback Generation

After identifying the best draft model D_{best} , the **LMM agent 5** decides on whether to finalize this model as the result D^* or proceed with refinement. In the latter case, the goal is to generate textual feedback F to guide enhancements for D_{best} . This decision hinges on whether the iteration count exceeds a maximum threshold T or if the agent believes no further modifications are needed.

$$F = f_{LMM}(D_{best}, X', m, p_{fb}) \quad (6)$$

Here, p_{fb} is the LMM prompt for feedback generation, and m denotes the Memory module (discussed in Sec. 3.7). This **LMM agent 6 for feedback generation** assesses discrepancies between D_{best} and the user IDEA X' , summarizing key inconsistencies. D_{best} is converted into multi-view images using the CM2I(*) function (from Sec. 3.2), aiding the LMM in pinpointing and suggesting specific enhancements. This step is crucial for refining the 3D model. Our experience suggests that clearly defining the aspects for review in p_{fb} significantly enhances the quality of the resultant 3D model.

3.6 Revised Prompt Generation

In the final stage of each iteration (noted as *iter*), the **LMM agent 1** comes to the stage again for *Revised Prompt Generation*. It uses textual feedback F and the memory module m to create N refined 3D model generation prompts $\{P_0^{iter+1}, P_1^{iter+1}, \dots, P_{n-1}^{iter+1}\}$. This step aims to enhance 3D models generated in the next iteration.

$$P_i^{iter+1} = f_{LMM}(F, m, X', p_{gen}) \quad (7)$$

Here, p_{gen} is the LMM prompt for I-2-3D prompt generation, which is the same with Eq. 3, despite the inputs are augmented with F and m . Note an LMM agent can readily handle different inputs by prompting it that there are two different cases: the initialization case without feedback and memory and the refinement case with feedback and memory. Agent 1 leverages the information stored in m and the previous iteration’s feedback F to generate improved prompts that effectively address the issues identified in F . For instance, if feedback F indicates specific visual inaccuracies in the best-to-date model, the revised prompts will focus on rectifying these details through enhanced descriptions.

3.7 Memory Module

The memory module m is integral to the *Idea23D* framework, serving as a repository for data accrued over the iterative process, as shown by the blue rectangle in Fig. 2. It stores feedback, selected draft 3D models, and corresponding text prompts in a structured 3D model-image-text sequence, enabling the LMM to leverage past experiences and insights gained during previous iterations.

$$m_t = \{P_0^*, D_0^*, F_0, \dots, P_{t-1}^*, D_{t-1}^*, F_{t-1}\} \quad (8)$$

In this representation, $P_{iter}^*, D_{iter}^*, F_{iter}$ represent the optimal text prompt, 3D model, and textual feedback from each *iter*th iteration. The memory module m aids the agent in identifying specific T-2-3D output traits, such as misunderstood keywords. This knowledge is then integrated into generating refined 3D model prompts, enhancing the precision and adaptability of model generation. If T-2-3D struggles with certain design aspects, m guides subsequent iterations to optimize prompts more effectively, ensuring continuous improvement in *Idea23D*, thereby increasing its alignment with complex IDEAs.

4 Experiments

4.1 Experimental Setup

In our early experiments, the LMMs used were GPT-4V (OpenAI, 2023a) and LLaVA (Liu et al., 2023b). Fig. 3, Fig. 7, Fig. 9, Fig. 10, and Fig. 11 use GPT-4V as the LMM, DALLE (OpenAI, 2023c) as the T-2-I model, and zero123 (Liu

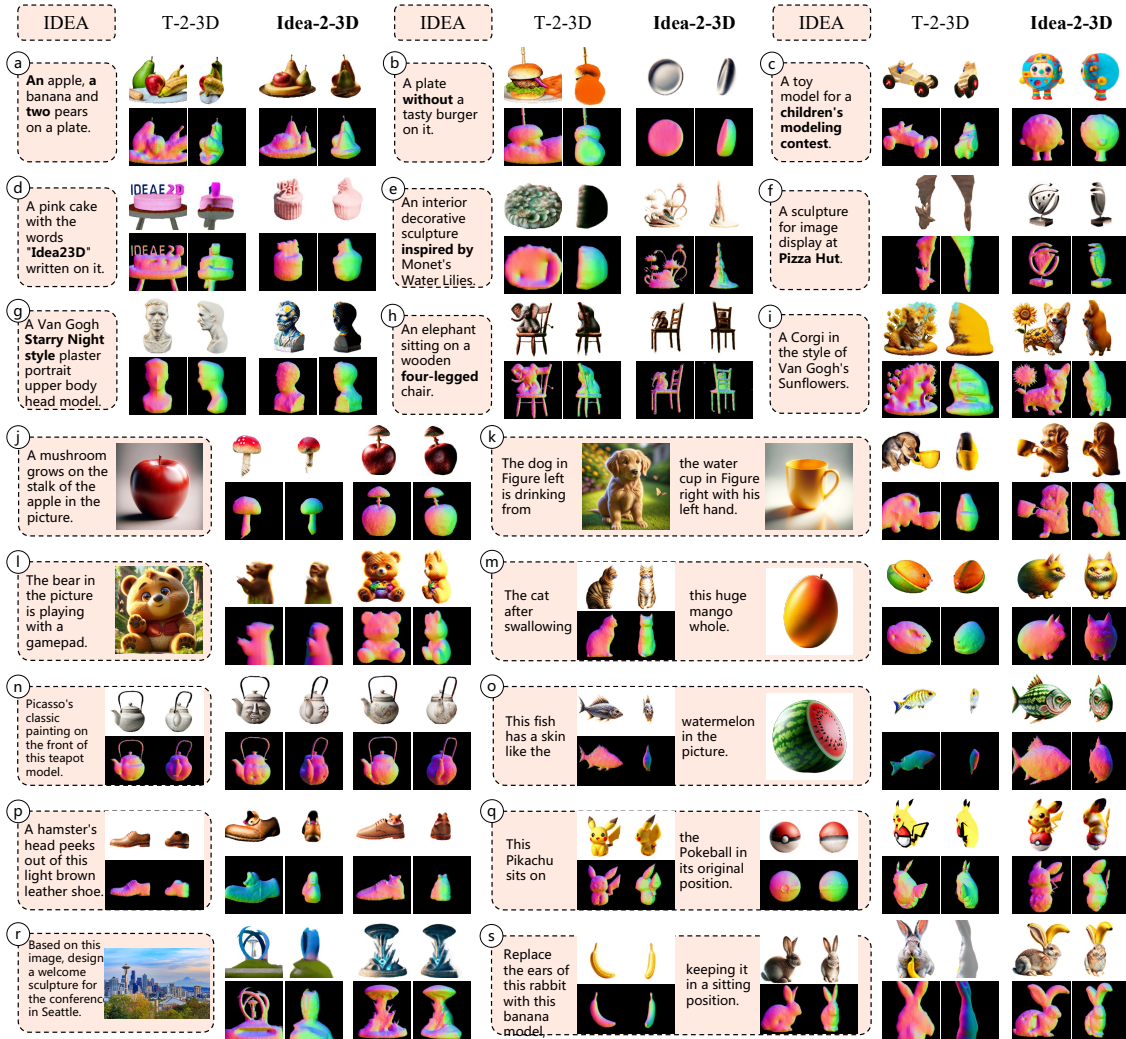


Figure 3: Overview of 3D models generated from various types of multimodal IDEA inputs supported by *Idea23D*. The light red box on the left is the user input IDEA containing text, images and 3D models. In the center are the baseline results generated directly from the same T-2-I model with caption-based T-2-I prompt (see Sec. 4.1). The model on the right is the result generated by iteratively self-refining the T-2-I prompts with *Idea23D*. Comparison with more existing methods is shown in Fig. 7.

et al., 2023d) as the I-2-3D model. The results of the user study, shown in Tab. 3, were obtained from these early experiments. With the emergence of new LMM, T-2-I, and I-2-3D models, we have also tested *Idea23D* with these new methods. Fig. 1 and Fig. 8 use GPT-4o as the LMM, FLUX as the T-2-I model, and InstantMesh as the I-2-3D model. The quantitative results of the ablation study, presented in Tab. 2, were also based on these new models. This further demonstrates that *Idea23D* exhibits excellent compatibility with the development of LMMs, T-2-I, and I-2-3D models.

T-2-3D baseline. Our first baseline is caption-based. Since no former 3D AIGC methods can be used for IDEA input, we convert image inputs and 3D model inputs (after multiview rendering) into

textual descriptions by captioning them with the LMM. In Fig. 3, Tab. 1 and Tab. 3, this baseline is called T-2-3D. All T-2-3D baselines only use LMM for generating captions and do not perform iterative refinement.

Ours w/o iterative refinement. To demonstrate the impact of an iterative self-refinement design, we construct an *Idea23D* variant with only one iteration. Compared with caption-based baselines, this one features multiple prompt generation and best 3D model selection. We presented the results of this part in our early user study, see Sec. 4.4.

4.2 Evaluation Dataset

In our experiments, we found that there is a lack of methods to align Text, Image, and 3D in the

Comparison stage: T-2-3D vs.			Avg. Iter.	CLIP \uparrow		ULIP-2 \uparrow	
LMM	T-2-I	I-2-3D		T-2-3D	Idea23D	T-2-3D	Idea23D
GT prompt	FLUX	InstantMesh	-	0.3152	-	0.3134	-
GPT-4o	FLUX	InstantMesh	1.49	0.3003	0.3078	0.2733	0.2917
InternVL2.5	FLUX	InstantMesh	2.56	0.2896	0.3035	0.2379	0.2756
LLaVA-CoT	FLUX	InstantMesh	4.67	0.2783	0.2826	0.2108	0.2125
Text Only	FLUX	InstantMesh	-	0.2745	-	0.2056	-
GT prompt	FLUX	Hunyuan3D	-	0.3085	-	0.3057	-
GPT-4o	FLUX	Hunyuan3D	1.41	0.2943	0.3010	0.2675	0.2869
InternVL2.5	FLUX	Hunyuan3D	2.68	0.2870	0.2970	0.2534	0.2783
LLaVA-CoT	FLUX	Hunyuan3D	4.48	0.2734	0.2768	0.2189	0.2273
Text Only	FLUX	Hunyuan3D	-	0.2700	-	0.2115	-
GT prompt	SDXL	InstantMesh	-	0.2972	-	0.2684	-
GPT-4o	SDXL	InstantMesh	2.02	0.2845	0.3001	0.2302	0.2599
InternVL2.5	SDXL	InstantMesh	3.49	0.2810	0.2822	0.2196	0.2211
LLaVA-CoT	SDXL	InstantMesh	4.53	0.2707	0.2726	0.1961	0.1991
Text Only	SDXL	InstantMesh	-	0.2680	-	0.1979	-
GT prompt	SDXL	Hunyuan3D	-	0.2941	-	0.2479	-
GPT-4o	SDXL	Hunyuan3D	2.14	0.2782	0.2903	0.2118	0.2404
InternVL2.5	SDXL	Hunyuan3D	3.94	0.2725	0.2828	0.2143	0.2198
LLaVA-CoT	SDXL	Hunyuan3D	4.32	0.2711	0.2735	0.1911	0.1956
Text Only	SDXL	Hunyuan3D	-	0.2663	-	0.1924	-

Table 1: We conducted experiments on the **Eval3DAIGC-198** dataset with the configuration of generating one image per prompt ($num_{img} = 1$), three prompts per round ($num_{draft} = 3$), and up to five iteration rounds ($max_{iters} = 5$). We used GPT-4o (OpenAI, 2024) instead of GPT-4-Vision-Preview (OpenAI, 2023a) from earlier studies. The models used include FLUX.1-dev (Labs, 2024), SD-XL 1.0 with refinement, Hunyuan3D-1.0 (Yang et al., 2024), InternVL2.5-78B (Chen et al., 2024b), and LLaVA-CoT-11B (Xu et al., 2024a).

evaluation of LMM’s 3D AIGC capabilities (Liu et al., 2023g; Andriushchenko et al., 2024; Guo et al., 2024; Wu et al., 2023b; Zhang et al., 2023a). Therefore, following the evaluation practices of Parti (Yu et al., 2022), we constructed a dataset for evaluating 3D AIGC tasks, called **Eval3DAIGC-198**, which involves 198 different multimodal interleaved inputs of IDEA, including examples shown in Fig. 1, Fig. 3 and Fig. 8. The distribution and examples of the dataset can be found in Appendix A. The cases in the evaluation dataset consist of a text prompt, which may include images or 3D models. Since the results of 3D AIGC are difficult to represent by a specific 3D model as a standard answer, we annotated a description of the 3D model to be generated for each case based on the textual instructions, which serves as the Ground Truth.

4.3 3D-Caption Quantitative Results

To ensure a fair and comprehensive comparison of the alignment between user inputs and the final generated 3D models, we employed various methods to convert the generated 3D models into textual descriptions. These textual features were then compared with the handwritten 3D annotations in the

Eval3DAIGC-198 evaluation dataset.

(1) CLIP (Radford et al., 2021) Metric. We render the 3D models from four views (front, back, left, right) and calculate the CLIP similarity between the text description and the rendered images. The CLIP similarity is:

$$S_{\text{CLIP}}(T, I) = \frac{1}{4} \sum_{i=1}^4 \frac{E_T \cdot E_{I_i}}{\|E_T\| \|E_{I_i}\|} \quad (9)$$

where E_T and E_{I_i} are the embeddings extracted from the CLIP model for the text description T and the rendered images I_i , respectively. Here, \cdot denotes the dot product and $\|\cdot\|$ is the L2 norm.

(2) ULIP-2 (Xue et al., 2024) Metric. ULIP-2 is a tri-modal pre-training framework that generates text descriptions for 3D shapes without human annotations. It evaluates the alignment between 3D models and texts by comparing 3D shape features with corresponding descriptions in the Eval3DAIGC-198 dataset.

The quantitative experimental results are shown in Tab. 1. The "GT prompt" row represents the scores from generating 3D models using manually annotated 3D captions with text-to-image and

image-to-3D methods. "Text Only" refers to using only textual instructions from the dataset as the baseline. The "T-2-3D" column shows results where LMM generates descriptions of images and 3D models, which are then combined with dataset instructions to generate 3D models. "Idea23D" represents our proposed method, and "Avg. Iter." shows the average number of optimization iterations for Idea23D.

The results demonstrate that Idea23D improves success rate and accuracy in 3D AIGC tasks, producing outputs closer to ground truth (GT prompt) with fewer iterations. Using GPT-4o as LMM, only two iterations are needed to generate realistic 3D models. Fig. 7 compares Idea23D with current commercial models, showing that Idea23D achieves new 3D AIGC capabilities not possible with existing models.

4.4 User Study

We presented the results of the caption-based T-2-3D baseline and *Idea23D* to the participants in our user study. Our evaluation results are detailed in Tab. 3, which compares the caption-based T-2-3D baseline with *Idea23D* using the model specified in each row. We asked users to assess which model (T-2-3D model, first-round results of *Idea23D*, and final-round results of *Idea23D*) was more satisfactory, and to evaluate whether each 3D model complied with the IDEA. Detailed User Study results and explanations can be found in Appendix C.

4.5 Visualization of Self-iterative Refinement

Fig. 9 showcase the evolution and refinement of 3D models at different *Idea23D* stages and how the caption-based baseline works. Fig. 10 shows the iterative self-improvement process of a case in *Idea23D*. These visual representations illustrate the framework’s effectiveness in aligning with user IDEAs and the incremental improvements achieved through its iterative process.

4.6 Ablation Study

We conducted ablation experiments using the first 38 cases from the dataset, with LMM using GPT-4o, T2I using FLUX, and I23D using InstantMesh. Iteration rounds and configurations are consistent with Tab. 1. Results are shown in Tab. 2. The case in Fig. 4 uses GPT-4V, SD-XL, and zero123.

As shown in the Pokemon case (Fig. 4), removing key modules slows convergence, while the full *Idea23D* model generates a satisfactory result

within 3 iterations. The quantitative study in Tab. 2 reveals: **(1)** The memory module prevents quality divergence, **(2)** The LMM feedback agent accelerates convergence, **(3)** Removing stored information from previous models lowers the quality limit after convergence. Additionally, ablation results show that removing the memory module leads to increasing deviation from user inputs during iterations. Feedback improves convergence speed, and retaining prior models accelerates the process. The ULIP score for 3D models generated with real prompts is 0.3213, with Idea23D reaching 0.3208 after five iterations, while a standard text-to-image pipeline achieves only 0.2830.

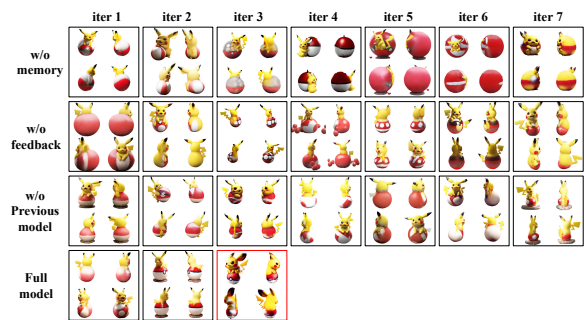


Figure 4: Key module ablation across iterations. Note that the full model reaches satisfactory results (judged by the LMM agent 5 in Fig. 2) at iteration 3. The illustrated experiment has no maximum iteration limit.

	Avg. Iter.	ULIP \uparrow				
		Iter 1	Iter 2	Iter 3	Iter 4	Iter 5
w/o memory	1.58	0.2953	0.2216	0.1902	0.2139	0.1960
w/o feedback	1.87	0.2930	0.2780	0.2759	0.2773	0.2811
w/o prev. model	1.78	0.2977	0.3031	0.3066	0.3012	0.2974
full model	1.49	0.3021	0.3056	0.3108	0.3146	0.3208
T-2-3D	-	0.2830	-	-	-	-
GT prompt	-	0.3213	-	-	-	-
Text Only	-	0.2717	-	-	-	-

Table 2: Ablation study results with the same configuration as in Tab. 1.

5 Conclusion

Idea23D, utilizing an LMM agent collaboration framework, revolutionizes 3D AIGC by automating the creation of models from high-level, interleaved multimodal user inputs (IDEAs). This innovative system excels in integrating text, images, and 3D models, underpinned by a unique iterative process that enhances model coherence and visual alignment with IDEAs. User studies underscore its superiority in satisfaction and comparative quality, marking *Idea23D* as a significant advancement in 3D AIGC and a benchmark for future design tools.

6 Limitations

Idea23D effectively improves the alignment between 3D generation models and user intent, but it still relies on LMM and I-2-3D models. As shown in Tab. 1, LMM significantly impacts the 3D generation results. However, commercial models like GPT-4V can already generate image prompts based on feedback very well. On the other hand, open-source LMMs like LLaVA still have significant shortcomings in image understanding capabilities.

According to our experiments, the main bottleneck of *Idea23D* at this stage lies in the Image-to-3D step. Even the most advanced Image-to-3D models can fail. In the worst-case scenario with a very low probability, the final model output when *Idea23D* reaches the maximum iteration may still not meet user requirements. Nonetheless, *Idea23D* can ensure that the final selected model is the most aligned with user input among all generated 3D models throughout the iterations.

Acknowledgments

This work was supported by the Fundamental Research Funds for the Central Universities under Grant Number KG16336301 and the China Postdoctoral Science Foundation under Grant Number 2024M764093.

References

- Maksym Andriushchenko, Alexandra Souly, Mateusz Dziemian, Derek Duenas, Maxwell Lin, Justin Wang, Dan Hendrycks, Andy Zou, Zico Kolter, Matt Fredrikson, et al. 2024. Agentharm: A benchmark for measuring harmfulness of llm agents. *arXiv preprint arXiv:2410.09024*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023b. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023c. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. 2023. Training diffusion models with reinforcement learning. In *The Twelfth International Conference on Learning Representations*.
- Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. 2023. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–10.
- Dave Zhenyu Chen, Yawar Siddiqui, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. 2023a. Text2tex: Text-driven texture synthesis via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18558–18568.
- Junhao Chen, Peng Rong, Jingbo Sun, Chao Li, Xiang Li, and Hongwu Lv. 2023b. *Soulstyler: Using large language model to guide image style transfer for target object*. *Preprint*, arXiv:2311.13562.
- Mingjin Chen, Junhao Chen, Xiaojun Ye, Huan-ang Gao, Xiaoxue Chen, Zhaoxin Fan, and Hao Zhao. 2024a. Ultraman: Single image 3d human reconstruction with ultra speed and detail. *arXiv preprint arXiv:2403.12028*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024b. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Zhipeng Chen, Kun Zhou, Beichen Zhang, Zheng Gong, Wayne Xin Zhao, and Ji-Rong Wen. 2023c. Chatcot: Tool-augmented chain-of-thought reasoning on chat-based large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14777–14790.
- Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. 2023. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4456–4465.
- Weixi Feng, Xuehai He, Tsu-Jui Fu, Varun Jampani, Arjun Reddy Akula, Pradyumna Narayana, Sugato Basu, Xin Eric Wang, and William Yang Wang. 2022. Training-free structured diffusion guidance for compositional text-to-image synthesis. In *The Eleventh*

- International Conference on Learning Representations*.
- Ran Gong, Qiuyuan Huang, Xiaojian Ma, Hoi Vo, Zane Durante, Yusuke Noda, Zilong Zheng, Song-Chun Zhu, Demetri Terzopoulos, Li Fei-Fei, et al. 2023. Mindagent: Emergent gaming interaction. *arXiv preprint arXiv:2309.09971*.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Nan Duan, Weizhu Chen, et al. 2023. Critic: Large language models can self-correct with tool-interactive critiquing. In *Second Agent Learning in Open-Endedness Workshop*.
- Jindong Gu, Zhen Han, Shuo Chen, Ahmad Beirami, Bailan He, Gengyuan Zhang, Ruotong Liao, Yao Qin, Volker Tresp, and Philip Torr. 2023a. A systematic survey of prompt engineering on vision-language foundation models. *arXiv preprint arXiv:2307.12980*.
- Yu Gu, Xiang Deng, and Yu Su. 2023b. Don't generate, discriminate: A proposal for grounding language models to real-world environments. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4928–4949.
- Hongcheng Guo, Wei Zhang, Junhao Chen, Yaonan Gu, Jian Yang, Junjia Du, Binyuan Hui, Tianyu Liu, Jianxin Ma, Chang Zhou, and Zhoujun Li. 2024. *Iw-bench: Evaluating large multimodal models for converting image-to-web*. Preprint, arXiv:2409.18980.
- Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. 2023. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>.
- Tanmay Gupta and Aniruddha Kembhavi. 2023. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14953–14962.
- Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava, and Ser-Nam Lim. 2024. Ma-lmm: Memory-augmented large multimodal model for long-term video understanding. *arXiv preprint arXiv:2404.05726*.
- Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. 2023. An embodied generalist agent in 3d world. *arXiv preprint arXiv:2311.12871*.
- Bu Jin, Xinyu Liu, Yupeng Zheng, Pengfei Li, Hao Zhao, Tong Zhang, Yuhang Zheng, Guyue Zhou, and Jingjing Liu. 2023. Adapt: Action-aware driving caption transformer. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7554–7561. IEEE.
- Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6007–6017.
- Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1931–1941.
- DeepFloyd Lab. 2023. *Deepfloyd if*. Accessed: October 28, 2023.
- Black Forest Labs. 2024. *Introducing flux.1 tools*. Accessed: 2024-12-15.
- Ariel Lee, Cole Hunter, and Nataniel Ruiz. 2023. Platypus: Quick, cheap, and powerful refinement of llms. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following*.
- Pengfei Li, Beiwen Tian, Yongliang Shi, Xiaoxue Chen, Hao Zhao, Guyue Zhou, and Ya-Qin Zhang. 2022. Toist: Task oriented instance segmentation transformer with noun-pronoun distillation. *Advances in Neural Information Processing Systems*, 35:17597–17611.
- Yang Li, Xiaoxue Chen, Hao Zhao, Jiangtao Gong, Guyue Zhou, Federico Rossano, and Yixin Zhu. 2023. Understanding embodied reference with touch-line transformer. In *ICLR*.
- Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309.
- Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. 2023a. Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv preprint arXiv:2304.11477*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023b. Improved baselines with visual instruction tuning. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction Following*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen, Chong Zeng, Jiayuan Gu, and Hao Su. 2023c. One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. *arXiv preprint arXiv:2311.07885*.

- Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 2023d. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309.
- Xinyu Liu, Beiwen Tian, Zhen Wang, Rui Wang, Kehua Sheng, Bo Zhang, Hao Zhao, and Guyue Zhou. 2023e. Delving into shape-aware zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2999–3009.
- Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. 2023f. Syncdreamer: Generating multiview-consistent images from a single-view image. In *The Twelfth International Conference on Learning Representations*.
- Zhiwei Liu, Weiran Yao, Jianguo Zhang, Le Xue, Shelby Heinecke, Rithesh Murthy, Yihao Feng, Zeyuan Chen, Juan Carlos Niebles, Devansh Arpit, et al. 2023g. Bolaa: Benchmarking and orchestrating llm-augmented autonomous agents. *arXiv preprint arXiv:2308.05960*.
- Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2023h. Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization. *arXiv preprint arXiv:2310.02170*.
- Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. 2023. [Wonder3d: Single image to 3d using cross-domain diffusion](#). *Preprint*, arXiv:2310.15008.
- Chongyang Ma, Haibin Huang, Alla Sheffer, Evangelos Kalogerakis, and Rui Wang. 2014. Analogy-driven 3d style transfer. In *Computer Graphics Forum*, volume 33, pages 175–184. Wiley Online Library.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2024. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36.
- Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. 2023. Latent-nerf for shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12663–12673.
- OpenAI. 2023a. [Chatgpt can now see, hear, and speak](#). Accessed: October 24, 2023.
- OpenAI. 2023b. [Dall-e 2](#). Accessed: October 27, 2023.
- OpenAI. 2023c. [Dall-e 3](#). Accessed: October 27, 2023.
- OpenAI. 2023d. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2023e. [Openai-gpt-3.5documentation](#). Accessed: October 21, 2023.
- OpenAI. 2024. [Hello gpt-4o](#). Accessed: 2024-12-15.
- Bo Pan and YuKai Ke. 2023. Efficient artistic image style transfer with large language model (llm): A new perspective. In *2023 8th International Conference on Communication and Electronics Systems (ICCES)*, pages 1729–1732. IEEE.
- Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. 2023. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188*.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*.
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*.
- Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. 2023. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. In *The Twelfth International Conference on Learning Representations*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR.
- Reworkd. 2023. Agentgpt. <https://github.com/reworkd/AgentGPT>.
- Toran Bruce Richards. 2023. Auto-gpt: An autonomous gpt-4 experiment.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.

- Mattia Segu, Margarita Grinvald, Roland Siegwart, and Federico Tombari. 2020. 3dsnet: Unsupervised shape-to-shape 3d style transfer. *arXiv preprint arXiv:2011.13388*.
- Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. 2023. Mvdream: Multi-view diffusion for 3d generation. In *The Twelfth International Conference on Learning Representations*.
- Zhan Shi, Xu Zhou, Xipeng Qiu, and Xiaodan Zhu. 2020. Improving image captioning with better use of caption. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7454–7464.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Cote, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. 2020. Alworld: Aligning text and embodied environments for interactive learning. In *International Conference on Learning Representations*.
- Dídac Surís, Sachit Menon, and Carl Vondrick. 2023. Vipergpt: Visual inference via python execution for reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11888–11898.
- Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. 2024. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. *arXiv preprint arXiv:2402.05054*.
- X Team. 2023. Xagent: An autonomous agent for complex task solving.
- Beiwen Tian, Mingdao Liu, Huan-ang Gao, Pengfei Li, Hao Zhao, and Guyue Zhou. 2023. Unsupervised road anomaly detection with language anchors. In *2023 IEEE international conference on robotics and automation (ICRA)*, pages 7778–7785. IEEE.
- Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. 2024. Triposr: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*.
- Christina Tsalicoglou, Fabian Manhardt, Alessio Tonioni, Michael Niemeyer, and Federico Tombari. 2023. Textmesh: Generation of realistic 3d meshes from text prompts. *arXiv preprint arXiv:2304.12439*.
- Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. 2024. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. *arXiv preprint arXiv:2403.12008*.
- Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. 2023. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024a. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. 2024b. Prolicfdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems*, 36.
- Yuxi Wei, Zi Wang, Yifan Lu, Chenxin Xu, Changxing Liu, Hao Zhao, Siheng Chen, and Yanfeng Wang. 2024. Editable scene simulation for autonomous driving via collaborative llm-agents. *arXiv preprint arXiv:2402.05746*.
- Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. 2023. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15943–15953.
- Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. 2023a. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*.
- Yue Wu, Xuan Tang, Tom M Mitchell, and Yuanzhi Li. 2023b. Smartplay: A benchmark for llms as intelligent agents. *arXiv preprint arXiv:2310.01557*.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*.
- Guowei Xu, Peng Jin, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. 2024a. *Llava-cot: Let vision language models reason step-by-step*. *Preprint*, arXiv:2411.10440.
- Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. 2024b. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*.
- Le Xue, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, et al. 2024. Ulip-2: Towards scalable multimodal pre-training for

- 3d understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27091–27101.
- Jianing Yang, Xuweiyi Chen, Shengyi Qian, Nikhil Madaan, Madhavan Iyengar, David Fouhey, and Joyce Chai. 2023a. Llm-grounder: Open-vocabulary 3d visual grounding with large language model as an agent. In *2nd Workshop on Language and Robot Learning: Language as Grounding*.
- Xianghui Yang, Huiwen Shi, Bowen Zhang, Fan Yang, Jiacheng Wang, Hongxu Zhao, Xinhai Liu, Xinzhou Wang, Qingxiang Lin, Jiaao Yu, Lifu Wang, Zhuo Chen, Sicong Liu, Yuhong Liu, Yong Yang, Di Wang, Jie Jiang, and Chunchao Guo. 2024. [Tencent hunyuan3d-1.0: A unified framework for text-to-3d and image-to-3d generation](#). *Preprint*, arXiv:2411.02293.
- Zhengyuan Yang, Jianfeng Wang, Linjie Li, Kevin Lin, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023b. Idea2img: Iterative self-refinement with gpt-4v (ision) for automatic image design and generation. *arXiv preprint arXiv:2310.08541*.
- Xiaoju Ye, Junhao Chen, Xiang Li, Haidong Xin, Chao Li, Sheng Zhou, and Jiajun Bu. 2024. [MMAD: multi-modal movie audio description](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11415–11428, Torino, Italia. ELRA and ICCL.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Benton C. Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. 2022. [Scaling autoregressive models for content-rich text-to-image generation](#). *Trans. Mach. Learn. Res.*, 2022.
- Xin Yu, Ze Yuan, Yuan-Chen Guo, Ying-Tian Liu, Jianhui Liu, Yangguang Li, Yan-Pei Cao, Ding Liang, and Xiaojuan Qi. 2024. Texgen: a generative diffusion model for mesh textures. *ACM Transactions on Graphics (TOG)*, 43(6):1–14.
- Xianfang Zeng, Xin Chen, Zhongqi Qi, Wen Liu, Zibo Zhao, Zhibin Wang, Bin Fu, Yong Liu, and Gang Yu. 2024. Paint3d: Paint anything 3d with lighting-less texture diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4252–4262.
- Baoli Zhang, Haining Xie, Pengfan Du, Junhao Chen, Pengfei Cao, Yubo Chen, Shengping Liu, Kang Liu, and Jun Zhao. 2023a. [ZhuJiu: A multi-dimensional, multi-faceted Chinese benchmark for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 479–494, Singapore. Association for Computational Linguistics.
- Hang Zhang, Xin Li, and Lidong Bing. 2023b. Video-llama: An instruction-tuned audio-visual language model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 543–553.
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023c. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847.
- Xiang Zhang, Zeyuan Chen, Fangyin Wei, and Zhuowen Tu. 2023d. Uni-3d: A universal model for panoptic 3d scene reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9256–9266.

A Appendix: Evaluation Dataset

These IDEAs span a range of complexities: 9 are text-only, 57 feature both text and image inputs, 68 include text and 3D model inputs, and 64 contain text, image, and 3D model inputs, as shown in Fig. 5. Each test case was meticulously designed to represent real-world scenarios. The dataset cases are the same as the IDEAs shown in Fig. 1, Fig. 3, and Fig. 8.




Figure 5: IDEA content distribution.

The dataset also includes a distribution of tags: 9 IDEAs contain 0 tags, 62 contain 1 tag, and 127 include 2 tags, highlighting the diversity of annotation complexities, as shown in Fig. 6.




Figure 6: Tag distribution in IDEAs.

B Appendix: Visualization Results

The results of the visualization are shown in the Fig. 1, Fig. 7 and Fig. 8. For Tab. 3, we evaluate three T-2-I models: DeepFloyd IF (Lab, 2023), DALL-E (OpenAI, 2023b), and SD-XL (Podell et al., 2023), five I-2-3D models: Zero123 (Guo et al., 2023; Liu et al., 2023d), Wonder3D (Long et al., 2023), TripoSR (Tochilkin et al., 2024), InstantMesh (Xu et al., 2024b) and LGM (Tang et al.,

2024), and two LMM agent options: GPT-4V (OpenAI, 2023d,d) and LLaVA (Liu et al., 2024, 2023b).

Our framework has **good compatibility**. After our initial experiments, advanced image-to-3D generation models such as TripoSR (Tochilkin et al., 2024), InstantMesh (Xu et al., 2024b), and LGM (Tang et al., 2024) have emerged. Tab. 1 demonstrates that using these state-of-the-art models as the base modules for our T-2-I and I-2-3D results in significant improvements in *Idea23D* generation.

C Appendix: User Study Results

Tab. 3 shows the results of our user preference study. Our user study was conducted by distributing an online web-based survey questionnaire, with over 200 users participating. This user study’s comparative analysis reveals the remarkable superiority of the *Idea23D* framework over existing caption-based T-2-3D methods. Due to the lack of multi-round iterations in the T-2-3D baseline, the empty parts in the "Iterative self-refined" column of the table are the same as those in the "Initial-round" column.

Participants were presented with various IDEAs alongside multiple 3D model outputs from both caption-based T-2-3D baselines and *Idea23D*. To aid decision-making, users viewed rotating video representations of each 3D model. The order of presentation for both the cases and the model outputs was randomized to avoid any sequence bias, ensuring an unbiased assessment of the users’ preferences.

Our evaluation, detailed in Tab. 3, compares caption-based T-2-3D baselines and *Idea23D* using models specified in each row. We ask users to evaluate which model (the T-2-3D model, the first round results of *Idea23D*, and the end round results of *Idea23D*) is more satisfying, as well as evaluating each 3D model for compliance with IDEA.

The results show that *Idea23D* markedly enhances user preference scores across a diverse range of T-2-I, I-2-3D, and LMM models. Notably, the *Idea23D* framework’s initial prompting stage significantly outperforms the caption-based T-2-3D results by effectively decomposing and interpreting the user’s multimodal IDEA, thereby selecting the most suitable 3D model. This improvement is further amplified in the iterative self-refinement stage of *Idea23D*.

For instance, in scenarios utilizing DALL-E,



Figure 7: Comparison with commercial T-2-3D models. Inputs are the same as Fig. 3. Case (x) corresponds to the case of Fig. 1.

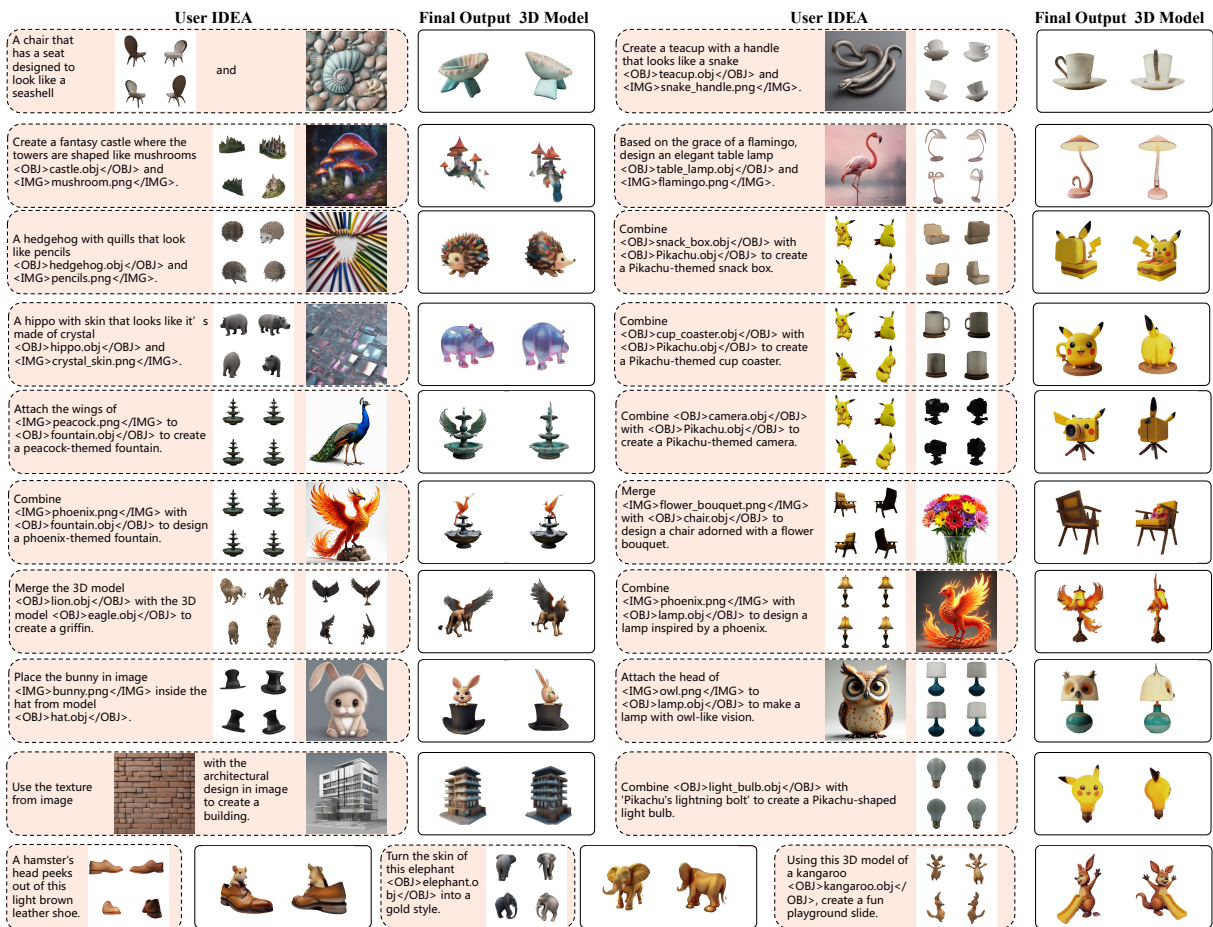


Figure 8: Results using GPT-4o, FLUX, InstantMesh as Idea2D components, cases using cases from the dataset.

Comparison stage: T-2-3D vs.			<i>Idea23D</i> (Initial-round)				<i>Idea23D</i> (Iterative self-refined)			
Evaluation			Which is better		Satisfying IDEA		Which is better		Satisfying IDEA	
Model			T-2-3D	<i>Idea23D</i>	T-2-3D	<i>Idea23D</i>	T-2-3D	<i>Idea23D</i>	T-2-3D	<i>Idea23D</i>
T-2-I	I-2-3D	LMM	T-2-3D	<i>Idea23D</i>	T-2-3D	<i>Idea23D</i>	T-2-3D	<i>Idea23D</i>	T-2-3D	<i>Idea23D</i>
SD-XL	TripoSR	GPT-4V	38.5%	61.5%	54.7%	80.2%	18.9%	81.1%	-	96.2%
SD-XL	InstantMesh	GPT-4V	43.2%	56.8%	53.1%	82.9%	20.1%	79.9%	-	96.4%
SD-XL	LGM	GPT-4V	41.3%	58.7%	56.2%	81.3%	12.5%	87.5%	-	94.5%
DALL-E	Zero123	GPT-4V	25.8%	74.2%	41.5%	78.2%	6.5%	93.5%	-	94.2%
DALL-E	Zero123	LLaVA	27.5%	72.5%	33.2%	70.6%	29.7%	70.3%	-	82.3%
DALL-E	Wonder3D	GPT-4V	18.8%	81.2%	47.9%	74.3%	3.5%	96.5%	-	91.1%
DALL-E	Wonder3D	LLaVA	16.3%	83.7%	28.3%	69.4%	10.3%	89.7%	-	76.2%
DeepFloyd IF	Zero123	GPT-4V	20.0%	80.0%	38.5%	64.1%	28.3%	71.7%	-	73.7%
DeepFloyd IF	Zero123	LLaVA	29.0%	71.0%	23.1%	57.6%	34.0%	66.0%	-	66.8%
DeepFloyd IF	Wonder3D	GPT-4V	29.3%	70.7%	32.6%	66.1%	14.8%	85.2%	-	75.9%
DeepFloyd IF	Wonder3D	LLaVA	26.0%	74.0%	19.3%	55.3%	18.3%	81.7%	-	64.5%

Table 3: Results of the user study.

Zero123, and GPT-4V, *Idea23D* models were preferred by users in 74.2% of the cases over T-2-3D in the initial round, demonstrating a higher IDEA satisfaction rate (78.2%). Conversely, T-2-3D models achieved only 41.5% satisfaction. In the iterative refinement comparisons, *Idea23D* models were favored even more (93.5% preference), achieving a remarkable 94.2% satisfaction rate.

These findings were consistent across various T-2-I, I-2-3D, and LMM configurations. Additionally, we observed that stronger T-2-I models, with enhanced language understanding capabilities, contributed to improved performance in *Idea23D*, suggesting that our framework may enjoy the development of off-the-shelf models it invokes.

D Appendix: Visualization of Self-iterative Refinement

Fig. 9 showcase the evolution and refinement of 3D models at different *Idea23D* stages and how the caption-based baseline works. Fig. 10 shows the iterative self-improvement process of a case in *Idea23D*.

E Appendix: Efficiency

Since generation from N prompts can run in parallel, the inference speed is agnostic of N . The inference speed of *Idea23D* mainly depends on the speed of the T-2-I and I-2-3D models, as the average number of iterations is 2 or 3. For example, with the *Idea23D* implemented using GPT-4V + zero123 + DALLE, generating an optimal result takes about 10 minutes (zero123 requires approximately 4 minutes to generate a 3D model from an image). In the Colab implementation mentioned in our abstract, we use GPT-4V + TripoSR + SDXL

to implement the *Idea23D* framework, and it takes about 5 minutes to generate a final 3D model (with 3 iterations). For reference, existing commercial methods typically generate a model in about 5-10 minutes. For example, 3DTopia³ and One-2-3-45++(Liu et al., 2023c) take about 5 minutes on average, while Luma⁴ and Meshy⁵ take about 10 minutes. Recent advanced methods such as TripoSR (Tochilkin et al., 2024), InstantMesh (Xu et al., 2024b), and LGM (Tang et al., 2024) have compressed the time to generate 3D models from images to within 1 minutes.

F Appendix: More Specific Domains

We tested our approach within the domain of automated design and modeling of **car**, **chair**, and **cloth**. The four views are plotted as a result of 3D modeling, and the text prompts above indicate the User Idea we entered. The result is shown in Fig. 11. The abstract IDEA input is above the picture.

³<https://github.com/3DTopia/3DTopia>

⁴<https://lumalabs.ai/genie?view=create>

⁵<https://www.meshy.ai/>

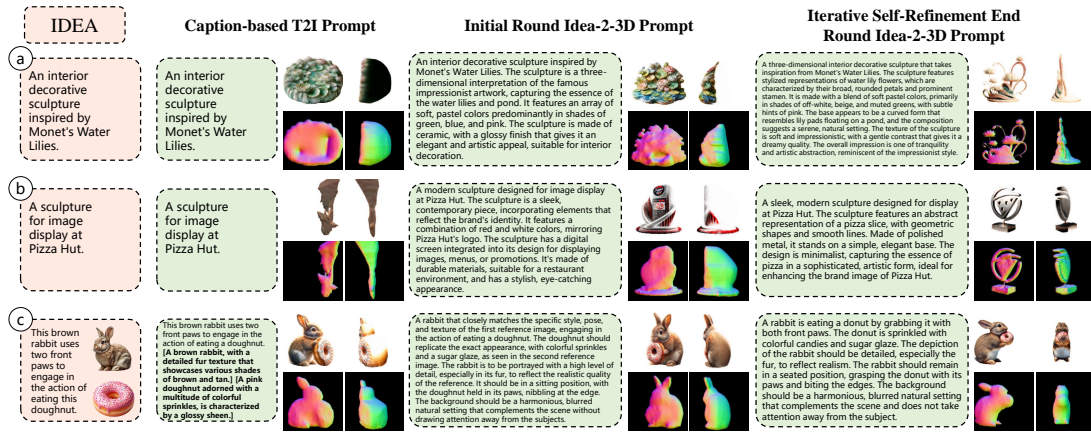


Figure 9: Comparison between prompts for "Caption-based T-2-I prompt for initial round", "Idea2-3-D prompt for initial round" and "Iterative self-refinement *Idea23D* end round" and the comparison between generated 3D models.

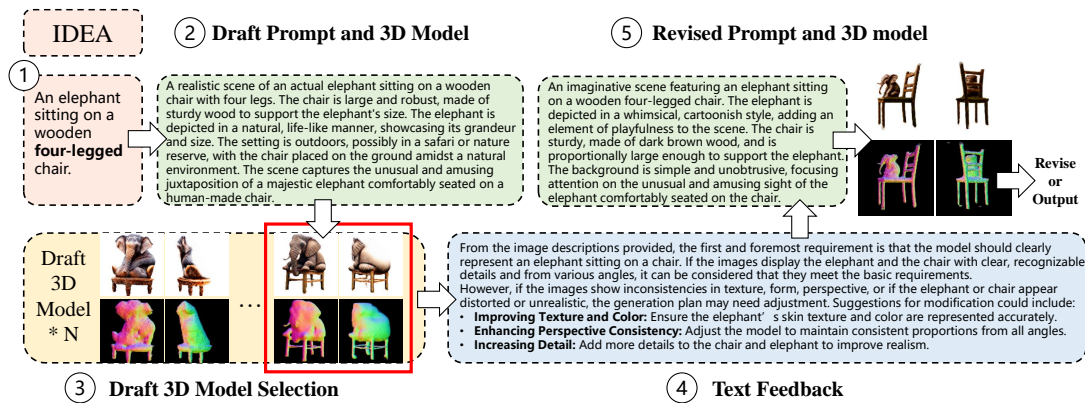


Figure 10: An example of the process for a 3D model generated through the *Idea23D* framework. An initial prompt is first generated based on the input IDEA and p_{gen} , then multiple images are generated from the T-2-I model, and then the corresponding draft 3D model is generated for each image separately. The most appropriate 3D model is selected based on p_{select} and feedback is generated. The results are output after several rounds of iterations. Green rectangles indicate text prompts for T-2-3D. blue rectangles indicate: modification suggestions given by agent based on draft 3D models and p_{fb} .

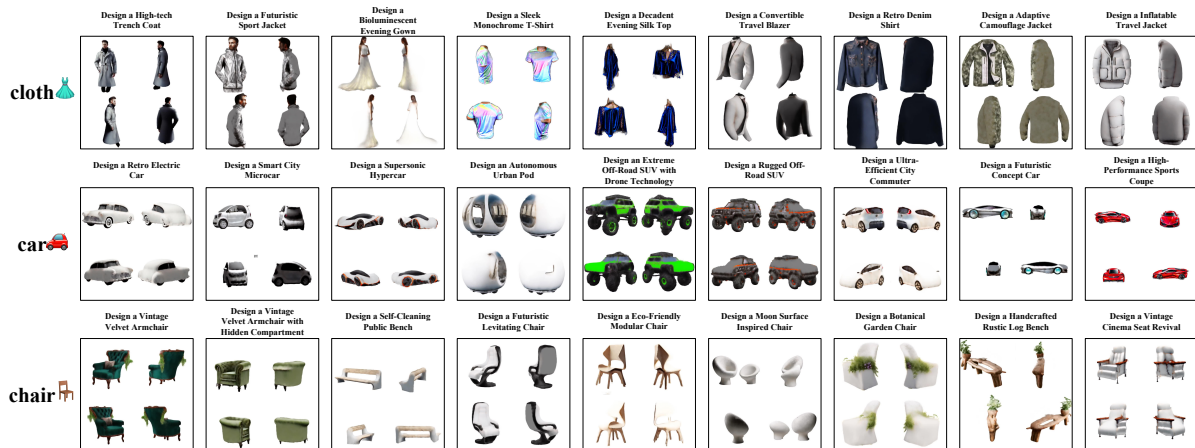


Figure 11: Domain-specific results of *Idea23D*.