RoomDreaming: Generative-AI Approach to Facilitating Iterative, Preliminary Interior Design Exploration

Shun-Yu Wang
National Taiwan University
Taipei, Taiwan
r11944013@csie.ntu.edu.tw

Wei-Chung Su
National Taiwan University
Taipei, Taiwan
r12944018@g.ntu.edu.tw

Serena Chen
University of California
San Diego, USA
sec022@ucsd.edu

Ching-Yi Tsai
National Taiwan University
Taipei, Taiwan
ching-yi.tsai@hci.csie.ntu.edu.tw

Marta Misztal
Queen Mary University of London
London, UK
marta.misztal00@gmail.com

Katherine M. Cheng
University of California
Berkeley, USA
katcheng@berkeley.edu

Alwena Lin
University of California
Los Angeles, USA
alwenalin1117@gmail.com

Yu Chen
National Taiwan University
Taipei, Taiwan
r11922026@ntu.edu.tw

Mike Y. Chen
National Taiwan University
Taipei, Taiwan
mikechen@csie.ntu.edu.tw

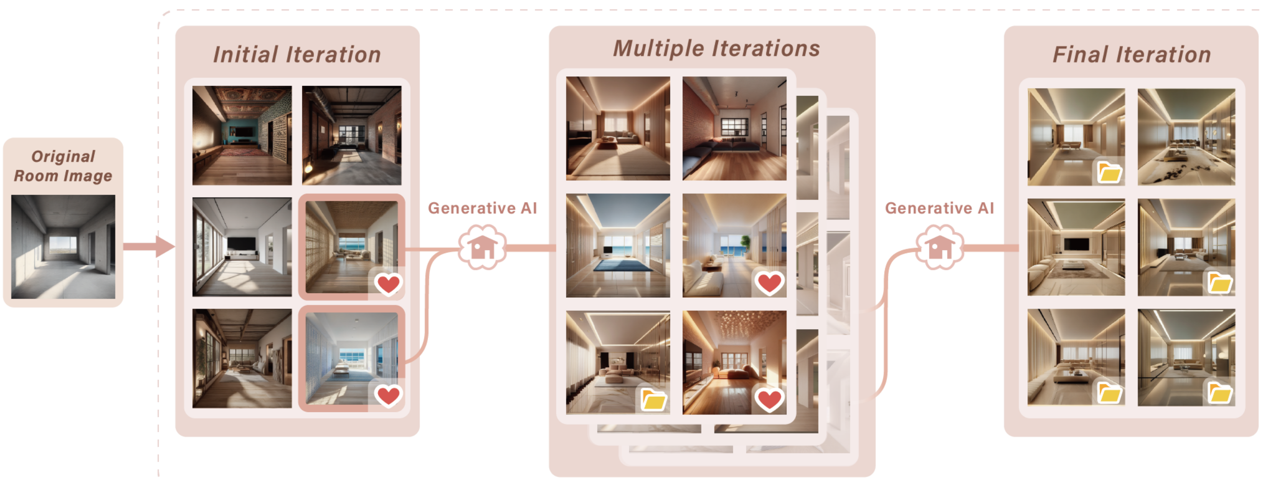

Figure 1: RoomDreaming, a generative AI tool designed to facilitate iterative, preliminary interior design exploration by creating photo-realistic designs based on the actual room layout and personal preferences indicated through *likes* and *bookmarks* with flexible creative control. The figure showcases samples from a homeowner-designer pair (G2) in our study, who used RoomDreaming for 11 iterations and reviewed 206 designs in one hour.

ABSTRACT

Interior design aims to create aesthetically pleasing and functional environments within an architectural space. For a simple room, the preliminary design exploration currently takes multiple meetings and days of work for interior designers to incorporate homeowners' personal preferences through layout, furnishings, form, colors, and materials. We present RoomDreaming, a generative AI-based approach designed to facilitate preliminary interior design exploration. It empowers owners and designers to rapidly and efficiently iterate through a broad range of AI-generated, photo-realistic design alternatives, each uniquely tailored to fit actual space layouts and individual design preferences. We conducted a series of formative and summative studies with a total of 18 homeowners and 20 interior designers to help design, improve, and evaluate RoomDreaming. Owners reported that RoomDreaming effectively increased the breadth and depth of design exploration with higher efficiency and

CHI '24, May 11–16, 2024, Hawai'i, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

satisfaction. Designers reported that one hour of collaborative designing with RoomDreaming yielded results comparable to several days of traditional owner-designer meetings, plus days to weeks worth of designer work to develop and refine designs.

KEYWORDS

generative-AI, interior design, architecture, human-centered AI

ACM Reference Format:

Shun-Yu Wang, Wei-Chung Su, Serena Chen, Ching-Yi Tsai, Marta Misztal, Katherine M. Cheng, Alwena Lin, Yu Chen, and Mike Y. Chen. 2024. RoomDreaming: Generative-AI Approach to Facilitating Iterative, Preliminary Interior Design Exploration. In *Proceedings of Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24).* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Interior design, *architecture*, and *landscaping* collaboratively create a harmonious and functional built environment to enhance the quality of life of people and connect them to the natural environment [10]. Specifically, interior design aims to create aesthetically pleasing and functional environments within an architectural space. It involves planning the layout and designing the furnishings, finishes, and lighting through characteristics such as form, shape, color, texture, and materials to reflect the desires and preferences of users of the space [9].

The interior design process comprises the following three iterative stages, as described by one of the authoritative textbooks of interior design, *Interior Design Illustrated* (Ching, 2018) [9]:

(1) *Programming* understands and analyzes user requirements, activity needs, furnishings requirements, original space, and desired qualities fitting to the architectural space.
(2) *Plan arrangement* develops and evaluates different design alternatives with specific furnishings, finishes, and lighting in 3D to iteratively progress from divergent possibilities to converge on a specific, final design. The arrangement of shapes and forms in space should respond to functional and aesthetical criteria by iteratively evaluating and refining the different design alternatives to decide on design characteristics (e.g., form, shape, color, texture, and material) for each design element.
(3) *Implementation* prepares detailed construction drawings, including floor plans, elevations, and sections, finalizes specifications for interior finishing materials, and physically completes the construction.

"While the initial stages of the design process encourage divergent thinking about the problem, the design development phase requires a convergent focus on a specific design solution." [9]

For the plan arrangement stage, homeowners explore design possibilities through *self-guided exploration*, *designer-assisted exploration*, or both. For *self-guided exploration*, owners collect ideas and reference images from sources such as Pinterest, search engines, designers' websites, and real-life experiences. However, existing reference designs do not match the actual space being designed and do not allow users to combine ideas to iteratively refine and explore the design space further.

With the advancement of generative-AI (Artificial Intelligence) for images and text, particularly the release of Stable Diffusion and ChatGPT in 2022, several generative-AI products have launched that allow a photo or 3D model of a space be used as input and then generate reference designs in a variety of styles, such as InteriorAI[4], RoomGPT[5], REimagineHome[6], SpacelyAI[7], and MagicRoomAI [8]. While generating images that match the actual space is a critical first step forward, existing approaches lack the ability for users to specify preferences to iterate further, which is necessary to help owners to explore the design space, make decisions, and for design exploration to converge toward a final design.

For *designer-assisted exploration*, interior designers must thoroughly understand owners' requirements and preferences, in order to develop design alternatives towards a final design. This is a time-consuming and labor-intensive process that typically starts with an initial owner-designer meeting to gather requirements and preferences for *programming* and *plan arrangement*, followed by multiple cycles of: 1) designers develop and propose design alternatives and 2) owner-designer design review meetings, which are repeated until converging on a final design. To improve the efficiency of developing design alternatives, researchers and commercial products have explored algorithms and AI to provide recommendations for specific design elements, characteristics, 2D floor plans, and 3D models [6–8, 29–31, 41, 44, 47, 51]. While these approaches only generate a specific design aspect of the entire space, they helped inspire the more recent generative-AI based approaches that generate entire designs for a space.

Even with all of the existing CAD tools, AI products, and research prototypes, the design iterations exploring the design alternatives for a simple space require days to weeks of designers' work plus owner-designer meetings for an overall typical time span of 6 to 15 weeks [16, 27, 28, 32, 33]. Furthermore, time and budget constraints limit the design exploration in terms of both the number of alternatives and the number of design iterations, resulting in final designs that may not fully reflect and satisfy owners' preferences.

To improve the efficiency and effectiveness of the early stages of the interior design process, we present RoomDreaming, a generative-AI approach to facilitate preliminary, iterative interior design exploration by generating photo-realistic designs based on the actual space layout and enabling users to iterate through vast design alternatives based on indicated preferences. We conducted a series of two formative studies and three summative studies with a combined total of 18 homeowners and 20 interior designers, shown in Table 1, to understand the needs of owners and interior designers, and iteratively improved the RoomDreaming system, which we developed using OpenAI GPT API, Stable Diffusion [1], and ControlNet [50]. Based on the feedback from these studies, our prototype uses *Likes* and *Bookmarks* to capture user preferences, and additionally supports *User Requirements* and *New Design Directions* for more precise control of the designs.

Compared to designer-assisted exploration that typically iterates 2~3 times through several design proposals over the span of several weeks, RoomDreaming's generative-AI approach enables users to rapidly iterate as many times as needed through hundreds of designs. As shown in Figure 1, one of the owner-designer pairs from our co-design study used RoomDreaming for 11 design iterations and reviewed 206 design alternatives in 1 hour. The interior designers from the study, who had an average of 8.3 years of professional

Table 1: The 5 user studies conducted for designing, improving, and evaluating this research and assessing the quality of AI-generated designs, with a combined total of 18 owners and 20 interior designers.

User Study	Duration (min)	Participants
1. Formative Study	90	3 Designers + 3 Owners
2. Assessment of AI-generated Design Quality	60	8 Designers
3. Self-guided Exploration	120	12 Owners
4. System Improvement	120	6 Designers
5. Co-design Exploration	120	3 Designers + 3 Owners

design experience, estimated that co-designing with RoomDreaming for 1 hour achieved the equivalent of several days of traditional owner-designer meetings, plus days to weeks of designer work to develop and refine designs.

Our key contributions include:

- Developing a generative AI approach to support iterative, preliminary interior design exploration.
- Designing, implementing, and iterative refining a generative-AI system, informed by a series of formative and summative studies with a combined total of 18 homeowners and 20 interior designers.
- Empowering owners and designers to rapidly iterate through a broad range of AI-generated, photo-realistic design alternatives, each uniquely tailored to fit actual space layouts and individual design preferences. This enhances both the breadth and depth of design exploration, as well as overall efficiency and satisfaction..

2 RELATED WORK

2.1 Generative-AI Interior Design Tools

With the recent, rapid advancement in AI, there has been growing discussion about human-AI interaction [2, 42], particularly with the release of AI image generators such as Stable Diffusion[1], Midjourney[2], and DALL-E[3] in 2022. As visual representation is critical for understanding personal preferences for interior design [11, 38], several generative AI products for interior design have been launched in 2023, including InteriorAI[4], RoomGPT[5], REimagineHome[6], SpacelyAI[7], and MagicRoomAI[8].

RoomGPT[5] takes an input image of a room and generates detailed renditions based on user style preferences. Though users can choose from a wide variety of styles for image regeneration, they cannot specify preferences regarding design characteristics or elements, and is only given one design at a time that is not based on the users' preference from prior generated designs. Interior AI[4] provides 4 transformation methods: Virtual Staging for detecting construction to avoid altering them, Interior Design for change in construction, Freestyle for randomization, and 360° Panorama for immersion. Users can select a style from the provided list, and is given a maximum of 9 rendered images per batch generation. Similar to RoomGPT, users are unable to specify design requirements. Additionally, though users can continuously regenerate new images, they are not based on user preferences. MagicRoomAI[8] offers theme and room type options, a designer's name to incorporate their design style, and a text description. Though regeneration is possible, users cannot influence the direction of the next regeneration. In SpacelyAI[7], users can render a 3D model or a basic sketch of a room. The style can be chosen from SpacelyAI's list, or the user can upload their own image of a style they would like to emulate, allowing for great flexibility. Users can also select their own color palette or replace objects in the image after generation for a different look. However, while users can upload their own stylistic references, they cannot state preferred design elements before generation.

For precise modification, REimagineHome[6] has a masking function to select which specific areas of the room to alter. They can also enter their own instructions regarding design characteristics and color preferences. While this system allows the user to customize their design, it is only focused on one design direction and does not allow users to explore other design variations and alternatives in the same design direction.

Although these generative-AI tools support the generation of initial designs based on a photo of a space, they do not support iterative design exploration, which RoomDreaming has been designed to empower. Furthermore, RoomDreaming provides users with control over the ratio of *New Design Directions*, which is fundamental to the divergence and convergence process of creative design exploration.

2.2 Computer Aided Design Tools for Interior Design Exploration

Computer-aided design (CAD) system has played a pivotal role in the modernization of interior design. These systems can be positioned on a spectrum that ranges from direct manipulation to fully automatic design. [48] Meanwhile, *Generative Design(GD)* is used to describe *computer-aided design (CAD)* systems that offer tools to modify designs beyond directly manipulating individual design elements. [48]

There are several research and commercial applications of CAD for interior design, facilitating the efficiency of the design development [13, 20, 21, 23, 36, 37, 43]. For example, AutoCAD[9] is one of the most popular CAD software that acts as a complete tool for automating graphical work (e.g. floor plans, sections, elevations, and construction drawings). It supports integration with other 3D modeling systems, broadening its applications that range from

[1]Stable Diffusion https://stablediffusionweb.com/
[2]Midjourney https://www.midjourney.com/
[3]DALL-E-2 https://openai.com/dall-e-2/
[4]InteriorAI https://interiorai.com/
[5]RoomGPT https://www.roomgpt.io/
[6]REimagineHome https://www.reimaginehome.ai/
[7]SpacelyAI https://www.spacely.ai/
[8]MagicRoomAI https://magicroom.ai/
[9]AutoCAD https://www.autodesk.com/products/autocad/overview/

static drawings to object interaction [23]. However, despite many methods proposed by researchers for the use of CAD in early-stage conceptual design, these still require significant user input and are mostly used in the final stages of design [48].

2.3 Generative Design

"Generative design (GD) as a rule-driven iterative design process is based on algorithmic and parametric modeling to automatically explore, iterate, and optimise design possibilities by defining high-level constraints and goals." [26] Different from the "rule-based" approach in the AI domain, "Generative Design" employs a "rule-driven" process. This involves setting high-level parameters and constraints to guide the automated exploration, iteration, and optimization of design possibilities. In contrast, "rule-based" AI relies on a fixed set of pre-defined rules to generate outputs.

Recently, researchers have explored the utilization of algorithms and AI based on generative design to provide interior design recommendations, covering the selection and arrangement of design elements and characteristics, as well as 2D floor plans, and 3D models [6–8, 29–31, 41, 44, 47, 51].

Because designers often have intuition and knowledge cultivated from experience, there are aspects of design that they consider and merge into a design that may not fall within an owner's ability or consideration. These aspects are often necessary to ensure a design that is aesthetically pleasing, harmonious, and other design principles. [9]. One aspect is selecting color attributes for characteristics in a design [7] and pairing these colors together [51], as colors are important to homeowners because of their effect on mood and emotion. While professional designers have the experience and intuition to pair colors aesthetically and comfortably, homeowners can find this difficult.

Chen et al. [7] used a statistical model to analyze color combinations in labeled interior design scenes. Their method compared favorably to random assignments and the established Magic Decorator tool in user studies. Zhu et al. [51] employed deep learning to learn from professional photos and renderings. Their system generated perceptually convincing color schemes, preferred by users over professional designers in most cases, and significantly faster. Both approaches show promise for efficient and user-friendly color recommendations for interior design.

Another aspect is determining compatibility between furniture pieces [44], which can help a scene look more harmonious but is a complex task that requires intuition born from experience. Using a deep learning network, style can be classified and modeled for style-compatible and consistent scenes. Two other aspects that can appear mysterious to homeowners are assigning textures and materials to elements [8], and color-material furnishing pairing [30]. The former is a task that designers complete based on experience, while the latter is referred to by the authors as "a 'black-box' for interior designers" because designers find it difficult to explain the rules behind their decisions that are fueled by intuition. Both tasks can be difficult for homeowners, who have less knowledge in interior design. To solve this problem, these two systems have "guidelines" or analysis framework that emulates a designer's experience and intuition to produce plausible and cohesive suggestions for users. Coming up with original ideas [29] is an aspect that is also difficult for homeowners because of their lack of experience, which prevents them from conveying the feeling they want their design to emanate. Using input from the user that specifies shape, material, and color, this system uses a machine learning algorithm to select and generate the appropriate interior style.

While these tools are effective for specific subsets of design, supporting users to explore all the design elements together in each design alternative is essential, as visual relationships among the design elements are shaped by principles like *proportion*, *scale*, *balance*, *harmony*, *unity and variety*, *rhythm*, and *emphasis*, arranging elements into recognizable patterns, allowing for visual order while accommodating function and purpose within the space [9]. These have inspired generative-AI based tools, which are capable of generating all the design elements of a space through training models on a large number of existing designs, which RoomDreaming uses.

3 STUDY #1: FORMATIVE STUDY

To gain insight into current interior design processes and challenges, we conducted a formative study by interviewing interior designers and homeowners who have recently collaborated with interior designers to complete residential projects.

3.1 Study Design, Procedure, and Participants

We designed a semi-structured interview focusing on the *"programming"* and *"plan arrangement"* stages of interior design, covering: 1) overall design process; 2) communication of owner requirements and preferences; and 3) design exploration and iteration.

We recruited a total of 6 participants, 3 homeowners and 3 interior designers, comprising 3 males and 3 females with ages ranging from 24 to 52. The 3 interior designers (D1~D3) have professional interior design experience of 4, 6, and 5 years, and specialize in residential, commercial, and workspace design, and the three homeowners (O1~O3) have completed between 1~3 residential projects.

Each participant was asked to bring their most-recently completed interior design projects, which included 5 residential and 1 commercial designs. Each interview took about 90 minutes.

3.2 Findings

The interviewers all mentioned owners' self-guided design exploration in addition to designer-guided design exploration, and the key frictions and pain points are as follows:

Significant effort needed for owners to collect preferred reference designs. Homeowners spent *"one week"* (O2) to *"10~15 days"* (O3) searching for reference designs through various sources, including designers' portfolio websites, Pinterest, Google Image Search, YouTube, interior design books, and personal photos of the interior designs they encountered, such as *"prepared lots of the photos I took in the several hostels and hotels around the world to assist the conversation with designers."* (O1)

Reference designs not matching the room layout. *"After discovering a design or material I like, I try to search with specific keywords for similar designs that align with my room layout with little success."* (O1) *"I spent lots of time trying to find reference designs with a similar room layout."* (O3)

Inability to explore the integration of multiple design ideas. Homeowners currently have to imagine how multiple design ideas integrate together. Furthermore, they are unable evaluate the feasibility and compatibility of their design ideas, such that the overall design satisfy design principles, like *proportion, scale, and harmony* [9]. As mentioned by one couple, *"we found many possible design ideas for our future home using Pinterest, but we were unsure how they fit together and couldn't make decisions."* (O3) Designers commented that *"clients commonly prepare several reference design images and ask us to merge them, but we need to explain that the integration may not be aesthetically pleasing or may conflicts with other preferences clients just mentioned."* (D1)

Verbal description alone, without visual, cannot precisely convey design preferences. In order for designers to understand owners' high-level design directions and preferences, designers utilize diverse approaches to assist homeowners in articulating their preferences and requirements, including thumbs up/down questions about reference designs, storytelling, creating draft drawings, and interviewing owners about their hobbies, habits, and daily life. In addition to the words in the answers, designers report the importance in observing their body language and emotions. Even with the above techniques, *"it's common that we are "guessing" clients' preference on the design material, lighting, and so on, based on their verbal descriptions."* (D3)

Furthermore, words often fail to convey design preferences, leading to errors in communication and understanding. One owner expressed frustration in that *"I didn't know the precise keywords to convey the design styles I liked, but once I saw the materials that the designer selected in the final rendering, I immediately knew that I didn't like some of them."* (O1)

Limited number of iterations and design alternatives (proposals) developed by designers. In the preliminary design process, two to three 2~3-hour owner-designer meetings took place with designers creating 3~9 design alternatives over 3~5 weeks. A design alternative in this phase included: 1) a floor plan, 2) existing reference design images, and 3) a small number of 3D images rendered for the project. Because 3D modeling is time-consuming and costly, reference images unrelated to the physical room layout are often used to convey ideas and to support the discussion. *"The average number of design proposals in the whole design process may vary depending on several factors, such as the scope of work. For this case, with a total of three rooms in this house, we met with the clients 3 times and provided a total of 6 design alternatives in three weeks to converge to the final design."* (D1) *"Regarding the meetings, if there are large changes in design directions, then the designer would redraw the designs. Otherwise, after 2 meetings with each 2~3 design proposals, the direction of the design was determined. Overall, I saw a total of about 3 rendered images."* (O2)

The cumulative impact of these factors makes it challenging for homeowners to to fully explore their design ideas and to convey their design preferences thoroughly to designers.

4 SYSTEM DESIGN AND IMPLEMENTATION

Our insight is to leverage generative-AI's ability to rapidly generate vast design alternatives, and to tailored it to support the inherent iterative nature of early-stage interior design exploration. Our goal is to empower users to efficiently broaden the scope and depth of their design exploration and to facilitate communication, with the following system design goals:

- **High-quality designs**: generating design alternatives based on the physical room layout that match user requirements, including structural, functional, and aesthetics. The generated images should be photo-realistic to help users experiment and assess design ideas.
- **Breadth**: expanding breadth of exploration by introducing new design directions, and by introducing variations within a preferred design direction.
- **Depth**: supporting rapid iteration and generation of new designs based on user-indicated preferences.

4.1 Web-based User Interface

For the user interface, we aim to achieve the "low floor" and "high ceiling" concepts proposed by Seymour Papert [19], that provides easy ways for novices to get started (low floor) but also ways for them to work on increasingly sophisticated designs over time (high ceiling):

- **Low-floor**: users only need to provide a photo of the room to start generating designs. User preferences are collected through familiar *Like* and *Bookmark* interactions.
- **High-ceiling**: optional guidance of AI generation through keywords and UI controls.

Figure 2 shows a screenshot of the RoomDreaming UI for browsing design alternatives, expressing preferences via Likes, saving designs via Bookmarks, and guidance of AI generation by providing requirement keywords and controlling the ratio of New Design Directions via sliders.

The interface is designed to enable users to easily browse and generate vast new design alternatives to gain a deeper understanding of their design preferences (both likes and dislikes) through successive iterations. This iterative process allows users to gradually refine their preferred designs while retaining the opportunity to learn from additional design suggestions.

4.2 System Overview

Figure 3 provides a high-level overview of the system, showing three main components: 1) web-based user interface; 2) backend to generate design alternatives, consisting of an *Image Analyzer*, *Prompt Composer*, and *Design Generator*; and 3) Large language model (LLM) for providing new design directions. The web UI generation is implemented using Gradio[10], a Python package for integrating machine learning models into web interfaces.

4.2.1 Understanding Room Elements and Spatial Information. To lower the barrier to start generating designs, RoomDreaming does not require a 3D model of the space, and can instead use a photo of the room as input. The room can be fully furnished spaces, like existing rooms, or empty spaces awaiting to be furnished. The Image Analyzer aims to understand *elements*, which are the existing element types in the room (e.g. window), and *spatial information*, which are the scale, shape, and relationship between each object

[10]Gradio https://www.gradio.app/

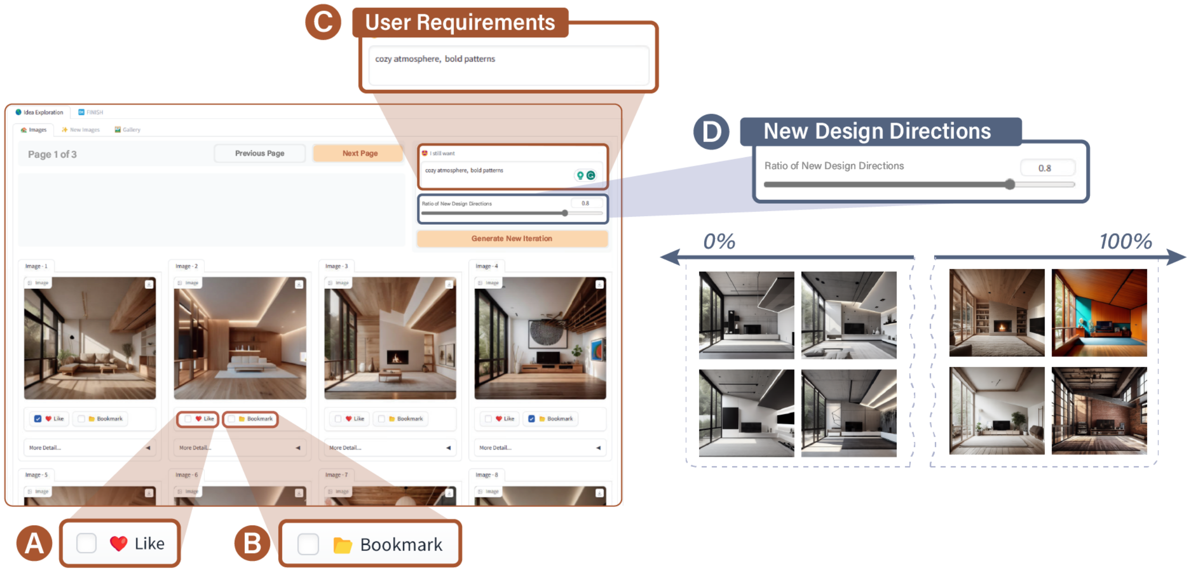

Figure 2: Screenshot of RoomDreaming's web-based user interface, enabling users to browse vast number of design alternatives and indicate preferences through A) *Likes* and B) *Bookmarks*. To provide additional control over the design generation, users can specify C) *Requirements* through keywords, and adjust D) ratio of *New Design Directions*.

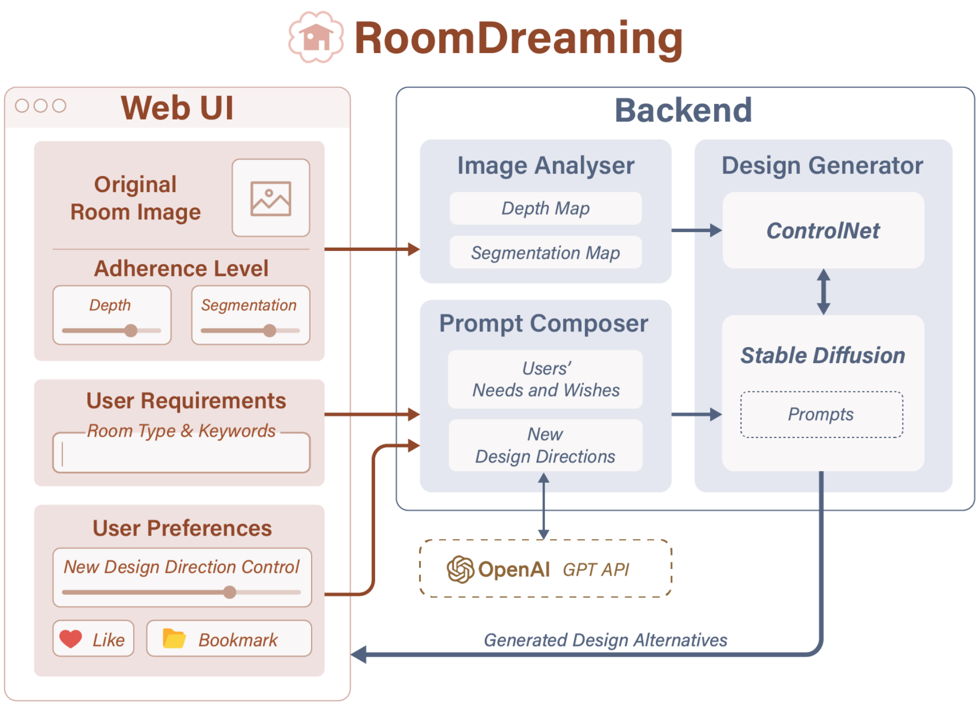

Figure 3: System architecture overview, showing 1) web-based user interface; 2) backend to generate design alternatives, consisting of an *Image Analyzer*, *Prompt Composer*, and *Design Generator*; and 3) Large language model (LLM), currently via API.

in the room (e.g. size and position of the window relative to the space).

To understand *elements*, we use the segmentation estimators UPerNet Model[11] based on Unified Perceptual Parsing (UPerNet) [46], which mimics human vision by categorizing and detecting objects

[11]UPerNet Model https://huggingface.co/docs/transformers/model_doc/upernet

within scenes, and the ADE20K image dataset[12], which is a comprehensive database with objects annotations to provide semantic information about the elements in the space. To understand *spatial information*, we use the estimator model res101[13]. This estimator is based on the widely adopted technique known as Monocular Depth Estimation[14], which calculates the distance of each pixel in a 2D image from the viewer's perspective, creating a depth map of the 3D space. Figure 4(a) shows an example of the segmentation map and spatial map generated from an input photo.

4.2.2 Generating Designs based on the Room. To generate photo-realistic interior design alternatives based on the segmentation and depth maps, we employed Stable Diffusion [1], a generative-AI model that produces unique images from text and image prompts, and ControlNet [50], a neural network structure capable of conditionally controlling diffusion models during image generation. Specifically, we integrate the depth[15] and segmentation[16] ControlNet models to facilitate multiple conditioning controls. These controls operate on the previous estimators' depth and segmentation maps throughout the diffusion model's image generation process to ensure the output images are based on the given room layout. For design image generation, we use a diffusion model [17] that has been fine-tuned for interior designs and highly rated on the AI community, CIVITAI [18].

To support independent design exploration of varying adherence levels to existing *elements* and *spatial information*, we allow users to independently specify the adherence levels, implemented using the `Control Weight` parameter in ControlNet that ranges from 0 to 1. A lower weight results in less adherence to the input room, while a higher weight generates designs more closely aligned with the input room. Figure 4(b) shows examples of low vs. high adherence levels.

4.2.3 Image Generation Latency. We evaluated the latency of the key components, by averaging over 100 trials. Local PC has an AMD R9 3950X 16-core CPU + Nvidia RTX 4070 GPU.

- 6.1s for image analysis of depth and segmentation map, a one-time computation at the beginning of each project on PC.
- 11.8s for GPT API to generate a batch of 5 new design directions prompts. Pre-fetching and caching can eliminate this latency.
- 3.21s/image for design image generation using Stable Diffusion and ControlNet on an Nvidia A10G (Large instance) on Huggingface, and 3.35s on a PC with Nvidia RTX 4070.

4.2.4 Expanding Breadth of Exploration. Studies of text-to-image generative AI have shown that users grappled with a "trial and error" approach, inefficiently modifying prompts and brainstorming to generate optimal descriptions for new images [15, 49]. In the context of interior design, homeowners from our formative study reported limited knowledge of possible design styles and limited terminology to express what they desire. Therefore, instead of generating design only within the scope of user-specified prompts, RoomDreaming's *Prompt Composer* leverages Large Language Models (LLM) to expand prompts to explore new design directions.

We use OpenAI GPT API, gpt-3.5-turbo[19] to generate prompts using GPT instructions based on prior work on instruction design [5, 14, 17] and for prompt structure suitable for interior design generation using Stable Diffusion [3, 18, 25, 45]. The prompts from GPT are appended after user specified requirements, as this prompt order gives the design directions a lower priority for Stable Diffusion image generation. To maintain divergence in the generated design directions, we leveraged GPT's conversation history and GPT API parameters with default `presence_penalty`, `temperature:1.2`, and `top_p:1` which level is needed to produce prompts suitable for Stable Diffusion and maintain the divergence in each descriptor.

4.2.5 Supporting Depth of Exploration. Users indicate their preference via *Likes* and *Bookmarks*, as shown in Figure 2(a)(b). The LLM portion of prompts corresponding to the liked and bookmarked designs are stored as *preferred prompts*. To enable users to control the exploration process, we provide a slider to control the ratio of *New Design Directions*, as shown in Figure 2(d). As an example, when the ratio of *New Design Directions* is set to 80% when generating the next batch of designs, 20% of prompts will be randomly sampled from the user's preferred prompts, if any, with the remaining 80% newly generated by LLM. When re-using a preferred prompt, we use a random `seed` to generate new design variations within the design direction.

5 STUDY #2: QUALITY ASSESSMENT OF AI-GENERATED INTERIOR DESIGNS

While numerous studies have assessed the quality and performance of Stable Diffusion with ControlNet [34, 35, 39, 50], no prior work has assessed the quality of generated images in the context of interior design.

Interior Design Illustrated (Ching, 2018) [9] outlined the following four key aspects of interior design relevant to preliminary design exploration:

- **Structural and Enclosure System**, assessing the integration of structural system (comprising vertical columns and horizontal beams) and enclosure systems (encompassing the building envelope, interior walls, partitions, and ceilings) in existing or proposed spaces.
- **User Requirements Compatibility**, evaluating the compatibility of user requirements with desired spatial quality.
- **Functional Criteria**, analyzing furniture layout and ergonomics for functional excellence, emphasizing a harmonious fit between the spatial form and dimensions and the human body.
- **Aesthetic Criteria**, careful attention to appropriate scale in relation to space function, visual grouping for unity with variety, figure-ground reading, 3D composition elements like rhythm, harmony, and balance, appropriate orientation toward light, view,

[12] ADE20K Website https://groups.csail.mit.edu/vision/datasets/ADE20K/
[13] Annotators-res101.pth https://huggingface.co/lllyasviel/Annotators/blob/main/res101.pth
[14] Monocular Depth Estimation https://paperswithcode.com/task/monocular-depth-estimation
[15] Depth ControlNet Model https://huggingface.co/lllyasviel/control_v11f1p_sd15_depth
[16] Segmentation ControlNet Model https://huggingface.co/lllyasviel/control_v11p_sd15_seg
[17] XSarchitectural-InteriorDesign-ForXSLora-V11 https://civitai.com/models/28112/xsarchitectural-interiordesign-forxslora
[18] CIVITAI https://civitai.com/
[19] OpenAI gpt-3.5-turbo https://platform.openai.com/docs/models/gpt-3-5

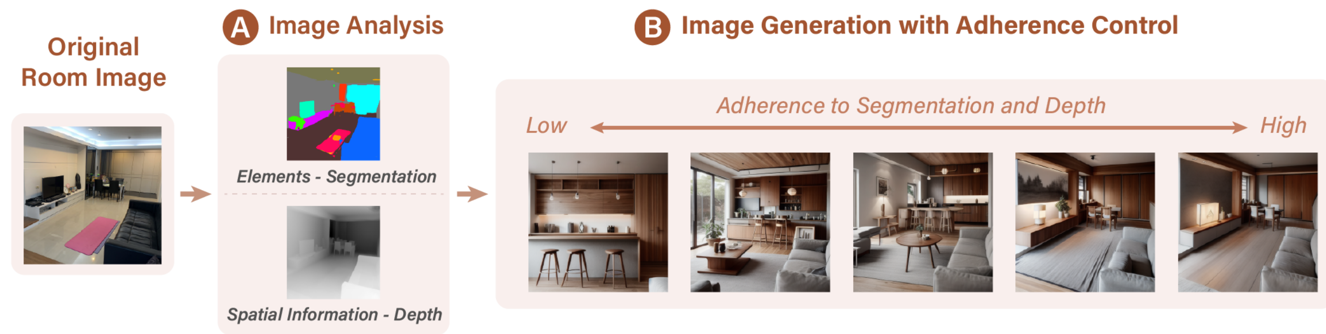

Figure 4: Illustration of the A) Image Analysis and B) Image Generation with Adherence Control pipeline. The system analyzes the user-input room image, employing depth and segmentation estimators to capture Elements and Spatial Information. The user can then control adherence to existing elements, as demonstrated in this example from the owner-designer co-design exploration study (G1)

		A				B															
		Overall				Structural and Enclosure System				User Requirements Compatibility				Functional Criteria				Aesthetical Criteria			
	Seg \ Depth	0.25	0.5	0.75	1	0.25	0.5	0.75	1	0.25	0.5	0.75	1	0.25	0.5	0.75	1	0.25	0.5	0.75	1
Empty Room	0.0625	66%	63%	65%	66%	0%	50%	70%	80%	100%	70%	15%	5%	75%	30%	95%	95%	90%	100%	80%	85%
	0.125	48%	73%	66%	66%	10%	65%	75%	70%	75%	75%	10%	15%	40%	70%	90%	90%	65%	80%	90%	90%
	0.1875	58%	73%	61%	63%	30%	65%	70%	80%	90%	80%	20%	5%	45%	70%	70%	85%	65%	75%	85%	80%
	0.25	55%	71%	53%	60%	25%	70%	70%	85%	85%	70%	15%	0%	40%	75%	50%	75%	70%	70%	75%	80%
With Furniture	0.0625	56%	55%	86%	80%	10%	50%	65%	70%	70%	70%	100%	80%	80%	55%	100%	100%	65%	45%	80%	70%
	0.125	56%	66%	81%	86%	25%	60%	75%	70%	75%	80%	90%	85%	75%	55%	85%	90%	50%	70%	75%	100%
	0.1875	61%	68%	88%	90%	30%	80%	75%	80%	85%	75%	95%	90%	50%	60%	100%	100%	80%	55%	80%	90%
	0.25	61%	74%	81%	90%	40%	70%	70%	80%	75%	75%	80%	85%	65%	65%	100%	100%	65%	85%	75%	95%

Figure 5: Percentage of quality ratings that are rated GOOD and VERY GOOD, for each of the 16 depth/segmentation parameter combinations, with each cell in the table representing 20 ratings, i.e. 10 images rated by 2 designers. (A) shows the overall, averaged percentage across the 4 key aspects shown in (B): Structural and Enclosure System, User Requirements Compatibility, Functional Criteria, and Aesthetic Criteria.

or internal focus, and the judicious use of shape, color, texture, and pattern.

Our primary goal is to assess the viability of the AI-generated images as design alternatives for interior design exploration, specifically in terms of the four key aspects of interior design. The secondary goal is to understand how the `Control Weight` parameters of depth and segmentation maps affect the quality of design, to help understand suitable ranges and tradeoffs.

5.1 Assessment by Interior Designers

Based on the interior design projects in the formative study, we created a typical project for a living room with design elements including a large window, cozy sofa, wood table, organized storage, and in minimalist style, and assessed the two common types of *site conditions*: 1) an empty room, and 2) a room with existing furniture.

We sampled the two `Control Weight` parameters for depth map and segmentation maps uniformly in 4 intervals, using suitable ranges based on feedback from a pilot study with 3 designers. For high segmentation weight values, the generated image were often overly influenced by the segmentation map, such as when an empty room image is inputted along with text prompts for new furnishings, yet the generated image still showed an empty room. Examples of generated designs corresponding to the segmentation weights from 0.25 to 1.0 are shown in Appendix A.1.

Overall, there are 16 combinations of the two parameters used for quality assessment, with 4 depth weights (ranging from 0.25 to 1 in 0.25 intervals) multiplied by 4 segmentation weights (ranging from 0.0625 to 0.25 in 0.0625 intervals). 10 images were generated with different seeds for each of the 16 combinations of control weights for each of the 2 site conditions, for a total of 320 images. The images were randomly divided into 4 sets of 80 images, and each image was graded by 2 designers then the ratings were averaged to reduce potential bias.

5.1.1 Participants and Procedure.

We recruited 8 designers (D4~D11), 4 males and 4 females, with ages ranging from 26 to 50. Their professional design experience ranged from 5-15 years (mean=8.375, SD=3.46). The designers were already familiar with *Interior Design Illustrated*. We first reviewed the four design aspects to be graded with the designers, then briefed them on the 5-point Likert scale for quality, ranging from Very Poor, Poor, Acceptable, Good, to Very Good, where Good and Very Good represent sufficient design quality that the designers would use for discussion with their clients. Each designer then evaluated the assigned set of 80 images over 4 rounds, with each round focusing on rating one of the four aspects. The assessment took about 60 minutes to complete.

5.2 Results and Discussion

Figure 5 shows the percentage of ratings that are rated as Good and Very Good for each of the 16 parameter combinations, across the two site conditions. Each cell in the table represents 20 ratings, i.e. 10 images rated by 2 designers. For completeness, example images corresponding to the the parameter combinations are shown for the two site conditions in Appendix A.2 (empty room) and Appendix A.3 (with furniture).

For the empty room, there is tradeoff between meeting user requirements and adhering to original building elements. The optimal depth weight is 0.5, as user requirements (e.g. sofa and wood table) are not met when depth >= 0.75. For the site condition with furniture, the optimal depth weight is >= 0.75, as it improves compatibility with user requirements, functional criteria (e.g. ergonomics), and aesthetics.

Overall, the quality assessment is promising as the overall quality ratings can achieve 70-90% of Good to Very Good design alternatives across the two site conditions. For a set of 20 AI-generated designs, this represents generating 14-18 good to very good design alternatives in about 1 minute on a single desktop PC GPU, compared to hours of designer work required to create a single, high-quality design alternative.

6 STUDY #3: SELF-GUIDED DESIGN EXPLORATION STUDY BY OWNERS

Divergent thinking, corresponding to *breadth* in creativity, involves generating a wide range of ideas, while *convergent thinking*, corresponding to *depth*, focuses on selecting and refining the most promising ones. These two modes of thinking are quite direct and complementary, with divergent thinking ensuring all possibilities are considered and convergent thinking guaranteeing the most viable solutions are developed. Together, they drive successful design exploration. [4, 22, 24, 40, 52]

To understand RoomDreaming's user experience for self-guided design exploration by owners, we conducted a study to compare RoomDreaming to two baselines. The first compares RoomDreaming to existing practices, i.e. participants can use any of their current approaches, such as image search engines. The second compares RoomDreaming vs. the same generative-AI capabilities of RoomDreaming *without* its support for iterative design process.

6.1 Study Design and Procedure

The study used a within-subjects design, with each participant comparing the experience of using RoomDreaming vs. one of the two baselines in counter-balanced ordering. For the baseline of existing online tools, participants freely chose their preferred tools. For the AI-generation baseline, we provided all the RoomDreaming UI and features, except the following two features that explicitly supported the iterative design process: 1) *Liking* and *Bookmarking* no longer affected the prompts, and 2) the *New Design Directions* slider was removed. In order to control for system response time when fewer prompts would be requested from the GPT API in RoomDreaming, we pre-fetched and cached 300 prompts from GPT at the beginning of each study session and use them for the two AI conditions for consistency.

For each conditions, participants spent 20 minutes exploring design alternatives for their project, followed by a 10-minute semi-structured interview. After completing the exploration with both conditions, we conducted a final 30-minute semi-structured interview. The entire study took about 120 minutes, with the first 30 minutes being an introduction to the study and becoming familiar with RoomDreaming.

The concluded interview delved into method comparisons and preferences, encompassing overall *efficiency* in understanding their design preferences, *satisfaction* with final designs from both methods and also exploring aspects like *breadth* and *depth*.

6.2 Participants

In order for the design exploration to be part of a real interior design project, we screened for participants who had already planned to design a residential or commercial space in the next 12 months, yet have not finalized their preliminary design directions.

We recruited a total of 12 owners, with 6 for each of the two baselines, comprising 6 males and 6 females with ages ranging from 25 to 53. Owners focused on their living rooms (x8), bedrooms (x3), and a psychological counseling studio (x1). 5 owners had already cooperated with interior designers for 1~2 months.

6.3 Results: RoomDreaming vs. Existing Tools

Participants used a variety of existing tools for the baseline condition: 4 used Pinterest, 4 used Google Image Search, 2 used YouTube, and 1 browsed a website of designers' portfolios.

6.3.1 Breadth and Depth of Exploration.

As shown in Figure 6, 4 participants preferred RoomDreaming for breath of exploration, with its automatic expansion of prompts to introduce new design directions. Also, *"unlike Google Images, where I struggle due to a lack of suitable keywords, RoomDreaming allows me to explore designs without typing any keywords and discover new designs simply by using "like" and "bookmark"."* (O4) On the other hand, one participant preferred existing tools, because *"the designs in the new iteration of RoomDreaming were too similar to previously liked and bookmarked images."* (O7) We have addressed this issue by improving prompt keyword ordering in version 2 of our system, as described in the next section.

5 participants preferred RoomDreaming for depth of exploration, as it empowered them to *"expanded on ideas I liked to test them."* (O6).

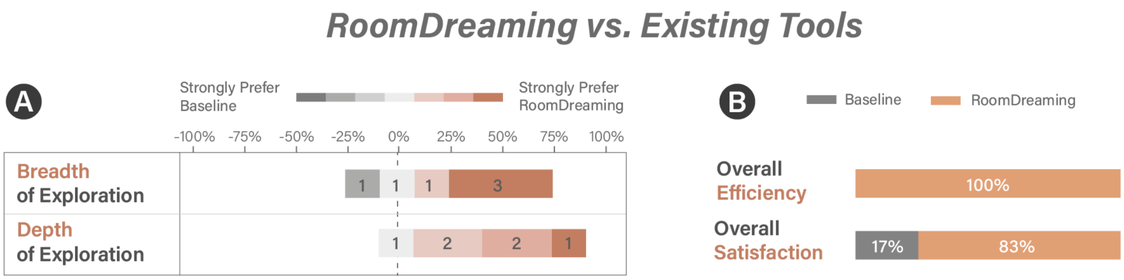

Figure 6: User preference for RoomDreaming vs. a baseline of current exploration tools: (A) Preference rating on a 7-point Likert scale for breadth and depth of exploration. (B) Overall preference for design exploration efficiency and satisfaction.

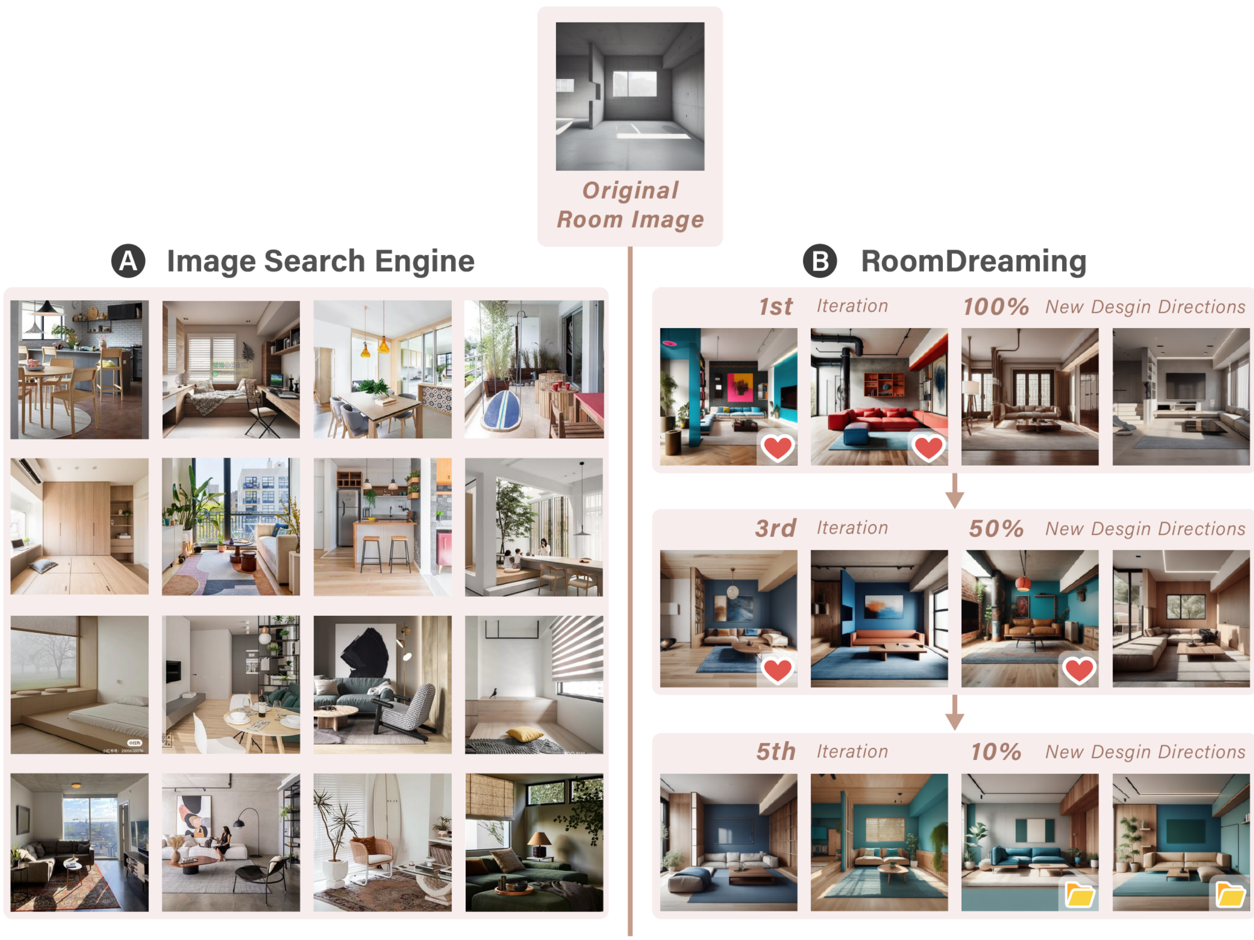

Figure 7: Actual images from one of the participants (O9) in the self-guided exploration study: (A) shows 16 of the 18 liked images collected using Pinterest, showing a wide range of design ideas not integrated and not matched to the participant's room; and (B) shows RoomDreaming designs that match the layout of the participant's room, with 4 examples out of 20 generated designs from each of the 1st, 3rd, and 5th iterations. The breadth in the 1st iteration helped the participant discover preference for bold colors, and iterated from *liking* designs in earlier iterations to *bookmarking* designs in the 5th iteration.

Participants also reported that the combination of breadth and depth was helpful: *"I originally thought I preferred a Japanese Zen style, but after using RoomDreaming, it generated images that combined Japanese and minimalist styles, reminding me that I had liked other design directions before but had forgotten about them."* (O5)

Figure 7 shows a case study of the actual design exploration of one of the participants (O9) from the study, showing sample images collected via Pinterest and RoomDreaming, from *liking* designs in earlier iterations to *bookmarking* designs in later iterations.

6.3.2 Overall Efficiency and Satisfaction. All 6 participants preferred RoomDreaming for overall efficiency, as each participant was able to explore between 100~120 designs in 20 minutes. Furthermore, 4 participants (O4, O5, O8, O9) specifically mentioned that while 20 minutes was insufficient for design exploration using current tools, it was adequate when using RoomDreaming *"as it incorporated the spatial layout of my room and generated new designs based on my preferences."* (O4). Furthermore, the designs based on room layout *"helped me made decisions more quickly and I plan to use these designs to discuss with my interior designer."* (O5)

For overall satisfaction, 5 participants preferred RoomDreaming, and reported that it *"allowed for a more effortless exploration"* (O6). The 1 participant who preferred existing tools mentioned that *"although the images generated by RoomDreaming were all beautiful, the rationality of size and layout configuration were better in photos of actual rooms."* (O7) Spatial rationality is indeed a key limitation of current generative-AI technologies for architecture design, and an active area of AI research.

6.4 Results: RoomDreaming vs. Generative-AI

6.4.1 Breadth and Depth of Exploration. As shown in Figure 8, 4 and 5 participants preferred RoomDreaming for breadth and depth of exploration, respectively. Participants mentioned that RoomDreaming is suitable for early exploration (O11, O12, O14), *"if users were unsure about their preferences, RoomDreaming was highly suitable to find the design directions they like."* (O14). Also, it offered improved control of design directions compared to the baseline (O10~O13, O15): *"RoomDreaming accurately presented what I desired during the iteration process based on my likes and bookmarks. It helped me confirm whether I genuinely liked that direction."* (O11)

However, one participant preferred the baseline for breadth *"because it consistently diverged during the exploration. At this stage, I needed lots of ideas and preferred having convergence in my own mind."* (O15) Note that in this case, the slider for adjusting the ratio of New Design Directions can always be set to 100%, which would provide the same breadth as the baseline condition.

Figure 9 shows a case study of the actual design exploration of one of the participant (O12) from the study. It shows four randomly sampled images for each of the 1st, 3rd, and 5th iteration for both conditions. While all generated images matched the physical room, the baseline AI diverged in design directions throughout, resulting in liked but no bookmarked designs. In contrast, RoomDreaming iteratively converged towards designs that the participant bookmarked, and *"some of the recommended designs are pleasant surprises that expand the acceptable designs that I like. For example, in the 5th iteration, RoomDreaming suggested a new light blue material that is unexpectedly well-suited for my room."* (O12)

6.4.2 Overall Efficiency and Satisfaction. 5 participants preferred RoomDreaming for overall efficiency. *"Thanks to the control slider that enabled generated designs to converge towards my desired direction, which helped me efficiently spend time exploring more possibilities."* (O13)

For overall satisfaction, 5 participants preferred RoomDreaming. However, one participant reported she preferred the overall efficiency and satisfaction in the baseline version as *"I felt like typing directly might be better because I had a pretty good idea of the design direction I want."* (O13) On the other hand, *"designs by RoomDreaming sometimes deviated from my initial, envisioned directions and sparked a curiosity to explore different styles beyond my original plans. I would like to continue to explore more using RoomDreaming!"* (O11)

To help assess whether participants were able to generate desired designs, we observed that the bookmarked designs totaled 6 for baseline AI vs. 66 for RoomDreaming. In particular, 4 participants (O10~O13) had not bookmarked any designs using the baseline AI, mentioning that it was harder to control the designs to their desired directions using prompts.

7 STUDY #4: SYSTEM IMPROVEMENT

The self-guided study provided valuable feedback for improvement from owners' perspective. To understand how RoomDreaming can better support co-design exploration, we conducted a study with 6 interior designers to collect their feedback and suggestions. The study design is based on the Self-guided Exploration study in the previous section, with a different set of semi-structured interview questions that focused on the co-design use case.

We recruited 6 interior designers (D12~D17), 3 males and 3 females, with ages ranging from 24 to 42. Their professional design experience ranged from 5-12 years (mean=8.35, SD=3.3). The design projects they provided, encompassed 3 residential design projects, which were living rooms (D12, D16, D17), and 3 commercial design projects, focusing on a clinic (D13), merchandise exhibition (D14), and a store (D15).

7.1 Feedback and RoomDreaming V2 Improvements

All designers immediately recognized that RoomDreaming would help improve their understanding of owners' preferences and dislikes, and lead to more concrete and efficient discussions. *"RoomDreaming is a bit like a personality test, helping homeowners explore the design they want and facilitating designers in understanding what they like and dislike."* (D13) At the same time, some mentioned concerns with the spatial rationality and ergonomics with AI-generated designs, which may mislead homeowner's expectations (D13, D14, D15).

Combining the feedback from the self-guided owner study and this study, we describe three RoomDreaming V1 limitations and how we addressed them in V2:

7.1.1 Generated designs being too similar to Likes and Bookmarks. Participants reported that the generated designs in the next iteration being *"too similar"* (O7), *"repetitive"* (D15), and *"converged too fast"* (D12). In addition to using random seeds, we further increased design variation within the same design direction, by shuffling the

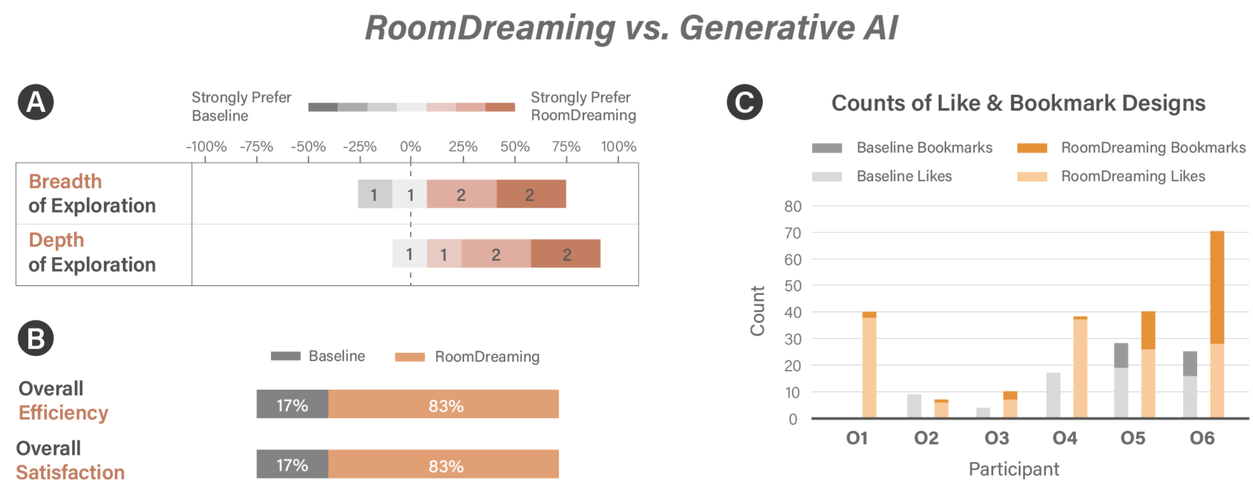

Figure 8: Design exploration using RoomDreaming vs. a baseline of generative-AI without support for iterative design process: (A) Preference rating on a 7-point Likert scale for breadth and depth of exploration; (B) Overall preference for efficiency and satisfaction. ; and (C) Total number of likes and bookmarks by each participant, showing higher number of likes and bookmarks for RoomDreaming.

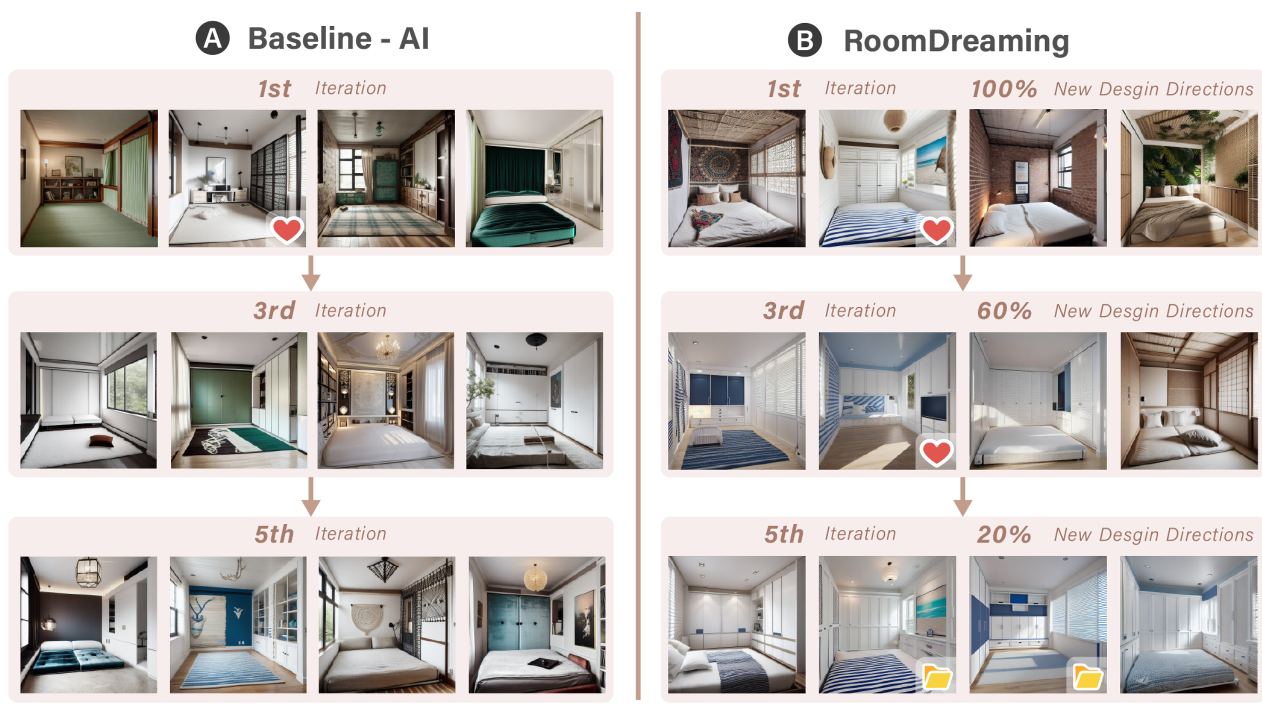

Figure 9: This figure showcased a real user case (Owner-12) in the study, comparing Baseline (AI-approach) and RoomDreaming. Over three iterations (1st, 3rd, and 5th), with four design alternatives sampled in each, the baseline AI continued to exhibit divergent design directions even in the 5th iteration. In contrast, RoomDreaming showed convergence to the owner's desired design direction by the 3rd iteration, and in the 5th iteration, suggested variations within the preferred design directions: *"a new light blue material that is unexpectedly well-suited for my room.""* (O12)

order of the descriptors in the liked prompts rather than using them as is.

7.1.2 Lack of support for negative user requirements. Participants mentioned the need to express specific negative requirements for colors, furnishing, etc. For example, *"I found images I liked in several iterations, but I didn't want that many cushions in my living room."* (O4) Also, *"during the initial exploratory phase, we generally want to see as much as possible, more like an 'addition' approach to design. However, when convergence begins, the design process shifts to a 'subtraction' approach."* (D14). In V1, the user requirements were passed to the default prompt of Stable Diffusion. We added a text field to the UI for negative requirements to utilize Stable Diffusion's negative prompt arguments.

7.1.3 Long batch generation time. In V1, users had to wait for the entire batch of 20 design alternatives to be generated before starting the next iteration, which took about 1 minute. Participants reported that the wait time was too long when they have *"clear ideas to try"* (O4) and *"test"* (D14). To make the system more responsive, we added the ability to start the next iteration by interrupting the previous batch generation.

8 STUDY #5: OWNER-DESIGNER CO-DESIGN EXPLORATION STUDY

Interior Design Illustrated (Ching, 2018) [9] describes the following two key collaborations between interior designers and owners, which we also learnt from our formative studies:

- **Identification of Owners' Needs**, owners convey their requirements and design preferences to the designers, and interior designers engage in understanding their expression.
- **Design Alternative Assessment with Owners**, the presentation of design proposals, including educating the owners about building systems and the assessment results of budget considerations, construction requirements, spatial rationality, and more.

While the previous works focused on the collaboration and communication between Humans and AI, there's one research [12] highlights communication issues between artists and clients, including ambiguity in artistic descriptions, challenges in interpreting artifact instances, managing client expectations, and the need for effective boundary objects in artistic communication.

Meanwhile, our study focused on assessing the potential improvement by RoomDreaming in addressing communication issues outlined in the formative study (referenced as 3.2) between homeowners and interior designers.

8.1 Study Design and Procedure

In this study, pairs of an owner and an interior designer collaboratively used RoomDreaming to aim to achieve a mutually satisfactory preliminary design direction, simulating the current initial discussion process in the stages of *"Programming"* and *"Plan Arrangement"*.

Because of the asymmetry in prior collaboration experiences, with designers having had worked with many owners vs. owners having limited, to no, prior experiences, our interviews additionally asked designers to compare using RoomDreaming to their current practices and extensive, prior co-design experiences.

After introducing the paired participants to each other, we introduced the study, and let the participants discuss basic background and initial requirements for 10 minutes. Owners and interior designers then co-designed using RoomDreaming for 60 minutes to explore owners' preferred interior designs, after which we conducted a semi-structured interview for their feedback. The study took about 120 minutes to complete.

8.2 Participants

We independently recruited 4 homeowners (O16~O19) and 4 interior designers (D18~D21), comprising 3 males and 5 females with ages ranging from 26 to 52, and randomly paired the owners and designers into groups of 2 (G1, G2, G3, G4). The 4 interior designers had professional design experience ranging from 5-15 years (mean=8.4, SD=3.2), specializing in residential, commercial, workspace, and architecture design. The 4 owners were interested in designing 3 residential (O16~O1) and 1 commercial projects (O19), and had not collaborated with interior designers.

8.3 Results

Unfortunately, the last owner-designer pair experienced GPT API downtime during the study, preventing their use of RoomDreaming. Consequently, the results reported will only cover the remaining three groups.

8.3.1 Identification of Owners' Needs. Designers mentioned that through observing owners using RoomDreaming, they could identify owners' needs and wishes *"faster and more accurately"* (D19) with *"less effort in guidance"* (D19) compared to existing methods. They also *"noticed that owners were able to explore specific design aspects more in-depth"* (D19) compared to current methods, and enabled *"owners to express their preferences more quickly and accurately because the images align closely with the original room layout. While the designs generated by RoomDreaming aren't perfect, and more like 80/100, they really were based on the client's preferences."* (D18)

One designer commented that co-designing using RoomDreaming *"helped me understand owner's thought processes more accurately, often revealing that owner initially emphasized certain elements verbally but prioritize differently."* (D20)

An unexpected behavior that we observed was that designers started to correctly guess which images in the new batch of 20 images the owners would Like/Bookmark, ahead of the owners doing so. We noted this behavior in the 3rd, 5th, and 4th iterations, which corresponded to about 30 minutes into using RoomDreaming.

8.3.2 Design Alternative Assessment with Owners. Designers commented on the accelerated pace to start discussing assessment and feasibility, including budget and construction, with owners much earlier than the current process. *"With RoomDreaming, feasibility issues arose quickly, allowing direct and concrete communication with owners in real-time...for example, in the 3rd iteration, I start to estimate the budget for owners"* (D18) *"When owner generated designs with costly materials and elements, I could directly ask them whether to substitute with other options. From my experience with similar cases, normally would need 8~15 weeks to have the same level of discussion."* (D19)

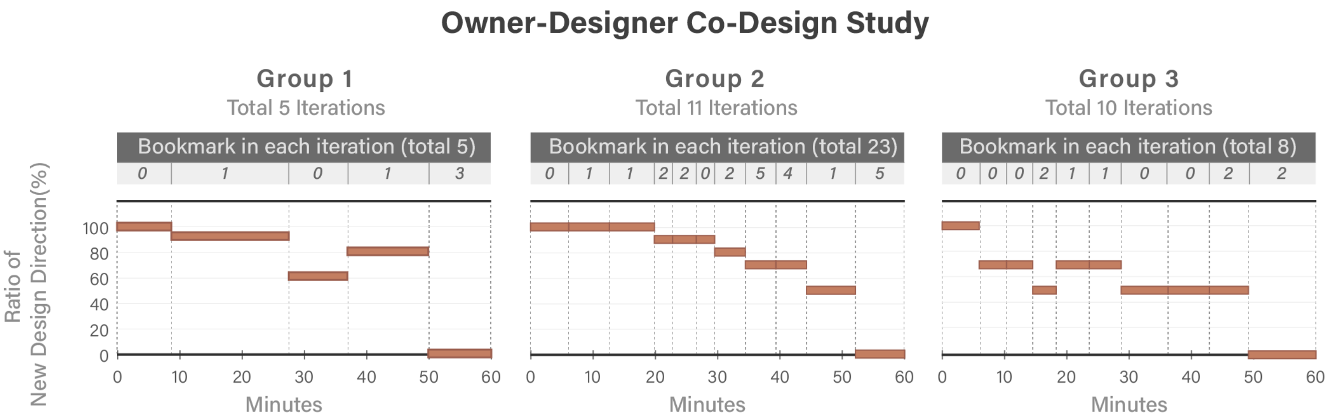

Figure 10: Number of bookmarks saved for each iteration and the ratio of New Design Directions during the 1-hour co-design exploration. 25 vs. 11 bookmarks were saved in the last 30 vs. the first 30 minutes, suggesting that participants were able to generate more desired designs over time. Also, the ratio for New Design Directions lowered over time as participants converged on their preferred design directions.

Table 2: Designers' estimate of the work time saved by co-designing using RoomDreaming for one hour in the study, for each of the two design stages: A) *Identification of Owners' Need*: understanding owners' design preferences and requirements; and B) *Develop and Refine Design*: developing plans, elevations, sections, and details.

Interior Design Project	**Identification of Owners' Needs** (Time Saved)	**Develop and Refine Design** (Time Saved)
Group1: 15m^2 living room with floor plan	≈ 3 working days	≈ 1 working day
Group2: 18m^2 empty living room	≈ 8~15 working days	≈ 14~16 working days
Group3: 10m^2 empty bedroom	≈ 2.5 working days	≈ 1.5~4 working days

8.3.3 Estimation of Time Saved. Designer commented on the time saved using RoomDreaming, saying *"previously, it took a week to create 3D models based on owner requirements, and meetings often required changes that takes another week to re-render. RoomDreaming is instant and more design iterations increase precision to preferences."* (D19)

Table 2 shows designers' estimate of the work time saved by co-designing using RoomDreaming for one hour in the study. RoomDreaming saved the equivalent of 2.5~15 working days of traditional owner-designer meetings and preparation, and 1~16 working days on developing and refining designs. The total time saved per project ranged from 4~31 working days.

9 DISCUSSION, LIMITATIONS, AND FUTURE WORK

9.1 Designing for Human + AI

Our goal through all this research and user studies has been learning how to best leverage generative AI, that generates designs quickly with inconsistent quality (at the moment), to *augment* human designers, who develop designs much more slowly but at consistently higher quality.

The insight we have learnt is that for use cases where the inconsistent quality has low costs in terms of user experience, generative-AI can significantly enhance the user experience. In the case of RoomDreaming, AI is at least 1000x faster in design generation (3 seconds vs. 3 hours), but currently can only produce 80% good designs, meaning that 20% or more of the generated designs are not acceptable. Nevertheless, because our browsing UI keeps the cost of seeing poor-quality designs low, users simply ignore and scroll past them. Thus, the tradeoff between speed and quantity vs. quality, works well for preliminary design exploration for both owners and designers.

We are happy to share that in addition to many owners wanting to continue using RoomDreaming after the user studies, 3 designers from the studies have inquired multiple times whether they could use RoomDreaming for their projects, including a design director who wants all 5 of their interior designers to start using RoomDreaming.

9.2 Tailoring to Region-specific Preferences

A designer noted RoomDreaming's limitation in recognizing region-specific preferences, leading to generated designs that, based on her experience, she knew would not appeal to owners in this particular city. Compared to RoomDreaming, she preferred the efficiency of using her library of interior designs, which she had curated over many year to match the popular owner preferences in our region. Even so, she found RoomDreaming *"helped to understand the logic and reasons behind each owner's preference."* (D20) This insight highlighted the opportunity to explore location-based tuning of the Prompt Composer, to generate designs that owners in the region are more likely to like. One challenge would be identifying optimal balance between breadth and region-specific preferences.

9.3 Creativity Control

Some designers felt that their creativity was constrained because the current designs generated by RoomDreaming reflect popular and common designs, *"RoomDreaming has difficulty achieving inspiring designs that are non-typical"* (D15), whereas they would like *"RoomDreaming to generate highly imaginative and unconventional ideas and challenge me to think about how to implement them."* (D14) In order to support higher and possibly extreme creativity, we are exploring ways to design prompts to create more imaginative designs, and to provide such control over creativity to users.

9.4 Element-specific Preference and Generation

RoomDreaming currently supports user preference of the entire design. Owners and designers have requested the ability to indicate preferences for specific elements in an image, such a lamp on a table, and also the ability to specify *dislikes* via the UI, rather than through negative keywords. In addition to support these, we are also exploring ways to support the ability to select and modify specific components within the image, such as showing 20 different styles of lamps on this table without affecting all other design elements.

9.5 Spatial Rationality and Multi-room Support

A key limitation of current generative AI technologies for architecture is spatial rationality and ergonomics. For example, currently, a bed that is aesthetic but too large for the bedroom may be rendered. While current technologies are helpful for preliminary design exploration, major progress on spatial rationality, which is currently a challenging and active topic for AI research, would be needed in order to further support the subsequent design process, such as floor plan generation and the implementation phase of construction and budget.

Beyond single-room design exploration, we are exploring multi-room support, such as the exploration of spatial proportion and designs of adjacent spaces (e.g. a bedroom and its connecting bathroom), to extend *RoomDreaming* into *HomeDreaming*.

10 CONCLUSION

We have proposed, designed, implemented, and evaluated RoomDreaming, a generative-AI approach aimed at facilitating iterative, preliminary interior design exploration. Inspired by advancements in generative-AI and the persistent challenges in existing design processes, we developed RoomDreaming to facilitate iterative and efficient exploration of design alternatives. Through an iterative design process and a series of formative and summative studies involving 18 homeowners and 20 interior designers (with a combined professional experience of 112 years), we have fine-tuned the system to align with users' needs and preferences. The results from our studies underscore the potential of RoomDreaming to accelerate the design process, enabling users to quickly explore a vast array of design alternatives more broadly and deeply, and improve communication between owners and designers.

ACKNOWLEDGMENTS

REFERENCES

[1] Stability AI. 2022. Stable Diffusion. https://ommer-lab.com/research/latent-diffusion-models/ Last accessed 18 June 2023.
[2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300233
[3] Andrew. 2022. Stable Diffusion prompt: a definitive guide. https://stable-diffusion-art.com/prompt-guide/
[4] Glory E Aviña, Christian D Schunn, Austin R Silva, Travis L Bauer, George W Crabtree, Curtis M Johnson, Toluwalogo Odumosu, S Thomas Picraux, R Keith Sawyer, Richard P Schneider, et al. 2018. The Art of Research: A Divergent/Convergent Thinking Framework and Opportunities for Science-Based Approaches. *Engineering a Better Future: Interplay between Engineering, Social Sciences, and Innovation* (2018), 167–186.
[5] Andrew Cantino. 2021. Prompt Engineering Tips and Tricks with GPT-3. https://blog.andrewcantino.com/blog/2021/04/21/prompt-engineering-tips-and-tricks/. Accessed: December 4, 2023.
[6] Stanislas Chaillou. 2020. Archigan: Artificial intelligence x architecture. In *Architectural Intelligence: Selected Papers from the 1st International Conference on Computational Design and Robotic Fabrication (CDRF 2019)*. Springer, Springer Nature Singapore, Singapore, 117–127.
[7] Guangming Chen, Guiqing Li, Yongwei Nie, Chuhua Xian, and Aihua Mao. 2016. Stylistic indoor colour design via Bayesian network. *Computers & Graphics* 60 (2016), 34–45.
[8] Kang Chen, Kun Xu, Yizhou Yu, Tian-Yi Wang, and Shi-Min Hu. 2015. Magic decorator: automatic material suggestion for indoor digital scenes. *ACM Transactions on graphics (TOG)* 34, 6 (2015), 1–11.
[9] F.D.K. Ching and C. Binggeli. 2012. *Interior Design Illustrated*. Wiley, New Jersey, US. https://books.google.com.tw/books?id=q3N07SkP8OYC
[10] Francis DK Ching. 2023. *Architecture: Form, space, and order*. John Wiley & Sons, New Jersey, US.
[11] Ji Young Cho and Joori Suh. 2020. Spatial Color Efficacy in Perceived Luxury and Preference to Stay: An Eye-Tracking Study of Retail Interior Environment. *Frontiers in Psychology* 11 (2020), 296. https://doi.org/10.3389/fpsyg.2020.00296
[12] John Joon Young Chung and Eytan Adar. 2023. Artinter: AI-Powered Boundary Objects for Commissioning Visual Arts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1997–2018. https://doi.org/10.1145/3563657.3595961
[13] Stephanie Clemons and Joan McLain-Kark. 1991. Computer–aided design in interior design programs: Status and challenges. *Journal of Interior Design Education and Research* 17, 2 (1991), 47–50.
[14] Jodie Cook. 2023. How To Write Effective Prompts For ChatGPT: 7 Essential Steps For Best Results. https://www.forbes.com/sites/jodiecook/2023/06/26/how-to-write-effective-prompts-for-chatgpt-7-essential-steps-for-best-results/?sh=3f4e57832a18. Accessed: December 4, 2023.
[15] Hai Dang, Lukas Mecke, Florian Lehmann, Sven Goller, and Daniel Buschek. 2022. How to Prompt? Opportunities and Challenges of Zero- and Few-Shot Learning for Human-AI Interaction in Creative Applications of Generative Models. arXiv:2209.01390 [cs.HC]
[16] CODDINGTON DESIGN. 2023. How Long Does an Interior Design Project Take? https://coddingtondesign.com/how-long-design-project-timelines/
[17] GPTBOT.io. 2023. Mastering ChatGPT: How to Craft Effective Prompts (Full Guide). https://gptbot.io/master-chatgpt-prompting-techniques-guide/. Accessed: December 4, 2023.
[18] Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. 2022. Optimizing Prompts for Text-to-Image Generation. arXiv:2212.09611 [cs.CL]
[19] Idit Ed Harel and Seymour Ed Papert. 1991. *Constructionism*. Ablex Publishing.
[20] Rahib Imamguluyev. 2021. Application of fuzzy logic model for correct lighting in computer aided interior design areas. In *Intelligent and Fuzzy Techniques: Smart and Innovative Solutions: Proceedings of the INFUS 2020 Conference, Istanbul, Turkey, July 21-23, 2020*. Springer, Springer International Publishing, Cham, 1644–1651.
[21] Ozge Sever Islamoglu and Kubra Ozlu Deger. 2015. The location of computer aided drawing and hand drawing on design and presentation in the interior

design education. *Procedia-Social and Behavioral Sciences* 182 (2015), 607–612.

[22] Syed Fahad Javaid and James Paul Pandarakalam. 2021. The association of creativity with divergent and convergent thinking. *Psychiatria danubina* 33, 2 (2021), 133–139.

[23] Alexey L Khoroshko et al. 2020. The Research of the Possibilities and Application of the AutoCAD Software Package for Creating Electronic Versions of Textbooks for" Engineering and Computer Graphics" Course. *TEM Journal* 9, 3 (2020), 1141–1149.

[24] Kyung Hee Kim and Robert A. Pierce. 2013. *Convergent Versus Divergent Thinking*. Springer New York, New York, NY, 245–250. https://doi.org/10.1007/978-1-4614-3858-8_22

[25] Vivian Liu and Lydia B Chilton. 2022. Design Guidelines for Prompt Engineering Text-to-Image Generative Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 384, 23 pages. https://doi.org/10.1145/3491102.3501825

[26] Wei Ma, Xiangyu Wang, Jun Wang, Xiaolei Xiang, and Junbo Sun. 2021. Generative design in building information modelling (BIM): approaches and requirements. *Sensors* 21, 16 (2021), 5439.

[27] mordorintelligence. 2023. INTERIOR DESIGN INDUSTRY SIZE & SHARE ANALYSIS - GROWTH TRENDS & FORECASTS (2023 - 2028). https://www.mordorintelligence.com/industry-reports/interior-design-services-market

[28] JD Institute of Fashion Technology. 2023. Interior design process – How long does it take? https://www.jdinstitute.edu.in/interior-design-process-how-long-does-it-take/

[29] Akihiro Ogino. 2017. A design support system for indoor design with originality suitable for interior style. In *2017 International Conference on Biometrics and Kansei Engineering (ICBAKE)*. IEEE, IEEE, Kyoto, Japan, 74–79.

[30] Bo Hyeon Park and Kyung Hoon Hyun. 2022. Analysis of pairings of colors and materials of furnishings in interior design with a data-driven framework. *Journal of Computational Design and Engineering* 9, 6 (2022), 2419–2438.

[31] Jelena Pejic and Petar Pejic. 2022. Linear kitchen layout design via machine learning. *AI EDAM* 36 (2022), e9. https://doi.org/10.1017/S089006042100038X

[32] HOUZZ PRO. 2023. Free Template: Interior Design Schedule & Guide. https://www.houzz.com/pro-learn/blog/startup-guide-interior-design-how-to-build-schedule-with-template

[33] HOUZZ PRO. 2023. How Many Hours Do Interior Designers Work? https://www.houzz.com/pro-learn/blog/startup-guide-interior-design-how-many-hours-work

[34] Sunil Ramlochan. 2023. Enhancing Stable Diffusion Models with ControlNet. https://promptengineering.org/enhancing-stable-diffusion-models-with-control-nets/. Accessed December 5, 2023.

[35] Ramsri Goutham Golla Rohit Ramesh. 2023. ControlNet - Adding control to Stable Diffusion's image generation. https://blog.segmind.com/what-is-stable-diffusion-controlnet/. Accessed December 5, 2023.

[36] Yujie Shu. 2021. Application of Computer Aided Design Software in Interior Design. In *Journal of Physics: Conference Series*. IOP Publishing, IOP Publishing, Bristol, BS2 OGR, UK, 022035.

[37] Nathanael Sitanggang, Putri Lynna Adelinna Luthan, and Felix Andika Dwiyanto. 2020. The Effect of Google SketchUp and Need for Achievement on the Students' Learning Achievement of Building Interior Design. *Int. J. Emerg. Technol. Learn.* 15 (2020), 4–19. https://api.semanticscholar.org/CorpusID:222261212

[38] Charles Spence. 2020. Senses of place: architectural design for the multisensory mind. In *Cognitive Research: Principles and Implications*. Springer Science and Business Media LLC, New York, NY, US. https://doi.org/10.1186/s41235-020-00243-4

[39] Steins. 2023. Stable Diffusion — ControlNet Clearly Explained! https://medium.com/@steinsfu/stable-diffusion-controlnet-clearly-explained-f86092b62c89. Accessed: December 5, 2023.

[40] Lubart Todd. 2016. Creativity and convergent thinking: Reflections, connections and practical considerations. *. :* 4 (2016), 7–15.

[41] Nobuyuki Umezu and Eriho Takahashi. 2017. Visualizing color term differences based on images from the web. *Journal of Computational Design and Engineering* 4, 1 (2017), 37–45.

[42] Dakuo Wang, Elizabeth Churchill, Pattie Maes, Xiangmin Fan, Ben Shneiderman, Yuanchun Shi, and Qianying Wang. 2020. From Human-Human Collaboration to Human-AI Collaboration: Designing AI Systems That Can Work Together with People. In *CHI EA '20: Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3334480.3381069

[43] Lisa K Waxman and Hong Zhang. 1995. Computer aided design training methods in interior design professional practice. *Journal of Interior Design* 21, 1 (1995), 21–29.

[44] Tomer Weiss, Ilkay Yildiz, Nitin Agarwal, Esra Ataer-Cansizoglu, and Jae-Woo Choi. 2020. Image-Driven Furniture Style for Interactive 3D Scene Modeling. In *Computer Graphics Forum*. Wiley Online Library, Hoboken, New Jersey, US, 57–68.

[45] Wiskkey. 2022. The maximum usable length of a Stable Diffusion text prompt. https://www.reddit.com/r/StableDiffusion/comments/wl4cn3/the_maximum_usable_length_of_a_stable_diffusion/?rdt=62942&onetap_auto=true

[46] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. 2018. Unified Perceptual Parsing for Scene Understanding. arXiv:1807.10221 [cs.CV]

[47] Fei Yu, Bang Liang, Bo Tang, and Hongrun Wu. 2023. An Interactive Differential Evolution Algorithm Based on Backtracking Strategy Applied in Interior Layout Design. *Algorithms* 16, 6 (2023), 275.

[48] Loutfouz Zaman, Wolfgang Stuerzlinger, Christian Neugebauer, Rob Woodbury, Maher Elkhaldi, Naghmi Shireen, and Michael Terry. 2015. GEM-NI: A System for Creating and Managing Alternatives In Generative Design. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1201–1210. https://doi.org/10.1145/2702123.2702398

[49] J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 437, 21 pages. https://doi.org/10.1145/3544548.3581388

[50] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. https://doi.org/10.48550/arXiv.2302.05543 arXiv:2302.05543 [cs.CV]

[51] Jie Zhu, Yanwen Guo, and Han Ma. 2017. A data-driven approach for furniture and indoor scene colorization. *IEEE transactions on visualization and computer graphics* 24, 9 (2017), 2473–2486.

[52] Weili Zhu, Siyuan Shang, Weili Jiang, Meng Pei, and Yanjie Su. 2019. Convergent thinking moderates the relationship between divergent thinking and scientific creativity. *Creativity Research Journal* 31, 3 (2019), 320–328.

A IMAGE QUALITY ASSESSMENT BY DESIGNERS

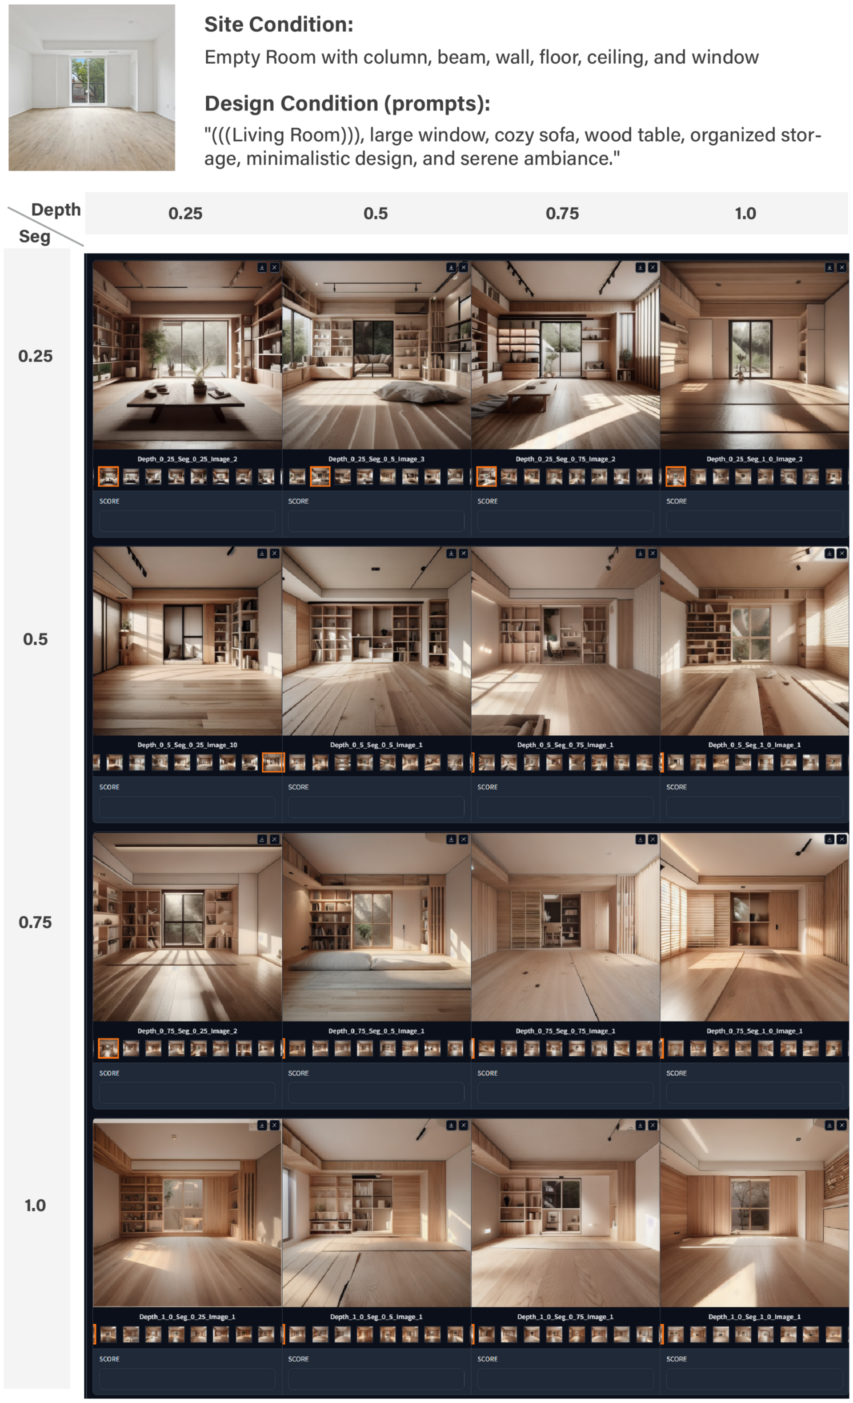

Figure A.1: Segmentation over 0.25

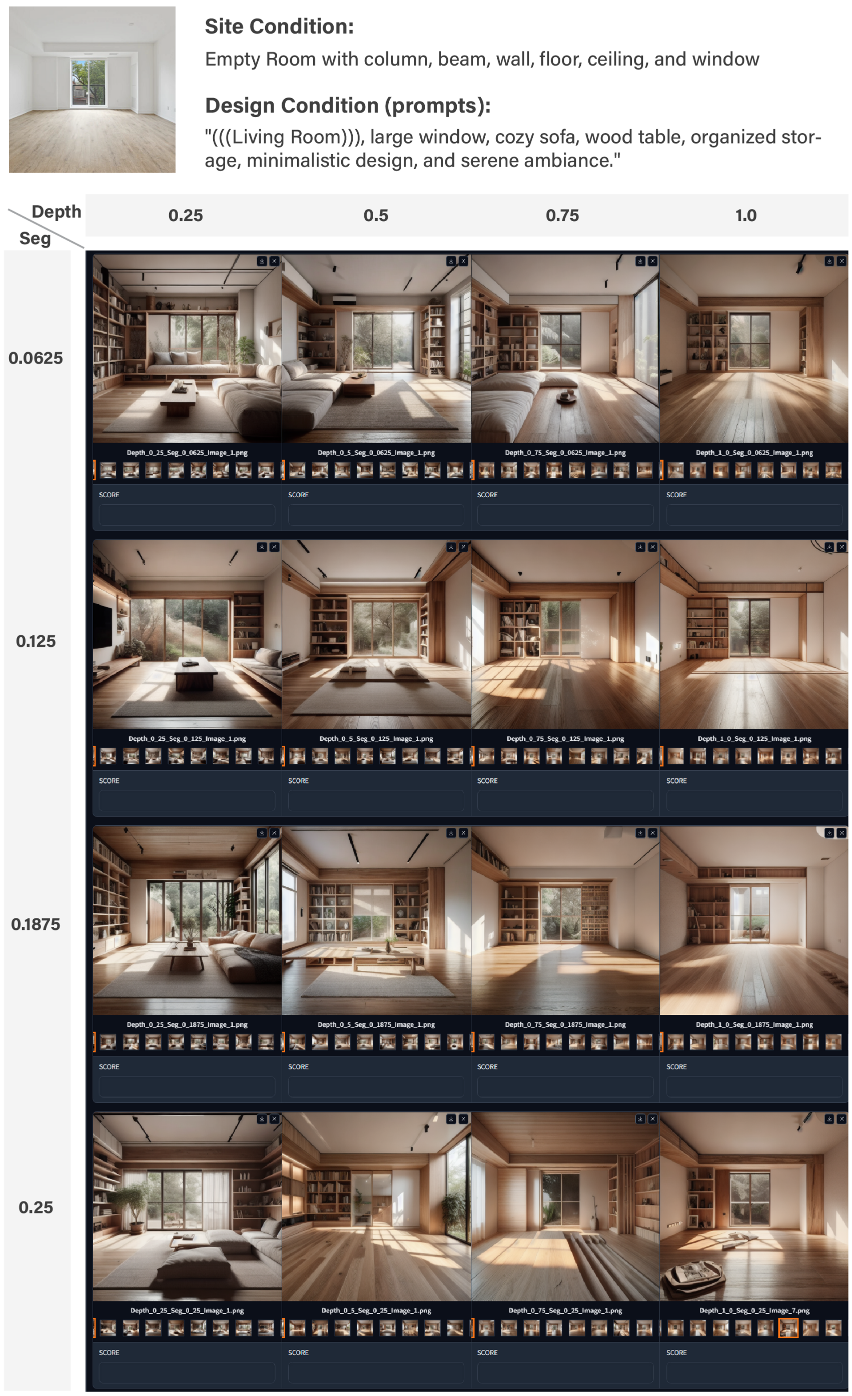

Figure A.2: Site Condition - Empty Room

Figure A.3: Site Condition - Room with exiting furniture