

Co-Layout: LLM-driven Co-optimization for Interior Layout

Chucheng Xiang¹, Ruchao Bao¹, Biyin Feng²,
Wenzheng Wu¹, Zhongyuan Liu³, Yirui Guan³, Ligang Liu^{1*}

¹University of Science and Technology of China

²Tsinghua University

³Tencent

{ccxiang, iambrc, wuwzh}@mail.ustc.edu.cn, fengby23@mails.tsinghua.edu.cn,
{lockliu, ezguan}@tencent.com, lgliu@ustc.edu.cn

Abstract

We present a novel framework for automated interior design that combines large language models (LLMs) with grid-based integer programming to jointly optimize room layout and furniture placement. Given a textual prompt, the LLM-driven agent workflow extracts structured design constraints related to room configurations and furniture arrangements. These constraints are encoded into a unified grid-based representation inspired by “Modulor”. Our formulation accounts for key design requirements, including corridor connectivity, room accessibility, spatial exclusivity, and user-specified preferences. To improve computational efficiency, we adopt a coarse-to-fine optimization strategy that begins with a low-resolution grid to solve a simplified problem and guides the solution at the full resolution. Experimental results across diverse scenarios demonstrate that our joint optimization approach significantly outperforms existing two-stage design pipelines in solution quality, and achieves notable computational efficiency through the coarse-to-fine strategy.

1 Introduction

Interior design significantly impacts our daily living environments. A certified interior designer must adeptly translate a client’s broad requirements into detailed plans for rooms, corridors, and furniture layouts. This process requires designers’ professional expertise and years of accumulated experience. Producing high-quality, detailed interior designs is both challenging and time-consuming, as it involves managing numerous design variables and balancing trade-offs between different design constraints. Developing automated design tools can significantly boost the productivity of interior designers while also stimulating their creativity.

Due to the complexity of interior design, existing approaches often separate the process into distinct phases: room layout (Nauata et al. 2021) and furniture layout (Littlefair, Dutt, and Mitra 2025). In practice, experienced interior designers require simultaneous consideration of both elements, as furniture determines the room’s functionality and the room influences the size and placement of furniture. Neglecting to co-optimize these aspects can lead to undesirable outcomes, such as insufficient space for movement.

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this work, we draw significant inspiration from the recent advancements in large language models (LLMs) (Wei et al. 2022), particularly their impressive performance in generative tasks. For the task of interior design, we discovered that advanced LLM agents can translate users’ high-level requirements into detailed plans involving both room and furniture layouts. However, one significant limitation of LLMs is their inability to generate interior designs with precise coordinates, often resulting in conflicts within the output. To address this, we leverage LLMs to produce semantic constraints between rooms and furniture, and subsequently develop a novel optimization approach to create feasible designs that satisfy these constraints.

Finding a plausible interior layout solution that satisfies LLM-generated constraints is technically challenging. First, modeling these constraints is non-trivial, particularly regarding corridor connectivity and room accessibility. We draw on the concept of “Modulor” from classical architecture theory (Corbusier 2000), which discretizes space into a grid system. Each resulting cell is assigned one room label (e.g., bedroom, living room) and associated furniture label (e.g., bed, sofa) to define spatial programming. Second, the output from LLMs often includes dozens of constraints, and after discretization, the optimization involves thousands of integer variables. To accelerate this process, we propose a coarse-to-fine optimization approach. Initially, the constraints are simplified to be solved on a coarse grid with reduced design variables, and this solution then serves as a warm start for solving the problem on a finer grid.

In summary, this paper makes three contributions which have not been previously demonstrated:

- We propose an automated interior design framework that simultaneously optimizes room layout and furniture placement by satisfying LLM-generated constraints.
- We introduce a grid-based representation to formalize LLM-derived constraints, transforming the design task into an integer programming problem.
- We develop a coarse-to-fine strategy to significantly accelerate the optimization process.

2 Related Work

Floorplan Representation. Recent methods represent floorplans either as rasterized images or as vectorized

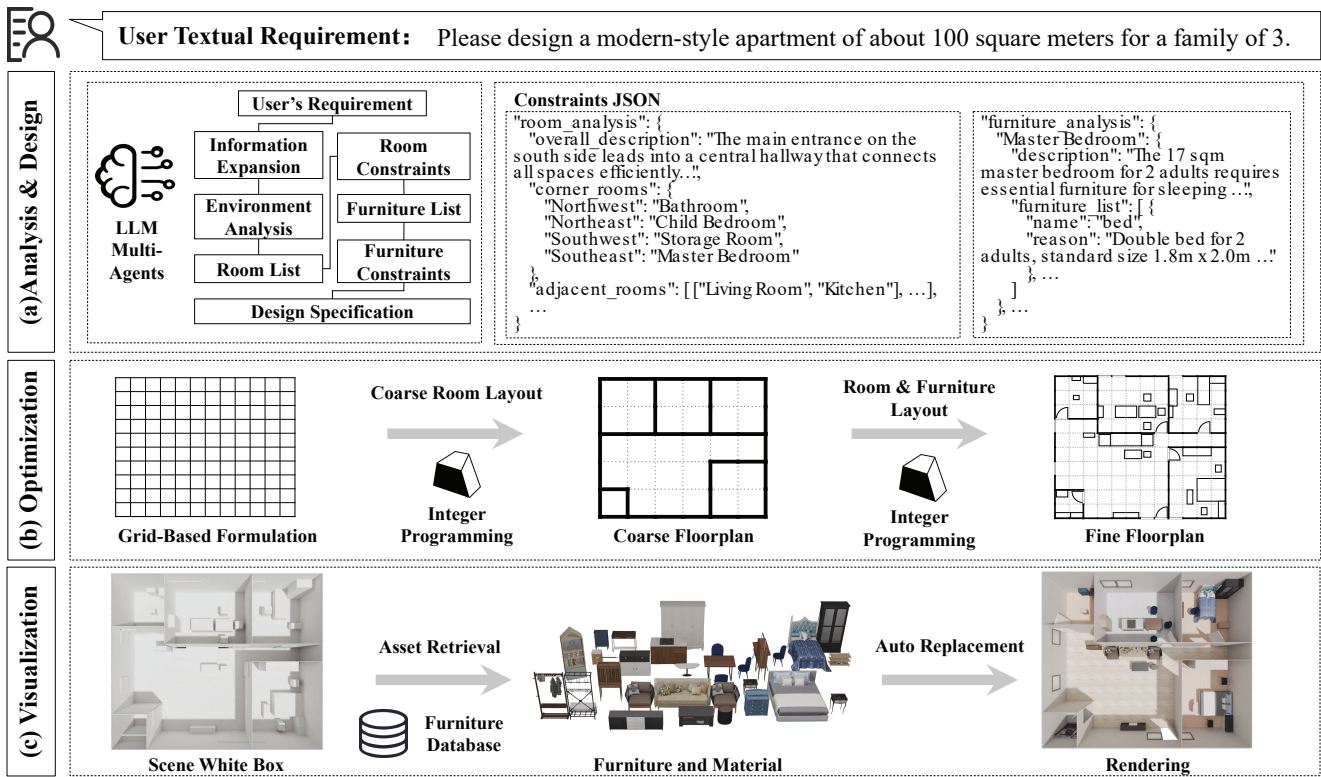


Figure 1: Overview of our automated framework. (a) First, the LLM-based workflow processes user requirements to generate spatial constraints for rooms and furniture, along with boundary conditions for the floor. (b) These constraints are then encoded into an integer programming model using our grid-based representation. The model is efficiently solved through a coarse-to-fine strategy to obtain the final layout. (c) The generated layout is converted into a scene white box in Blender (Community 2025). Suitable assets are then retrieved from the 3D-FUTURE (Fu et al. 2021) dataset and a curated library of over 2,000 assets via semantic embedding and size matching to populate the scene.

graphs. Rasterized approaches treat floorplans as image-like pixel matrices, and characterize rooms as segmentation masks. Representative works include RPLAN (Wu et al. 2019), HouseGAN (Nauata et al. 2020), HouseGAN++ (Nauata et al. 2021) and Graph2Plan (Hu et al. 2020), etc. These representations are limited by resolution. On the other hand, vectorized approaches represent rooms by geometric parameters (e.g., the vertices of the room’s boundary, or the position and size of the room rectangle) (Wu et al. 2018; Para et al. 2021; Liu et al. 2022; Shabani, Hosseini, and Furukawa 2023). This representation better preserves geometric and topological integrity, and supports constraint modeling. Some methods adopt pure optimization pipelines (Wu et al. 2018) or autoregressive generation frameworks (Para et al. 2021) under such representations. However, both rasterized and vectorized representations struggle to describe high-level constraints such as corridor connectivity and room accessibility in our framework. In contrast, grid-based representation (Drira, Pierrel, and Hajri-Gabouj 2007; Peng, Yang, and Wonka 2014; Hua et al. 2019) discretizes space into unit cells, enabling explicit modeling of such spatial constraints. Moreover, it aligns well with the concept of “Modular” (Corbusier 2000),

making it practical for real-world design scenarios.

Automatic Floorplan Generation. Floorplan generation can be categorized based on input modalities, including bubble diagrams and text descriptions. Bubble diagrams (Merrell, Schkufza, and Koltun 2010) offer geometric parameters and constraints, enabling the derivation of optimal layouts. Textual inputs can be processed by LLMs to extract optimization settings (Zong, Zhan, and Tan 2024; Chen et al. 2020; Qin et al. 2024). Models can also be trained to directly map text to floorplans (Leng et al. 2023). The generation approaches can be divided into data-driven methods (Mostafavi et al. 2024; Wu et al. 2019; Hu et al. 2020; Nauata et al. 2020; Wang et al. 2021; Para et al. 2021; Nauata et al. 2021; Sun et al. 2022; Shabani, Hosseini, and Furukawa 2023; Hong et al. 2024) and optimization-based techniques (Wu, Tojo, and Umetani 2024; Wu et al. 2018). However, existing methods often overlook complex constraints, such as downstream tasks like furniture placement, as well as requirements like corridor connectivity and room accessibility. These factors are often handled manually or through post-processing (Wu et al. 2018), instead of incorporating them within the generation framework.




Figure 2: The floor is modeled as a 2D grid. Each cell (i, j) has binary variables indicating if it belongs to a corridor ($p_{i,j}$), room r_k ($x_{i,j}^k$), or is occupied by the l -th furniture in the k -th room $s_{k,l}$ ($f_{i,j}^{k,l}$). Each furniture $s_{k,l}$ has dimensions $W_{k,l} \times L_{k,l}$. Its orientation $\nu_{k,l}$, chosen from 4 axis-aligned directions, is set by two binary variables $\sigma_{k,l}$ and $\mu_{k,l}$.

Indoor Furniture Layout. This task typically involves selecting appropriate furniture based on input room information, and predicting their placement. Some methods are based on optimization, such as defining constraints according to design principles and specifying optimization settings through interactions (Merrell et al. 2011), using physics-based rules (Weiss et al. 2018), and maximizing ergonomic comfort (Yu et al. 2011). Some data-driven methods aim to learn spatial patterns directly from layout datasets, using models such as convolutional networks (Ritchie, Wang, and Lin 2019; Wang et al. 2018), Transformers (Paschalidou et al. 2021), or diffusion models (Tang et al. 2024; Leimer et al. 2022) to predict furniture arrangements. Recent advancements utilize LLMs to interpret natural language prompts, identify relevant scene components and constraints, and guide furniture layout generation accordingly (Yang et al. 2024; Feng et al. 2023; Çelen et al. 2024; Littlefair, Dutt, and Mitra 2025; Fu et al. 2024). However, in interior design, furniture layout and floorplan generation are inherently interdependent. Existing approaches typically treat them as separate stages, making it difficult to achieve synergistic consideration across both levels of design.

3 Overview

Our method converts text input into a detailed interior layout design through two phases, as illustrated in Fig. 1. We first introduce basic notations of a grid-based representation for floorplan modeling. Then, an LLM-based preprocessor is proposed to transfer user input into a structural scene graph. Based on the notations and the scene graph, we formulate the joint optimization of room layout and furniture arrangement under various constraints as an integer programming problem. Finally, we propose a coarse-to-fine strategy to solve the large-scale problem efficiently.

Notations. We adopt a unified grid-based representation for floorplan and furniture arrangement, see Fig. 2 for an illustration. The entire design space for a single floor is discretized into a 2D grid $\mathcal{G} = \{(i, j) : 0 \leq i < W, 0 \leq j < L\}$, where W and L are the width and length of the floor rectangle (with the “Modulor” as the basic unit). In order to

support non-rectangular boundaries, we denote the indoor space as the subset $\mathcal{G}' \subseteq \mathcal{G}$, excluding cells pre-allocated for outdoor space. This indoor space \mathcal{G}' is required to be partitioned into rooms $R = \{r_k\}_{k=1}^N$, corridors, and furniture items $S = \{s_{k,l} : 1 \leq k \leq N, 1 \leq l \leq N_k\}$ within each room. We assign binary variables $x_{i,j}^k$, $p_{i,j}$, and $f_{i,j}^{k,l}$ to each grid cell (i, j) to indicate which spatial components occupy the design space. The orientation of the furniture $s_{k,l}$ is denoted as a vector $\nu_{k,l}$. It only takes the four directions aligned with coordinate axes, identified as $\nu_{k,l}(\sigma_{k,l}, \mu_{k,l}) = (\sigma_{k,l}(1 - 2\mu_{k,l}), (1 - \sigma_{k,l})(1 - 2\mu_{k,l}))$, where $(\sigma_{k,l}, \mu_{k,l})$ are two binary variables. $W_{k,l}$ and $L_{k,l}$ represent the width and length of $s_{k,l}$, which are constants determined by LLMs. For a cell (i, j) , its neighborhood on the grid is $\mathcal{N}(i, j) = \{(i, j + 1), (i, j - 1), (i + 1, j), (i - 1, j)\}$.

LLM-based Preprocessor. Our optimization part takes a scene graph as input, which comprises various constraints on room layouts and furniture arrangements. However, manually designing such a scene graph is labor-intensive. To address this, we leverage LLMs to automate the generation of scene graphs such that users only need to provide a textual idea as input. Specifically, we design tailored prompts for an LLM-based workflow to output structured room and furniture lists R, S along with associated constraints. In this stage, fundamental floor parameters, environmental conditions, room configurations, and furniture constraints are comprehensively considered.

4 Formulation and Optimization

The core of our method is an IP model that co-optimizes room layout and furniture arrangement, based on interior design principles and the grid-based representation. Notably, the positions of doors and windows are not optimization variables; instead, they are determined via post-processing.

4.1 Key Constraints

Our optimization model organizes constraints hierarchically, from basic spatial exclusivity to complex inter-room and furniture placement rules. This structure enables the model to capture essential design interdependencies while remaining computationally efficient.

Non-Overlapping Constraint. This fundamental constraint ensures the spatial exclusivity of elements within the floor. Each cell $(i, j) \in \mathcal{G}'$ must be uniquely assigned to either a room or a corridor segment. This constraint is mathematically formulated as:

$$p_{i,j} + \sum_{k=1}^N x_{i,j}^k = 1, \quad \forall (i, j) \in \mathcal{G}'.$$

The summation guarantees mutual exclusivity and complete coverage of the assignable floor area, which underpins the optimization.

Corridor Connectivity Constraint. Open rooms are defined as rooms without walls, judged by LLMs, such as certain living rooms, which are treated as part of the corridor. To

ensure that the corridor and open rooms form a single connected network, we employ a flow-based formulation. Let $\mathcal{D} = \{E, W, S, N\}$ denote the set of cardinal directions, and let $\text{adj}(i, j, d)$ return the neighboring cell of (i, j) in direction $d \in \mathcal{D}$ if it exists within \mathcal{G}' , and \emptyset otherwise. We define variables $\text{flow}_{i,j}^d \geq 0$ representing the flow from cell (i, j) in direction d . Let R_O denote the set of open rooms. For each cell $(i, j) \in \mathcal{G}'$, we define

$$q_{i,j} = p_{i,j} + \sum_{k \in R_O} \left(x_{i,j}^k - \sum_{l=1}^{N_k} f_{i,j}^{k,l} \right).$$

The connectivity is ensured through flow capacity constraints and conservation principles. Flow can only occur between corridor or open room cells:

$$\text{flow}_{i,j}^d \leq M \cdot q_{i,j}, \quad \forall (i, j) \in \mathcal{G}', d \in \mathcal{D},$$

$\text{flow}_{i,j}^d \leq M \cdot q_{i',j'}, \forall (i, j) \in \mathcal{G}', d \in \mathcal{D}, (i', j') = \text{adj}(i, j, d)$, where M is a sufficiently large constant. We define the outflow and inflow for each cell $(i, j) \in \mathcal{G}'$ as:

$$\text{outflow}_{i,j} = \sum_{d \in \mathcal{D}} \text{flow}_{i,j}^d, \quad \text{inflow}_{i,j} = \sum_{\substack{d \in \mathcal{D}, (i', j') \in \mathcal{G}' \\ (i', j') = \text{adj}(i, j, d^{-1})}} \text{flow}_{i',j'}^d$$

where d^{-1} denotes the opposite direction of d (i.e., $E^{-1} = W$, $S^{-1} = N$). The flow balance constraint below ensures that the entrance cell (i_e, j_e) acts as a source supplying flow equal to the total number of corridor and open room cells, while each such cell consumes exactly one unit of flow:

$$\text{outflow}_{i,j} - \text{inflow}_{i,j} = \begin{cases} \sum_{(i', j') \in \mathcal{G}'} q_{i',j'} & \text{if } (i, j) = (i_e, j_e), \\ -q_{i,j} & \text{if } (i, j) \neq (i_e, j_e). \end{cases}$$

Room Constraints. Room constraints are crucial for generating functional and geometrically sound floorplans.

Accessibility. Each room r_k must be accessible from the corridor or open rooms, ensuring that there exists a path from the entrance to each room. To guarantee unobstructed access, room cells serving as interface points must remain free of furniture placement. We introduce binary auxiliary variables $\alpha_{i,j}^k$ to indicate whether cell (i, j) of room r_k serves as an access point:

$$\sum_{(i,j) \in \mathcal{G}'} \alpha_{i,j}^k \geq 1; \quad \sum_{(i',j') \in \mathcal{N}(i,j) \cap \mathcal{G}'} q_{i',j'} \geq \alpha_{i,j}^k; \quad \alpha_{i,j}^k \leq x_{i,j}^k - \sum_{l=1}^{N_k} f_{i,j}^{k,l}.$$

Notably, the last constraint ensures that an access point can only be established at a room cell (i, j) that belongs to room r_k but is not occupied by any furniture.

Adjacency. For specified room pairs (r_k, r_m) that require spatial adjacency, we enforce the presence of at least one shared boundary. This is modeled using auxiliary binary variables $\beta_{i,j}^{k,m}$ that indicate whether cell (i, j) of room r_k forms an adjacency point with room r_m :

$$\sum_{(i,j) \in \mathcal{G}'} \beta_{i,j}^{k,m} \geq 1; \quad \sum_{(i',j') \in \mathcal{N}(i,j) \cap \mathcal{G}'} x_{i',j'}^m \geq \beta_{i,j}^{k,m}; \quad \beta_{i,j}^{k,m} \leq x_{i,j}^k.$$

Bounding Box. We employ standard big-M formulations to define the axis-aligned bounding box for each room r_k . Let $\gamma_k^{\text{min}_i}, \gamma_k^{\text{min}_j}$ denote the minimum coordinates (along the i and j axes, respectively) and $\gamma_k^{\text{len}_i}, \gamma_k^{\text{len}_j}$ denote the dimensions of the bounding box for room r_k . For all $(i, j) \in \mathcal{G}'$:

$$\begin{aligned} \gamma_k^{\text{min}_i} &\leq i + M \cdot (1 - x_{i,j}^k), \\ \gamma_k^{\text{min}_i} + \gamma_k^{\text{len}_i} - 1 &\geq i - M \cdot (1 - x_{i,j}^k), \\ \gamma_k^{\text{min}_j} &\leq j + M \cdot (1 - x_{i,j}^k), \\ \gamma_k^{\text{min}_j} + \gamma_k^{\text{len}_j} - 1 &\geq j - M \cdot (1 - x_{i,j}^k). \end{aligned}$$

Note that these constraints do not enforce a tight fit which is encouraged later on via the room rectangularity objective.

Furniture Constraints. Furniture constraints ensure the proper arrangement of furniture within room layouts while maintaining functional rationality and spatial efficiency.

Containment. Any grid cell (i, j) occupied by the furniture $s_{k,l}$ is also part of room r_k :

$$f_{i,j}^{k,l} \leq x_{i,j}^k, \quad \forall (i, j) \in \mathcal{G}'.$$

Non-Overlap. Any grid cell (i, j) can be occupied by at most one piece of furniture within the same room:

$$\sum_{l=1}^{N_k} f_{i,j}^{k,l} \leq 1, \quad \forall (i, j) \in \mathcal{G}'.$$

Shape and Size. Each furniture item $s_{k,l}$ is a rectangle with LLM-generated dimensions, which are rounded to the nearest grid cells, denoted by $\ell_{k,l}^i$ and $\ell_{k,l}^j$. We define its bounding box using minimum coordinates $\xi_{k,l}^{\text{min}_i}, \xi_{k,l}^{\text{min}_j}$ and dimensions $\ell_{k,l}^i, \ell_{k,l}^j$, with constraints similar to those for rooms. In addition, the total number of cells occupied by the furniture must equal its area:

$$\sum_{(i,j) \in \mathcal{G}'} f_{i,j}^{k,l} = \ell_{k,l}^i \cdot \ell_{k,l}^j.$$

4.2 Objective Functions

The overall objective function E combines weighted penalty terms to guide the search toward high-quality layouts:

$$E = \sum_s \omega_s \cdot E_s(\mathbf{x}, \mathbf{p}, \mathbf{f}),$$

where E_s represents individual penalty functions and ω_s denotes their respective weights. The objectives are hierarchically organized into three categories based on their impacts.

Geometric Quality Objectives. These primary objectives ensure geometrically sound and aesthetically pleasing room layouts by promoting regular and compact shapes. They also prevent rooms from being fragmented into multiple disconnected components:

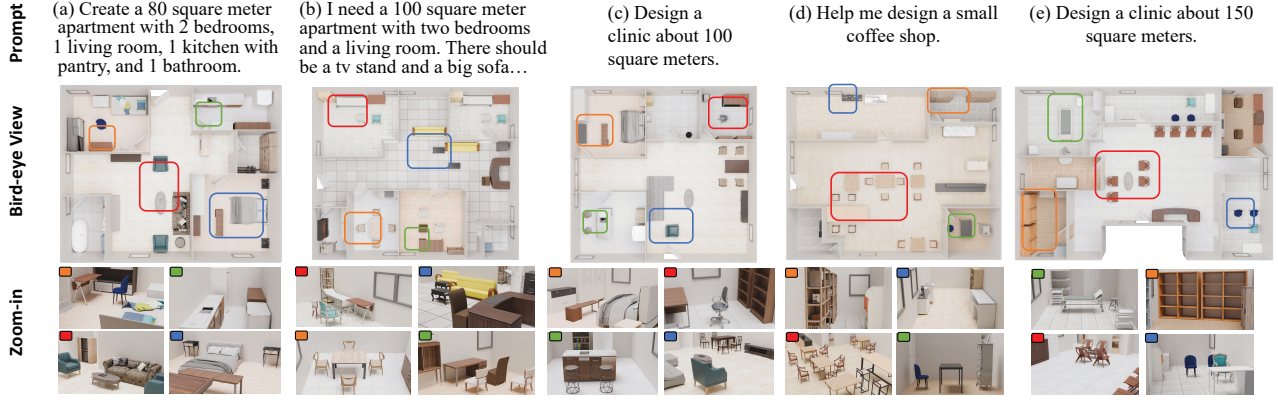


Figure 3: Various examples generated by our method, demonstrating its capability to handle diverse inputs including residential and non-residential spaces. Users can specify requirements such as building type, floor area, room functions, and required furniture. Top: Input text. Middle: Bird’s-eye view of the scene. Bottom: Zoom-in views highlighting key details.

Room Rectangularity. We minimize shape irregularity by penalizing the deviation between each room’s bounding box area and its actual cell count:

$$E_{\text{rect}} = \sum_{k=1}^N \left(\gamma_k^{\text{len}_i} \cdot \gamma_k^{\text{len}_j} - \sum_{(i,j) \in \mathcal{G}'} x_{i,j}^k \right).$$

Room Perimeter. To encourage compact room shapes, we penalize the exposed perimeter of each room:

$$E_{\text{perim}} = \sum_{k=1}^N \sum_{(i,j) \in \mathcal{G}'} x_{i,j}^k \cdot |\mathcal{N}_{\text{out}}(i,j)|,$$

where $|\mathcal{N}_{\text{out}}(i,j)|$ denotes the number of neighboring cells not assigned to room r_k .

Room Functional Objectives. These objectives incorporate interior design principles and user-specified preferences to ensure functional and livable spaces.

Area Target. Specified room area requirements are treated as soft constraints, so we penalize deviations from target areas A_k^{target} :

$$E_{\text{area}} = \sum_{k=1}^N \left| A_k^{\text{target}} - \sum_{(i,j) \in \mathcal{G}'} x_{i,j}^k \right|.$$

Aspect Ratio Control. Extreme room proportions are discouraged through minimizing the difference between room length and width:

$$E_{\text{aspect}} = \sum_{k=1}^N \left| \gamma_k^{\text{len}_i} - \gamma_k^{\text{len}_j} \right|.$$

Furniture Placement Objectives. These objectives ensure functional and rational furniture arrangements by governing spatial relationships and distribution balance.

Spatial Relationships. For furniture pairs with relative positioning requirements, we minimize deviation from target

arrangements:

$$E_{\text{rel}} = \sum_{k=1}^N \sum_{l_1, l_2 \in \mathcal{A}_{\text{rel}}^k} \|\mathbf{c}_{k,l_1} - \mathbf{c}_{k,l_2} - \boldsymbol{\delta}_{l_1,l_2}\|_1,$$

where $\mathcal{A}_{\text{rel}}^k$ denotes the set of furniture pairs with specified relative positions in room r_k , $\mathbf{c}_{k,l}$ denotes the centroid of furniture $s_{k,l}$ and $\boldsymbol{\delta}_{l_1,l_2}$ is the target offset vector between the furniture centroids.

Distribution Balance. To promote balanced furniture layouts, we minimize the distance between the area-weighted furniture centroid and room geometric centroid:

$$E_{\text{bal}} = \sum_{k=1}^N \left\| \hat{\mathbf{C}}_k - \mathbf{C}_k \right\|_1,$$

where $\hat{\mathbf{C}}_k$ is the area-weighted centroid of all furniture in room r_k and \mathbf{C}_k is the room’s geometric centroid.

4.3 Coarse-to-Fine Optimization Strategy

The joint room-furniture optimization exhibits exponential computational complexity with respect to grid resolution, making large scale instances computationally intractable. To address this challenge, we utilize a coarse-to-fine strategy.

Coarse Phase. The coarse-level formulation leverages spatial down-sampling via grid aggregation. For example, when down-sampling from a 12×10 grid to a 6×5 grid, each coarse cell represents a 2×2 block of fine cells, reducing the problem scale by 75%. At this level, we focus solely on room layout, deliberately deferring furniture placement to facilitate efficient exploration of high-level spatial organization.

Coarse to Fine Mapping. We need to transfer the coarse solution back to the original fine-grained grid. For each fine-level grid cell (i,j) , we determine its corresponding coarse-level grid cell (i',j') . Assuming a down-sampling factor of s_i along the i -axis and s_j along the j -axis, the mapping can be defined as:

$$(i',j') = \left(\left\lfloor \frac{i-1}{s_i} \right\rfloor + 1, \left\lfloor \frac{j-1}{s_j} \right\rfloor + 1 \right), \quad \forall (i,j) \in \mathcal{G}_{\text{fine}},$$

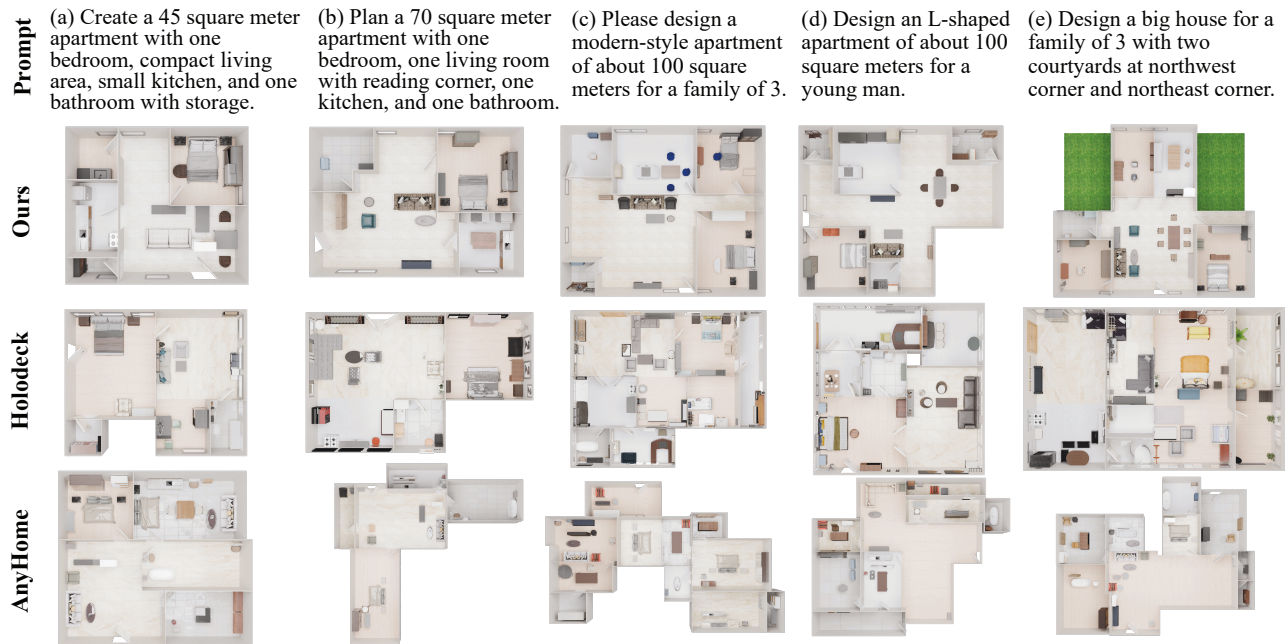


Figure 4: Comparison with baselines. Our approach consistently generates well-structured layouts with full accessibility and clear circulation. In contrast, baseline methods often produce designs with critical flaws, such as illogical circulation paths that violate privacy (e.g., Holodeck in a, d) and unreachable spaces or impractical room shapes (e.g., AnyHome in c, e).

where $\mathcal{G}_{\text{fine}}$ is the set of fine grid cells. If coarse cell (i', j') is assigned to room r_k (i.e., $x_{i', j'}^k = 1$), the corresponding fine cell (i, j) is provisionally assigned to the same room.

Fine Phase. This stage performs the full optimization on the original grid, incorporating both room planning and furniture placement variables. Firstly, the mapped coarse solution provides a *warm-start initialization* for the room assignment variables $x_{i, j}^k$ on the fine grid. However, note that this warm-start is not always feasible; in some cases, it may lead to infeasible solutions. Therefore, we introduce an additional mechanism. Secondly, we add penalty terms to the fine-level objective function to encourage consistency with the coarse-level room assignments. Specifically, for each fine grid cell (i, j) , we penalize if it is not assigned to the same room $k = \pi(i, j)$ (the coarse assignment) by:

$$E_{\text{ref}} = \sum_{(i, j) \in \mathcal{G}_{\text{fine}}} (1 - x_{i, j}^{\pi(i, j)}) \cdot z_{i, j},$$

where binary variable $z_{i, j}$ indicates whether the corresponding coarse cell was assigned to any room or not.

5 Experiments

5.1 Experiment Settings

Environment. We implement our algorithm in Python, using the integer programming solver GUROBI (Gurobi Optimization, LLC 2025), version 12. Experiments were mainly conducted on a laptop equipped with an 8-core Ryzen processor and 16 GB of RAM.

Metrics. We assess the generated floorplans in three key dimensions: physical plausibility, image quality/aesthetics, and text-image alignment.

- **Physical Plausibility:** Measured by the object overlap rate (OOR) and out of boundary rate (OOB). OOR quantifies the percentage of overlapping area between objects, while OOB calculates the proportion of objects extending beyond room boundaries.
- **Image Quality & Aesthetics:** Evaluated using the Q-Align model (Wu et al. 2023) for both image quality assessment (IQA) and image aesthetic assessment (IAA).
- **Text-Image Alignment:** Assessed by the CLIP (Radford et al. 2021) similarity score between the input text and the rendered floorplan image.

Baselines. We compare our co-optimization approach with state-of-the-art two-stage baselines—Holodeck (Yang et al. 2024) and AnyHome (Fu et al. 2024)—which sequentially generate room layouts before populating furniture. For fair comparison, we use their methods solely to generate layout information, convert the results into a unified format, and render them using identical visualization scripts.

5.2 Quantitative Evaluations

The quantitative comparison based on our three evaluation metrics is presented in Table 1. The evaluation is conducted on the five examples shown in Fig. 4. Our method achieves superior performance across all metrics, particularly eliminating object overlaps and boundary violations.

Method	OOD \downarrow	OOB \downarrow	IQA \uparrow	IAA \uparrow	CLIP \uparrow
Holodeck	0.82	2.33	4.03	3.32	25.15
AnyHome	0.00	0.04	4.10	3.32	25.75
Ours	0.00	0.00	4.17	3.35	26.50

Table 1: Quantitative evaluation results. Our method achieves the best performance in terms of physical plausibility, image quality/aesthetics, and text-image alignment.

5.3 Qualitative Evaluations

Method Demonstration. We present representative results of our method in Fig. 3, demonstrating the open-vocabulary generation capability. The gallery begins with two apartments (a) and (b), showcasing our method’s capability in generating residential layouts with detailed requirements. Then, a 100m² clinic (c) demonstrates the method’s capability in professional spaces. A small coffee shop (d) showcases adaptability to commercial scenarios. Finally, a 150m² clinic (e) further validates the generation ability for larger professional spaces. These diverse examples highlight the flexibility of our approach in handling various inputs.

Comparative Analysis. We compare our method with Holodeck and AnyHome in Fig. 4. While baseline methods can generate a rich variety of rooms and furniture, their layouts often exhibit randomness, indicating a lack of robust understanding of room function, privacy, and connectivity. In contrast, our co-optimization approach consistently produces well-structured designs with logical spatial hierarchy and clear circulation paths, where every room is connected to a common corridor. The baselines frequently violate this principle; for instance, Holodeck’s layouts may require passing through a bedroom to reach the main living area (a, d), while AnyHome can generate entirely unreachable rooms (c, e). Furthermore, our method ensures practical room shapes, whereas AnyHome often generates irregular or narrow rooms, and Holodeck’s results lack coherent circulation. In summary, our joint optimization of rooms and furniture ensures the resulting layouts are not only geometrically sound and fully accessible but also functionally coherent, demonstrating a deeper understanding of architectural logic.

5.4 User Study

To evaluate the quality of the generated floorplans, we conducted a user study with 71 participants. After filtering invalid submissions (e.g., completion times under two minutes or uniform high/low scores), we retained 64 valid responses for analysis. The study involved two tasks. First, participants rated the floorplans shown in Fig. 4 on a 5-point scale based on three criteria: semantic alignment, layout rationality, and path clearance. Second, for each prompt, participants ranked the outputs from our method and the baselines. We calculated the average scores for the rating criteria and the Mean Reciprocal Rank (MRR) (Voorhees et al. 1999) for the ranking task. The aggregated results are presented in Table 2, which shows that our method achieved the highest scores across all metrics.

Method	Semantic. \uparrow	Layout. \uparrow	Path. \uparrow	MRR. \uparrow
Holodeck	3.43	3.12	3.06	0.59
AnyHome	3.07	2.59	2.80	0.45
Ours	3.77	3.23	3.41	0.80

Table 2: User study results. Our method achieved the highest scores across all metrics.

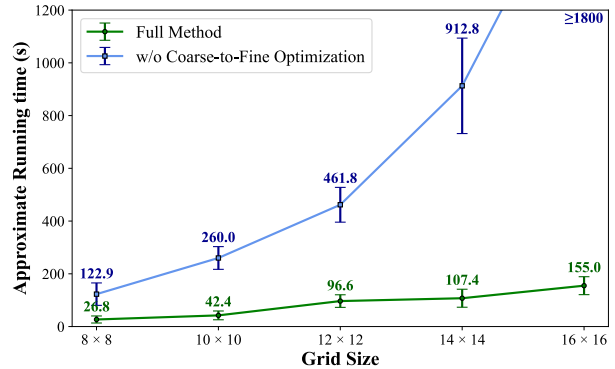


Figure 5: Ablation result of the coarse-to-fine strategy.

5.5 Ablation Study

Ablation of Coarse-to-Fine Strategy. To evaluate the effectiveness of our coarse-to-fine strategy, we conducted an ablation study on its impact on computational efficiency. We tested five times with a fixed prompt, varied the grid resolution, and calculated the mean and standard deviation of the running time required to obtain a feasible solution with a similar optimization gap. As illustrated in Fig. 5, the results show that our strategy significantly reduces computation time as the grid resolution increases, while maintaining comparable solution quality.

6 Conclusion

We present a novel framework that integrates LLMs with integer programming to automate interior layout design, jointly optimizing room layout and furniture placement. By leveraging LLMs to extract structured scene graphs from high-level user input and formulating the design problem on a grid-based representation, our method ensures corridor connectivity, room accessibility, and spatial coherence, outperforming two-stage pipelines in various metrics. A coarse-to-fine strategy further enhances computational efficiency.

Our work still holds limitations that open avenues for future research. First, the scope of furniture is currently limited to floor-based objects. Our framework could be extended to support wall-mounted items and tabletop objects. Second, LLMs may occasionally produce conflicting constraints, degrading solution quality. Iterative refinement or human-in-the-loop feedback could mitigate such situations. Finally, our framework could be extended to design multi-story buildings, by incorporating inter-floor constraints.

Acknowledgments

We thank Dr. Ziqi Wang from The Hong Kong University of Science and Technology for his valuable guidance and advice on this work. We also thank the anonymous reviewers for their appreciation of our work. This work is supported by the National Natural Science Foundation of China (No. 62025207).

References

- Celen, A.; Han, G.; Schindler, K.; Van Gool, L.; Armeni, I.; Obukhov, A.; and Wang, X. 2024. I-design: Personalized llm interior designer. *arXiv preprint arXiv:2404.02838*.
- Chen, Q.; Wu, Q.; Tang, R.; Wang, Y.; Wang, S.; and Tan, M. 2020. Intelligent home 3d: Automatic 3d-house design from linguistic descriptions only. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12625–12634.
- Community, B. O. 2025. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam.
- Corbusier, L. 2000. *The modulor: a harmonious measure to the human scale, universally applicable to architecture and mechanics*, volume 1. Birkhäuser.
- Drira, A.; Pierreval, H.; and Hajri-Gabouj, S. 2007. Facility layout problems: A survey. *Annual reviews in control*, 31(2): 255–267.
- Feng, W.; Zhu, W.; Fu, T.-j.; Jampani, V.; Akula, A.; He, X.; Basu, S.; Wang, X. E.; and Wang, W. Y. 2023. Lay-outgpt: Compositional visual planning and generation with large language models. *Advances in Neural Information Processing Systems*, 36: 18225–18250.
- Fu, H.; Jia, R.; Gao, L.; Gong, M.; Zhao, B.; Maybank, S.; and Tao, D. 2021. 3d-future: 3d furniture shape with texture. *International Journal of Computer Vision*, 129(12): 3313–3337.
- Fu, R.; Wen, Z.; Liu, Z.; and Sridhar, S. 2024. Anyhome: Open-vocabulary generation of structured and textured 3d homes. In *European Conference on Computer Vision*, 52–70. Springer.
- Gurobi Optimization, LLC. 2025. Gurobi Optimizer Reference Manual.
- Hong, S.; Zhang, X.; Du, T.; Cheng, S.; Wang, X.; and Yin, J. 2024. Cons2Plan: Vector Floorplan Generation from Various Conditions via a Learning Framework based on Conditional Diffusion Models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 3248–3256.
- Hu, R.; Huang, Z.; Tang, Y.; Van Kaick, O.; Zhang, H.; and Huang, H. 2020. Graph2plan: Learning floorplan generation from layout graphs. *ACM Transactions on Graphics (TOG)*, 39(4): 118–1.
- Hua, H.; Hovestadt, L.; Tang, P.; and Li, B. 2019. Integer programming for urban design. *European Journal of Operational Research*, 274(3): 1125–1137.
- Leimer, K.; Guerrero, P.; Weiss, T.; and Musialski, P. 2022. Layoutenhancer: Generating good indoor layouts from imperfect data. In *SIGGRAPH Asia 2022 Conference Papers*, 1–8.
- Leng, S.; Zhou, Y.; Dupty, M. H.; Lee, W. S.; Joyce, S. C.; and Lu, W. 2023. Tell2design: a dataset for language-guided floor plan generation. *arXiv preprint arXiv:2311.15941*.
- Littlefair, G.; Dutt, N. S.; and Mitra, N. J. 2025. FlairGPT: Repurposing LLMs for interior designs. In *Computer Graphics Forum*, e70036. Wiley Online Library.
- Liu, J.; Xue, Y.; Duarte, J.; Shekhawat, K.; Zhou, Z.; and Huang, X. 2022. End-to-end graph-constrained vectorized floorplan generation with panoptic refinement. In *European Conference on Computer Vision*, 547–562. Springer.
- Merrell, P.; Schkufza, E.; and Koltun, V. 2010. Computer-generated residential building layouts. In *ACM SIGGRAPH Asia 2010 Papers*, SIGGRAPH ASIA '10. New York, NY, USA: Association for Computing Machinery. ISBN 9781450304399.
- Merrell, P.; Schkufza, E.; Li, Z.; Agrawala, M.; and Koltun, V. 2011. Interactive furniture layout using interior design guidelines. *ACM transactions on graphics (TOG)*, 30(4): 1–10.
- Mostafavi, F.; van Engelenburg, C.; Khademi, S.; and Vrachliotis, G. 2024. Floor plan generation: The interplay among data, machine, and designer. *International Journal of Architectural Computing*, 14780771241290649.
- Nauata, N.; Chang, K.-H.; Cheng, C.-Y.; Mori, G.; and Furukawa, Y. 2020. House-gan: Relational generative adversarial networks for graph-constrained house layout generation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, 162–177. Springer.
- Nauata, N.; Hosseini, S.; Chang, K.-H.; Chu, H.; Cheng, C.-Y.; and Furukawa, Y. 2021. House-gan++: Generative adversarial layout refinement network towards intelligent computational agent for professional architects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13632–13641.
- Para, W.; Guerrero, P.; Kelly, T.; Guibas, L. J.; and Wonka, P. 2021. Generative layout modeling using constraint graphs. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6690–6700.
- Paschalidou, D.; Kar, A.; Shugrina, M.; Kreis, K.; Geiger, A.; and Fidler, S. 2021. Atiss: Autoregressive transformers for indoor scene synthesis. *Advances in Neural Information Processing Systems*, 34: 12013–12026.
- Peng, C.-H.; Yang, Y.-L.; and Wonka, P. 2014. Computing layouts with deformable templates. *ACM Transactions on Graphics (TOG)*, 33(4): 1–11.
- Qin, S.; He, C.; Chen, Q.; Yang, S.; Liao, W.; Gu, Y.; and Lu, X. 2024. ChatHouseDiffusion: Prompt-Guided Generation and Editing of Floor Plans. *arXiv preprint arXiv:2410.11908*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmlR.

- Ritchie, D.; Wang, K.; and Lin, Y.-a. 2019. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6182–6190.
- Shabani, M. A.; Hosseini, S.; and Furukawa, Y. 2023. Housediffusion: Vector floorplan generation via a diffusion model with discrete and continuous denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5466–5475.
- Sun, J.; Wu, W.; Liu, L.; Min, W.; Zhang, G.; and Zheng, L. 2022. Wallplan: synthesizing floorplans by learning to generate wall graphs. *ACM Transactions on Graphics (TOG)*, 41(4): 1–14.
- Tang, J.; Nie, Y.; Markhasin, L.; Dai, A.; Thies, J.; and Nießner, M. 2024. Diffuscene: Denoising diffusion models for generative indoor scene synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20507–20518.
- Voorhees, E. M.; et al. 1999. The trec-8 question answering track report. In *Trec*, volume 99, 77–82.
- Wang, K.; Savva, M.; Chang, A. X.; and Ritchie, D. 2018. Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)*, 37(4): 1–14.
- Wang, K.; Xu, X.; Lei, L.; Ling, S.; Lindsay, N.; Chang, A. X.; Savva, M.; and Ritchie, D. 2021. Roominoes: Generating novel 3d floor plans from existing 3d rooms. *Computer Graphics Forum*, 40(5): 57–69.
- Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Weiss, T.; Littenecker, A.; Duncan, N.; Nakada, M.; Jiang, C.; Yu, L.-F.; and Terzopoulos, D. 2018. Fast and scalable position-based layout synthesis. *IEEE Transactions on Visualization and Computer Graphics*, 25(12): 3231–3243.
- Wu, H.; Zhang, Z.; Zhang, W.; Chen, C.; Liao, L.; Li, C.; Gao, Y.; Wang, A.; Zhang, E.; Sun, W.; et al. 2023. Q-align: Teaching llms for visual scoring via discrete text-defined levels. *arXiv preprint arXiv:2312.17090*.
- Wu, W.; Fan, L.; Liu, L.; and Wonka, P. 2018. Miqp-based layout design for building interiors. *Computer Graphics Forum*, 37(2): 511–521.
- Wu, W.; Fu, X.-M.; Tang, R.; Wang, Y.; Qi, Y.-H.; and Liu, L. 2019. Data-driven interior plan generation for residential buildings. *ACM Transactions on Graphics (TOG)*, 38(6): 1–12.
- Wu, X.; Tojo, K.; and Umetani, N. 2024. Free-form floor plan design using differentiable Voronoi diagram. In *Pacific Graphics Conference Papers and Posters*, volume 43.
- Yang, Y.; Sun, F.-Y.; Weihs, L.; VanderBilt, E.; Herrasti, A.; Han, W.; Wu, J.; Haber, N.; Krishna, R.; Liu, L.; et al. 2024. Holodeck: Language guided generation of 3d embodied ai environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16227–16237.
- Yu, L. F.; Yeung, S. K.; Tang, C. K.; Terzopoulos, D.; Chan, T. F.; and Osher, S. J. 2011. Make it home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics (TOG)-Proceedings of ACM SIGGRAPH 2011*, v. 30,(4), July 2011, article no. 86, 30(4).
- Zong, Z.; Zhan, Z.; and Tan, G. 2024. HouseLLM: LLM-Assisted Two-Phase Text-to-Floorplan Generation. *arXiv preprint arXiv:2411.12279*.