

iPLAN: Interactive and Procedural Layout Planning

Feixiang He
 University of Leeds, UK
 scfh@leeds.ac.uk

Yanlong Huang
 University of Leeds, UK
 y.l.huang@leeds.ac.uk

He Wang *
 University of Leeds, UK
 h.e.wang@leeds.ac.uk

Abstract

*Layout design is ubiquitous in many applications, e.g. architecture/urban planning, etc, which involves a lengthy iterative design process. Recently, deep learning has been leveraged to automatically generate layouts via image generation, showing a huge potential to free designers from laborious routines. While automatic generation can greatly boost productivity, designer input is undoubtedly crucial. An ideal AI-aided design tool should automate repetitive routines, and meanwhile accept human guidance and provide smart/proactive suggestions. However, the capability of involving humans into the loop has been largely ignored in existing methods which are mostly end-to-end approaches. To this end, we propose a new human-in-the-loop generative model, iPLAN, which is capable of automatically generating layouts, but also interacting with designers throughout the whole procedure, enabling humans and AI to co-evolve a sketchy idea gradually into the final design. iPLAN is evaluated on diverse datasets and compared with existing methods. The results show that iPLAN has high **fidelity** in producing similar layouts to those from human designers, great **flexibility** in accepting designer inputs and providing design suggestions accordingly, and strong **generalizability** when facing unseen design tasks and limited training data.*

1. Introduction

Layout generation has recently spiked research interests in computer vision/graphics, aiming to automate the design process and boost the productivity. The traditional design process follows a diagram of iteratively adjusting/finalizing details from coarse to fine and global to local, which imposes repetitive and laborious routines on designers. Very recently, it has been shown that automatic image generation of such designs (with minimal human input) is possible through learning from data [12, 26, 27, 36]. This new line of research combines deep learning with design and has

demonstrated a new avenue for AI-aided design.

To achieve full automation, current research tends to learn from existing designs in an end-to-end fashion, and then to generate new ones with qualitative similarity and sufficient diversity. Taking floorplan as an example, automated generation can be based on simple human input, such as the boundary of the floor space [36], the relations among rooms [26, 27], or both [12]. While fully automated generation is important, design is in nature a procedural process, which involves alternations between repetitive routines and creative thinking at multiple intermediate stages [29]. Therefore, an ideal AI-aided system should automate the routine part while allowing the designer to impart creativity. This requires the system to be able to *interact* with the designer, in the sense that it should accept designer’s guidance, then actively suggest possible solutions accordingly, completing a feed-back loop. So far, the human-in-the-loop element is largely missing, which prevents a closer integration of AI and existing design practice.

Designing such an AI model faces several intrinsic challenges. In practice, learning how to interact with the designer requires a full observation of the decisions made at every intermediate stage. However, existing datasets, such as RPLAN [36] and LIFULL [22], usually only include the final designs, without the stage-to-stage design process. One potential solution to overcome this issue is to reverse-engineer intermediate stages from final designs, which however leads to another difficulty: the order of stages depends on the specific task/goal and could vary dramatically even for the same final design. Further, the order uncertainty is exacerbated by the strong personal styles and preferences of designers. Thus, how to design an AI system that can account for the above factors is a key research question, which is under-explored to date.

In this paper, we propose a new human-in-the-loop generative model for layout generation, which is referred to as *interactive planning* (iPLAN). Unlike previous work, iPLAN is equipped with a user-friendly interaction mechanism, which is achieved by letting the AI model learn the multi-stage design process, aiming to accommodate free-form *user interactions* and propose *design suggestions* at

*Corresponding author

every stage. This allows designer inputs at different stages across a wide range of levels of detail, while offering the capability of fully automated generation. To address the challenge of missing procedural design data, we reverse-engineer the final design to obtain the stage-to-stage process, based on principles that are widely adopted by professional designers [29]. This enables us to design a Markov chain model to capture the full design procedure. Since there is more than one way to reverse-engineer the final designs (*i.e.*, the stage order can vary), our model is designed with the capacity of accepting inputs with an arbitrary order, and consequently can learn the style variations implicitly from the data.

While iPLAN is general, we focus on floorplan design in this paper. iPLAN has been validated on two large-scale benchmark datasets, *i.e.*, RPLAN [36] and LIFULL [22], under diverse scenarios. The experiments show that our model is *versatile* in accepting designer inputs at various levels of detail, from minimal input and automatic generation, to stage-to-stage human guidance and interactive design. By learning from designs augmented by reverse-engineered processes, our model exhibits *high fidelity* in generating new designs with close style similarity and sufficient diversity. Finally, our model is highly *flexible* and *generalizable* when trained on varying amounts of data and facing unseen spaces and design requirements that are categorically different from the training data.

Contributions: (i) We propose a novel human-in-the-loop generative model iPLAN which respects design principles and mimics the design styles of professional designers implicitly. (ii) We demonstrate a successful fine-grained stage-to-stage generative model for floorplan, as opposed to existing end-to-end approaches. (iii) We show a variety of design scenarios, including fully automated generation, interactive planning with user instructions, and generalization for unseen tasks; (iv) We conduct extensive evaluations on diverse benchmark datasets and demonstrate that iPLAN outperforms the state of the art under multiple metrics.

2. Related Work

Layout generation has been an active research area in computer vision, *e.g.*, indoor scene synthesis [6, 24, 37, 39, 40] and floorplan generation [2, 9, 12, 25, 28, 36], image composition [1, 15], *etc.* Existing approaches can be generally grouped into two categories: handcrafted rule-based methods and data-driven methods. We mainly review the latter as they are closely related to our research.

Indoor scene synthesis. The synthesis of indoor scenes typically involves the placement of furniture models from an existing database into a given room. Convolutional neural networks can be trained to iteratively insert one object at a time into a room for indoor scene generation [31, 35]. High-level scene semantics can also be employed,

e.g., scene graphs, as a prior for a more controlled generation [34]. The biggest difference between indoor scene synthesis and floorplan generation is their requirements on the space partitioning. While indoor scene synthesis places objects in a room where the room itself does not need to be divided, floorplan normally requires explicit space division for different functionalities.

Image composition from scene graphs. Another related field is image composition from scene graphs, where the task is to derive the scene from a layout graph that describes the locations and features of the objects. Such generation can be achieved by Generative Adversarial Networks (GANs) based on graph convolution [15]. Further improvements can be obtained by separating the layout from the appearance of the objects [1]. For more controllability, Li *et al.* [21] synthesize images from a scene graph and the corresponding image crops. In contrast to the floorplan, the challenge in image composition is how to compose different objects into an image rather than partitioning the space.

Floorplan generation Floorplan generation can be formulated as an image synthesis problem which is one active research area in computer vision. Due to the surge of deep learning, the most promising approaches are GANs [3, 7, 16–19, 38]. Image-based GANs have been proven effective in floorplan generation [4, 13, 14, 30, 41–43]. Graph-based GANs can also produce floorplans by only taking spacial constraints, such as room connections, room types, in the form of graph [26, 27]. However, all these methods are end-to-end approaches, and therefore provide limited interactivity to the designer.

More recently, some human-in-the-loop approaches are proposed. Wu *et al.* [36] propose a two-phase approach to produce floorplans of residential buildings. The model successively predicts locations of rooms and walls given a building boundary, and converts the predicted layout into a vector graphics format. Graph2Plan [12] combines the topology information in the form of graphs with spatial constraints, to instantiate rooms accordingly. These methods enable human interactions at certain stages, *e.g.*, modifying room locations and retrieving the graph. Different from the existing methods, we propose a fine-grained generative model which enables interactions with the designer at different levels, from providing high-level design requirements to low-level instructions at a step.

3. Methodology

The overview of iPLAN is presented in Fig. 1. Without loss of generality, our model takes the space boundary as input, and decomposes the design procedure into: acquiring room types, locating rooms, and finalizing room partitions. Such workflow aims to mimic human designers [29] and accept designer inputs at any stage. The workflow is modeled as a joint probability distribution of all the afore-




Figure 1. Overview of our framework iPLAN. Room types are predicted at a time, while room locations and partitions are predicted iteratively.

mentioned factors, which is then factorized into stages and later formulated as a Markov chain (Sec. 3.1). Next, each factorized distribution provides a flexible entry point to incorporate user input or can be used for automatic generation (Sec. 3.2–Sec. 3.3).

3.1. Problem Formulation

A dataset of H layouts is denoted by $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^H$ with the i -th layout $\mathcal{D}_i = (\mathbf{B}_i, \mathbf{R}_i, \mathbf{T}_i, N_i, \mathbf{C}_i)$. $\mathbf{B}_i \in \mathbb{R}^{128 \times 128}$ is the boundary, N_i denotes the total number of rooms. We use j to index a specific room. $\mathbf{R}_i = \{\mathbf{r}_{i,j}\}_{j=1}^{N_i}$ denotes the room regions with $\mathbf{r}_{i,j} \in \mathbb{R}^4$ indicating the top-left and bottom-right corners of the bounding box of the room. $\mathbf{T}_i = \{t_{i,j}\}_{j=1}^{N_i}$ is a set of room types $t_{i,j} \in \mathbb{Z}^+$. $\mathbf{C}_i = \{\mathbf{c}_{i,j}\}_{j=1}^{N_i}$ is a set of room centers $\mathbf{c}_{i,j} \in \mathbb{R}^2$. Given \mathcal{D} , we aim to design a generative model for $\mathcal{P}(\mathcal{D}) = \prod_{i=1}^H \mathcal{P}(\mathcal{D}_i)$.

Similarly to existing methods, we also formulate floorplan design as an image generation problem, but with our focus on proposing a fine-grained generative model to enable human-in-the-loop interaction, through decomposing $\mathcal{P}(\mathcal{D})$ appropriately. From a mathematical perspective, there are many ways to decompose $\mathcal{P}(\mathcal{D})$. In order to allow human inputs across various levels of detail, we rely on design principles and widely adopted practice [29] to mimic the human design workflow, and naturally divide a floorplan design procedure into several stages. First, the desired number of rooms and their types are determined. Next, locations of areas with specific functionality (e.g., living rooms, bedrooms) are roughly estimated. Finally, room partitioning is conducted to finalize the design. This design diagram serves as a strong inductive bias in our model, following which $\mathcal{P}(\mathcal{D})$ is decomposed into:

$$\begin{aligned} \mathcal{P}(\mathcal{D}) &= \prod_{i=1}^H \mathcal{P}(\mathcal{D}_i) = \prod_{i=1}^H \mathcal{P}(\mathbf{R}_i, \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i) \\ &= \prod_{i=1}^H \mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i) \mathcal{P}(\mathbf{C}_i | \mathbf{T}_i, N_i, \mathbf{B}_i) \\ &\quad \mathcal{P}(\mathbf{T}_i, N_i | \mathbf{B}_i) \mathcal{P}(\mathbf{B}_i) \end{aligned} \quad (1)$$

where $\mathcal{P}(\mathbf{B}_i)$ accounts for the boundary known *a priori*; $\mathcal{P}(\mathbf{T}_i, N_i | \mathbf{B}_i)$ is to infer the desired number of rooms and

room types; $\mathcal{P}(\mathbf{C}_i | \mathbf{T}_i, N_i, \mathbf{B}_i)$ and $\mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i)$ respectively correspond to the coarse and fine designs of the layout, where the former estimates the room locations while the latter predicts the exact partitions. A visualization of Eq. (1) is provided in Figure 1, where the second, third and fourth blocks correspond to $\mathcal{P}(\mathbf{T}_i, N_i | \mathbf{B}_i)$, $\mathcal{P}(\mathbf{C}_i | \mathbf{T}_i, N_i, \mathbf{B}_i)$ and $\mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i)$, respectively.

Inspired by indoor scene synthesis where objects are placed iteratively [31, 35], we assume rooms are designed one by one. Formally, we model $\mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i)$ and $\mathcal{P}(\mathbf{C}_i | \mathbf{T}_i, N_i, \mathbf{B}_i)$ as Markov chains, *i.e.*, designs are conducted in a step-wise manner and early decisions will affect later ones, which allows a designer to focus on one room at a time and give guidance at any step:

$$\mathcal{P}(\mathbf{C}_i | \mathbf{T}_i, N_i, \mathbf{B}_i) = \prod_{j=1}^{N_i} \mathcal{P}(c_{i,j} | c_{i,<j}, t_{i,j}, N_i, \mathbf{B}_i) \quad (2)$$

$$\mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i) = \prod_{j=1}^{N_i} \mathcal{P}(r_{i,j} | r_{i,<j}, c_{i,j}, t_{i,j}, N_i, \mathbf{B}_i) \quad (3)$$

where $r_{i,<j} = \{r_{i,1}, \dots, r_{i,j-1}\}$ and $c_{i,<j} = \{c_{i,1}, \dots, c_{i,j-1}\}$ denote the set of allocated room partitioning and centers before the j -th room, respectively.

Connections to existing research. Eq. (1)–(3) is a generalization of existing methods and is more fine-grained. Graph2Plan [12] simultaneously determines \mathbf{C}_i , \mathbf{T}_i and N_i given \mathbf{B}_i , and then $\mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i)$ is predicted. RPLAN [36] estimates $\mathcal{P}(\mathbf{C}_i, \mathbf{T}_i, N_i | \mathbf{B}_i)$ via predicting $\mathcal{P}(c_{i,j}, t_{i,j} | c_{i,<j}, t_{i,<j}, \mathbf{B}_i)$ consecutively, and then indirectly estimates room areas \mathbf{R}_i by locating walls. In contrast, we further decompose $\mathcal{P}(\mathbf{C}_i, \mathbf{T}_i, N_i | \mathbf{B}_i)$ into $\mathcal{P}(\mathbf{T}_i, N_i | \mathbf{B}_i)$ and $\mathcal{P}(\mathbf{C}_i | \mathbf{T}_i, N_i, \mathbf{B}_i)$, and further decompose the latter using Eq. (2). Moreover, $\mathcal{P}(\mathbf{R}_i | \mathbf{C}_i, \mathbf{T}_i, N_i, \mathbf{B}_i)$ is also decomposed into multiple steps by Eq. (3). Our decompositions bring a more fine-grained procedural generative model that allows for user interactions at arbitrary steps. This enables more flexible and closer human-AI interactions. For the sake of brevity, we omit the subscripts i of $\mathbf{T}_i, N_i, \mathbf{B}_i$ and \mathbf{C}_i in the following sections.

3.2. The Number and Types of Rooms

T and N are normally given beforehand. However, given a specific \mathbf{B} , the design might not be unique, *e.g.*, the same space can be designed as a 2-bed or 3-bed flat. In other words, there exists a distribution of possible designs. Thus, we propose to learn their distributions $\mathcal{P}(T, N | \mathbf{B})$ from real designs by professional designers to enable automatic exploration. $\{T, N\}$ can be replaced by a random variable $\mathbf{Q} = \{q_k\}_{k=1}^K$, where K denotes the number of room types in \mathcal{D} and q_k corresponds to the number of rooms




Figure 2. Prediction of a room center at one step

under the k -th type. So, we model $P(\mathbf{Q}|\mathbf{B})$ instead.

We propose a boundary-conditioned Variational Autoencoder (BCVAE) based on VAE [32] where \mathbf{B} serves as a condition. The model consists of an embedding module \mathcal{F}_{ed} , an encoder \mathcal{F}_{en} and a decoder \mathcal{F}_{de} . By feeding \mathcal{F}_{ed} with \mathbf{B} , it outputs an embedded vector $\gamma \in \mathbb{R}^{128}$. Also, $\{\mu, \Sigma\} = \mathcal{F}_{en}(\mathbf{Q}, \gamma)$ where $\mu \in \mathbb{R}^{32}$ and $\Sigma \in \mathbb{R}^{32 \times 32}$ are the mean and covariance of a Gaussian distribution. Furthermore, a latent variable z is sampled from $\mathcal{N}(\mu, \Sigma)$ using the reparameterization trick [20]. Given z and γ , we reconstruct \mathbf{Q} by $\hat{\mathbf{Q}} = \mathcal{F}_{de}(z, \gamma)$. We employ a standard VAE loss for training. The detailed network architecture and the training loss are discussed in the supplementary material.

During inference, for a new boundary \mathbf{B} , we sample z from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and predict $\hat{\mathbf{Q}}$ by $\hat{\mathbf{Q}} = \mathcal{F}_{de}(z, \mathcal{F}_{ed}(\mathbf{B}))$, which is then used to recover $\hat{\mathbf{T}}$ and $\hat{\mathbf{N}}$.

3.3. Locating Rooms

Locating room regions $\mathcal{P}(\mathbf{C}|\mathbf{T}, N, \mathbf{B})$ plays an essential role in the design process. Similarly to [36], we model room region prediction as a step-wise classification task. At each step, given a multi-channel image representation of the current design state and the next desired room type t_j , we predict the center of the next room. The multi-channel image encodes the boundary \mathbf{B} and all previously predicted room centers $c_{<j}$. The image consists of $K + 4$ binary channels, three of which label \mathbf{B} , i.e., the boundary, the front door and the interior area pixels. K channels represent the predicted room centers, with each channel corresponding to a room type. For each room, its center is represented by a 9×9 square of pixels with value 1. We also use a channel to summarize the centers of all predicted rooms. Regarding the desired room type t_j , we convert it into a one-hot vector.

We feed the multi-channel representation to a Resnet-18 [10] network and the one-hot vector of the desired room type to an embedding network to extract features. The embedding network has three fully-connected layers, followed by four convolution blocks, a Batch Normalization Layer and a LeakyReLU Layer. The extracted features are concatenated and subsequently fed to a decoder module that contains 4 atrous spatial pyramid pooling (ASPP) [5], a Convolutional block and a Deconvolutional block. The output is $\mathbf{O} \in \mathbb{R}^{(K+3) \times 128 \times 128}$, giving a probability vector of

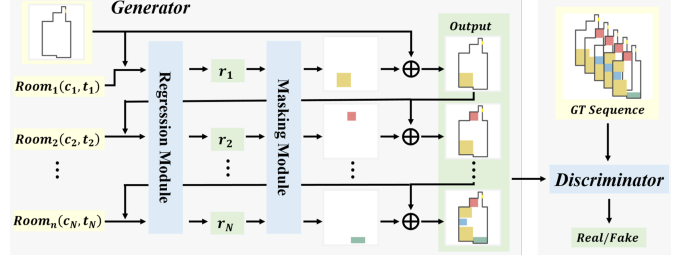


Figure 3. Procedural room shape generation

all pixels. For each pixel, we predict $K + 3$ labels, i.e., K room types and EXISTING, FREE, and OUTSIDE, where EXISTING, FREE and OUTSIDE indicate whether a pixel belongs to an existing room, a free space (in the interior of the boundary) and the exterior, respectively. The whole model is shown in Figure 2.

For training, we predict one room at a time, by decomposing the final design into a series of design states with one room added at each stage. However, there are multiple possible sequences for a given final design and the ground-truth step-to-step decisions are unavailable. Therefore, we employ a stochastic training process to learn all possible sequences by randomly removing rooms from a design. We propose a pixel-wise cross-entropy loss:

$$\mathcal{L} = \sum_{h=1}^{128} \sum_{w=1}^{128} -\omega_y \log \frac{\exp(O_{y,h,w})}{\sum_{k=1}^{K+3} \exp(O_{k,h,w})}, \quad (4)$$

where y is the ground-truth class index for the pixel located at (h, w) . ω_y represents the weight of y -th label, we set it to be 2 for K room types and 1.25 for the other three labels.

During the inference phase, for a sequence of room types $\mathbf{T} = \{t_1, t_2, \dots, t_N\}$ we first predict the room center for t_1 , i.e., $\mathcal{P}(c_1|t_1, N, \mathbf{B})$, then we update the current design state using this predicted center \hat{c}_1 and continue to predict the next room center for t_2 , i.e., $\mathcal{P}(c_2|\hat{c}_1, t_2, N, \mathbf{B})$. This procedure is repeated until all centers $\hat{\mathbf{C}}$ are determined. Note that when the order of elements in \mathbf{T} varies, its corresponding layout can also change, thus providing design diversity. Also, at any intermediate step, user input (e.g., adjust a room center) can be incorporated into our model before it predicts the next room center.

3.4. Predicting Room Partitioning

After predicting the room centers $\hat{\mathbf{C}}$, we are ready for detailed room partitioning $\mathcal{P}(\mathbf{R}|\hat{\mathbf{C}}, \hat{\mathbf{T}}, \hat{\mathbf{N}}, \mathbf{B})$, where we need to consider the room size and the room shape. The room size is directly related to the specific functionality of the room, e.g., living rooms are normally larger than bathrooms for social interactions. The shape is affected by the functionality too but also strongly affected by the boundary geometry and the inner walls. Further, there is usually no ‘optimal’ solution but a distribution of near-optimal solutions

when looking at the final designs. Since the distribution can be arbitrary, we propose a new Generative Adversarial Network (GAN) to model the step-wise prediction of room partitioning. The general model is shown in Figure 3. Instead of directly predicting the shape of each room, we first predict its bounding box. This allows the bounding box predicted at a later stage to grab the space that is already allocated. We find this is a straightforward yet effective way to generate non-box rooms.

Generator: The generator begins with a bounding-box regressor \mathcal{F}_b that outputs the top-left and bottom-right corners of the box. \mathcal{F}_b consists of six convolutional units and two fully connected layers, each unit with a Convolution layer, a Layer-Normalization layer and a ReLU layer. \mathcal{F}_b enables the designer to interact with the system easily, *e.g.*, relocating a room, modifying a room bounding box. However, only predicting the corners of bounding boxes is not convenient. We want to label all the pixels within the box so that we can work in the image space consistently. So we design a tailor-made masking module \mathcal{F}_m to map each room bounding box into a room mask image. \mathcal{F}_m contains six convolution layers with 3×3 kernel, followed by a Sigmoid activation function.

For a given sequence of room centers and types $\{(c_1, t_1), (c_2, t_2), \dots, (c_N, t_N)\}$, we use \mathcal{F}_b to predict the coordinates \hat{r}_1 of the bounding box of the first room specified by c_1 and t_1 , where $\hat{r}_1 = \mathcal{F}_b(\mathbf{S}_0, c_1, t_1)$ with $\mathbf{S}_0 = \mathbf{B}$. Furthermore, we obtain the room mask using $\hat{M}_1 = \mathcal{F}_m(\hat{r}_1) \in \mathbb{R}^{128 \times 128}$. The output of the generator at the first step is computed by

$$\hat{\mathbf{S}}_1 = \mathbf{S}_0 \times (\mathbf{1} - \hat{M}_1) + \hat{M}_1 \times t_1, \quad (5)$$

where $\mathbf{1} \in \mathbb{R}^{128 \times 128}$ represents a matrix with all elements being 1. During the prediction in the second step, the generator takes as input $\hat{\mathbf{S}}_1, c_2, t_2$ and outputs $\hat{\mathbf{S}}_2 = \hat{\mathbf{S}}_1 \times (\mathbf{1} - \hat{M}_2) + \hat{M}_2 \times t_2$ with $\hat{M}_2 = \mathcal{F}_m(\hat{r}_2) = \mathcal{F}_m(\mathcal{F}_b(\hat{\mathbf{S}}_1, c_2, t_2))$. The same procedure is iterated (for totally \hat{N} steps) until the entire design is accomplished.

Discriminator: The discriminator retains the similar backbone as \mathcal{F}_b in the generator, except that the last two convolution layers are dropped. The discriminator is employed to distinguish whether a sequence of design states comes from the generator or the ground truth. Specifically, a sequence of predicted design states $\hat{\mathbf{S}} = \{\hat{\mathbf{S}}_1, \hat{\mathbf{S}}_2, \dots, \hat{\mathbf{S}}_{\hat{N}}\}$ should be recognized as ‘FALSE’, while a sequence from the ground truth with the same order of room types (*i.e.*, $\{t_1, t_2, \dots, t_{\hat{N}}\}$) should be predicted as ‘TRUE’.

For training, we define a loss on top of WGAN-GP [8] loss, *i.e.*,

$$\begin{aligned} \mathcal{L} = & \mathbb{E}_{\hat{\mathbf{S}} \sim \mathbb{P}_g} [D(\hat{\mathbf{S}})] - \mathbb{E}_{\mathbf{S} \sim \mathbb{P}_r} [D(\mathbf{S})] \\ & + \lambda_1 \mathbb{E}_{\hat{\mathbf{S}} \sim \mathbb{P}_g} (\|\nabla_{\hat{\mathbf{S}}} D(\hat{\mathbf{S}})\|_2 - 1)^2 + \lambda_2 \mathcal{L}_s, \end{aligned} \quad (6)$$

where \mathbb{P}_g and \mathbb{P}_r represent the generated and real data distribution, respectively. $\hat{\mathbf{S}} \sim \mathbb{P}_{\hat{\mathbf{S}}}$ denotes a random interpolation between $\hat{\mathbf{S}}$ and \mathbf{S} , and is used for a gradient penalty with $\lambda_1 = 10$. $\|\cdot\|_2$ is the ℓ_2 -norm. The last term is introduced to explicitly regularize the bounding boxes with $\lambda_2 = 100$:

$$\mathcal{L}_s = \sum_{j=1}^N l_j, \quad l_j = \begin{cases} 0.5(\mathbf{r}_j - \hat{\mathbf{r}}_j), & \text{if } \|\mathbf{r}_j - \hat{\mathbf{r}}_j\|_1 < 1 \\ \|\mathbf{r}_j - \hat{\mathbf{r}}_j\|_1 - 0.5, & \text{otherwise} \end{cases} \quad (7)$$

where \mathbf{r}_j is the ground truth and $\|\cdot\|_1$ is the ℓ_1 -norm.

At the inference stage, our procedural model predicts room areas in a step-wise manner, which hence allows user input at an arbitrary intermediate step, such as modifying room types, changing room centers or even the already predicted room areas. Gaps may exist between the predicted rooms. Thus, we employ a simple post-processing step as in [33] and detail it in the supplementary material.

4. Experiments

4.1. Datasets

We conduct experiments on two commonly used datasets, RPLAN [36] and LIFULL [22]. **RPLAN** is collected from the real-world residential buildings in the Asian real estate market, which contains over 80k floorplans and 13 types of rooms¹. All floorplans in RPLAN are axis-aligned and pre-processed to the same scale. The training-validation-test split of the dataset is 70%–15%–15% [12]. **LIFULL HOME**’s dataset offers approximately five million apartment floorplans from the Japanese house market. The original dataset is given in the form of images, but a subset has been parsed into vector format by [23]. We select a subset with 4–10 rooms. This specific dataset consists of approximately 54k floorplans and 9 room types², among which 85% (randomly sampled) of the data serves as the training set while the remaining for testing.

4.2. Metrics

Although each floorplan only contains one design (*i.e.*, one boundary corresponds to one design), our model can predict the distribution of plausible designs, offering design diversity and alternative choices. To evaluate the predicted distribution, we employ the *Fréchet Inception Distance* (FID) [11] to calculate the distance between two distributions, which was also used in floorplan generation [26,27]. Built on FID, we introduce three metrics FID_{img} , FID_{area} and FID_{type} : (i) FID_{img} : computed on rendered images to evaluate the distributional differences of

¹LivingRoom, MasterRoom, Kitchen, Bathroom, DiningRoom, ChildRoom, StudyRoom, SecondRoom, GuestRoom, Balcony, Entrance, Storage, Wall-in.

²LivingRoom, Kitchen, Bedroom, Bathroom, Office, Balcony, Hallway, OtherRoom.

Dataset	Method	FID_{img}	FID_{area}	FID_{type}
RPLAN	HouseGAN++	51.33	1.36×10^8	0.038
	Rplan	4.1	2.29×10^5	0.58
	Our _{III}	1.22	3.13×10^4	0.05
	Our _{II}	0.72	1.09×10^4	0.03
	Rplan*	0.11	8.62×10^3	6.4×10^{-3}
	Graph2Plan	0.62	8.82×10^3	2.70×10^{-4}
	Our _I	0.16	4.89×10^2	4.44×10^{-6}
LIFULL	Rplan	50.19	4.29×10^6	5.15
	Our _{III}	37.35	9.81×10^5	2.52
	Our _{II}	32.65	7.75×10^5	2.14
	Rplan*	1.43	5.62×10^5	0.064
	Graph2Plan	0.64	2.87×10^3	2.63×10^{-5}
	Our _I	0.38	2.07×10^3	3.59×10^{-6}

Table 1. FID-based metrics on RPLAN and LIFULL.

the generated and true images. (ii) FID_{area} : to evaluate the distributional differences of room areas. Each layout is represented by a $1 \times K$ vector $area_i$, with its k -th element $area_{i,k}$ representing the average area of the k -th type of rooms in the i -th floorplan. K denotes the number of room types. (iii) FID_{type} : to calculate the distributional differences of room numbers against room types. Each layout is represented by a $1 \times K$ vector $type_i$, whose element $type_{i,k}$ represents the number of rooms under the k -th type in the i -th floorplan. It is worthwhile to mention that the above three metrics are not biased towards our method as they are not involved in the training process at all.

4.3. Baselines

We choose three methods as baselines. Rplan³ [36] is an image-based method, which outputs a floorplan given a boundary. We further use a variant of Rplan, named Rplan*, which also takes as input the room types and centers. HouseGAN++ [27] is a graph-constrained approach, which treats rooms as nodes and requires node types and connectivity as input. It generates mask images for all nodes according to the graph, then blends them to form the layout. Graph2plan [12] requires a boundary and a room relation graph as input, where the graph includes information on room size, room center, room type and their connections. All floorplans are rendered in the same way as [12], including door and window placements.

4.4. Quantitative Evaluations

Ablation Study. iPLAN is capable of fully automated generation while allowing for user interactions at different stages. Based on the amount of information received from the designer, three variants are evaluated: (v1) Our_I takes as input the boundary, room types, and room centers; (v2) Our_{II} takes as input the boundary and room

³We rename the approach RPLAN as Rplan in order to distinguish it from the dataset RPLAN.

types; (v3) Our_{III} takes as input only the boundary. The variants show different levels of automation/user interactivity, demonstrating iPLAN’s flexibility when the depth of human involvement varies, from little human input (e.g., iPLAN freely explores designs), to step-to-step guidance (e.g., where to put each room). The implementation details are in the supplementary document. Table 1 shows the results. iPLAN achieves good results on all three settings, among which Our_I achieves the best results and Our_{II} is better than Our_{III} . This is natural as more prior information helps for better prediction. Considering RPLAN and LIFULL are significantly different in terms of the overall shapes of boundaries and rooms, iPLAN can indeed handle different data distributions well.

Comparisons on RPLAN iPLAN in general outperforms all baseline methods by large margins shown in Table 1. Looking closely, HouseGAN++ achieves worse results, except FID_{type} with a small margin, than other methods. The reason is that its layouts are not bounded by boundaries. By taking the boundary as input, both Rplan and Our_{III} predict room centers and regions successively, but Our_{III} outperforms Rplan on all three metrics. Furthermore, when both the boundary and room types are given, Our_{II} outperforms the first three methods. Finally, Rplan*, Graph2Plan and Our_I are given the full information (boundary, room types and centers). Our_I further improves on all metrics, where Rplan* achieves a slightly better score on FID_{img} . We observe that the step-wise prediction in iPLAN is less likely to generate partitioning ambiguity (allocating a space to multiple rooms), while is observed in Rplan*. Also, learning conditional probabilities (as iPLAN does) as oppose to a joint probability of multiple factors from final designs (as Rplan* and Graph2Plan do) is easier, where the decomposition of final designs can be seen as a data augmentation strategy.

Comparisons on LIFULL Similar comparisons have been done on LIFULL (Table 1). We skip HouseGAN++ because it needs door locations in the input but there is no such data in LIFULL. Compared with Rplan*, Our_I and Graph2Plan, room centers are not provided as priors for Rplan, Our_{III} and Our_{II} , leading to a large performance gap between these two groups. It is understandable because LIFULL is more challenging/heterogeneous with multi-scale samples and nested rooms. Next, Our_{II} outperforms Our_{III} on all three metrics because more prior knowledge (i.e., room types) is given, and both are better than Rplan. Furthermore, when the room centers are provided, inaccurate wall prediction affects Rplan*. Different from the results on RPLAN, Graph2Plan outperforms Rplan* on LIFULL. This is because Graph2Plan takes spacial room relations as input which is more robust for such a multi-scale dataset. Our_I is still better than Rplan* and Graph2Plan. Therefore, we can conclude that, when the

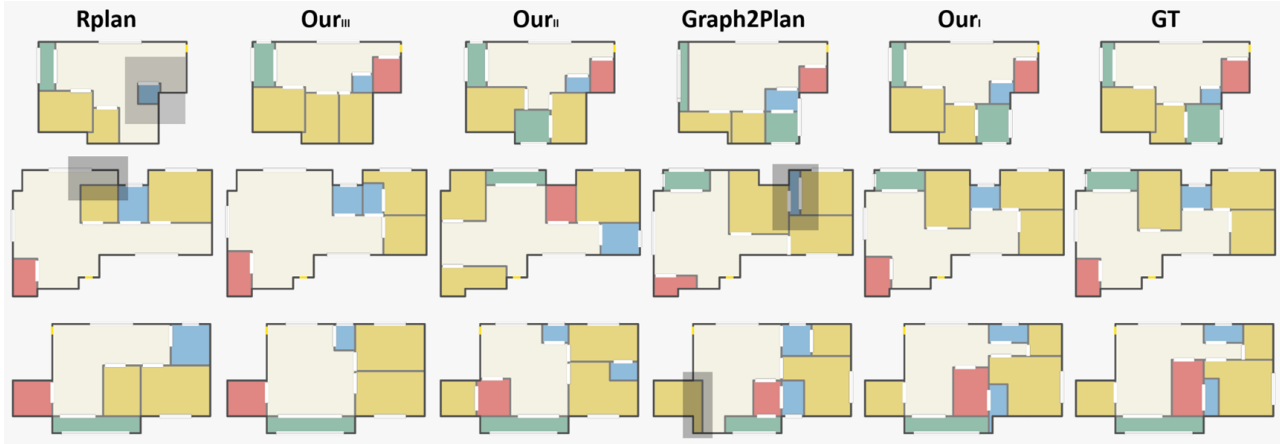


Figure 4. Qualitative comparisons. Shaded areas indicate design choices that are questionable and unobserved in data.

same amount of prior knowledge is given, iPLAN achieves the best performance.

4.5. Qualitative Evaluation

To show intuitive results, qualitative comparisons are given in Figure 4. Each color represents one room type (see color indication in Fig. 5). The shaded boxes indicate questionable design areas. Specifically, in the first row, Rplan designs an isolated bathroom at a corner in the first image, which is a strange design and not observed in the data. In the second row, the first image by Rplan allocates an unreasonable space (shaded) to the public area, which however should be assigned to the neighboring bedroom. Similarly, the layout produced by Graph2Plan consists of 3 bedrooms (second row), but the bathroom is embedded among them (shaded), which means people need to pass through one bedroom to go to the toilet. Besides, the bathroom is too small. As for the last row, the room area is not properly predicted by Graph2plan because the region (shaded) at the bottom-left corner is too narrow.

By comparing Our_I , Our_{II} and Our_{III} , it is apparent that Our_I is closest to GT as a result of the strong input constraints. In contrast, when less constraints are imposed, the diversity is significant, as depicted in Our_{II} and Our_{III} . More results and analysis are shown in the supplementary material.

4.6. User Interactions

The performance of iPLAN under different levels of human interactions is showcased in Figure 5. Figure 5(a) is fully automated generation without human input other than a boundary. iPLAN can generate diverse layouts, with a varying number of rooms and different room types/areas. Figure 5(b) shows the user can decide the room types, where the final layouts are different when the order of room types varies, showing the flexibility of iPLAN. We further provide iPLAN with more priors in Fig. 5(c), *i.e.*, the bound-

Dataset	Method	FID_{img}	FID_{area}	FID_{type}
RPLAN	Our_{II}^8	1.2	1.36×10^4	0.06
	Our_I^8	0.2	2.69×10^2	1.25×10^{-5}
	$Our_{II}^{7 \sim 8}$	2.28	4.87×10^4	0.18
	$Our_I^{7 \sim 8}$	0.24	6.89×10^2	1.54×10^{-5}

Table 2. Generalization results on RPLAN.

ary, room types and centers, leading to a planned layout that is nearly the same as the ground truth. We also evaluate iPLAN by introducing a user input at an intermediate step (Fig. 5(d)), where the balcony is moved to the left and consequently a different layout is obtained. Figure 5 demonstrates the interactive ability of iPLAN, which is crucial in settings where the designer leads the design.

4.7. Generalizability

Quantitative results To further push iPLAN, we quantitatively evaluate the generalizability of iPLAN by setting up two groups of experiments on RPLAN. In the first group, we select 51322 layouts consisting of 4 ~ 7 rooms for training and randomly choose 12,000 8-room layouts for testing. In the second group, we consider a more challenging task by randomly selecting 26574 layouts containing 4 ~ 6 rooms for training and 12000 layouts containing 7 ~ 8 rooms for testing. Note that in both experiments, the layouts for test include more rooms than the ones used for training.

Quantitative results are reported in Table 2. Again, we can observe that Our_I outperforms Our_{II} , which is consistent with our intuition since more prior knowledge is provided to Our_I . Further, the results in first group are better than those in the second group, which is reasonable because the generalization setting in the second group is more challenging. However, both groups achieve comparable performance to Table 1 and still outperform the baselines.

Qualitative Results An interesting test is to see if iPLAN trained on RPLAN can be generalized to unseen types of boundaries. The qualitative results are also pro-

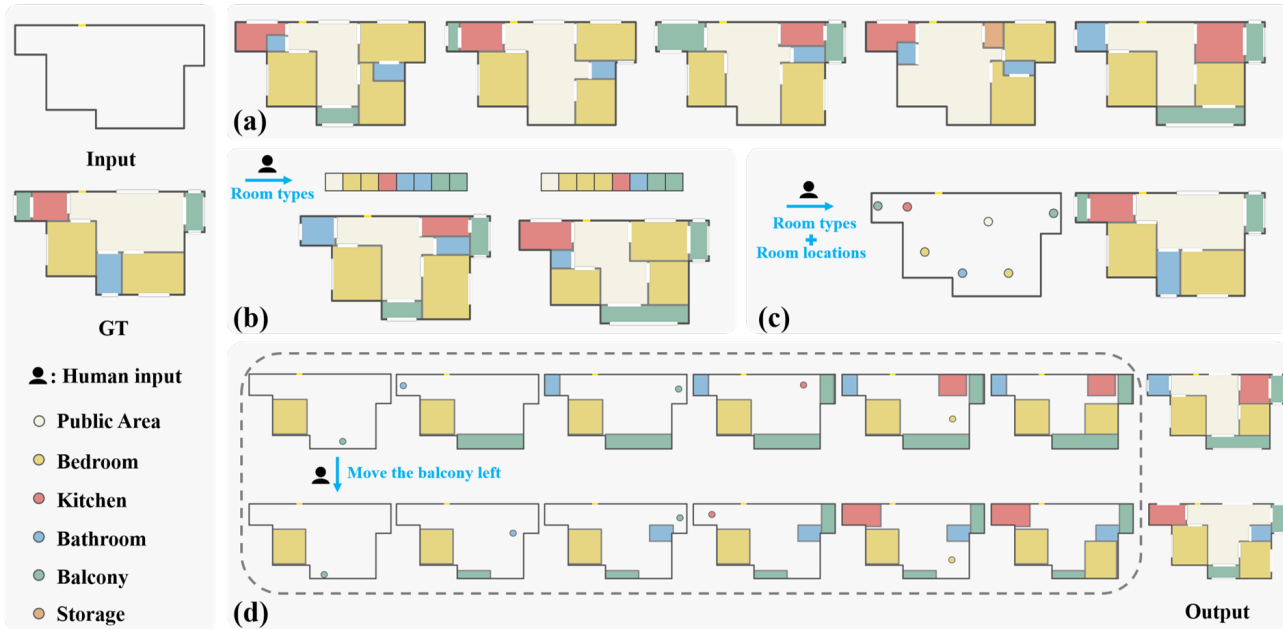


Figure 5. User interactions. (a)–(c) indicate interactions via providing different levels of human input; (d) shows fine-grained interaction in a step-wise generation.




Figure 6. Layouts with non-axis aligned edges. Left: successful examples. Right: Rplan fails but our model succeeds.

vided to verify the generalizability of iPLAN. Note that RPLAN only contains layouts with straight axis-aligned edges in the boundaries, so we consider the boundaries with non-axis aligned edges and curves. Since HouseGAN++ cannot design a layout for a specific building boundary and Graph2Plan is restricted to boundaries with axis-aligned edges, we only compare iPLAN with Rplan.

Figure 6 shows several examples of non-axis aligned design produced by Rplan and iPLAN. The left lists four successful cases on both approaches and the right shows two examples where Rplan fails but iPLAN succeeds. Even when successful, Rplan often predicts isolated rooms which divide the public areas into strange shapes and reduce their usability (the first two examples), while iPLAN utilizes the space better. As for the failures of Rplan, we give two examples of room center estimation and wall prediction, where it fails to reasonably fill and partition the space, a problem from which iPLAN does not suffer. Please refer to the supplementary document for more examples.

5. Limitation & Conclusion

While iPLAN is able to handle some irregular boundaries, it cannot cope with extremely irregular ones. One important future direction would be to handle complex environments, e.g., non-axis aligned boundaries and compound and nested rooms, which is crucial to generalize iPLAN to public spaces such as train stations and shopping malls.

In this paper, we proposed a novel human-in-the-loop generative model iPLAN to learn professional designs in a stage-to-stage fashion while respecting design principles. While being capable of fully automated generation, iPLAN allows close interactions with humans by accepting user guidance at every stage and automatically suggesting possible designs accordingly. Comprehensive evaluations on two benchmark datasets RPLAN and LIFULL show that iPLAN outperforms the state-of-the-art methods, both quantitatively and qualitatively. Importantly, iPLAN has exhibited strong generalization capability to unseen design tasks and boundary inputs.

Acknowledgements: In this paper, we used "LIFULL HOME'S Snapshot Data of Rentals" provided by LIFULL Co., Ltd. via IDR Dataset Service of National Institute of Informatics. We thank Jing Li for her input on the design practice. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 899739 CrowdDNA and the Marie Skłodowska-Curie grant agreement No 101018395. Feixiang He has been supported by UKRI PhD studentship [EP/R513258/1, 2218576].

References

- [1] Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4561–4569, 2019. 2
- [2] Fan Bao, Dong-Ming Yan, Niloy J Mitra, and Peter Wonka. Generating and exploring good building layouts. *ACM Transactions on Graphics (TOG)*, 32(4):1–10, 2013. 2
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 2
- [4] Stanislas Chaillou. Archigan: Artificial intelligence x architecture. In *Architectural intelligence*, pages 117–127. Springer, 2020. 2
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 4
- [6] Matthew Fisher, Daniel Ritchie, Manolis Savva, Thomas Funkhouser, and Pat Hanrahan. Example-based synthesis of 3d object arrangements. *ACM Transactions on Graphics (TOG)*, 31(6):1–11, 2012. 2
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [8] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017. 5
- [9] Mikako Harada, Andrew Witkin, and David Baraff. Interactive physically-based manipulation of discrete/continuous models. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 199–208, 1995. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 5
- [12] Ruizhen Hu, Zeyu Huang, Yuhan Tang, Oliver Van Kaick, Hao Zhang, and Hui Huang. Graph2plan: Learning floor-plan generation from layout graphs. *ACM Transactions on Graphics (TOG)*, 39(4):118–1, 2020. 1, 2, 3, 5, 6
- [13] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 172–189, 2018. 2
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2
- [15] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018. 2
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 2
- [17] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *arXiv preprint arXiv:2006.06676*, 2020. 2
- [18] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 2
- [19] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 2
- [20] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 4
- [21] Yikang Li, Tao Ma, Yeqi Bai, Nan Duan, Sining Wei, and Xiaogang Wang. Pastegan: A semi-parametric method to generate image from scene graph. *Advances in Neural Information Processing Systems*, 32:3948–3958, 2019. 2
- [22] Ltd. LIFULL Co. Lifull home’s snapshot data of rentals, nov 2015. 1, 2, 5
- [23] Chen Liu, Jiajun Wu, Pushmeet Kohli, and Yasutaka Furukawa. Raster-to-vector: Revisiting floorplan transformation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2195–2203, 2017. 5
- [24] Paul Merrell, Eric Schkufza, Zeyang Li, Maneesh Agrawala, and Vladlen Koltun. Interactive furniture layout using interior design guidelines. *ACM transactions on graphics (TOG)*, 30(4):1–10, 2011. 2
- [25] Pascal Müller, Peter Wonka, Simon Haegler, Andreas Ulmer, and Luc Van Gool. Procedural modeling of buildings. In *ACM SIGGRAPH 2006 Papers*, pages 614–623. 2006. 2
- [26] Nelson Nauata, Kai-Hung Chang, Chin-Yi Cheng, Greg Mori, and Yasutaka Furukawa. House-gan: Relational generative adversarial networks for graph-constrained house layout generation. In *European Conference on Computer Vision*, pages 162–177. Springer, 2020. 1, 2, 5
- [27] Nelson Nauata, Sepidehsadat Hosseini, Kai-Hung Chang, Hang Chu, Chin-Yi Cheng, and Yasutaka Furukawa. House-gan++: Generative adversarial layout refinement networks. *arXiv preprint arXiv:2103.02574*, 2021. 1, 2, 5, 6
- [28] Chi-Han Peng, Yong-Liang Yang, and Peter Wonka. Computing layouts with deformable templates. *ACM Transactions on Graphics (TOG)*, 33(4):1–11, 2014. 2
- [29] Roberto J Rengel. *The interior plan: Concepts and exercises*. A&C Black, 2011. 1, 2, 3
- [30] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding

- in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021. 2
- [31] Daniel Ritchie, Kai Wang, and Yu-an Lin. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6182–6190, 2019. 2, 3
- [32] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28:3483–3491, 2015. 4
- [33] Chun-Yu Sun, Qian-Fang Zou, Xin Tong, and Yang Liu. Learning adaptive hierarchical cuboid abstractions of 3d shape collections. *ACM Transactions on Graphics (TOG)*, 38(6):1–13, 2019. 5
- [34] Kai Wang, Yu-An Lin, Ben Weissmann, Manolis Savva, Angel X Chang, and Daniel Ritchie. Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. 2
- [35] Kai Wang, Manolis Savva, Angel X Chang, and Daniel Ritchie. Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 2, 3
- [36] Wenming Wu, Xiao-Ming Fu, Rui Tang, Yuhan Wang, Yu-Hao Qi, and Ligang Liu. Data-driven interior plan generation for residential buildings. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 38(6), 2019. 1, 2, 3, 4, 5, 6
- [37] Lap Fai Yu, Sai Kit Yeung, Chi Keung Tang, Demetri Terzopoulos, Tony F Chan, and Stanley J Osher. 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), 2011. 2
- [38] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International conference on machine learning*, pages 7354–7363. PMLR, 2019. 2
- [39] Xi Zhao, Ruizhen Hu, Paul Guerrero, Niloy Mitra, and Taku Komura. Relationship templates for creating scene variations. *ACM Transactions on Graphics (TOG)*, 35(6):1–13, 2016. 2
- [40] Xi Zhao, He Wang, and Taku Komura. Indexing 3d scenes using the interaction bisector surface. *ACM Trans. Graph.*, 33(3):22:1–22:14, June 2014. 2
- [41] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2
- [42] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Multimodal image-to-image translation by enforcing bi-cycle consistency. In *Advances in neural information processing systems*, pages 465–476, 2017. 2
- [43] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5104–5113, 2020. 2