

Geometry-Based Layout Generation with Hyper-Relations AMONG Objects

Shao-Kui Zhang
Tsinghua University

zhangsk18@mails.tsinghua.edu.cn

Wei-Yu Xie
Beijing Institute of Technology

ervinxie@qq.com

Song-Hai Zhang*
Tsinghua University

shz@tsinghua.edu.cn

Abstract

Recent studies show increasing demands and interests in automatic layout generation, while there is still much room for improving the plausibility and robustness. In this paper, we present a data-driven layout generation framework without model formulation and loss term optimization. We achieve and organize priors directly based on samples from datasets instead of sampling probabilistic distributions. Therefore, our method enables expressing relations among three or more objects that are hard to be mathematically modeled. Subsequently, a non-learning geometric algorithm is proposed to arrange objects considering constraints such as positions of walls and windows. Experiments show that the proposed method outperforms the state-of-the-art and our generated layouts are competitive to those designed by professionals.¹

1. Introduction

3D scenes are becoming fundamental in many areas of computer graphics, e.g., photo-realistic rendering, virtual reality (VR), providing datasets for computer vision [9], etc. However, the progressive development of computer graphics requires better modeling of 3D scenes and generation of layouts. Therefore, we have been investigating techniques of automatically generating scene layouts.

Generating scene layouts benefits various applications. First, it saves the effort of manually placing objects in video games or industrial designs². Li et al. [11] generate various layouts for better simulations of wheelchair training. Handa et al. [9] generate multi-view images from much fewer 3D scenes.

Existing works already show the progress of scene synthesis [31], where scene layouts focus on their plausibility and aesthetic, i.e., visual identifications given generated layouts. Existing works are divided into neural network based techniques and others. The former trains several neural net-

works for different steps such as placing objects, rotating objects, deciding termination of arrangements [25]. The latter formulates a set of mathematical models including graphs, and typically optimize a shuffled area based on e.g., Markov Chain Monte Carlo (MCMC) [30, 19], since the models are too complicated to be solved. Nevertheless, algorithmic methods have not been investigated as far as we reviewed, because similarly we have to embed layout rules into an algorithm so that it operates properly. However, layout rules are innumerable. A qualitative comparison of existing techniques is beyond the scope of this paper. Despite underlying technical details, this paper focuses on the final results, i.e., improving the plausibility and aesthetic of generated layouts.

In this paper, we propose an algorithmic framework for generating room layouts as shown in Figure 1. Our framework is split into a data-driven phase: coherent grouping and a non-data-driven phase: geometric arranging. In the coherent grouping phase, objects are clustered into several coherent groups (section 3), where priors are learnt for suggesting layouts within each coherent group. We directly use correct and denoised samples extracted from datasets as priors. This brings two benefits. First, we no longer hypothesize distributions of layout rules between/among objects, especially the mathematically inexpressible relations. Second, we could easily formulate and represent relations among three or more objects since we only have to load samples of real distributions. Similar to “hyper-graphs” where an edge connects to more than two vertices, we name our learnt relations among objects “hyper-relations” (section 4). Thus, several objects of the same coherent group are arranged in $O(1)$ time by sampling from their hyper-relations. In the geometric arranging phase, given independent coherent groups where objects of the same group are already properly arranged with respect to each other, a geometric algorithm is proposed to assign positions and orientations of each group. Since layout rules among objects are applied during the coherent grouping phase, geometry phase concentrates on much fewer rules related to walls, windows, etc (section 5).

Current technologies of synthesizing 3D scenes include selecting a set of appropriate objects and generating plausible layouts for the objects. We do “layout” while tech-

¹Code is publicly available at <https://github.com/Shao-Kui/3DScenePlatform>, including the proposed framework (algorithm) and a 3D scene platform (toolbox).

²<https://planner5d.com/>

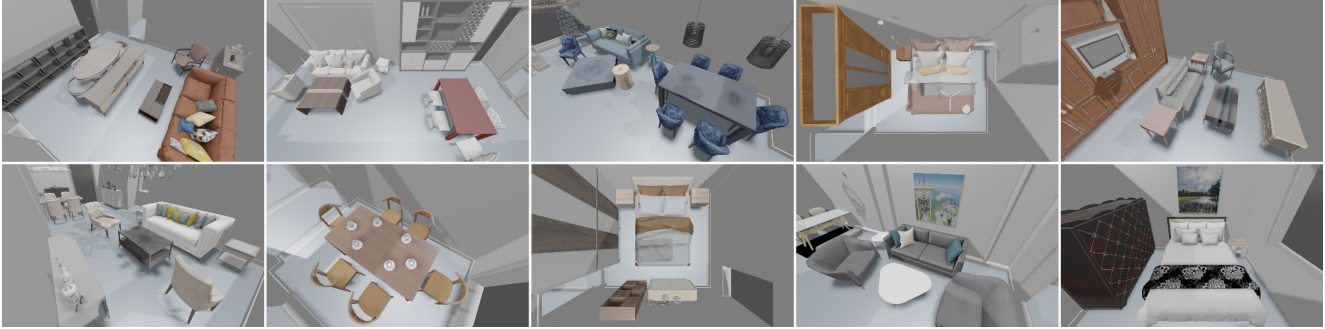


Figure 1: Our framework uniformly layouts objects, e.g., small objects on a surface are arranged concurrently instead of another layout problem. In addition to the overall plausibility, we emphasize reasonableness among objects related to each other, i.e., coherent groups. Ours is also friendly to objects hung on walls.

niques for selecting objects are easily incorporated such as [10]. Note that in this paper, we prefer instance-based priors to category-based priors, e.g., we consider a spatial relation between a specific coffee table to a specific chair, where both of them have unique textures and geometries. If categories are being based, distinct features of objects are lost. As shown in Figure 2, different shapes of several armchairs have their own priors to the same coffee table.

In this paper, we make the following contributions:

1. We first introduce and learn hyper-relations among three or more objects, which increases the aesthetics and plausibility of arranging objects in same coherent groups and only requires $O(1)$ time to sample layouts of each group, e.g., a coffee table surrounded with several distinct sofas and a TV stand, thus increasing the overall performance.
2. We propose a new scalable geometry-based framework for layout generation, which considers detailed aspects of room layout, e.g., doors, windows, wall decorations, small objects, etc. In coordination with hyper-relations, more plausible and robust layouts are generated.
3. We develop an open-source platform for manipulating 3D scenes, where operations such as rendering, exploring and modifying scenes are supported, thus allowing researches to focus thoroughly on algorithms.

2. Related Works

3D scene synthesis is to select a set of appropriate objects and transform them plausibly [31]. Earlier works of synthesizing 3D scenes are mainly based on hand-crafted design rules, e.g., [8, 16] or data-driven priors. For the former, designing rules are mathematically formulated as a set of constraints followed by optimizations [18, 26]. For the latter, since learnt distributions are too complicated to be

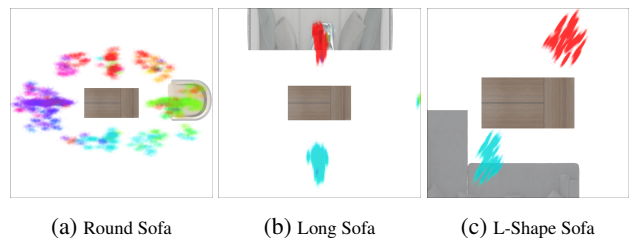


Figure 2: Extracting priors based on instances results in finer priors. This figure shows the priors of the same coffee table with respect to three sofa instances of different shapes and geometries.

differentiated, MCMC is assembled to solve such situation by attempting proposals [30, 14, 13, 19, 29].

Some of them present a framework including both object selection and layout generation, while the rest focus on layouts, though it may also focus on selecting objects [10]. Our method focuses on generations of layouts, i.e., we contribute mainly on how to make layouts more plausible and robust.

With the continuous study of the deep learning technique, several works based on convolutional or graph neural networks are proposed [25, 24, 20], including the current state-of-art work PlanIT [24] which serves as the baseline in this paper. One feature of network-based works is that they couple selections and layouts, i.e., selecting an object depends on pending layouts, vice versa. In contrast earlier works aforementioned separate two stages. Literature based on other techniques does exist, e.g., human-centric assessments [7]. Please refer to a more insightful survey on 3D indoor scene synthesis [31].

Several works also synthesize 3D scenes with input other than 3D scenes. Xu et al. recovery 3D scenes from hand sketch [28]. Luo et al [15] generate 3D scenes from scene graphs. [1, 3, 22] generate room layouts based on RGB-D images or 3D scans. [27, 5] generate scenes based on in-



Figure 3: Three coherent groups where white dots denote respective dominant objects.

put examples. [2, 17] translate human language to 3D scene configurations. However, different input results in different constraints, frameworks and even applications, so these works are beyond the scope of this paper.

3. Definitions

Given a list of objects with the positions of doors and windows and a room shape, we formulate its corresponding graph $G = \langle V, E \rangle$ where each object $o \in V$. E is the set of edges which are also simply relations between/among objects. Note that in this paper, we assume that a relation may involve more than two instances i.e., a hyper-relation among objects (section 4).

A coherent group g is a list containing objects where one object connects to at least one another object in the same group. In other words, two coherent groups never have an edge between their instances respectively. Conceptually, generating coherent groups $g \in G$ is equivalent to formulating maximal connected subgraph of G , given priors as connections. When generating layout given input, we always initially group objects into several $g_i \subset V$ even though a group may contain one object, such as a wardrobe, a picture frame, a kitchen cabinet, etc. Coherent groups are hierarchical as shown in Figure 3, where visual edges are pairwise between parents and children.

A transformation of an object includes its translation (x, y, z) and Y-axis rotation θ where floors align with XoZ plain. In this paper, we do not re-scale objects. The same is true of coherent groups. Additionally, transformations of coherent groups are propagated to their subordinate objects.

Priors are used to group objects into coherent groups and suggest layouts within each coherent group. A prior set $P_{o_1, o_2, o_3, \dots, o_n}$, abbreviated as P_O , involves two or more objects. A single prior $p_O^k \in P_O$ suggests a set of plausible transformations for all objects involved. Each prior set contains a dominant object such as o_1 , and other secondary objects. For example, if a dining table is surrounded with several chairs and supports a plant, “dinning table” is the dominant object in this scenario and remaining objects are secondary objects. If only two objects are involved in P_O , P_O is a “pairwise relation” between the two objects (section 4.1). If more than two objects are involved and all



(a) Pairwise Relation. (b) Pattern Chain. (c) Hyper-Relation.

Figure 4: Three types of priors in this paper. Links with the same color suggest same secondary objects. 4a: a pairwise “one-to-one” relation between a desk and a chair; 4b: a pairwise “one-to-many” relation between a table and several identical chairs; 4c: a hyper-relation among several different objects dominated by a coffee table.

secondary objects derives from the same instance, P_O is a “pattern chain set”. Otherwise, P_O is a “hyper-relation” 4.3.

4. Priors

In this section we show how we extract relations among objects. We start by extracting traditional pairwise relations, e.g., a desk with respect to a chair. Then, we present pre-computed pattern chains which generalize one-to-one relations to one-to-many relations, e.g., a dining table surrounded by several identical chairs. Finally, based on pairwise relations and pattern chains, we further generalize and formulate hyper-relations among objects, i.e, relations “among” more than two instances. Figure 4 suggests the differences between them. In this section, we show how priors are represented and generated, and the usage is shown in section 5.

Theoretically, pairwise relations and pattern chains are both special forms of hyper-relations. The reason for discussing them separately is because directly learning hyper-relations is difficult. As a result, we first introduce pairwise relations which derive more general pattern chains thus enabling forming hyper-relations.

4.1. Pairwise Relation

A Pairwise relation is a set of priors P_{ab} from a dominant object a to a secondary object b . Given a pairwise relation P_{ab} , we can sample a prior $p_{ab,k} \in P_{ab}$ that is directly a transformation of b with respect to a . Note that pairwise relations are directional, and sampled transformations are only relative between two objects involved i.e., global transformations are still required (section 5).

We extract discrete pairwise priors by utilizing density peak clustering (DPC) [21], which firstly calculate $\rho_k = \sum_{k'} I_{\{d \leq d_c\}}(d_{k,k'})$, $d_c = d_{(0.015K^2)}$ and $\delta_k = \min_{k': \rho_k < \rho_{k'}}(d_{k,k'})$ for all points. In our situation, $d_{k,k'}$ denoting the Euclidean distance from the transformation of dominant object k to the transformation of secondary object k' . A transformation includes translations and rotations. d_c is a hyper-parameter and rho_k is the number of $d_{k,k'}$ that

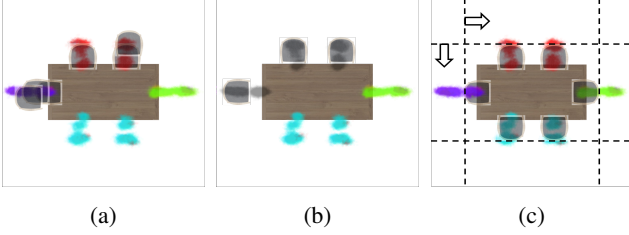


Figure 5: **5a**: Directly sampling a pairwise relation without pre-computed pattern results in obvious implausibility. **5b**: Recursively formulating a pattern chain. **5c**: Additional constraints are optional if e.g., well-aligned layouts are required.

is lower than d_c . The selection of d_c follows [21], i.e., the $0.015K^2$ th greatest $d_{k,k'}$ among all k^2 relative distances. δ_k seeks a minimal $d_{k,k'}$ among all $d_{k,k'}$ with higher $\rho_{k'}$ than ρ_k . Please refer to [21] for more details about this algorithm. Although DPC is typically used for clustering, it does anomaly detection for eliminating noises, i.e., removing points with low values of ρ and high values of δ . Cluster centers and ordinary points are treated equally since they are already reasonable transformations in this paper.

After elimination, remaining “points” are plausible relations directly from datasets (human designers) where each “point” become a single pairwise prior $p_{ab,k} \in P_{ab}$ for locally arranging a dominant object and its secondary object. Typical dominant objects include desk, dining table, coffee table, bed, etc. We manually label a set of instances that are capable of being dominant objects according to their categories.

4.2. Pre-Computed Pattern Chain

Commonly, a dominant object has several secondary copies of the same instance, e.g., a dining table with several identical chairs. If we sample them twice or more as shown in Figure 5a, aforementioned pairwise relations do not guarantee the plausibility of “one-to-many” relations. Thus, we solve it by presenting pattern chains.

A pattern chain set C_{ab} is a prior set between object a and b . Each $c_{ab}^j = \{j_1, j_2, \dots, j_n\}$, $c_{ab}^j \subset \mathbb{N}$ is a list of indices to its pairwise relation P_{ab} , e.g., j_x indexes to the x -th pairwise relation p_{ab,j_x} in P_{ab} . Generating one pattern chain c_{ab}^j is a recursive process. First, a $p_{ab,j_1} \in P_{ab}$ is randomly selected from P_{ab} . As discussed, p_{ab,j_1} gives a plausible transformation between a and b . Second, we traverse all $p_{ab,i} \in P_{ab}$. If a copy of object b with the transformation of $p_{ab,i}$ do not collide with another copy with the transformation of p_{ab,j_1} , $p_{ab,i}$ is included in a new subset $P'_{ab} \subset P_{ab}$. Third, we would like to place another copy of b , so p_{ab,j_2} is randomly selected from P'_{ab} and the above procedure is executed recursively until P'_{ab} is empty.

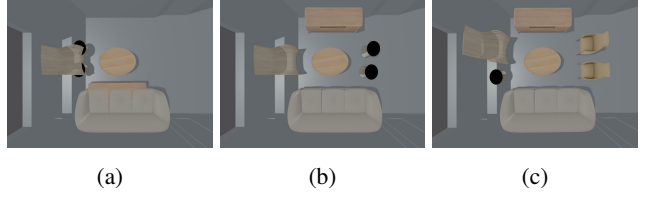


Figure 6: **6a**: Using only pairwise relations results in implausibilities among secondary objects in a coherent group. **6b**: Using hyper-relations results possibilities among all objects involved. **6c**: A different object set requires another hyper-relation, since we can not assume “as many objects as possible”.

As shown in Figure 5b, after three iterations, placing three chairs around a table filters out a subset of their pairwise priors (gray). Therefore, a fourth chair can only be placed in the remaining pigmented areas. When a chain is generated, we can optionally adjust it, e.g., Figure 5c suggests “horizontals and verticals” to make the chain well-aligned.

Note the above process generates one pattern chain $c_{ab}^j = \{j_1, j_2, \dots\}$. In theory, a P_{ab} of $O(n)$ size has $O(n!)$ undetermined pattern chains. In practice, we only generate one pattern chain for each $p_{ab,k} \in P_{ab}$, to make sure each pairwise relation is used at least once, instead of figuring out the entire pattern chain set. Otherwise, it requires $O(n!)$ time and space to compute only a single set, which also slows down online arrangement by restricting prior loading.

4.3. Hyper-Relation

A hyper-relation H_O is a prior set among several objects $O = \{o_{dom}, o_{sec1}, o_{sec2}, \dots\}$. A dominant object o_{dom} exists in H_O such as a coffee table and secondary objects relate to each other, e.g., chairs on a rug, armchairs beside a long sofa. Purely sampling pairwise prior sets results in scenarios such as Figure 6a, where secondary objects are only plausible with respect to their dominant object. Hyper-relation is essentially different from pattern chains. Pattern chain sets are still one-to-one relations and a pattern chain assumes incorporating as many secondary objects as possible. In contrast, a hyper-relation has a definite list of objects, i.e., we can not assume what instances are included and how many copies each instance has in a specific hyper-relation, because areas are limited. As shown in Figure 6b and 6c, different numbers and instances of seats derives two distinct hyper-relations.

To generate hyper-relations, we do not hypothesize and learn concrete distributions because real distributions are too complicated to be expressed, solved and sampled [12]. Instead, we try achieving as many exact samples as possible. Given a set of objects O and its dominant object $o_{dom} \in O$, we randomly select a secondary object $o_{sec} \in O$

and randomly sample a prior from the pairwise relation between o_{dom} and o_{sec} . Thus, o_{sec} is transformed with respect to o_{dom} . Next, similar to generating pattern chains, we filter the remaining pairwise relations between o_{dom} and other secondary objects $o_{sec1}, o_{sec2}, \dots \in O$, to ensure “collision free”. With multiple instances, additional rules are required. We use “tiers”, which as far as we studied is firstly terminology in [30], for finer filtering. For example, rugs are placed on the ground where objects such as tables and beds can be put on top of it. Merely detecting collisions would mistakenly filter plausible priors. Not detecting collisions between objects of different tiers alleviates such situations. After filtering the remaining pairwise relations, recursively, we randomly select another secondary object and repeat the above steps until all secondary objects are placed appropriately with no implausibilities. After that, a single hyper-prior is generated with transformations of all secondary objects. We iteratively re-run the entire process to enrich the pending hyper-relation.

Yet, the above steps still require a definite list of objects. Nevertheless, figuring out all undetermined lists is almost equivalent to exhaustively traverse all combinations of objects. To address this, we systematically optimize extractions. After forming coherent groups (section 5), we examine their hierarchies. If a parent has two or more children, we try to assemble the hyper-relation for them. If the hyper-relation does not exist, a new thread is started to generate it in background. In other words, we either load existing hyper-relations if they are already generated or establish a thread for generating them when we need them. Alternatively, users can manually suggest their own lists of objects to generate their hyper-relation.

5. Geometry-Based Layout Generation

5.1. Coherent Grouping

We show how we arrange objects in this section. First, objects are decomposed into several coherent groups $g_i \subset G$ based on finding maximal connected subgraphs using pairwise relations between objects as shown in Figure 7, where whether or not two objects are connected depends on existence of pairwise relations between objects.

One secondary object can have at most one dominant object. If multiple available dominant objects exist with respect to a secondary object o_{sec} , we randomly select a dominant object and discard relations between o_{sec} and other dominant objects. Each dominant object also has finite lengths of copies of secondary instances guided by lengths of respective pattern chains. This makes our framework more flexible, e.g., given only one chair but a dressing table and a desk in a bedroom, we randomly assign the chair to either the dressing table or the desk, which gives more variance to the generated results.

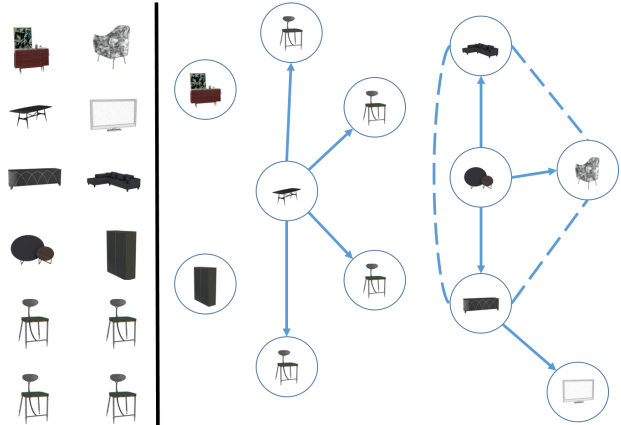


Figure 7: Coherent Grouping. Dotted dashes denote hyper-relations of secondary objects. Given a list of objects to generate their layout, we first group them into several coherent groups. For example, a coffee table relates to two sofas and a TV stand and the TV stand relates to a TV, so they form one coherent group. Two cabinets have no relation to others, so each of them form their own groups.

After that, input objects are distributed in coherent groups. As discussed in section 4, within a specific coherent group, we can directly sample a set of transformations for all objects locally within the group. As shown in Figure 7, if a parent has two or more descendants and each descendants are different, the hyper-relation is assembled or started to be generated in background, e.g., coffee table with respect to two sofas and a TV stand. If the descendants are identical, the pattern-chain set is sampled, e.g., dining table and four chairs. Otherwise, we use pairwise priors, e.g., TV stand and TV. Therefore, the final process is to transform several coherent groups properly in the room.

5.2. Geometric Arranging

Eventually, we assign transformations to each coherent group and propagate transformations to objects. Since priors already arrange objects sufficiently within groups, three more constraints are required to make layouts physically plausible among groups: 1, all groups should be inside a room; 2, all groups should not overlap each other; 3, clear paths should exist for windows and doors.

Placing a set of shapes (coherent groups) in another larger polygon (room) is an NP-hard problem [4] in computational geometry. Thus, we geometrically simplify coherent groups as cuboids, consider doors and windows as fixed (pre-arranged) blocks, and do heuristic attempts as shown in algorithm 1.

We first sort coherent groups according to their area occupied from largest to smallest, since bigger groups usually represent more central functionality of rooms, e.g., a bed-

room is call a “bedroom” due to a coherent group dominated by a bed. Then, coherent groups are placed with regard to this order, whereas random positions are assigned to them along the inner side of the targeting room, since the definition of coherent groups indicates the relations among different coherent groups are weak (section 3). After placing a group, we check potential collisions between this group and other groups or blocks. If collided, we discard the transformation and randomly re-select a new transformation. To enhance the performance, we used exponentially increasing sampling density. If a proper transformation fails at a density of d for the pending coherent group, we increase d to $2d$ to find more possible positions. But if it still collides after several times of increasing density, we discard the group and conduct the next one. To increase the plausibility, we add more heuristic rules: 1, we initially attempt to transform groups at corners of rooms and sides of other existing coherent groups. During collision detection, we take the height into consideration. So it is possible that some furniture with lower height is placed in front of windows. Finally, “liftings” L_f are assigned to groups. If $L_f = 0$, a groups is placed against walls. If L_f equals to half the length of the room, a group is placed in the middle of the room.

6. Experiments

6.1. Setup

We utilize a recent 3D scene dataset “3D-Front³” [6] with 70,000+ layouts and 9,992 3D models. To roam and render 3D scenes, we develop an open-source 3D scene platform as shown in Figure 8, where we can add, delete, modify and search for objects. We can orbitally control the perspective camera for selecting better views. By clicking “layout”, a configuration of the current room is generated by our proposed framework. We render 3D scenes using Three.js⁴ and the algorithm is mainly implemented by PyTorch and NumPy. Several results are shown in Figure 9. Please refer to our supplementary materials for more details.⁵

6.2. Plausibility and Aesthetic

We compare our framework with the state-of-the-art PlanIT [24]. PlanIT not only performs object arrangements, but also object selections. However, since we focus on arranging objects, we show better plausibility and aesthetic achieved using our framework by re-arranging results of PlanIT, i.e, we generate layouts given objects and room shape selected by PlanIT.

³<https://tianchi.aliyun.com/dataset/dataDetail?dataId=65347>

⁴<http://threejs.org/>

⁵We also run our framework on SUNCG [23] before this dataset became unavailable. We include results of SUNCG optionally in our supplementary materials only to verify the effectiveness of our framework.

Algorithm 1 Geometric Arranging

Input:

- 1: Polygon of room’s inner side P_r ;
- 2: List of rectangles of coherent groups with height A_{rec} ;
- 3: List of rectangles of windows and doors;

Output: Transformations of rectangles T_{rec} ;

```

4: function CHECKOK( $A$ )
5:   if  $A$  does not overlap with existing groups and
     blocks then
6:     return True
7:   else
8:     return False
9:   end if
10: end function
11: function APPLYTRANSFORM( $A, t$ )
12:   apply transformation  $t$  to  $A$ 
13:   return  $A$ 
14: end function
15: function INSERTRECTANGLE( $A$ )
16:   Let  $T$  be array of transformations //For heuristic
17:   for  $edge \in P_r$  and  $p \in$  existing polygons do
18:     Push heuristic transformation of  $edge$  or  $p$  to  $T$ 
19:   end for
20:   for  $t \in T$  do
21:     if CheckOK(ApplyTransform( $A, t$ )) then
22:       return  $t$ ;
23:     end if
24:   end for
25:   //For random
26:   Clear  $T$ 
27:   for  $n = 1 \rightarrow max$  sampling density do
28:     for  $edge \in P_r$  do
29:       Push  $2^n * len(edge)$  random transforma-
30:       tions on  $edge$  to  $T$ 
31:     end for
32:     Shuffle  $T$ 
33:     for  $t \in T$  do
34:       if CheckOK(ApplyTransform( $A, t$ )) then
35:         return  $t$ ;
36:       end if
37:     end for
38:     Clear  $T$ 
39:   end for
40:   for  $a \in A_{rec}$  do
41:     Push InsertRectangle( $a$ ) to  $T_{rec}$ ;
42:   end for

```

Qualitatively, as shown in Figure 11, ours is friendly for layouts among objects with strong relations, i.e., “coherent groups” in this paper. For example, a TV stand and a sofa

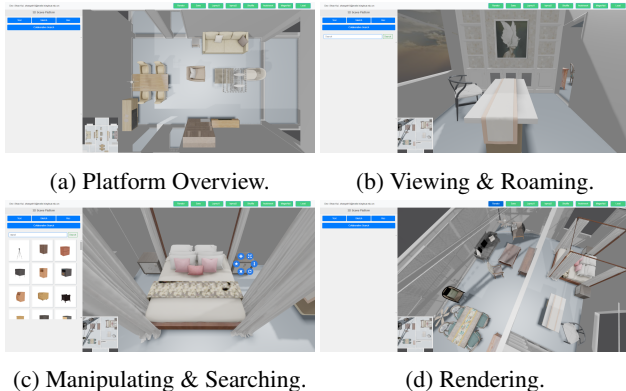


Figure 8: We develop an open-source 3D scene platform allowing adding, deleting, modifying, searching objects and rendering, saving scenes. Users can explore given 3D scenes by orbital control. Our platform is embedded with the proposed algorithm.

are strongly related to a coffee table. Ours makes sure they are plausibly arranged with respect to each other. Additionally, ours does not block paths to doors and windows. Quantitatively, we also conduct a user study as shown in Figure 10a. 43 subjects are invited. Subjects are university students, workers, housewives, interior designers, etc.⁶ Each subject is given 20 questions. In each question, a layout generated by ours and one by PlanIT are shown in random order. For each question, a subject compares two layouts and marks them respectively. Scores ranged from 0 (very poor) to 4 (very plausible). All subjects are taught how to use the user study system before experiencing. In Figure 10a, the Chinese characters are rendered as “there are two room layouts below, please compare the two layouts, considering aesthetic, plausibility and reasonableness, thus marking them respectively.”, “0: totally unreasonable, in-aesthetic. It may never appear in the real world layout.” and “5: very aesthetic and plausible. I will refer to this layout in the real world.” For example, Figure 10a shows two scenes. One subject may dislike the one on the left and prefer the one on the right. Therefore, the subject marks 2 and 4 respectively. Table 1 compares average score (standard deviation) of the two methods on various types of rooms.

6.3. Robustness

In this section, we compare our generated layouts with those designed by professionals (ground truths) to verify that ours is competitive to human designers. Subjects are the same from section 6.2. Each subject is required to choose a most plausible layout from ten alternative layouts as shown in Figure 10b, where one layout is designed by a human designer and the remaining nine layouts are gener-

⁶Few subjects preserve privacy.

Table 1: User study: results of comparing PlanIT with ours.

| Room Type | PlanIT | Ours |
|--------------|---------------|---------------|
| Bedrooms | 1.847 (1.336) | 2.66 (1.125) |
| Living Rooms | 1.749 (1.327) | 2.572 (1.266) |
| Bathrooms | 1.028 (1.2) | 2.553 (1.314) |
| Kitchens | 1.549 (1.342) | 2.651 (1.167) |
| Total | 1.543 (1.341) | 2.609 (1.221) |

ated by ours. Subjects can zoom in layouts by right clicks such as Figure 10c. All subjects are taught before experiencing and manuals are available. Ground truths are randomly selected from 3D-Front. In Figure 10b, the Chinese characters are rendered as “there are ten layouts below and please select your favorite one considering aesthetic, plausibility and reasonableness”, “left-click for selections and right click for zooming in” and “after selecting, press submit for the next question”.

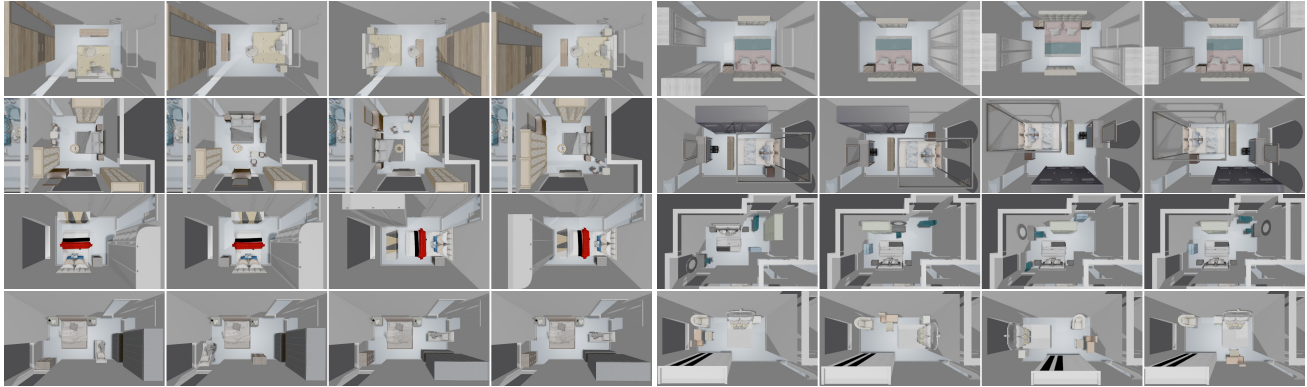
Results are shown in Figure 12. Two distributions are plotted for bedrooms and “living-dinning” rooms respectively, i.e., each line is averaged distributions of user selections of its room type, where “0” denotes ground truth. Although human-designed layouts outperform ours, generated layouts are still favored competitively as shown in Figure 12.

6.4. Efficiency

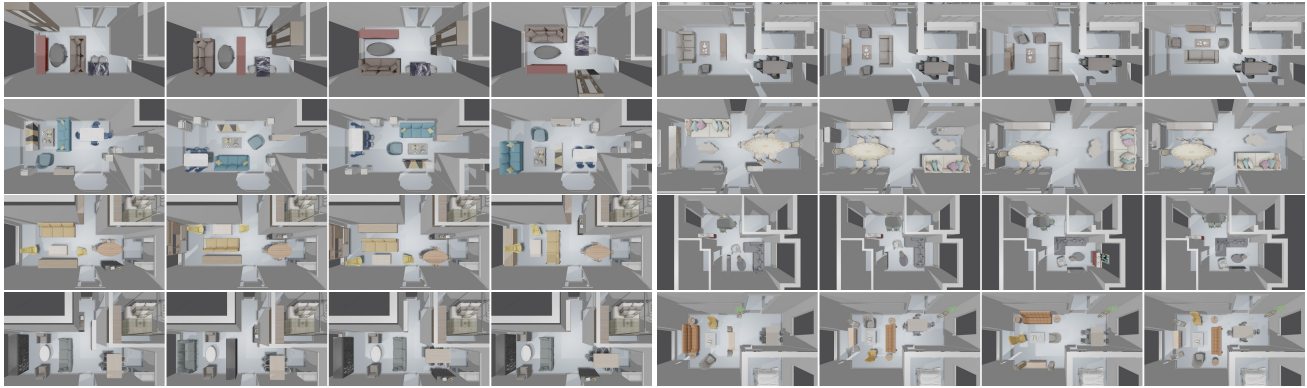
We run our framework on a PC with AMD 2700X (GHz), GTX 970, and WD20EZR. Time consumption of layouts depend on degrees of crowding, i.e., ratio of the total area of coherent groups to the area of the room. Higher degrees result in more discards during geometry-based arrangements (section 5.2), thus slowing down generations.

To generate layouts based on 3D-Front such as Figure 9, if priors are cached, our framework consume within 3.5 seconds for a layout. If corresponding priors of several objects are not loaded, additional IO is required up to 2 seconds for a layout. For non-crowded rooms, with cached priors, our framework generates layouts in real time.

We also run the state-of-the-art PlanIT [24] on a server with GTX 1080ti. According to our experiments, generating a layout requires more than a minute. Nevertheless, this includes both object selection and object arrangement and the two are interleaved with each other. Testing the exact time consumption of “layout generations” of PlanIT is beyond the scope of this paper. Furthermore, [26] is not a data-driven framework. Therefore, it is hard to conclude “better efficiency” as a contribution.

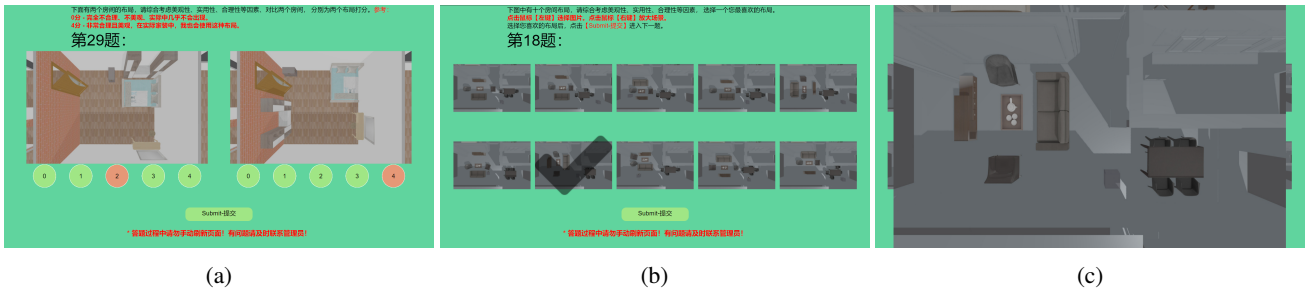


Bedroom.



Living Room & Dining Room.

Figure 9: Results. Please zoom in for more details. More results are included in the supplementary files.



(a)

(b)

(c)

Figure 10: User studies. **10a**: Marking ours and PlanIT [24] respectively; **10b**: Selecting the most plausible layout from ten alternative scenes where one scene is generated by human designers; **10c**: subjects can zoom in a particular layouts for better cognition.

7. Conclusions and Future Works

In this paper, we present a new framework of generating room layouts and we experimentally verify the plausibility and robustness of the proposed method. The code of this framework and a toolbox platform is publicly available. We hope this could benefit the community. However, this work still suffers from the following weaknesses.

The biggest difficulty we encountered is arranging “chains” of objects around walls. For independent objects

such as wardrobes, transformations of them have high degree of freedom since we find appropriate places for them with no collision and implausibility. However, for groups of objects such as kitchen cabinets and ovens, they are frequently placed next to each other as shown in Figure 13. Firstly, orders of a chain should be carefully considered. For example, commonly we place similar cabinet next to each other. Otherwise, layouts are not aesthetic as shown in Figure 13b. Secondly, an L-shape chain should somehow turn at corners, especially when we have L-shape objects such

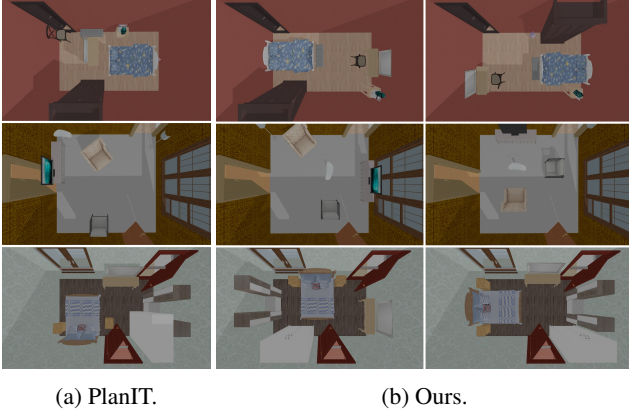


Figure 11: Qualitatively comparing PlanIT with ours.

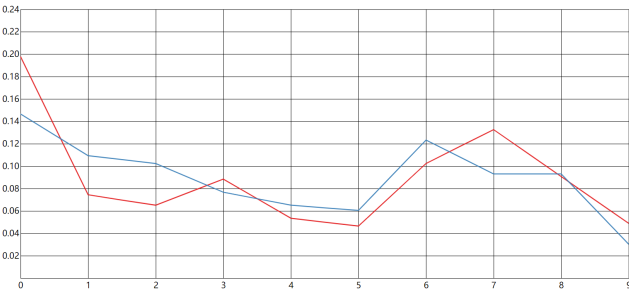


Figure 12: Distributions of user selected layouts of bedrooms (BLUE) and “living-dinning rooms” (RED).

as L-shape cabinets which are frequently treated as “corner objects” as shown in Figure 13c. Thirdly, doors and windows are also challenges for arranging chains. In our framework, if we treat a chain as an entire group, currently we do not have plans for sampling such priors. On the other hand, if we treat a chain as individual objects, complicated rules are required but we also do not have a plan for formulating the rules. As a result, we demonstrate this weakness in detail and we would try fixing it in future. Fortunately, in real-world decoration, most cabinets are fixed on walls.

The storage and loading of priors may require further system-level optimizations. Currently, all priors are structured in “json” format, which is inefficient if a prior of a coherent group is too large. When arranging objects online, loading priors may consume up to few seconds for loading corresponding priors into the memory. Although this only affects the first attempt, since priors are cached after that, it is still a concern in practice. Eventually, the way of extracting patterns is robust to shapes and textures instead of incorporating them. This is useful if we want objects arranged more compacted.

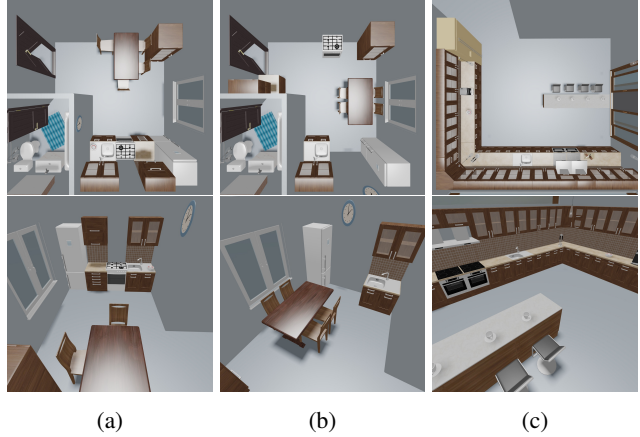


Figure 13: The problem of “chains” of objects around walls. 13a: The ground truth; 13b: A failure case of ours; 13c: An L-shape “chain” with L-shape objects.

Acknowledgements

This work was supported by the National Key Technology R&D Program (Project Number 2017YFB1002604), the National Natural Science Foundation of China (Project Numbers 61772298, 61832016), Research Grant of Beijing Higher Institution Engineering Research Center, and Tsinghua–Tencent Joint Laboratory for Internet Innovation Technology. We also thank all reviewers for their constructive comments and Yuan-Chen Guo for a thorough proof-reading.

References

- [1] A. Avetisyan, M. Dahnert, A. Dai, M. Savva, A. X. Chang, and M. Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. *arXiv preprint arXiv:1811.11187*, 2018. 2
- [2] A. Chang, W. Monroe, M. Savva, C. Potts, and C. D. Manning. Text to 3d scene generation with rich lexical grounding. *arXiv preprint arXiv:1505.06289*, 2015. 3
- [3] K. Chen, Y. Lai, Y.-X. Wu, R. R. Martin, and S.-M. Hu. Automatic semantic modeling of indoor scenes from low-quality rgb-d data using contextual information. *ACM Transactions on Graphics*, 33(6), 2014. 2
- [4] M. De Berg, M. Van Kreveld, M. Overmars, and O. Schwarzkopf. Computational geometry. In *Computational geometry*, pages 1–17. Springer, 1997. 5
- [5] M. Fisher, D. Ritchie, M. Savva, T. Funkhouser, and P. Hanrahan. Example-based synthesis of 3d object arrangements. *ACM Transactions on Graphics (TOG)*, 31(6):135, 2012. 2
- [6] H. Fu, R. Jia, L. Gao, M. Gong, B. Zhao, S. Maybank, and D. Tao. 3d-future: 3d furniture shape with texture. *arXiv preprint arXiv:2009.09633*, 2020. 6
- [7] Q. Fu, H. Fu, H. Yan, B. Zhou, X. Chen, and X. Li. Human-centric metrics for indoor scene assessment and synthesis. *Graphical Models*, 110:101073, 2020. 2

- [8] T. Germer and M. Schwarz. Procedural arrangement of furniture for real-time walkthroughs. In *Computer Graphics Forum*, volume 28, pages 2068–2078. Wiley Online Library, 2009. 2
- [9] A. Handa, V. Patraucean, V. Badrinarayanan, S. Stent, and R. Cipolla. Understanding real world indoor scenes with synthetic data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4077–4085, 2016. 1
- [10] Y. He, Y. Cai, Y.-C. Guo, Z.-N. Liu, S.-K. Zhang, S.-H. Zhang, H.-B. Fu, and S.-Y. Chen. Style-compatible object recommendation for multi-room indoor scene synthesis. *arXiv preprint arXiv:2003.04187*, 2020. 2
- [11] W. Li, J. Talavera, A. G. Samayoa, J.-M. Lien, and L.-F. Yu. Automatic synthesis of virtual wheelchair training scenarios. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 539–547. IEEE, 2020. 1
- [12] Y. Li, J. Zhang, Y. Cheng, K. Huang, and T. Tan. Df 2 net: Discriminative feature learning and fusion network for rgb-d indoor scene classification. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 4
- [13] Y. Liang, F. Xu, S.-H. Zhang, Y.-K. Lai, and T. Mu. Knowledge graph construction with structure and parameter learning for indoor scene design. *Computational Visual Media*, 4(2):123–137, 2018. 2
- [14] Y. Liang, S.-H. Zhang, and R. R. Martin. Automatic data-driven room design generation. In *International Workshop on Next Generation Computer Animation Techniques*, pages 133–148. Springer, 2017. 2
- [15] A. Luo, Z. Zhang, J. Wu, and J. B. Tenenbaum. End-to-end optimization of scene layout. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3754–3763, 2020. 2
- [16] G. H. Lyons. *Ten Common Home Decorating Mistakes & How to Avoid Them*. Blue Sage Press, 2008. 2
- [17] R. Ma, A. G. Patil, M. Fisher, M. Li, S. Pirk, B.-S. Hua, S.-K. Yeung, X. Tong, L. Guibas, and H. Zhang. Language-driven synthesis of 3d scenes from scene databases. In *SIGGRAPH Asia 2018 Technical Papers*, page 212. ACM, 2018. 3
- [18] P. Merrell, E. Schkufza, Z. Li, M. Agrawala, and V. Koltun. Interactive furniture layout using interior design guidelines. In *ACM transactions on graphics (TOG)*, volume 30, page 87. ACM, 2011. 2
- [19] S. Qi, Y. Zhu, S. Huang, C. Jiang, and S.-C. Zhu. Human-centric indoor scene synthesis using stochastic grammar. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5899–5908, 2018. 1, 2
- [20] D. Ritchie, K. Wang, and Y.-a. Lin. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6182–6190, 2019. 2
- [21] A. Rodriguez and A. Laio. Clustering by fast search and find of density peaks. *Science*, 344(6191):1492–1496, 2014. 3, 4
- [22] T. Shao, W. Xu, K. Zhou, J. Wang, D. Li, and B. Guo. An interactive approach to semantic modeling of indoor scenes with an rgb-d camera. *ACM Transactions on Graphics (TOG)*, 31(6):136, 2012. 2
- [23] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017. 6
- [24] K. Wang, Y.-A. Lin, B. Weissmann, M. Savva, A. X. Chang, and D. Ritchie. Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)*, 38(4):132, 2019. 2, 6, 7, 8
- [25] K. Wang, M. Savva, A. X. Chang, and D. Ritchie. Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)*, 37(4):70, 2018. 1, 2
- [26] T. Weiss, A. Litteneker, N. Duncan, M. Nakada, C. Jiang, L.-F. Yu, and D. Terzopoulos. Fast and scalable position-based layout synthesis. *arXiv preprint arXiv:1809.10526*, 2018. 2, 7
- [27] G. Xiong, Q. Fu, H. Fu, B. Zhou, G. Luo, and Z. Deng. Motion planning for convertible indoor scene layout design. *IEEE Transactions on Visualization and Computer Graphics*, 2020. 2
- [28] K. Xu, K. Chen, H. Fu, W.-L. Sun, and S.-M. Hu. Sketch2scene: sketch-based co-retrieval and co-placement of 3d models. *ACM Transactions on Graphics (TOG)*, 32(4):123, 2013. 2
- [29] Y.-T. Yeh, L. Yang, M. Watson, N. D. Goodman, and P. Hanrahan. Synthesizing open worlds with constraints using locally annealed reversible jump mcmc. *ACM Transactions on Graphics (TOG)*, 31(4):56, 2012. 2
- [30] L.-F. Yu, S. K. Yeung, C.-K. Tang, D. Terzopoulos, T. F. Chan, and S. Osher. Make it home: automatic optimization of furniture arrangement. *ACM Trans. Graph.*, 30(4):86, 2011. 1, 2, 5
- [31] S.-H. Zhang, S.-K. Zhang, Y. Liang, and P. Hall. A survey of 3d indoor scene synthesis. *Journal of Computer Science and Technology*, 34(3):594, 2019. 1, 2