

DStruct2Design: Data Structure Driven Generative Floor Plan Design

Anonymous authors

Paper under double-blind review

Abstract

Text conditioned generative models for images have yielded impressive results. Text conditioned floorplan generation as a special type of raster image generation task has also received particular attention. However there are many use cases in floorplan generation where numerical properties of the generated result are more important than the aesthetics. For instance, one might want to specify sizes for certain rooms in a floorplan and compare the generated floorplan with given specifications. Current approaches, datasets and commonly used evaluations do not support these kinds of constraints. As such, an attractive strategy is to use as input a data structure which directly encodes those constraints and can therefore be conditioned on to generate the final floorplan. To explore this setting we (1) convert popular image based floorplan datasets into text format for this data-structure to data-structure formulation of floorplan generation using two popular image based floorplan datasets RPLAN and ProcTHOR-10k, and provide the tools to convert further procedurally generated ProcTHOR floorplan data into our format. (2) We explore the task of floorplan generation given a partial or complete set of constraints and we design a series of metrics and benchmarks to enable evaluating how well samples generated from models respect the constraints. (3) We create multiple baselines by finetuning a large language model (LLM), Llama3, and demonstrate the feasibility of using floorplan data structure conditioned LLMs for the problem of floorplan generation respecting numerical constraints. We hope that our language-based approach to this image-based design problem and our newly developed benchmarks will encourage further research on different ways to improve the performance of LLMs and other generative modelling techniques for generating designs where quantitative constraints are only partially specified, but must be respected.

1 Introduction

Generative modelling has the potential to accelerate and improve design tasks but in order for them to achieve widespread use by real-world practitioners they need to be flexible and exhibit a sufficiently high degree of controllability. For instance, in floorplan generation, users will want to be able to specify a set of constraints detailing how rooms should be connected and what their dimensions are. The resulting model should also be flexible in the sense that it performs equally well on as few or as many constraints as needed. Just as there are many possible constraints that the user can specify, there also needs to be a rigorously-defined set of evaluation metrics to probe how well the model performs (or underperforms) on them, which we contribute in this work.

Most existing approaches to floorplan-based generative models focus on 2D top-down renders. These methods typically specify room connectivity via bubble diagrams as well as other information such as room geometries. However, these methods fall short in various areas, for instance some produce raster-based renderings of the floorplan (e.g. vision-based generative models) which make it difficult to directly access metadata. Others may only support the specification of adjacencies instead of also supporting room geometry, limiting the controllability of the model. In this work, we instead opt to represent generated floorplans as a JSON-based data structure specifying all the required metadata of the floorplan as well as specifying each room as a list of polygons (similar to Shabani et al. (2022)). These floorplans are a result of consolidating two existing

large floorplan datasets, namely RPLAN and ProcTHOR. Furthermore, we validate the use of such a dataset by training an LLM to produce such a structure conditioned on constraints specified in natural language. Therefore, our contributions are as follows:

1. We contribute a unified floorplan dataset by merging existing popular datasets (RPLAN and ProcTHOR) and postprocessing them into a format which is amenable to language-based generative models which take as input and output structured text.
2. We design an array of metrics to benchmark the performance of the generated floorplans, in particular how well they respect the constraints set by the user.
3. We train and evaluate an LLM on our task setting to demonstrate the feasibility of our proposed problem formulation.

2 Related Work

Datasets. One of the most commonly used datasets for floorplan generation is RPLAN (Wu et al., 2019). Since we make use of the dataset in this work, we defer its description to Section 3.2. LIFULL (National Institute of Informatics, 2015) is an extremely large-scale dataset consisting of floorplans sourced from a Japanese realtor company. Floorplans are provided in raster format, however a subset of vectorized versions exist (Liu et al., 2017). One common shortcoming is that there is a lack of well-annotated floorplan data (e.g. vectorised) which is amenable to training. Architext (Galanos et al., 2023) address this issue by generating synthetic but diverse floorplans via a CAD script for Rhinoceros 3D. In a similar vein, we leverage ProcTHOR (Deitke et al., 2022) (also used and described in more detail in Section 3.2) which is a fully open source simulator of buildings (and by extension floorplans). Unlike Rhinoceros 3D this simulator is completely open source.

Generative Models. *House-GAN* and *House-GAN++* (Nauata et al., 2020a; 2021a) are a family of GAN-based methods which learn to generate floorplans using convolutional graph-based networks. Notably however this line of work only conditions on a bubble graph. *FloorplanGAN* (Luo and Huang, 2022) proposes a self-attention-based GAN which takes as input (for each room) room centers, desired areas (encoded as relative proportions), as well as room type for each room. Because the GAN is tasked with also refining the initial input constraints, the output may not respect the original ones. While the use of a differentiable rasteriser to compute losses in pixel space opens up possibilities, the core method does not appear to support a partial specification of constraints or polygons. *HouseDiffusion* (Shabani et al., 2023) is a diffusion-based method which directly predicts a list of polygons for each room, utilising a transformer architecture which also conditions on a bubble diagram.¹ Lastly, *Architext* (Galanos et al., 2023) also leverages LLMs to generate output floorplans but it appears the prompts are only limited to natural language descriptions rather than geometry.

3D Scene Generation. Recently, full 3D scene generation methods have shown impressive results. AnyHome (Wen et al., 2023) and Holodeck (Yang et al., 2024) are able to generate floor plans, windows, doors, furniture and meaningful placement in 3D all from a natural language query (e.g. ‘a 1b1b apartment’). In our method, we focus only on the floor plan aspect, but we allow for specification of room dimensions and areas as well as total floor plan area, which neither method does. Since our proposed dataset is partially comprised of ProcTHOR (a 3D-based floorplan simulator), in principle a future version of our proposed dataset could support the placement of furniture and other props. An overview of how our method compares to the others can be found in Tab. 1.

3 Our Text-Based DStruct2Design Floorplan Dataset

In this work, we view the problem of floor plan generation from a text-based data structure perspective. Specifically, we wish to utilize the data structure to enable, through language input, a unified method that

¹While this method could in principle be conditioned on an initial polygonal room specification (via partial noising and subsequent denoising), from the point of view of pure inference (that is, starting from pure noise) *HouseDiffusion*’s only form of conditioning is via a bubble diagram.

Method	Input	Output	Can specify Geometry?	Absolute Coordinates?
House-GAN, House-GAN++, HouseDiffusion	Bubble diagram (room adjacency)	Room masks	no	no
ArchiText	Natural Language	Room polygon coords	no	yes
FloorplanGAN	Fixed-Length Tensor	Room masks	yes	no
AnyHome, Holodeck	Natural Language	Furnished 3D models	no	yes
Ours	Structured Language	Room polygon coords	yes	yes

Table 1: **Comparison to related works.** Our method focuses on generation of floor plans as opposed to full 3D scenes, and it allows the user to specify individual room geometry (height, width, and/or area) and global apartment geometry in meters and sq. meters.

not only allows users to apply numerical constraints but also retains the ability to condition on graphs such as bubble diagrams, which have been the input in the traditional floorplan generation task. As such, the choice of data representation for the floorplans is very important. In our work, we carefully design a JSON based data structure to be used as a new representation for the floorplans. As prior work all use datasets of floorplans in 2D image or 3D scenes, we create a new dataset by converting existing image and scene data into our data structure for the use of our task.

3.1 Data Structure

In this structure, we define crucial numerical data such as the number of rooms present, the total area and type of each room that appears in each floor plan. Each room within a floor plan is further defined with room specific fields. Importantly, each room’s location is defined through a set of vertices that forms a polygon. This ensures that the language model cannot cheat by outputting vague locations, and instead it has to predict the exact coordinates of the vertices. The specific structure is presented in Table 2, and a full example of floorplan in this structure is presented in the Appendix ??.

This structure also has several additional advantages: **1.** Using numerical values gives us the ability to set numerical constraints in the input, and it allows for clear and explicit evaluation on how well the generation process adheres to these constraints. **2.** This structure is designed to have sufficient information to be transformed into higher level representations such as floor plan images, allowing seamless integration with traditional visualization techniques. **3.** The format is highly extensible, enabling the inclusion of additional information in the generation process, such as the placement and attributes of objects within the floorplan. This flexibility ensures that our approach can be further adapted to various task expansions in the future.

3.2 Datasets & Preparation for Our New Task Formulation

ProcTHOR-10k (Deitke et al., 2022) is a dataset of 12,000 procedural generated, fully interactive 3D houses designed for research in Embodied AI. We clean and process the raw data from each house, focusing on the geometric properties of the rooms. We employ the shoelace formula to calculate areas and determined room dimensions based on their x and y coordinates. Additionally, we remove redundant points and apply rounding to the coordinates for consistency. Next, we categorize and count the types of rooms, and compute the total area for each house layout. The processed data is organized into the new JSON structure explained in Table 2.

RPLAN (Wu et al., 2019) is a manually collected dataset of 80,788 real world floor plans of buildings in Asia. Each floor plan in RPLAN is stored as a $256 \times 256 \times 4$ vector image. Channels 1 and 2 store interior and exterior boundary information; channel 3 contains room information where each pixel value denotes which room it belongs to; channel 4 has extra information to distinguish rooms with the same room type

"room_count"	▷ number of rooms
"total_area"	▷ total area in square units
"room_types" []	▷ list of string of room types present
"rooms" []	▷ list of individual room dictionary containing room specifics
"area"	▷ area of this room
"floor_polygon" []	▷ list of vertices that defines this room's layout
"x"	▷ x coordinate of the vertex
"y"	▷ y coordinate of the vertex
"is_regular"	▷ flag that indicates if the room's shape is rectangular
"height"	▷ the y-axis length of the rectangle bounding box enclosing the room
"width"	▷ the x-axis length of the rectangle bounding box enclosing the room
"id"	▷ unique id of this room
"room_type"	▷ the type of this room
"edges" []	▷ defines the connection between rooms for bubble gram generation

Table 2: Processed Floorplan Structure

value in channel 3. To convert this 4 channel image into a JSON structure with well-defined room location, we first extract all pixel coordinates for all of the rooms. For each room, we locate the pixels that make up its perimeter. Then, by tracing the perimeter in one direction we are able to capture all of the vertices in an order that allows recreation of the room polygon. We take the room's type and deduce all the other fields listed in Table 2 to complete the data structure. We convert 80,315 floorplans from RPLAN.

3.3 Bubble Diagrams

Traditionally, much of the prior work has assumed that at the beginning of the floor plan design process a bubble diagram is used to conceptualize the layout of the floor plan. As shown in Figure 2, a bubble diagram is used to represent different room and their spacial relationship with one another as a guidance and constraint for the floor plan design process to follow. Formally, a bubble diagram is a graph $\mathcal{G} = (N, E)$ where each node n_k in $N = \{n_j\}_1^R$ represents the k^{th} room in the floor plan with R rooms, and each edge $e_i = (n_p, n_q)$ in $E = \{e_j\}_1^M$ denotes a connection between room p and q .

In this work, we enable the use bubble diagram as an additional conditioning in our floor plan generation process. To obtain the bubble diagram we check pair-wise proximity of the rooms in each floor plan. If the Manhattan distance between the rooms' boundaries are within a threshold, they are counted as connected. This threshold is set to 8 pixels for the RPLAN dataset in line with prior work, and to 2 pixels for ProcTHOR to account for the length unit differences between the two datasets. This adjacency information is stored in the `edges` field in our data representation. During training, we use \mathcal{G} as conditioning for our floor plan generation.




Figure 1: Model Pipeline Overview

4 Our DStruct2Design (DS2D) LLM

We train an LLM by prompting it with structured numerical constraints and or graph constraints (Bubble Diagram) and asking it to predict the converted floor plan, as depicted in Figure 1.

Numerical Constraints $\mathcal{C} = \{c_j\}_1^T$ are a set of T conditions that users may impose on the final floor plan. In the task of floor plan design and generation, these can include, but are not limited to, total square footage of the entire floor, the number and the types of rooms present, the size of each room. These conditions are mostly inherently numerical, and thus we can take advantage by directly using it in its data structure form. For example, the constraint of having a total square footage of 990 can be transformed into `{"total_area": 990}`. This allows a direct match between the input prompt and output JSON string, which may help the model understand the structure and its relationship better.

Bubble Diagrams are represented differently as conditioning to our model. We chose to pass it in as tuples of connecting rooms. For an edge $e_i = (n_p, n_q)$ connecting nodes p and q , it is formatted as `(room_p, room_q)`. Due to the possible existence of multiple rooms of the same type, which can be clearly distinguished in a graph but not easily in text, we decide to represent `room_p` and `room_q` by their room IDs in addition to their room types.

We process the constraint set and the bubble diagram into "constraint string" and "adjacency string". They are then combined with an instruction phrase that is used throughout the training process to create the final prompt. An example of the full prompt is shown in Appendix C

5 Our Metrics and Benchmarks

To assess floorplan quality, we focus on the numerical consistency of the generation. Specifically, we design two groups of metrics to evaluate what we call self-consistency and prompt-consistency:

- **Self Consistency** metrics measures how numbers agree with each other in the generated floor plan. For instance, this includes a metric to check if the area defined by the polygon vertices is the same as the area number presented in the `room_area` field.
- **Prompt Consistency** metrics evaluate how consistent the generation is to the constraints used in the prompt. An example of this type of metrics is one that measures if the generated number of rooms adds up to the number of room requested in the prompt.

In addition to the two main groups of metrics, we also incorporate **Compatibility** metrics which is used in past work (Shabani et al., 2022; Nauata et al., 2020b; 2021b; Johnson et al., 2018; Ashual and Wolf, 2019) to measure similarity between the input bubble diagram and the output floor plan. It is by definition the graph edit distance (Abu-Aisheh et al., 2015) between the input bubble diagram and the output diagram extracted from the output JSON. The extraction method is the same one used to generate the bubble diagrams in the first place as described in Section 3.3. The full array of metrics and their explanations are presented in Table 3.

6 Experiments

All of our experiments are ran on the LLaMA3-8B-Instruct variant of the LLaMA model family (Touvron et al., 2023). We train our models by running 8-bit quantization along with LoRA (Hu et al., 2021). Each of our experiments are also run on single RTX8000 GPU, they take less than a day to complete. The exact training parameters and data split details can be found in Appendix D

6.1 Model Variants

We train six variants of our model on the converted ProcTHOR dataset and four on converted RPLAN dataset. Each of the four models on RPLAN – **5-R**, **6-R**, **7-R**, **8-R** – is trained without floorplan data of a certain room count. 5-R is trained only on floorplans with 6, 7, and 8 rooms, etc. This follows prior work

Metric Type	Metric	Explanation
Self Consistency (SC)	† Total Area	percentage difference between stated total area vs. the sum of all rooms' state area
	‡ Polygon Area (P. Area)	percentage difference between a room's stated area vs. area calculated from polygon
	‡ Overlap	boolean check for existence of polygon overlap
	‡ Room ID (ID)	boolean check for duplicate room id
	‡ Room Count (R. Count)	boolean check to see if "room_count" field number is equal to the number of rooms in the floor plan.
Both SC and PC	◁ Room Height (R. H)	percentage difference between height obtained from polygon vs. stated height in generation (SC) or requested height (PC).
	▷ Room Width (R. W)	percentage difference between width obtained from polygon vs. stated width in generation (SC) or requested width (PC).
Prompt Consistency (PC)	• Total Area	percentage difference between the sum of room polygon areas vs. requested total area
	♣ Num. Room (Num. R.)	percentage difference between number of rooms in the floorplan vs. requested number of rooms in the prompt.
	♠ Room ID (ID)	precision and recall of room ids in the floorplan with respect to the ids present in the prompt
	♡ Room Area (R. Area)	percentage difference between a room's polygon area vs. requested room area in the prompt
	◇ Room Type (Type)	precision and recall of room types in the floorplan with respect to the room types present in the prompt
	○ ID-Type Match (IDvsType)	percentage difference between the number of rooms with correctly matching type and id in the floor plan vs. the total number of rooms present in the prompt
Bubble Graph	Compatibility	graph edit distance (GED) between the input bubble diagram and one extracted from the output floor plan.

Table 3: Metrics and Associated Acronyms Used to Evaluate Generation Quality

	F	M	ProcTHOR Models				RPLAN Models			
			PM	F+BD	M+BD	PM+BD	5-R	6-R	7-R	8-R
Bubble Diagram	-	-	-	✓	✓	✓	✓	✓	✓	✓
Masking	-	random	preset	-	random	preset	random	random	random	random

Table 4: Summary of model variants

(Nauata et al., 2020b; 2021b; Shabani et al., 2023). On the other hand, the six models trained on ProcTHOR are divided into two main variants: bubble diagram enabled model, and numerical constraint only model. In the **bubble diagram enabled model (BD)**, both the constraint string and the adjacency string described in Section 4 are used as part of the input. In contrast, in the **numerical constraint only model**, only the constraint string is used. We further train three sub-variants for each of the two main variants to investigate how robust the models are to missing information:

- **Full-Prompt (F)** model takes advantage of every attribute available in the constraint set. During training, every constraint attribute is part of the constraint string.

- **Mask (M)** model uses a subset of constraint set by applying a 50% masking on every single possible constraint. In the case of **rooms** which is an attribute containing a list, the masking is applied independently to individual item in the list. (**rooms** is only dropped out if the entire list becomes empty.) As a safety measure, we always keep at least 1 constraint in the set.
- **Preset Mask (PM)** randomly selects one of four preset constraint sets with varying degree of missing information. The idea is to not just have IID random masking, but a hierarchy of attributes from general to specific which are detailed in Table 10 in Appendix F.

The summary of our six variants and their abbreviations are shown in Table 4

6.2 Generation Prompts

To test the model, we run four different generation prompts with varying amount of constraints in the prompt. This simulates the real world problem of floorplan design, where the designer often receives only partial criteria. For instance, the user might only ask the living room and bedroom to be a certain size, and leaves the designer with the freedom to imagine layouts with varying sizes for the other rooms. The four generation prompts are listed in order of decreasing amount of constraints:

- **Specific (S)** is one where we use all the possible constraints
- **All Room Area (AR)** is one where we pass in the area of all the rooms (total area can be inferred from this information).
- **Partial Room Area (PR)** is similar to Total Area. In addition to the total area of the floor plan, the area of some of the rooms are also passed as conditions in the constraint string.
- **Total Area (TA)** is one where we only use the total area of the floor plan as constraint.

For all of the Bubble Diagram model variants, Bubble Diagrams are used in each of these types of prompts.

7 Results




Figure 2: Pink background: ground truth floor plan and ground truth bubble diagram. Green background: generated floor plan using Specific prompts, and extracted bubble diagram. Blue background: generated floor plan using Total Area prompts, and extracted bubble diagrams.

	Self-consistency Benchmark						
	P. Area [‡]	Total Area [‡]	P. Overlap [‡]	R. Count [‡]	ID [‡]	R. H [◁]	R. W [▷]
F _S	0.94±0.07	1.00±0.03	0.16±0.37	1.00	1.00	0.99±0.02	0.99±0.02
M _{AR}	0.90±0.07	0.99±0.03	0.12±0.33	1.00	1.00	0.99±0.02	0.99±0.01
PM _{PR}	0.92±0.06	0.97±0.07	0.15±0.35	0.96	1.00	1.00±0.01	1.00±0.01
PM _{TA}	0.93±0.05	0.99±0.04	0.14±0.35	0.95	1.00	1.00±0.01	1.00±0.01
5-R _{AR}	0.90±0.16	0.98±0.05	0.37±0.48	0.97	1.00	1.00±0.01	0.98±0.02
6-R _S	0.93±0.05	0.98±0.06	0.36±0.48	0.99	1.00	0.99±0.01	0.97±0.02
7-R _S	0.93±0.04	0.98±0.05	0.42±0.49	0.97	1.00	0.99±0.01	0.98±0.02
8-R _{TA}	0.82±3.50	0.97±0.07	0.53±0.50	0.89	1.00	0.99±0.01	0.98±0.02

	Prompt-consistency Benchmark						
	Num. R [♣]	Total Area [♣]	R. Area [♡]	ID [♠]	IDvsType [◦]	R. H [◁]	R. W [▷]
F _S	1.00±0.00	0.95±0.07	0.94±0.07	1.00±0.03	1.00±0.00	0.99±0.02	0.99±0.02
M _{AR}	1.00±0.00	-	0.90±0.07	1.00±0.00	1.00±0.00	-	-
PM _{PR}	1.00±0.00	0.93±0.07	0.91±0.11	1.00±0.03	1.00±0.00	-	-
PM _{TA}	1.00±0.00	0.95±0.06	-	-	-	-	-
5-R _{AR}	-	-	0.92±0.04	1.00±0.01	1.00±0.00	-	-
6-R _S	1.00±0.00	0.94±0.07	0.94±0.06	1.00±0.05	1.00±0.00	0.92±0.23	0.94±0.14
7-R _S	1.00±0.00	0.94±0.06	0.94±0.06	1.00±0.02	1.00±0.00	0.93±0.17	0.94±0.29
8-R _{TA}	-	0.93±0.07	-	-	-	-	-

Table 5: Our self and prompt-consistency benchmark results. Subscripts under model variants are generation prompts.

We test floor plan generation quality on ProcTHOR-trained six model variants with the four different generation prompts with self and prompt consistency metrics. Results from our best model variants are listed in Table 5, and all results on all 48 sets of experiments are listed in Appendix G. We also perform the same evaluation on the four RPLAN-trained model variants. Because each model variant trained on RPLAN uses a slightly different set of data (following prior work), we show results for each model variant in Table 5. Results for the full 24 experiments are shown in Appendix I.

Analysis Our LLM models demonstrate a high level of competence in generating floorplans that are largely mathematically consistent, and are also consistent to input numerical constraint. As seen in Figure 2, when given full specifications, the generated rooms largely resembles ground truth in sizes, width and height. This quality is accurately reflected by our Total Area, R.H and R.W metrics. On the other hand, our metric suggests that generations can produce overlapping room layouts. An example of this is found in Figure 2, in the second row’s green background’s right-most floorplan: the top-left bathroom is overlapping on top of the bedroom. This problem is especially apparent in the RPLAN-trained models, and is accurately reflected by the P. Overlap metrics in Table 5. This limitation can perhaps be tackled in future work.

To compare with prior work that focus on using bubble diagrams as input on the RPLAN dataset, we prompt our RPLAN-trained models with bubble diagram. For evaluation we compute the standard compatibility metric obtained from graph edit distance calculation between the input bubble diagrams and the ones extracted from generations. Results are shown in Table 6. We also evaluate compatibility metrics on the ProcTHOR-trained BD model variants. The results are shown in Table 7.

Analysis Overall, our RPLAN-trained model generates results that are competitive with the state-of-the-art image based algorithms, even though our focus here has been on numerical accuracy. One caveat is that as the complexity of the generations increases (for example, the number of rooms increases, and the generations becomes much longer), the fine-tuned model starts to create some extreme floorplans. It would generate floorplans with a lot more overlaps causing the bubble diagram to be very different than the one used as constraints. In this sense, when there are more rooms in a floorplan, overlapping becomes a bigger issue, and its effect on the bubble diagram becomes more apparent. This causes our method to have a much higher

Model	Number of Rooms in Floorplan			
	5	6	7	8
Ashual and Wolf (2019)	7.5 ± 0.0	9.2 ± 0.0	10.0 ± 0.0	11.8 ± 0.0
Johnson et al. (2018)	7.7 ± 0.0	6.5 ± 0.0	10.2 ± 0.0	11.3 ± 0.1
House-GAN (Nauata et al., 2020a)	2.5 ± 0.1	2.4 ± 0.1	3.2 ± 0.0	5.3 ± 0.0
House-GAN++ (Nauata et al., 2021a)	1.9 ± 0.3	2.2 ± 0.3	2.4 ± 0.3	3.9 ± 0.5
HouseDiffusion (Shabani et al., 2023)	1.5 ± 0.0	1.2 ± 0.0	1.7 ± 0.0	2.5 ± 0.0
Our BD Model w/ BD prompt	0.46 ± 0.73	0.79 ± 0.98	1.27 ± 1.28	2.50 ± 1.97

Table 6: Comparison of Compatibility \downarrow score (Graph Edit Distance) on the RPLAN dataset.

variance in the compatibility metrics compared to existing image-based approaches. Table 7 suggests our model has a harder time following the input bubble diagram as the number of room increases, or when given room area information. This is reflected in Figure 2 where the second row’s bubble diagrams show more variations. This limitation might be tackled by using a better data structure to represent bubble diagrams, and perhaps also to be included as part of the floorplan data representation. This way, the bubble diagram will simply be part of the constraint set, and this may help the model understand the input better. This can be explored in future work in the same task setting. Additionally, in the future, our data can be expanded to account for more floorplan attributes such as doors and walls and our benchmark can include additional metrics to account for more obscure numerical checks.

Model	Generation	Number of Rooms				
		3	4	5	6	7
F+BD	Specific	0.15 ± 0.41	0.33 ± 0.58	1.35 ± 1.34	2.24 ± 1.76	4.23 ± 1.89
	All Room Area	5.91 ± 0.29	9.75 ± 0.54	13.79 ± 1.15	18.09 ± 1.71	22.35 ± 2.24
	Partial Room Area	3.80 ± 1.20	5.73 ± 2.03	8.19 ± 2.72	10.15 ± 2.88	11.42 ± 4.52
	Total Area	0.19 ± 0.40	0.62 ± 0.78	3.48 ± 1.83	3.41 ± 1.50	6.23 ± 2.20
M+BD	Specific	0.16 ± 0.45	0.47 ± 0.67	2.06 ± 1.33	2.85 ± 1.79	5.23 ± 2.14
	All Room Area	5.91 ± 0.29	9.62 ± 0.56	12.90 ± 1.27	17.48 ± 1.82	22.13 ± 1.94
	Partial Room Area	3.85 ± 1.11	6.35 ± 1.95	8.44 ± 2.48	11.56 ± 3.42	14.13 ± 3.90
	Total Area	0.20 ± 0.40	0.54 ± 0.73	2.40 ± 1.03	3.09 ± 2.07	6.26 ± 2.03
PM+BD	Specific	0.15 ± 0.39	0.58 ± 0.96	2.42 ± 1.33	2.97 ± 1.62	5.97 ± 1.25
	All Room Area	5.88 ± 0.33	9.77 ± 0.49	13.17 ± 1.23	18.56 ± 1.58	21.97 ± 2.36
	Partial Room Area	3.87 ± 1.13	6.42 ± 1.97	8.99 ± 2.67	12.06 ± 3.05	13.68 ± 3.71
	Total Area	0.34 ± 0.93	0.46 ± 0.84	2.08 ± 1.29	2.65 ± 1.87	5.84 ± 1.61

Table 7: Compatibility \downarrow score on ProcTHOR-trained models.

Analysis We further present few-shot results with GPT4o-mini. This is done by prompting the LLM with multiple example specifications along with output floorplan Jsons. The results are shown in Table 8. The generated floorplan has less overlapping polygons compared to the ones generated by the fine-tuned LLAMA3 models, however, they are less accurate in following the size, height and width specifications.

Analysis To encourage generation of realistic floorplans, we further add windows, doors, and open concept rooms to our data structure. The evaluation results are shown in Table 9, and some generation results are shown in Figure 3. We can see from the figure that the numerical constraints of the original ground truth floor plan is well respected. Every room has the same total area as the original.

8 Conclusions

We have motivated the need for new datasets and benchmarks for real-world use case scenarios for floorplan generation. We have developed and explored a Llama 3 based LLM for the problem and it yields SOTA results on the previous compatibility based evaluation and along with our proposed metrics, it highlights different use cases scenarios where improvements are possible. In particular we see from Table 7 that when

	Self-consistency Benchmark						
	P. Area [‡]	Total Area [‡]	P. Overlap [‡]	R. Count [‡]	ID [‡]	R. H [‡]	R. W [‡]
Few-Shot _S	0.99±0.00	1.00±0.00	0.04±0.20	1.00	1.00	0.93±0.02	0.93±0.01
Few-Shot _{AR}	0.99±0.01	0.96±0.07	0.01±0.11	1.00	1.00	0.88±0.07	0.93±0.00
Few-Shot _{PR}	1.00±0.00	0.94±0.12	0.00±0.04	0.70	1.00	0.86±0.07	0.93±0.00
Few-Shot _{TA}	1.00±0.00	1.00±0.00	0.00±0.00	1.00	1.00	0.93±0.00	0.93±0.00

	Prompt-consistency Benchmark						
	Num. R [♣]	Total Area [♣]	R. Area [♡]	ID [♣]	IDvsType [◊]	R. H [‡]	R. W [‡]
Few-Shot _S	0.82±0.39	1.00±0.00	0.99±0.00	0.82±0.39	1.00±0.00	0.93±0.02	0.93±0.01
Few-Shot _{AR}	0.95±0.22	-	0.99±0.01	0.95±0.22	1.00±0.00	-	-
Few-Shot _{PR}	0.98±0.14	1.00±0.00	1.00±0.00	0.98±0.14	1.00±0.00	-	-
Few-Shot _{TA}	1.00±0.03	1.00±0.00	-	-	-	-	-

Table 8: Our self- and prompt-consistency on few-shot results

	Self-consistency Benchmark						
	P. Area [‡]	Total Area [‡]	P. Overlap [‡]	R. Count [‡]	ID [‡]	R. H [‡]	R. W [‡]
Window+Doors	0.99±0.01	1.00±0.02	0.23±0.42	1.00	1.00	0.93±0.00	0.93±0.01
Window+Door _{AR}	0.99±0.00	0.99±0.03	0.18±0.39	1.00	1.00	0.93±0.00	0.93±0.00
Window+Door _{PR}	0.99±0.00	0.98±0.05	0.16±0.37	0.99	1.00	0.93±0.00	0.93±0.00
Window+Door _{TA}	0.99±0.00	0.99±0.03	0.16±0.37	1.00	1.00	0.93±0.00	0.93±0.00

	Prompt-consistency Benchmark						
	Num. R [♣]	Total Area [♣]	R. Area [♡]	ID [♣]	IDvsType [◊]	R. H [‡]	R. W [‡]
Window+Doors	1.00±0.00	0.99±0.00	0.99±0.00	1.00±0.02	1.00±0.00	0.93±0.01	0.93±0.01
Window+Door _{AR}	1.00±0.00	-	0.99±0.00	1.00±0.02	1.00±0.00	-	-
Window+Door _{PR}	1.00±0.00	0.99±0.00	0.99±0.00	1.00±0.02	1.00±0.00	-	-
Window+Door _{TA}	1.00±0.00	0.99±0.00	-	-	-	-	-

Table 9: Our windows and doors augmented floorplan results

conditioning on bubble diagrams and room area information this model struggles to respect the bubble diagram constraints. From Table 5 we see that the model does well, but is far from perfect at generating rooms with correct areas as computed from the generated polygons. We also see that the biggest weakness of this otherwise SOTA model is linked to issues with generating overlapping rooms.

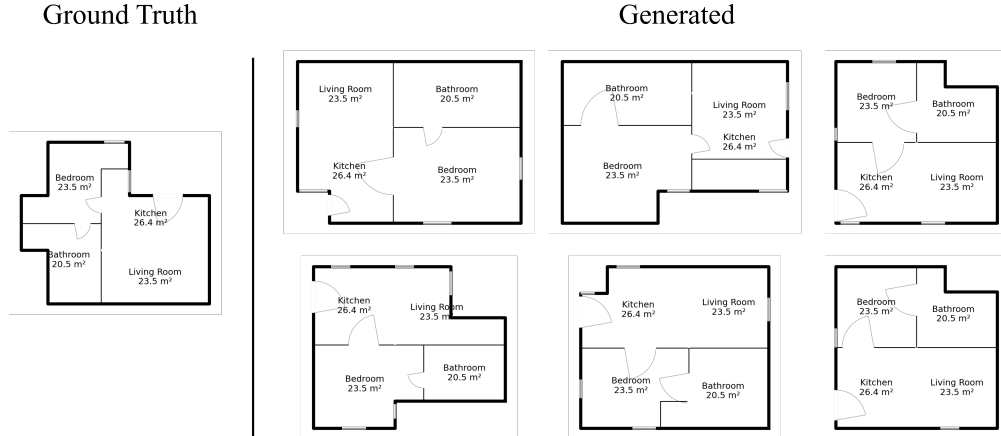


Figure 3: Generated realistic floor plans with windows, doors and open concept rooms.

We hope that our data, simulated data generation procedure and benchmarks will stimulate further developments for this new problem formulation for floorplan generation.

References

- Mohammad Amin Shabani, Sepidehsadat Hosseini, and Yasutaka Furukawa. Housediffusion: Vector floorplan generation via a diffusion model with discrete and continuous denoising, 2022.
- Wenming Wu, Xiao-Ming Fu, Rui Tang, Yuhang Wang, Yu-Hao Qi, and Ligang Liu. Data-driven interior plan generation for residential buildings. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 38(6), 2019.
- National Institute of Informatics. Lifull homes dataset, 2015.
- 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.
- Theodoros Galanos, Antonios Liapis, and Georgios N Yannakakis. Architext: Language-driven generative architecture design. *arXiv preprint arXiv:2303.07519*, 2023.
- Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Jordi Salvador, Kiana Ehsani, Winson Han, Eric Kolve, Ali Farhadi, Aniruddha Kembhavi, and Roozbeh Mottaghi. ProcTHOR: Large-Scale Embodied AI Using Procedural Generation. In *NeurIPS*, 2022. Outstanding Paper Award.
- 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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 162–177. Springer, 2020a.
- Nelson Nauata, Sepidehsadat Hosseini, Kai-Hung Chang, Hang Chu, Chin-Yi Cheng, and Yasutaka Furukawa. House-gan++: Generative adversarial layout refinement network towards intelligent computational agent for professional architects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13632–13641, 2021a.
- Ziniu Luo and Weixin Huang. Floorplangan: Vector residential floorplan adversarial generation. *Automation in Construction*, 142:104470, 2022. ISSN 0926-5805. doi: <https://doi.org/10.1016/j.autcon.2022.104470>. URL <https://www.sciencedirect.com/science/article/pii/S0926580522003430>.
- Mohammad Amin Shabani, Sepidehsadat Hosseini, and Yasutaka Furukawa. Housediffusion: Vector floorplan generation via a diffusion model with discrete and continuous denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5466–5475, 2023.
- Zehao Wen, Zichen Liu, Srinath Sridhar, and Rao Fu. Anyhome: Open-vocabulary generation of structured and textured 3d homes. *arXiv preprint arXiv:2312.06644*, 2023.
- Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, et al. Holodeck: Language guided generation of 3d embodied ai environments. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2024)*, volume 30, pages 20–25. IEEE/CVF, 2024.
- 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, 2020b.
- Nelson Nauata, Sepidehsadat Hosseini, Kai-Hung Chang, Hang Chu, Chin-Yi Cheng, and Yasutaka Furukawa. House-gan++: Generative adversarial layout refinement networks, 2021b.
- Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs, 2018.
- Oron Ashual and Lior Wolf. Specifying object attributes and relations in interactive scene generation, 2019.

Zeina Abu-Aisheh, Romain Raveaux, Jean-Yves Ramel, and Patrick Martineau. An exact graph edit distance algorithm for solving pattern recognition problems. In *4th International Conference on Pattern Recognition Applications and Methods 2015*, 2015.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021.

Huggingface generation. https://huggingface.co/docs/transformers/v4.40.2/en/main_classes/text_generation#transformers.GenerationMixin.generate.

A Dataset License and Privacy

The **ProcTHOR-10k** dataset (Deitke et al., 2022) is procedurally generated under Apache License 2.0 and distributed with the Allen Institute for AI’s prior package.

The **RPLAN** dataset Wu et al. (2019) is created with anonymized floorplan data that eliminates user and privacy information. Furthermore, the original paper states that the floorplans used are ensured to have no copyright issues. Besides, researchers are not permitted to redistribute the downloaded data, in whole or in part, through other media.

Due to redistribution restrictions of the **RPLAN**, we cannot directly share our new dataset. However, after downloading the **RPLAN** from the author’s source, the code provided in the supplementary materials can be used to create our new dataset.

B Possible Negative Impacts

Job Displacement The automation of floorplan generation through advanced technology threatens job displacement for professionals in architecture, real estate, and related fields. As these tools become more sophisticated, human expertise may be undervalued or rendered obsolete.

Mitigation To address this issue, it is important to promote technology as a means to augment human capabilities rather than replace them. By providing training programs, professionals can adapt to new tools and workflows, ensuring they remain relevant and can leverage technology to enhance their work rather than be supplanted by it.

Quality and Safety There is a significant risk that generated floorplans might not meet established safety standards or quality expectations, leading to potential hazards in construction or living conditions. Poorly designed floorplans could result in unsafe buildings and legal liabilities and reduce occupants’ overall quality of life.

Mitigation To mitigate these risks, it is essential to incorporate rigorous safety and quality checks into the floorplan generation process. Collaborating with experts in architecture and engineering can help ensure that generated floorplans adhere to all relevant standards and regulations, thus maintaining the integrity and safety of the built environment.

C Prompt Structure

```
<|start_header_id|>system<|end_header_id|>
you are to generate a floor plan in a JSON structure where each room
is defined by polygon vertices, make sure to not overlap the polygons.
you have to satisfy the adjacency constraints given as pairs of
neighboring rooms; two connecting rooms, room1 and room2, are presented
as (room1_type/"room1_id", room2_type/"room2_id"). you have to also
match the specifications passed by the user in a JSON structure when
they exist. when room area and total area requirements exist, make sure
the polygon areas add up to the required number.
<|eot_id|><|start_header_id|>user<|end_header_id|>
adjacency constraints: (Bedroom/"room|4", Bathroom/"room|5"),
(Bedroom/"room|4", Kitchen/"room|6"), (Bedroom/"room|4",
LivingRoom/"room|7"), (Bathroom/"room|5", Kitchen/"room|6"),
(Bathroom/"room|5", LivingRoom/"room|7"),
(Kitchen/"room|6", LivingRoom/"room|7"). specifications:
{'room_count\': 4, \'total_area\': 146.8, \'rooms\': [
{'area\': 41.3, \'id\': \'room|4\', \'room_type\': \'Bedroom\'},
```

```
{\`area\`: 27.5, \`id\`: \`room|7\`, \`room_type\`: \`LivingRoom\`}}}  
<|eot_id|><|start_header_id|>assistant<|end_header_id|>
```

D Training Details

D.1 Dataset Details

For floorplan data converted from **RPLAN** dataset, we follow the original data split. RPLAN dataset is divided into 5 parts according to the number of rooms present in the floorplan. Specifically, each part respectively contains floorplans with 4, 5, 6, 7, 8 rooms. Each part is further divided into a training, validation, and a test split. Following prior work, we use only training data with 5, 6, 7, 8 rooms. 4 different models are trained, each has never seen one of the 4 part – when training the model for 5 room floorplan generation, training is done with training splits of 6, 7, 8 room floorplans etc. During test time, the model is asked to generate only 4 room floorplans even though it has never seen a 4 room floorplan in the training set. This is in line with prior work that uses the RPLAN dataset.

For floorplan data converted from **ProcTHOR** dataset, the validation and test split is randomly chosen, with each of validation and test split being 10% of the training split. The training split is used to train the model; the validation split used to perform early stopping; and the test split is used to generate prompts for generation and evaluation.

D.2 Hyper-parameters

We follow the suggest training parameters from Llama-recipes (<https://github.com/meta-llama/llama-recipes/tree/main>), and we do not perform any hyper-parameter search.

Random Seed used for generation is chosen to be 12345 to select 100 (sampling generation) or 1000 (greedy generation) random test set data for generation. This seed is simply chosen and never changed or compared with any other random seed.

For **Sampling** generation, we generate 20 output floorplans for each prompt using nucleus sampling. We choose $p = 0.8$ for sampling as it is a common value for this parameter.

For **Greedy** generation, we use the default parameter provided by HuggingFace hug.

E Bubble Diagram Conversion

We follow RPLAN’s definition of room adjacency when creating bubble diagrams. Every room is first outlined by a bounding box defined by the room’s highest and lowest xy coordinates. Then, two rooms are considered adjacent if their corresponding bounding boxes are within 9 pixels at the closet point. Note that each floorplan pixel in the RPLAN dataset translates to 18/256 meters. We round this threshold to the nearest pixel in order to create bubble diagrams for the floorplans in ProcTHOR to ensure maximum comparability.

F Specific Attributes Sets in Preset Masking

Here we present the attributes used in the preset masking variant of our model. The four presets have varying amount of attributes from general to specific.

G Self and Prompt Consistency Metrics on ProcTHOR

In total, we have 6 model variants trained on converted ProcTHOR data. 3 of them are trained with bubble diagram input and 3 without. For each of the 6 model variants, we run 4 different sets of generations with different prompts as explained in Section 6.2; and for each of the generation prompt, we run both greedy and sampling algorithm to obtain 2 sets of different floorplan generations.

Preset	Constraint Attributes
1	room_count, room_types, total_area
2	room_count, room_types, total_area, partial list of (rooms[id], rooms[type], rooms[area])
3	room_count, room_types, partial list of (rooms[id], rooms[type], rooms[area])
4	room_count, room_types, full list of (rooms[id], rooms[type], rooms[area])

Table 10: Different presets of constraint set. Partial list represents information for only some of the rooms: dropout is applied on the individual room in the list of all rooms

Therefore, in total, we run 48 sets of experiments and evaluations. We present the full evaluation using our designed metrics here, organized into 4 different sections according to the generation prompt. Each section has 2 main tables, one for self consistency metrics, and one for prompt consistency metrics.

G.1 Generation Prompt: Specific

Greedy Generation							
Self Consistency Metrics							
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.94±0.07	0.16±0.37	1.00	0.99±0.02	1.00	0.99±0.02	1.00±0.03
F+BD	0.95±0.06	0.22±0.41	1.00	0.99±0.02	1.00	0.99±0.04	1.00±0.00
M	0.93±0.07	0.18±0.39	1.00	0.99±0.04	1.00	0.98±0.06	1.00±0.00
M+BD	0.93±0.07	0.19±0.39	1.00	0.99±0.04	1.00	0.99±0.03	1.00±0.03
PM	0.92±0.06	0.12±0.32	0.98	0.98±0.07	1.00	1.00±0.01	1.00±0.02
PM+BD	0.90±0.09	0.17±0.37	1.00	1.00±0.03	1.00	1.00±0.01	1.00±0.00

Sampling Generation							
Self Consistency Metrics							
	P. Area	Overlap	R. Count	R. H	RID	R. W	Total Area
F	0.95±0.06	0.17±0.38	1.00	0.99±0.02	1.00	0.99±0.02	1.00±0.00
F+BD	0.95±0.06	0.22±0.41	1.00	0.99±0.03	1.00	0.99±0.03	1.00±0.00
M	0.93±0.08	0.16±0.37	1.00	0.99±0.04	1.00	0.98±0.07	1.00±0.00
M+BD	0.93±0.07	0.17±0.38	1.00	0.99±0.04	1.00	0.99±0.04	1.00±0.02
PM	0.91±0.07	0.14±0.35	0.97	0.96±0.09	1.00	1.00±0.02	1.00±0.02
PM+BD	0.90±0.08	0.20±0.40	1.00	0.99±0.06	1.00	1.00±0.00	1.00±0.00

Table 11: Evaluation of **self consistency metrics** on generated results from **Specific Prompt** (Section 6.2). Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

Greedy Generation								
Prompt Consistency Metrics								
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	0.95±0.07	0.94±0.07	0.99±0.02	1.00±0.03	1.00±0.00	1.00±0.00	0.99±0.02
F+BD	1.00±0.00	0.96±0.05	0.95±0.06	0.99±0.03	1.00±0.00	1.00±0.00	1.00±0.00	0.99±0.04
M	1.00±0.00	0.94±0.07	0.93±0.07	0.97±0.07	1.00±0.00	1.00±0.00	1.00±0.00	0.97±0.08
M+BD	1.00±0.00	0.95±0.07	0.93±0.07	0.97±0.07	1.00±0.03	1.00±0.00	1.00±0.00	0.97±0.08
PM	1.00±0.00	0.74±0.39	0.92±0.06	0.80±0.17	1.00±0.00	1.00±0.00	1.00±0.00	0.80±0.19
PM+BD	1.00±0.00	0.86±0.25	0.90±0.09	0.83±0.19	1.00±0.00	1.00±0.00	1.00±0.00	0.81±0.19

Sampling Generation								
Prompt Consistency Metrics								
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	0.95±0.06	0.95±0.06	0.99±0.03	1.00±0.00	1.00±0.00	1.00±0.00	0.99±0.02
F+BD	1.00±0.00	0.96±0.06	0.95±0.06	0.99±0.03	1.00±0.00	1.00±0.00	1.00±0.00	0.99±0.03
M	1.00±0.00	0.95±0.07	0.93±0.08	0.97±0.07	1.00±0.00	1.00±0.00	1.00±0.00	0.96±0.09
M+BD	1.00±0.00	0.95±0.06	0.93±0.07	0.96±0.07	1.00±0.02	1.00±0.00	1.00±0.00	0.97±0.08
PM	1.00±0.00	0.75±0.37	0.91±0.07	0.78±0.17	1.00±0.01	1.00±0.00	1.00±0.00	0.77±0.20
PM+BD	1.00±0.00	0.86±0.26	0.90±0.08	0.79±0.19	1.00±0.00	1.00±0.00	1.00±0.00	0.79±0.19

Table 12: Evaluation of **prompt consistency metrics** on generated results from **Specific Prompt**. Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

G.2 Generation Prompt: Total Area

Greedy Generation							
Self Consistency Metrics							
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.79±0.13	0.07±0.26	0.96	1.00±0.01	1.00	1.00±0.01	0.97±0.07
F+BD	0.83±0.09	0.19±0.39	0.97	1.00±0.02	1.00	0.99±0.03	0.97±0.09
M	0.92±0.06	0.09±0.28	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.02
M+BD	0.91±0.06	0.14±0.35	1.00	1.00±0.01	1.00	1.00±0.01	1.00±0.02
PM	0.93±0.05	0.14±0.35	0.95	1.00±0.01	1.00	1.00±0.01	0.99±0.04
PM+BD	0.90±0.07	0.13±0.34	0.98	1.00±0.01	1.00	1.00±0.01	0.98±0.12

Sampling Generation							
Self Consistency Metrics							
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.78±0.13	0.14±0.34	0.96	0.99±0.02	1.00	0.99±0.02	0.98±0.06
F+BD	0.81±0.10	0.26±0.44	0.96	0.99±0.02	1.00	0.99±0.02	0.96±0.15
M	0.90±0.06	0.14±0.35	1.00	1.00±0.01	1.00	1.00±0.01	1.00±0.03
M+BD	0.90±0.06	0.15±0.35	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.03
PM	0.92±0.06	0.18±0.38	0.97	1.00±0.01	1.00	1.00±0.01	0.99±0.03
PM+BD	0.88±0.07	0.15±0.35	0.96	1.00±0.01	1.00	1.00±0.01	0.96±0.18

Table 13: Evaluation of **self consistency metrics** on generated results from **Specific Prompt** (Section 6.2). Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

Greedy Generation								
Prompt Consistency Metrics								
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	0.83±0.14	-	-	-	-	-	-
F+BD	1.00±0.00	0.87±0.12	-	-	-	-	-	-
M	1.00±0.00	0.94±0.06	-	-	-	-	-	-
M+BD	1.00±0.00	0.94±0.05	-	-	-	-	-	-
PM	1.00±0.00	0.95±0.06	-	-	-	-	-	-
PM+BD	1.00±0.00	0.91±0.13	-	-	-	-	-	-

Sampling Generation								
Prompt Consistency Metrics								
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	0.82±0.15	-	-	-	-	-	-
F+BD	1.00±0.00	0.84±0.16	-	-	-	-	-	-
M	1.00±0.00	0.93±0.06	-	-	-	-	-	-
M+BD	1.00±0.00	0.93±0.06	-	-	-	-	-	-
PM	1.00±0.00	0.94±0.05	-	-	-	-	-	-
PM+BD	1.00±0.00	0.88±0.18	-	-	-	-	-	-

Table 14: Evaluation of **prompt consistency metrics** on generated results from **Total Area Prompt**. Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

G.3 Generation Prompt: Partial Room Area

	Greedy Generation						
	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.81±0.11	0.03±0.17	0.69	1.00±0.01	1.00	1.00±0.02	0.86±0.21
F+BD	0.79±0.10	0.17±0.37	0.76	1.00±0.02	1.00	0.99±0.03	0.92±0.15
M	0.90±0.07	0.11±0.31	0.94	0.99±0.02	1.00	1.00±0.01	0.97±0.10
M+BD	0.90±0.06	0.13±0.34	0.99	1.00±0.01	1.00	1.00±0.01	0.99±0.04
PM	0.92±0.06	0.15±0.35	0.96	1.00±0.01	1.00	1.00±0.01	0.97±0.07
PM+BD	0.90±0.07	0.14±0.34	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.03

	Sampling Generation						
	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.79±0.13	0.06±0.25	0.75	0.99±0.02	1.00	0.99±0.03	0.87±0.22
F+BD	0.78±0.11	0.23±0.42	0.81	0.99±0.02	1.00	0.99±0.02	0.94±0.13
M	0.89±0.07	0.14±0.34	0.94	1.00±0.01	1.00	1.00±0.02	0.97±0.11
M+BD	0.89±0.07	0.14±0.35	0.99	1.00±0.01	1.00	1.00±0.01	0.99±0.03
PM	0.91±0.07	0.18±0.39	0.98	1.00±0.01	1.00	1.00±0.01	0.98±0.04
PM+BD	0.88±0.07	0.17±0.38	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.04

Table 15: Evaluation of **self consistency metrics** on generated results from **Specific Prompt** (Section 6.2). Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

	Greedy Generation							
	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	0.48±0.40	0.82±0.14	-	0.99±0.09	1.00±0.02	0.99±0.07	-
F+BD	1.00±0.00	0.74±0.21	0.79±0.14	-	0.97±0.15	1.00±0.00	0.99±0.06	-
M	1.00±0.00	0.91±0.10	0.80±0.35	-	0.99±0.07	0.89±0.27	0.99±0.09	-
M+BD	1.00±0.00	0.93±0.06	0.87±0.13	-	1.00±0.05	0.99±0.11	1.00±0.00	-
PM	1.00±0.00	0.93±0.07	0.91±0.11	-	0.99±0.05	1.00±0.00	1.00±0.04	-
PM+BD	1.00±0.00	0.92±0.07	0.89±0.10	-	1.00±0.02	1.00±0.00	1.00±0.01	-

	Sampling Generation							
	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.02	0.50±0.39	0.80±0.15	-	0.99±0.07	1.00±0.00	1.00±0.06	-
F+BD	1.00±0.00	0.73±0.25	0.75±0.23	-	0.96±0.16	1.00±0.02	0.99±0.07	-
M	1.00±0.00	0.90±0.12	0.80±0.26	-	0.99±0.07	0.88±0.29	0.99±0.07	-
M+BD	1.00±0.00	0.93±0.07	0.87±0.14	-	0.99±0.06	0.98±0.15	1.00±0.00	-
PM	1.00±0.00	0.94±0.06	0.90±0.11	-	1.00±0.03	1.00±0.00	1.00±0.01	-
PM+BD	1.00±0.00	0.90±0.08	0.87±0.13	-	1.00±0.02	1.00±0.00	1.00±0.01	-

Table 16: Evaluation of **prompt consistency metrics** on generated results from **Partial Room Area Prompt**. Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

G.4 Generation Prompt: All Room Area

	Greedy Generation						
	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.82±0.10	0.16±0.37	1.00	1.00±0.02	1.00	1.00±0.01	0.93±0.17
F+BD	0.76±0.11	0.25±0.43	1.00	0.99±0.02	1.00	0.99±0.02	0.90±0.20
M	0.90±0.07	0.12±0.33	1.00	0.99±0.02	1.00	0.99±0.01	0.99±0.03
M+BD	0.91±0.06	0.14±0.35	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.02
PM	0.91±0.06	0.14±0.35	1.00	1.00±0.01	1.00	1.00±0.01	0.97±0.09
PM+BD	0.90±0.07	0.19±0.39	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.05

	Sampling Generation						
	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
F	0.80±0.10	0.18±0.39	0.99	0.99±0.02	1.00	0.99±0.04	0.95±0.14
F+BD	0.76±0.11	0.31±0.46	1.00	0.99±0.02	1.00	0.99±0.03	0.93±0.15
M	0.89±0.07	0.17±0.37	1.00	1.00±0.01	1.00	0.99±0.02	0.99±0.03
M+BD	0.90±0.07	0.15±0.36	1.00	1.00±0.01	1.00	1.00±0.01	0.99±0.04
PM	0.91±0.07	0.18±0.38	1.00	1.00±0.01	1.00	1.00±0.01	0.96±0.10
PM+BD	0.88±0.07	0.19±0.40	1.00	1.00±0.01	1.00	1.00±0.01	1.00±0.00

Table 17: Evaluation of **self consistency metrics** on generated results from **Specific Prompt** (Section 6.2). Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

	Greedy Generation							
	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	-	0.81±0.10	-	1.00±0.06	1.00±0.00	1.00±0.06	-
F+BD	1.00±0.00	-	0.76±0.11	-	1.00±0.00	1.00±0.00	1.00±0.00	-
M	1.00±0.00	-	0.90±0.07	-	1.00±0.00	1.00±0.00	1.00±0.00	-
M+BD	1.00±0.00	-	0.91±0.06	-	1.00±0.01	1.00±0.00	1.00±0.00	-
PM	1.00±0.00	-	0.91±0.06	-	1.00±0.00	1.00±0.00	1.00±0.00	-
PM+BD	1.00±0.00	-	0.90±0.07	-	1.00±0.00	1.00±0.00	1.00±0.00	-

	Sampling Generation							
	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
F	1.00±0.00	-	0.79±0.10	-	0.99±0.09	1.00±0.00	0.99±0.09	-
F+BD	1.00±0.00	-	0.76±0.11	-	1.00±0.00	1.00±0.00	1.00±0.00	-
M	1.00±0.00	-	0.89±0.07	-	1.00±0.00	1.00±0.00	1.00±0.00	-
M+BD	1.00±0.00	-	0.90±0.07	-	1.00±0.04	1.00±0.00	1.00±0.00	-
PM	1.00±0.00	-	0.91±0.07	-	1.00±0.01	1.00±0.00	1.00±0.01	-
PM+BD	1.00±0.00	-	0.88±0.07	-	1.00±0.02	1.00±0.00	1.00±0.00	-

Table 18: Evaluation of **prompt consistency metrics** on generated results from **All Room Area Prompt**. Top half are results on single greedy generations from 1000 prompts from 1000 different test set floorplans; bottom half are results on sampled generations on 100 different prompt with 20 generations from each prompt.

H Compatibility Metrics on ProcTHOR

We further run compatibility evaluations on all of the Bubble Diagram model variants to test how well bubble diagrams are followed. The results are presented here:

Model	Generation	Number of Rooms				
Variant	Prompt	3	4	5	6	7
F+BD	Specific	0.15 \pm 0.41	0.33 \pm 0.58	1.35 \pm 1.34	2.24 \pm 1.76	4.23 \pm 1.89
	Total Area	0.19 \pm 0.40	0.62 \pm 0.78	3.48 \pm 1.83	3.41 \pm 1.50	6.23 \pm 2.20
	Partial Room Area	3.80 \pm 1.20	5.73 \pm 2.03	8.19 \pm 2.72	10.15 \pm 2.88	11.42 \pm 4.52
	All Room Area	5.91 \pm 0.29	9.75 \pm 0.54	13.79 \pm 1.15	18.09 \pm 1.71	22.35 \pm 2.24
M+BD	Specific	0.16 \pm 0.45	0.47 \pm 0.67	2.06 \pm 1.33	2.85 \pm 1.79	5.23 \pm 2.14
	Total Area	0.20 \pm 0.40	0.54 \pm 0.73	2.40 \pm 1.03	3.09 \pm 2.07	6.26 \pm 2.03
	Partial Room Area	3.85 \pm 1.11	6.35 \pm 1.95	8.44 \pm 2.48	11.56 \pm 3.42	14.13 \pm 3.90
	All Room Area	5.91 \pm 0.29	9.62 \pm 0.56	12.90 \pm 1.27	17.48 \pm 1.82	22.13 \pm 1.94
PM+BD	Specific	0.15 \pm 0.39	0.58 \pm 0.96	2.42 \pm 1.33	2.97 \pm 1.62	5.97 \pm 1.25
	Total Area	0.34 \pm 0.93	0.46 \pm 0.84	2.08 \pm 1.29	2.65 \pm 1.87	5.84 \pm 1.61
	Partial Room Area	3.87 \pm 1.13	6.42 \pm 1.97	8.99 \pm 2.67	12.06 \pm 3.05	13.68 \pm 3.71
	All Room Area	5.88 \pm 0.33	9.77 \pm 0.49	13.17 \pm 1.23	18.56 \pm 1.58	21.97 \pm 2.36

Table 19: Compatibility (\downarrow) score on ProcTHOR experiments. The best score in each category (by number of rooms in a floor plan) is bolded, and the 2nd best score has an asterisk.

I Self and Prompt Consistency Metrics on RPLAN

This section contains the full self- and prompt-consistency metrics evaluated on the 4 models (Section 6.1) trained on the converted RPLAN dataset.

Similar to the ProcTHOR-trained model experiments, we run the 4 different generation prompts on each of the RPLAN-trained model variants, resulting in 16 different experiments. They are organized by generation prompt and shown here.

I.1 Generation Prompt: Specific

	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
5-Room	0.89 \pm 0.31	0.34 \pm 0.47	0.40	1.00 \pm 0.01	1.00	0.98 \pm 0.03	0.92 \pm 0.08
6-Room	0.93 \pm 0.05	0.36 \pm 0.48	0.99	0.99 \pm 0.01	1.00	0.97 \pm 0.02	0.98 \pm 0.06
7-Room	0.93 \pm 0.04	0.42 \pm 0.49	0.97	0.99 \pm 0.01	1.00	0.98 \pm 0.02	0.98 \pm 0.05
8-Room	0.93 \pm 0.07	0.62 \pm 0.49	0.89	0.99 \pm 0.01	1.00	0.98 \pm 0.02	0.95 \pm 0.08

Table 20: **Self consistency** metrics on RPLAN generations using **Specific** generation prompts.

	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
5-Room	1.00 \pm 0.00	0.92 \pm 0.06	0.93 \pm 0.09	0.94 \pm 0.11	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	0.95 \pm 0.09
6-Room	1.00 \pm 0.00	0.94 \pm 0.07	0.94 \pm 0.06	0.92 \pm 0.23	1.00 \pm 0.05	1.00 \pm 0.00	1.00 \pm 0.00	0.94 \pm 0.14
7-Room	1.00 \pm 0.00	0.94 \pm 0.06	0.94 \pm 0.06	0.93 \pm 0.17	1.00 \pm 0.02	1.00 \pm 0.00	1.00 \pm 0.00	0.94 \pm 0.29
8-Room	1.00 \pm 0.00	0.92 \pm 0.08	0.94 \pm 0.07	0.95 \pm 0.10	0.99 \pm 0.06	1.00 \pm 0.00	1.00 \pm 0.00	0.96 \pm 0.09

Table 21: **Prompt consistency** metrics on RPLAN generations using **Specific** generation prompts.

I.2 Generation Prompt: Total Area

	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
5-Room	0.89±0.41	0.31±0.46	0.19	1.00±0.01	1.00	0.98±0.03	0.95±0.06
6-Room	0.93±0.05	0.32±0.47	1.00	0.99±0.01	1.00	0.98±0.02	0.99±0.03
7-Room	0.94±0.03	0.38±0.48	1.00	0.99±0.01	1.00	0.98±0.02	0.99±0.02
8-Room	0.82±3.50	0.53±0.50	0.89	0.99±0.01	1.00	0.98±0.02	0.97±0.07

Table 22: **Self consistency** metrics on RPLAN generations using **Total Area** generation prompts.

	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
5-Room	-	0.92±0.07	-	-	-	-	-	-
6-Room	-	0.94±0.06	-	-	-	-	-	-
7-Room	-	0.94±0.06	-	-	-	-	-	-
8-Room	-	0.93±0.07	-	-	-	-	-	-

Table 23: **Prompt consistency** metrics on RPLAN generations using **Total Area** generation prompts.

I.3 Generation Prompt: Partial Room Area

	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
5-Room	0.77±1.93	0.25±0.44	0.49	1.00±0.01	1.00	0.97±0.10	0.95±0.08
6-Room	0.93±0.07	0.34±0.47	0.98	0.99±0.04	1.00	0.97±0.02	0.97±0.08
7-Room	0.93±0.04	0.40±0.49	0.96	0.99±0.02	1.00	0.98±0.02	0.96±0.07
8-Room	0.91±0.34	0.54±0.50	0.73	0.99±0.02	1.00	0.97±0.02	0.93±0.11

Table 24: **Self consistency** metrics on RPLAN generations using **Partial Room Area** generation prompts.

	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
5-Room	-	0.92±0.08	0.91±0.14	-	1.00±0.05	1.00±0.04	1.00±0.00	-
6-Room	-	0.93±0.08	0.86±0.72	-	0.99±0.06	0.99±0.07	1.00±0.02	-
7-Room	-	0.91±0.09	0.88±0.19	-	1.00±0.03	0.99±0.06	1.00±0.01	-
8-Room	-	0.89±0.11	0.73±0.48	-	0.98±0.09	0.94±0.13	0.99±0.05	-

Table 25: **Prompt consistency** metrics on RPLAN generations using **Partial Room Area** generation prompts.

I.4 Generation Prompt: All Room Area

	Self Consistency Metrics						
	P. Area	Overlap	R. Count	R. H	ID	R. W	Total Area
5-Room	0.90±0.16	0.37±0.48	0.97	1.00±0.01	1.00	0.98±0.02	0.98±0.05
6-Room	0.93±0.06	0.31±0.46	1.00	0.99±0.01	1.00	0.97±0.03	0.95±0.04
7-Room	0.92±0.15	0.38±0.49	0.93	0.99±0.02	1.00	0.98±0.02	0.98±0.03
8-Room	0.93±0.10	0.61±0.49	0.96	0.99±0.02	1.00	0.98±0.02	0.91±0.07

Table 26: **Self consistency** metrics on RPLAN generations using **All Room Area** generation prompts.

	Prompt Consistency Metrics							
	Num. R.	Total Area	R. Area	R. H	ID	IDvsType	Type	R. W
5-Room	-	-	0.92±0.04	-	1.00±0.01	1.00±0.00	1.00±0.00	-
6-Room	-	-	0.93±0.05	-	1.00±0.02	1.00±0.00	1.00±0.00	-
7-Room	-	-	0.93±0.07	-	1.00±0.02	1.00±0.00	1.00±0.00	-
8-Room	-	-	0.93±0.05	-	0.99±0.04	1.00±0.00	1.00±0.00	-

Table 27: **Prompt consistency** metrics on RPLAN generations using **All Room Area** generation prompts.

J Model Usage

Input with full numerical constraints on all aspects of the floorplan:

```
<|start_header_id|>system<|end_header_id|>
you are to generate a floor plan in a JSON structure where each room
is defined by polygon vertices, make sure to not overlap the polygons.
you have to satisfy the adjacency constraints given as pairs of
neighboring rooms; two connecting rooms, room1 and room2, are presented
as (room1_type/"room1_id", room2_type/"room2_id"). you have to also
match the specifications passed by the user in a JSON structure when
they exist. when room area and total area requirements exist, make sure
the polygon areas add up to the required number.
<|eot_id|><|start_header_id|>user<|end_header_id|>
adjacency constraints:
(Bedroom/"room|4", Bathroom/"room|5"), (Bedroom/"room|4", Kitchen/"room|6"),
(Bedroom/"room|4",LivingRoom/"room|7"), (Bathroom/"room|5", Kitchen/"room|6"),
(Bathroom/"room|5", LivingRoom/"room|7"), (Kitchen/"room|6", LivingRoom/"room|7").
specifications:
{'room_count': 4, 'room_types': ['Bedroom', 'Bathroom', 'Kitchen', 'LivingRoom'],
'rooms': [{'area': 27.1, 'height': 5.5, 'id': 'room|4', 'is_regular': 0, 'room_type':
'Bedroom', 'width': 5.5}, {'area': 13.6, 'height': 3.7, 'id': 'room|5', 'is_regular':
1, 'room_type': 'Bathroom', 'width': 3.7}, {'area': 30.5, 'height': 5.5, 'id':
'room|6', 'is_regular': 1, 'room_type': 'Kitchen', 'width': 5.5}, {'area': 20.3,
'height': 5.5, 'id': 'room|7', 'is_regular': 1, 'room_type': 'LivingRoom', 'width':
3.7}], 'total_area': 91.5}]}
```

Model Output:

```
{"room_count": 4, "total_area": 91.5, "room_types": ["Bedroom", "Bathroom", "Kitchen",
"LivingRoom"], "rooms": [{"area": 27.1, "floor_polygon": [{"x": 0.0, "y": 0.0},
{"x": 0.0, "y": 5.5}, {"x": 5.5, "y": 5.5}, {"x": 5.5, "y": 1.8}, {"x": 3.7, "y": 1.8},
{"x": 3.7, "y": 0.0}], "height": 5.5, "id": "room|4", "is_regular": 0, "room_type":
"Bedroom", "width": 5.5}, {"area": 13.6, "floor_polygon": [{"x": 0.0, "y": 5.5},
{"x": 0.0, "y": 9.2}, {"x": 3.7, "y": 9.2}, {"x": 3.7, "y": 5.5}], "height": 3.7,
"id": "room|5", "is_regular": 1, "room_type": "Bathroom", "width": 3.7}, {"area": 30.5,
"floor_polygon": [{"x": 3.7, "y": 5.5}, {"x": 5.5, "y": 9.2}, {"x": 5.5, "y": 11.0}, {"x": 9.2, "y": 11.0}, {"x": 9.2, "y": 5.5}, {"x": 5.5, "y": 5.5}],
"height": 5.5, "id": "room|6", "is_regular": 0, "room_type": "Kitchen", "width": 5.5},
{"area": 20.3, "floor_polygon": [{"x": 5.5, "y": 1.8}, {"x": 5.5, "y": 5.5},
{"x": 9.2, "y": 5.5}, {"x": 9.2, "y": 0.0}, {"x": 3.7, "y": 0.0}, {"x": 3.7, "y": 1.8}],
"height": 5.5, "id": "room|7", "is_regular": 0, "room_type": "LivingRoom", "width": 5.5}],
"doors": [{"id": "door|4|5", "position": [{"x": 2.4, "y": 5.5}, {"x": 0.4, "y": 5.5}]},
{"id": "door|4|7", "position": [{"x": 5.5, "y": 3.5}, {"x": 5.5, "y": 2.5}]},
{"id": "door|1|6", "position": [{"x": 9.2, "y": 8.7}, {"x": 9.2, "y": 6.8}]}],
"windows": [{"id": "window|4|0", "position": [{"x": 5.2, "y": 1.8}, {"x": 4.0, "y": 1.8}]},
{"id": "window|7|0", "position": [{"x": 3.7, "y": 0.3}, {"x": 3.7, "y": 1.5}]},
{"id": "window|7|1", "position": [{"x": 9.2, "y": 2.5}, {"x": 9.2, "y": 1.6}]},
{"id": "window|6|0", "position": [{"x": 4.6, "y": 9.2}, {"x": 5.5, "y": 9.2}]}],
"open_walls": [{"id": "wall|6|5.53|5.53|9.22|5.53", "position": [{"x": 5.5, "y": 5.5},
{"x": 9.2, "y": 5.5}]}, {"id": "wall|7|5.53|5.53|9.22|5.53", "position":
[{"x": 5.5, "y": 5.5}, {"x": 9.2, "y": 5.5}]}]}
```

Rendering of the generated floorplan:




Figure 4: Rendering of the generated floorplan text

Input with partial numerical constraints:

```
<|start_header_id|>system<|end_header_id|>
you are to generate a floor plan in a JSON structure where each room
is defined by polygon vertices, make sure to not overlap the polygons.
you have to satisfy the adjacency constraints given as pairs of
neighboring rooms; two connecting rooms, room1 and room2, are presented
as (room1_type/"room1_id", room2_type/"room2_id"). you have to also
match the specifications passed by the user in a JSON structure when
they exist. when room area and total area requirements exist, make sure
the polygon areas add up to the required number.
<|eot_id|><|start_header_id|>user<|end_header_id|>
adjacency constraints:
(Bedroom/"room|4", Bathroom/"room|5"), (Bedroom/"room|4", Kitchen/"room|8"),
(Bedroom/"room|4", LivingRoom/"room|9"), (Bathroom/"room|5", Kitchen/"room|8"),
(Bathroom/"room|5", LivingRoom/"room|9"), (Bedroom/"room|6", Kitchen/"room|8"),
(Bedroom/"room|6", LivingRoom/"room|9"), (Kitchen/"room|8", LivingRoom/"room|9").
specifications:
{'room_count': 5, 'total_area': 102.4, 'rooms': [{'area': 12.8, 'id':
'room|5', 'room_type': 'Bathroom'}, {'area': 22.4, 'id': 'room|6', 'room_type':
'Bedroom'}, {'area': 19.2, 'id': 'room|9', 'room_type': 'LivingRoom'}]}
```

Model Output:

```
{
  "room_count": 5,
  "total_area": 102.4,
  "room_types": ["Bedroom", "Bathroom", "Bedroom", "Kitchen", "LivingRoom"],
  "rooms": [
    {
      "area": 19.2,
      "floor_polygon": [
        {"x": 0.0, "y": 0.0}, {"x": 0.0, "y": 3.6}, {"x": 5.4, "y": 3.6}, {"x": 5.4, "y": 0.0}
      ],
      "height": 3.6,
      "id": "room|4",
      "is_regular": 0,
      "room_type": "Bedroom",
      "width": 5.4,
      "area": 12.8,
      "floor_polygon": [
        {"x": 0.0, "y": 3.6}, {"x": 0.0, "y": 7.2}, {"x": 1.8, "y": 7.2}, {"x": 1.8, "y": 5.4}, {"x": 5.4, "y": 5.4}, {"x": 5.4, "y": 3.6}
      ],
      "height": 3.6,
      "id": "room|5",
      "is_regular": 0,
      "room_type": "Bathroom",
      "width": 5.4,
      "area": 22.4,
      "floor_polygon": [
        {"x": 9.0, "y": 1.8}, {"x": 9.0, "y": 7.2}, {"x": 14.6, "y": 7.2}, {"x": 14.6, "y": 3.6}, {"x": 10.8, "y": 3.6}, {"x": 10.8, "y": 1.8}
      ],
      "height": 5.4,
      "id": "room|6",
      "is_regular": 0,
      "room_type": "Bedroom",
      "width": 5.4,
      "area": 28.8,
      "floor_polygon": [
        {"x": 5.4, "y": 0.0}, {"x": 5.4, "y": 5.}, {"x": 9.0, "y": 5.4}, {"x": 9.0, "y": 1.8}, {"x": 10.8, "y": 1.8}, {"x": 10.8, "y": 0.0}
      ],
      "height": 5.4,
      "id": "room|8",
      "is_regular": 0,
      "room_type": "Kitchen",
      "width": 5.4,
      "area": 19.2,
      "floor_polygon": [
        {"x": 5.4, "y": 5.4}, {"x": 5.4, "y": 10.8}, {"x": 9.0, "y": 10.8}, {"x": 9.0, "y": 5.4}
      ],
      "height": 5.4,
      "id": "room|9",
      "is_regular": 1,
      "room_type": "LivingRoom",
      "width": 3.6,
      "doors": [
        {
          "id": "door|4|5",
          "position": [
            {"x": 4.1, "y": 3.6}, {"x": 3.0, "y": 3.6}
          ]
        },
        {
          "id": "door|4|8",
          "position": [
            {"x": 5.4, "y": 1.6}, {"x": 5.4, "y": 3.6}
          ]
        },
        {
          "id": "door|6|8",
          "position": [
            {"x": 9.0, "y": 4.0}, {"x": 9.0, "y": 3.0}
          ]
        },
        {
          "id": "door|1|9",
          "position": [
            {"x": 6.1, "y": 10.8}, {"x": 7.1, "y": 10.8}
          ]
        }
      ],
      "windows": [
        {
          "id": "window|4|0",
          "position": [
            {"x": 5.2, "y": 0.0}, {"x": 3.6, "y": 0.0}
          ]
        },
        {
          "id": "window|8|0",
          "position": [
            {"x": 10.8, "y": 1.5}, {"x": 10.8, "y": 0.3}
          ]
        },
        {
          "id": "window|9|0",
          "position": [
            {"x": 9.0, "y": 9.5}, {"x": 9.0, "y": 8.0}
          ]
        }
      ],
      "open_walls": [
        {
          "id": "wall|8|5.39|5.39|9.02|5.39",
          "position": [
            {"x": 5.4, "y": 5.4}, {"x": 9.0, "y": 5.4}
          ]
        },
        {
          "id": "wall|9|5.39|5.39|9.02|5.39",
          "position": [
            {"x": 5.4, "y": 5.4}, {"x": 9.0, "y": 5.4}
          ]
        }
      ]
    }
  ]
}
```

Rendering of the generated floorplan:




Figure 5: Rendering of the generated floorplan text