

Tell2Design: A Dataset for Language-Guided Floor Plan Generation

Sicong Leng^{1,*}, Yang Zhou^{2,*†}, Mohammed Haroon Dupty³,
Wee Sun Lee³, Sam Conrad Joyce⁴, Wei Lu¹

¹StatNLP Research Group, Singapore University of Technology and Design

²Institute of High Performance Computing (IHPC), A*STAR Singapore

³School of Computing, National University of Singapore

⁴Meta Design Lab, Singapore University of Technology and Design

{sicong_leng, sam_joyce, luwei}@sutd.edu.sg

zhou_yang@ihpc.a-star.edu.sg {dmharoon, leews}@comp.nus.edu.sg

Abstract

We consider the task of generating designs directly from natural language descriptions, and consider floor plan generation as the initial research area. Language conditional generative models have recently been very successful in generating high-quality artistic images. However, designs must satisfy different constraints that are not present in generating artistic images, particularly spatial and relational constraints. We make multiple contributions to initiate research on this task. First, we introduce a novel dataset, *Tell2Design* (T2D), which contains more than 80k floor plan designs associated with natural language instructions. Second, we propose a Sequence-to-Sequence model that can serve as a strong baseline for future research. Third, we benchmark this task with several text-conditional image generation models. We conclude by conducting human evaluations on the generated samples and providing an analysis of human performance. We hope our contributions will propel the research on language-guided design generation forward¹.

1 Introduction

Recently, text-conditional generative AI models (Nichol et al., 2022; Saharia et al., 2022b; Ramesh et al., 2022; Dhariwal and Nichol, 2021; Ho et al., 2022) have demonstrated impressive results in generating high-fidelity images. Such models generally focus on understanding high-level visual concepts from sentence-level descriptions, and the generated images are valued for looking realistic and being creative, thereby being more suitable for generating artwork. However, besides less constrained generation like artworks, generating designs that meet various requirements specified

* Equal contribution † Most work done at NUS

¹Code and dataset are available at <https://github.com/LengSicong/Tell2Design>.

Language Instructions:

The *balcony*, located on the Eastern side of the unit, North of the bathroom, juts out from the rest of the unit. It is 4' wide from West to East, and the entire balcony is located Eastward of the Easternmost wall of the bedroom, common area, and bathroom. The balcony spans 11' from North to South, and its Northern wall is shared with the Southern wall of the kitchen. The kitchen and balcony both jut out approximately 4' to the East of the unit.

The *bathroom* is located on the central-eastern side of the unit between the common area and the master bedroom. While the Northern wall of the master bedroom and the Southern wall of the common area are both roughly 12' wide ...

Floor Plan:



Living Room Common Room Kitchen
Bath Room Master Room Balcony

Figure 1: A data sample from our *Tell2Design* dataset.

in natural languages is also much needed in practice (Stiny, 1980; Seneviratne et al., 2022; Zhang et al., 2022; Wei et al., 2022). In particular, a design process always involves interaction between users/clients, who define objectives, constraints, and requirements that should be met, and designers, who need to develop various solutions with domain-specific experiences and knowledge. For example, users may dictate their house design requirements in text and expect expert architects to perform the floor plan generation.

Previous research in layout generation aims to automate the process of layout design in different domains such as scientific documents, mobile UIs, indoor scenes, etc (Zhong et al., 2019; Deka et al., 2017; Song et al., 2015; Janoch et al., 2013; Xiao et al., 2013; Silberman et al., 2012; Cao et al., 2022; Wu et al., 2019). Most of them perform the generation either based on several hand-crafted constraints or by using unconstrained generation. In practice, it can be more convenient for users to indicate their preferences in natural language.

Among various design tasks, floor plan² design, as shown in Figure 1 is of moderate complexity. However, it still intrinsically involves multiple rounds of communications between clients and de-

²Architectural floor plans, i.e., interior building layouts, are documents that indicate room types, room connections, room sizes, etc. They play a crucial role while designing, understanding, or remodeling indoor spaces (Liu et al., 2017).

signers for specifying requirements, and requires a high level of precision and alignment to detail. AI systems that can learn to generate practically useful floor plan designs directly from natural languages will go a long way in reducing the protracted design process and making Generative AI directly usable for design by the end users.

To allow people without expertise to participate and further enhance the design process, we aim to enable users to design by “telling” instructions, with a specific focus on the floor plan domain as the initial area of research. This sets forth a new machine learning task where the model learns to generate floor plan designs directly from language instructions. However, this task brings up two technical challenges. First, a floor plan is a structured layout that needs three intrinsic components to be valid: (1) *Semantics*, which describes the functionality of rooms (e.g., for living or bathing); (2) *Geometry*, which indicates the shape and dimension of individual rooms; (3) *Topology*, which defines the connectivity among different rooms (Pizarro et al., 2022). Second, these instructions are expressed in natural languages, which, besides the diversity of expressions, inherently suffer from ambiguity, misleading information, and missing descriptions for intrinsic components.

To address the above challenges, we make multiple contributions to initiate research on the task of *language-guided floor plan generation*. First, we contribute a novel dataset, *Tell2Design* (T2D), to the research community. The T2D dataset contains more than 80k real floor plans from residential buildings. Each floor plan is associated with a set of language instructions that describes the intrinsic components of every room in a plan. An example from the dataset is illustrated in Figure 1. Second, we propose a Sequence-to-Sequence (Seq2Seq) approach as a solution to this task which also serves as a strong baseline for future research. Our approach is strengthened by a new strategy to explicitly incorporate the floor plan boundary constraint by transforming the outline into a box sequence. Third, in order to benchmark this novel task and evaluate our proposed approach, we implement strong baselines in text-conditional image generation on our T2D dataset and ask humans to perform the same task. The generation alignment with language instructions is evaluated both quantitatively and qualitatively. Finally, we discuss several future directions that are worth exploring based on our

experimental results.

In summary, our main contributions are:

- We introduce a novel *language-guided floor plan generation* task along with the T2D dataset consisting of both natural human-annotated and large-scale artificially generated language instructions (Section 3).
- We propose a new approach that formulates the floor plan generation task as a Seq2Seq problem (Section 4).
- We provide adequate quantitative evaluations on all baselines and qualitative analysis of human evaluations and performances (Section 5).

2 Related Work

Text-Conditioned Image Generation Image generation is a well-studied problem, and the most popular techniques have been applied for both unconditional image generation and text-conditional settings. Early works apply auto-regressive models (Mansimov et al., 2015), or train GANs (Xu et al., 2018; Zhu et al., 2019; Tao et al., 2022; Zhang et al., 2021; Ye et al., 2021) with publicly available image captioning datasets to synthesize realistic images conditioned on sentence-level captions. Other works have adopted the VQ-VAE technique (Van Den Oord et al., 2017) to text-conditioned image generation by concatenating sequences of text tokens with image tokens and feeding them into auto-regressive transformers (Ramesh et al., 2021; Ding et al., 2021; Aghajanyan et al., 2022). More recently, some works have applied diffusion models (Ho et al., 2020; Nichol and Dhariwal, 2021; Saharia et al., 2022c; Dhariwal and Nichol, 2021; Ho et al., 2022; Saharia et al., 2022a; Rombach et al., 2022; Nichol et al., 2022; Saharia et al., 2022b; Ramesh et al., 2022) and received wide success in image generation, outperforming other approaches in fidelity and diversity, without training instability and mode collapse issues (Brock et al., 2018; Dhariwal and Nichol, 2021; Ho et al., 2022). However, these models operate on extracting high-level visual concepts from the short text and produce artwork-like images that are expected to be realistic and creative, thereby not suitable for generating designs that must satisfy various user/client requirements.

Layout Generation Layout generation is essentially a design process that requires meeting domain-specific constraints, where the desirable

layouts could be documents, natural scenes, mobile phone UIs, and indoor scenes. For example, PubLayNet (Zhong et al., 2019) is proposed to generate machine-annotated scientific documents with five different element categories *text*, *title*, *figure*, *list*, and *table*. RICO (Deka et al., 2017) is introduced to develop user interface designs for mobile applications, which contains *button*, *toolbar*, etc. SUN RGB-D (Song et al., 2015) presents a combined scene-understanding task, including indoor scenes from three other datasets (Janoch et al., 2013; Xiao et al., 2013; Silberman et al., 2012). Moreover, ICVT (Cao et al., 2022) aims to produce advertisement poster layouts automatically, where the image background is given as input. The above methods are designed for different layout domains and cannot be directly applied to floor plan design. Moreover, none of them has considered generating the layout design directly from languages.

Floor Plan Generation Several methods have been proposed to generate floor plan designs automatically (Wu et al., 2018; Liu et al., 2013; Merrell et al., 2010; Hua, 2016; Wu et al., 2019; Chen et al., 2020; Chaillou, 2020). Most of these methods generate floor plans conditioned on certain constraints, such as room types, adjacencies, and boundaries. For example, Merrell et al. (2010) generate buildings with interior floor plans for computer graphics applications using Bayesian networks without considering any human preferences. Liu et al. (2013) present an interactive tool to generate desired floor plan following a set of manually defined rules. Hua (2016) particularly focus on generating floor plans with irregular regions. Wu et al. (2018) cast the generation as a mixed integer quadratic programming problem where some floor plan components are formulated into a set of inequality constraints. More recently, Wu et al. (2019) propose a CNN-based method to determine the location of different rooms given boundary images as a constraint. Chen et al. (2020) provide a small amount of template-based artificial verbal commands and manually parses them to scene graphs for guiding the generation.

In summary, existing methods represent the intrinsic components of floor plans in several specific formats as generation constraints. Some formats are straightforward, such as boundary images, but they only specify limited constraints and lead to less controllable generation. Other formats, such as scene graphs and inequalities, can incorporate more information but require specific domain knowledge

and extra-human efforts in pre-processing. We instead provide a unified and natural way of conditioning the floor plan generation with a set of language instructions that is much more flexible and user-friendly to characterize floor plans with various constraints.

3 Tell2Design Dataset

In this section, we introduce how we construct our T2D dataset, followed by the data analysis and a discussion of the main dataset challenges.

3.1 Task Definition

Given a set of language instructions describing a floor plan’s intrinsic components, our aim is to generate reasonable 2D floor plan designs that comply with the provided instructions.

Input & Output For each data sample, the input is a set of natural language instructions that characterize the key components of the corresponding floor plan design, which include: (1) *Semantics* specifies the type and functionality of each room. For example, a room as *Kitchen* is for cooking. (2) *Geometry* specifies the shape and dimension of each room. For residential buildings, it involves the room’s general orientation (e.g., the north, south, northeast, southwest), area in square feet, aspect ratio, etc. (3) *Topology* describes the relationships among different rooms. It can be divided into three categories: relative location, connectivity, and inclusion³. The desirable output is a structured interior layout that aligns with the input language instructions.

3.2 Floor Plan Collection

We use floor plans from RPLAN⁴ (Wu et al., 2019) to construct our Tell2Design dataset. We remove floor plans with rarely-appeared rooms and merge similar room types such as *Second Room* and *Guest Room*. As a result, 8 different room types (i.e., common room, bathroom, balcony, living room, master room, kitchen, storage, and dining room) and 80,788 floor plans are selected for collecting language instructions. Each floor plan is converted into a 256×256 image where different pixel values

³As a result, language instructions specifying the above features for a floor plan lead to a document-level description in natural language. We compare the T2D dataset with several document-level NLP datasets in Appendix B.

⁴<http://staff.ustc.edu.cn/~fuxm/projects/DeepLayout/index.html>

	Human	Artificial
Avg. # words per instance	200.30	260.47
Avg. # sent. per instance	11.89	23.46
Avg. # words per room	29.48	38.44
Avg. # sent. per room	1.75	3.46

Table 1: Language instruction statistics.

indicate different room types, from which we extract room-type labels and bounding boxes of each room to construct our dataset.

3.3 Language Instruction Collection

Human Instructions To collect real human language instructions, we hire crowdworkers from Amazon Mechanical Turk (MTurk)⁵ and ask them to write a set of instructions for each room according to a given floor plan image. The requested instructions should reflect the *Semantic*, *Geometric*, and *Topological* information of the floor plan, such that designers could ideally reproduce the floor plan layout according to the instructions. In particular, turkers are encouraged to include (but are not limited to) attributes such as room types, locations, sides, and relationships in their instructions. The definitions of these attributes are given as follows: The room type (e.g., bathroom and kitchen) specifies the functionality of a room. The room location specifies the global location of a room in the floor plan and can be described by phrases such as “north side” and “southeastern corner”. The room sides specify the length and width of a room (e.g., “8 feet wide and 10 feet long”). The room relationships specify the relative position of a room with other rooms such as “next to”, “between”, and “opposite”⁶.

Due to the noisy nature of crowdsourcing annotations, we discard some low-quality annotations to ensure the overall quality of our datasets. To this end, we manually review each annotation and discard human instructions with: (1) incoherence, grammatical errors; (2) insufficient attributes; or (3) irrelevance to the given floor plan. As a result, we collect human instructions for 8, 220 floor plans, and 5, 051 of them are finally accepted after

⁵<https://www.mturk.com/>

⁶To mimic the real-world scenarios, we do not provide any bounding box information and ask crowdworkers to make rough estimations of the room size from the given floor plan image only. We also do not restrict the format of text descriptions or require all the above attributes to be mentioned, leading to unstructured and diverse instructions.

Human Instructions:

#1: The balcony, located on the Eastern side of the unit, North of the bathroom, juts out from the rest of the unit. It is 4' wide from West to East, and the entire balcony is located Eastward of the Easternmost wall of the bedroom, common area, and bathroom. The balcony spans 11' from North to South, and its Northern wall is shared with the Southern wall of the kitchen. The kitchen and balcony both jut out approximately 4' to the East of the unit. (*expression diversity*)

#2: The balcony is located in the South after the master room. The room size is 10x4. (*ambiguity*)

#3: The balcony is located in the North middle. (*information missing*)

Artificial Instructions:

#1: Can we have a balcony? Can we have a balcony to be on the north side? It would be great to have a balcony approx 50 sqft with an aspect ratio of 3 over 1. The balcony should be next to the master room. (*complete information with structured expression*)

Figure 2: Human instructions vs. artificial instructions.

manual assessment to construct our dataset⁷.

Artificial Instructions In addition to the human-written instructions, we also generate language instructions artificially for the remaining 75, 737 floor plans from pre-defined templates. To ensure that the artificial instructions are as informative as human-written ones and include all the required components, we ask 5 educated volunteers with natural language processing (NLP) backgrounds to write language instructions for each room that appeared in the given floor plan. We then summarize their instructions into multiple templates and ask expert architectural designers for proofreading. Hence, each instruction template is ensured to be informative, grammatically correct, and coherent.

In summary, our T2D dataset consists of 5, 051 human-annotated and 75, 737 artificially-generated language instructions⁸.

3.4 Data Analysis

In this section, we analyze various aspects of Tell2Design to provide a more comprehensive understanding of the dataset.

Language Instructions Table 1 shows the statistics of the language instructions in our dataset. For each floor plan, the human instructions are organized in nearly 11 sentences consisting of 200 words on average. This includes around 30 words used to describe each room in more than 2 sentences. The artificially generated instructions follow a similar pattern with slightly more words.

To show the connections and differences between human and artificial instructions, we com-

⁷Our human instruction collection involves 5, 109 different workers with 1, 723 working hours, and each worker receives full compensation in alignment with the established standards of MTurk.

⁸More details on the language instruction collection can be found in Appendix C.

pare them for the same room type, *Balcony*, in Figure 2. Artificial instructions always exhibit complete information, including all three key components of a floor plan. They are also formatted in a structured expression, such as “on the ** side”, “** sqft with an aspect ratio of **”, and “next to”. However, human instructions are more diverse in expression but suffer from ambiguity and missing components.

Dataset Comparison In order to see our T2D dataset in perspective, we note its main differences with respect to other related datasets used for similar generation tasks⁹. T2D differs from other datasets in several perspectives: (1) T2D is the first large-scale dataset that aims to generate designs (i.e., floor plans) from direct user input natural language; (2) T2D has much longer text annotations (i.e., 256 words per instance) compared with other text-conditional generation datasets; (3) All text in T2D is written by humans or generated artificially, instead of being crawled from the internet.

3.5 Dataset Challenges

In this section, we discuss three main challenges of our collected T2D dataset. We hope this dataset can facilitate the research on both design generation and language understanding.

Design Generation under Constraints The first challenge is to perform the design generation under much stricter constraints compared with artwork-like text-conditional image generation. Most works in text-conditional image generation operate on generating realistic and creative images that align with the main visual concepts represented by the short input text. However, creating a design from languages has much stricter requirements on precision and alignment to text details. In particular, the generated floor plan design should comply with constraints such as room type, location, size, and relationships, which are specified by users using natural languages. Our main results in Section 5 comparing different baselines demonstrate that existing text-conditional image generation techniques fail to follow detailed user requirements on this design task.

Fuzzy & Entangled Information The second challenge is to understand the big picture of the entire floor plan from document-level unstructured text with fuzzy and entangled information. Be-

⁹We provide detailed comparisons with tables in Appendix B.

sides the general abilities required for language understanding, such as entity recognition, coreference resolution, relation extraction, etc, models also need to collaborate with fuzzy individual room attributes and reason over entangled relationships among different rooms to understand the entire floor plan. Specifically, one language instruction usually either specifies fuzzy descriptions for a room’s *Semantic* and *Geometric* information such as “on the north side” and “at the southeast corner”, or indicates the relationship of one specific room with others like “next to” and “between”. The provided information in such instructions is coarse and relative, rather than complete and precise information like numerical coordinates. As a result, to determine the location of all rooms and design a reasonable floor plan, models must collaborate with fuzzy and entangled information residing in multiple instructions, and incorporate the boundary information. Human evaluations in Section 5.4 demonstrate that room relationships described in language instructions are the most challenging component to be understood and aligned with.

Noisy Human Instructions The third challenge comes from the ambiguous, incomplete, or misleading information in human instructions. As introduced in Section 3.3, the artificial instructions are template-based so that they always contain precise and coherent information. However, for human-written language instructions, ambiguous or noisy information always exists. For example, during human instruction collection, workers are asked to write natural sentences estimating some numeric-related attributes like room size and aspect ratio referring to the floor plan image, which may sometimes be inaccurate. Moreover, as previously discussed in Figure 2, other than the expression diversity, human instructions also exhibit ambiguous phrasing and incomplete information. It is thus more challenging for models to retrieve accurate, complete, and consistent information from human instructions.

4 T2D Model

In this section, we propose a simple yet effective method for *language-guided floor plan generation*. Unlike existing floor plan generation methods (Wu et al., 2019; Chen et al., 2020) that use a regression head to generate the bounding box of each room one at a time, we cast the floor plan genera-

tion task as a Seq2Seq problem under the encoder-decoder framework, where room bounding boxes are re-constructed into a *structured* target sequence. This way, our method can easily deal with various lengths of instructions for floor plans with different numbers of rooms.

4.1 Target Sequence Construction

Recall that our aim is to generate a floor plan layout from language instructions, where each room can be represented by a room-type label (e.g., bathroom and kitchen) and a bounding box. One bounding box can be determined by four values (x, y, h, w) , which indicate the x and y coordinate of the center point, height (h), and width (w), respectively. To solve language-guided floor plan generation as a Seq2Seq problem, we treat the instructions as the input sequence and consider bounding boxes of rooms as the target sequence. Specifically, each of the continuous values (x, y, h, w) is discretized into integers between $[0, 255]$, and the room type is given by the plain text in natural language, so that they can be naturally represented as a sequence of tokens. The target sequence is then constructed by grouping the room type and the bounding box together with certain special tokens. For example, the target sequence for a *Balcony* with the bounding box $(87, 66, 18, 23)$ is given as follows:

[Balcony | x coordinate = 87 | y coordinate = 66 | height = 18 | width = 23],

where the special tokens “[” and “]” are used to indicate the start and end of the target sequence for one room and “|” is used to separate different target components. We have also added semantic prefixes such as “x coordinate =” and “height =” before the values of bounding boxes to assist the target sequence generation. Finally, we concatenate the target sequences of all the rooms in a floor plan and add an $\langle \text{eos} \rangle$ token at the end to indicate the end of the overall target sequence.

4.2 Boundary Information Incorporation

The outline/boundary of a floor plan is one of the most important constraints in floor plan generation, which directly affects where each room should be placed and how different rooms should be aligned with the floor plan boundary. However, it is non-trivial to incorporate such boundary information into floor plan generation. Previous methods either fail to take the floor plan outline into account (Wu




Figure 3: An illustration of how to transform visual floor plan boundaries into boxes.

et al., 2018; Liu et al., 2013; Merrell et al., 2010; Hua, 2016; Chen et al., 2020) or only consider the boundary image, ignoring all other constraints (Wu et al., 2019; Chaillou, 2020), leading to less controllable floor plan design.

In this work, we propose a novel approach to incorporate boundary information by representing the irregular outline as a set of boxes. The idea is to encode the boundary information by an enclosing box that is the minimum bounding region containing the entire floor plan and several exterior boxes that are inside the enclosing box but excluded from the floor plan. Figure 3 illustrate how the floor plan boundary can be characterized by the enclosing (in red) and exterior boxes (in yellow). This way, we have an enclosing box represented by $(x^{en}, y^{en}, h^{en}, w^{en})$ and M exterior boxes by $(x_i^{ex}, y_i^{ex}, h_i^{ex}, w_i^{ex})$. Then we adopt a similar strategy in Section 4.1 to represent the enclosing and exterior boxes in a sequence as follows:

$$+ x^{en} y^{en} h^{en} w^{en} - x_1^{ex} y_1^{ex} h_1^{ex} w_1^{ex} \dots - x_M^{ex} y_M^{ex} h_M^{ex} w_M^{ex},$$

where the coordinates of the enclosing and exterior box are following the tokens “+” and “-”, respectively.

Finally, the above sequence is added after the input language instructions for training. Our experimental results in Section 5 show that the proposed boundary information incorporation strategy is effective in enhancing our Seq2Seq method to generate valid rooms that align well with the floor plan boundary.

4.3 Architecture, Objective and Inference

Treating the target sequences that we construct from floor plans as a text sequence, we turn to recent architectures and objective functions that have been effective in Seq2Seq language modeling.

Architecture We use the popular Transformer-based (Vaswani et al., 2017) encoder-decoder structure to build our Seq2Seq model for floor plan generation. The model is initialized by a pre-trained language model T5 (Raffel et al., 2020) for better language understanding abilities¹⁰.

Objective Similar to language modeling, our T2D model is trained to predict the next token, given an input sequence and preceding tokens, with a maximum likelihood objective function, i.e.,

$$\max_{\theta} \sum_{j=1}^L \log P_{\theta}(\tilde{y}_j | \mathbf{x}, \mathbf{y}_{1:j-1}), \quad (1)$$

where \mathbf{x} is a set of instructions in natural language concatenated with the previously defined boundary sequence, \mathbf{y} is the target bounding box sequence, and L is the target sequence length.

Inference At inference time, we sample¹¹ tokens one by one from the model likelihood, i.e., $P(\tilde{y}_j | \mathbf{x}, \mathbf{y}_{1:j-1})$. The sequence generation ends once the <eos> token is sampled, and it is straightforward to parse the target sequence into predicted floor plans.

5 Experiments

5.1 Baselines

Since our T2D dataset is the first to consider language-guided floor plan generation, existing layout generation methods are not applicable to this task. To further illustrate the challenge of the design generation task and the difference with the existing text-conditional image generation problem, we adapt several state-of-the-art text-conditional image generation methods as baselines for comparison. In particular, we compare our method with the following:

- Obj-GAN (Li et al., 2019) is an object-driven attentive generative adversarial network that follows a two-step generation process. We apply the first-step box generator module, which takes the language as input and generates target objects’ bounding boxes (with class labels).

¹⁰We have also tried other pre-trained language models like Bart (Lewis et al., 2020), and preliminary experiments indicate that initializing our model with T5 leads to better performances.

¹¹There are several common sampling strategies like *Greedy Search*, *Beam Search*, and *Nucleus Sampling* (Holtzman et al., 2019). We apply *Greedy Search* since it leads to better generation quality in our preliminary experiments.

- CogView (Ding et al., 2021) applies pre-trained VQ-VAE to transform the target image into a sequence of image tokens. Then the text and image tokens are concatenated together and fed to a Transformer decoder (i.e., GPT (Brown et al., 2020; Radford et al., 2019)) to generate text-conditional images.
- Imagen (Saharia et al., 2022b) is one of the state-of-the-art text-to-image generation models that build upon both large language models (e.g., T5) for text understanding and diffusion models for high-fidelity image generation.

5.2 Experimental Settings

Model Training For model training, we consider a Warm-up + Fine-tuning pipeline (Goyal et al., 2017), where the model is first warmed up on 75,737 artificial instructions, and then fine-tuned on 2,743 human instructions. To evaluate how floor plan generation methods generalize to unseen instructions, we use the remaining 2,308 human instructions as the test set, such that there is no overlapping between annotators of the training set and the test set¹².

Evaluation Metrics For testing, we use macro and micro Intersection over Union (IoU) scores between the ground-truth (GT) and generated floor plans at pixel level as the evaluation metrics, whose definitions are given as follows:

$$\text{Micro IoU} = \frac{\sum_{r=1}^R I_r}{\sum_{r=1}^R U_r}, \text{Macro IoU} = \frac{1}{R} \sum_{r=1}^R \frac{I_r}{U_r},$$

where the I_r and U_r , respectively, denote the intersection and union of the ground-truth and predicted rooms labeled as the r -th room type in a floor plan. R is the total number of room types. Macro IoU calculates the average IoU over different types of rooms, and Micro IoU calculates the global IoU by aggregating all rooms.

Since Obj-GAN and our T2D model generate bounding boxes rather than images, we use a simple strategy to transform the outputs of Obj-GAN and the T2D model into images without any further refinement for a fair comparison. Specifically, we paint each room in descending order in terms of the total area of the room type¹³ and different colors

¹²We provide more implementation details in Appendix A.

¹³Total area of the room type is computed by adding up the specific room type area across all floor plans in our dataset. This gives us the following order: living room, common room,

Models	Micro IoU	Macro IoU
Training on artificial instructions only		
Obj-GAN	15.74	11.12
CogView	10.01	8.31
Imagen	14.74	15.57
T2D (w/o bd)	6.46	4.01
T2D	9.13	6.06
Training on human instructions only		
Obj-GAN	10.72	8.29
CogView	13.48	11.26
Imagen	9.29	6.64
T2D (w/o bd)	32.22	26.24
T2D	42.93	38.48
Warm up on artificial + fine-tune on human		
Obj-GAN	10.68	8.44
CogView	13.30	11.43
Imagen	12.17	14.96
T2D (w/o bd)	35.95	29.95
T2D	54.34	53.30

Table 2: IoU scores between ground-truth and generated floor plans for the T2D model and other baselines.

refer to different room types. Previous colors will be replaced by the subsequent paintings if there is an overlapping between bounding boxes (rooms). For image-based approaches (e.g., CogView and Imagen), we compute the maximized IoU scores by shifting the floor plan central point in the generated image.

5.3 Main Results

Table 2 shows the floor plan generation results on the T2D dataset, where T2D (w/o bd) indicates the T2D model without incorporating boundary information¹⁴. The T2D model achieves the highest IoU scores with a micro IoU of 54.34 and a macro IoU of 53.30, outperforming other baselines by a large margin. These can be attributed to our Seq2Seq model in controlling the target box sequence generation based on salient information extracted from language instructions. In contrast, text-conditional image generation methods fail to perform well. This is probably because those models are designed to generate artwork-like images with high-level visual concepts from the short text, instead of following multiple instructions with various constraints for a specific design.

master room, balcony, bathroom, kitchen, storage, dining room.

¹⁴We provide baseline generation samples in Appendix D.

Alignment	GT ratings	T2D ratings
Room type	4.99	4.71
Room location	4.86	3.67
Room size	4.75	3.89
Relationships	4.89	3.65
Meet all %	85%	38%

Table 3: Human evaluation results.

When training only on artificial instructions while testing on human-written ones, our method cannot perform well. This indicates there is a language distribution gap between artificial and human instructions. Nevertheless, when artificial instructions are used for warming up before training on human instructions, the performance of our method is significantly improved with over 10 IoU scores increment. This suggests that despite the language gap, artificial and human instructions are mutually beneficial data portions during training.

In addition, in all the training settings, representing the floor plan boundary as a sequence of boxes consistently improves the performance of our Seq2Seq approach. This demonstrates that this strategy could be one of the possible solutions to incorporate the floor plan boundary.

5.4 Result Analysis

It is worth noting that the quantitative results *indirectly* evaluate how well the generated floor plans align with the language instructions since IoU scores essentially measure the overlap between generated and ground-truth floor plan layouts. Due to the complexity of our task, it is possible for the same language instruction to map to multiple floor plan designs. Therefore, a low IoU score does not necessarily mean a bad generation.

Human Evaluations To *directly* evaluate the alignment between generated floor plans and language instructions, we conduct human evaluations on a subset of the T2D test set, which consists of 100 randomly sampled instructions written by different annotators. For this purpose, we invite 5 volunteers with NLP backgrounds to evaluate the degree of alignment between source language instructions and target floor plans.

Specifically, we consider four partial alignment criteria in terms of room types, locations, sizes, and relationships. Each volunteer is asked to provide four ratings on a scale of 1 to 5, according to the

Models	Micro IoU	Macro IoU
T2D	55.10	55.16
Human	64.67	62.32

Table 4: The T2D model vs. human performance.

above-mentioned alignment criteria, respectively¹⁵. Besides, we also consider global alignment and ask our volunteers to justify whether the generated floor plan meets *all* the specifications in the instructions. We perform the above subjective evaluations for both T2D-generated and ground-truth floor plan designs.

Table 3 shows the human evaluation results. As can be seen, ground-truth floor plans get high ratings for all the partial alignment criteria, and 85% of them meet all the requirements specified in the instructions. This indicates our dataset contains high-quality human instructions that align well with the ground-truth floor plan designs. On the other hand, our T2D model receives no rating less than 3.5, indicating that at least 50% rooms with respect to their locations, sizes, and relationships can be correctly predicted. However, our method still has a gap with ground truth designs, especially in room location and relationships, which indicates the potential for improvements.

Human Performance To study human performance on the T2D task, we further ask our volunteers to design floor plans on 100 instances of the same subset used for human evaluations¹⁶. Table 4 reports the IoU scores for our T2D model and human performance. Humans generally achieve better IoU scores. However, even if human-generated floor plans intrinsically have much better alignment with the instructions, they only obtain around 63% IoU scores with the ground truths. This exposes the nature of design diversity, i.e., a set of language instructions can map to multiple plausible floor plan designs. Figure 4 provides a real example that both our method and humans follow the same instructions¹⁷ but generate different floor plans.

¹⁵For example, a rating of r with respect to room locations indicates that $(r - 1) \times 20\%$ to $r \times 20\%$ rooms in the floor plan follow their location specifications in the instructions.

¹⁶Specifically, volunteers are asked to draw a bounding box for each room, given the floor plan outline and the language instructions.

¹⁷The language instructions for this sample are shown in Figure 10 in Appendix D.




Figure 4: Floor plan layouts from our T2D model, the ground-truth, and human annotators based on the same language instructions. Although all of them satisfy the given instructions, they are different in detail, especially for the upper right corner of the floor plan.

6 Future Research

In the future, the following directions may be worth exploring to promote the performance or extend our task: (1) How to build robust language understanding models that can adapt to the presence of noise in human instructions or even locate and refine potentially inconsistent information? (2) How to explicitly incorporate the nature of design diversity and develop techniques for diverse floor plan design? (3) How to extend the language-guided floor plan generation task to more domains or more practical but challenging scenarios, where designs should be refined according to feedback from users/clients?

7 Conclusion

In this paper, we initiate the research of a novel language-guided design generation task, with a specific focus on the floor plan domain as a start. We formulate it as *language-guided floor plan generation* and introduce *Tell2Design* (T2D), a large-scale dataset that features floor plans with natural language instructions in describing user preferences. We propose a Seq2Seq model as a strong baseline and compare it with several text-conditional image generation models. Experimental results demonstrate that the design generation task brings up several challenges and is not well-solved by existing text-conditional image generation techniques. Human evaluations assessing the degree of alignment between text and design, along with the human performance on the task, expose the challenge of understanding fuzzy and entangled information, and the nature of design diversity in our task. We hope this paper will serve as a foundation and propel future research on the task of the language-guided design generation.

Limitations

The proposed T2D dataset has several limitations, which could be addressed in future work. First, it only considers and collects language instructions for the floor plan domain. Future work could extend this language-guided design generation task to other design domains such as documents, mobile UIs, etc. Second, it is limited in the scope of languages where we only collect instructions written in English. Future work could assess the generalizability of the T2D dataset to other languages. Third, although generating floor plan designs from languages exhibit diversity, we do not consider improving generation diversity at this moment. Future works could consider building frameworks that specifically aim at design diversity.

Ethics Statement

In this section, we discuss the main ethical considerations of *Tell2Design* (T2D): (1) Intellectual property protection. The floor plans of the T2D dataset are from the RPLAN (Wu et al., 2019) dataset. Our dataset should be only used for research purposes. (2) Privacy. The floor plan data sources are publicly available datasets, where private data from users and floor plans have been removed. Language instructions are either generated artificially or collected from Amazon Mechanical Turk, a legitimate crowd-sourcing service, and do not contain any personal information. (3) Compensation. During the language instruction collection, the salary for annotating each floor plan is determined by the instruction quality and Mturk labor compensation standard.

Acknowledgements

We would like to thank the anonymous reviewers, our meta-reviewer, and senior area chairs for their constructive comments and support of our work. We also gratefully acknowledge the support of NVIDIA AI Technology Center (NVAITC) for our research. This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Program (AISG Award No: AISG2-RP-2020-016).

References

Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis,

and Luke Zettlemoyer. 2022. *Cm3: A causal masked multimodal model of the internet*. *ArXiv*.

Andrew Brock, Jeff Donahue, and Karen Simonyan. 2018. *Large scale gan training for high fidelity natural image synthesis*. *ArXiv*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. *Language models are few-shot learners*. In *Proc. of NIPS*.

Yunning Cao, Ye Ma, Min Zhou, Chuanbin Liu, Hongtao Xie, Tiezheng Ge, and Yuning Jiang. 2022. *Geometry aligned variational transformer for image-conditioned layout generation*. In *Proc. of ACM Multimedia*.

Stanislas Chaillou. 2020. *Archigan: Artificial intelligence x architecture*. In *Architectural intelligence*.

Qi Chen, Qi Wu, Rui Tang, Yuhan Wang, Shuai Wang, and Mingkui Tan. 2020. *Intelligent home 3d: Automatic 3d-house design from linguistic descriptions only*. In *Proc. of CVPR*.

Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hibsman, Daniel Afegan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. 2017. *Rico: A mobile app dataset for building data-driven design applications*. In *Proc. of UIST*.

Prafulla Dhariwal and Alexander Nichol. 2021. *Diffusion models beat gans on image synthesis*. In *Proc. of NIPS*.

Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. 2021. *Cogview: Mastering text-to-image generation via transformers*. In *Proc. of NIPS*.

Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. *Accurate, large minibatch sgd: Training imagenet in 1 hour*. *ArXiv*.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. *Denoising diffusion probabilistic models*. In *Proc. of NIPS*.

Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. 2022. *Cascaded diffusion models for high fidelity image generation*. In *Proc. of JMLR*.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. *The curious case of neural text degeneration*. In *Proc. of FAIR*.

Ruizhen Hu, Zeyu Huang, Yuhan Tang, Oliver Van Kaick, Hao Zhang, and Hui Huang. 2020. *Graph2plan: Learning floorplan generation from layout graphs*. *ACM Transactions on Graphics*.

- Hao Hua. 2016. [Irregular architectural layout synthesis with graphical inputs](#). *Automation in Construction*.
- Allison Janoch, Sergey Karayev, Yangqing Jia, Jonathan T Barron, Mario Fritz, Kate Saenko, and Trevor Darrell. 2013. [A category-level 3d object dataset: Putting the kinect to work](#). In *Consumer Depth Cameras for Computer Vision*.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *ArXiv*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proc. of ACL*.
- Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. 2019. [Object-driven text-to-image synthesis via adversarial training](#). In *Proc. of CVPR*.
- Chen Liu, Jiajun Wu, Pushmeet Kohli, and Yasutaka Furukawa. 2017. [Raster-to-vector: Revisiting floorplan transformation](#). In *Proc. of ICCV*.
- Han Liu, Yong-Liang Yang, Sawsan AlHalawani, and Niloy J Mitra. 2013. [Constraint-aware interior layout exploration for pre-cast concrete-based buildings](#). *The Visual Computer*.
- Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. 2015. [Generating images from captions with attention](#). *ArXiv*.
- Paul Merrell, Eric Schkufza, and Vladlen Koltun. 2010. [Computer-generated residential building layouts](#). In *Proc. of SIGGRAPH*.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022. [Glide: Towards photorealistic image generation and editing with text-guided diffusion models](#). In *Proc. of ICML*.
- Alexander Quinn Nichol and Prafulla Dhariwal. 2021. [Improved denoising diffusion probabilistic models](#). In *Proc. of ICML*.
- Pablo N Pizarro, Nancy Hitschfeld, Ivan Sipiran, and Jose M Saavedra. 2022. [Automatic floor plan analysis and recognition](#). *Automation in Construction*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). In *Proc. of JMLR*.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. [Hierarchical text-conditional image generation with clip latents](#). *ArXiv*.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. [Zero-shot text-to-image generation](#). In *Proc. of ICML*.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. [High-resolution image synthesis with latent diffusion models](#). In *Proc. of CVPR*.
- Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022a. [Palette: Image-to-image diffusion models](#). In *Proc. of SIGGRAPH*.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. 2022b. [Photorealistic text-to-image diffusion models with deep language understanding](#). *ArXiv*.
- Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. 2022c. [Image super-resolution via iterative refinement](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Sachith Seneviratne, Damith Senanayake, Sanka Rasnayaka, Rajith Vidanaarachchi, and Jason Thompson. 2022. [Dalle-urban: Capturing the urban design expertise of large text to image transformers](#). *ArXiv*.
- Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. 2012. [Indoor segmentation and support inference from rgb-d images](#). In *Proc. of ECCV*.
- Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. 2015. [Sun rgb-d: A rgb-d scene understanding benchmark suite](#). In *Proc. of CVPR*.
- George Stiny. 1980. [Introduction to shape and shape grammars](#). *Environment and planning B: planning and design*.
- Ming Tao, Hao Tang, Songsong Wu, Nicu Sebe, Xiao-Yuan Jing, Fei Wu, and Bingkun Bao. 2022. [Df-gan: Deep fusion generative adversarial networks for text-to-image synthesis](#). In *Proc. of CVPR*.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. [Neural discrete representation learning](#). In *Proc. of NIPS*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proc. of NIPS*.
- Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Zhentao Tan, Lu Yuan, Weiming Zhang, and Nenghai Yu. 2022. [Hairclip: Design your hair by text and reference image](#). In *Proc. of CVPR*.

- Wenming Wu, Lubin Fan, Ligang Liu, and Peter Wonka. 2018. *Miqp-based layout design for building interiors*. In *Proc. of CGF*.
- Wenming Wu, Xiao-Ming Fu, Rui Tang, Yuhan Wang, Yu-Hao Qi, and Ligang Liu. 2019. *Data-driven interior plan generation for residential buildings*. *ACM Transactions on Graphics*.
- Jianxiong Xiao, Andrew Owens, and Antonio Torralba. 2013. *Sun3d: A database of big spaces reconstructed using sfm and object labels*. In *Proc. of ICCV*.
- Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. 2018. *Attngan: Fine-grained text to image generation with attentional generative adversarial networks*. In *Proc. of CVPR*.
- Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. *Docred: A large-scale document-level relation extraction dataset*. In *Proc. of ACL*.
- Hui Ye, Xiulong Yang, Martin Takac, Rajshekhar Sunderraman, and Shihao Ji. 2021. *Improving text-to-image synthesis using contrastive learning*. In *Proc. of BMVC*.
- Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. 2021. *Cross-modal contrastive learning for text-to-image generation*. In *Proc. of CVPR*.
- Xujie Zhang, Yu Sha, Michael C Kampffmeyer, Zhenyu Xie, Zequn Jie, Chengwen Huang, Jianqing Peng, and Xiaodan Liang. 2022. *Armani: Part-level garment-text alignment for unified cross-modal fashion design*. In *Proc. of ACM Multimedia*.
- Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. 2019. *Publaynet: largest dataset ever for document layout analysis*. In *Proc. of ICDAR*.
- Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. 2019. *Dm-gan: Dynamic memory generative adversarial networks for text-to-image synthesis*. In *Proc. of CVPR*.

A Implementation Details

T2D Parameters In practice, we initialize all weights of our proposed baseline method from T5-base¹⁸. In training, we use Adam (Kingma and Ba, 2014) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e - 08$ to update the model parameters. We fine-tune our model on 3 RTX 8000 GPUs with batch size 12 and learning rate $5e - 4$ for 20 epochs.

¹⁸<https://huggingface.co/t5-base>

Baseline Implementation For the mentioned baselines, only Obj-GAN and CogView are open-sourced. Therefore, we adapt and implement the models from their official GitHub repositories¹⁹. However, as Imagen’s source codes are not published, we implement it from the most starred GitHub repository²⁰ (i.e., 5.9k stars until writing this paper) and adapt it to our T2D dataset. We use the process floor plan images in Graph2Plan (Hu et al., 2020) for training. Although all baselines have provided pre-trained checkpoints for fine-tuning, our preliminary experiments indicate that training those baselines from scratch on the T2D dataset will obtain better performances. One most probable reason is the huge discrepancy between the data distributions of the baseline pre-training corpus and our T2D dataset. Those baseline checkpoints are mostly trained with real-life images with various objects and backgrounds. But our T2D dataset only focuses on the floor plan domain.

Specifically, for Obj-GAN, we adopt and freeze the pre-trained text encoder, and train the rest of the networks (e.g., LSTMs) from scratch. For CogView, we freeze the pre-trained VQ-VAE and initialized the main backbone, decoder-only transformer, from GPT (Radford et al., 2019; Brown et al., 2020). During training, only the parameters of the transformer backbone will be updated. For Imagen, we import the T5-large model’s encoder from *Hugging Face* for text encoding and freeze all its parameters during training. The rest U-nets for diffusion will be updated according to the loss propagation.

B Dataset Analysis

Floor Plan Statistics As shown in Tabel 5 and Figure 5, we present the statistics on the occurrence of each room type and the number of rooms per floor plan. There are 8 types of rooms in total. More than 92% of floor plans include at least 6 distinct rooms, and the most frequent room types are *Common Room*, *Bathroom*, *Balcony*, *Living Room*, *Master Room*, and *Kitchen*.

Dataset Comparison As shown in Table 8, compared with other-related layout generation datasets, our T2D dataset is the first to have language annotations and aims to generate layout designs directly

¹⁹<https://github.com/jamesli1618/Obj-GAN>; <https://github.com/THUDM/CogView>

²⁰<https://github.com/lucidrains/imagen-pytorch>

Room type	#Floor plan
CommonRoom	100,847
Bathroom	97,113
Balcony	86,545
LivingRoom	80,788
MasterRoom	80,466
Kitchen	77,768
Storage	3,351
DiningRoom	1,312

Table 5: Room type occurrences in the T2D dataset.

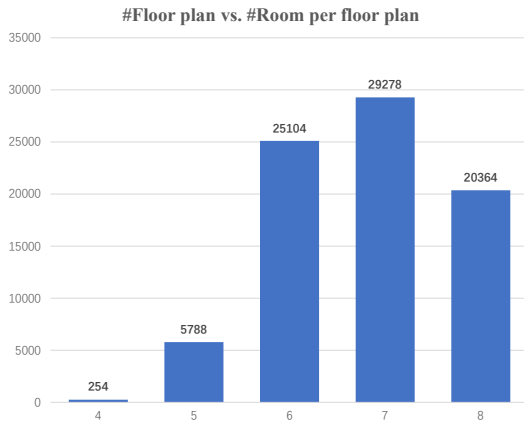


Figure 5: Number of floor plans vs. number of rooms per floor plan in the T2D dataset.

from languages. Since generating floor plan designs from language instructions can be naturally formulated into a text-conditional image generation problem, we compare our dataset with two benchmarking text-conditional image generation datasets in Table 6. We observe that our dataset is with a similar number of images with MSCOCO and Flickr30K but contains far longer text annotation for each image (i.e. T2D has 256 words on average describing each floor plan image). Moreover, as a set of language instructions for a floor plan results in a document-level text description, we compare our dataset with other document-level NLP datasets in Table 7. We hope that our dataset can also propel the research on document-level language understanding. It is shown that our dataset has a comparable total number of samples and words with the largest DocRED (Yao et al., 2019). More importantly, our “documents” are either human-annotated or artificially generated, instead of being crawled from the internet.

Dataset	# Img.	Avg. # Words
MS COCO	82,783	11.3
Flickr30K	31,000	11.8
T2D (ours)	80,788	256.7

Table 6: Comparisons between our T2D dataset and text-conditional image generation datasets.

Dataset	# Doc.	# Word	# Sent.
SCIERC	500	60755	2217
BC5CDR	1,500	282k	11,089
DocRED (Human)	5,053	1,002k	40,276
DocRED (Distantly)	101,873	21,368k	828,115
T2D (Human)	5,051	1,011k	60,057
T2D (Artificial)	75,737	19,727k	1,776k

Table 7: Comparisons between our T2D dataset and document-level NLP datasets.

C Dataset Collection Details

Human instruction We employ Amazon Mechanical Turk (MTurk)²¹ to let annotators write language instructions for a given RGB 2D floor plan. Amazon considers this web service “artificial intelligence,” and it is applied in various fields, including data annotation, survey participation, content moderation, and more. The global workforce (called “turkers” in the lingo) is invited for a small reward to work on “Human Intelligence Tasks” (HITs), created from an XML description of the task from business companies or individual sponsors (called “requesters”). HITs can display a wide variety of content (e.g., text and images) and provide many APIs, e.g., buttons, checkboxes, and input fields for free text. In our case, turkers are required to fill the blank input fields in HITs with language instructions for each room, following our guidelines. A screenshot of one of our HITs is displayed in Figure 6. We also show a full example of human instructions in Figure 7.

Artificial Instruction The artificial instructions in our T2D dataset are generated from scripts with several pre-defined templates. We carefully select volunteers with natural language processing backgrounds for drafting templates. Before participating in the annotation process, each annotator was required to undergo a qualification round consisting of a series of test annotations. We illustrate how we generate an instruction to describe one room’s as-

²¹<https://www.mturk.com/>

Instructions
Shortcuts

⊗

Instructions ✕

Background

Assume that you were writing to a designer or an architect. If the architect were to see your descriptions, they should know what the floor plan looks like and ideally be able to reproduce the floor plan given the house outline.

Task Instructions

- Given a floor plan image, describe each room (bounding box) in natural language with the following attributes:
 - Location:** The absolute location of a room in the floor plan.
Examples: North side, north west corner, north east corner, south side, south west corner, south east corner.
 - Relationship:** The relative location of a room with other rooms.
Examples: near, next, between, in front of, opposite, en suite.
 - Size:** The numerical estimation of the room size in one of the following measurements:
 - Sides (Recommended):** The length and width of a room in the floor plan.
Examples: 8 feet width and 10 feet length.
 - Aspect ratio:** A proportional relationship between a room's width and length.
Examples: 2 over 3, 16:9.
 - Area:** The area of a room in the floor plan.
Examples: 20 sqft feet.
- All of the three attributes should be mentioned in your descriptions.
- Please arrange the room descriptions in the **descending order** of importance (i.e., describe the most important room first).
- Please make sure your descriptions are **grammatically correct and coherent**, and clearly indicate the **name** of each room appeared in the floor plan image.

Payment

Your results will be reviewed and evaluated by the requesters in terms of precision, coherence, and cohesion in describing the floor plan image.

We will pay a base reward of \$0.50/HIT for every accepted HIT.

[More Instructions](#)

IMPORTANT:

Please do **NOT** work on this HIT without carefully reading the detailed instructions (please click the top left tab) first. Results that do not satisfy the task instructions will be rejected.

Instructions:

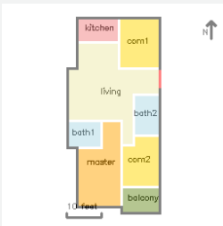
Given a floor plan image, write a few lines for each room to describe the following attributes (see the detailed instructions for more information):

- Location:** The absolute location of a room in the floor plan.
- Relationship:** The relative location of a room with other rooms.
- Size:** The **numerical** estimation of the room size. To estimate the room size, you may use the **scale bar** at the bottom of the floor plan image, whose length represents 10 feet in real size.

Your descriptions for each room should be **informative, grammatically correct, and coherent**. Poor quality works (see the following examples) are likely rejected.

- Incomplete, incorrect, or incoherent sentences:** "north facing", "SOUTH CORNER", "between kitchen and master", "right corner up side master room", "master room top side corner right side", "the master room in balcony side", "SOUTH EAST CORNER ROOM, 15 X20 FEET, COMM 2 ROOM AND BATH ROOM CENTER SIDE" (incoherent descriptions).
- Less informative descriptions:** "The Master room is located at the left side." (no relation and size information), "Master room is behind the living room" (no location and size information), "Master room is down south corner of the image. The absolute location of a room in the floor plan." (no size information).
- Inaccurate size information:** "The living room in the center position. it is 5feet."

Floor Plan Image:



Room Descriptions:

Please **drag and drop** the descriptions of each room in the **descending order** of importance

living room: (A paragraph describing living. Please DO mention the FULL NAME of the room in your descriptions.)

master room: (A paragraph describing master. Please DO mention the FULL NAME of the room in your descriptions.)

kitchen: (A paragraph describing kitchen. Please DO mention the FULL NAME of the room in your descriptions.)

Submit

Figure 6: A screenshot of our HITs GUI.

Dataset	Domain	Basic objects	Object annotations	Other annotations
PubLayNet	scientific documents	{text,title,figure,...}	bounding boxes	None
RICO	mobile UIs	{button,tool bar,...}	bounding boxes,interactions	animations,hierarchies
SUB RGB-D	3D indoor scenes	{chair,table,pillow,...}	bounding boxes	2D&3D polygons
ICVT	poster designs	{text,logo,...}	bounding boxes,substrates	background images
T2D (ours)	floor plans	{kitchen,balcony,...}	bounding boxes	language instructions

Table 8: Comparisons between our T2D dataset and other layout generation datasets.

Human Instruction Example:

Balcony one is located at the northwestern point of the floorplan. It is approximately 50 square feet in size. It can be accessed through the kitchen, bathroom and common room 2.

Balcony 2 is located in the most southern point of the floorplan, just east of the master room. It is roughly 12 feet in length and five feet in width. Access points include the livingroom and master room.

The **bathroom** is located south of the kitchen. It can be accessed through the east wall of the livingroom, as well as the southern kitchen wall. Common room 2 has access to the bathroom through its eastern wall. It is the smallest room in the floorplan. At approximately 25 square feet in size, it is just a bit smaller than balcony 1 and half the size of balcony 2.

Common room one is located on the western portion of the floor plan. It can be accessed through common room 2 at the north, the master room at the south and the livingroom from the east. Common room one is about 100 square feet, 10 feet in length and 10 feet in width, making it the second largest room in the floor plan.

Common room 2 is just a bit smaller than Common room 1, approximately 90 square feet. It is located just north of Common room 1, with access points including the bathroom at the northeast, balcony at the north, and livingroom at the east.

The **kitchen** is located in the most northern point of the floorplan. It is roughly 50 square feet, 10 feet in length and 5 feet in width. The bathroom can be accessed at the southern point of the kitchen, the first balcony toward the western kitchen wall and the livingroom at the eastern wall. The kitchen is relatively closest in size to the balconies.

The **livingroom** is located in the northeastern portion of the floorplan. It is entered through the front entry door, and is approximately 480 square feet. The bathroom, kitchen, common rooms one and two, the master bedroom and second balcony can all be accessed through the livingroom.

The **master room** is located in the southern end of the floorplan. It is roughly 200 square feet, 10 feet in width and 20 feet in length. The second balcony, livingroom and common room one can be accessed through the master room. The master room is about the size of the two common rooms combined.

Figure 7: An example of human-written language instructions from the T2D dataset.

pect ratio in Figure 8. We also show a full example of artificial instructions in Figure 9.

D Baseline Generation Samples

To better understand and compare different baselines, we provide a case study of generated samples for the same language instructions from all baselines shown in Figure 10. Obj-GAN (Li et al., 2019) has difficulties in capturing salient information from the given language instructions, resulting in generating rooms with incorrect attributes and relationships. One possible reason could be that it does not utilize any pre-trained large language model and thus suffers from understanding the given document-level instructions. CogView (Ding et al., 2021) instead auto-regressively generates the image tokens conditioned on all input instructions with a pre-trained GPT as the backbone. However, the image tokens sampled near the end of the generation show confusing information, resulting in an incomplete design. This is probably because presenting the whole floor plan design as a sequence of image tokens hinders the potential connections among different elements in the floor

plan. Imagen (Saharia et al., 2022b) exhibits its strong ability to generate realistic images in the target domain. However, it also fails to meet various design requirements specified in language instructions, indicating its limitation for design generation under multiple stricter constraints.

Instruction Backbone	Approximate Phrase	Values
"Make the aspect ratio of $\{room.type\}$ "		
"The aspect ratio of $\{room.type\}$ should be "	"about"	
"I would like to have the aspect ratio of $\{room.type\}$ " +	"around" +	$\{\{room.aspect_ratio\}\}$ "
"Can you make the aspect ratio of $\{room.type\}$ "	"approx"	
"Can we have the aspect ratio of $\{room.type\}$ to be "		
"It would be great to have the aspect ratio of $\{room.type\}$ "		

Figure 8: Illustration on generating artificial instructions describing room’s aspect ratio.

Artificial Instruction Example:

It would be good to have a **common room** . I would like to place common room at the north side of the apartment. The common room should be around 200 sqft with the aspect ratio of 3 over 4. The common room should have an en-suite bathroom. The common room should be next to the bathroom, kitchen, balcony.

The **bathroom** should be considered. Place bathroom at the south side of the apartment. Make bathroom around 50 sqft with the aspect ratio of 7 over 8. The bathroom can be used by guest. The bathroom connects to the common room, master room, living room.


Make a **kitchen** . The kitchen should be at the south side of the apartment. Make kitchen approx 50 sqft with the aspect ratio of 7 over 4. The kitchen attaches to the common room, balcony, master room, living room.

Can you make a **balcony** ? I would like to place balcony at the south side of the apartment. Can you make balcony around 50 sqft with the aspect ratio of 5 over 2? The balcony is private. The balcony connects to the common room, kitchen, master room.

The **master room** should be considered. The master room should be at the south side of the apartment. Make master room approx 150 sqft with the aspect ratio of 4 over 5. The master room should have an en-suite bathroom. The master room should be next to the bathroom, kitchen, balcony.

It would be great to have a **living room** . Make living room around 650 sqft with the aspect ratio of 1 over 2.

Figure 9: An example of artificially-generated language instructions from the T2D dataset.



Human Instructions:

The north side of this home is not complete without the **balcony**. Access to the approximately 16 sq ft area can be made through the living room or through the common room beside it.

Bathroom 1 is in the eastern section of the home. It is located next to the living room and is approximately 15 sq ft.

The larger of the two, **Bathroom 2**, is approximately 30 sq ft. It is between the master bedroom and common area 2, along the western side of the house.

Common room 1 occupies the northeast corner of the property. At roughly 80 sq ft it is conveniently located next to the balcony.

Common room 2 is nearly 100 sq ft. Occupying the northwest corner, it is easily accessible from the kitchen beside it, or the shared access from the living area.

The **kitchen** is positioned on the north side of the house, between the living room and second common area. It measures about 50 sq ft.

The **living room** is conveniently located in the southeast corner of the home. It spans approximately 250 sq ft while offering access to almost every room in the house.

Located in the southwest corner of the home is the **master bedroom**. This space is approximately 120 sq ft and is positioned next to the living room.

Figure 10: Generated samples from different baselines according to the same human-written language instructions.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
The individual section named Limitations after section 8.
- A2. Did you discuss any potential risks of your work?
The individual section named Ethics Statement
- A3. Do the abstract and introduction summarize the paper’s main claims?
section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section3, Section5

- B1. Did you cite the creators of artifacts you used?
Section3, Section5
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section Ethics Statement, Section Appendix B
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section Ethics Statement
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Section Ethics Statement
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 3
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 3, Section 4

C Did you run computational experiments?

Section 5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
No response.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- ✓ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 5, Section Appendix B
 - ✓ C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 5
 - ✓ C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 5, Section Appendix B
- D** ✓ **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 3
- ✓ D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Section 3, Appendix D
 - ✓ D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Section 3, Appendix D
 - ✓ D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Ethics Statement
 - ✓ D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Section 3
 - ✓ D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Section 3