

BLT: Bidirectional Layout Transformer for Controllable Layout Generation

Xiang Kong^{2*}, Lu Jiang¹, Huiwen Chang¹, Han Zhang¹,
Yuan Hao¹, Haifeng Gong¹, Irfan Essa^{1,3}

¹ Google² LTI, Carnegie Mellon University, ³ Georgia Institute of Technology
xiangk@cs.cmu.edu, lujiang@google.com

Abstract. Creating visual layouts is a critical step in graphic design. Automatic generation of such layouts is essential for scalable and diverse visual designs. To advance conditional layout generation, we introduce BLT, a bidirectional layout transformer. BLT differs from previous work on transformers in adopting non-autoregressive transformers. In training, BLT learns to predict the masked attributes by attending to surrounding attributes in two directions. During inference, BLT first generates a draft layout from the input and then iteratively refines it into a high-quality layout by masking out low-confident attributes. The masks generated in both training and inference are controlled by a new hierarchical sampling policy. We verify the proposed model on six benchmarks of diverse design tasks. Experimental results demonstrate two benefits compared to the state-of-the-art layout transformer models. First, our model empowers layout transformers to fulfill controllable layout generation. Second, it achieves up to 10x speedup in generating a layout at inference time than the layout transformer baseline. Code is released at <https://shawnkx.github.io/blt>.

Keywords: Design, Layout Creation, Transformer, Non-autoregressive.

1 Introduction

Graphic layout dictates the placement and sizing of graphic components, playing a central role in how viewers interact with the information provided [24]. Layout generation is emerging as a new research area with a focus of generating realistic and diverse layouts to facilitate design tasks. Recent works show promising progress for various applications such as graphic user interfaces [2, 18], presentation slides [13], magazines [45, 43], scientific publications [1], commercial advertisements [24, 36, 14], computer-aided design [41], indoor scenes [4], layout representations [28, 42], *etc.*

Previous work explores neural models for layout generation using Generative Adversarial Networks (GANs) [10, 25] or Variational Autoencoder (VAEs) [22, 19, 34, 24]. Currently, layout transformers hold the state-of-the-art performance

* Work done during their research internship at Google.

for layout generation [15, 1]. These transformers represent a layout as a sequence of objects and an object as a (sub)sequence of attributes (See Fig. 1a). Layout transformers predict the attribute sequentially based on previously generated output (*i.e.* autoregressive decoding). Like other vision tasks, by virtue of the powerful self-attention [38], transformer models yield superior quality and diversity than GAN or VAE models for layout generation [15, 1].

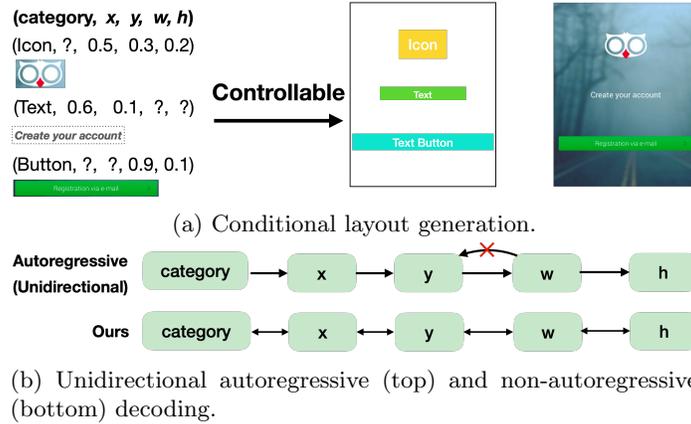


Fig. 1: **(a) Conditional layout generation.** Each object is modeled by 5 attributes ‘category’, ‘ x ’, ‘ y ’, ‘ w ’ (width) and ‘ h ’ (height). In conditional generation, attributes are partially given by the user and the goal is to generate the unknown attributes, *e.g.* putting the icon or button on the canvas. **(b) Illustration of immutable dependency chain in autoregressive decoding.**

Unlike Layout VAE (or GAN) models that are capable of generating layouts considering user requirements, layout transformers, however, have difficulties in conditional generation as a result of an acknowledged limitation discussed in [15] (*c.f.* order of primitives). Fig. 1a illustrates a scenario in which a designer has objects with partially known attributes and hopes to generate the missing attributes. Specifically, each object is modeled by five attributes ‘category’, ‘ x ’, ‘ y ’, ‘ w ’ (width) and ‘ h ’ (height). The designer wants the layout model to 1) place the “icon” and “button” with known sizes onto the canvas (*i.e.* generating x , y from w , h , and ‘category’), and 2) determines the size of the centered “text object” (*i.e.* generating w , h from x , y , and ‘category’).

Such functionality is currently missing in the layout transformers [15, 1] due to *immutable dependency chain*. This is because autoregressive transformers follow a pre-defined generation order of object attributes. As shown in Fig. 1b, attributes must be generated starting from the category c , then x and y , followed by w and h . The dependency chain is immutable *i.e.* it cannot be changed at decoding time. Therefore, autoregressive transformers fail to perform con-

ditional layout generation when the condition disagrees with the pre-defined dependency, *e.g.* generating position y from the known width w in Fig. 1b.

In this work, we introduce Bidirectional Layout Transformer (or BLT) for controllable layout generation. Different from the traditional transformer models [15, 1], BLT enables controllable layout generation where every attribute in the layout can be modified, with high flexibility, based on the user inputs (*c.f.* Fig. 1a). During training, BLT learns to predict the masked attributes by attending to attributes in two directions (*c.f.* Fig. 2a). At inference time, BLT adopts a non-autoregressive decoding algorithm to refine the low-confident attributes iteratively into a high-quality layout (*c.f.* Fig. 2b). We propose a simple hierarchical sampling policy that is used both in training and inference to guide the mask generation over attribute groups.

BLT eliminates a critical limitation in the prior layout transformer models [15, 1] that prevents transformers from performing controllable layout generation. Our model is inspired by the autoregressive work in NLP [9, 11, 5, 12]. However, we find directly applying the non-autoregressive translation models [9, 23] to layout generation only leads to inferior results than the autoregressive baseline. Our novelty lies in the proposed simple yet novel hierarchical sampling policy, which, as substantiated by our experiments in Section 5.4, is essential for high-quality layout generation.

We evaluate the proposed method on six layout datasets under various metrics. These datasets cover representative design applications for graphic user interface [2], magazines [45] and publications [46], commercial ads [24], natural scenes [27] and home decoration [8]. Experiments demonstrate two benefits to several strong baseline models [15, 1, 19, 19]. First, our model empowers transformers to fulfill controllable layout generation and thereby outperforms the previous conditional models based on VAE (*i.e.*, LayoutVAE [19] and NDN [24]). Even though our model is not designed for unconditional layout generation, it achieves quality on-par with the state-of-the-art. Second, our new method reduces the time complexity in [15, 1] while achieving 4x-10x speedups in layout generation.

To summarize, we make the following contributions:

1. We address a critical limitation in state-of-the-art layout transformers [15, 1] and hence empower transformers to fulfill controllable layout generation.
2. Though our idea is inspired by the autoregressive work in NLP [9, 11, 5, 12], a novel hierarchical mask sampling policy is introduced in training and decoding, which is essential for high-quality layout generation.
3. Extensive experiments validate that our method performs favorably against state-of-the-art models in terms of realism, alignment, and semantic relevance on six diverse layout benchmarks.

2 Related Work

Layout synthesis: Recently, automatic generation of high-quality and realistic layouts has fueled increasing interest. Unlike early work [32, 31, 33, 35, 44, 29, 7,

40, 6, 39], recent data-driven methods rely on deep generative models such as GANs [10] and VAE [22]. For example, LayoutGAN [25] uses a GANs-based framework to synthesize semantic and geometric properties for scene elements. During inference time, LayoutGAN generates layouts from the Gaussian noise. Afterwards, LayoutGAN is extended to attribute-conditioned design tasks [26]. LayoutVAE [19] introduces two conditional VAEs. The first aims to learn the distribution of category counts which will be used during layout generation. The second produces layouts conditioning on the number and category of objects generated from the first VAE or ground-truth data. Recently, various VAE models are proposed [34, 20, 24]. Among them, Neural Design Networks (NDN) [24] is a competitive VAEs-based model for conditional layout generation, which focuses on modeling the asset relations and constraints by graph convolution. Our work is different from LayoutVAE and NDN in modeling layout and user inputs by the transformer, which, as shown in Table 1, perform more favorably thanks to the transformer architecture. Our finding is consistent with [1] where Arroyo *et al.* find VAEs underperforming transformers for unconditional layout generation [1].

Currently, the state-of-the-art for layout generation is held by the transformer models [38]. In particular, [15] employs the standard autoregressive Transformer decoder with unidirectional attention. They find out that self-attention is able to explicitly learn relationships between objects in the layout, resulting in superior quality compared to prior work. Furthermore, to increase the diversity of generated layout, [1] incorporates the standard autoregressive Transformer decoder into a VAE framework and [30] employs multi-choice prediction and winner-takes-all loss. Despite the superior performance, this work addresses a critical limitation acknowledged in [15] that prevents transformers from performing controllable layout generation. Following LayoutGAN [25], [20] proposes a Transformer based layout GAN model, LayoutGAN++. In this framework, the input is a set of asset labels and randomly generated code and the output is the location and size of these asset. Different from the LayoutGAN++, the input to our proposed model is more flexible and can support unconditional generation and various types of conditional generation tasks.

Bidirectional transformer and non-autoregressive decoding: The classic Transformer [38] decoder uses the unidirectional self-attention mechanism to generate the sequence token-by-token from left to right, leaving the right-to-left contexts unexploited. Several NLP works [9, 23, 37] are proposed to investigate language generation tasks by non-autoregressive generation with bidirectional Transformers, which allow representations to attend in both directions [3]. However, non-autoregressive decoding process leads to an apparent performance drop compared to the autoregressive decoding algorithm [11, 12, 9]. In this work, we find that applying the non-autoregressive NLP model [9] to layout generation also leads to inferior results than the autoregressive baseline. To this end, we propose a simple yet effective hierarchical sampling policy which is essential for high-quality layout generation.

3 Problem Formulation

Following [15], we use 5 attributes to describe an object, *i.e.*, (c, x, y, w, h) , in which the first element $c \in C$ is the object category such as the logo or button, and the remainder details the bounding box information *i.e.* the center location $(x, y) \in \mathbb{R}^2$ and the width and height $(w, h) \in \mathbb{R}^2$. Furthermore, float values in bounding box information is discretized using 8-bit uniform quantization. For instance, the x -coordinate after the quantization becomes $\{x|x \in \mathbb{Z}, 0 \leq x \leq 31\}$. A layout l of K assets is hence denoted as a flattened sequence of integer indices:

$$l = [\langle \text{bos} \rangle, c_1, x_1, y_1, w_1, h_1, c_2 \cdots, h_K, \langle \text{eos} \rangle] \quad (1)$$

where $\langle \text{bos} \rangle$ and $\langle \text{eos} \rangle$ are special tokens to denote the start and the end of sequence. We use a shared vocabulary and represent each element in l as an integer index or equivalently as a one-hot vector with the same length. It is trivial to extend the attribute dimension to model more complex layouts.

Issues To train the model, prior work [15, 1] estimates the joint likelihood of observing a layout as $p(l) = \prod_{i=1}^{|l|} p(l_i | l_{1:i})$.

During training, an autoregressive Transformer model is learned to maximize the likelihood using ground-truth attribute as input (*i.e.* teacher forcing). At inference time, the transformer model predicts the attribute sequentially based on previously generated output (*i.e.* autoregressive decoding), starting from the begin-of-sequence or $\langle \text{bos} \rangle$ token until yielding the end-of-sequence token $\langle \text{eos} \rangle$. The generation must follow a fixed conditional dependency. For example, Eq. (1) defines an immutable generation order $x \rightarrow y \rightarrow w \rightarrow h$. And in order to generate the height h for an object, one must know its x - y coordinates and width w .

There are two issues with autoregressive decoding for the conditional generation. First, it is infeasible to process user conditions that differ from the dependency order used in training. For instance, the model using Eq. (1) is not able to generate x - y coordinates from width and height, which corresponds to a practical example of placing an object with given size. This issue is exacerbated by complex layouts that require more attributes to represent an object. Second, the autoregressive inference is not parallelizable, rendering it inefficient for the dense layout with a large number of objects or attributes.

4 Approach

Our goal is to design a transformer model for controllable layout generation. We propose a method to learn non-autoregressive transformers. Unlike existing layout transformers [15, 1], the new layout transformer is bidirectional and can generate all attributes simultaneously in parallel, which allows not only for flexible conditional generation but also more efficient inference. In this section, we first discuss the model and training objective; then detail a novel hierarchical sampling policy for training and parallel decoding.

4.1 Model and Training

The BLT backbone is the multi-layer bidirectional Transformer encoder [38] as shown in Fig. 2. We use the identical architecture as in the existing autoregressive layout transformers [1, 15] but a bidirectional attention mechanism.

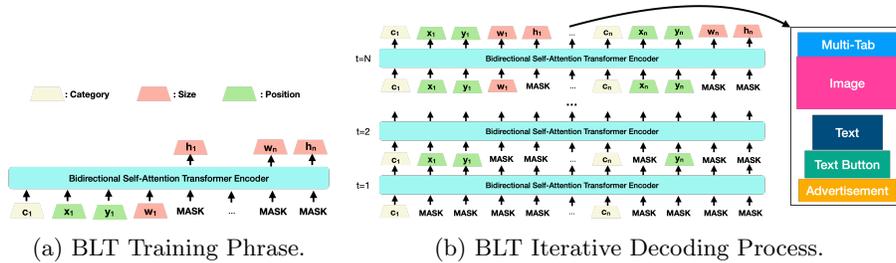


Fig. 2: The training (left) and decoding (right) stages of the proposed Bidirectional Layout Transformer (BLT).

Inspired by BERT [3], during training, we randomly select a subset of attributes in the input sequence, replace them with a special “[MASK]” token, and optimize the model to predict the masked attributes. For a layout sequence l , let \mathcal{M} denote a set of masked positions. Replacing attributes in l with “[MASK]” at \mathcal{M} yields the masked sequence $l^{\mathcal{M}}$.

Given a layout set \mathcal{D} , the training objective is to minimize the negative log-likelihood of the masked attributes:

$$\mathcal{L}_{mask} = - \mathbb{E}_{l \in \mathcal{D}} \left[\sum_{i \in \mathcal{M}} \log p(l_i | l^{\mathcal{M}}) \right], \quad (2)$$

The masking strategy greatly affects the quality of the masked language model [3]. BERT [3] applies random masking with a fixed ratio where a constant 15% masks are randomly generated for each input. Similarly, we find masking strategy is important for layout generation, but the random masking used in BERT does not work well. We propose to use a new sampling policy. Specifically, we divide the attributes of an object into semantic groups, *e.g.* Fig. 2 showing 3 groups: category, position, and size. First, we randomly select a semantic group. Next, we dynamically sample the number of masked tokens from a uniform distribution between one and the number of attributes belonging to the chosen group, and then randomly mask that number of tokens in the selected group. As such, it is guaranteed that the model only predicts attributes of the same semantic meaning each time. Therefore, given the hierarchical relations between these groups, we call this method as the hierarchical sampling. We will discuss how to apply the hierarchical sampling policy to decoding in the next subsection.

4.2 Parallel Decoding by Iterative Refinement

In BLT, all attributes in the layout are generated simultaneously in parallel. Since generating layouts in a single pass is challenging [9], we employ a parallel language model. The core idea is to generate a layout iteratively in a small number of steps where parallel decoding is applied at each.

Algorithm 1 Decoding by Iterative Attribute Refinement

Require: Sequence l with partially-known attributes. Constant T for the number of iterations.

- 1: **for** g in $[C, S, P]$ **do** ▷ Loop over semantic group
- 2: **for** $i \leftarrow 1$ to $T/3$ **do**
- 3: $p, l^i = \text{BLT}(l)$
- 4: $\gamma_i = \frac{T-3i}{T}$ ▷ Compute mask ratio
- 5: $n_i = \lfloor \gamma_i \times |g| \rfloor$ ▷ $|g|$: # attributes in g
- 6: $\mathcal{M} = \arg_{k=n_i} \text{top-k}(-p)$ ▷ Get mask indices
- 7: Obtain l by masking l^i with respect to \mathcal{M}
- 8: **end for**
- 9: **end for**
- 10: **return** l

Algorithm 1 presents the non-autoregressive decoding algorithm. The procedure is also illustrated in Fig. 2b. The input to the decoding algorithm is a mixture sequence of known and unknown attributes, where the known attributes are given by the user inputs, and the model aims at generating the unknown attributes denoted by the [MASK] token. Like in training, we employ the hierarchical sampling policy to generate attributes of three semantic groups: category (C), size (S), and position (P). For each iteration, one group of attributes is sampled. In Step 3 of Algorithm 1, the model makes parallel predictions for all unknown attributes, where p denotes the prediction scores. Step 6 samples the attributes that belong to the selected group and have the lowest prediction scores. Finally, it masks low-confident attributes on which the model has doubts. The prediction probabilities from the softmax layer are used as the confidence scores. These masked attributes will be re-predicted in the next iteration of decoding conditioning on all other ascertained attributes so far. The masking ratio calculated in Step 4 decreases with the number of iterations. This process will repeat T times until all attributes of all objects are generated (*c.f.* Fig. 2b).

Our model is inspired by the autoregressive models in NLP [9, 23]. It is noteworthy that Algorithm 1 differs from the non-autoregressive NLP models [9, 23] in the proposed hierarchical sampling. This paper finds applying [9, 23] to layout generation only leads to inferior results than our autoregressive baseline. We hypothesize that it is because layout attributes, unlike natural language, have apparent structures, and the non-autoregressive models designed for word sequences [9, 23] might not sufficiently capture the complex correlation between

layout attributes. We empirically demonstrate that Algorithm 1 outperforms our non-autoregressive NLP baselines in Section 5.4.

Algorithm 1 can be extended to unconditional generation. In this case, the input is a layout sequence of only “[MASK]” tokens, and the same algorithm is used to generate all attributes in the layout. Unlike conditional generation, we need to know the sequence length in advance, *i.e.* the number of objects to be generated. Here, we can use the prior distribution obtained on the training dataset. During decoding, we obtain the number of objects through sampling from this prior distribution.

5 Experimental Results

This section verifies the proposed method on six diverse layout benchmarks under various metrics to examine realism, alignment, and semantic relevance. The results show our model performs favorably against the strong baselines and achieves a 4x-10x speedup than autoregressive decoding in layout generation.

5.1 Setups

Datasets We employ six datasets that cover representative graphic design applications. *RICO* [2] is a dataset of user interface designs for mobile applications. It contains 91K entries with 27 object categories (button, toolbar, list item, *etc.*). *PubLayNet* [46] contains 330K examples of machine annotated scientific documents crawled from the Internet. Its objects come from 5 categories: text, title, figure, list, and table. *Magazine* [45] contains 4K images of magazine pages and six categories (texts, images, headlines, over-image texts, over-image headlines, backgrounds). *Image Ads* [24] is the commercial ads dataset with layout annotation detailed in [24]. *COCO* [27] contains ~ 100 K images of natural scenes. We follow [1] to use the Stuff variant, which contains 80 things and 91 stuff categories, after removing small bounding boxes ($\leq 2\%$ image area), and instances tagged as “iscrowd”. *3D-FRONT* [8] is a repository of professionally designed indoor layouts. It contains around 7K room layouts with objects belonging to 37 categories, *e.g.*, the table and bed. Different from previous datasets, objects in 3D-FRONT are represented by 3D bounding boxes. The maximum number of objects in our experiments is 25 in the RICO dataset and 22 in the PubLayNet dataset.

Evaluation metrics We employ five common metrics in the literature as well as a user study to validate the proposed method’s effectiveness. Specifically, *IOU* measures the intersection over the union between the generated bounding boxes. We use an improved perceptual IOU (see more discussions in the Appendix). *Overlap* [25] measures the total overlapping area between any pair of bounding boxes inside the layout. *Alignment* [24] computes an alignment loss with the intuition that objects in graphic design are often aligned either by center or edge. *FID* [16] measures the distributional distance of the generated layout to

the real layout. Following [24], we compute FID using a binary layout classifier to discriminate real layouts. We employ a 2-layer Transformer to train the classifier. Notice that the lower, the better for all IOU, Overlap, Alignment, and FID.

The above metrics ignore the input condition. For conditional generation, we employ a metric called *Similarity* [34] and a user study, where the former compares the generated layout with the ground-truth layout under the same input. Following [34], *DocSim* is used to calculate the similarity between two layouts. The user study is used to further evaluate human’s perception about the conditionally-generated layouts.

Generation settings We examine three layout generation scenarios (2 conditional and 1 unconditional).

- Conditional on **Category**: only object categories are given by users. The model needs to predict the size and position of each object.
- Conditional on **Category + Size**: the object category and size are specified. The model needs to predict the positions, *i.e.* placing objects on the canvas.
- **Unconditional** Generation: no information is provided by users. Prior layout transformer work focuses on this setting.

In unconditional generation, the model generates 1K samples from the random seed. The test split of each dataset is used for conditional generation.

Implementation details The model is trained for five trials with random initialization and the averaged metrics with standard deviations are reported. All models including ours have the same configuration, *i.e.*, 4 layers, 8 attention heads, 512 embedding dimensions and 2,048 hidden dimensions. Adam optimizer [21] with $\beta_1 = 0.9$ and $\beta_2 = 0.98$ is used. Models are trained on 2×2 TPU devices with batch size 64. For conditional generation, we randomly shuffle objects in the layout. For unconditional generation, to improve diversity, we use the nucleus sampling [17] with $p = 0.9$ for the baseline Transformers and the top-k sampling ($k = 5$) for our model. Greedy decoding method is used for conditional generation. Please refer to the Appendix for more detailed hyperparameter configurations.

5.2 Quantitative Comparison

Conditional generation The results are shown in Table 1 and Table 2. State-of-the-art layout transformers are compared *i.e.* **LayoutTransformer (Trans.)** [15] and **Variational Transformer Network (VTN)** [1]. In addition, two representative VAEs for conditional generation: **LayoutVAE (L-VAE)** [19] and **Neural Design Network (NDN)** [24] are also compared on the large datasets of RICO and PubLayNet. Two conditional generation tasks are examined *i.e.* Conditioned on Category and Conditioned on Category + Size (Column “+ Size”). The same model is used for both conditional cases and “-” indicates the baseline models fail to process the condition “Category + Size”. The results

RICO							
Model	Conditioned on Category					+ Size	
	IOU↓	Overlap↓	Alignment↓	FID↓	Sim.↑	Sim.↑	FID↓
L-VAE [19]	0.41±1.5%	0.39±2.3%	0.38±1.9%	122±19	0.13±1.5%	0.19	76
NDN [24]	0.37±1.7%	0.36±1.9%	0.41±1.6%	97±21	0.15±2.3%	0.21	63
Trans. [15]	0.31±0.2%	0.33±0.8%	0.30±0.8%	76±24	0.20±0.1%	-	-
VTN [1]	0.30 ±0.1%	0.30±0.3%	0.32±0.9%	82±23	0.20±0.1%	-	-
Ours	0.30 ±0.4%	0.23 ±0.2%	0.20 ±1.1%	70 ±29	0.21 ±0.2%	0.30	26
PubLayNet							
Model	Conditioned on Category					+ Size	
	IOU↓	Overlap↓	Alignment↓	FID↓	Sim.↑	Sim.↑	FID↓
L-VAE	0.45±1.3%	0.15±0.9%	0.37±0.7%	513±26	0.07±0.3%	0.09	239
NDN	0.34±1.8%	0.12±0.8%	0.39±0.4%	425±37	0.06±0.3%	0.09	178
Trans.	0.19 ±0.3%	0.06±0.3%	0.33±0.3%	127 ±29	0.11 ±0.1%	-	-
VTN	0.21±0.6%	0.06±0.2%	0.33±0.4%	159±21	0.10±0.1%	-	-
Ours	0.19 ±0.2%	0.04 ±0.1%	0.25 ±0.7%	134±24	0.11 ±0.2%	0.18	87

Table 1: Conditional layout generation on two settings (Category and Category+Size) on the large datasets of RICO and PubLayNet.

COCO				Magazine			
Model	Conditioned on Category		+ Size	Model	Conditioned on Category		+ Size
	IOU↓	Sim.↑	Sim.↑		IOU↓	Sim.↑	Sim.↑
Trans. [15]	0.60±0.4%	0.20±0.2%	-	Trans.	0.20±0.8%	0.15±0.3%	-
VTN [1]	0.63±0.4%	0.22±0.1%	-	VTN	0.18 ±1.8%	0.15±0.9%	-
Ours	0.43 ±0.5%	0.24 ±0.1%	0.44	Ours	0.18 ±0.6%	0.18 ±0.4%	0.27
Ads				3D-FRONT			
Model	Conditioned on Category		+ Size	Model	Conditioned on Category		+ size
	IOU↓	Sim.↑	Sim.↑		Sim.↑	Sim.↑	Sim.↑
Trans. [15]	0.19±0.1%	0.30±0.1%	-	Trans.	0.04±0.7%	-	-
VTN [1]	0.18±0.2%	0.30±0.1%	-	VTN	0.04±0.4%	-	-
Ours	0.10 ±0.4%	0.31 ±0.1%	0.41	Ours	0.06 ±0.7%	0.10	0.10

Table 2: Category (+ Size) conditional layout generation on four datasets.

are aggregated on independently trained models, where the mean and standard deviation over five trails are reported.

Because of the non-autoregressive decoding, our model is able to conduct conditional generation on category + size while the baseline transformer models (Trans. [15] and VTN [1]) fail. Our model also outperforms VAE-based conditional layout models (L-VAE [19] and NDN [24]) across all metrics in Table 1 by statistically significant margins. This result is consistent with the prior finding in [1] that transformers outperform VAEs for unconditional layout generation.

Unconditional Generation Although our model is not designed for this task, we compare it to the models [19, 15, 1] on unconditional layout generation. From Table 3, our model outperforms LayoutVAE [19] and achieves comparable performance with two autoregressive transformers (Trans. [15] and VTN [1]).

User study We conduct user studies on RICO and PubLayNet to assess generated layouts for conditional generation. We randomly select 50 generated layouts under both conditional settings specified in Section 5.1 and collect their golden

Methods	RICO			PubLayNet			COCO		
	IOU↓	Overlap↓	Alignment↓	IOU↓	Overlap↓	Alignment↓	IOU↓	Overlap↓	Alignment↓
LayoutVAE [19]	0.193	0.400	0.416	0.171	0.321	0.472	0.325	2.819	0.246
Trans. [15]	0.086	0.145	0.366	0.039	0.006	0.361	0.194	1.709	0.334
VTN [1]	0.115	0.165	0.373	0.031	0.017	0.347	0.197	2.384	0.330
Ours	0.127	0.102	0.342	0.048	0.012	0.337	0.227	1.452	0.311

Table 3: Unconditional layout generation comparison to the state-of-the-art on three benchmarks. Results of baselines are cited from [1] and our scores are calculated following the same method described in [1].

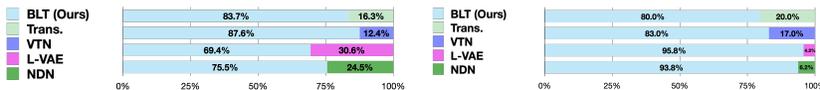


Fig. 3: We conduct a user study to compare the quality of generated samples from our model and baseline models on RICO (left) and PubLayNet (Right).

layouts. For each trial, we present Amazon Mechanical Turk workers two layouts generated by different methods along with the golden layout for reference, and ask “which layout is more similar to the true reference layout?”. There are 75 unique workers participating in the study. Qualitative comparison is shown in the Appendix. The results, which are plotted in Fig. 3, verify that the proposed model outperforms all baseline models for conditional layout generation.

5.3 Qualitative Result

We show some generated layouts, along with the rendered examples for visualization, in Fig. 4. The setting is conditional generation on category and size for three design applications, including the mobile UI interface, scientific paper, and magazine. We observe that our method yields reasonable layouts, which facilitates generating high-quality outputs by rendering.

Next, we explore the home design task on the 3D-Front dataset [8]. The goal is to place the furniture with the user-given category and length, height, and width information. Examples are shown in Fig. 5. Unlike previous tasks, the model needs to predict the position of the 3D bounding box. The result suggests the feasibility of our method extending to 3D object attributes. The low similarity score on this dataset indicates that housing design layout is still a challenging task that needs future research.

To further understand what relationships between attributes BLT has learned, we visualize the patterns in how our model’s attention heads behave. We choose a simple layout with two objects and mask their positions (x, y) . The model needs to predict these masked attributes from other known attributes. Examples of heads exhibiting these patterns are shown in Fig. 6. We use $\langle \text{layer} \rangle$ - $\langle \text{head number} \rangle$ to denote a particular attention head. For the head 0-2, $[\text{MASK}]_{y_2}$ specializes to attending on its category (c2) and especially, its height information (h2), which is reasonable because y -coordinate is highly relevant to the height of

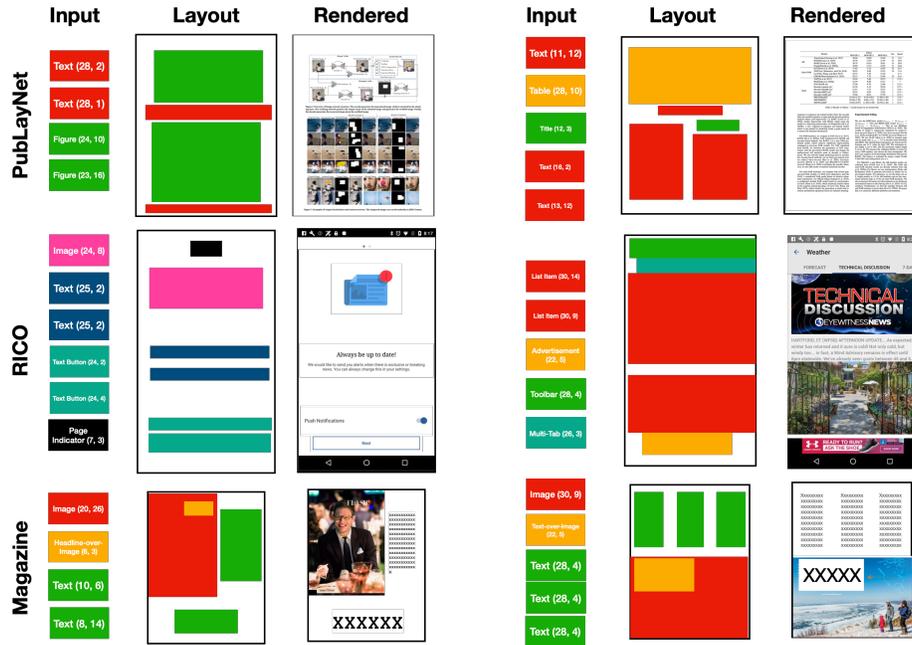


Fig. 4: Conditional layout generation for scientific papers, user interface, and magazine. The user inputs are the object category and their size (width, height). We present the rendered examples constructed based on the generated layouts.

the object. Furthermore, for heads 2-4 and 3-2, $[\text{MASK}]_{x_1}$ focuses on the width of not only the first but the second object as well. Given this contextual information from other objects, the model is able to predict the position of these objects more accurately. The similar pattern is also found at head 3-2 for $[\text{MASK}]_{x_2}$.

5.4 Ablation Study

Decoding speed We compare the inference speed of our model and the autoregressive transformer models [15, 1]. Specifically, all models generate 1,000 layouts with batch size 1 on a single GPU. The average decoding time in millisecond is reported. The result is shown in Fig.7, where the x -axis denotes the number of objects in the layout. It shows that autoregressive decoding time grows with $\#$ objects. On the contrary, the decoding speed of the proposed model appears not affected by $\#$ objects. The speed advantage becomes evident when producing dense layouts. For example, our fastest model obtains a 4x speedup when generating around 10 objects and a 10x speedup for 20 objects.

Hierarchical sampling This experiment investigates the effectiveness of the hierarchical sampling strategy used in training (Section 4.1) and non-autoregressive decoding (Section 4.2). Specifically, we compare with the non-autoregressive



Fig. 5: 3D-FRONT sample layouts.

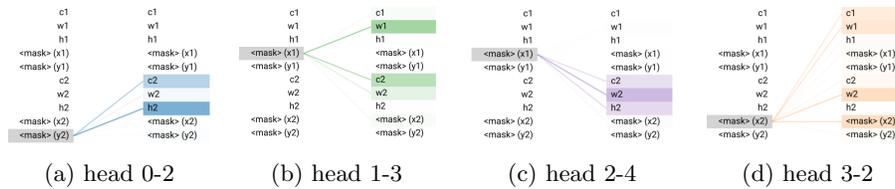


Fig. 6: Examples of attention heads exhibiting the patterns for masked tokens. The darkness of a line indicates the strength of the attention weight (some attention weights are so low they are invisible). We use (layer)-<head number> to denote a particular attention head.

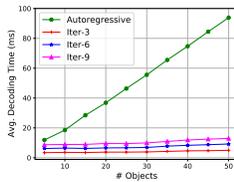


Fig. 7: Decoding speed versus number of generated assets. ‘Autoregressive’ denote the autoregressive Transformer-based model [15]. ‘Iter-*’ shows the proposed model with various number of iterations.

method [9] in NLP on the large datasets of RICO and PubLayNet in Table 4. Autoregressive transformer results [15] are also included for reference but notice that autoregressive methods [15] have difficulties with conditional generation.

The results in Table 4 show that the non-autoregressive baseline yields inferior results than the autoregressive one. We hypothesize that it is because the non-autoregressive models designed for word sequences [9, 23] might not sufficiently capture the apparently-structural correlation between layout attributes. The proposed method with hierarchical sampling significantly outperforms the

RICO		IoU ↓	Overlap↓	Align.↓	FID↓	Sim. ↑
Autoregressive [15]		0.30	0.33	0.30	76	0.20
Non-autoregressive [9]		0.37	0.33	0.24	104	0.17
Non-autoregressive + HSP (Ours)		0.30	0.23	0.20	70	0.21
PubLayNet		IoU ↓	Overlap↓	Align.↓	FID↓	Sim. ↑
Autoregressive [15]		0.19	0.06	0.33	127	0.11
Non-autoregressive [9]		0.16	0.12	0.32	217	0.09
Non-autoregressive + HSP (Ours)		0.19	0.04	0.25	134	0.11

Table 4: Comparison with the non-autoregressive method [9] in NLP on the RICO and PubLayNet datasets. Autoregressive results are included for reference. HSP denotes hierarchical sampling policy proposed in this work.

Order	IoU ↓	Overlap↓	Alignment↓
C→S→P	0.127	0.102	0.342
C→P→S	0.129	0.107	0.344
S→C→P	0.147	0.109	0.351
S→P→C	0.162	0.121	0.357

Table 5: Layout generation results with different iteration group orders on the RICO dataset. C, S, and P denote category, size, and position attribute groups, respectively.

non-autoregressive NLP baseline, which suggests the necessity of the proposed hierarchical sampling strategy. We also explore the effect of hierarchical sampling order. In Algorithm 1, we prespecify an order of attribute groups, *i.e.*, Category (C) → Size (S) → Position (P). Here, more orders are explored in Table 5. It seems better to first generate the category and afterward determine either location or size.

6 Conclusion and Future Work

We present BLT, a bidirectional layout transformer capable of empowering the transformer-based models to carry out conditional and controllable layout generation. Moreover, we propose a hierarchical sampling policy during BLT training and inference processes which has been shown to be essential for producing high-quality layouts. Thanks to the high computation parallelism, BLT achieves 4-10 times speedup compared to the autoregressive transformer baselines during inference. Experiments on six benchmarks show the effectiveness and flexibility of BLT. A limitation of our work is content-agnostic generation. We leave this out to have a fair and lateral comparison to our baselines which do not use visual information either. In the future, we will explore using rich visual information.

Acknowledgement

The authors would like to thank all anonymous reviewers and area chairs for helpful comments.

References

1. Arroyo, D.M., Postels, J., Tombari, F.: Variational transformer networks for layout generation. In: CVPR (2021)
2. Deka, B., Huang, Z., Franzen, C., Hibsichman, J., Afergan, D., Li, Y., Nichols, J., Kumar, R.: Rico: A mobile app dataset for building data-driven design applications. In: UIST (2017)
3. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAAACL (2019)
4. Di, X., Yu, P.: Multi-agent reinforcement learning of 3d furniture layout simulation in indoor graphics scenes. arXiv preprint arXiv:2102.09137 (2021)
5. Donahue, J., Dieleman, S., Binkowski, M., Elsen, E., Simonyan, K.: End-to-end adversarial text-to-speech. In: ICLR (2020)
6. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 605–613 (2017)
7. Fisher, M., Ritchie, D., Savva, M., Funkhouser, T., Hanrahan, P.: Example-based synthesis of 3d object arrangements. ACM Transactions on Graphics (TOG) **31**(6), 1–11 (2012)
8. Fu, H., Cai, B., Gao, L., Zhang, L., Li, C., Zeng, Q., Sun, C., Fei, Y., Zheng, Y., Li, Y., Liu, Y., Liu, P., Ma, L., Weng, L., Hu, X., Ma, X., Qian, Q., Jia, R., Zhao, B., Zhang, H.: 3d-front: 3d furnished rooms with layouts and semantics. arXiv preprint arXiv:2011.09127 (2020)
9. Ghazvininejad, M., Levy, O., Liu, Y., Zettlemoyer, L.: Mask-predict: Parallel decoding of conditional masked language models. In: EMNLP-IJCNLP (2019)
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. NeurIPS **27** (2014)
11. Gu, J., Bradbury, J., Xiong, C., Li, V.O., Socher, R.: Non-autoregressive neural machine translation. In: ICLR (2018)
12. Gu, J., Kong, X.: Fully non-autoregressive neural machine translation: Tricks of the trade. In: Findings of ACL-IJCNLP (2021)
13. Guo, M., Huang, D., Xie, X.: The layout generation algorithm of graphic design based on transformer-cvae. arXiv preprint arXiv:2110.06794 (2021)
14. Guo, S., Jin, Z., Sun, F., Li, J., Li, Z., Shi, Y., Cao, N.: Vinci: an intelligent graphic design system for generating advertising posters. In: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. pp. 1–17 (2021)
15. Gupta, K., Lazarow, J., Achille, A., Davis, L.S., Mahadevan, V., Shrivastava, A.: Layouttransformer: Layout generation and completion with self-attention. In: ICCV (2021)
16. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. NeurIPS (2017)
17. Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. In: ICLR (2019)
18. Jiang, Z., Sun, S., Zhu, J., Lou, J.G., Zhang, D.: Coarse-to-fine generative modeling for graphic layouts (2022)
19. Jyothi, A.A., Durand, T., He, J., Sigal, L., Mori, G.: Layoutvae: Stochastic scene layout generation from a label set. In: CVPR (2019)
20. Kikuchi, K., Simo-Serra, E., Otani, M., Yamaguchi, K.: Constrained graphic layout generation via latent optimization. In: MM (2021)

21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
22. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
23. Kong, X., Zhang, Z., Hovy, E.: Incorporating a local translation mechanism into non-autoregressive translation. arXiv preprint arXiv:2011.06132 (2020)
24. Lee, H.Y., Jiang, L., Essa, I., Le, P.B., Gong, H., Yang, M.H., Yang, W.: Neural design network: Graphic layout generation with constraints. In: ECCV (2020)
25. Li, J., Yang, J., Hertzmann, A., Zhang, J., Xu, T.: Layoutgan: Generating graphic layouts with wireframe discriminators. In: ICLR (2018)
26. Li, J., Yang, J., Zhang, J., Liu, C., Wang, C., Xu, T.: Attribute-conditioned layout gan for automatic graphic design. *IEEE Transactions on Visualization and Computer Graphics* **27**(10), 4039–4048 (2020)
27. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
28. Manandhar, D., Ruta, D., Collomosse, J.: Learning structural similarity of user interface layouts using graph networks. In: European Conference on Computer Vision. pp. 730–746. Springer (2020)
29. Merrell, P., Schkufza, E., Li, Z., Agrawala, M., Koltun, V.: Interactive furniture layout using interior design guidelines. *ACM transactions on graphics (TOG)* **30**(4), 1–10 (2011)
30. Nguyen, D.D., Nepal, S., Kanhere, S.S.: Diverse multimedia layout generation with multi choice learning. In: MM (2021)
31. O’Donovan, P., Libeks, J., Agarwala, A., Hertzmann, A.: Exploratory font selection using crowdsourced attributes. *ACM Transactions on Graphics (TOG)* **33**(4), 1–9 (2014)
32. O’Donovan, P., Agarwala, A., Hertzmann, A.: Learning layouts for single-pagegraphic designs. *IEEE transactions on visualization and computer graphics* **20**(8), 1200–1213 (2014)
33. Pang, X., Cao, Y., Lau, R.W., Chan, A.B.: Directing user attention via visual flow on web designs. *ACM Transactions on Graphics (TOG)* **35**(6), 1–11 (2016)
34. Patil, A.G., Ben-Eliezer, O., Perel, O., Averbuch-Elor, H.: Read: Recursive autoencoders for document layout generation. In: CVPR Workshops (2020)
35. Qi, S., Zhu, Y., Huang, S., Jiang, C., Zhu, S.C.: Human-centric indoor scene synthesis using stochastic grammar. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5899–5908 (2018)
36. Qian, C., Sun, S., Cui, W., Lou, J.G., Zhang, H., Zhang, D.: Retrieve-then-adapt: Example-based automatic generation for proportion-related infographics. *IEEE TVCG* **27**(2), 443–452 (2020)
37. Stern, M., Chan, W., Kiros, J., Uszkoreit, J.: Insertion transformer: Flexible sequence generation via insertion operations. In: ICML. pp. 5976–5985. PMLR (2019)
38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
39. Wang, K., Lin, Y.A., Weissmann, B., Savva, M., Chang, A.X., Ritchie, D.: Planit: Planning and instantiating indoor scenes with relation graph and spatial prior networks. *ACM Transactions on Graphics (TOG)* **38**(4), 1–15 (2019)
40. Wang, K., Savva, M., Chang, A.X., Ritchie, D.: Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)* **37**(4), 1–14 (2018)
41. Willis, K.D., Jayaraman, P.K., Lambourne, J.G., Chu, H., Pu, Y.: Engineering sketch generation for computer-aided design. In: CVPR (2021)

42. Xie, Y., Huang, D., Wang, J., Lin, C.Y.: Canvasemb: Learning layout representation with large-scale pre-training for graphic design. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 4100–4108 (2021)
43. Yamaguchi, K.: Canvasvae: Learning to generate vector graphic documents. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5481–5489 (2021)
44. Yu, L.F., Yeung, S.K., Tang, C.K., Terzopoulos, D., Chan, T.F., Osher, S.J.: Make it home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics (TOG)-Proceedings of ACM SIGGRAPH 2011*, v. 30,(4), July 2011, article no. 86 **30**(4) (2011)
45. Zheng, X., Qiao, X., Cao, Y., Lau, R.W.: Content-aware generative modeling of graphic design layouts. *TOG* **38**(4), 1–15 (2019)
46. Zhong, X., Tang, J., Yepes, A.J.: Publaynet: largest dataset ever for document layout analysis. In: *IEEE ICDAR* (2019)