

# LayoutDiffusion: Improving Graphic Layout Generation by Discrete Diffusion Probabilistic Models

Junyi Zhang<sup>1\*</sup>   Jiaqi Guo<sup>2</sup>   Shizhao Sun<sup>2</sup>   Jian-Guang Lou<sup>2</sup>   Dongmei Zhang<sup>2</sup>  
<sup>1</sup>Shanghai Jiao Tong University   <sup>2</sup>Microsoft Research Asia  
 junyizhang@sjtu.edu.cn   {jiaqiguo, shizsu, jlou, dongmeiz}@microsoft.com

## Abstract

Creating graphic layouts is a fundamental step in graphic designs. In this work, we present a novel generative model named LayoutDiffusion for automatic layout generation. As layout is typically represented as a sequence of discrete tokens, LayoutDiffusion models layout generation as a discrete denoising diffusion process. It learns to reverse a mild forward process, in which layouts become increasingly chaotic with the growth of forward steps and layouts in the neighboring steps do not differ too much. Designing such a mild forward process is however very challenging as layout has both categorical attributes and ordinal attributes. To tackle the challenge, we summarize three critical factors for achieving a mild forward process for the layout, i.e., legality, coordinate proximity and type disruption. Based on the factors, we propose a block-wise transition matrix coupled with a piece-wise linear noise schedule. Experiments on RICO and PubLayNet datasets show that LayoutDiffusion outperforms state-of-the-art approaches significantly. Moreover, it enables two conditional layout generation tasks in a plug-and-play manner without re-training and achieves better performance than existing methods. Project page: <https://layoutdiffusion.github.io>.

## 1. Introduction

Graphic layout, i.e., the sizes and positions of elements, is important to the interaction between the viewer and the information. Recently, layout generation attracts growing research interest. Leading approaches [11, 16, 17, 20] often represent a layout as a sequence of elements and leverage Transformer [40] to model element relationships. As the placement of one element could depend on any part of a layout, *global context modeling* plays a critical role in layout generation. However, there is no satisfactory solution to it. Some studies simply consider biased context [1, 11, 16, 17]. They generate layout sequences autoregressively, where the generation order for elements is pre-

\*Work done during an internship at Microsoft Research Asia.



Figure 1. Comparison of different forward corruption processes.

We sample the layouts at the timesteps 0, 1/6, 2/6, 3/6, 4/6, 5/6, and 1 of the total timestep. The blank page is used when the format of the layout sequence is destroyed.

defined and the placement of one element only depends on a certain part of layout. A few other studies try to utilize global context by non-autoregressive generation [22]. Unfortunately, they fail to improve the generation quality significantly since it is too challenging to generate a sequence in a single pass [7].

Meanwhile, the emerging diffusion probabilistic model (DDPM) [13, 38] achieves amazing performance on many generation tasks [12, 21, 33–36, 44]. It consists of multiple rounds, each of which gradually denoises the latent variables towards the desired data distribution. This sort of process seems to be a promising solution to layout generation. First, the layout generated in the last round could serve as the global context for the generation in the next round. Second, by multiple rounds of denoising, a layout could be refined iteratively, overcoming the challenge of single-pass generation from non-autoregressive models.

To this end, we propose *LayoutDiffusion* to improve graphic layout generation. As a layout is represented as a sequence of discrete tokens [11, 16, 20], we formulate layout generation as a discrete diffusion process. Roughly speaking, it samples a layout by reversing a forward process. The forward process corrupts the real data into a sequence of increasingly noisy latent variables by a fixed Markov Chain. The reverse process starts from noise and denoises it step by step via learning the posterior distribution.

To ease the estimation of the posterior distribution, it is critical to design a *mild* forward corruption process [29], in which latent variables in neighboring steps do not differ too much and become increasingly chaotic with the growth of forward steps (see Fig. 1a). However, designing such a process for layout is non-trivial, due to the *heterogeneous* nature of the layout sequence, where the tokens representing element types are *categorical* while the tokens representing element coordinates are *ordinal*. Existing discrete forward processes hardly consider heterogeneous tokens. Directly applying them to layout data often leads to harsh corruptions, where a layout is changed dramatically at each step (see Figs. 1b and 1c). For example, the uniform process in Fig. 1c will transition an element type token to a coordinate token, drastically violating the layout semantics.

To realize a mild corruption process for layout, we make three important observations. (i) *Legality*. The transition between type tokens and coordinate tokens will lead to an illegal layout sequence, resulting in an unpredictable change between forward steps. Hence, it is vital to impose legality during the corruption process. (ii) *Coordinate Proximity*. Coordinate tokens are ordinal, and thus transitioning a coordinate token to its proximal tokens (e.g., from 0 to 1) will introduce a milder change to a layout compared with transitioning to distant ones (e.g., from 0 to 127). (iii) *Type Disruption*. Unlike coordinate tokens, type tokens are categorical and do not have particular proximity. Simply transitioning one type to another may cause abrupt semantic changes to a layout (e.g., from a button to a background image).

Motivated by the above observations, we propose a block-wise transition matrix coupled with a piece-wise linear noise schedule in *LayoutDiffusion*. The transition matrix is designed as follows. First, to achieve legality, we only allow the internal transition between coordinate tokens and that between type tokens. Second, regarding coordinate proximity, we leverage discretized Gaussian [2], where the transition between more proximal tokens takes a higher probability, for the transition between coordinate tokens. Third, as for type disruption, we introduce absorbing state [2]. Each type token either stays the same or transitions to the absorbing state. To further alleviate type disruption, we propose a piece-wise linear noise schedule to make the transition for element types only occur in the late stage of the forward process. With above techniques, *LayoutDif-*

*fusion* achieves the mild forward process shown in Fig. 1a.

Our design also enables *LayoutDiffusion* to perform certain conditional layout generation tasks in a *plug-and-play* manner without re-training, which has never been explored by previous work. Specifically, owing to the mild forward process achieved by *LayoutDiffusion*, its reverse process is to iteratively improve a layout, which naturally supports the task of layout refinement [32]. Besides, as the transition of element types only occurs in the late forward process, *LayoutDiffusion* will determine the element types in a layout quickly in the reverse process. Thus, it can perform generation conditioned on types by simply keeping the types fixed and running the reverse process.

In summary, this work makes four key contributions:

1. We formulate layout generation as a discrete diffusion process, which addresses biased context modeling by iterative refinement from a non-autoregressive model.
2. We design a new diffusion process based on the heterogeneous nature of layout sequence (legality, coordinate proximity and type disruption). It not only better suits layout data but also showcases a promising way of applying diffusion models to other heterogeneous data.
3. We enable certain conditional layout generation tasks in a plug-and-play manner without re-training.
4. We make extensive experiments and user studies. *LayoutDiffusion* outperforms existing methods on all the tasks in terms of most evaluation metrics, even if it is not re-trained for conditional generation tasks.

## 2. Related Work

**Graphic Layout Generation.** Early work on graphic layout generation has explored classical optimization approaches [30, 31], as well as generative models such as Generative Adversarial Networks (GANs) [19, 23] and Variational Autoencoders (VAEs) [1, 17, 18, 42].

Recently, inspired by the success of NLP, masking strategies [20], language models [11], and encoder-decoder architectures [16] have been studied. These approaches represent the layout as a sequence of elements and use Transformer [40] as the basic model architecture. As the placement of one element can depend on any part of a layout, one critical issue in layout generation is *global context modeling*. Some previous studies introduce unnatural biases and fail to model global context effectively [1, 11, 16, 17, 42]. They generate the layout sequence in an autoregressive manner, where there is a predefined generation order and the placement of one element can only depend on the generated part of the layout. On the other hand, a few other studies consider global context but do not achieve significantly better performance [19, 20, 23]. They generate the layout sequence in a non-autoregressive manner, where there is no predefined generation order, and all the tokens are generated in parallel. However, generating a sequence in a single pass

is too challenging [8]. BLT [20] explored an iterative refinement mechanism to alleviate the difficulty. However, it relies on heuristic rules instead of being learned from the data. The above limitation motivates us to seek a better model for layout generation. We think diffusion models are well-suited. By multiple rounds of denoising, it naturally takes the layout in the last step as the global context and generates a layout iteratively instead of by a single pass.

Another branch of studies has explored incorporating diverse user constraints into the layout generation [22, 24, 32, 45, 47]. They treat layout generation tasks with different constraints separately, which introduces repetitive training and hinders knowledge sharing across different tasks. By utilizing the flexible forward process of diffusion models [3], we enable some conditional generation tasks without re-training for the first time, which can be potentially extended to handle more conditional generation tasks.

**Diffusion Models for Discrete Data.** Diffusion models on continuous data have achieved outstanding results [43]. Recently, diffusion models on discrete data are also emerging. They can be grouped into two categories. The first category [5, 9, 25] maps discrete data to continuous state space via a learnable or fixed embedding, and then utilizes techniques from classical continuous diffusion models. These approaches enable simple technology migration from continuous diffusion models, but make the fine-grained control of the forward corruption process much difficult. Another category [2, 4, 10, 15, 39, 41] chooses to directly perform diffusion in discrete state space by modeling the forward corruption process as a random walk between different states. This category makes it easy to incorporate domain-dependent structure to the transition matrices and thus enables flexible control of the forward process. Different to the discrete data (e.g., images and texts) explored by previous work, the layout data studied in this work is heterogeneous by nature. Thus, we fully consider such characteristic and propose a new transition matrix coupled with noise schedules to achieve a mild corruption process.

### 3. Problem Formulation

**Graphic Layout.** A graphic layout  $x$  is composed of a set of graphic elements  $\{e_i\}_{i=1}^N$ , where  $N$  denotes the number of elements. Each element  $e_i$  has an element type  $c_i$  and a bounding box indicating its left  $l_i$ , top  $t_i$ , right  $r_i$ , and bottom  $b_i$  coordinates. Following the advanced layout generation methods [1, 11, 17, 27, 32, 42], we represent an element as a sequence with 5 discrete tokens, i.e.,  $e_i = \{c_i l_i t_i r_i b_i\}$ , where the continuous bounding box coordinates are uniformly discretized into integers between  $[0, K)$ . Then, we represent a layout as a concatenation of element sequences:

$$\mathbf{x} = \{\langle \text{sos} \rangle c_1 l_1 t_1 r_1 b_1 \| \dots \| c_N l_N t_N r_N b_N \langle \text{eos} \rangle\}, \quad (1)$$

where  $\langle \text{sos} \rangle$  and  $\langle \text{eos} \rangle$  are special tokens indicating the start and end of a sequence, and token  $\|$  indicates the separator between any two elements. Obviously, the layout sequence is *heterogeneous*. The element type tokens are *categorical*, while the coordinate tokens are *ordinal*.

**Graphic Layout Generation.** In this work, we primarily focus on unconditional layout generation. Specifically, we learn a generative model  $p_\theta(\mathbf{x})$  parameterized by  $\theta$ , which synthesizes diverse and high-quality graphic layouts.

### 4. LayoutDiffusion

We formulate the layout generation problem as a discrete denoising diffusion process (see Fig. 2). It consists of two Markov chains, where the forward process is hand-designed and fixed while the reverse process is parameterized.

Give a real layout  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ , the *forward* process corrupts it into a sequence of increasingly noisy latent variables  $\mathbf{x}_{1:T} = \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T$ ,

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad (2)$$

$$q(x_t|x_{t-1}) = x_t \mathbf{Q}_t x_{t-1}. \quad (3)$$

Here,  $x_t$  denotes the one-hot version of a single discrete token in the layout sequence  $\mathbf{x}_t$ .  $\mathbf{Q}_t$  is the transition matrix, where  $[\mathbf{Q}_t]_{ij} = q(x_t = j|x_{t-1} = i)$  represents the probabilities that  $x_{t-1}$  transitions to  $x_t$ . Due to the property of Markov chain, the cumulative probability of  $x_t$  at arbitrary timestep from  $x_0$  can be derived as  $q(x_t|x_0) = \mathbf{x}_t \bar{\mathbf{Q}}_t \mathbf{x}_0$ , where  $\bar{\mathbf{Q}}_t = \mathbf{Q}_1 \mathbf{Q}_2 \dots \mathbf{Q}_t$  (refer to [2] for details).

To generate a layout, the *reverse* process starts with a random noise  $\mathbf{x}_T$  and gradually recovers it relying on the learned posterior distribution  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ ,

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t). \quad (4)$$

In the following, we will introduce how to construct a mild forward process  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$  for layout generation (Sec. 4.1), and how to learn the generative model  $p_\theta(\mathbf{x}_0)$  in the reverse process (Sec. 4.2).

#### 4.1. Forward Process

In LayoutDiffusion, we propose a block-wise transition matrix  $\mathbf{Q}_t$  and a piece-wise linear noise schedule to realize a mild forward process, in which layouts in the neighboring steps do not differ too much and become increasingly disordered as the forward step grows (see Fig. 1a).

The design of the transition matrix and noise schedule stems from our three important observations. (i) *Legality*. As defined in Sec. 3, layout sequence has a rigorous format. Any transition between element type tokens and coordinate

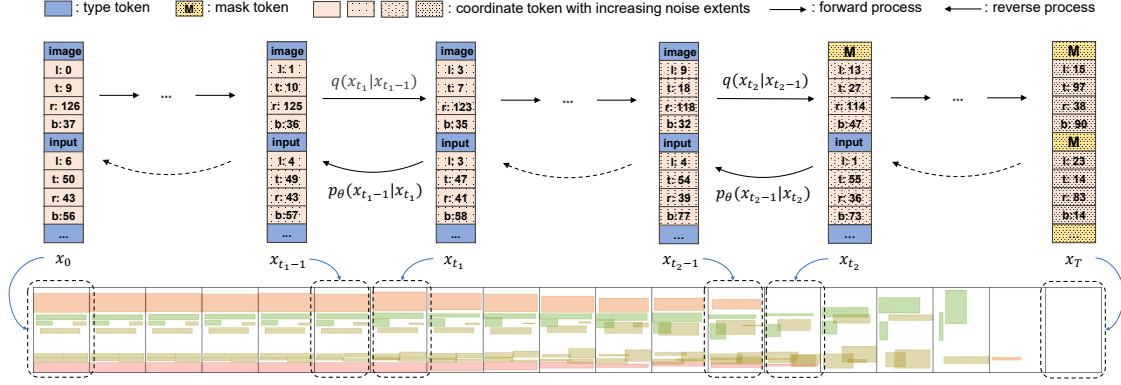


Figure 2. An illustration for **LayoutDiffusion**. In the forward process, the coordinates are mildly corrupted into stationary distribution, and the element types are absorbed into MASK in the late stage. In the reverse process, the element types are first recovered, and then the rough coordinates are gradually refined. For brevity, only two elements are shown, while the other elements and the special tokens are omitted.

tokens will lead to an illegal layout sequence, resulting in a disruptive change between forward steps. Hence, it is vital to impose sequence legality in the transition matrix. (ii) *Coordinate Proximity*. Coordinate tokens in layout sequence are ordinal and have a meaningful proximity. Transitioning a coordinate token to its proximal tokens (e.g., from 0 to 1) will introduce a milder change to a layout, compared with transitioning to distant ones (e.g., from 0 to 127). Thus, it is helpful to encode the proximity prior in the transition matrix. (iii) *Type Disruption*. Type tokens are categorical and do not present particular proximity. Each type of element has its unique coordinate distribution. For example, a background image tends to have a large size, while a button has a small size. Transitioning a type to another type may produce an abnormal element (e.g., a button has a large size and is placed at the top-left corner), leading to abrupt changes in layout. This is also consistent with the observation from diffusion models on other categorical data, e.g., latent code and text [2, 10]. Therefore, it is beneficial to alleviate type disruption in the transition matrix and noise schedule.

**Transition Matrices.** There are three kinds of tokens in the layout sequence, including type tokens (i.e.,  $c_i$ ), coordinate tokens (i.e.,  $l_i$ ,  $t_i$ ,  $r_i$  and  $b_i$ ) and special tokens (i.e.,  $\langle \text{sos} \rangle$ ,  $\langle \text{eos} \rangle$ ,  $\|$ , and  $\text{PAD}$ ). Denote the number of different coordinate tokens and type tokens as  $K$  and  $C$ . Then, the transition matrix is denoted as  $\mathbf{Q}_t \in \mathbb{R}^{V \times V}$ , where  $V = K + C + 4$ .

To achieve the legality of the layout sequence, we only allow the internal transition within each kind of tokens. Thus,  $\mathbf{Q}_t$  can be reduced to a block-wise diagonal matrix,

$$\mathbf{Q}_t = \begin{bmatrix} \mathbf{Q}_t^{\text{coord}} & & \\ & \mathbf{Q}_t^{\text{type}} & \\ & & \mathbf{Q}_t^{\text{spec}} \end{bmatrix}, \quad (5)$$

where  $\mathbf{Q}_t^{\text{coord}}$ ,  $\mathbf{Q}_t^{\text{type}}$  and  $\mathbf{Q}_t^{\text{spec}}$  depicts the probabilities of the internal transition within coordinate tokens, type tokens and special tokens, respectively.

For  $\mathbf{Q}_t^{\text{coord}}$ , to encode the ordinal proximity, we introduce the discretized Gaussian matrix [2] for coordinate tokens, which assigns a higher probability to the transition between more proximal tokens,

$$[\mathbf{Q}_t^{\text{coord}}]_{ij} = \begin{cases} \frac{\exp\left(-\frac{4|i-j|^2}{(K-1)^2\beta_t}\right)}{\sum_{n=-(K-1)}^{(K-1)} \exp\left(-\frac{4n^2}{(K-1)^2\beta_t}\right)}, & i \neq j \\ 1 - \sum_{l=0, l \neq i}^{(K-1)} [\mathbf{Q}_t^{\text{coord}}]_{il}, & i = j \end{cases} \quad (6)$$

where the parameters  $\beta_t$  influence the variance of the forward process distributions.

For  $\mathbf{Q}_t^{\text{type}}$ , to alleviate the type disruption, we choose to transit a type token to a special MASK token instead of another meaningful type token. Therefore, we introduce the absorbing state transition matrix [2] for type tokens,

$$\mathbf{Q}_t^{\text{type}} = \begin{bmatrix} 1 - \gamma_t & 0 & \cdots & 0 \\ 0 & 1 - \gamma_t & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_t & \gamma_t & \cdots & 1 \end{bmatrix}, \quad (7)$$

where  $\gamma_t$  indicates the probability that a token is absorbed into a MASK token, and  $1 - \gamma_t$  is the probability that a token stays unchanged.

For  $\mathbf{Q}_t^{\text{spec}}$ , as special tokens describe the structure of the layout sequence, any transition between them will lead to an invalid layout sequence. Therefore, we choose to disable any transition between them,

$$\mathbf{Q}_t^{\text{spec}} = \mathbf{I}, \quad (8)$$

where  $\mathbf{I}$  is an identity matrix.

**Noise Schedules.** An early absorbing of type tokens (i.e., transitioning to MASK token) will bring an abrupt change to the layout. Hence, to further eliminate type disruption, we choose to make the element type begin to change only

in the late stage of the forward process. Specifically, we design  $\bar{\gamma}_t = 1 - \prod_{i=1}^t (1 - \gamma_i)$  for the cumulative probability  $q_t(x_t|x_0)$  as a piece-wise linear function,

$$\bar{\gamma}_t = \begin{cases} 0, & t < \tilde{T} \\ (t - \tilde{T}) / (T - \tilde{T}), & t \geq \tilde{T} \end{cases} \quad (9)$$

Here,  $\tilde{T}$  is the timestep where the absorbing is enabled, and  $T$  is the terminal timestep.

Besides, although existing work often uses linear schedule for Gaussian transition process, we choose to use  $\beta_t = g / (T - t + \epsilon)^h$  for the transition of coordinate tokens  $\mathbf{Q}_t^{\text{coord}}$ . Here  $g$  and  $h$  are hyper-parameters, and  $\epsilon$  denotes a small positive quantity. It is generalized from a commonly used noise schedule  $1 / (T - t + 1)$  [2, 37]. We find that with  $h > 1$ , it achieves a slower and more smooth corruption to the layout in the early forward process, which helps the model in the reverse process better learn the posterior distribution.

## 4.2. Reverse Process

To reverse the forward process, we optimize the generative model  $p_\theta(\mathbf{x}_0)$  to fit the data distribution  $q(\mathbf{x}_0)$  by minimizing the variational lower bound (VLB) [2],

$$\mathcal{L}_{\text{VLB}} = -\log p_\theta(\mathbf{x}_0|\mathbf{x}_1) + D_{\text{KL}}(q(\mathbf{x}_T|\mathbf{x}_0) \| p(\mathbf{x}_T)) \quad (10)$$

$$+ \sum_{t=2}^T D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)).$$

Following recent work [2, 10], we predict  $p_\theta(\mathbf{x}_0|\mathbf{x}_t)$  instead of  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ , and encourage good predictions of  $\mathbf{x}_0$  at each step by combining  $\mathcal{L}_{\text{VLB}}$  with an auxiliary objective,

$$\mathcal{L} = \mathcal{L}_{\text{VLB}} - \lambda \log p_\theta(\mathbf{x}_0|\mathbf{x}_t). \quad (11)$$

Specifically, we leverage Transformer encoder [40] to learn  $p_\theta(\mathbf{x}_0|\mathbf{x}_t)$ . Denote the embedding of  $i$ -th token in the layout sequence  $\mathbf{x}_t$  as  $\text{emb}(x_{t,i})$  and its positional embedding as  $p_i$ . Denote the embedding of the timestep  $t$  as  $\text{emb}(t)$ . Then, Transformer takes the aggregation of them, i.e.,  $\{\text{emb}(x_{t,i}) + p_i + \text{emb}(t)\}_{i=1}^M$ , as the input and predicts a new layout sequence  $\tilde{\mathbf{x}}_0 = \{\tilde{x}_{0,i}\}_{i=1}^M$  as the output.

In practice, we set an  $N$  as the maximum number of elements. During inference, we first sample an element count  $n$  from the training set’s prior distribution. For constructing  $\mathbf{x}_T$ , we assign MASK tokens for the type and random coordinate tokens for bounding boxes of the first  $n$  elements. For the remaining  $(N - n)$  elements, PAD tokens are utilized to ensure a consistent length. By performing denoising from timestep  $T$  to 0, we derive the layout  $\mathbf{x}_0$ .

## 4.3. Enabling Conditional Layout Generation in a Plug-and-Play Manner

Although LayoutDiffusion is trained for unconditional layout generation, it can handle some conditional gener-

ation tasks without re-training, which has never been explored by previous work. Such a plug-and-play feature of conditional generation is enabled by the design of transition matrices and noise schedules. In the following, we introduce how LayoutDiffusion achieves it.

**Refinement** is a user-oriented layout generation task first posed in RUIE [32], and is recently studied by LayoutFormer++ [16]. Its goal is to take a user given flawed layout as input and provide a high-quality layout for the user while maintaining the original design style. With the proposed transition matrices and noise schedules, a layout is gradually corrupted in the forward process. With such a forward process, the reverse process learned by LayoutDiffusion is to iteratively improve a layout, which naturally enables refinement. Specifically, in LayoutDiffusion, we achieve refinement by feeding the flawed layout into the model and then running reverse process from a certain timestep. Here the timestep is related to how noisy the input layout is.

**Generation Conditioned on Types (Gen-Type)** is also a widely studied conditional layout generation task [16, 20, 22] to satisfy the needs of user. It aims to generate layouts with the given element types. In LayoutDiffusion, there is no transition between coordinates and types (see Eq. (5)). Besides, with the noise schedule in Eq. (9), the change of the types only occurs in the late forward process. With the above two mechanisms, LayoutDiffusion will determine the element types in the early reverse steps very quickly and then continue to improve the coordinates in the remaining reverse steps without changing the types (see the transformation of the layout from right to left in Fig. 2). In other words, the generation for coordinates and that for element types are approximately decoupled. Thus, in LayoutDiffusion, we achieve Gen-Type by feeding in the element types in the early stage and running the reverse process.

# 5. Experiments

## 5.1. Setups

**Datasets.** We employ two widely-used public datasets of graphic layouts. *RICO* [6] is a dataset of user interface designs for mobile applications, which contains 66K+ UI layouts with 25 element types. *PublayNet* [46] includes 360K+ annotated scientific document layouts with 5 element types. Both datasets contain a few over-length entries. We filter out the layouts longer than 20 elements as in LayoutFormer++ [16]. Then, we split the filtered data into a training, validation, and test set by 90%, 5%, and 5%.

**Baselines.** First, we compare LayoutDiffusion with leading approaches for layout generation. Specifically, we compare against *LayoutTransformer* [11], *VTN* [1], *Coarse2Fine* [17], and *LayoutFormer++* [16] on unconditional generation (UGen); against *NDN-none* [22], *LayoutGan++* [19], *BLT* [20], and *LayoutFormer++* [16]

Subtasks	Methods	RICO				PubLayNet			
		mIoU ( $\uparrow$ )	Overlap ( $\rightarrow$ )	Align. ( $\rightarrow$ )	FID ( $\downarrow$ )	mIoU ( $\uparrow$ )	Overlap ( $\rightarrow$ )	Align. ( $\rightarrow$ )	FID ( $\downarrow$ )
Un-Gen	LayoutFormer	0.587	<u>0.542</u>	0.037	24.320	0.359	<u>0.0045</u>	0.067	30.048
	VTN	0.336	0.561	0.477	88.115	0.312	0.221	0.207	105.909
	Coarse2Fine	0.360	0.676	<u>0.128</u>	46.483	0.361	0.142	0.221	50.854
	LayoutFormer++	<u>0.634</u>	0.546	0.051	20.198	0.401	0.0010	<b>0.028</b>	47.082
	Diffusion-LM $\diamond$	<b>0.662</b>	0.631	0.184	11.448	<b>0.439</b>	0.0125	0.076	11.895
	D3PM (absorbing) $\diamond$	0.585	0.619	0.157	<u>4.985</u>	0.401	0.0427	0.075	12.218
	D3PM (uniform) $\diamond$	0.595	0.658	0.229	5.576	0.405	0.0571	0.099	<u>11.212</u>
LayoutDiffusion $\diamond$ (ours)	0.620	<b>0.502</b>	<b>0.069</b>	<b>2.490</b>	<u>0.417</u>	<b>0.0030</b>	<u>0.065</u>	<b>8.625</b>	
Gen-Type	NDN-none	<u>0.35</u>	0.55	0.56	13.76	0.31	0.17	0.35	35.67
	LayoutGan++	0.298	0.620	0.261	5.954	0.297	0.148	0.124	14.875
	BLT	0.216	0.983	0.150	25.633	0.140	0.196	0.036	38.684
	LayoutFormer++	<b>0.377</b>	<u>0.537</u>	<u>0.124</u>	<u>2.483</u>	0.333	<u>0.009</u>	<b>0.025</b>	10.151
	Diffusion-LM $\diamond$	0.324	0.574	0.199	6.530	0.316	0.026	0.046	<u>7.396</u>
	D3PM (absorbing) $\diamond$	0.337	0.594	0.192	3.506	0.333	0.058	0.057	8.858
	D3PM (uniform) $\diamond$	0.317	0.621	0.218	5.771	<b>0.351</b>	0.063	0.064	8.275
LayoutDiffusion $\diamond$ (ours)	0.345	<b>0.491</b>	<b>0.124</b>	<b>1.557</b>	<u>0.343</u>	<b>0.005</b>	<u>0.029</u>	<b>3.731</b>	
Refinement	RUITE	<u>0.658</u>	<u>0.492</u>	0.177	7.926	0.637	0.0375	0.073	7.890
	LayoutFormer++	0.656	0.503	<u>0.141</u>	<u>3.666</u>	<u>0.642</u>	<u>0.0126</u>	<u>0.042</u>	<u>2.937</u>
	Diffusion-LM $\diamond$	0.621	0.499	0.181	8.578	0.573	0.0408	0.140	13.985
	D3PM (uniform) $\diamond$	0.568	0.526	0.266	8.407	0.564	0.0515	0.110	14.553
LayoutDiffusion $\diamond$ (ours)	<b>0.719</b>	<b>0.469</b>	<b>0.102</b>	<b>0.549</b>	<b>0.660</b>	<b>0.0079</b>	<b>0.035</b>	<b>2.045</b>	
Real Data	-	0.466	0.093	-	-	0.0031	0.022	-	

Table 1. Quantitative results. Methods with  $\diamond$  are diffusion-based, which achieve conditional generation (i.e., Gen-Type and Refinement) in a *plug-and-play* manner, while other methods require re-training for each subtask. The best and the second best values of each metric are **bold** and underlined respectively. For mIoU, the higher the score, the better the performance (indicated by  $\uparrow$ ). For Overlap and Align, the closer to real data, the better (indicated by  $\rightarrow$ ). For FID, the lower the score, the better the performance (indicated by  $\downarrow$ ).

on generation conditioned on type (Gen-Type); against *RUITE* [32], and *LayoutFormer++* [16] on refinement. Moreover, we compare *LayoutDiffusion* with the existing diffusion models that do not consider the characteristics of layouts. *Diffusion-LM* [25] maps discrete data to continuous state space, while *D3PM (uniform)* [2] and *D3PM (absorbing)* [2] perform diffusion in discrete state space using different transition matrices.

**Implementation Details.** We set the weight of auxiliary loss as  $\lambda = 0.0001$  (see Eq. (11)). In the training, we set the timestep as  $T = 200$ ; in the inference, we set the timesteps  $T_{\text{UGen}} = 200$ ,  $T_{\text{Gen-Type}} = 160$ , and  $T_{\text{Refine}} = 50$  for different generation tasks. For the layout sequence (see Eq. (1)), we arrange elements in the alphabetical order of the type, and each token is embedded with  $d = 128$  dimensions. For the denoising network  $p_{\theta}(\mathbf{x}_0|\mathbf{x}_t)$ , we apply a 12-layer Transformer encoder with 12 attention heads. We train the model using AdamW optimizer [26] with 2  $\sim$  4 NVIDIA V100 GPUs. We also employ importance timestep sampling [29] during the training. See Supplementals for more details.

**Evaluation Metrics.** We adopt four metrics to measure the performance comprehensively. Among them, Fréchet Inception Distance (**FID**) measures the overall performance, while Maximum Interaction over Union (**mIoU**), Alignment (**Align.**) and **Overlap** measure the quality from a specific aspect. Specifically, *FID* computes the distance between the distribution of the generated layouts and that of real layouts. Following the previous practice [19, 22], we train a classification-based neural network to get the feature

embedding for the layout. *mIoU* calculates the maximum IoU between bounding boxes of the generated layouts and those of the real layouts with the same type set [19]. *Align.* measures whether the elements in a generated layout are well-aligned, either by center or by edges. In addition to the original implementation [22], a normalization over the number of elements is applied. *Overlap* measures the overlapping area between elements in the generated layout. Following *LayoutFormer++* [16], we ignore normal overlaps, e.g., elements on top of the background.

## 5.2. Comparison with Existing Approaches for Layout Generation

**Quantitative Analysis.** In Tab. 1, the methods without the symbol  $\diamond$  are existing approaches for layout generation.

First, we compare FID as it is an overall metric for generation performance. *LayoutDiffusion* achieves significantly better FID scores than all other methods. For example, on Un-Gen, *LayoutDiffusion* achieves **2.490** and **8.625** on RICO and PubLayNet datasets, respectively, while the best existing work only achieves 20.198 and 30.048.

Furthermore, we examine individual metrics, including mIoU, Overlap, and Align., each of which measures the quality from a specific aspect. *LayoutDiffusion* is in the top two on almost every metric and frequently achieves the best performance. On the contrary, existing approaches may perform well on a certain metric but fail on the other individual metrics and the overall metric (i.e., FID), indicating that *LayoutDiffusion* is a well-rounded approach. For ex-

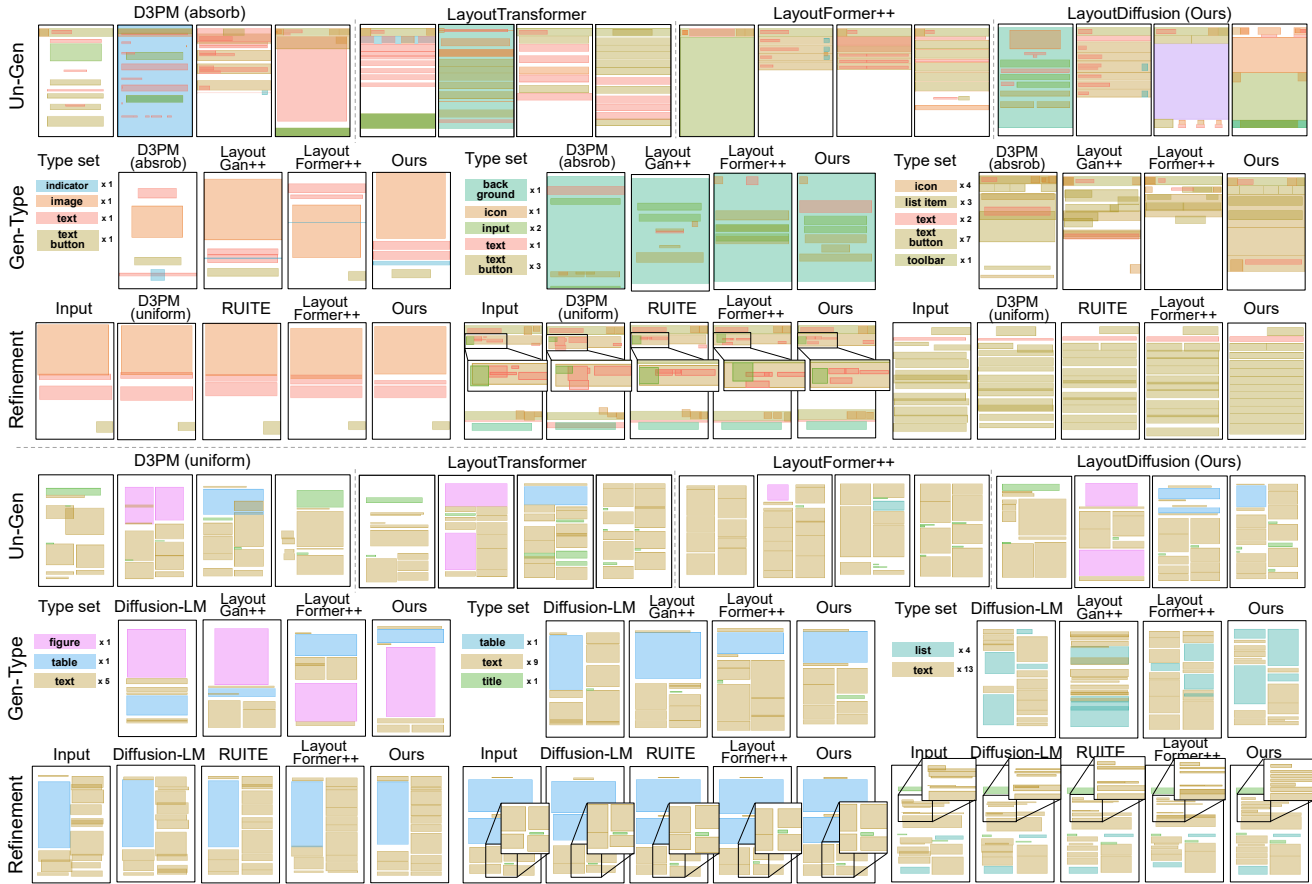


Figure 3. Qualitative comparison against strongest baselines selected by FID (better view in color and 2× zoom). The first three row is for RICO and the last three is for PubLayNet. LayoutDiffusion generates high-quality and diverse layouts. Layouts from LayoutFormer++ either lack diversity (Un-Gen) or are flawed (Gen-Type and Refinement). Layouts from other methods misalign and overlap frequently.

ample, on Un-Gen, LayoutTransformer has a good overlap score, but it does not perform well on Align., mIoU and FID; LayoutFormer++ has the best Align. for PubLayNet, but underperforms in mIoU, overlap and FID.

Moreover, on conditional generation tasks (i.e., Gen-Type and Refinement), the above observations still hold, even though LayoutDiffusion is *not re-trained* on these tasks (see Sec. 4.3) while existing approaches are re-trained.

**Qualitative Analysis.** Fig. 3 shows qualitative results. On Un-Gen, LayoutDiffusion generates diverse and high-quality layouts. In contrast, LayoutTransformer mainly suffers from incorrect spacing and overlap, and LayoutFormer++ is deficient in diversity. For example, for LayoutFormer++, most layouts on RICO contain a top toolbar and several list items, and most layouts on PubLayNet are double-columned and have many texts. Besides, on Gen-Type and Refinement, LayoutDiffusion outperforms other methods (e.g., alignment, overlap and spacing) while it is not re-trained. For more qualitative results, please refer to Supplemental Sec. E.

**User Study.** On each task, we select the best two baselines by FID for the user study. We design two kinds of evalua-

tion. One is *quality* evaluation. We show three layouts from three models respectively (two from baselines and one from LayoutDiffusion) and invite the user to choose which one has the best quality (e.g., more plausible overall structure and pleasing details). Another one is *diversity* evaluation. We show three sets of layouts from three models respectively, where each set contains five layouts from the same model. Then, we invite the user to choose which set has the most diverse layouts. For Refinement, we do not conduct diversity evaluation as it is not necessary for this scenario.

Fig. 4 shows the results. Across different datasets, tasks and evaluation modes, there are 10 groups of user studies in total, in each of which we invite 15 people and everyone labels 50 groups of layouts. The user study shows that LayoutDiffusion outperforms other methods significantly.

### 5.3. Comparison with Traditional Diffusion Models

**Quantitative Analysis.** In Tab. 1, methods marked with  $\diamond$  are traditional diffusion models. They are originally proposed for other generation tasks (e.g., image and text). We adapt them for layout generation. On Un-Gen, LayoutDiffusion achieves the best performance on most metrics. On

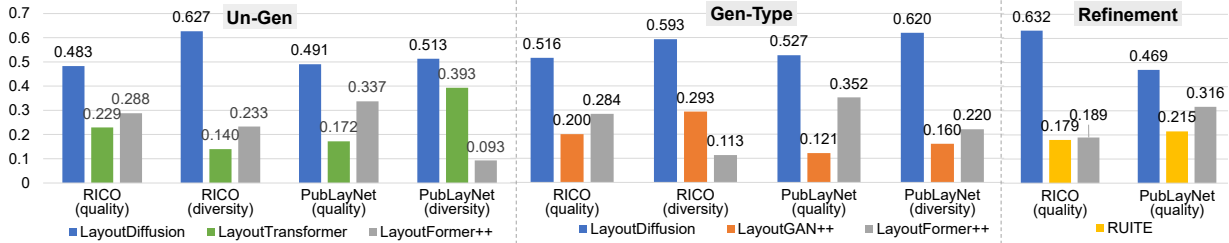


Figure 4. Results of the user study. For each model, we count how many people prefer the layouts generated from this model. The study shows that the results generated by LayoutDiffusion were favored by users over the other methods, particularly in terms of diversity.

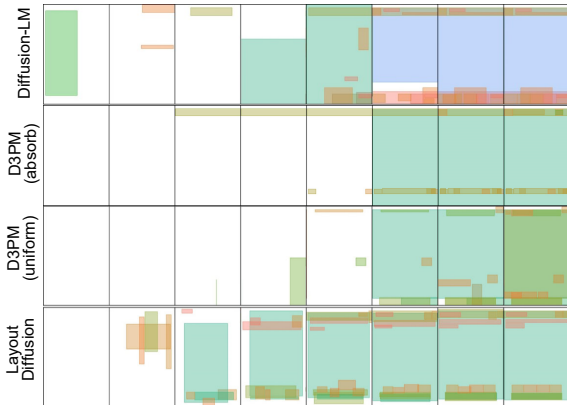


Figure 5. Reverse denoising process for unconditional generation on RICO (from left to right). Each row is for one model. The blank page is used when the generated layout sequence is invalid.

Gen-Type and Refinement, while traditional diffusion models can be used for conditional generation tasks without re-training, their performance is usually worse than existing methods for layout generation (e.g., LayoutFormer++), not to mention LayoutDiffusion. These observations demonstrate that our consideration of the heterogeneous nature of layout data is critical for both achieving good performance and realizing plug-and-play conditional generation.

**Qualitative Analysis.** Fig. 3 shows a qualitative comparison with the best traditional diffusion model (selected by FID). LayoutDiffusion consistently generates better layouts, e.g., better alignment and less overlap. Moreover, Fig. 5 compares the reverse denoising processes of different diffusion models. LayoutDiffusion quickly generates a draft layout and then gradually refines it to a pleasing layout, while other diffusion models take many steps to generate a rough layout and fewer steps for iterative refinement, which may limit the modeling of precise relationships between elements, such as strict alignment and no overlap.

#### 5.4. Ablation Studies and Discussions

**Transition Matrices.** Our transition matrices are designed by considering three critical factors, i.e., legality, coordinate proximity and type disruption (see Sec. 4.1). We remove the technique corresponding to each factor. First, without considering the legality, the techniques for the other two factors cannot be applied. Thus, LayoutDiffusion degrades to

Method	mIoU $\uparrow$	Overlap $\rightarrow$	Align. $\rightarrow$	FID $\downarrow$
Uniform $\mathbf{Q}_t^{\text{coord}}$	0.614	0.633	0.119	3.174
Absorbing $\mathbf{Q}_t^{\text{coord}}$	0.636	0.607	0.114	3.672
Uniform $\mathbf{Q}_t^{\text{type}}$	0.593	0.478	0.139	2.534
Linear $\bar{\gamma}_t$	0.580	0.522	0.156	2.846
Linear $\beta_t$	0.589	0.517	0.202	2.598
LayoutDiffusion	0.620	0.502	0.069	2.490
Real Data	-	0.466	0.093	-

Table 2. Ablation studies on RICO with unconditional generation.

FID $\downarrow$	Number of Steps				
	100	200	500	1000	2000
Diffusion-LM	17.759	17.164	11.984	11.741	11.448
D3PM (absorbing)	7.464	6.102	6.091	4.985	5.110
D3PM (uniform)	6.986	6.910	5.351	5.575	5.239
LayoutDiffusion	3.875	2.490	2.387	2.295	2.360

Table 3. Ablation study on timesteps for diffusion models. All the experiments are on RICO with unconditional generation. The training and inference steps are set as the same. 2000 and 1000 are default settings used in Diffusion-LM [25] and D3PMs [2].

D3PM (absorbing or uniform). Tab. 1 shows that LayoutDiffusion achieves better performance. Second, to ignore coordinate proximity, we use uniform or absorbing transition for coordinate tokens, denoted as Uniform  $\mathbf{Q}_t^{\text{coord}}$  and Absorbing  $\mathbf{Q}_t^{\text{coord}}$  in Tab. 2. LayoutDiffusion outperforms these two variations on most metrics, especially Overlap and FID. Third, to study type disruption, we use uniform transition for type tokens, denoted as Uniform  $\mathbf{Q}_t^{\text{type}}$  in Tab. 2. LayoutDiffusion outperforms this variation, where the improvement of mIoU is most significant.

**Noise Schedules.** To make the corruption of element types occur throughout the forward process, we set  $\hat{T}$  as 0 for  $\bar{\gamma}_t$  (see Eq. (9)), which results in a linear schedule (denoted as Linear  $\bar{\gamma}_t$  in Tab. 2). We replace the noise schedule for  $\beta_t$  with the original linear schedule in previous work [2] (denoted as Linear  $\beta_t$  in Tab. 2). LayoutDiffusion outperforms both of them on most metrics, especially mIoU and Align.

**Timesteps.** Tab. 3 shows FID of diffusion models with different timesteps. While 200 steps are enough for LayoutDiffusion, the performance of Diffusion-LM, D3PM (absorbing), and D3PM (uniform) saturates at 500, 1000, and 500 timesteps respectively. Besides, even the fastest version of LayoutDiffusion surpasses all the other diffusion models.



**Additional Analyses.** For a detailed discussion on the diversity of the generated layouts, see Supplemental Sec. B. Experiments on additional settings of conditional generation are available in Sec. C. Ablation studies concerning conditional generation tasks, noise schedules, and sequence ordering can be found in Sec. D.

## 6. Conclusion

In this work, we propose *LayoutDiffusion* to improve graphic layout generation by discrete diffusion models. The core of our method lies in realizing a mild forward process by considering the heterogeneous characteristics of the layout. Our method also enables two conditional generation tasks without re-training. Experiments demonstrate the superiority of *LayoutDiffusion* over leading approaches for layout generation and existing diffusion models. In the future, we plan to incorporate diverse conditions [14, 28] in *LayoutDiffusion*. Besides, we will also explore how to extend *LayoutDiffusion* to handle other heterogeneous data.

## 7. Acknowledgement

We would like to thank Zhaoyun Jiang for helpful discussions. We also appreciate Zhaoyun Jiang’s support in providing the pre-processed datasets and pretrained models for the FID evaluations.

## References

- [1] Diego Martin Arroyo, Janis Postels, and Federico Tombari. Variational transformer networks for layout generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13642–13652, 2021. 1, 2, 3, 5
- [2] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021. 1, 2, 3, 4, 5, 6, 8
- [3] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022. 3
- [4] Andrew Campbell, Joe Benton, Valentin De Bortoli, Tom Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *arXiv preprint arXiv:2205.14987*, 2022. 3
- [5] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022. 3
- [6] Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hibschman, Daniel Afegan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. Rico: A mobile app dataset for building data-driven design applications. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, pages 845–854, 2017. 5
- [7] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. Mask-predict: Parallel decoding of conditional masked language models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6112–6121, 2019. 1
- [8] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*, 2019. 3
- [9] Shanshan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*, 2022. 3
- [10] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10696–10706, 2022. 3, 4, 5
- [11] Kamal Gupta, Justin Lazarow, Alessandro Achille, Larry S Davis, Vijay Mahadevan, and Abhinav Shrivastava. Layout-transformer: Layout generation and completion with self-attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1004–1014, 2021. 1, 2, 3, 5
- [12] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 1
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 1
- [14] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 9
- [15] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Towards non-autoregressive language models. *CoRR*, abs/2102.05379, 2021. 3
- [16] Zhaoyun Jiang, Jiaqi Guo, Shizhao Sun, Huayu Deng, Zhongkai Wu, Vuksan Mijovic, Zijiang James Yang, Jian-Guang Lou, and Dongmei Zhang. Layoutformer++: Conditional graphic layout generation via constraint serialization and decoding space restriction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18403–18412, 2023. 1, 2, 5, 6
- [17] Zhaoyun Jiang, Shizhao Sun, Jihua Zhu, Jian-Guang Lou, and Dongmei Zhang. Coarse-to-fine generative modeling for graphic layouts. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(1):1096–1103, Jun. 2022. 1, 2, 3, 5
- [18] Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. Layoutvae: Stochastic scene layout generation from a label set. In *Proceedings of the IEEE/CVF*

- International Conference on Computer Vision*, pages 9895–9904, 2019. [2](#)
- [19] Kotaro Kikuchi, Edgar Simo-Serra, Mayu Otani, and Kota Yamaguchi. Constrained graphic layout generation via latent optimization. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 88–96, 2021. [2](#), [5](#), [6](#)
- [20] Xiang Kong, Lu Jiang, Huiwen Chang, Han Zhang, Yuan Hao, Haifeng Gong, and Irfan Essa. Blt: Bidirectional layout transformer for controllable layout generation. *arXiv preprint arXiv:2112.05112*, 2021. [1](#), [2](#), [3](#), [5](#)
- [21] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020. [1](#)
- [22] Hsin-Ying Lee, Lu Jiang, Irfan Essa, Phuong B Le, Haifeng Gong, Ming-Hsuan Yang, and Weilong Yang. Neural design network: Graphic layout generation with constraints. In *European Conference on Computer Vision*, pages 491–506. Springer, 2020. [1](#), [3](#), [5](#), [6](#)
- [23] Jianan Li, Jimei Yang, Aaron Hertzmann, Jianming Zhang, and Tingfa Xu. Layoutgan: Generating graphic layouts with wireframe discriminators. In *International Conference on Learning Representations*, 2018. [2](#)
- [24] Jianan Li, Jimei Yang, Jianming Zhang, Chang Liu, Christina Wang, and Tingfa Xu. Attribute-conditioned layout gan for automatic graphic design. *IEEE Transactions on Visualization and Computer Graphics*, 27(10):4039–4048, 2020. [3](#)
- [25] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *arXiv preprint arXiv:2205.14217*, 2022. [3](#), [6](#), [8](#)
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [6](#)
- [27] David D Nguyen, Surya Nepal, and Salil S Kanhere. Diverse multimedia layout generation with multi choice learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 218–226, 2021. [3](#)
- [28] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. [9](#)
- [29] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [2](#), [6](#)
- [30] Peter O’Donovan, Aseem Agarwala, and Aaron Hertzmann. Designscape: Design with interactive layout suggestions. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 1221–1224, 2015. [2](#)
- [31] Peter O’Donovan, Aseem Agarwala, and Aaron Hertzmann. Learning layouts for single-page graphic designs. *IEEE transactions on visualization and computer graphics*, 20(8):1200–1213, 2014. [2](#)
- [32] Soliha Rahman, Vinoth Pandian Sermuga Pandian, and Matthias Jarke. Ruite: Refining ui layout aesthetics using transformer encoder. In *26th International Conference on Intelligent User Interfaces-Companion*, pages 81–83, 2021. [2](#), [3](#), [5](#), [6](#)
- [33] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#)
- [34] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. [1](#)
- [35] 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. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [1](#)
- [36] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. [1](#)
- [37] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. [1](#), [5](#)
- [38] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. [1](#)
- [39] Zhicong Tang, Shuyang Gu, Jianmin Bao, Dong Chen, and Fang Wen. Improved vector quantized diffusion models. *arXiv preprint arXiv:2205.16007*, 2022. [3](#)
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#), [5](#)
- [41] Pan Xie, Qipeng Zhang, Zexian Li, Hao Tang, Yao Du, and Xiaohui Hu. Vector quantized diffusion model with codeunet for text-to-sign pose sequences generation. *arXiv preprint arXiv:2208.09141*, 2022. [3](#)
- [42] Kota Yamaguchi. Canvasvae: Learning to generate vector graphic documents. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5481–5489, 2021. [2](#), [3](#)
- [43] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*, 2022. [3](#)
- [44] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. *arXiv preprint arXiv:2210.06978*, 2022. [1](#)
- [45] Xinru Zheng, Xiaotian Qiao, Ying Cao, and Rynson WH Lau. Content-aware generative modeling of graphic design

- layouts. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. 3
- [46] Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. Publaynet: largest dataset ever for document layout analysis. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1015–1022. IEEE, 2019. 5
- [47] Min Zhou, Chenchen Xu, Ye Ma, Tiezheng Ge, Yuning Jiang, and Weiwei Xu. Composition-aware graphic layout gan for visual-textual presentation designs. *arXiv preprint arXiv:2205.00303*, 2022. 3